

IBM System Storage DR550 V4.5 Setup and Implementation




New machine type model (2233)



Enhanced RAS features



Storage enhancements



Bertrand Dufrasne
Frank Boerner
Deeporn Beardsley
Torsten Protze
Alexander Safonov
Rene Wuellenweber

Redbooks



International Technical Support Organization

**IBM System Storage DR550 V4.5
Setup and Implementation**

July 2008

Archived

Note: Before using this information and the product it supports, read the information in “Notices” on page ix.

Archived

Sixth Edition (July 2008)

This edition applies to Version 4.5 of the DR550 (product number 2233-DR1 and 2233-DR2).

© Copyright International Business Machines Corporation 2004, 2005, 2006, 2007, 2008. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	ix
Trademarks	x
Preface	xi
The team that wrote this book	xi
Become a published author	xiii
Comments welcome	xiii
Summary of changes	xv
Chapter 1. Introduction	1
1.1 Business environment	2
1.2 Retention-managed data	3
1.3 Storage and data characteristics	4
1.4 IBM strategy and key products	4
1.5 IBM System Storage DR550	7
Chapter 2. DR550 and File System Gateway overview	9
2.1 DR550 Model DR2 overview (2233-DR2)	10
2.1.1 Hardware components	12
2.1.2 Software components	18
2.2 DR550 Model DR1 overview (2233-DR1)	21
2.2.1 Hardware components	22
2.2.2 Software components	25
2.3 DR550 File System Gateway overview	25
2.3.1 Hardware components	26
2.3.2 Software components	27
2.4 Positioning the offerings	28
2.5 What is preconfigured in the DR550	29
2.5.1 DR550 SSAM Server	29
2.5.2 Components security	30
2.5.3 Network configuration	33
2.5.4 HACMP configuration (DR550 Model DR2 dual-node only)	38
2.5.5 DR550 SAN Switch configuration	39
2.5.6 Storage configuration	45
2.5.7 Disk space allocation	52
2.6 What is preconfigured on the optional FSG	56
2.6.1 Hardware components	56
2.6.2 Software components	56
2.6.3 FSG security	57
2.6.4 FSG Network configuration and preset IP addresses	57
2.6.5 Storage configuration	59
2.6.6 Cluster configuration	60
2.7 DR550 offerings summary	60
Chapter 3. DR550 and FSG planning and installation	63
3.1 Planning the installation	64
3.1.1 Planning for the DR550 DR1	66
3.1.2 Planning for the DR550 DR2	68

3.1.3 Planning the optional DR550 FSG	71
3.2 Physical installation	72
3.2.1 Installing the DR550 DR1	72
3.2.2 Installing the DR550 DR2	73
3.2.3 Installing the optional DR550 File System Gateway (FSG)	76
3.3 Power-on and power-off sequence	77
3.3.1 Power-on sequence	77
3.3.2 Power-off sequence	79
3.4 Accessing the DR550 SSAM Servers	81
3.4.1 Accessing the DR550	82
3.5 Accessing the FSG nodes	86
3.6 Required configuration tasks for the DR550	87
3.6.1 Basic setup	88
3.6.2 Attaching the DR550 DR1 to the network	91
3.6.3 Attaching the DR550 DR2 to the IP network	93
3.6.4 Configuring the IP network for a DR550 DR1 and DR2 single-node	94
3.6.5 Configuring the IP address and HACMP for a dual-node DR550 DR2	96
3.6.6 SSAM and data retention policy setup	112
3.7 Required configuration tasks for FSG	113
3.7.1 Prepare DR550 SSAM for FSG attachment	113
3.7.2 Attaching the DR550 File System Gateway to the network	114
3.7.3 Configuring the network settings at the DR550 FSG	116
3.7.4 Time zone configuration and NTP setup at the DR550 FSG node	122
3.7.5 Check network connectivity	124
3.7.6 Customize the FSG software	125
3.7.7 FSG post-installation and initial setup	125
3.7.8 Testing connectivity and initial backup	131
3.8 Changing passwords on base DR550	133
3.9 Changing passwords on the FSG	135
Chapter 4. Operating the DR550 and FSG	137
4.1 Starting and stopping HACMP cluster services	138
4.1.1 Cluster services	138
4.1.2 Starting cluster services	140
4.1.3 Stopping cluster services	144
4.2 Starting and stopping IBM System Storage Archive Manager	148
4.2.1 Starting the System Storage Archive Manager server	148
4.2.2 Stopping the System Storage Archive Manager server	149
4.3 Starting/Stopping the FSG	151
4.3.1 Clustered FSG configuration	151
4.4 DR550 check procedures summary	152
4.4.1 Check software levels	152
4.4.2 AIX basic check	153
4.4.3 Tivoli Storage Manager basic check	153
4.4.4 HACMP basic check	154
4.4.5 Tivoli Storage Manager API checkout	154
4.5 Check FSG status	154
Chapter 5. IBM System Storage Archive Manager	157
5.1 IBM System Storage Archive Manager overview	158
5.1.1 IBM System Storage Archive Manager architecture overview	158
5.1.2 IBM System Storage Archive Manager basic concepts	163

5.2 IBM System Storage Archive Manager	169
5.2.1 Archive copy group retention parameters	170
5.2.2 Chronological archive retention	173
5.2.3 Event-based retention policy	173
5.2.4 Deletion hold and release	175
5.2.5 IBM System Storage Archive Manager prerequisites	175
5.2.6 Data retention protection	177
5.2.7 Expiration processing	178
5.2.8 Encryption	179
5.2.9 Tape drive encryption	180
5.2.10 Data shredding	183
5.2.11 Device support for data retention	188
5.2.12 IBM System Storage Archive Manager and IBM System Storage DR550	188
5.2.13 IBM System Storage Archive Manager policy examples	201
5.3 Explore archive retention features	204
5.3.1 SSAM/Tivoli Storage Manager BA Client V5.5	204
5.3.2 SSAM/Tivoli Storage Manager API (using the sample application dapismp)	214
5.4 IBM Tivoli Storage Manager operational reporting	223
Chapter 6. IBM DR550 File System Gateway	229
6.1 File System Gateway overview	230
6.2 DR550 FSG architecture overview	230
6.2.1 Physical architecture	230
6.2.2 Logical architecture	232
6.2.3 Deletion protection	235
6.2.4 DRG profiles	237
6.3 File System Gateway administration and operations	238
6.3.1 Integrating the DR550 FSG with SSAM Server	238
6.3.2 Enabling deletion protection	239
6.3.3 Editing DRG profiles	241
6.4 Integrating client applications with DR550 FSG	245
6.4.1 Create local users and groups on the DR550 FSG	245
6.4.2 Create and configure shares for NFS exports	248
6.4.3 Create and configure shares for CIFS	253
6.4.4 Client access to the DR550 FSG	257
6.5 DR550 FSG maintenance activities	257
6.5.1 Updating the network configuration	257
6.5.2 Changing and resetting passwords	258
6.5.3 Changing file share customizations	259
6.6 DR550 FSG data archiving and expiration scenarios	260
6.7 DR550 File System Gateway support	266
6.7.1 Failover and data access	266
6.7.2 Manual failover in a high availability cluster	267
6.7.3 Clearing replication errors after failover	267
6.7.4 Recovery from a failed file system	269
6.7.5 Identify orphaned files	270
6.7.6 Wait for pending files to complete replication (HA cluster only)	270
6.7.7 Restore DR550 FSGs to their original roles (HA cluster only)	271
6.7.8 DR550 File System Gateway replacement	272
Chapter 7. Centralized user management and FSG scenarios	275
7.1 Introduction to directories and LDAP	276
7.1.1 Directory components	276

7.1.2 Directory and directory services	276
7.2 FSG configuration for open LDAP	277
7.2.1 Environment for our NFS scenario	277
7.2.2 Installing required LDAP packages	277
7.2.3 Defining users and groups in LDAP	278
7.2.4 Configuring the LDAP client	280
7.3 FSG configuration for Active Directory	283
7.3.1 Environment for our CIFS scenario	283
7.3.2 Preparing Active Directory	284
7.3.3 Preparing the FSG for Active Directory	288
Chapter 8. Using IBM System Storage DR550	303
8.1 Enterprise content management systems and DR550	304
8.2 IBM Enterprise Content Management portfolio	304
8.2.1 IBM Content Manager	306
8.2.2 IBM Content Manager OnDemand	307
8.2.3 IBM FileNet Content Manager	308
8.3 Integrating Content Manager with DR550 SSAM Server	310
8.4 Integrating Content Manager OnDemand with DR550	330
8.5 IBM Optim overview	340
8.5.1 IBM Optim Archive	341
8.6 IBM Optim integration with IBM System Storage DR550	343
8.6.1 Storage Profile	343
8.6.2 Tivoli tab	343
8.6.3 File Retention tab	345
Chapter 9. DR550 call home features	349
9.1 Call home functions	350
9.2 Electronic Service Agent for AIX	351
9.2.1 How eSA for AIX works	351
9.2.2 Configuring eSA for AIX	352
9.3 IBM Director and eSA for System x	359
9.3.1 IBM Director overview	359
9.3.2 IBM Director ISS Extensions for DR550	361
9.3.3 Electronic Service Agent for System x	364
9.4 RSM for DR550	372
9.4.1 Introduction to RSM for DR550	373
9.4.2 Connecting the Remote Support Manager	374
9.4.3 Server installation and configuration	375
9.4.4 RSM installation and configuration	376
9.4.5 RSM users and terminology	379
9.4.6 RSM modem connectivity	381
9.5 IBM Director agent for AIX	384
9.6 CIM agent for DR5550 Storage Controller	385
9.6.1 Introduction	385
9.6.2 Storage Management Initiative - Specification	385
9.6.3 Prerequisites for using SMI-S	385
9.6.4 Installing CIM agent for IBM DS4000	385
9.6.5 Engenio SMI-S Provider service availability	386
9.6.6 IBM Director Storage implementation	386

Chapter 10. DR550 SNMP monitoring	389
10.1 Simple Network Management Protocol (SNMP) overview	390
10.2 Implementation scenario example	393
10.2.1 DR550 SSAM Server error notification and monitoring	395
10.2.2 SNMP setup for error notification on DR550 SSAM Server	395
10.3 Setting up trap handlers on RSM Server for DR550	396
10.3.1 Initial configuration	397
10.3.2 Setting up the trap handler for AIX	398
10.3.3 Setting up the trap handler for HACMP	400
10.3.4 Setting up a trap handler for SSAM	405
10.3.5 Setting up a trap handler for DR550 Storage Controller	409
10.3.6 Setting up a trap handler for SAN switches	413
10.4 Configure SNMP monitoring for DR550 SSAM Server	418
10.4.1 Configure SNMP monitoring for AIX	418
10.4.2 Configure SNMP monitoring for HACMP	426
10.4.3 Configure SNMP monitoring for SSAM	426
10.5 Configure SNMP monitoring for DR550 Storage Controller	435
10.6 Configure SNMP monitoring for SAN Switch	436
10.7 Configure SNMP monitoring for FSG	440
Chapter 11. Tape attachment	447
11.1 Tape device attachment	448
11.2 Planning tape attachment	450
11.2.1 Hardware	452
11.2.2 Software	453
11.3 IBM tape libraries and drives: examples	454
11.3.1 IBM System Storage TS1120 Tape Drive	455
11.3.2 IBM System Storage Enterprise 3592 WORM Tape Cartridges	456
11.3.3 IBM System Storage Enterprise Automated Tape Library (3494)	456
11.3.4 IBM System Storage 3588 Model F3A Ultrium 3 WORM Tape Drive	459
11.3.5 IBM System Storage 3589 Ultrium 3 WORM Tape Cartridge	460
11.3.6 IBM System Storage TS3500 Tape Library	460
11.3.7 Device driver verification for 3592 and 3494	464
11.3.8 Device driver verification for 3588 and TS3500	468
11.4 Integrating devices into System Storage Archive Manager	469
11.4.1 Using the IBM SSAM Administration Center	469
11.4.2 Using the SSAM command line	470
11.4.3 Defining 3592 and 3494 devices to SSAM	470
11.4.4 Integrating 3592 and 3494 into SSAM storage hierarchy	472
11.4.5 Defining 3588 and TS3500 devices to SSAM	472
11.4.6 Integrating 3588 and TS3500 into SSAM storage hierarchy	480
Chapter 12. DR550 backup and restore	491
12.1 What disaster recovery is	492
12.2 Disaster Recovery Manager	492
12.3 Recovery strategy for the server	495
12.4 Using Disaster Recovery Manager on the DR550	497
12.5 Back up and restore the DR550 AIX environment on DVD	501
12.6 File System Gateway (FSG) backup and restore	511
12.6.1 Metadata backup and restore	512
12.6.2 Back up and restore the FSG configuration	512
12.6.3 Full FSG backup and restore	513

Chapter 13. DR550 and Enhanced Remote Mirroring	521
13.1 Enhanced Remote Mirroring overview	522
13.1.1 Requirements	524
13.1.2 ERM terminology	525
13.2 Primary and secondary logical drives	526
13.2.1 Logical drive roles	526
13.2.2 Host accessibility of secondary logical drive	527
13.2.3 Mirrored logical drive controller ownership	527
13.3 Mirror repository logical drives	527
13.4 Mirror relationship	528
13.4.1 Remote Mirror status	529
13.5 Data replication process	531
13.5.1 Metro Mirroring (synchronous mirroring)	531
13.5.2 Global Copy (asynchronous mirroring without write consistency group)	532
13.5.3 Global Mirroring (asynchronous mirroring with consistency group)	533
13.5.4 Data resynchronization process	534
13.5.5 Data synchronization priority	536
13.6 SAN fabric connectivity	536
13.7 Enhanced Remote Mirroring on DR550: Step-by-step	538
13.7.1 Establishing network access to the DR550 Storage Controllers	540
13.7.2 Merging the primary and secondary SAN fabrics	542
13.7.3 Enabling and activating Enhanced Remote Mirroring	546
13.7.4 Creating Enhanced Remote Mirroring relationships	550
13.7.5 Viewing Enhanced Remote Mirroring properties and status	556
13.7.6 Mapping a secondary drive	558
13.7.7 Suspend and resume a mirror relationship	558
13.8 ERM and disaster recovery	562
13.8.1 Role reversal concept	562
13.8.2 Reestablishing Remote Mirroring after failure recovery	565
13.9 Bringing up the secondary DR550: Step by step	566
13.9.1 Single-node DR550	566
13.9.2 Dual-node (HACMP) DR550	567
13.10 ERM maintenance	569
13.11 Performance considerations	571
Related publications	573
IBM Redbooks	573
Other publications	573
Online resources	574
How to get IBM Redbooks	575
Help from IBM	575
Index	577

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.


COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>.

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX 5L™	FlashCopy®	RDN™
AIX®	HACMP™	Redbooks®
Alerts®	i5/OS®	Redbooks (logo)  ®
AS/400®	IBM®	SysBack™
Asset ID™	iSeries®	System p™
BladeCenter®	Lotus Notes®	System p5™
DB2 Universal Database™	Lotus®	System Storage™
DB2®	MVS™	System x™
Domino®	NetView®	Tivoli Enterprise Console®
DPI®	Notes®	Tivoli®
DS4000™	OmniFind™	TotalStorage®
Electronic Service Agent™	POWER™	VideoCharger™
eServer™	POWER5™	xSeries®
FICON®	POWER5+™	
FileNet®	Predictive Failure Analysis®	

The following terms are trademarks of other companies:

FileNet, and the FileNet logo are registered trademarks of FileNet Corporation in the United States, other countries or both.

Novell, SUSE, the Novell logo, and the N logo are registered trademarks of Novell, Inc. in the United States and other countries.

Oracle, JD Edwards, PeopleSoft, Siebel, and TopLink are registered trademarks of Oracle Corporation and/or its affiliates.

SAP R/3, SAP, and SAP logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries.

Java, Power Management, RSM, Ultra, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Active Directory, Internet Explorer, Microsoft, SQL Server, Win32, Windows Server, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel Xeon, Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

Retention-managed data is defined as data that needs to be stored in a non-erasable, non-rewritable format, either because it represents a company asset or to comply with government or industry regulations (data retention for regulatory compliance). As organizations address regulatory and legal compliance requirements, they must ensure that the necessary documents and records are available in their original formats and accessible when needed.

Well-suited for archiving e-mail, digital images, database applications, instant messages, account records, contracts or insurance claim documents, and a range of other information, the IBM® System Storage™ DR550 model DR2 (2233-DR2) and the entry level model DR1 (2233-DR1) offer archival and retention capabilities, as well as synchronous and asynchronous replication capabilities to help organizations address emerging governmental and industry regulatory requirements and corporate governance practices.

This IBM Redbooks® publication explores the characteristics of the various DR550 models with details about how the different elements are initially configured. It explains and illustrates the additional configuration tasks required to deploy the 2233-DR1 and 2233-DR2 products in an existing network and storage environment.

The book also presents key features of IBM System Storage Archive Manager, which is the core of the DR550 and covers the DR550 File System Gateway option, which is designed to provide NFS and CIFS file system access to applications. We also explain how the DR550 integrates with IBM Enterprise Content Management products, to deliver an efficient, fully managed data retention environment.

This book also explains how to complement the solution with tape attachment and the use of WORM tape media for migration and backup purposes. Finally, the book discusses disaster recovery considerations and has a chapter dedicated to the Enhanced Remote Mirroring option.

This edition also covers various call home features and illustrates the use of SNMP and IBM director to monitor and manage the solution.

This book is intended for those who want an understanding of the DR550 (machine type 2233) and targets readers who need detailed advice about how to configure and deploy it.

The team that wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

Bertrand Dufrasne is a Certified Consulting IT Specialist and Project Leader for disk storage products at the International Technical Support Organization, San Jose Center. He has worked at IBM in various IT areas. Before joining the ITSO, he worked for IBM Global Services as an IT Application Architect. He holds a MS degree in Electrical Engineering.

Frank Boerner is an IT Specialist in IBM Germany. He has 14 years of experience with AS/400® and IBM eServer™ iSeries® server, integrated operating environments, and more than eight years of experience with Windows® and Citrix Metaframe integration. His areas of expertise include planning, implementation, and maintenance of complex IBM solutions. He works in the Solution Competence Center and provides DR550 and GAM support for EMEA, AP, and LA.

Deeporn Beardsley is a Storage Service IT Specialist with US IBM System and Technology Group Lab Services currently on a two year temporary assignment to Shanghai, China. Deeporn holds a degree in Computer Science and has over 16 years of experience in development, test, and implementation services. Her areas of expertise include virtualization, SAN, Tivoli® Storage Manager, UNIX/AIX®, and IBM System Storage disk subsystems.

Torsten Protze is an IT Specialist in IBM Germany. He has eight years of experience with IBM Business Continuity and Resiliency Services (BCRS). Within BCRS delivery, he is responsible for Storage, SAN, and Network planning and implementation. Furthermore, he supports customers with their Disaster Recovery tests.

Alexander Safonov is a Senior IT Specialist with System Sales Implementation Services, IBM Global Technology Services Canada. He has over 15 years of experience in the computing industry, with the last 10 years spent working on Storage and UNIX® solutions. He holds multiple product and industry certifications, including Tivoli Storage Manager, AIX, and SNIA. Alexander spends most of his client contracting time working with Tivoli Storage Manager, data archiving, storage virtualization, replication and migration of data. He holds an honors degree in Engineering from the National Aviation University of Ukraine.

Rene Wuellenweber is a System Specialist for DS4000™ and DR550 working for IBM Germany in Leipzig. Rene has seven years of experience as a hardware customer engineer and three years of experience supporting Midrange DASD products. He works in the Solution Competence Center and provides DR550 and GAM support for EMEA, AP, and LA.



Figure 1 The team: Rene, Alex, Frank, Torsten, Deeporn, and Bertrand

Thanks to the following people for their contributions to this project:

Regina Pope-Ford, Funda Eceral, Dean Underwood, Suri Desai, Richard Liwski, Jason Auvenshine, Tony Pacheco, Dan Sivilli, Linda Benhase, Craig Schultz, Peter Suchodolski, Dave Dehaan, Tim Laurence
IBM US

Nils Haustein and Andreas Feldner
IBM Germany

Kai Nunnemann
Becom Informationssysteme, GmbH

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks in one of the following ways:

- Use the online **Contact us** review IBM Redbooks form found at:

ibm.com/redbooks

- Send your comments in an e-mail to:

redbooks@us.ibm.com

- Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Archived

Summary of changes

This section describes the technical changes made in this edition of the book and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes
for SG24-7091-05
for IBM System Storage DR550 V4.5 Setup and Implementation

July 2008, Sixth Edition

This revision reflects the addition, deletion, or modification of new and changed information described below.

New information:

- ▶ Introducing a new Machine Type: 2233 and models DR1 and DR2
- ▶ New alerting and call home support with RSM™ for the DS4200
- ▶ New call home support with Electronic Support Agent for System p™
- ▶ New call home support with the Electronic Support Agent for System x™ (an IBM Director extension)
- ▶ Monitoring using IBM Director and Net-SNMP (for illustration only; Net-SNMP is not officially supported as part of the DR550.)
- ▶ CIM agent for DS4200
- ▶ New Content Management products (Optim)

Changed information:

- ▶ New disk storage (DS4200 and EXP420)
- ▶ New configurations and available disk capacity
- ▶ Modified cabling, including an internal Ethernet network
- ▶ Various software updates (SSAM, Storage Manager, and FSG)
- ▶ HMC removal

Archived

Introduction

This chapter discusses the need for storage products, such as the IBM System Storage DR550 model DR2 (2233-DR2) and model DR1 (2233-DR1), that can address data retention and other regulatory compliance requirements. With data shredding and support for tape hardware encryption, the DR550 Version 4.5 expands security capabilities to help small, medium, and large enterprises address the growing challenge of retention-managed data. The optional DR550 File System Gateway (FSG) enables the DR550 to also support archiving applications that rely on CIFS or NFS protocols, such as the Picture Archiving Communication Solutions (PACS) used by the health care industry.

This chapter discusses the overall business, legal, and regulatory climate, which is the underlying driving force behind the growth in retention-managed data.

We define terms such as *retention-managed data*, *compliance data*, *long-term storage*, just to name a few, and review the concept of *Information Lifecycle Management* (ILM).

Finally, this chapter briefly presents the IBM approach to ILM and key products for data retention solutions.

1.1 Business environment

Information or data is essential to any business and for the most part can be considered a company asset. With that understanding, companies see a potential value in aggregating large amounts of data.

In addition to the sheer growth of data, laws and regulations governing the storage and secure retention of business and client information are increasingly becoming part of the business landscape, making data retention a major challenge to any institution. Businesses must comply with these laws and regulations, including, for example, the Sarbanes Oxley Act, the USA Patriot Act, HIPAA, the European Union Data Protection Directive, Securities Exchange Commission (SEC) Rule 17a-4, and many more. Regulated information can include e-mail, instant messages, business transactions, accounting records, contracts, or insurance claims processing, all of which might need to be retained for varying periods of time. Some of this data might be kept several years. Some data might also be kept forever.

Moreover, some data must be kept just long enough and not any longer. Indeed, content is an asset when it needs to be kept; however, kept past its mandated *retention period*, it could also become a liability. Furthermore, the retention period can change due to factors such as litigation.

All these factors mandate tight coordination and the need for Information Lifecycle Management (ILM). ILM is a process for managing information through its life cycle, from conception until disposal, in a manner that optimizes storage and access at the lowest cost.

ILM is not just hardware or software: It includes processes and policies to manage the information. It is designed based on the recognition that different types of information can have different values at different points in their life cycle. Predicting storage needs and controlling costs can be especially challenging as the business grows. The overall objectives of managing information with Information Lifecycle Management are to help reduce the total cost of ownership (TCO) and help implement data retention and compliance policies. In order to effectively implement ILM, owners of the data need to determine how information is created, how it ages, how it is modified, and if or when it can safely be deleted.

Directives and regulations

It is beyond the scope of this book to enumerate and explain the regulations in existence today. For illustration purposes only, we list some of the major regulations in Table 1-1, summarizing their intent and applicability.

Table 1-1 Major regulations affecting many enterprises

Major regulation	Intent of regulation	Applicability
SEC/NASD	Prevents securities fraud	All financial institutions and companies regulated by the SEC
Sarbanes Oxley Act	Ensures accountability for public firms	All public companies trading on an U.S. Exchange
HIPAA	Privacy and accountability for health care providers and insurers	Health care providers and insurers, both human and veterinarian
US DoD 5015.2	Electronic record keeping	U.S. government, particularly DoD
21 CFR 11	Approval accountability	FDA regulation of pharmaceutical and biotechnology companies

Table 1-2 lists several requirements found in SEC 17a-4 with which financial institutions and broker-dealers must comply. It is this type of data that is referred to as compliance data, a subset of retention-managed data.

Table 1-2 Example of SEC/NASD requirements

Requirement	Met by
Capture all correspondence (unmodified). [17a-4(f)(3)(v)].	Capturing incoming and outgoing e-mail before reaching users
Store in non-rewritable, non-erasable format. [17a-4(f)(2)(ii)(A)]	Write Once Read Many (WORM) storage of all e-mail and all documents
Verify automatically recording integrity and accuracy. [17a-4(f)(2)(ii)(B)]	Validated storage to magnetic media and WORM
Duplicate data and index storage. [17a-4(f)(3)(iii)]	Mirrored or duplicate storage servers (copy pools)
Enforce retention periods on all stored data and indexes. [17a-4(f)(3)(iv)(c)]	Structured records management
Search and retrieve all stored data and indexes. [17a-4(f)(2)(ii)(D)]	High-performance search retrieval

Again, this is just for illustration purposes and is not an exhaustive list of requirements.

1.2 Retention-managed data

Beyond laws and regulations, data often needs to be archived and managed simply because it represents a critical company asset.

Examples of such data include contracts, CAD/CAM designs, aircraft build and maintenance records, and e-mail, including attachments, instant messaging, insurance claim processing, presentations, transaction logs, Web content, user manuals, training material, digitized information (such as check images, medical images, historical documents, and photographs), and much more.

The characteristics of such data can be very different in their representation, size, and industry segment. It becomes apparent that the most important attribute of this kind of data is that it needs to be retained and managed, so it is called *retention-managed data*.

Retention-managed data is data that is written once and is read rarely (sometimes never). Other terms abound to describe this type of data, such as reference data, archive data, content data, or other terms that imply that the data cannot be altered.

Retention-managed data is data that needs to be kept (retained) for a specific (or unspecified) period of time, usually years.

Retention-managed data applies to many types of data and formats across all industries. The file sizes can be small or large, but the volume of data tends to be large (multi-terabyte to petabytes). It is information that might be considered of high value to an organization; therefore, it is retained near-line for fast access. It is typically read infrequently and thus can be stored on economical disk media such as SATA disks; depending on its nature, it can be migrated to tape after some period.

It is also important to recognize what does not qualify as retention-managed data. It is not the data that changes regularly, known as *transaction data* (account balance, inventory status, and orders today, for example). It is not the data that is used and updated every business cycle (usually daily), or the backup copy of this data. The data mentioned here changes regularly, and the copies used for backup and disaster recovery are there for exactly those purposes, meaning backup and disaster recovery. They are there so that you can restore data that was deleted or destroyed, whether by accident, a natural or human-made disaster, or intentionally.

1.3 Storage and data characteristics

When considering the safekeeping of retention-managed data, companies also need to consider storage and data characteristics that differentiate it from transactional data.

Storage characteristics of retention-managed data include:

- ▶ Variable data retention periods: Usually a minimum of a few months, up to forever.
- ▶ Variable data volume: Many customers are starting with 5 to 10 TB of storage for this kind of application (archive) in an enterprise. It also usually consists of a large number of small files.
- ▶ Data access frequency: Write once read rarely or read never. See data life cycle in the following list.
- ▶ Data read/write performance: Write handles volume; read varies by industry and application.
- ▶ Data protection: Pervasive requirements for non-erasability, non-rewritability, and destructive erase (data shredding) when the retention policy expires.

Data characteristics of retention-managed data include:

- ▶ Data life cycle: Usage after capture, 30 to 90 days, and then near zero. Some industries have peaks that require access, such as check images in the tax season.
- ▶ Data rendering after long-term storage: Ability to view or use data stored in a very old data format (say after 20 years).
- ▶ Data mining: With all this data being saved, we think there is intrinsic value in the content of the archive that could be exploited.

1.4 IBM strategy and key products

Regulations and other business imperatives, as we just briefly presented, stress the need for an Information Lifecycle Management process and tools to be in place.

The unique experience of IBM with the broad range of ILM technologies and its broad portfolio of offerings and solutions can help businesses address this particular need and provide them with the best solutions to manage their information throughout its life cycle.

IBM provides a comprehensive and open set of solutions to help. IBM has products that provide content management, data retention management, and sophisticated storage management, along with the storage systems to house the data.

To specifically help companies with their risk and compliance efforts, the IBM Risk and Compliance framework is another tool designed to illustrate the infrastructure capabilities needed to help address the myriad of compliance requirements. Using the framework,

organizations can standardize the use of common technologies to design and deploy a compliance architecture that might help them deal more effectively with compliance initiatives.

Important: The IBM offerings are intended to help customers address the numerous and complex issues relating to data retention in regulated and non-regulated business environments. Nevertheless, each customer's situation is unique, and laws, regulations, and business considerations impacting data retention policies and practices are constantly evolving. Customers remain responsible for ensuring that their information technology systems and data retention practices comply with applicable laws and regulations, and IBM encourages customers to seek appropriate legal counsel to ensure their compliance with those requirements. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer is in compliance with any law.

Key products of the IBM data retention and compliance solutions are IBM System Storage Archive Manager and IBM Enterprise Content Management products, along with any needed disk-based and tape-based storage:

- ▶ IBM software for enterprise content management integrates and delivers critical business information that offers new business value on demand. Enterprise content management software and solutions support multiple information types, such as images, documents, e-mail, and e-records, and provide the appropriate content, based on user intent and relevancy. The IBM Enterprise Content Management portfolio is designed to help transform business with improved productivity and streamlined compliance.
 - IBM Content Manager for Data Retention Compliance is a comprehensive software platform combining IBM DB2® Content Manager, DB2 Records Manager, DB2 CommonStore, and services that are designed to help companies address the data retention requirements of SEC, NASD, and other regulations.
 - Also part of the IBM Enterprise Content Management portfolio is the FileNet® P8 family of products, which includes back-end services, development tools, and applications that address enterprise content and process management requirements. IBM FileNet Content Manager is one of the core products in the FileNet P8 family. IBM FileNet Content Manager provides full content life cycle and extensive document management capabilities for digital content. It combines document management with workflow and process capabilities to automate and drive content-related tasks and activities.
 - The IBM Optim Archive provides everything you need to create and manage archives of relationally intact data from databases with any number of tables, interconnected with any number of DBMS and application-managed relationships, regardless of their complexity. After creating an Archive File, Archive selectively removes data from the production database, according to your instructions, to maximize database performance and response time. Optim Archive can use the DR550 as its storage device.

Refer to 8.2, “IBM Enterprise Content Management portfolio” on page 304 for more details.

- ▶ IBM System Storage Archive Manager offers expanded policy-based data retention capabilities that are designed to provide non-rewritable, non-erasable storage controls to prevent deletion or alteration of data stored using IBM System Storage Archive Manager. These retention features are available to any application that integrates with the open IBM System Storage Manager API.

Key technologies for IBM data retention and archiving solutions include:

- ▶ IBM System Storage DS4200 Storage Server with the EXP420 Storage Expansion Unit: This storage server and the disk expansion enclosure use Serial Advanced Technology Attachment (SATA) disk drives to provide near-line storage at an affordable price. It is also capable of Enhanced Remote Mirroring to a secondary site.
- ▶ With Write Once Read Many (WORM) media technology, the IBM System Storage TS1120 (with WORM cartridges), the IBM 358x family of LTO libraries (with WORM cartridges), and the new TS1040 LTO4 libraries are ideal tape solutions for data retention. Once written, data on the cartridges cannot be overwritten (to delete the data, the tape must be physically destroyed). This capability is of particular interest to clients that need to store large quantities of electronic records to meet regulatory and internal audit requirements.
- ▶ The DR550 File System Gateway is designed to offer file archiving capability without requiring any application enabling, and to provide NFS and CIFS file system access to applications.

Figure 1-1 shows where the DR550 (2233-DR2 or 2233-DR1), built on some of these key products and technologies, fits.

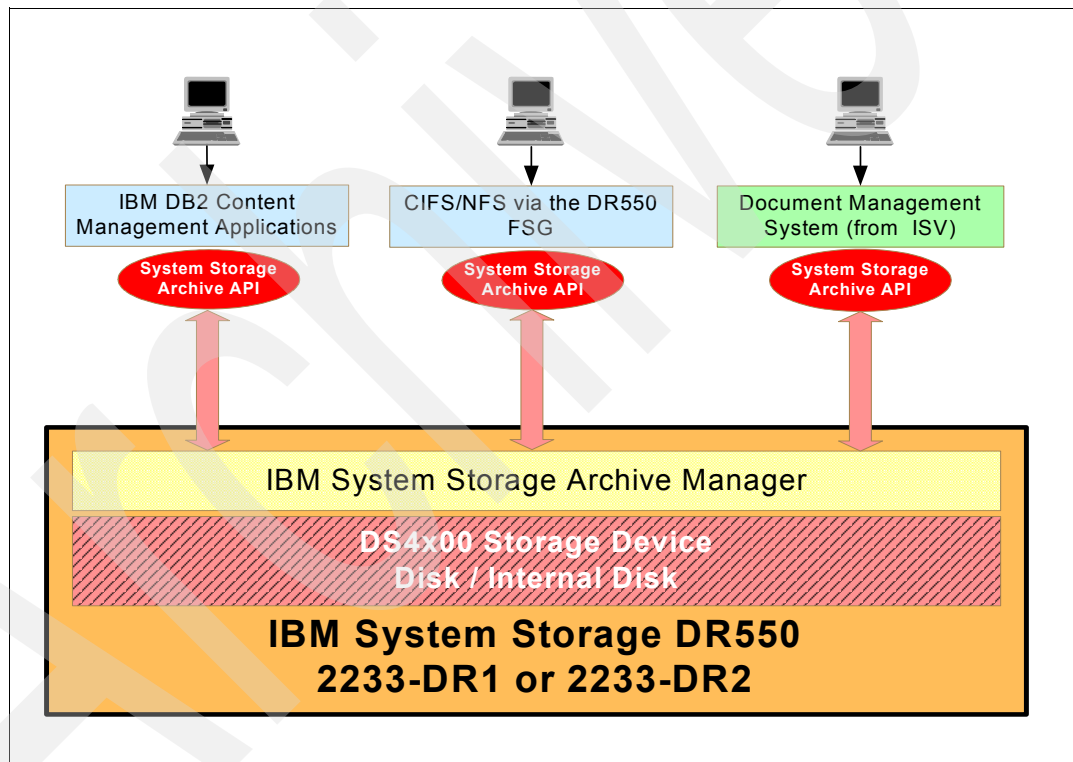


Figure 1-1 IBM DR550 context diagram

The main focus of the IBM System Storage DR550 is to provide for a secure storage system, where deletion or modification of data is completely disallowed except through a well-defined retention and expiration policy.

The IBM System Storage DR550 is the repository for regulated business information. It does not create the information. A complete solution includes applications that gather information, such as e-mail and instant messaging, content management applications, such as IBM DB2 Content Manager, IBM FileNet Content Manager, or other document management systems from independent software vendors that can index, archive, and later present this information

to compliance reviewers, and finally storage systems that retain the data in accordance with regulations.

The remainder of this book offers a comprehensive review of the DR550, discusses the components of the solution, explains how to deploy it, and illustrates its usage.

1.5 IBM System Storage DR550

IBM System Storage DR550 (2233-DR1 and 2233-DR2), one of the IBM Data Retention offerings, is an integrated archiving system for clients that need to retain and preserve electronic business records. It is designed to help store, retrieve, manage, and retain regulated and non-regulated data. In other words, it is not just an offering for compliance data, but can also be an archiving solution for other types of data.

The optional IBM System Storage DR550 File System Gateway enables many additional content-management applications to access the DR550 by adding NFS and CIFS network file access. Previously, only content-management applications that support the IBM System Storage Archive Manager Application Program Interface (SSAM API) were able to access the DR550. The File System Gateway takes files that are sent using network file protocols and associates these files with an IBM System Storage Archive Manager management policy. This is done by using customer-configurable path and file naming pattern matching rules. The File System Gateway sends these files with their associated policies to the DR550 using the SSAM application program interface. Files are retrieved using the same name and directory as they were stored under.

By integrating IBM System p5™ servers (using POWER5™ processors) with IBM System Storage hardware products and IBM System Storage Archive Manager software, the DR550 system is specifically designed to provide a central point of control to help manage growing compliance and data retention needs.

The DR550 brings together off-the-shelf IBM hardware and software products. The hardware comes already mounted in a secure rack. The software is preinstalled and to a large extent preconfigured. The system's compact design can help with fast and easy deployment and incorporates an open and flexible architecture.

Archived

DR550 and File System Gateway overview

The IBM System Storage DR550 product offering consist of two models: the DR550 Model DR2 (2233-DR2) and the DR550 Model DR1 (2233-DR1). Both models offer various configuration and capacity options. All models are shipped as preconfigured and integrate servers, storage, software, and starting with DR550 V4.0, an optional DR550 File System Gateway (FSG). They are aimed at helping customers preserve and retain electronic business records, either to comply with government and industry regulations, or simply because there is a business need for retaining data.

This chapter presents an overview of DR550 Model DR1 (2233-DR1) and Model DR2 (2233-DR2), including the optional DR550 File System Gateway (2229-FSG). First, we review the DR550 as a whole and its intended usage. This is followed by a description of each of the elements, hardware, and software, with detailed information about how they are initially packaged, installed, and configured.

Note: For a brief summary of the DR550 models and options, refer to Table 2-14 on page 60.

2.1 DR550 Model DR2 overview (2233-DR2)

Note: This section only applies to the DR550 Model DR2 single-node or dual-node systems. For the DR550 Model DR1, refer to 2.2, “DR550 Model DR1 overview (2233-DR1)” on page 21.

The DR550 brings together off-the-shelf IBM hardware and software products. The hardware comes premounted in a secure rack. The software is preinstalled and to a large extent preconfigured. This offering is designed to be easy to deploy.

The DR550 Model DR2 is available in several configurations with storage from 6 TB to 168 TB, and is available with Enhanced Remote Mirroring (both Metro Mirroring and Global Mirroring). The DR550 includes FC HBA ports for external tape but does not include cables or tape drives or tape libraries. You should acquire and attach tape drives to be able to back up your DR550 configuration and data (refer to Chapter 11, “Tape attachment” on page 447).

A standard DR550 DR2 (2233-DR2) single node consists of:

- ▶ One DR550 DR2 rack (7014 T00 rack - 36U (ETA) tall)
- ▶ One DR550 SSAM Server (which is an IBM System p5 52A, product number 9131-52A). It is a 4-EIA (4U), 19-inch rack mounted server. The 52A is configured as a 2-core system with 2.1 GHz processors. The total system memory is 2 GB. It also includes standard dual power supplies.
- ▶ One console kit (16-TF3 with Keyboard, Video, Mouse) and (optional) KVM switch
- ▶ One (optional) DR550 SAN Switch (2005-B16 FC Switch)
- ▶ Two DR550 Ethernet Switches (Alphanetworks 16-port Ethernet switches)
- ▶ One or two DR550 Storage Controllers (IBM System Storage DS4200s)
- ▶ One or more DR550 Expansion Drawers (IBM System Storage EXP420)

The recommended DR550 Model DR2 (2233-DR2) dual-node consists of:

- ▶ One DR550 DR2 rack (7014 T00 rack - 36U (ETA) tall)
- ▶ Two DR550 SSAM Servers (System p5 52A) configured in an HACMP™ active/passive cluster
- ▶ Two DR550 SAN Switches (2005-B16 FC Switches)
- ▶ One console kit (16-TF3 with Keyboard, Video, Mouse) and KVM switch
- ▶ One or two DR550 Storage Controllers (IBM System Storage DS4200)
- ▶ One or more DR550 Expansion Drawers (IBM System Storage EXP420)

With the DR550 (2233-DR2), the dual-node configurations gradually increase in capacity going from 6 TB to 168 TB (raw capacity, using 750 GB SATA drives). For up to 48 TB configurations, the storage comes in 6 TB increments. Thereafter, the storage comes in 12 TB increments. For more than 84 TB, the configuration requires a second DR550 DR2 rack (7014-T00) to accommodate the second DR550 Storage Controller (DS4200) and the additional DR550 Expansion Drawers (EXP420s).

The storage is preconfigured as RAID 5. With the DR550 Version 4.5, you also have the option to order a RAID 6 ready configuration.

An Enhanced Remote Mirroring (ERM) option is available for configurations up to 84 TB. Note that this option requires two DR550s (same capacity), preferably at different sites.

The software bundle includes the IBM AIX 5L™ Version 5.3 operating system with HACMP cluster software (for dual-node only), IBM System Storage Archive Manager Version 5.5, and DS4000 Storage Manager for DR550, customized for additional protection, all running on the DR550 SSAM Servers (p52A servers).

In addition, the configurations can also include the optional DR550 File System Gateway (2229-FSG). The FSG is available as stand-alone or high availability 2-node cluster option. This optional gateway provides the interface for NFS or CIFS protocol support. An FSG node uses a dual core Xeon based processor and includes 2 GB of memory (see 2.3, “DR550 File System Gateway overview” on page 25).

Figure 2-1 shows the DR550 with the console kit (monitor and keyboard).



Figure 2-1 DR550 (2233-DR2) dual node with console kit

2.1.1 Hardware components

In this section, we provide additional details for the DR550 (2233-DR2) hardware components.

DR550 DR2 rack

The DR550 DR2 rack is a 7014-T00 rack that stacks all the components vertically. The rack come with doors in the front and in the back and include the Rack Security Kit to secure physical access to any of the DR550 components.

The servers and switches are placed in the bottom half of the rack. The storage units start from the bottom, populating toward the top as the storage capacity installed increases (see Figure 2-7 on page 16). The FSG node(s) (machine type FSG-2229) are housed at the very top. At the high end, the DR550 comes in two vertical racks with the second rack holding the second DR550 Storage Controller and DR550 Expansion Drawers.

Note: When the rack doors are closed, most of the display LEDs and panels normally visible on the different components (storage, switches, and servers) are hidden. These displays show critical status and other information about the health of the components, and it is important that the person monitoring the system can quickly get access to the key when service is required.

DR550 SSAM Server

The DR550 includes one or optionally two DR550 SSAM Servers (IBM System p5, product number 9131-52A) running the IBM AIX 5L Version 5.3 operating system. In the case of the dual-node configuration, both servers have identical hardware, and they are configured as nodes of an HACMP cluster. The DR550 SSAM Server (referred to hereafter as the *p52A*, *p5 52A*, *node*, or *engine*) is a cost-effective, high-performance, and space-efficient server that uses advanced IBM technology. The p52A uses POWER5+™ microprocessor, and is designed for use in LAN clustered environments.

The p52A is a member of the symmetric multiprocessor (SMP) UNIX servers from IBM, and, physically, it is a 4-EIA (4U), 19-inch, rack-mounted server. The p52A is configured as a dual-core system with 2.1 GHz processors.

The total system memory installed is 2 GB. The p52A includes six hot-plug PCI-X slots, an integrated dual channel Ultra320 SCSI controller, two 10/100/1000 Mbps integrated Ethernet controllers, and eight front-accessible disk bays supporting hot-swappable disks. The server shipped in the DR550 Model DR2 includes dual power supplies and has two 73 GB disk drives configured as RAID 1 (mirrored) and mounted in bays position 4 and 5. These disks are for use by the operating system. Figure 2-2 on page 13 shows the DR550 SSAM Server.

For additional details on the DR550 SSAM Server and what adapters (located in the PCI slots) are installed, refer to 2.5.1, “DR550 SSAM Server” on page 29.



Figure 2-2 The DR550 SSAM Server (9131 p5 model 52A)

The redundant hot-plug cooling fans are another reliability and availability feature. Along with the hot-plug components, the p52A is designed to provide an extensive set of reliability, availability, and serviceability (RAS) features that include improved fault isolation, recovery from errors without stopping the system, avoidance of recurring failures, and predictive failure analysis. For more information about the p52A servers, refer to *IBM System p5 520 and 520Q Technical Overview and Introduction*, REDP-4137.

Note: The DR550 SSAM Server now comes standard with two power supplies.

Figure 2-3 shows you the back view of a DR550 SSAM Server.




Figure 2-3 Back view of DR550 SSAM Server

DR550 Storage Controller (DS4200) and DR550 Expansion Drawer (EXP420)

The disk storage used in the DR550 is the IBM System Storage DS4200 Express. Additional storage is added by including DR550 Expansion Drawers. The DR550 Expansion Drawer is an IBM System Storage EXP420. Each EXP420 enclosure packaged with the DR550 includes eight or sixteen 750 GB Serial ATA (SATA) disk drive modules, offering up to 12 TB of raw physical capacity. The DR550 Model DR2 (2233-DR2) can consist of up to two DR550 Storage Controllers (DS4200s) and up to 12 DR550 Expansion Drawers (EXP420s) for a maximum of 168 TB of raw storage.

Attention: Some reconfiguration is required when adding capacity to the DR550.

The DR550 Storage Controller (DS4200) also supports online controller firmware upgrades to help provide better performance and functionality. For further information about the IBM DS4200 Storage Server, see:

<http://www.ibm.com/servers/storage/disk/ds4000/ds4200/index.html>

Restriction: This link is only for access to DS4200 related documentation. To download any firmware or software, you must refer to the DR550 support Web site.

The DS4200 and EXP420s used in the DR550 configurations utilize Serial Advanced Technology Attachment (SATA) disk drives. With the DR550, users get the advanced features of the DS4200 Storage Controller with the cost-effective advantage of SATA disks that are well-suited for fixed content, sparingly accessed storage applications. Figure 2-4 shows the front of the DR550 Storage Controller (DS4200), while Figure 2-5 shows its rear. The EXP420 looks identical from the front except for the label on the front bezel.




Figure 2-4 DR550 Storage Controller(DS4200)




Figure 2-5 DR550 Storage Controller - Rear view (DS4200)

The EXP420 has two hot-swappable Environmental Service Modules (ESMs), two power supplies, and two fan units that provide for sufficient redundancy and availability. The DS4200 and the EXP420 also have hot-swappable drives. The hot-swap drive bays are preinstalled in drive trays. This drive and carrier assembly called a *customer replaceable unit*

(CRU) includes the drive tray, SATA drive, and hard disk drive interposer card; they are installed in the 16 drive bays on the front of the unit. Each of these can be replaced as a unit. Figure 2-6 shows the rear view of DR550 Expansion Drawer (EXP420).




Figure 2-6 DR550 Expansion Drawer(EXP420) - rear view

The storage used in the DR550 comes in several capacities. The minimum capacity in the DR550 model DR2 is 6 TB of raw disk space that is built with eight (half a drawer) 750 GB disks in the DR550 Storage Controller (DS4200) enclosure. The storage is configured as RAID 5 whereby a half drawer, if it is the only storage drawer, is formatted 2+p and 3+p with a global spare; otherwise it is formatted 6+p with one global spare and a full drawer is formatted 6+p and 7+p with one global spare.

Note: You can also order the DR550 Storage in a RAID 6 configuration.

The DR550 comes in configurations with capacities in increments of 6 TB up to 168 TB. A second DR550 Storage Controller (DS4200) is required for configurations including more than six EXP420s (up to a maximum of twelve EXP420s). Those configurations have two vertical racks. One rack contains the DR550 SSAM Servers (p52A servers), the Fibre Channel switches, the Ethernet switches, the console kit, the first DR550 Storage Controller (DS4200), along with at most three EXP420s, and the optional File System Gateways. The other rack contains the second DR550 Storage Controller (DS4200) with up to nine DR550 Expansion Drawers (EXP420s).

Figure 2-7 shows the rack population with the FSG.

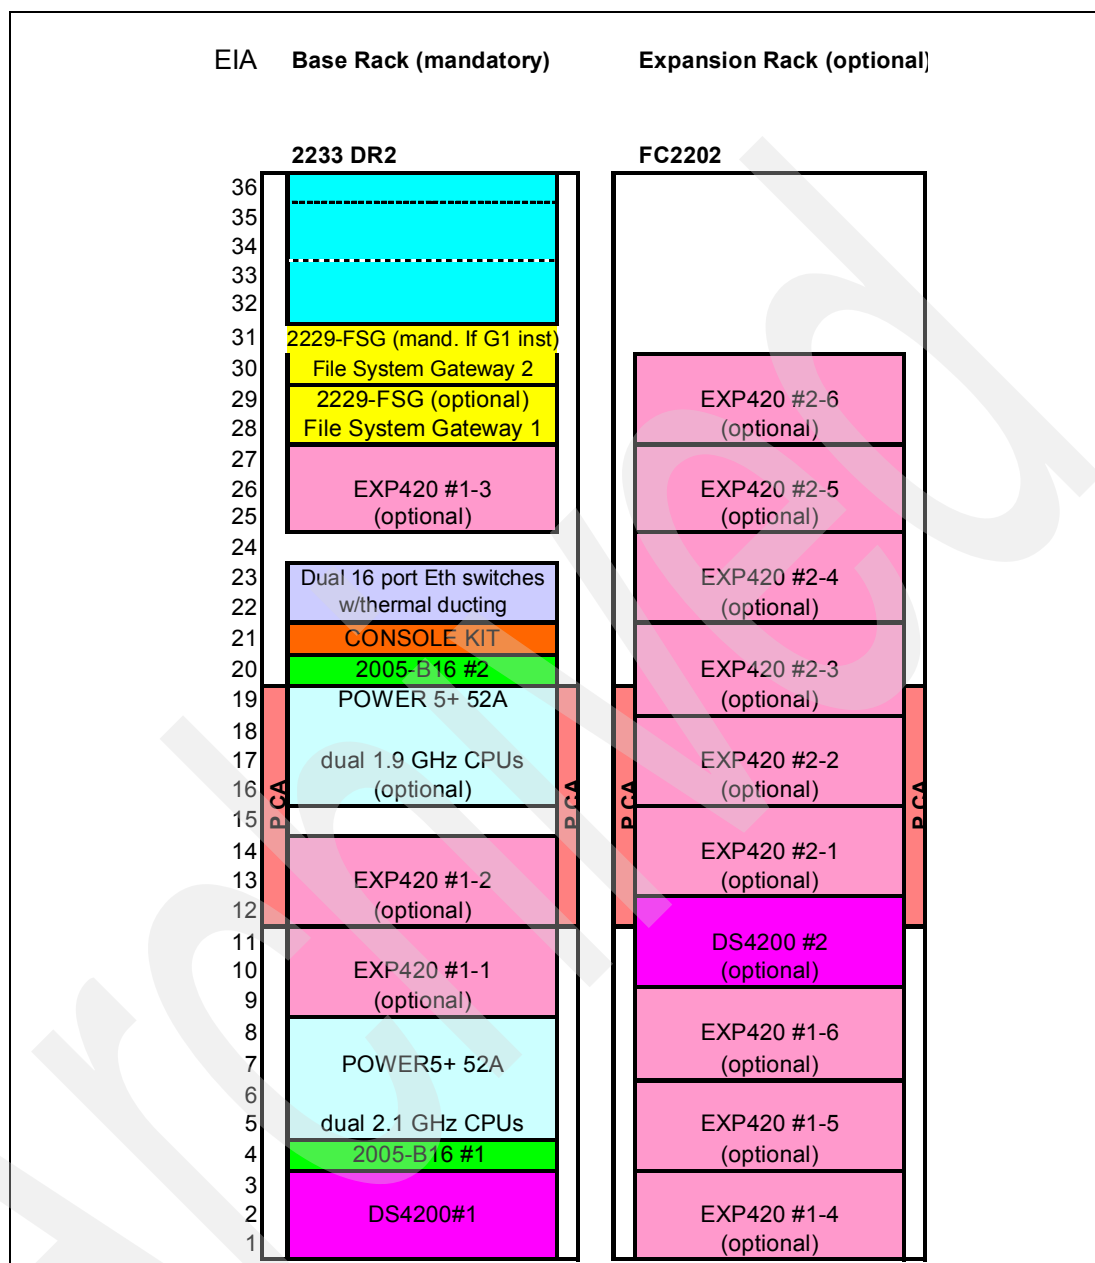


Figure 2-7 Rack population for DR550 Model DR2 (including optional dual node FSG)

DR550 SAN Switch

The DR550 SAN Switch is an IBM System Storage SAN Switch 2005-B16 (B16). It is used to interconnect the DR550 SSAM Servers (p52A servers) with the DR550 Storage Controller (DS4200) and Fibre Channel-based tape solution. The dual-node system includes two DR550 SAN Switches.

The DR550 SAN Switch is a 16-port high performance auto-sensing Fibre Channel switch. With next-generation switch technology, these switches are designed to provide improved availability capabilities, fully non-blocking performance, and advanced intelligence features.

The DR550 SAN Switch is well-suited to address small and medium business (SMB) customer requirements for infrastructure simplification and improved business continuity. For the DR550, it is a reasonable choice to address the requirements of the archiving offering.

The DR550 SAN Switch provides 1, 2, or 4 Gbps link speed. The port speeds can be set to any of these values or can be set to auto-negotiate the highest speed that the attaching devices support. In the DR550, the DR550 SSAM servers (p52A), and DR550 Storage Controller (DS4200) ports operate at 4 Gbps. Figure 2-8 shows the DR550 SAN Switch.



Figure 2-8 DR550 SAN Switch (2005-B16)

In the dual-node configuration, the switches are configured to operate two distinct fabrics independently for better availability and redundancy. Within single-node configurations, only one switch is used.

The required Fibre Channel cabling between the DR550 SSAM Servers and the DR550 Storage Controller units is done in manufacturing. The customer does not have to perform any reconfiguration at installation time unless the Enhanced Remote Mirroring option is purchased later. In addition, the zoning definitions are done in the switch or switches at the factory, with separate zones for disks and tape. For further details about Fibre Channel cabling and zoning, refer to “DR550 Model DR2 Fibre Channel connections” on page 40 and “Zone configurations” on page 43.

Note: Although technically possible, we do not recommend that you share the DR550 fabric or fabrics with other fabrics for attaching non-DR550 components, such as external servers or storage devices. Doing so compromises the security concepts of the DR550 and could have implications on third-party compliance certification, for example.

Console kit

The DR550 Console kit is an IBM 7316-TF3 rack-mounted flat panel console kit consisting of the following:

- ▶ One 17-inch (337.9 mm x 270.3 mm) flat panel color monitor
- ▶ One rack keyboard tray
- ▶ One IBM travel keyboard (English only) with integrated mouse
- ▶ One KVM (Netbay-LCM) switch

The KVM switch is packaged as a 1U kit and is mounted in the rack along with the other DR550 components. The LCM switch is mounted in the same rack space, located behind the flat panel monitor. The IBM Travel Keyboard is configured for English. An integrated mouse is included in the keyboard. The DR550 SSAM Servers and the DR550 FSGs are connected to the LCM switch, so that the monitor and keyboard can access all of the servers in the DR550.

Alternating Current (AC) power cabling

The DR550 DR2 is configured with two Power Control Assemblies (PCAs). Each Power Control Assembly (PCA) is located at the rear of the rack in the left and right rack side pockets and receives power through the rack power cord that is plugged into the client's power source. The on/off switch is located at the front of the rack and controls each PCA. To enhance the availability of the rack components, each rail in a set can be connected to a different AC power feed when that is available.

Ethernet cabling

The DR550 DR2 has all the required internal IP connections preconfigured and wired.

The DR550 Model DR2 includes an internal Ethernet network for some connections between the DR550 SSAM Servers (p52a), the DR550 Storage controller (DS4200), and the optional DR550 FSG nodes. These connections are made through two Ethernet switches and are completely isolated from the customer network (for details, refer to "DR550 Model DR2 Single-node Ethernet connectivity" on page 34 and "DR550 Model DR2 dual-node Ethernet connectivity" on page 35).

Fibre Channel cabling

All the cabling needed for the DR550 SSAM Servers to access the storage (DR550 Storage Controllers and DR550 Expansion Drawers) is done at the factory. The host bus adapters (HBAs) in the DR550 SSAM Servers are connected to the Fibre Channel switches, and the switches in turn are connected to the DR550 Storage Controllers. The DR550 Storage Controller connects to the first DR550 Expansion Drawer, which then is daisy-chained to the following DR550 Expansion Drawers. See 2.5.5, "DR550 SAN Switch configuration" on page 39 and 2.5.6, "Storage configuration" on page 45 for details.

Management console cabling

A USB cable (carrying keyboard and mouse signals) and the graphics adapter (slot 2) cable of the DR550 SSAM Server node 1 are combined on a CAT5 conversion option cable, which is daisy-chained to the CAT5 conversion option cable of DR550 SSAM Server node 2 and from there to the central LVM switch (dual-node). In a single-node configuration, the conversion cable of node 1 is directly connected to the switch. The optional File System Gateways are also connected to the KVM switch. The flat panel monitor, keyboard, and mouse cables are also linked to that switch, thus allowing you to switch to the desired engine. The communication between the p5 nodes and the optional FSGs are shown in Figure 3-4 on page 83.

2.1.2 Software components

This section describes the software components of the DR550 *without* the optional File System Gateway (FSG).

Table 2-1 shows the software levels as installed by manufacturing on a DR550 dual-node V4.5. The various software components are described in subsequent sections.

Table 2-1 Software levels for DR550 V4.0

Components	Software levels
DS4200	Firmware 07.10.x
DS4000 Storage Manager Client	Version 10.10.xx (Special DR550 Edition)
2005-B16	Firmware 5.3.0 or later

Components	Software levels
AIX	5300-06-04-078 A tape driver 10.7.3.0 Atltd 6.5.5.0 AIX filesets bos.adt.debug 5.3.0.61 AIX filesets bos.alt_disk_install.ret 5.3.0.62
HACMP	5.4.2
IBM Director Agent - AIX	5.20.2
System Storage Archive manager (SSAM)	5.5
Tivoli Integrated Solutions Console (ISC)	6.0.1.1
Tivoli AdminCenter	5.5
SSAM API Client	5.3

Note that some early shipments of the DR550 V4.5 may include Version 6.60.xx of the DS4200 firmware. If that is the case, remember that you must use a special Firmware Upgrade Utility program to upgrade to level 7.10.xx of the firmware.

AIX and HACMP

Data retention is a critical business process. The DR550, therefore, provides layers of redundancy, minimizing single points-of-failure. The key elements driving this high availability environment are the capabilities of AIX and High Availability Cluster Multi-Processing (HACMP) with dual System p servers, redundant networks, and multiple paths to storage. However, for customers with lower requirements in terms of accessibility and availability, two single-node configurations without the enhanced protection of HACMP are available.

The DR550 SSAM Server (p52A server) or Servers in the DR550 run AIX 5L Version 5.3. AIX 5L is a high-performance, UNIX-based, multiuser, and multiprocessing operating system with a wide range of systems and network management capabilities. AIX 5L is one of the world's most open UNIX operating systems and includes functions to improve usability, security, system availability, and performance. These include improved availability of mirrored data and enhancements to AIX Workload Manager, which help solve challenges of mixed workloads by quickly and dynamically providing resource availability to critical applications.

With the DR550, AIX security has been customized for additional protection. We discuss this in further detail in 2.5.2, "Components security" on page 30.

HACMP ensures that critical resources, such as applications, are available for processing. HACMP takes measures to ensure that applications in the DR550, such as System Storage Archive Manager, remain operational even if a component in the cluster fails. In the case of a component failure, HACMP will move the application along with the resources from the active server to the standby (passive) server. In the DR550, the two cluster nodes communicate over the cluster networks. Each node sends heartbeats over this network to check on the health of the other cluster node. If the passive standby node detects no heartbeats from the active node, the active node is considered "failed", and resources are automatically transferred to the standby node, which consequently becomes the active node. When the failed node recovers, it takes over again as the primary node.

DS4000 Storage Manager for DR550

The DS4000 Storage Manager for DR550 software (hereafter referred to as *Storage Manager*) is installed in the DR550. This *special* version of Storage Manager is used to support centralized management of the DR550 Storage Controllers in the DR550.

Generally speaking, Storage Manager enables administrators to quickly configure and monitor their DR550 Storage Controller from either a command-line interface or a Java™-based graphical user interface. It is designed to enable storage administrators to customize and change settings, configure new volumes, define mappings, handle routine maintenance, and dynamically add new enclosures and capacity to existing volumes without interrupting user access to data. It is also used to configure, monitor, and maintain Enhanced Remote Mirroring. Failover drivers, performance-tuning routines, and cluster support are also standard features of Storage Manager.

Important: Before upgrading the DR550 Storage Controller or DR550 Expansion Drawer firmware, make sure it is compatible with the Storage Manager version currently installed on your DR550 SSAM Server.

To upgrade from a V6.x of the firmware, you *must* use the Storage Manager Firmware Upgrade Utility (downloadable from the DS4000 support Web site)

The special version of Storage Manager (DS4000 Storage Manager client) installed with the DR550 is for exclusive use with the DR550 and is available for download from the Web. This enhanced version provides additional protection against accidental data deletion. For this reason, we also recommend that you upgrade the Storage Manager client software only with this special version.

IBM System Storage Archive Manager

IBM System Storage Archive Manager (SSAM) is designed to provide archive services and to prevent the loss of critical data as well as to prevent critical data from being erased or overwritten. IBM System Storage Archive Manager is available for more than 30 operating platforms, covering mobile, desktop, and server systems over the entire distributed world. IBM System Storage Archive Manager supports many different storage devices. IBM System Storage Archive Manager is used to provide and manage retention (archiving) of data. It is not meant to be a backup offering.

For applications that use the IBM Tivoli Storage Manager API, policy-based data management capabilities are already available. Now, it will also prevent data deletion before the retention criteria is satisfied. Content management and archive applications (such as DB2 Records Manager integrated with DB2 Content Manager, as detailed in Chapter 8, “Using IBM System Storage DR550” on page 303 or such as the new optional DR550 FSG) can use the IBM System Storage Archive Manager client API to apply business policy management for ultimate deletion of archived data at the appropriate time.

Tip: IBM System Storage Manager Archive Manager (SSAM) is the same software as the IBM Tivoli Storage Manager, but with the archiveretentionprotection attribute set to on.

IBM System Storage Archive Manager includes the following functionality and device support:

- Data retention protection

Designed to prevent deliberate or accidental deletion of data until the retention criterion is met.

- ▶ Event-based retention policy

This addresses the case where the retention period must be based on the occurrence of a certain event. In other words, the retention time for data is based on a specific event and not on the initial storage of that data. For example, to delete a customer record when the customer account is closed.

- ▶ Expiration or deletion suspension (deletion hold)

Designed to prevent deletion of data when the regulatory retention period has passed, but other requirements (such as a court order) mandate that the data continues to be maintained. The data cannot be deleted until the deletion hold is released.

- ▶ Data Encryption

128-bit Advanced Encryption Standard (AES) is available with the archive API client. Data can be encrypted before it is transmitted to the DR550, and it will then be stored on the disk or tape in an encrypted format.

- ▶ Support for WORM tape and DVD-ROM.

- ▶ The Integrated Service Console (ISC) and the Tivoli Storage Manager AdminCenter are installed on the system. You must configure these products before you can use them.

- ▶ SSAM V5.5 supports hardware tape drive encryption, which can be used with IBM TS1120-E05 or IBM TS1040 LTO4 tape drives.

- ▶ Also with SSAM V5.5, a data shredding function was made available to guarantee a safe wipeout of the data on the media.

For more information about the System Storage Archive Manager, refer to Chapter 5, “IBM System Storage Archive Manager” on page 157.

2.2 DR550 Model DR1 overview (2233-DR1)

The IBM System Storage DR550 Model DR1 has many of the features and benefits of the DR550 Model DR2, but at a much lower cost, which makes it ideal for small to medium businesses. It is packaged in a 25U (ETA) tall rack, which is 49 inches tall. Figure 2-9 on page 22 shows the DR550 Model DR1.

To help clients better respond to changing business environments as they transform their infrastructure, the DR550 Model DR1 ships with 0.88 TB of physical capacity using the internal SCSI disk of a DR550 SSAM Server (six non-removable 146 GB Ultra™ SCSI 3 disk drives in this System p52a server). The first two disk bays are equipped with 73 GB disks and are reserved for operating system use. The DR550 Model DR1 is tape ready. It also includes standard dual power supplies

Restriction: The DR550 Model DR1 does not include SAN switches.

Storage can be expanded, in 6 TB (half drawer) increments, to up to 36 TB using a DR550 Storage Controller populated with two DR550 Expansion Drawers (EXP420) populated with 750 GB drives.

Important: Remember that if you order the optional DS4200 with your DR550-DR1, the ERM feature is not supported with that configuration. If you need ERM, you must order a DR2 single or dual node.



Figure 2-9 DR550 Model DR1(2233-DR1)

The server shipped in the DR550 Model DR1 has two 73 GB disk drives in RAID 1 (mirrored) configuration in two of the bays for use by the operating system.

The software bundle includes the IBM AIX 5L Version 5.3 operating system, IBM System Storage Archive Manager Version 5.5, and the DS4000 Storage Manager client, customized for additional protection, all running on the DR550 SSAM Server. Document management applications used with the DR550 Model DR1 must be capable of communicating with the IBM System Storage Archive Manager Client API to archive data on the DR550 Model DR1.

In addition, with the DR550 Model DR1, configurations can also include the optional DR550 File System Gateway (2229-FSG). The FSG is available as a stand-alone or a high availability 2-node cluster option. This optional gateway provides the interface for NFS and CIFS protocol support. An FSG node uses a dual core Intel® based processor and includes 2 GB of memory (see 2.3, “DR550 File System Gateway overview” on page 25).

2.2.1 Hardware components

The hardware components for the DR550 Model DR1 consist of the following elements.

DR550 DR1 Rack IBM

The 1.3 meter-high Model S25 rack (IBM 7014-S25 rack) satisfies many light-duty requirements for organizing smaller rack-mounted servers. Front and rear rack doors include locks and keys, helping to keep your DR550 Model DR1 secure. This 25U rack can accept server and expansion units up to 28-inches deep. There is only one PCA in the rack located

at the rear of the rack at EIA position 1 and it receives power through the rack power cord that is plugged into the client's power source.

The front door is reversible so that it can be configured for either a left or right opening. The rear door is split vertically in the middle and hinges on both the left and right sides. The DR550 DR1 rack has the following specifications:

- ▶ Width: 605 mm (23.8 in.) with side panels
- ▶ Depth: 1001 mm (39.4 in.) with front door
- ▶ Height: 1344 mm (49.0 in.)
- ▶ Weight: 100.2 kg (221.0 lb.)

The DR550 DR1 rack stacks all the components vertically. The servers and the optional storage units start from the bottom, populating toward the top as the storage capacity installed increases (see Figure 2-10).

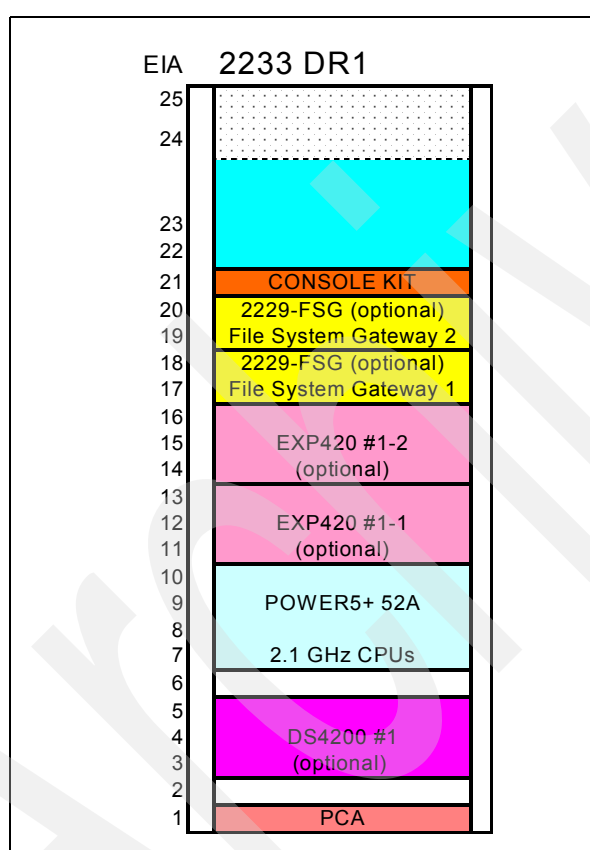


Figure 2-10 Rack population for DR550 Model DR1

DR550 SSAM Server

The DR550 Model DR1 includes one DR550 SSAM Server (IBM System p p5 52A (9131-52A)) running the IBM AIX 5L Version 5.3 operating system. The DR550 SSAM Server (referred to hereafter as the p5 52A, p52A, or engine) is a cost-effective, high performance, and space-efficient server that uses advanced IBM technology. The p52A uses the 64-bit, copper-based POWER5 microprocessor technology and is designed for use in LAN clustered environments.

The p5 52A is a member of the symmetric multiprocessor (SMP) UNIX servers from IBM. The p5 52A server (product number 9131-52A) is a 4-EIA (4U), 19-inch rack-mounted server. The p5 52A is configured with a dual core 2.1 GHz processor.

The total system memory installed is 2 GB. The p5 52A includes six hot-plug PCI-X slots, an integrated dual channel Ultra320 SCSI controller, two 10/100/1000 Mbps integrated Ethernet controllers, and eight front-accessible disk bays supporting hot-swappable disks. The server shipped in the DR550 Model DR1 has two 73 GB disk drives in RAID 1 (mirrored) configuration in two of the bays for use by the operating system. The remaining six bays are populated with 146 GB UltraSCSI 3 disk drives

In addition, there are also three media bays available:

- ▶ Media - dev0 - not used in the DR550
- ▶ Media - dev1 - Slimline DVD-RAM (FC 1993)
- ▶ SCSI tape drive (not included)

It also now includes standard dual power supplies.

Figure 2-11 shows the DR550 SSAM Server.



Figure 2-11 The DR550 SSAM Server used in the DR550 Model DR1

Console kit

DR550 Model DR1 base order comes with a tray and display for the console. The KVM switch and keyboard are a separate feature. The IBM 7316-TF3 is a rack-mounted flat panel console kit consisting of a 337.9 mm x 270.3 mm (17 in.) flat panel color monitor, a rack keyboard tray, an IBM travel keyboard (English only), and the KVM Switch. This is packaged as a 1U kit and can be mounted in the rack along with the other DR550 Model DR1 components. The KVM Switch is mounted in the same rack space, located behind the flat panel monitor. The IBM Travel Keyboard is configured for English. An integrated mouse is included in the keyboard. The POWER5 52A server is connected to the KVM Switch so that the monitor and keyboard can access the server.

DR550 Storage Controller and DR550 Expansion Drawer

You can upgrade the storage by adding a DR550 Storage Controller (DS4200, with dual controllers) if you want to expand the storage of the DR550 Model DR1. This is done by adding a half-populated DR550 Storage Controller (DS4200) consisting of eight 750 GB SATA drives for 6 TB or a fully-populated DR550 Storage Controller consisting of sixteen 750 GB SATA drives for 12 TB. More storage can be added by installing DR550 Expansion Drawers (EXP420s) in half-populated (6 TB) or fully-populated (12 TB) increments. For the DR550 Model DR1, the maximum number of DR550 Expansion Drawers is two, resulting in a maximum possible storage capacity of 36 TB.

The storage used in the DR550 comes in several capacities. The minimum capacity in the DR550 model DR1 is 0.88 TB of raw disk space that is built with six 146 GB drives in to the DR1 System p server. The minimum capacity in the DR550 model DR1 with DS4200 controller enclosure is 6 TB of raw disk space that is built with eight (half a drawer) 750 GB disks in the DS4200 controller enclosure.

The RAID type is RAID 5, whereby a half drawer, if it is the only storage drawer, is formatted 2+p and 3+p with a global spare, otherwise it is formatted 6+p with one global spare, and a full drawer is formatted 6+p and 7+p with one global spare.

See “DR550 Storage Controller (DS4200) and DR550 Expansion Drawer (EXP420)” on page 13 for more information about DR550 Storage Controller.

2.2.2 Software components

This section describes the software components of the DR550 Model DR1.

AIX

The DR550 SSAM Server (p5 52A server) in the DR550 Model DR1 runs AIX 5L Version 5.3-06-04. AIX 5L is a high-performance, UNIX-based, multiuser, and multiprocessing operating system with a wide range of systems and network management capabilities. AIX 5L is one of the world's most open UNIX operating systems and includes functions to improve usability, security, system availability, and performance. These include improved availability of mirrored data and enhancements to AIX Workload Manager, which helps solve problems of mixed workloads by quickly and dynamically providing resource availability to critical applications.

With the DR550 Model DR1, AIX security has been customized for additional protection. We discuss this in further detail in 2.5.2, “Components security” on page 30.

DS4000 Storage Manager for DR550

The DS4000 Storage Manager for DR550 Client software is installed in the DR550 Model DR1. This software will only be used if the DR550 Storage Controller is attached.

See “DS4000 Storage Manager for DR550” on page 20 for more information about DS4200 Storage Manager for DR550.

IBM System Storage Archive Manager

See “IBM System Storage Archive Manager” on page 20 for more information about IBM System Storage Archive Manager.

2.3 DR550 File System Gateway overview

The IBM System Storage DR550 File System Gateway (FSG) provides an NFS and CIFS file system interface to applications. The FSG provides the mount point for the NFS/CIFS oriented applications and then transfers the data to the SSAM application within the DR550 (using the SSAM API). This enables the DR550 to also support archiving applications that rely on CIFS or NFS protocols such as the Picture Archiving Communication Solutions (PACS) used by the healthcare industry.

Important: Before ordering the File System Gateway, make sure that your application can properly work using file shares provided by the FSG. This is because the FSG modifies the file share behavior by disallowing deletions or modification to any file stored on the FSG file share.

Optional for *new* DR550 orders, the File System Gateway comes in stand-alone or clustered high availability configurations.

For DR550 Model DR1 and single-node Model DR2, the File System Gateway is offered in either a stand-alone or a clustered high availability configuration. For dual-node DR550 Model DR2, File System Gateway is offered in a clustered high availability configuration only.

2.3.1 Hardware components

The DR550 File System Gateway consists of one or two IBM 2229-FSG servers. The 2229-FSG is based on an IBM System x 3650 that comes preloaded with Novell®'s SUSE® Linux® Enterprise System 10 and the FSG application (see 2.3.2, "Software components" on page 27 for details on the preloaded software components).

The File System Gateway (FSG) offers excellent performance using a dual core Intel based processor.

Figure 2-12 shows the File System Gateway.



Figure 2-12 IBM DR550 File System Gateway (2229-FSG)

It is preconfigured with the following hardware:

- ▶ Dense, 2U, rack-mounted packaging.
- ▶ Single dual core 3.2 GHz Intel Xeon® processor.
- ▶ 2 GB of fast DDR2 fully buffered memory.
- ▶ Two PCI-X Ethernet adapters (your choice of SX or TX connections) located in slots 1 and 2 (slots 3 and 4 are empty). These adapters are used for connection to the customer LAN.
- ▶ Two (optional) integrated Ethernet ports used for cluster heartbeat if there is a second File System Gateway (clustered).
- ▶ Two 73 GB SAS disk drives (3.5 inch).
- ▶ Four 300 GB SAS disk drives (3.5 inch).
- ▶ RAID 6 implementation.
- ▶ Redundant power supplies.
- ▶ DVD-CD/RW drive.

The FSG features a Lightpath diagnostics panel that provides information about a failing component without interrupting system operation and expedites hardware repairs to dramatically reduce service time.




Figure 2-13 FSG-2229 rear view

2.3.2 Software components

The DR550 File System Gateway comes preloaded with SUSE Linux Enterprise Server (SLES) 10 and the DR550 File System Gateway software.

The DR550 File System Gateway software enables client applications to store and retrieve data from the DR550 by means of the Common Internet File System (CIFS) or Network File System (NFS) file system protocols. The FSG interacts with the other DR550 components through the System Storage Archive Manager API.

As files are sent to the File System Gateway, they are cached locally and forwarded to the archive within the DR550 for persistent storage. The FSG software maintains the file system directory tree locally. The actual data is stored within the DR550 repository (SSAM archive pools), making the file system appear to have a very high capacity, even though the DR550 File System Gateway server has relatively modest local resources. When a client application requests an object that must be retrieved from the DR550, the FSG software requests the object from the DR550. The object is retrieved from the DR550 storage pool, sent back to the FSG, which stores it in its local disk storage, and finally sends the object to the requesting application.

Configuration of the File System Gateway is achieved by the editing of various XML files (refer to Chapter 6, “IBM DR550 File System Gateway” on page 229 for details).

The IBM Director Agent is preinstalled and can be used to monitor the FSG (refer to 10.7, “Configure SNMP monitoring for FSG” on page 440).

Table 2-2 summarizes the software levels on the FSG nodes.

Table 2-2 Software levels for FSG nodes in DR550 V4.5

Components	Software levels
File System Gateway BIOS	1.08
File System Gateway Application	V1.1.1
SLES 10	kernel 2.6.16.21-0.8-smp
IBM Director Agent - Linux	5.20.2
SSAM API Client	5.3

2.4 Positioning the offerings

Because the IBM DR550 Models DR1 and DR2 are made of off-the-shelf hardware and software components that can be used in many contexts, it is important to stress the intended use of the DR550.

What it is

For enterprises, data retention and managing retention with well-defined policies are becoming critical parts of overall data management. More important, compliance with data retention stipulations of various regulatory bodies is a new challenge that needs to be addressed.

Key areas of concern in managing such data for regulatory compliance are:

- ▶ Assurance that data content is *kept* and not altered or deleted for a specified amount of time.
- ▶ Content is an asset when kept for the required amount of time, but it can become a liability when kept beyond that time.
- ▶ Content retention can vary during its retention life due to events such as litigation.

All these factors mandate that a well-managed and coordinated retention life cycle is in place. The DR550 is the IBM end-to-end offering to address all of these data retention and content management concerns.

The DR550 is an offering designed and built primarily to address the issue of compliance with government and industry regulations, in particular, the SEC and NASD regulations as they pertain to retaining critical data. It is also an ideal offering for several business types, such as banking and insurance companies, where data retention and implementing and managing data retention with the appropriate retention and expiration policies are critical business needs. The main focus of this offering is to provide a secure storage system where deletion or modification of data is completely disallowed except under a well-defined retention and expiration policy.

What it is not

The DR550 is *not* meant to and cannot be used as a standard backup and restore system, even though many of the components of the offering (IBM Tivoli Storage Manager in particular) have been deployed as general purpose backup and restore solutions.

The DR550 Storage Controller (DS4200), by itself, is a standard external enterprise storage system with a wide range of applications such as workgroup and departmental storage, including on demand storage access. It is a highly scalable storage system that supports multiple host systems. However, the implementation used in this offering uses the SATA technology, which, unlike the Fibre Channel technology, is not well-suited for many enterprise applications that are mission-critical. The design of the offering also does not allow adding multipatform hosts. The only access to the storage is through the two DR550 SSAM Servers (p5 52A) provided with the DR550 and using the Tivoli Storage Manager API or using the optional DR550 File System Gateway (FSG) which provides NFS/CIFS access to the DR550.

2.5 What is preconfigured in the DR550

The DR550 Models DR1 and DR2 as shipped come almost completely cabled and configured, and there is very little in terms of installation and configuration that you have to do before deploying this system in your environment. In the sections that follow, we explain what is already configured for the elements of the DR550.

2.5.1 DR550 SSAM Server

DR550 SSAM Server is an IBM System p5 52A (9131052A). Each DS550 SSAM Server is configured as follows:

- ▶ **Media bays**

Three media bays are available. Only one is used and has one Fibre Channel (FC) 1993 slim-line DVD-RAM device installed.

- ▶ **PCI-X slots**

The p52A has six hot-plug PCI-X slots:

- Two long 64-bit 133 MHz (slots 5 and 6).
- One short 64-bit 266 MHz (slot 4).
- Two short 32-bit 66 MHz (slots 2 and 3).
- One short 64-bit 133 MHz (slot 1).

In the DR550, the PCI-X slots are populated as follows:

- Slot 1: Dual FC 4 Gbps HBA: Port 1 for connection to the internal SAN for disk attachment, and port 2 for connection to the internal SAN for external tape attachment.
- Slot 2: One POWER™ GXT135P Graphics Adapter, connected to the console kit.
- Slot 3: LVD SCSI: FC3560 (optional) or second external host TX connections (mixed media case option).
- Slot 4: Dual FC 4 Gbps HBA: Port 1 for connection to the internal SAN for disk attachment, and port 2 for connection to the internal SAN for external tape attachment.
- Slot 5: Dual 1 Gb Ethernet (TX) for connections to the internal Ethernet network.
- Slot 6: Dual 1 Gb FC 3550 (SX) or FC3551 (TX) (mutually exclusive options), for connections to the customer network.

- ▶ **I/O ports (on system board)**

- Two 10/100/1000 Ethernet ports (T-5 and T-6).
- Two serial ports (RS232).
- Two USB ports. The upper one (T7) is used to connect to the keyboard and mouse.
- Two RIO ports, which are not used by DR550.
- Two Hardware Management Console (HMC) ports, which are not used by the DR550.
- Two SPCN ports, which are not used by DR550.

Figure 2-14 shows the DR550 SSAM Server rear view with the I/O adapters and PCI slot locations.




Figure 2-14 DR550 SSAM Server (p5 52A) rear view with I/O adapters and PCI slot locations

The detailed network and Fibre Channel configuration for the adapters installed on these servers are described in 2.5.3, “Network configuration” on page 33 and 2.5.5, “DR550 SAN Switch configuration” on page 39.

2.5.2 Components security

The DR550 SSAM Servers (p5 52A) are shipped with a preinstalled, standard AIX 5L Version 5.3 operating system. However, there are certain modifications to the AIX security environment to make the DR550 more secure, mainly for compliance with data retention regulations. For example, these modifications will not allow remote administration tasks initiated through commands, such as **telnet**, and File Transfer Protocol (**ftp**). The only exception is Secure Shell (**ssh**). We recommend access for management activities through the integrated console.

Login and user accounts

To provide a greater level of security, the DR550 Models DR1 and DR2 are configured for limited access. The following restrictions are built into the DR550 Models DR1 and DR2:

- ▶ Limited user definitions
- ▶ Limited access to commands from certain accounts
- ▶ No remote access with authority to make changes (except for System Storage Archive Manager administration)

The following user accounts have been created with specific roles and restrictions. All passwords should be changed in accordance with company policies and guidelines. To enhance security, certain user accounts do not have any change authority, and other accounts can only be accessed from the integrated console:

- ▶ AIX users dr550adm and ibmce must log in with Secure Shell (ssh).
- ▶ User dr550 can only access the engines locally at the management console and is the only user that can switch to the AIX “root” account (the root account has full system privileges). The purpose of this design is to prevent unauthorized persons without physical access to the system to log in to the DR550 Models DR1 and DR2.
- ▶ FSG accounts are discussed in 2.6.3, “FSG security” on page 57.
- ▶ The System Storage Archive Manager administration account admin can be used locally or remotely. Locally applied, it is running in the security context of the AIX login.
- ▶ The System Storage Archive Manager account hacmpadm (DR550 dual-node only, not included in DR550 Model DR1 or Model DR2 single-node), as the name suggests, is required for HACMP to perform internal cluster services tasks related to the System Storage Archive Manager server application and is not meant for regular use.
- ▶ The DR550 SAN Switch (IBM SAN Switch B16) can, for administration or monitoring purposes, be temporarily connected and accessed through its Ethernet port. Log in with the user admin.

Table 2-3 through Table 2-6 on page 32 outline the registered accounts on the DR550 and their roles and restrictions.

Table 2-3 AIX user IDs

AIX user ID	Role and characteristics	Management Console access	Network access (SSH)	su root ability
dr550	For System Storage Archive Manager administration tasks Only user who can su to root Home directory /home/dr550	Yes	No	Yes
root	For AIX and HACMP administration tasks Has all system privileges Home directory /	Not directly, only through dr550 user	No	N/A
dr550adm	For viewing log files and performing Tivoli Storage Manager admin client tasks Home directory /home/dr550adm	Yes	Yes	No
ibmce	For viewing log files and performing Tivoli Storage Manager client tasks Home directory /home/ibmce	Yes	Yes	No
esaadmin	Electronic Service Agent™ for AIX, administration user ID	Yes	Yes	No

Table 2-4 Default user accounts on the FSG

User account	Password
root	dR550fsg
fsgadm	C1fsNFS
drg	drg

User fsgadm has to be used locally at the console as well; after you are logged on as fsgadm, you can in turn switch to user 'root'.

Table 2-5 Tivoli Storage Manager user IDs⁵ 52A Flexible Service Processor (FSP)

Tivoli Storage Manager user ID	Roles and characteristics
admin	For System Storage Archive Manager administrative tasks
hacmpadm	Used by HACMP scripts on the DR550 dual-node only

Table 2-6 SAN Switch 2005-B16 user ID

2005-B16 user ID	Role and characteristic
admin	Privileged user for 2005-B16 administrative tasks

Important: The System Storage Archive Manager account hacmpadm is only used in case of a node failure by HACMP to perform a shutdown of the System Storage Archive Manager server. The default password for hacmpadm is set to expire after a 5.5 year period (2000 days). We recommend that you change the default password according to the company policies and guidelines. Each time the password for hacmpadm is changed, you are required to update the script /usr/bin/stopserver to reflect the new password. If this is not done, HACMP cannot stop the System Storage Archive Manager server and a successful failover cannot be accomplished.

It is the customer's responsibility to change the default passwords and to record and protect the new passwords. Changing the passwords should be done during the initial installation. The password for root (on AIX) must be changed during initial installation.

Other security restrictions include:

- ▶ Server processes or daemons to remotely access the engines with facilities, such as **telnet**, **ftp**, **rsh**, and **rlogin**, are either removed, disabled, or will be rejected; only **ssh** is authorized and only for users dr550adm and ibmce.
- ▶ Insecure services, such as **telnet** and **ftp**, have been disabled so that they cannot be executed. The reason these are disabled is to effectively disallow attempts to access any remote machine *from* the servers with **telnet** or **ftp**. In contrast, **rlogin**, **rsh**, and **sftp** are functional, because they are secure services.
- ▶ Ports for **telnet**, **ftp**, **echo**, **rsh**, **timed**, and **chargen** have been blocked out.
- ▶ Services, such as spooler and sendmail, have been removed.

2.5.3 Network configuration

The DR550 Models DR1 and DR2 and the optional File System Gateway (FSG) has TCP/IP addresses, net masks, and other network parameters preconfigured. Depending on the user's network environment, you can attach the DR550 Models DR1 or DR2 or FSG to a 10 Mbps/100 Mbps or Gigabit copper-based or fiber optics-based Ethernet network. The main difference between the connection to a copper-based and a fiber optics-based network is the p5 and the System x Ethernet adapter and cables you use with these networks.

DR550 Model DR1 Ethernet connectivity

The DR550 Model DR1 p52A server connects to the existing network using either 10/100/1000 copper connections or a Gigabit fiber connection. Only one Ethernet connection is needed. It is the lower port (*en0*) of the Gigabit Ethernet adapter in slot 6. It is preconfigured with IP address 192.168.6.11. This address has to be changed during installation to match the customer's network environment. The customer supplies both Ethernet switches and cables. If the fiber optic Ethernet option is ordered, a Gigabit Ethernet-SX PCI-X fibre optic adapter is used.

Figure 2-15 shows a DR550 Model DR1 copper-based Ethernet connection.




Figure 2-15 DR550 Model DR1 copper-based Ethernet connection

If you order the optional DR550 Storage Controller (DS4200) with DR550 Model DR1, the two onboard Ethernet ports of p52A server (T5 and T6) are used to establish the crossover Ethernet connection with the controllers A and B of DR550 Storage Controller (DS4200). The upper port T5 of onboard adapter is directly connected to the first Ethernet port of the DS4200 Controller A. The lower port T6 of onboard adapter is directly connected to the first Ethernet port of the DS4200 Controller B.

Figure 2-16 shows a DR550 Model DR1 with DR550 Storage Controller (DS4200)




Figure 2-16 DR550 Model DR1 with DR550 Storage Controller (DS4200)

DR550 Model DR2 Single-node Ethernet connectivity

The DR550 Model DR2 p52A server is connected to the existing network using either 10/100/1000 copper connections or a Gigabit fiber connection. Only one Ethernet connection is needed. It is the first (lower) port (*en0*) of the Gigabit Ethernet adapter in slot 6. It is preconfigured with IP address 192.168.6.11. This address has to be changed during installation to match the customer's network environment. The customer supplies both external Ethernet switches and cables.

The DR550 Model DR2 single-node comes with two internal Ethernet switches. They are preinstalled and preconfigured. Figure 2-17 on page 35 shows the Ethernet connections that include two DR550 Storage Controllers (but not the optional FSG). For a single DR550 Storage Controller scenario, ports 8 and 9 of both internal switches are not used.

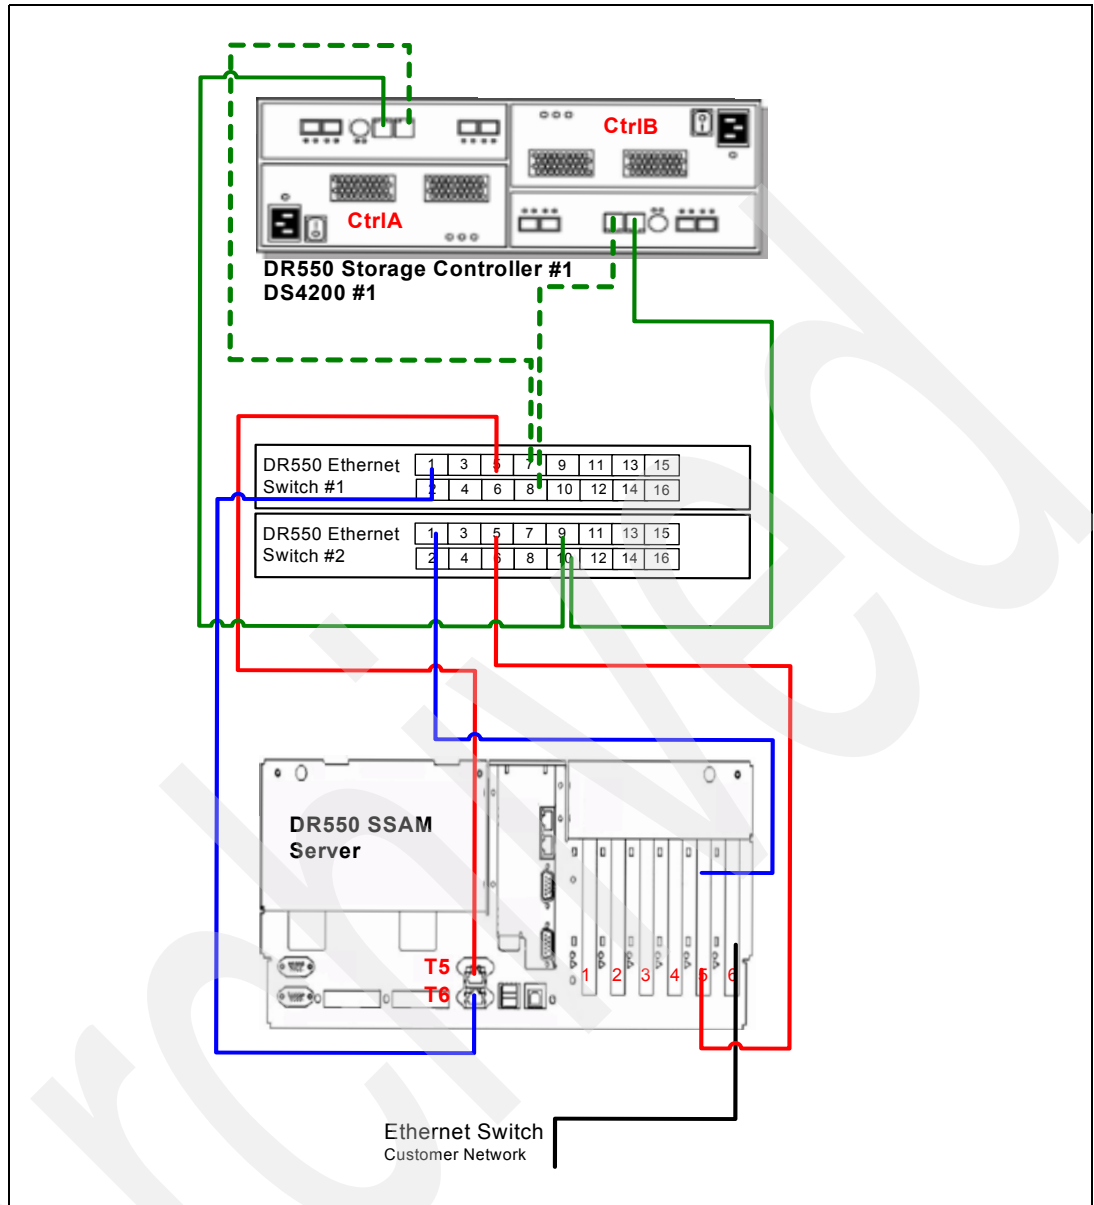


Figure 2-17 Ethernet connections of a single-node DR550 Model DR2 (no FSG)

Do not change the internal IP addresses. What specific port is being used for connecting to the internal Ethernet switch is not important.

DR550 Model DR2 dual-node Ethernet connectivity

Figure 2-18 on page 36 and Figure 2-19 on page 37 shows the Ethernet cabling for a dual-node configuration. It is expected that the customer will connect the two p52A servers to their existing network using either 10/100/1000 copper connections or a Gigabit fiber connection. Only one Ethernet connection from each p52A server is needed. It is the first (lower) port (*en0*) of the Gigabit Ethernet adapter in slot 6. The factory-set IP addresses have to be changed during installation to match the customer's network environment. The customer supplies both external Ethernet switches and cables. If the fiber optic Ethernet option is ordered, a Gigabit Ethernet-SX PCI-X fibre optic adapter is used instead of the onboard copper-based Gigabit Ethernet port.

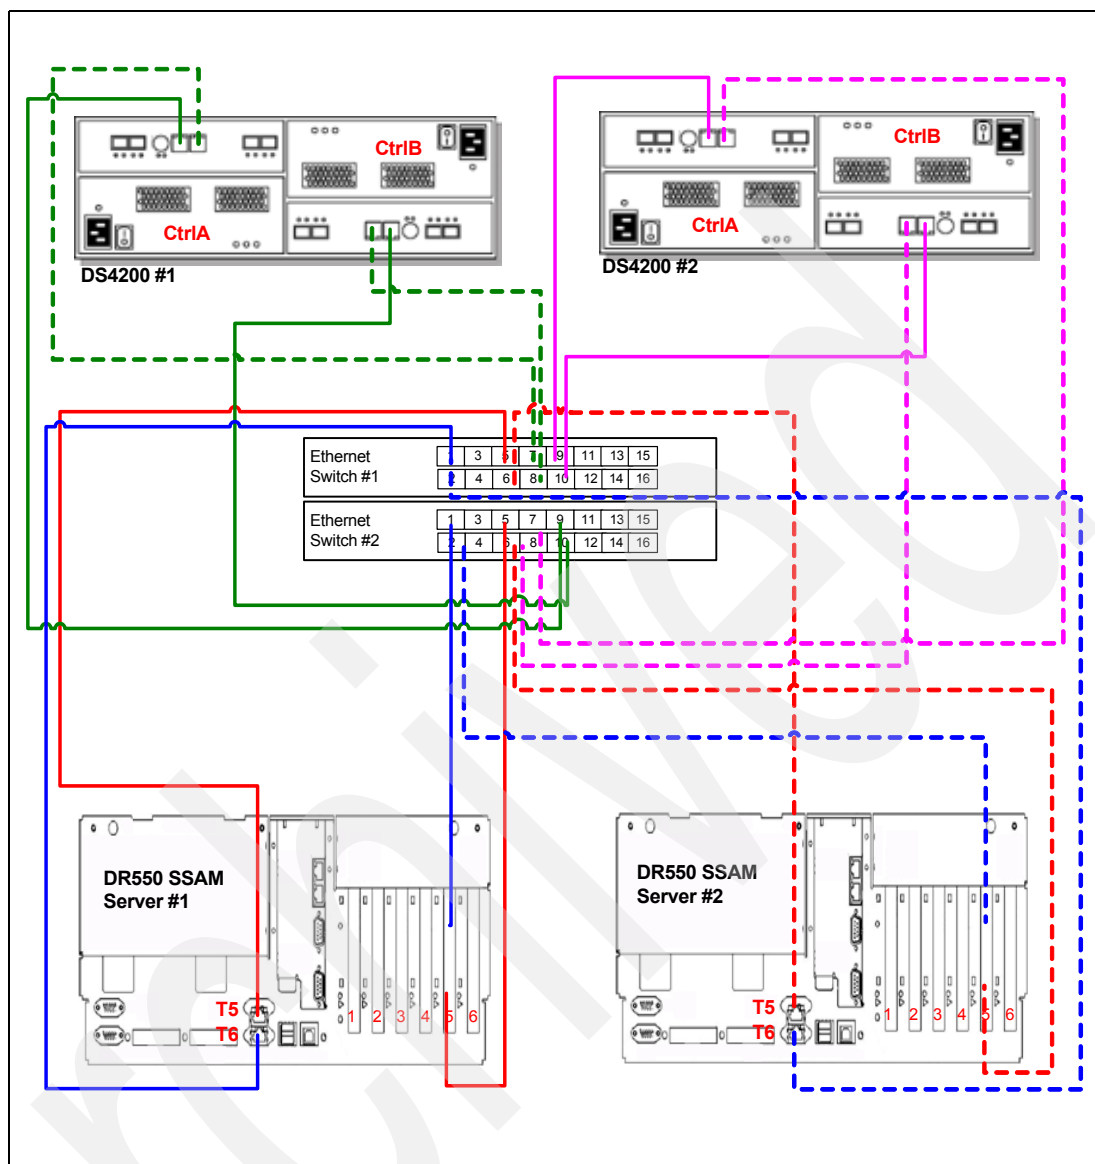


Figure 2-18 DR2 dual node internal Ethernet Network (no FSG)

The DR550 Model DR2 dual-node comes with two internal Ethernet switches. They are preinstalled and preconfigured. Figure 2-19 on page 37 shows the Ethernet connections that include two DS4200 Storage Servers. For a single DS4200 Storage Server scenario, ports 8 and 9 of both internal switches are not used, as there will not be DR550 Storage Controller #2.

Important: You should not connect any other devices or adapters to the internal Ethernet switches. The only exception is for connecting the optional RSM for DR550 server.

Internal IP addresses are preassigned and there is no need to change them.

For the additional Ethernet connections used by the optional 2229-FSG, refer to 2.6.4, “FSG Network configuration and preset IP addresses” on page 57.

Summary: Before you connect the DR550 to the LAN, make sure that the boot and standby addresses have been changed to addresses available and suitable for that network.

Refer to 3.6.5, “Configuring the IP address and HACMP for a dual-node DR550 DR2” on page 96 for detailed instructions about setting the boot and standby IP addresses.

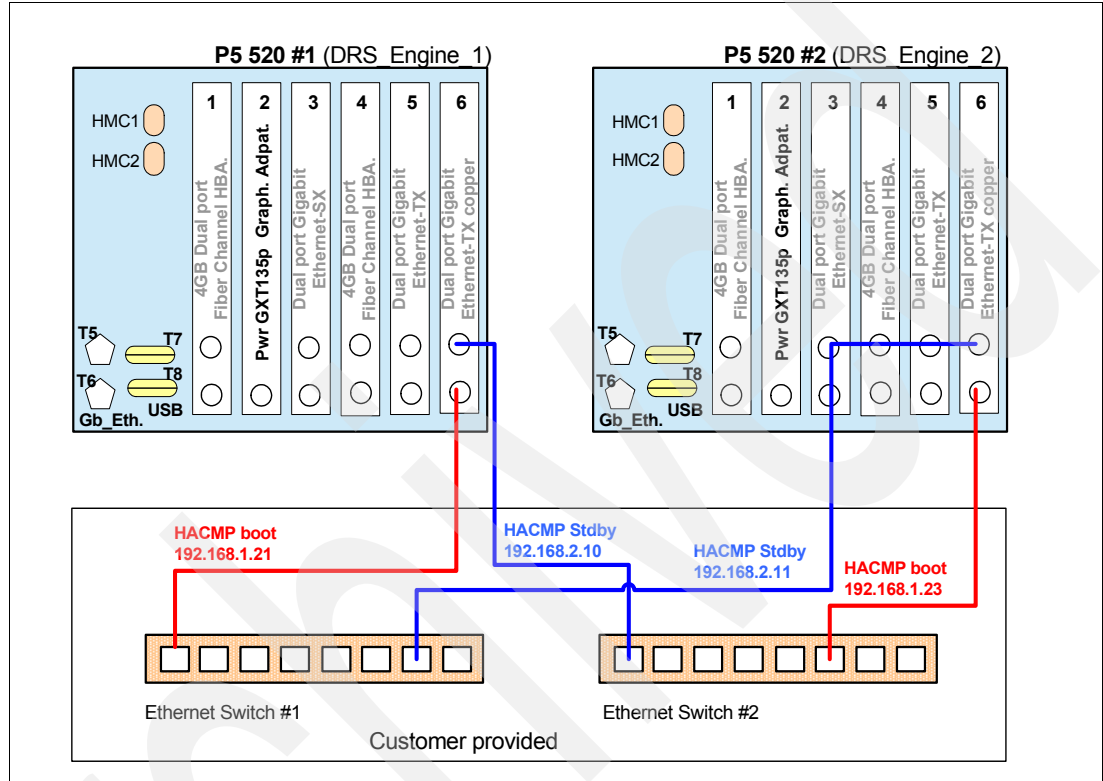


Figure 2-19 Ethernet connections of a dual-node customer network attachment

Host names and TCP/IP addresses

The DR550 must be connected to the customer’s IP network using the Ethernet adapter in slot 6 of the DR550 SSAM Servers (p52A servers). The preconfigured network addresses, along with the symbolic host names and the devices to which they connect, are shown Table 2-7 for a single-node configuration and Table 2-8 on page 38 for a dual-node setup. These IP addresses must be changed upon connection to customer network.

Table 2-7 Preconfigured TCP/IP addresses in a single-node configuration

IP address	Host name	Device	Configuration change
192.168.6.11	drs_engine	en2	Yes

Table 2-8 Preconfigured TCP/IP addresses in a dual-node configuration

IP address	Host name	Device	Configuration change
engine_1			
192.168.1.21	drs_engine1_boot	en5	No
192.168.6.21	drs_engine1_ext_boot	en2	Yes
192.168.2.10	drs_engine1_stdbby	en0	No
192.168.7.10	drs_engine1_ext_stdbby	en3	Yes
192.168.6.22	drs_cluster_ext_svc	Dynamically assigned	Yes
192.168.1.22	drs_cluster_svc	en5	No
192.168.4.24	drs_engine1_FAStT1_ctrlA	en4	No
192.168.5.26	drs_engine1_FAStT1_ctrlB	en1	No
engine_2			
192.168.1.23	drs_engine2_boot	en5	No
192.168.6.23	drs_engine2_ext_boot	en2	Yes
192.168.2.11	drs_engine2_stdbby	en0	No
192.168.7.11	drs_engine2_ext_stdbby	en3	Yes
192.168.4.23	drs_engine2_FAStT1_ctrlA	en4	No
192.168.5.25	drs_engine2_FAStT1_ctrlB	en1	No

Do not change the factory set IP addresses of DR550 Storage Controller (DS4200) or the SAN switches.

2.5.4 HACMP configuration (DR550 Model DR2 dual-node only)

For the HACMP configuration, you need to understand the network topology and which of the adapters and interfaces are being used for the boot address, standby, and service (or cluster) IP address. These are, for the most part, already set in the DR550.

The host names for the two servers are *drs_engine1* and *drs_engine2*, and these map to the boot addresses; they are also aliased as *drs_engine1_ext_boot* and *drs_engine2_ext_boot*, respectively. These are described in 2.5.3, “Network configuration” on page 33. The cluster address (the default is 192.168.6.22 and it needs to be changed during installation) maps to a host name of *drs_cluster_ext_svc*. The cluster or service address is the one that is used by any application to access the server. HACMP ensures that this address is available as long as one of the cluster nodes is available. In fact, HACMP dynamically applies this service address onto the active Ethernet port of the active node. For example, when HACMP assigns the service address to port 1 of the Gigabit adapter in node 1, the cluster can be reached through that address, and the original address 192.168.6.21 becomes ineffective for as long as HACMP does not failover to another port or another node.

HACMP information is exchanged between the two nodes using the common disk assigned to both of these nodes. During cluster operation, this is mainly the heartbeat that is sent between the nodes, allowing the standby node to detect a potential communication loss to the active server resulting in a node failover. Other information is exchanged during configuration of the cluster, such as discovering another node and getting its resource information.

Important: Remember that with Version 4.5 of the DR550 DR2 dual-node, there is also an internal HCAMP network and cluster service address defined. Do not change these internal network settings. Refer to Figure 3-8 on page 97 for details.

2.5.5 DR550 SAN Switch configuration

The DR550 Model DR1 does not include any SAN switches. The DR550 Model DR2 includes one (single-node) or two (dual-node) DR550 SAN Switches (IBM 2005 model B16 Fibre Channel Switches). Each of the switches is cabled with all the required Fibre Channel connections and configured for operations. No additional customer configuration is required unless purchased with the Enhanced Remote Mirroring option (See 13.7.2, “Merging the primary and secondary SAN fabrics” on page 542).

There are no cables running from the switch’s Ethernet ports. However, the ports come configured with preset IP addresses. They are 192.168.1.31 and 192.168.1.32 for switches 1 and 2, respectively. The IP address of the B16 included with the DR550 Model DR1 is also 192.168.1.31 by default. If there is a need to access the DR550 SAN Switches, these addresses can be changed to match the subnet to which they will be connected, using an ASCII terminal connected to the serial port. You can also use a Web browser on a notebook to do the configuration. Although under normal circumstances there is no need to access the DR550 SAN Switches, there can be instances where this might become necessary. This would, for example, be the case when a service representative needs to do problem determination and needs to gather information from the switch.

To access the DR550 SAN Switch with an ASCII terminal or a workstation running a terminal emulation program, such as Hyperterm, connect a serial cable to the 9-pin serial port of the switch. The terminal or emulation needs to be configured to 9600 bps, 8N1, and a flow control of “none”.

Figure 2-20 shows the rear view of DR550 SAN Switches.




Figure 2-20 Rear view of DR550 SAN Switch (IBM SAN Switch B16)

For further details, see *IBM System Storage SAN16B-2 Installation, Service and User's Guide*, GC26-7753, available at:

http://www-1.ibm.com/support/docview.wss?rs=1135&context=STCUSC3&dc=DA400&uid=ssg1S7001367&loc=en_US&cs=utf-8&lang=en

DR550 Model DR2 Fibre Channel connections

The first two ports of the first quad of SFPs (ports 0-1) of the switch are used for DR550 SSAM Server (p52A server) connectivity.

- For DR550 Model DR2 single node, these ports are connected to the upper ports of PCI slots 1 and 4 of DR550 SSAM Server.
- For DR550 Model DR2 dual nodes, the same ports of the second switch are connected to the upper ports of PCI slots 4 and 5 of the second DR550 SSAM Server.

The remaining two ports of the first quad (ports 2 and 3) and the first two ports of the third quad (ports 8 and 9) are used for DR550 Storage Controller (DS4200) connectivity.

For single-node configuration, when there is no mirroring, switch ports 2 and 3 are connected to the first ports of controllers A and B of DR550 Storage Controller (see Figure 2-21 on page 41), and, when a second DR550 Storage Controller is present, switch ports 8 and 9 are connected to the controllers A and B of the second DR550 Storage Controller.

When there is mirroring, either port 10 (SW) of the second quad switch SFP is used to connect to the secondary DR550 or port 12 (LW) of the fourth quad switch SFP is used to connect to the remote DR550 (see Figure 2-21). If the remote switch is less then 300 meters away, then you can use short wave SFPs; if it is more then 300 meters away, you have to use long wave SFPs.

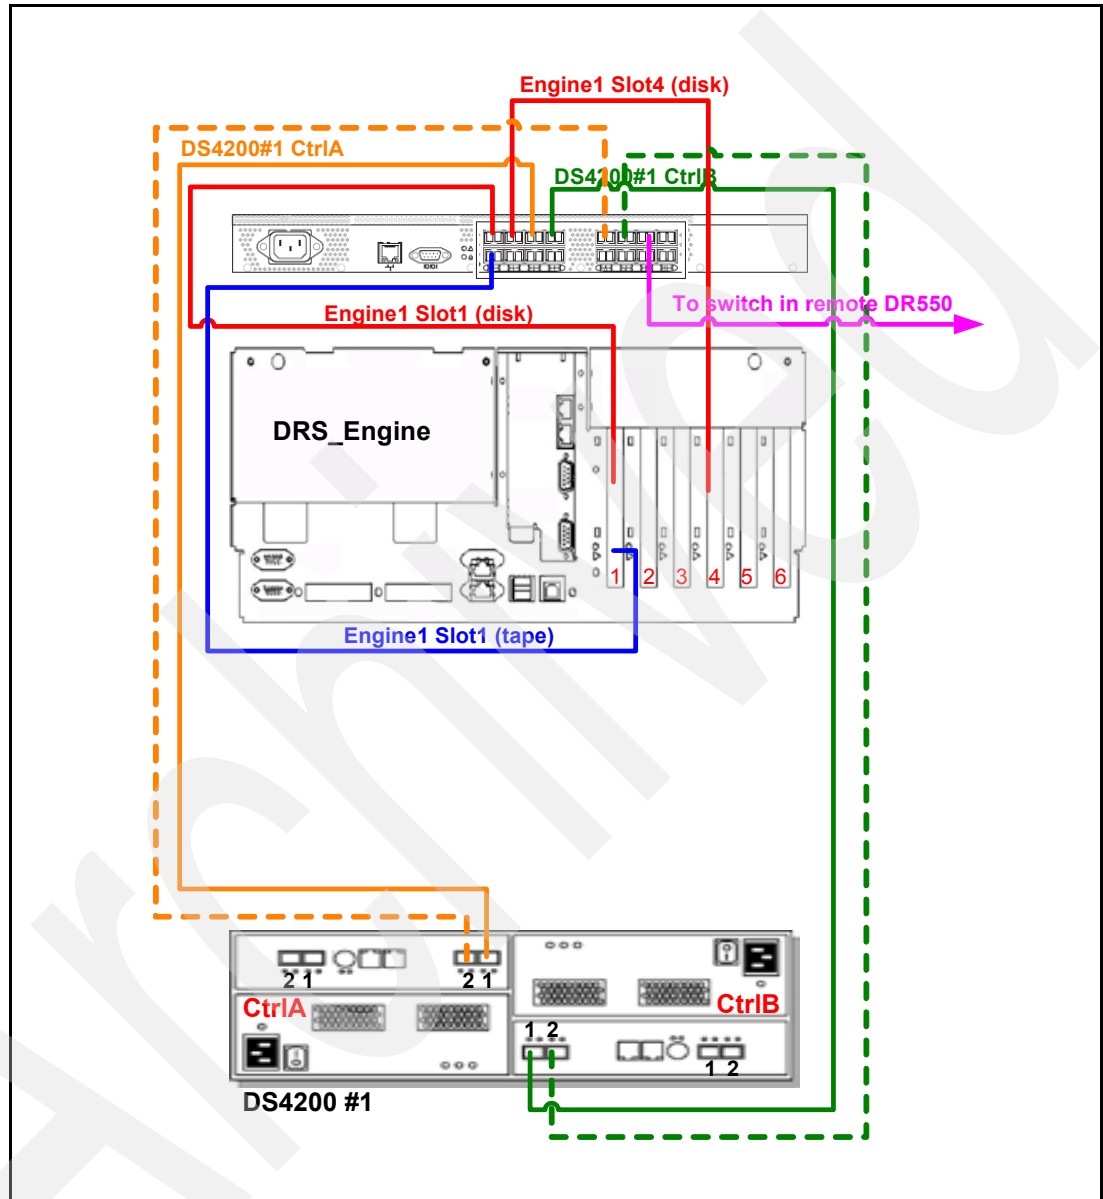


Figure 2-21 DR 550 Model DR2 FC connections: Single node from one DR550 Storage Controller (DS4200) to one DR550 SAN Switch to one DR550 SSAM Server (p5 server) with SW mirroring

In dual node configuration where there is only one DR550 Storage Controller (DS4200), ports 1 of controller A is connected to port 2 of the first switch, ports 1 of controller B is connected to ports 2 of the second switch (see Figure 2-22).

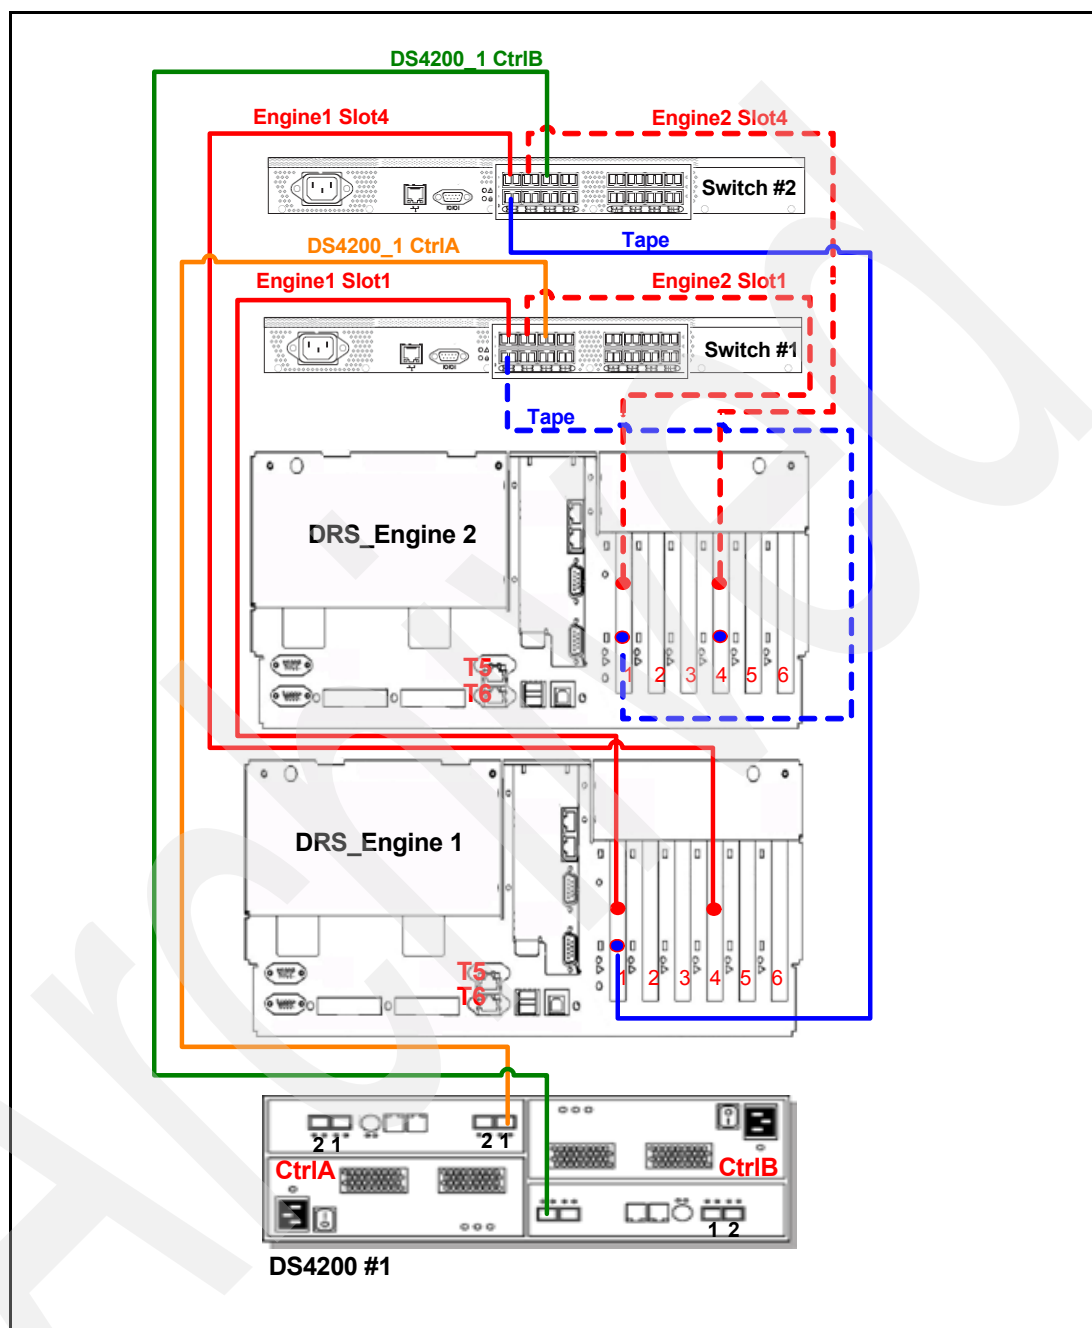


Figure 2-22 DR 550 Model DR2 FC connections: Dual node from two DR550 Storage Controllers (DS4200) to two DR550 SAN Switches to two DR550 SSAM Server (p5 servers) without mirroring

Note: Although technically possible, we do not recommend that you share the DR550 fabric or fabrics with other fabrics or attach non-DR550 related components, such as external servers or storage devices. Doing so compromises the security concepts of the DR550 and could have implications, for example, regarding compliance certification.

Zone configurations

Zoning is a licensed fabric management service for the 2005-B16 switches, which can be used to create logical subsets of devices within a storage area network and enable partitioning of resources for management and access control purposes. It is a recommended *best practice* in SANs to implement zoning. In particular, in implementations connecting DS4200 to AIX, which is the case in the DR550 offering, zoning is a must.

The SAN Switch 2005-B16 can be zoned two ways: by ports or by worldwide names. Zoning by ports allows easy replacement of connecting devices without having to change zones as long as they are connected to the same ports on the switch. In the DR550, the switches are preconfigured with a port-based zoning (see Table 2-10 and Table 2-11 on page 44).

Table 2-9 lists the zone configuration for the single SAN Switch B1 in the single-node configurations.

Table 2-9 Zone configuration for single-node DR550 DR2

Zone	Ports	Zone name	Usage
1	Short wave ports 0, 2, 8	zone.Engine1_Slot1_port1_to_DS4200_CtrlA: Engine1_Slot1_port1; DS4200_1_Ctrl_A1; DS4200_2_Ctrl_A1 alias.Engine1_Slot1_port1:1,0 alias.DS4200_1_Ctrl_A1:1,2 alias.DS4200_2_Ctrl_A1:1,8	SSAM Server #1 Storage Controller A #1 Storage Controller A #2 (Optional)
2	Short wave ports 1, 3, 9	zone.Engine1_Slot4_port1_to_DS4200_CtrlB: Engine1_Slot4_port1; DS4200_1_Ctrl_B1; DS4200_2_Ctrl_B1 alias.Engine1_Slot4_port1:1,1 alias.DS4200_1_Ctrl_B1:1,3 alias.DS4200_2_Ctrl_B1:1,9	SSAM Server #1 Storage Controller B #1 Storage Controller B #2 (Optional)
3	Short wave ports 4, 5, 6, 7	zone.Engine1_to_Tape:1,4;1,5;1,6;1,7	Tape devices

Table 2-10 and Table 2-11 on page 44 list the zone configurations for the SAN switches in the dual-node configurations.

Table 2-10 Zone configuration for switch 1 in dual-node DR550 DR2

Zone	Ports	Zone name	Usage
1	Short wave ports 0, 2, 8	zone.Engine1_to_DS4200: Engine1_Slot1_port1; DS4200_1_Ctrl_A1; DS4200_2_Ctrl_A1 alias.Engine1_Slot1_port1:1,0 alias.DS4200_1_Ctrl_A1:1,2 alias.DS4200_2_Ctrl_A1:1,8	SSAM Server #1 Storage Controller A #1 Storage Controller A #2 (Optional)
2	Short wave ports 1, 2, 8	zone.Engine2_to_DS4200: Engine2_Slot1_port1; DS4200_1_Ctrl_A1; DS4200_2_Ctrl_A1 alias.Engine2_Slot1_port1:1,1 alias.DS4200_1_Ctrl_A1:1,2 alias.DS4200_2_Ctrl_A1:1,8	SSAM Server #2 Storage Controller A #1 Storage Controller A #2 (Optional)

Zone	Ports	Zone name	Usage
3	Short wave ports 4, 5, 6, 7	zone.Engine1_to_Tape:1,4;1,5;1,6;1,7	Tape devices

Table 2-11 Zone configuration for switch 2 in dual-node DR550 DR2

Zone	Ports	Zone name	Usage
1	Short wave ports 0, 2, 8	zone.Engine1_to_DS4200: Engine1_Slot4_port1; DS4200_1_Ctrl_B1; DS4200_2_Ctrl_B1 alias.Engine1_Slot4_port1:1,0 alias.DS4200_1_Ctrl_B1:1,2 alias.DS4200_2_Ctrl_B1:1,8	SSAM Server #1 Storage Controller B #1 Storage Controller B #2 (Optional)
2	Short wave ports 1, 2, 8	zone.Engine2_to_DS4200: Engine2_Slot4_port1; DS4200_1_Ctrl_B1; DS4200_2_Ctrl_B1 alias.Engine2_Slot4_port1:1,1 alias.DS4200_1_Ctrl_B1:1,2 alias.DS4200_2_Ctrl_B1:1,8	SSAM Server #2 Storage Controller B #1 Storage Controller B #2 (Optional)
3	Short wave ports 4, 5, 6, 7	zone.Engine2_to_Tape:1,4;1,5;1,6;1,7	Tape devices

Note: The defined alias names might differ; however, this does not affect the communication between the devices. The important part is that the appropriate switch ports are combined in the distinct zones.

Figure 2-23 and Figure 2-24 on page 45 illustrate the zoning configuration for the single-node and dual-node DR550 model DR2.

For configurations with the Enhanced Remote Mirroring option (ERM), ERM is explained in detail in Chapter 13, “DR550 and Enhanced Remote Mirroring” on page 521.

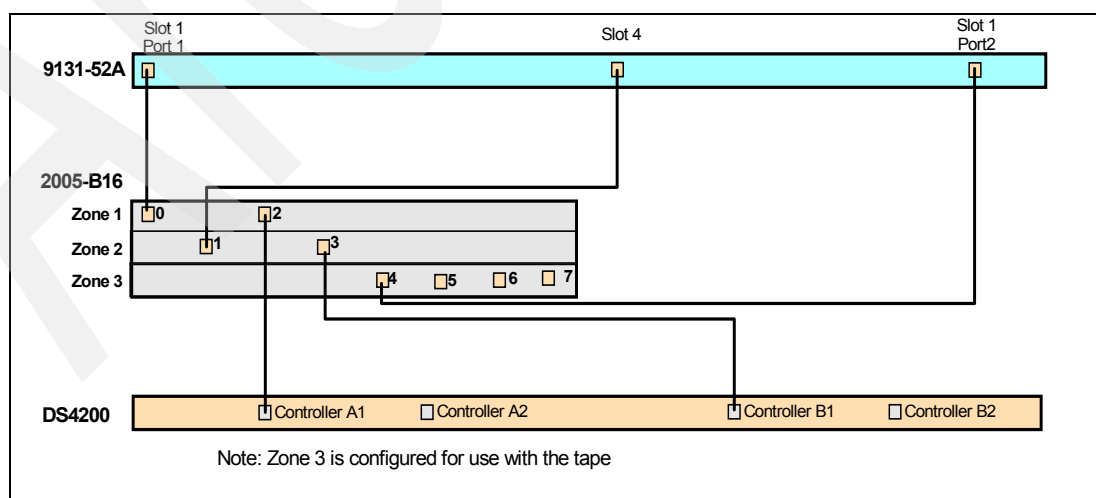


Figure 2-23 Zoning illustration for single-node configurations

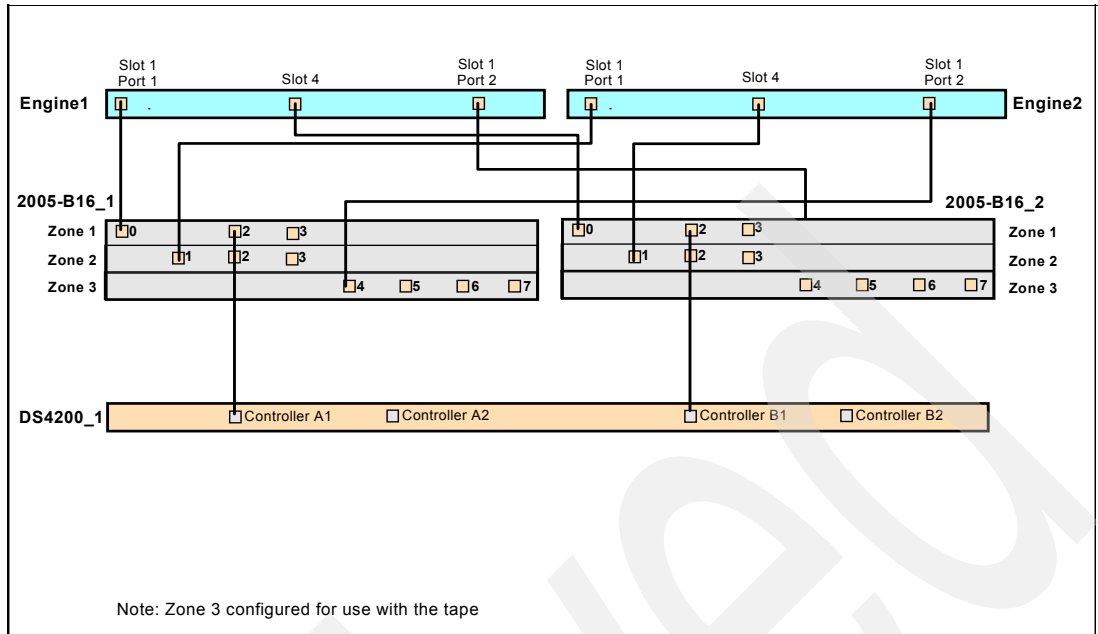


Figure 2-24 Zoning illustration for dual-node configurations

2.5.6 Storage configuration

The following applies to the DR550 Storage Controller (DS4200).

DR550 Storage Controller configuration and management

The DR550 Models DR1 and DR2 (for this topic, we will refer to both models as DR550) are shipped with the DS4000 Storage Manager for DR550 (SMclient) installed on the DR550 SSAM Servers (p5 52A servers). The SMclient provided with the DR550 has been enhanced to provide additional security so that no deletion of data when using the SMclient can take place either by accident or by malicious intent.

As you can see in Figure 2-25, the delete logical drive and delete array functions, for example, are not visible on the menu because they have been disabled.

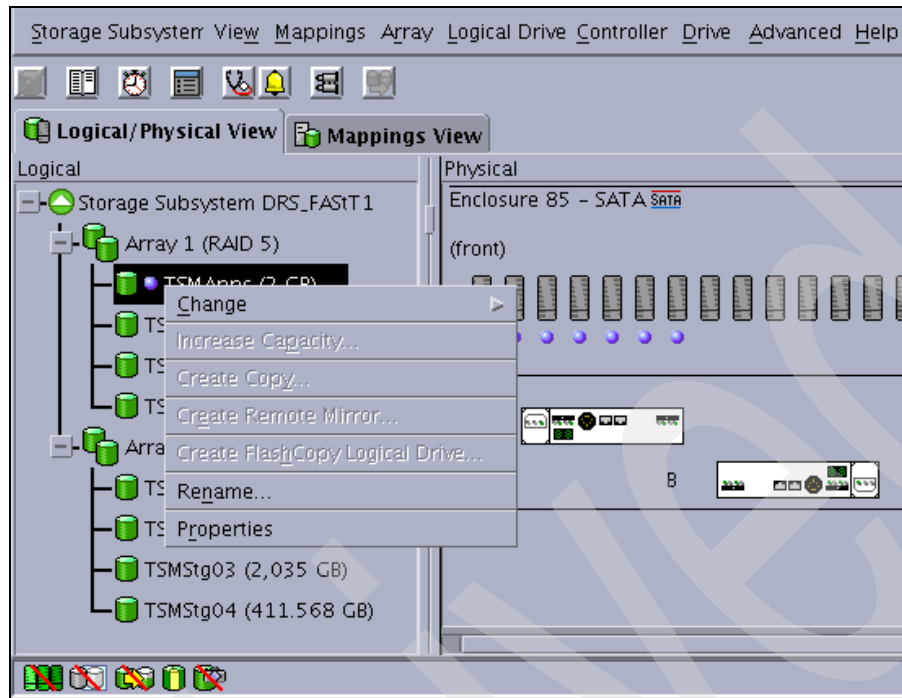


Figure 2-25 Storage Subsystem Management: No delete logical drive or array function

The SMclient graphical user interface must be started directly at the DR550 console (root access privileges are normally required to run the AIX SMclient). Connect the flat panel monitor to a DR550 SSAM Server (p5 engines) by pressing the Print Screen key (alternatively, press Ctrl twice) on the keyboard and selecting the appropriate entry from the window. Log in to AIX as dr550 and su to root; Enter startx to start the X-Server environment if not already done, then enter SMclient.

After a few seconds, the Storage Manager initial window will open (Figure 2-26 on page 47), followed shortly by the main window (Figure 2-27 on page 47).



Figure 2-26 Initial window of DS4000 Storage Manager for DR550

Important: You need first to log on as the dr550 user, then switch user to root, run `su -`, and then you can run `startx` and run `SMclient`.

For compliance with most regulations, the AIX root user in the DR550 is not able to work from a remote Ethernet connection. This means that using `SMclient` is only possible at the local console.

To open the subsystem management window, left-click a subsystem in the navigation pane on the left (for example, Storage Subsystem DRS_DS4200_Secondary), right-click, and select **Manage Device** from the menu (or just double-click the subsystem name).

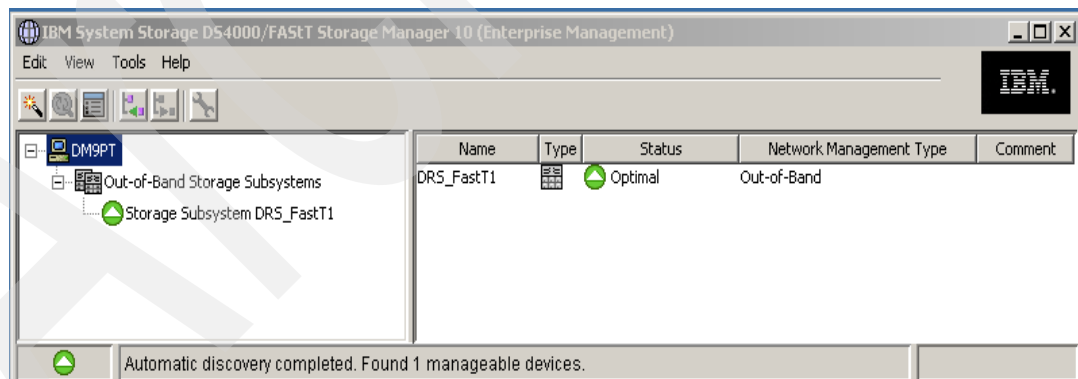


Figure 2-27 Main window of IBM DS4000 Storage Manager for DR550

To end the graphical (CDE) session at the console, close the Storage Manager application. Then, hold down the right button, select **End Session**, and then **Quit** from the pop-up menu. Confirm by clicking **OK** in the next window, and the CDE session will close. You are returned to the ASCII-prompt, where you can type `exit` to log out. Alternatively, pressing the Ctrl key and Alt key simultaneously and then the Backspace key will also close the CDE session.

Storage configuration and partitioning for DR550 Storage Controller

This section introduces common management concepts and basics associated with storage configuration using DR550 Storage Controller (DS4200). In parallel, we explain default configuration settings used for the DR550.

Storage configuration on the DR550 Storage Controller is accomplished by means of storage arrays and logical drives. An *array* is a set of drives that the controller groups logically together to provide one or more logical drives to an application host or cluster.

A *logical drive (or volume)* is a logical structure you create on a storage subsystem for data storage. Creating arrays and logical drives is one of the most basic steps and is required before you can start using the physical disk space, that is, you divide your disk drives into arrays and create one or more logical drives inside each array.

Enhanced Remote Mirroring (ERM) (only for DR550)

The Enhanced Remote Mirroring (ERM) option is a premium feature of the IBM DS4000. This feature is only available for DR550 Model DR2.

The ERM option is used for online, real-time replication of data between data retention subsystems at different locations. In the event of a disaster or an unrecoverable error at one data retention subsystem, you can promote the second data retention subsystem to take over responsibility for normal I/O operations. Refer to Chapter 13, “DR550 and Enhanced Remote Mirroring” on page 521 for details.

Important: Note that if you implement ERM with a DR550 equipped with the optional FSG, data cached on the FSG disks and local metadata (that is not yet backed up in the storage of DR550 Storage Controller) are not mirrored.

RAID levels and array configuration

Redundant Array of Independent Disks (RAID) is a method of configuring multiple disk drives in a storage subsystem for high availability or high performance, or a combination of both. These goals are sometimes mutually exclusive and are attained by technologies called *striping* (performance enhancer) and *mirroring* (redundancy and availability). Striping of data means that data is split into chunks and spread across multiple drives such that all drives in an array are accessed in parallel, so a larger amount of contiguous data is written or read at a time. Mirroring is merely making exact copies of an entire drive, so in case of failures, data is still available on the mirrored disk. There are various RAID levels that implement combinations of these technologies.

Therefore, one of the most important decisions that you must make before creating an array is determining the RAID level.

For reasons of performance, fault tolerance, capacity, and storage efficiency, RAID 5 is the level of choice in the DR550. The DR550 Models DR1 and DR2 also provide a RAID 6 option.

Note: The DR550 uses RAID 5 arrays.

For the DR550 Model DR1 ordered with DS4200 Storage Server, the storage comes in several capacities. The minimum capacity is 6 TB of raw disk space that is built with eight (half a drawer) 750 GB disks in the DS4200 RAID controller enclosure. Up to two EXP420 units can be used to provide a maximum of 36 TB of storage. The RAID type is RAID 5, whereby a half drawer, if it is the only storage drawer, is formatted 2+P and 3+P with a global spare, otherwise it is formatted 6+P with one global spare. and a full drawer is formatted 6+P and 7+P with one global spare.

For the DR550 Model DR2, the storage comes also in several capacities. The minimum capacity is 6 TB of raw disk space that is built with eight (half a drawer) 750 GB disks in the DS4200 RAID controller enclosure. Up to two DS4200 Storage Servers and 12 EXP420 units can be used to provide a maximum of 168 TB of storage. The RAID type for DR550 V4.5 is RAID 5, whereby a half drawer, if it is the only storage drawer, is formatted 2+p and 3+p with a global spare, otherwise it is formatted 6+p with one global spare. and a full drawer is formatted 6+p and 7+p with one global spare.

Figure 2-28 shows an array configuration on a DR550 V4.5

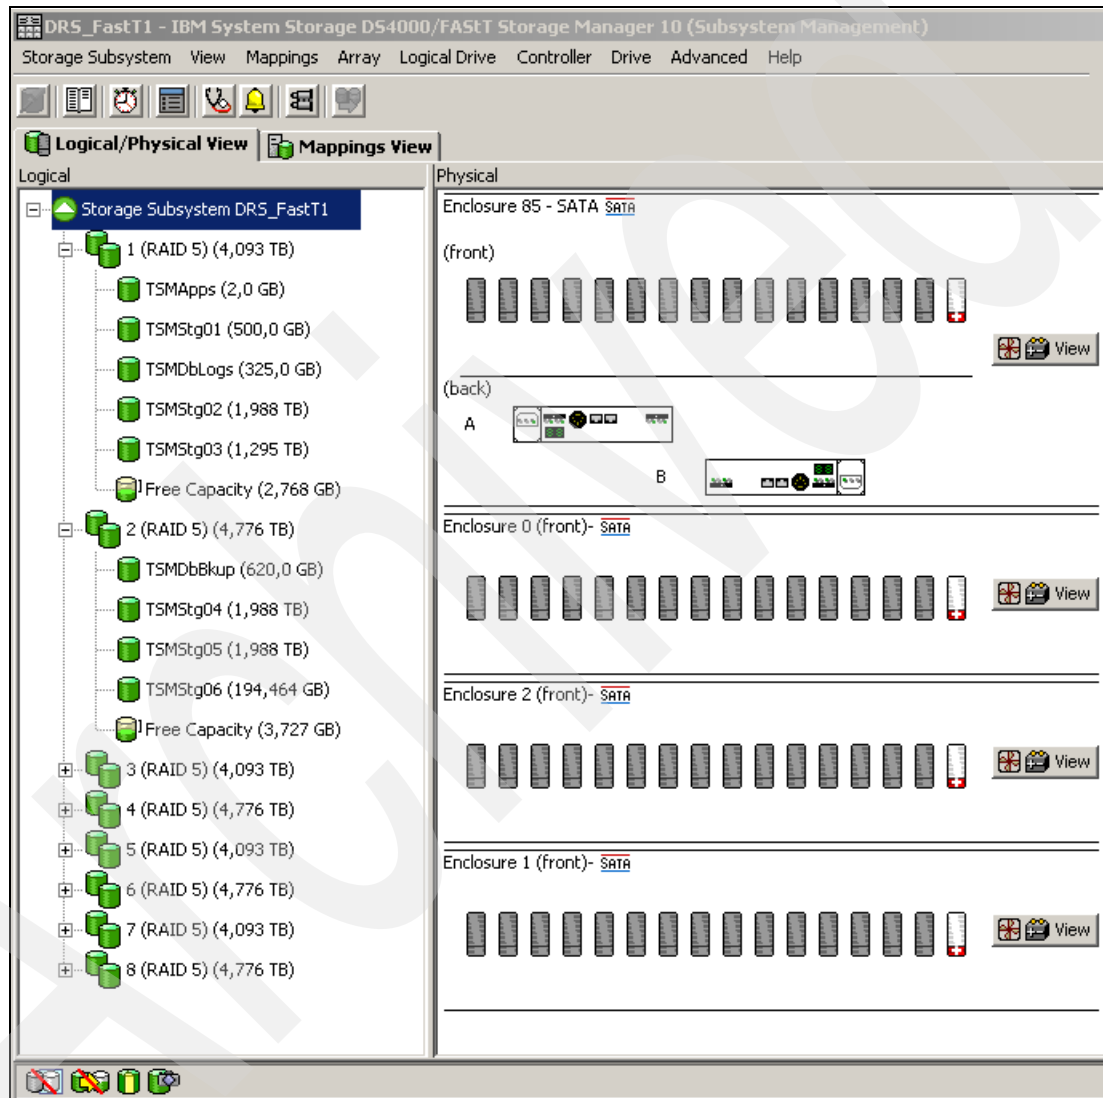


Figure 2-28 Array configuration on a DR550 V4.5

We discuss the mapping of the DR550 Storage Controller (DS4200) logical drives to the AIX Logical Volume Manager and the volume group design in detail in 5.2.12, "IBM System Storage Archive Manager and IBM System Storage DR550" on page 188.

Hot-spare drive

Hot-spare drives provide additional protection that might be essential in case of a disk drive fault. A hot-spare drive is similar to a standby replacement drive. The data from the failed disk drive is automatically rebuilt by the controller to the hot-spare drive, and the spare takes the place of the failed one. When the failed drive is eventually replaced with a new one, the data from the hot-spare drive is copied back to the new drive, and the hot-spare drive goes back to its role as a replacement drive. Alternatively, you can also keep the spare as part of the array (in that case, no copyback takes place) and in turn, define the replacement drive as a new hot spare. It is important to note that the DS4000 series uses global hot-spares, meaning that they can take over for any failed drive regardless of its enclosure.

Note: The DR550 Model DR1 ordered with DR550 Storage Controller (DS4200) and the DR550 Model DR2 come with the arrays and volumes predefined.

Segment size

The choice of a *segment size* can have a major influence on performance in both I/Os per second and throughput. Larger segment sizes provide better I/Os per second, and smaller segment sizes give better data throughput. The DR550 uses the default segment size of 64 KB, which is set by the Storage Manager.

Volume mapping

As you can see from the mappings view of the DS4000 Storage Manager for DR550 (shown in Figure 2-29 on page 51), all volumes are mapped to the default group, meaning all logical drives that are created on the arrays are available to both DR550 SSAM Servers (p52A servers) attached to the DR550 Storage Controller (DS4200). This is necessary in HACMP and other cluster environments.

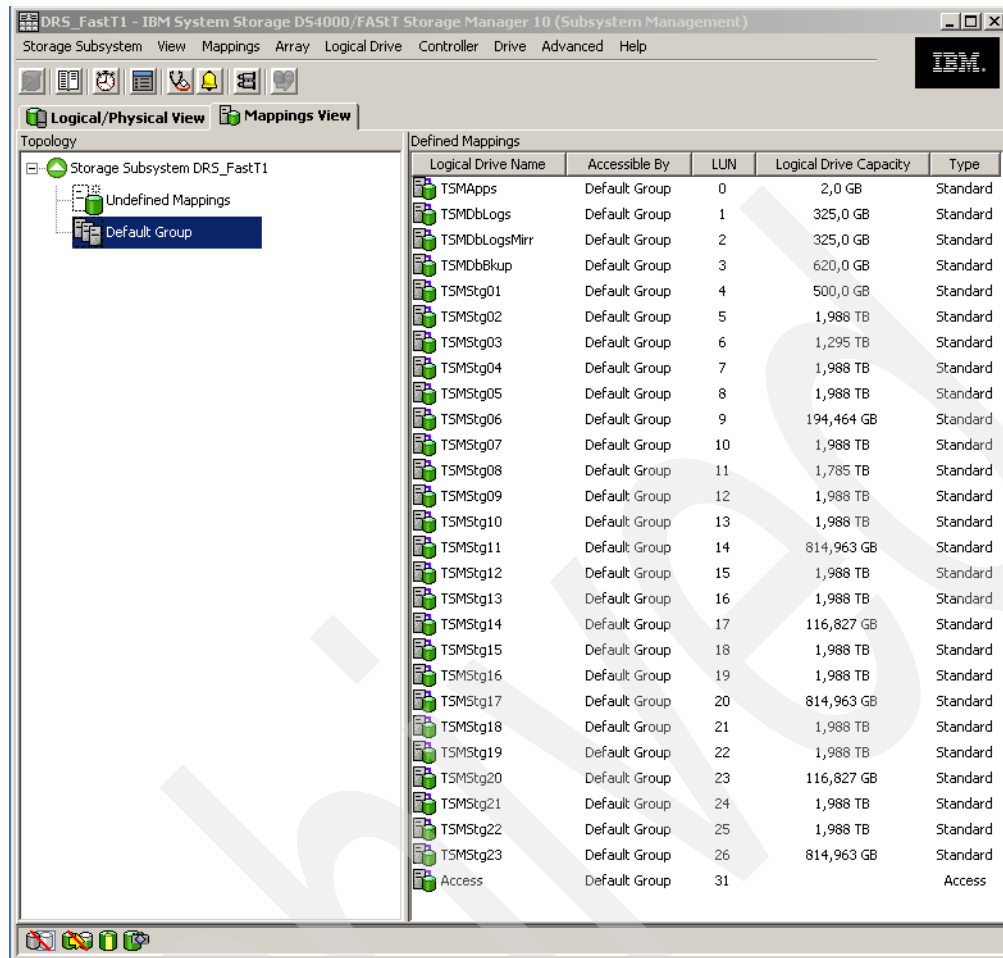


Figure 2-29 Volume mapping of a DR550 V4.5

Preferred path

The DR550 Storage Controller (DS4200) has two controllers (A and B) for redundancy. All logical drives created on the DR550 Storage Controller are accessible from either of the two controllers, but have a preferred owner. Each DR550 SSAM Server (p 52A server) in the DR550 has two Fibre Channel host bus adapters (FC HBAs) installed. Each FC HBA has one or more paths to Controller A of the DR550 Storage Controller. Similarly, the other FC HBA has one or more paths to Controller B. In case of a path failure, meaning either a FC HBA failure, switch failure, SFP, fiber link failure, or even a DS4200 controller failure, the logical drives are accessible on the remaining path(s). For performance reasons, the preferred paths are distributed between the controllers automatically.

Figure 2-30 shows the distribution of the logical drives to the controllers (Controller A shown).

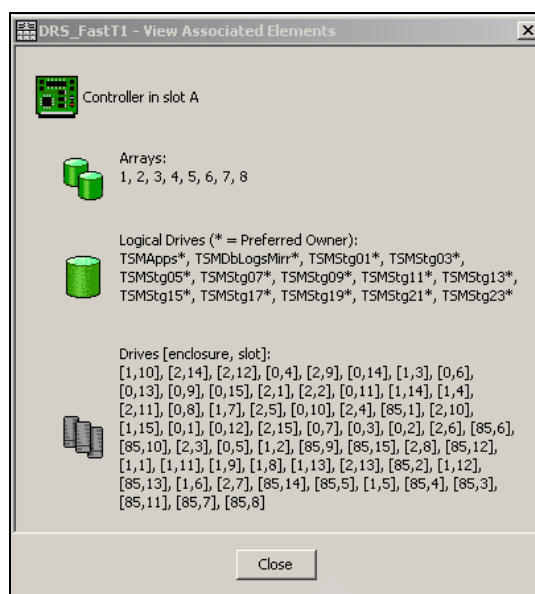


Figure 2-30 Channel distribution on a DR550 V4.5

2.5.7 Disk space allocation

System Storage Archive Manager (SSAM) is installed on each delivered DR550 SSAM Server (p5 52A server). The SSAM server executable files are installed locally on each of the DR550 SSAM Servers in a single or dual-node system. This is the standard configuration for a Tivoli Storage Manager server running within an HACMP cluster in the case of a dual-node configuration.

For the DR550 Model DR1, all configuration files, the SSAM database, the recovery log, and the storage pool volumes are configured on specific file systems within the internal SCSI drives in the DR550 SSAM Server. This will differ from time to time as the option to purchase an external disk is available, and then the storage pool volumes will be configured on the external disk, not the internal disk.

For the DR550 Model DR2, all configuration files, the SSAM database, the recovery log, and the storage pool volumes are configured on specific logical volumes residing on the DR550 Storage Controller (DS4200). With the dual-node configuration, all those files or resources are shared in the HACMP cluster.

DS4200 configuration

It is possible to order the DR550 Model DR1 with additional storage and in that case, this section will apply. If you did not order the DR550 Model DR1 with additional storage, this section will not apply to the DR550 Model DR1.

From the SMclient graphical interface view for the logical and physical configuration of the DR550 Storage Controller (DS4200), you can see the volume sizes and the positioning of the logical drive TSMDBKup, which is used for Tivoli Storage Manager database backups.

Tip: Starting with DR550 V4.0, the Tivoli Storage Manager database volumes and the Tivoli Storage Manager recovery log volumes are mirrored across Array1 and Array2 of the DR550 Storage Controller (DS4200).

DR550 Storage Controller configuration summary

- ▶ DR550 Storage Controller base unit has two RAID 5 arrays configured. The preferred path for array 1 goes to controller A and the preferred path of array 2 goes to controller B.
- ▶ When the DR550 Storage Controller base unit is fully populated, it includes one hot spare to recover single disk failures for the RAID 5 arrays. Each fully populated DR550 Expansion Drawer unit, configured as RAID 5, includes a global hot spare.
- ▶ 500 MB of disk storage space is reserved for the definition of the ERM Repository Logical drives, if needed.
- ▶ The Tivoli Storage Manager database and the recovery log volumes are mirrored across array1 and array2 in the DR550 Storage Controller base unit.
- ▶ There is also additional space available for emergency expansion of the database files if needed. This space is in the same volume group as the existing database and log files. 10 GB is reserved for the primary database, 10 GB reserved for the secondary copy, and 20 GB is reserved for the backup space. If you expand the primary database, be sure to expand the secondary copy and the backup space as well.

Figure 2-31 illustrates the array and the volume configuration of a DR550 Storage Controller with 750 GB SATA disk drives.

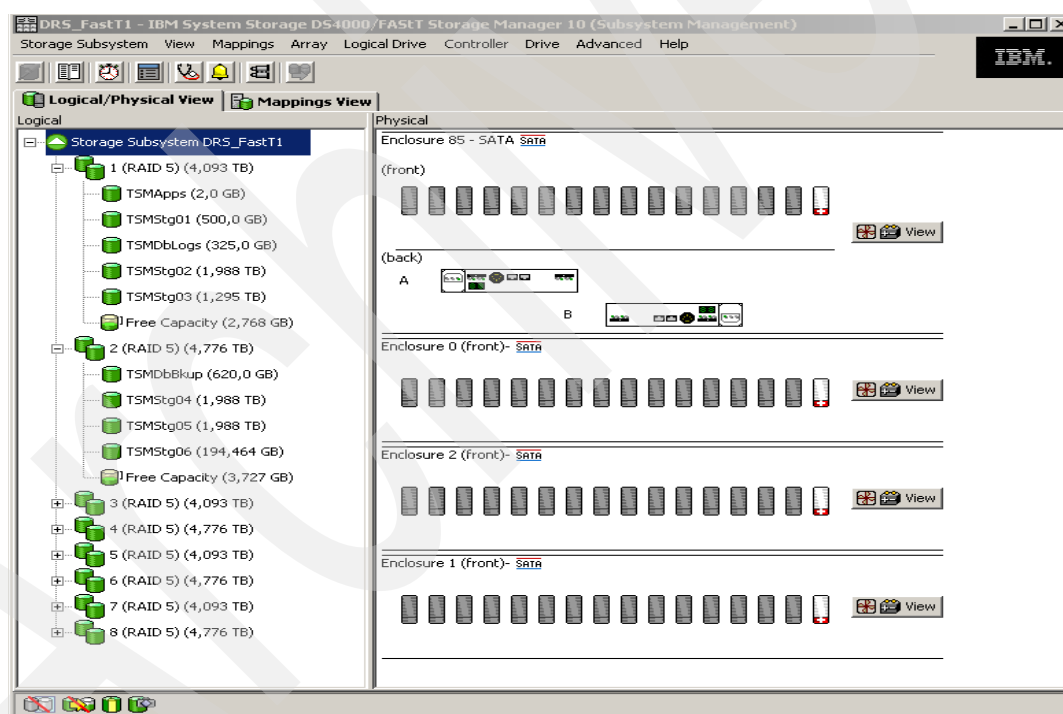


Figure 2-31 Physical configuration of the DR550 Storage Controller

AIX Logical Volume Manager (LVM) configuration

DR550 Model DR1 ships with internal disks in two RAID configurations. The first two disks are on a RAID-1 array and are used for the operating system. Drives 3 through 8 are in a RAID 5 array (5+P). All data (SSAM database, logs, database backup space, and archive spare) resides on the RAID 5 array within the internal disks.

For the DR550 Model DR1 ordered with a DR550 Storage Controller (DS4200), and for DR550 Model DR2, LUNs are configured based on the total storage capacity ordered. The whole disk capacity is allocated for use by SSAM. No disk capacity is set aside for other applications or servers.

For detailed information about the LUN configurations set at the default, refer to Appendix C, “DS4000 Logical Volume configuration”, of the *IBM System Storage DR550 Version 4.5 Problem Determination and Service Guide*, GA32-0576.

Figure 2-32 shows the complete AIX LVM layout and volume group design for the DR550 single and dual engines, including the internal SCSI hard disk drives and the DR550 Storage Controller logical drives directly relating to AIX physical volumes (hdisks). On these hdisks, the AIX logical volumes are configured.

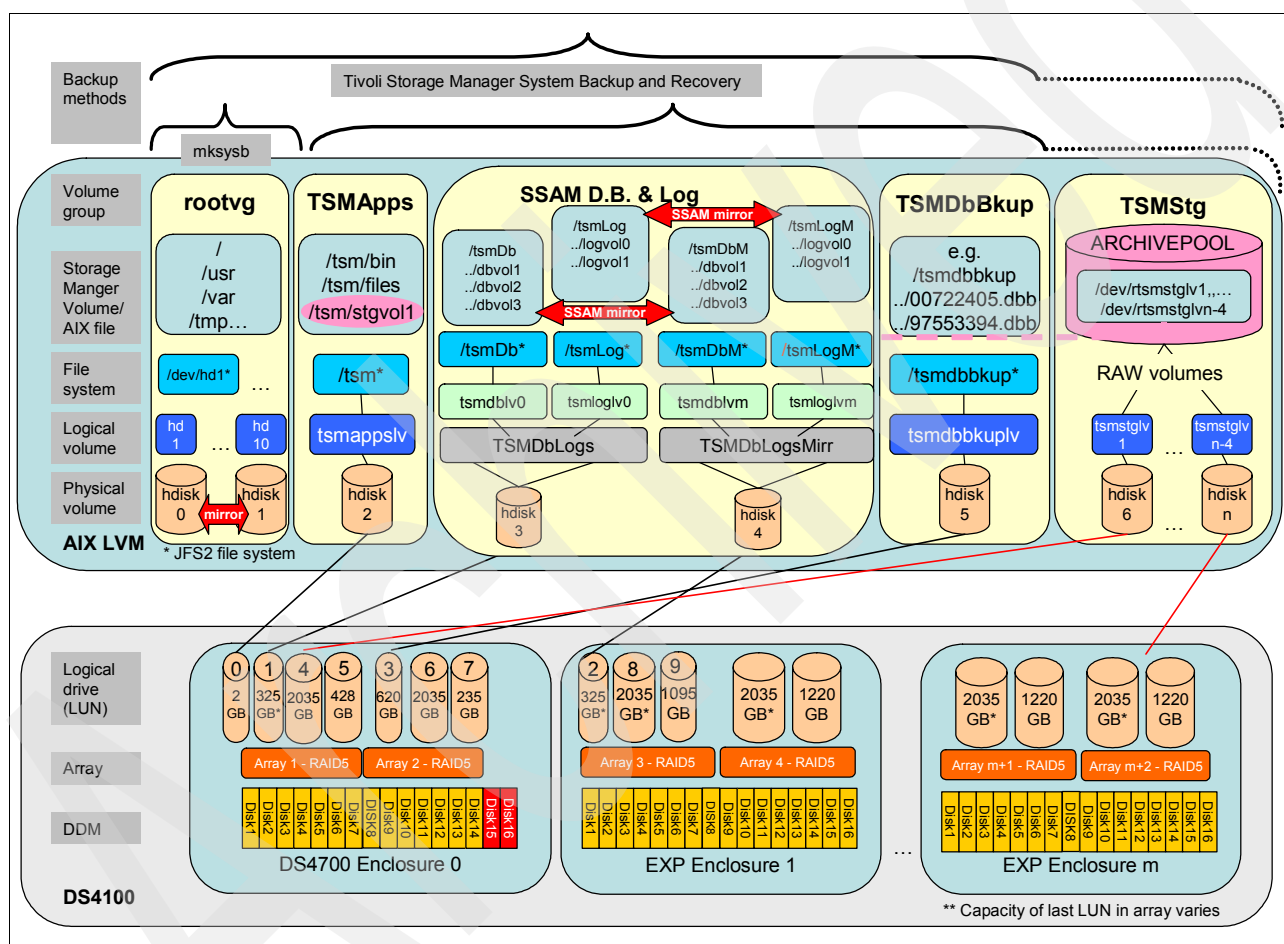


Figure 2-32 AIX LVM design and DS4700 configuration for DR550

The logical volumes can be formatted with a file system as is the case for the volume groups rootvg, TSMApps, TSMDBLogs, and TSMDBBkup (JFS2). Both the SSAM database and log file are mirrored to other logical volumes using SSAM mirroring. System Storage Archive Manager volumes defined in these file systems are basically AIX files residing in the directory structure of the file system. Contrary to that, the logical volumes within the TSMStg volume group do not carry a file system, but are accessed as raw logical volumes by System Storage Archive Manager and make up the primary pool ARCHIVEPOOL.

Predefined primary storage pools and storage pool volumes

Note: A comparably small file of 5 MBs called /tsm/stgvol1 from the volume group TSMApps is also allocated to the Tivoli Storage Manager ARCHIVEPOOL. This is the only System Storage Archive Manager volume in that pool not residing in the volume group TSMStg.

All remaining space on the DR550 Storage arrays is divided into 2,000 GB LUNs (the last LUN of each array is always smaller because of the remaining space of the array) and mapped to the AIX servers. Once mapped, they are configured as AIX physical volumes (*hdisks*). On top of these *hdisks*, logical volumes are configured with the AIX Logical Volume Manager (LVM). Each of these logical volumes has been defined to the primary storage pool archivepool as a raw logical volume. These volumes together make up the total capacity of the primary storage pool.

SSAM is configured with one storage pool volume (500 GB) set at the factory. The remaining volumes are predefined in AIX volume group TSMStg, but not assigned as storage pool volumes to the System Storage Archive Manager.

To add the predefined volumes to the SSAM storage pool, stop HACMP if applicable and SSAM (on both DR550 SSAM Servers in dual node systems). From the AIX command line, issue the **/usr/bin/addstg** command. Note that this must be run from DR550 SSAM Server #1 on dual node systems. The script creates /tmp/bldstg. This SSAM script lists all storage pool volumes that can be added to the SSAM Storage Pool. If you do not want to add the whole capacity, you can edit the bldstg SSAM script and remove the logical volume lines.

To execute the bldstg script, log in to SSAM and run the macro /tmp/bldstg. Now the storage pool volumes will be added to the ARCHIVEPOOL within SSAM. This may take a few minutes to complete. After the script completes, you will need to restart HACMP if applicable and SSAM (see 4.1, "Starting and stopping HACMP cluster services" on page 138).

DR550 Model DR1 without a DR550 Storage controller (DS4200)

The DR550 Model DR1 without DR550 Storage Controller uses only the internal storage in the DR550 SSAM Server (p5 52A). The server is equipped with eight 146 GB UltraSCSI 3 disk drives.

The DR550 Model DR1 storage has two arrays; the first two disks are mirrored in RAID 1 for the operating system. Disks 3 to 8 are configured as a RAID 5 system. There is no spare disk.

There is no space available for other applications or for any modifications. The DR550 Model DR1 without DR550 Storage controller should not be changed to accommodate other applications or workloads.

The disk space is used as follows:

- ▶ The AIX operating system occupies two disks (hdisk0).
- ▶ The IBM System Storage Archive Manager (SSAM) application code uses 1 GB.
- ▶ The System Storage Archive Manager database and database logs use 40 GBs.
- ▶ The backup space for the database and logs uses 130 GBs.
- ▶ The archivepool uses 509 GB (all of the remaining usable space).

2.6 What is preconfigured on the the optional FSG

The optional File System Gateway (FSG) is available as a stand-alone system or as a high availability (two-node) cluster.

Figure 2-33 shows an overview diagram of the FSG architecture.

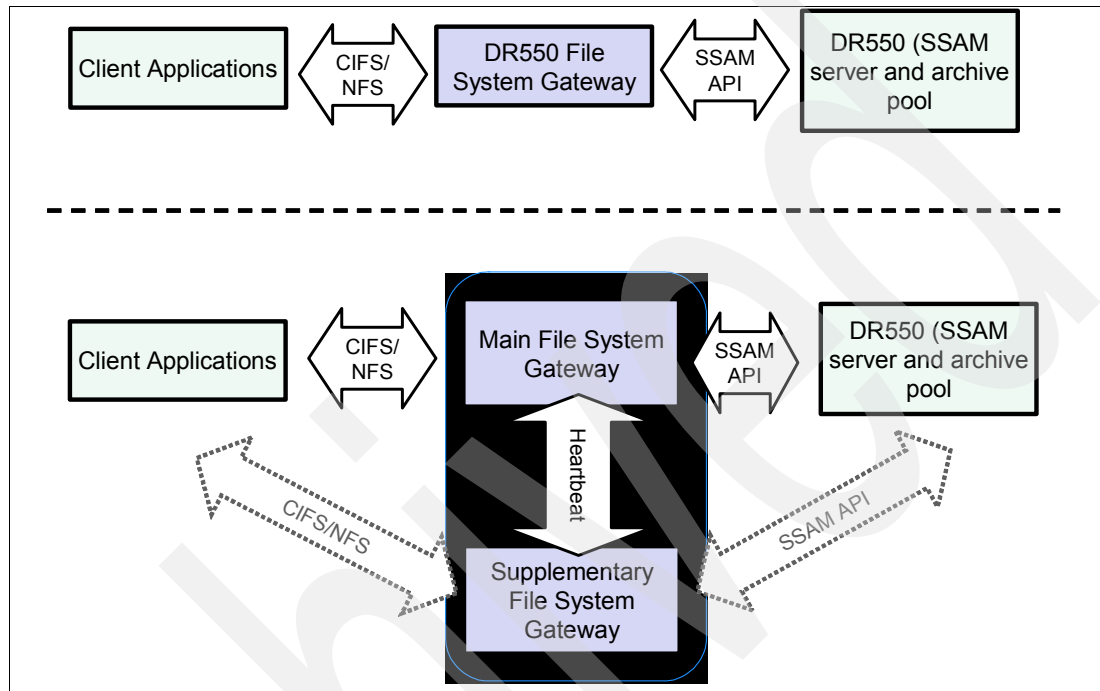


Figure 2-33 Overview File System Gateway - 2229-FSG

2.6.1 Hardware components

The File System Gateway (FSG) has its own type and model number, 2229-FSG. The FSG is based on an IBM System x 3650 with 2 GB RAM and is configured to fit the requirements for the File System Gateway software and for the customers. Each FSG node comes with six Ethernet ports. For details on the Ethernet port configuration, refer to 2.6.4, "FSG Network configuration and preset IP addresses" on page 57.

The 2229-FSG has six disks. Two of them, with a capacity of 73 GB each, are configured as RAID 1, and host the SLES operating system and the FSG application code. The remaining four disks, each offering a 300 GB capacity, are configured as RAID 6. For details on the storage configuration, refer to 2.6.5, "Storage configuration" on page 59.

2.6.2 Software components

The File System Gateway software is based on SUSE Linux Enterprise Server (SLES) Version 10.

The following software components are already installed on the FSG:

- ▶ BIOS: 1.03, System x 3650
- ▶ Operating System: SLES 10, Kernel 2.6.16.21-0.8-smp
- ▶ Level 2 Agent (IBM Director): 5.20.1

- ▶ SSAM (Tivoli Storage Manager) Archive client: 5.5.0.0
- ▶ Application: File System Gateway V1.1.1

2.6.3 FSG security

Table 2-12 shows the user accounts predefined for the FSG.

Table 2-12 Default user accounts on the FSG

User account	Password
root	dR550fsg
fsgadm	C1fsNFS
drg	drg

User fsgadm has to be used locally at the console as well; after you are logged on as fsgadm, you can in turn switch to user 'root'.

The SLES 10 firewall is disabled. All other SLES 10 settings are still defaults.

Important: In the FSG Version 1 software, the SLES firewall has to be disabled.

2.6.4 FSG Network configuration and preset IP addresses

The File System Gateway FSG-2229 comes equipped with six Ethernet ports.

- ▶ The two on-board adapters are only used in the HA cluster FSG for cluster heartbeat.
- ▶ Dual port Gigabit Ethernet over copper adapters are installed in PCI slots 3 and 4. One port on each adapter is used for communicating with DR550 SSAM servers either using internal switches (DR550 DR2 only) or using customer supplied external Ethernet switch.
- ▶ Optional fiber Ethernet adapters are installed in PCI slot 1 and 2 and will be used for attachment to the customer fiber Ethernet network.

Figure 2-34 shows the FSG cabling.

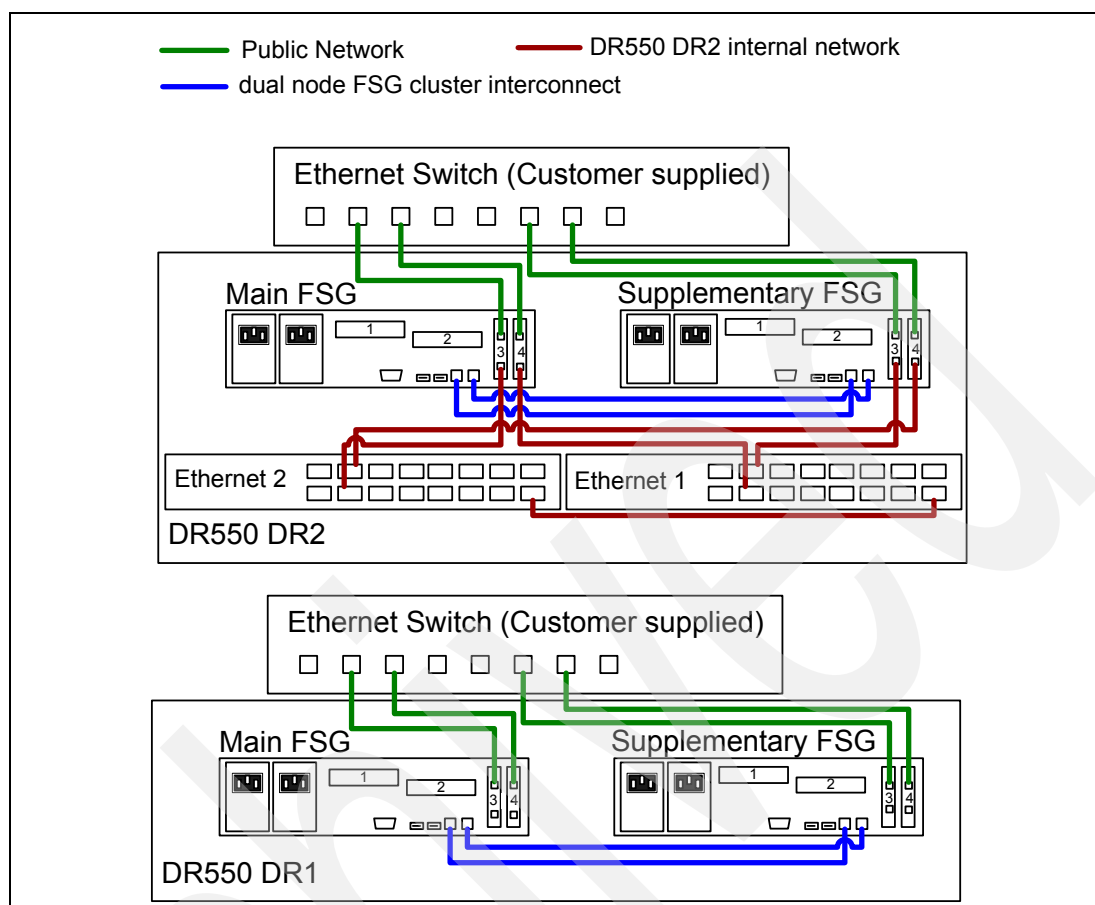


Figure 2-34 FSG cabling

Important: Do not change the main and supplementary host names. Changing those names would make the FSG inoperable.

Table 2-13 shows the FSG IP address settings.

Table 2-13 File System Gateway (FSG) - IP address settings

IP address	Description	Device	Host name	IP address change
192.168.1.121	Public port 1	eth0	Main	No
192.168.1.122	Public port 2	eth1	Main	No
169.254.1.1	heartbeat_1	eth2	Main	No
169.254.1.5	heartbeat_2	eth3	Main	No
192.168.1.120	Bonded IP address (public)	ifcfg-bond0 (/etc/sysconfig/network)	Main	Yes
192.168.1.131	Public port 1	eth0	Supplementary	No
192.168.1.132	Public port 2	eth1	Supplementary	No

IP address	Description	Device	Host name	IP address change
169.254.1.2	heartbeat_1	eth2	Supplementary	No
169.254.1.6	heartbeat_2	eth3	Supplementary	No
192.168.1.140	Bonded IP address (public)	ifcfg-bond0 (/etc/sysconfig/network)	Supplementary	Yes
192.168.1.140	Virtual IP address (cluster)	Configured in the DRGC.xml file	-	Yes

For additional details about the network configuration of the File System Gateway, refer to 3.7.3, “Configuring the network settings at the DR550 FSG” on page 116.

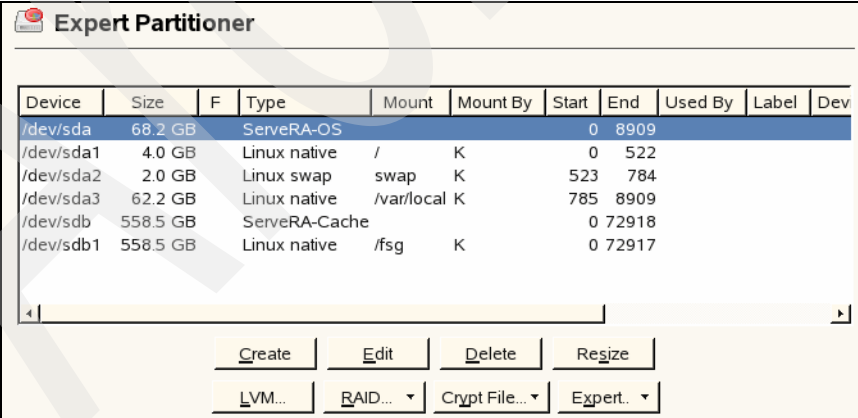
2.6.5 Storage configuration

As described before, the SUSE Linux (SLES 10) is installed and preconfigured on a RAID 1 made of the two 73 GB hard disks. Linux sees this RAID 1 logical drive as its first SCSI disk /dev/sda.

The other four disks installed in the FSG are preconfigured as a RAID 6 logical drive. It is seen by the Linux operating system as device /dev/sdb.

RAID 6 can be thought of as “RAID 5, but more”. It stripes blocks of data and parity across an array of drives like RAID 5, but calculates and maintains two sets of parity information for each block of data. The goal of this duplication is solely to improve fault tolerance; RAID 6 can handle the failure of any two drives in the array while other single RAID levels can handle at most only one faulty drive. Performance-wise, RAID 6 is generally slightly worse than RAID 5 in terms of writes due to the added impact of more parity calculations, but might be slightly faster in random reads due to spreading of data over one more disk.

Figure 2-35 shows the preconfigured partitioning for the File System Gateway.



Device	Size	F	Type	Mount	Mount By	Start	End	Used By	Label	Dev
/dev/sda	68.2 GB		ServerRA-OS			0	8909			
/dev/sda1	4.0 GB		Linux native	/	K	0	522			
/dev/sda2	2.0 GB		Linux swap	swap	K	523	784			
/dev/sda3	62.2 GB		Linux native	/var/local	K	785	8909			
/dev/sdb	558.5 GB		ServerRA-Cache			0	72918			
/dev/sdb1	558.5 GB		Linux native	/fsg	K	0	72917			

Expert Partitioner

Create Edit Delete Resize

LVM... RAID... Crypt File... Expert...

Figure 2-35 File System Gateway (FSG) - Partition table

2.6.6 Cluster configuration

It is important, before you order your File System Gateway, to decide whether you need the stand-alone or high availability cluster FSG. Settings will be set by manufacturing depending upon the configuration ordered.

Important: If you want to change the role of the FSG from main to supplementary, you have to reinstall the FSG from scratch (operating system and FSG application code).

The stand-alone FSG configuration is similar to the configuration of the primary node in an HA cluster, with the exception of the DRGC.xml configuration file (see 3.7.7, “FSG post-installation and initial setup” on page 125).

When you order the dual-node FSGs, the Ethernet crossover cables for the cluster communication (heartbeat) are preinstalled.

2.7 DR550 offerings summary

Table 2-14 is a summary of all the DR550 offerings.

Table 2-14 Summary of offerings

	DR1	DR2 single-node	DR2 dual-node
7014-T00 Rack	-	Yes	Yes
7014-S25 Rack	Yes	-	-
9131-52A	Yes	Yes	Yes
7316 TF3 console	Yes	Yes	Yes
KVM Switch	Optional	Optional	Yes
First DS4200	Optional	Yes	Yes
EXP420s	Optional	Optional	Optional
Second DS4200	-	Optional	Optional
First 2005-B16	-	Optional	Yes
Second 2005-B16	-	-	Yes
Internal Ethernet switches	-	Yes	Yes
Stand-alone FSG	Optional	Optional	-
High Availability FSGs	Optional	Optional	Optional
AIX 5L V5.3	Yes	Yes	Yes
HACMP V5.4.2	-	-	Yes
SSAM V5.5	Yes	Yes	Yes
Tivoli Integrated Solution Console (ISC)	Yes	Yes	Yes

	DR1	DR2 single-node	DR2 dual-node
Tivoli Admin Center	Yes	Yes	Yes
DS4000 Storage Manager for DR550	Optional	Yes	Yes
Remote Mirroring	-	Optional	Optional

Archived

DR550 and FSG planning and installation

This chapter provides information about the planning and installation of the IBM System Storage DR550 model DR1 (2233-DR1) and model DR2 (2233-DR2), with the optional DR550 File System Gateway (2229-FSG). This includes the site preparation, the sequence of steps for starting and stopping the system, and the required configuration settings for deployment in an existing environment.

For the DR550, the chapter explains how to use the keyboard-video-mouse (KVM) console kit, and how to get local and remote graphical output. For the dual-node DR550 DR2 configuration, the chapter also discusses the HACMP cluster configuration and shows how to start and stop the cluster services. For the DR550 File System Gateway, the chapter explains how to access and set up the FSG.

New capabilities of DR550 Storage Controller (DS4200 disk subsystem) are also discussed in this chapter.

3.1 Planning the installation

The very first step when planning for a DR550 is to make sure that you order the model that best fits your requirements. When you have determined the configuration you need, you can start planning for the physical installation and the initial setup of the DR550 and, if applicable, the DR550 File System Gateway. It is also recommended to plan for the installation of the additional monitoring servers (IBM Director server and RSM). These servers are not shipped as part of the DR550 and have to be installed separately. DR550 supports the following types of monitoring servers:

- ▶ IBM Director

IBM Director provides monitoring capabilities for operating systems and application components of DR550, and call home support for the FSG systems, through the Electronic Service Agent (eSA) for System x extension.

More information is available at:

<http://www-03.ibm.com/systems/management/director/>

- ▶ Remote Support Manager

Remote Support Manager (RSM) is designed to allow problem reporting and remote access (from IBM support) to the DR550 Storage Controller (DS4200). There is a specific version of RSM for the DR550.

More information about RSM in general is available at:

<http://www-03.ibm.com/systems/storage/disk/rsm/index.html>

General planning considerations

Before ordering a DR550, make sure you know your requirements for:

- ▶ High availability and disaster recovery
 - ▶ Archival capacity of the system and RAID type
 - ▶ The type of interface your archiving or content management application requires
- Specifically, does your application support the SSAM APIs, or do you need a CIFS or NFS protocol?

Data Center environment conditions

- ▶ Does your data center provide enough floor load capacity and service clearance?
- ▶ Do you have enough capacity left on your power source or do you want to attach the DR550 to an uninterruptible power supply (UPS)?
- ▶ Do you have a climate controlled environment?

Network

- ▶ Do you have enough free Ethernet ports and cables?
- ▶ Do you have enough unused IP addresses?
- ▶ Do you know your firewall settings?
- ▶ Do you have an SNMP server (such as the IBM Director)?

Alerting / Monitoring

- ▶ The call home function of the DR550 SSAM servers (System p p52A), is based on eSA for System p (included). You have to provide an internet connection.

- ▶ To monitor the optional File System Gateway (FSG), you should implement an SNMP server such as the IBM Director. Call home for the FSG servers is provided through the eSA for System x extension. This extension must be installed on the IBM Director server.
- ▶ To monitor your DR550 Storage Controller (DS4200), you should set up an RSM server.

The various monitoring methods are discussed in Chapter 9, “DR550 call home features” on page 349.

File System Gateway (FSG)

Depending on your availability requirements you can order the FSG as stand-alone or dual-node cluster FSG.

Disk space for archiving

- ▶ Estimate the amount of archive data required now, and estimate the growth for the next few years to size your required disk capacity.
- ▶ The decision of what RAID type will be used for data protection on DR550 Storage Controller (DS4200) has to be made at the time the system is ordered. Currently, there are two types of RAID available: RAID 5 and RAID 6.

Attention: The DR550 Storage Controller (DS4200) can be shipped as RAID 6 or with RAID 5 configured arrays and one spare disk pre-allocated for every array. If you get it shipped as RAID 5, then you would have to unassign the hotspares, add one of the disks to each array, and then you are able to convert to RAID 6. Be aware that it can take a long time for the conversion to complete. In our test lab, it took six days on average to convert one array from RAID 5 to RAID 6.

- ▶ If you want a copy of all archived data, or if you want to store some data at an offsite location, you should consider the implementation of a tape library or DS4000 Enhanced Remote Mirroring (ERM).

Disaster recovery

- If you have special requirements for availability and disaster recovery or if you need remote mirroring, you can use the diagram in Figure 3-1 to discover the possible configurations and options of the various DR550 models.




Figure 3-1 DR550 configuration overview

3.1.1 Planning for the DR550 DR1

DR550 DR1 hardware is shipped pre-assembled, and the software is installed and to a large extent configured. However, some planning is still required before deployment.

This section does not include planning for the optional FSG. This is addressed separately in 3.1.3, "Planning the optional DR550 FSG" on page 71.

The following aspects have to be taken into account prior to the physical installation of the equipment and attaching it to the network:

- Physical site preparation

Adequate site planning before the hardware is installed will reduce the risk of installation issues. Site planning has to cover equipment location specifications, air-conditioning and electrical requirements, raised and non-raised floor determinations, and determination of cable lengths.

Consider the following general planning guidelines:

- The rack space required by the equipment rack:
 - Height of the 7014-S25 IBM 25U rack.
 - Space needed for an external tape library (Tape libraries inside a DR550 rack are not supported in the current version of DR550.)
- The power and environmental requirements
- ▶ Power supply

Two 100 to 127 or 200 to 240 V AC (auto-ranging) are required for the DR550 DR1: one for the p5 52A (600 watts maximum) and another for the console (150 watts maximum).
- ▶ Climate control

Climate control is needed for the DR550 SSAM Server (p52A) and DR550 Storage Controller (DS4200), if attached. For details, refer to Chapter 2, “Physical Planning for the DR550”, of the *DR550 Version 4.5 Introduction and Planning Guide*, GA32-0577.
- ▶ The network environment
 - TCP/IP addresses

You need a unique TCP/IP address to attach the DR550 DR1 to the Ethernet network (because the network address is preconfigured, check for potential address conflicts before attaching to the network).
 - Ethernet switch or switches

You must order the DR550 DR1 with the appropriate Ethernet option for your network environment, that is, copper or fiber cable support. The p5 52A has two on board Gigabit Ethernet ports (copper) or you can order one dual port Gigabit Ethernet-SX PCI-X adapter (fiber cable). Your Ethernet switch or switches must support the Ethernet option you select.
- ▶ Cables and connectors

To attach the DR550 DR1 to the network, you must provide Ethernet cables depending on the Ethernet network option ordered:

 - On board Gigabit Ethernet port (copper)

The port provides 10/100/1000 Mbps connectivity. The port conforms to the IEEE 802.3ab 1000Base-T standard. Use CAT-5 twisted pair bulk cables (we recommend TIA/EIA 568A).
 - Dual-port Gigabit Ethernet-SX PCI-X adapter (fiber optics)

The SX adapter provides 1 Gbps throughput on a standard shortwave (850nm) 50/62.5 micron multimode optical cable and conforms to the IEEE 802.3z standard and supports distances of 260m for 62.5 µm MMF and 550m for 50 µm MMF. You can use an optional LC-SC 62.5 micron converter cable, IBM part number 11P1374, FC2459. For 50 micron LC-SC connections, use converter cable IBM part number 11P1373, FC2456.
 - RSM server 10/100/1000 Mbps Ethernet adapter (copper)

The ports provide 10/100/1000 Mbps connectivity. The port conforms to the IEEE 802.3ab 1000Base-T standard. Use CAT-5 twisted pair bulk cables (we recommend TIA/EIA 568A). RSM server needs at least three Ethernet connections: two connections using CAT-5 crossover cables directly to DS4200 controllers and one normal cable to a customer switch.

Two Fibre Channel Host Bus Adapter ports are pre-installed for tape attachment. There is also an LVD SCSI adapter for tape or optical storage.

- General environment

Security: Default passwords are established for various administrative user IDs (AIX and System Storage Archive Manager) defined for the DR550 DR1. It is critical that you change the default passwords immediately after installation.

- DR550 Storage Controller (DS4200) monitoring with Remote Support Manager (RSM)

If you plan to use RSM to monitor your DR550 Storage Controller (DS4200) or to use the call home facility of RSM to forward hardware problems with the DS4200 automatically to IBM, you must provide the following equipment:

- An IBM System x server (low end x servers, for example x306m or x3250)
- SUSE Linux (SLES 10SP1)
- Generic 56 Kbps V.90/V.92 external modem (for call home function)
- Analog phone line (for call home function)
- Ethernet cables to connect RSM server to DS4200 controllers and customer switch

For more information about RSM, refer to the Chapter 9, “DR550 call home features” on page 349.

- SSAM administration

To manage SSAM using a graphical user interface, you can use the Administration Center and the Integrated Solution Console (ISC). Both software components are preinstalled on the DR550 engines. Alternatively, you can install and run the Administration Center and the ISC on a separate, customer provided workstation, which can be an IBM System x with Linux or Windows as the operating system. Deploying ISC and Administration Center on a separate server can be beneficial from the memory and CPU utilization point of view. Segregating the workloads will allow you to allocate more resources for SSAM application. See 5.1.1, “IBM System Storage Archive Manager architecture overview” on page 158 for more information.

- Requirements for data retention

The DR550 gives you a wide range of options to define the retention criteria and retention period. You should consider the data retention requirements for the different application servers and type of data that you will archive in your environment. This will facilitate the definition of the SSAM policies. Refer to Chapter 5, “IBM System Storage Archive Manager” on page 157 for more information about retention policies.

- Security features

The DR550 provides some features like data shredding, encrypted data transfer, or tape drive encryption to comply with security requirements and regulations compliance. Refer to Chapter 5, “IBM System Storage Archive Manager” on page 157 for more information about these features.

3.1.2 Planning for the DR550 DR2

Although the DR550 hardware and software elements are already assembled, and the software is installed and configured, some planning is still required before deployment.

This section does not include planning for the optional FSG. This is addressed separately in 3.1.3, “Planning the optional DR550 FSG” on page 71.

Remember that any configuration that includes more than three DR550 Expansion Drawers (EXP420) storage expansion enclosures consists of two racks. Before physically installing the racks in your environment and attaching to the network, plan or validate the following:

► Physical site preparation

Adequate site planning before the hardware is installed will reduce the risk of installation issues. Site planning has to cover equipment location specifications, air-conditioning and electrical requirements, raised and non-raised floor determinations, and determination of cable lengths.

Refer to the IBM publication *IBM eServer 7014 Series Model T00 and Model T42 System Rack Installation Guide*, SA38-0641 for more information. You can obtain a copy at:

http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base/hardware_docs/pdf/380641.pdf

Consider the following general planning guidelines:

- The specifications and requirements of the rack or racks.
- The size of the floor area required by the equipment:
 - Floor-load capacity.
 - Space needed for expansion.
 - Location of columns.
- The power and environmental requirements.
- Create a floor plan to check for clearance problems:
 - Make a full-scale template (if necessary) of the rack and carry it along the access route to check for potential clearance problems through doorways and passage ways, around corners, and in elevators.
 - Provide space for storage cabinets, card files, desks, communication facilities, daily storage of tapes, and other supplies.

Information about these topics and detailed instructions for positioning, leveling, and powering up the rack can be found in *IBM eServer 7014 Series Model T00 and Model T42 System Rack Installation Guide*, SA38-0641.

► Power supply

Multiple 200 volts or 220 volts and 30 Amp power service feeds are required:

- Two feeds for single rack configurations (includes single and dual-node).
- Four feeds for two rack configurations.

We also recommend that you split (half and half) the feeds on two distinct AC power sources (distinct circuit breakers) for availability.

► Climate control

Climate control is needed for the entire DR550. For details, refer to Chapter 2 “Physical Planning for the DR550”, of the *DR550 Version 4.5 Introduction and Planning Guide*, GA32-0577.

► The network environment

- TCP/IP addresses.

You need a unique TCP/IP address (single node, no FSG) or addresses to attach the DR550 DR2 to the Ethernet network (because the network addresses are already set at the factory for each of the p5 52A nodes, check for potential address conflicts before attaching to the network).

- Ethernet switch or switches.

You must order the DR550 with the appropriate Ethernet adapter for your network environment, that is, copper or fiber cable support. You can order one 2-port Gigabit Ethernet-TX PCI-X adapter (copper, FC 5706) or one 2-port Gigabit Ethernet-SX PCI-X adapter (fiber optics, FC 5707). Your Ethernet switch or switches must support the Ethernet adapter selected.

► Cables and connectors

To attach the DR550 DR2 to the network, you must provide Ethernet cables depending on the Ethernet network adapter of your DR550:

- Two-port Gigabit Ethernet-TX PCI-X adapter (copper, FC 5706 verify).

The adapter presents one electrical load but appears as two independent devices to the software. It provides 10/100/1000 Mbps connectivity over four pairs of standard CAT-5 cable up to 100m for each port. The adapter conforms to the IEEE 802.3ab 1000Base-T standard. Use CAT-5 twisted pair bulk cables (we recommend TIA/EIA 568A).

- Two-port Gigabit Ethernet-SX PCI-X adapter (fiber optics, FC 5707 verify).

The adapter presents one electrical load but appears as two independent devices to the software. Dual-SX provides 1000 Mbps throughput on a standard shortwave (850nm) 50/62.5 micron multimode optical cable and conforms to the IEEE 802.3z standard and supports distances of 260 m for 62.5 µm MMF and 550 m for 50 µm MMF. You can use an optional LC-SC 62.5 micron converter cable, IBM part number 11P1374, FC2459. For 50 micron LC-SC connections, use converter cable IBM part number 11P1373, FC2456.

- RSM server 10/100/1000 Mbps Ethernet adapter (copper).

The ports provide 10/100/1000 Mbps connectivity. The port conforms to the IEEE 802.3ab 1000Base-T standard. Use CAT-5 twisted pair bulk cables (we recommend TIA/EIA 568A). The RSM server needs at least three Ethernet connections: two connections to internal DR550 Ethernet switches, one to each switch, and one to a customer switch.

Two Fibre Channel Host Bus Adapter ports are pre-installed for tape attachment. For more information about tape attachment, see Chapter 11, “Tape attachment” on page 447.

► General environment

Security: Default passwords are established for various administrative user IDs (AIX, System Storage Archive Manager, DS4000 Storage Manager, and DR550 SAN switches) defined for the DR550. It is critical that you change the default passwords immediately after installation.

► DR550 Storage Controller (DS4200) monitoring with Remote Support Manager (RSM)

If you plan to use the RSM software to monitor your DR550 Storage Controller (DS4200) or to use the call home facility of RSM to forward hardware problems with the DS4200 automatically to IBM, you must provide the following equipment:

- An IBM System x server (low end System x servers, for example, x306m or x3250).
- SUSE Linux (SLES 10 SP1).
- Generic 56Kbps V.90/V.92 external modem (for call home function).
- Analog phone line (for call home function).
- Ethernet cables to connect RSM server to internal DR550 Ethernet switches and to a customer switch.

For more information about RSM, refer to Chapter 9, “DR550 call home features” on page 349.

- ▶ **SSAM administration**

To administrate the SSAM using a graphical user interface, you can use the Administration Center and the Integrated Solution Console (ISC). Both software components are preinstalled on the DR550 engines. Alternatively, you can install and run the Administration Center and the ISC on a separate, customer provided workstation, which can be an IBM System x with Linux or Windows as the operating system. Deploying ISC and Administration Center on a separate server can be beneficial from the memory and CPU utilization point of view. Segregating the workloads allows you to allocate more resources for the SSAM application. See 5.1.1, “IBM System Storage Archive Manager architecture overview” on page 158 for more information.

- ▶ **Requirements for data retention**

The DR550 gives you a wide range of possibilities to define the retention criteria and retention period. You should consider the data retention requirements for the different application servers and type of data that you will archive in your environment. This will facilitate the definition of the SSAM policies. Refer to Chapter 5, “IBM System Storage Archive Manager” on page 157 for more information about retention policies.

- ▶ **Security features**

The DR550 provides some features like data shredding, encrypted data transfer, or tape drive encryption to comply with security requirements and regulation for retention. Refer to Chapter 5, “IBM System Storage Archive Manager” on page 157 for more information about these features.

3.1.3 Planning the optional DR550 FSG

The FSG hardware is already assembled and placed in the rack (when the feature is selected as part of the DR550 initial order). The software (operating system and FSG application) is installed. Some planning is still required before deployment:

- ▶ **Network environment**

- **TCP/IP addresses for stand-alone FSG configurations**

You need an available TCP/IP address to attach the FSG to the customer Ethernet network (because the network addresses are already set at the factory for the FSG node, check for potential address conflicts before attaching to the network). This IP address is bonded to two physical network interfaces in the FSG node.

- **TCP/IP addresses for dual-node FSG configurations**

You need three available TCP/IP addresses to attach the FSG to the customer Ethernet network (because the network addresses are already set at the factory for the FSG node, check for potential address conflicts before attaching to the network). Two IP addresses are bonded to two physical network interfaces in the FSG. The third IP address is used for cluster services.

- **Network Time Protocol (NTP) Services for dual-node configurations**

To synchronize the clocks in the FSG cluster, the nodes need access to a time server through the NTP protocol.

- Cables and connectors

To attach the FSG to the customer network, you must provide two Ethernet cables per node depending on the type of Ethernet network adapter of your FSG:

- 10/100/1000 Mbps TX Ethernet adapter (copper)

It provides 10/100/1000 Mbps connectivity to the network. Use CAT-5 twisted pair bulk cables (we recommend TIA/EIA 568A).

- 1 Gbps SX Ethernet adapter (fibre optics)

The 1 Gbps-SX adapter provides 1000 Mbps throughput on a standard shortwave (850nm) 50/62.5 micron multimode optical cable and conforms to the IEEE 802.3z standard and supports distances of 260 m for 62.5 µm MMF and 550m for 50 µm MMF.

- File system structure at the FSG

Data written to the FSG cannot be deleted until the specified retention period has passed. The expiration period depends on the definition in the DRG profile. This applies to files and directories. Therefore, it is important to plan the file system layout, naming space structure, file and directory naming conventions, and the necessary access permissions for all CIFS and NFS shares of the FSG before deployment.

- User Management for CIFS and NFS shares

If you plan to integrate FSG managed NFS and CIFS shares into an environment with centralized user management utilizing tools like Active Directory® or Lightweight Directory Access Protocol (LDAP), you need to do an extensive planning of the required user and group authorities prior to the deployment.

3.2 Physical installation

The DR550 DR1 comes pre-installed in a single *IBM 7014 Series S25 System Rack*.

The 0.88 TB capacity option does not require the DR550 Storage Controller (DS4200). DR550 DR1 model is pre-configured with 6 x 146 GB internal drives installed in the DR550 SSAM Server (p52A) disk bays. Optional DR550 Storage Controller (DS4200) adds 6 TB of raw capacity and up to 36 TB of combined raw capacity with two DR550 Expansion Drawers (EXP420).

3.2.1 Installing the DR550 DR1

The following applies to the DR550 DR1. This section does not include planning for the optional FSG, which is covered separately in 3.2.3, “Installing the optional DR550 File System Gateway (FSG)” on page 76.

Power supply connections

For DR550 DR1, each device in the rack is connected to the AC power rails inside the rack. Sets of AC power rails (left and right) are placed vertically in the rear of the rack cabinet. Each rail in a set should be connected to a different AC power feed to enhance the availability of the rack components.

Fibre Channel connections

The cabling of the internal devices in the DR550 DR1 is done before shipment to clients. DR550 DR1 is not equipped with a DR550 SAN switch and therefore all fiber attachments are done directly to DR550 SSAM Server (p52A). When attaching tape devices to the DR550 DR1 as recommended, you need to establish additional Fibre Channel connections between

one or two Fibre Channel Host Bus Adapter ports on the DR550 SSAM Server (p52A) and the tape devices. Alternatively, you can attach a tape device using the LVD SCSI adapter installed in slot 3. For more information about tape attachment, see Chapter 11, “Tape attachment” on page 447.

Network connections

The DR550 DR1 can be connected to a 10/100/1000 Base Ethernet using copper or fiber cables. You have to provide the appropriate network cables and switches (copper or fiber, as described in 3.1.1, “Planning for the DR550 DR1” on page 66).

Refer to 3.6.2, “Attaching the DR550 DR1 to the network” on page 91 for instructions about how to connect the network cables according to your environment. Do not connect to the network before you have completed the tasks described in 3.3, “Power-on and power-off sequence” on page 77 and 3.4, “Accessing the DR550 SSAM Servers” on page 81.

3.2.2 Installing the DR550 DR2

The following applies to the DR550 DR2 single-server and dual-server models. This section does not include planning for the optional FSG, which is covered separately in 3.2.3, “Installing the optional DR550 File System Gateway (FSG)” on page 76.

Power supply connections

When shipped from the factory, each rack has its AC power switch set to the OFF position. The optional DR550 SAN Switch (2005-B16) does not have an AC power switch; it is turned on or off by plugging the power cable in or out.

Each device in the rack or racks is connected to the AC power rails inside the rack or racks. Sets of AC power rails (left and right) are placed vertically in the rear of the rack cabinet. Each rail in a set should be connected to a different AC power feed to enhance the availability of the rack components. Power cords are shipped with each rack (one per rail) to enable cabling to AC power above or below each rack.

You have to connect the AC power rails installed in the rack or racks to an appropriate power outlet (220 volts or 200V rated at 30 amps).

Fibre Channel connections

All cables between the DR550 Storage Controller (DS4200) or Servers and the DR550 Expansion Drawer (EXP420) expansion units (if applicable) are set in place by manufacturing. Furthermore, the cables from the DR550 SSAM Server or servers to the DR550 SAN switch or switches and from the DR550 SAN switch to the DS4200 Storage Server or Servers are installed.

The cabling of internal devices in the DR550 DR2 is done before shipping to clients. The DR550 DR2 in a single-server configuration can be ordered without a SAN switch, in which case the wiring is done similar to DR550 DR1 with a DR550 Storage Controller (DS4200) attached directly to the DR550 SSAM Server (p52A). A single SAN switch will be added to the configuration if Enhanced Remote Mirroring or a second DR550 Storage Controller (DS4200) is ordered. Two DR550 SAN switches will be included in the dual-server DR550 DR2 configuration regardless of whether Enhanced Remote Mirroring or a second DR550 Storage Controller (DS4200) is ordered.

However, there are two types of situations where you need to establish your external Fibre Channel connections:

- Initial setup of the configuration that requires more than three DR550 Expansion Drawers (EXP420).

The first rack can hold up to three DR550 Expansion Drawers (EXP420). If there is a need to install more expansion enclosures, all additional units will have to be installed in a second rack. The cabling between the first and the second racks is accomplished with two LC-LC Fibre Channel cables provided with the DR550.

The second optional DR550 Storage Controller (DS4200) is also installed in the second rack. Fiber cables between the second DR550 Storage Controller and expansion enclosures is done at the factory; however, the cabling between the second DR550 Storage Controller and DR550 SAN switch(es) will have to be done at the time of the installation.

Refer to the diagram in Figure 3-2 on page 75 to establish the Fibre Channel connections between the two racks during the initial setup of the configuration that requires more than three DR550 Expansion Drawers (EXP420).

The Fibre Channel cables need to be connected as indicated in Table 3-1.

Table 3-1 The two racks configuration

What	From		To	
Fibre Channel cable	Rack 2	DR550 Storage Controller (DS4200) #2 Ctrl A	Rack 1	2005-B16 #1 Port 8
Fibre Channel cable	Rack 2	DR550 Storage Controller (DS4200) #2 Ctrl B	Rack 1	2005-B16 #2 Port 8

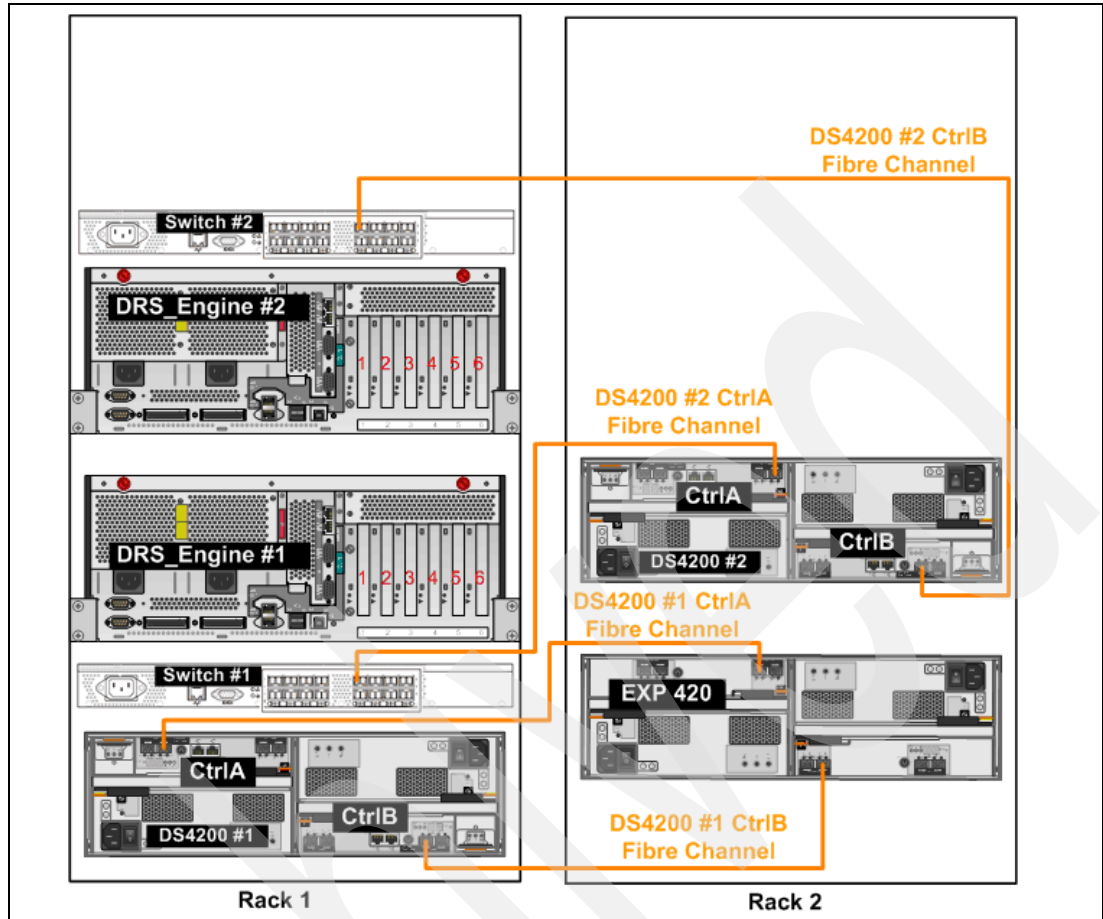


Figure 3-2 Fiber connections required between racks for two rack configuration

The procedure is described in detail in 2.5, “What is preconfigured in the DR550” on page 29. Note that if you have additional EXP420s in the second rack, there will be additional crossover cables between the racks.

► Tape attachment.

You need to establish additional Fibre Channel connections when you are attaching tape devices to the DR550 DR2. Tape drives must be cabled to the DR550 SAN switch or switches of the DR550. For more information about tape attachment, see Chapter 11, “Tape attachment” on page 447.

Network connections

The DR550 DR2 can be connected to a 10/100/1000 Base-T Ethernet using copper or to a 1000BASE-SX using fiber cables. You have to provide the appropriate network cables and switches (copper or fiber, as described in 3.1.2, “Planning for the DR550 DR2” on page 68).

Refer to 3.6.3, “Attaching the DR550 DR2 to the IP network” on page 93 for instructions about how to connect the network cables according to your environment. Do not connect to the network before you have completed the tasks described in 3.3, “Power-on and power-off sequence” on page 77 and 3.4, “Accessing the DR550 SSAM Servers” on page 81.

3.2.3 Installing the optional DR550 File System Gateway (FSG)

The following applies to the FSG stand-alone and dual-node configurations.

Power supply connections

Each FSG is equipped with two power supplies and they should be connected to different power rails in the rack to enhance the availability of the FSG. Note that the DR1 has only one power rail; in this case, a bifurcated power cable is used.

Network connections

The FSG can be connected to a 10/100/1000 Base-TX Ethernet using copper or to a 1000 BASE-SX Ethernet using fiber cables. Optional fiber Ethernet adapters are installed in PCI slot 1 and 2 and will be used for attachment to the customer fiber Ethernet network. Dual port Gigabit Ethernet over copper adapters are installed in PCI slots 3 and 4. One port on each adapter is used for communicating with DR550 SSAM servers either using internal switches (DR550 DR2 only) or using a customer supplied external Ethernet switch. You have to provide the appropriate network cables and switches (copper or fiber, as described in 3.1.3, “Planning the optional DR550 FSG” on page 71

Dual-node network connections

Two additional Ethernet connections are required in the dual-node FSG configuration. The purpose for these connections is to maintain a heartbeat between the two FSG nodes. Two crossover Ethernet copper cables are connected between the on board Ethernet adapters of the FSG. The DR550 ships with pre-installed Ethernet cables for the dual node FSG cluster connections. Refer to the diagram in Figure 3-3.

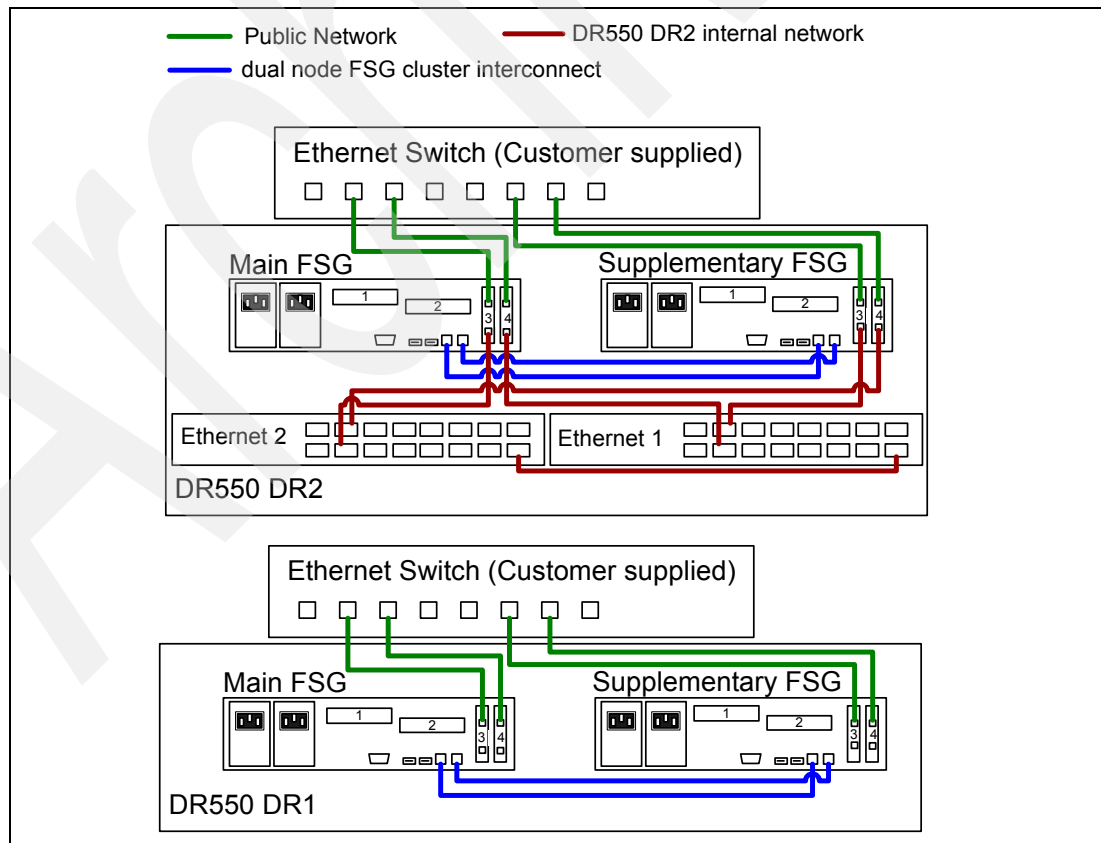


Figure 3-3 Network cabling for a dual-node FSG configuration DR2 and DR1 models

3.3 Power-on and power-off sequence

Powering on all the devices in the DR550 makes the system available for use. For the devices, different programs, and processes running on these devices to function properly, it is necessary to power on the devices in a specific order. This is true not only for the very first boot of the DR550, but also for all consecutive power-on resets of the system.

3.3.1 Power-on sequence

Refer to the section that applies to your model.

DR550 Model DR1

The devices in the rack should be powered on in the following sequence:

1. DR550 DR1 Rack
2. KVM console kit (if applicable)
3. DR550 Expansion Drawer (EXP420) (if applicable)
4. DR550 Storage Controller (DS4200)
5. DR550 SSAM Server (System p5)
6. Optional Main-FSG (if applicable)
7. Optional Supplementary FSG (if applicable)

For DR550 DR2, the sequence depends on the configuration, single-node, or dual-node.

DR550 Model DR2 single-node configuration

For the single-node configuration, the devices in the rack must be powered on in the following sequence:

1. DR550 DR2 Rack / 16-port Ethernet switches.
2. KVM console kit (if applicable).
3. DR550 SAN Switch (2005-B16). Before proceeding to step 3, confirm that the SAN Switch has completed its power-on sequencing. This is indicated with a green light showing for power.
4. DR550 Expansion Drawer (EXP420) (if applicable).
5. DR550 Storage Controller (DS4200).
6. DR550 SSAM Server (System p5).
7. Optional Main-FSG (if applicable).
8. Optional Supplementary-FSG (if applicable).

The detailed procedure is as follows:

1. DR550 DR2 Rack (7014) / 16-port Ethernet switches. Each device in the rack is connected to the AC Power rails. When shipped from the factory, each component will have the AC power switch in the OFF position. The 16-port Ethernet switches do not have an AC switch, and thus will be powered up as soon as the rails become energized.
2. KVM console kit (if applicable). Power on the KVM console kit by pressing the power button.
3. DR550 SAN Switch 2005-B16. The SAN Switches do not have an AC switch, and thus will be powered up as soon as the rails become energized.
4. DR550 Expansion Drawers (EXP420) (if applicable). Power on the expansion units using the two black AC power switches on the back of the devices.

5. DR550 Storage Controller (DS4200). As soon as the SAN switch and the EXP420, if applicable, have completed their power-on sequences, power on the DR550 Storage Controller (DS4200) by using the two black AC power switches on the back of the device.
6. DR550 SSAM Server (System p5). Power on the DR550 SSAM Server with the power button located on the operator panel. A solid green light on the operator panel of the DR550 SSAM Server is an indication of a power-on status. Use the system console to monitor the AIX startup process. The AIX startup process is complete when you see a login prompt. Please note that Main FSG can only be started when the SSAM application is operational. The SSAM startup can take several minutes. Before starting the Main FSG, make sure that SSAM application is up and running (you can access the SSAM server through the dsmadm administrative interface).
7. DR550 Main FSG (if applicable). Power on the FSG with the white power button on the front of the FSG. The Main FSG is always the lower one in a dual-node FSG configuration.
8. DR550 Supplementary FSG (if applicable). The boot process of the Main FSG must be completed before powering on the Supplementary FSG. Power on the FSG with the white power button on the front of the FSG. The Supplementary FSG is always the upper one in a dual-node FSG configuration.

DR550 Model DR2 dual-node configuration

For the dual-node configuration, the devices in the rack or racks should be powered on in the following sequence:

1. DR550 DR2 Rack(s) / 16-port Ethernet switches
2. KVM console kit
3. DR550 SAN Switches (2005-B16)
4. DR550 Expansion Drawers (EXP420) (if applicable)
5. DR550 Storage Controller (DS4200) – 1
6. DR550 Storage Controller (DS4200) – 2 (if applicable)
7. DR550 SSAM Server 1 (System p5 – 1)
8. DR550 SSAM Server 2 (System p5 – 2)
9. DR550 Main-FSG (if applicable)
10. DR550 Supplementary-FSG (if applicable)

The detailed procedure is as follows:

1. DR550 DR2 Rack (7014) / 16-port Ethernet switches. Each device in the rack is connected to the AC power rails. When shipped from the factory, each component will have the AC power switch in the OFF position. 16-port Ethernet switches do not have an AC switch, and thus will be powered up as soon as the rails become energized.
2. KVM console kit (if applicable). Power on the KVM console kit by pressing the power button.
3. IBM TotalStorage® SAN Switch 2005-B16. (Before proceeding to step 3, confirm that the DR550 SAN Switch has completed its power-on sequence. This is indicated with a green light showing for power.)
4. DR550 Expansion Drawers (EXP420) (if applicable). Power on the expansion units using the two black AC power switches on the back of the devices.
5. DR550 Storage Controller (DS4200). As soon as the DR550 SAN switch and the DR550 Expansion Drawers (EXP420) (if applicable) have completed their power-on sequences, power on the DR550 Storage Controller by using the two black AC power switches on the back of the device.
6. DR550 SSAM Server 1 (System p5-1). Power on the DR550 SSAM Server with the power button located on the operator panel. A solid green light on the operator panel of the DR550 SSAM Server is an indication of a power on status. Use the system console to

monitor the AIX startup process. The AIX startup process is complete when you see a login prompt. Please note that Main FSG can only be started when the SSAM application is operational. The SSAM startup can take several minutes. Before starting the Main FSG, make sure that the SSAM application is up and running (you can access the SSAM server through the dsmadmc administrative interface).

7. DR550 SSAM Server 2 (System p5-2). Power on the DR550 SSAM Server with the power button located on the operator panel. A solid green light on the operator panel of the DR550 SSAM Server is an indication of a power-on status. Use the system console to monitor the AIX startup process. The AIX startup process is complete when you see a login prompt.
8. DR550 Main FSG (if applicable). Power on the FSG with the white power button on the front of the FSG. The Main FSG is always the lower one in a dual-node FSG configuration.
9. DR550 Supplementary FSG (if applicable). The Boot process of the Main FSG must be completed before you power on the Supplementary FSG. Power on the FSG with the white power button on the front of the FSG. The Supplementary FSG is always the upper one in a dual-node FSG configuration.

3.3.2 Power-off sequence

Refer to the section that applies to your model.

DR550 Model DR1

The devices in the rack should be powered off in the following sequence:

1. DR550 Supplementary-FSG (if applicable)
2. DR550 Main-FSG (if applicable)
3. DR550 SSAM Server (System p5)
4. DR550 Storage Controller (DS4200)
5. DR550 SAN Switch (2005-B16)
6. KVM console kit (if applicable)
7. DR550 DR1 Rack

Refer to “DR550 Model DR2 single-node configuration” on page 79 for the detailed shutdown and power-off procedure.

DR550 Model DR2 single-node configuration

For the single-node configuration, the devices in the rack must be powered off in the following sequence:

1. DR550 Supplementary-FSG (if applicable)
2. DR550 Main-FSG (if applicable)
3. DR550 SSAM Server (System p5)
4. DR550 Storage Controller (DS4200)
5. DR550 Expansion Drawer (EXP420) (if applicable)
6. KVM console kit
7. DR550 SAN Switch (2005-B16)
8. DR550 DR2 Rack / 16-port Ethernet switches

Before shutting down your system, first halt the SSAM server. The next step is to power off the DR550 SSAM Server, placing it in standby mode (the AIX operating system is shutdown). Before removing power from the DR550 rack power distribution units, ensure that the AIX shutdown process is complete. Failure to do so can result in the loss of data. The following sequence is used to power off the devices in the rack *after* all the applications are stopped:

1. DR550 Supplementary-FSG (if applicable). Shut down the FSG node with the root user by running the Linux **shutdown -h now** command.
2. DR550 Main-FSG (if applicable). Shut down the FSG node with the root user by running the Linux **shutdown -h now** command.
3. DR550 SSAM Server. Shut down the operating system of the DR550 SSAM Server as the root user by running the AIX **shutdown -F** command. Under certain conditions, it might be necessary to force a system power down by pressing power button on the control panel (in a situation when AIX becomes unresponsive, for example). This procedure can be used only after all attempts to shut down AIX gracefully have failed. The indication of a power down state of DR550 SSAM Server is a flashing green light on the control panel.
4. DR550 Storage Controller (DS4200). Power off the device by using the two black AC power switches on the back of the device.
5. DR550 Expansion Drawer (EXP420). Power off the expansion unit (if applicable) using the two black AC power switches on the back of each device.
6. DR550 SAN Switch (2005-B16). To turn the unit off, unplug the power cord at the front of the unit. Alternatively, the switch will be powered off when the power is removed from the rack power distribution units.
7. KVM console kit (if applicable). Power off the KVM console kit by pressing the power button.
8. DR550 DR2 Rack (7014) / 16-port Ethernet switches. The 16-port Ethernet switches do not have an AC switch, and thus will be powered down as soon as the power is removed from the rack power distribution units.

We recommend that individual devices are not powered off as though they were stand-alone devices unless you have a good understanding of the interdependency issues. For example, you should not power off the SAN Switches (in order to reset them) unless you know that this will not affect ongoing data flow.

DR550 Model DR2 dual-node configuration

For the dual-node configuration, the devices in the rack must be powered off in the following sequence:

1. DR550 Supplementary-FSG (if applicable)
2. DR550 Main-FSG (if applicable)
3. DR550 SSAM Server 1 (System p5 – 2)
4. DR550 SSAM Server 2 (System p5 – 1)
5. DR550 Storage Controller (DS4200) – 2 (if applicable)
6. DR550 Storage Controller (DS4200) – 1
7. DR550 Expansion Drawers (EXP420)
8. DR550 SAN Switches (2005-B16)
9. KVM console kit (if applicable)
10. DR550 DR2 Rack / 16-port Ethernet switches

Before shutting down your system, stop HACMP on both cluster nodes. An HACMP stop on an active cluster node has to be performed without initiating HACMP takeover, otherwise the SSAM application will be re-started on a standby cluster node.

Note: Do not halt the SSAM server manually. The SSAM server process is being constantly monitored by the HACMP application monitor and if HACMP detects that the SSAM server process is down, it will automatically initiate its restart. The SSAM server process will be shut down by HACMP when it is stopped.

The next step is to power off both DR550 SSAM Servers, placing them in standby mode (the AIX operating system is shut down). Before removing power from the DR550, ensure that the AIX shutdown process is complete. Failure to do so can result in a loss of data.

The following sequence is used to power off the devices in the racks *after* all the applications are stopped:

1. DR550 Supplementary-FSG (if applicable). Shut down the FSG node with the root user by running the Linux **shutdown -h now** command.
2. DR550 Main-FSG (if applicable). Shut down the FSG node with the root user by running the Linux **shutdown -h now** command.
3. DR550 SSAM Server 2 (System p5 - 2). Shut down the operating system of the DR550 SSAM Server #2 as the root user by running the AIX **shutdown -F** command. Under certain conditions it might be necessary to force a system power down by pressing the power button on the operator panel (in the situation when AIX became unresponsive, for example). This procedure can be used only after all attempts to shut down AIX gracefully have failed. The indication of a power down state of the DR550 SSAM Server is a flashing green light on operator panel.
4. DR550 SSAM Server 1 (System p5 - 1). Use the same power-off procedure as for DR550 SSAM Server 2.
5. DR550 Storage Controller (DS4200). Power off all DR550 Storage Controller (DS4200) by using the two black AC power switches on the back of each device.
6. DR550 Expansion Drawer (EXP420). Power off all expansion units using the two black AC power switches on the back of each enclosure.
7. DR550 SAN Switches (2005-B16). To turn the unit off, unplug the power cord at the front of the unit. Alternatively, the switch will be powered off when the power is removed from the rack power distribution units.
8. KVM console kit. Power off the KVM console kit by pressing the power button.
9. DR550 DR2 Rack (7014) / 16-port Ethernet switches. The 16-port Ethernet switches do not have an AC switch, and thus will be powered down as soon as the power is removed from the rack power distribution units.

We recommend that individual devices are not powered off as though they were stand-alone devices unless you have a good understanding of the interdependency issues. For example, you should not power off the DR550 SAN Switches (in order to reset them) unless you know that this will not affect ongoing data flow.

3.4 Accessing the DR550 SSAM Servers

The DR550 SSAM Servers (also referred to as the p5 nodes) can be accessed using the following methods:

- ▶ Using the keyboard-video-mouse (KVM) console kit and switch or system console
- ▶ Using IBM Director (optional)
- ▶ Using the local area network (LAN)

The KVM console kit and switch are included in any DR550 configuration that has either more than a single System DR550 SSAM Server or an FSG server.

From the DR550 console keyboard, you can also start a management session for the DR550 SSAM Servers. The system console supports both text and graphical modes for accessing AIX.

The local area network can be used to access AIX on the DR550 SSAM Servers or applications running on the servers with an application specific access method. For example, the Secure Hyper Text Transfer Protocol (HTTPS) can be used to manage SSAM server remotely. Also, the LAN session can be used to receive graphical output from the operating system running on DR550 SSAM Servers.

An explanation for the access methods listed above is provided in the following sections.

3.4.1 Accessing the DR550

To administer the DR550, you can connect (locally) directly to the DR550 SSAM Server or access it remotely over the LAN for any administrative tasks that do not require superuser access privileges.

Using Keyboard Video Mouse (KVM)

The monitor, the keyboard, and the mouse are directly attached to the DR550 SSAM Server in the DR550 DR1 configuration. The KVM console kit and switch are included in the DR550 DR1 configuration when an optional FSG is ordered. The KVM console kit is used to access the DR550 SSAM Server and the FSG node(s) through a single console.

From the console keyboard of the DR550 DR1, you can start a local management session for the DR550 SSAM Server. Using the KVM attached console gives you the same level of access as though you had a dedicated console. The KVM console kit allows you to switch between console sessions of individual servers and maintain the state of all the sessions so when you switch back to a console session of DR550 SSAM Server, it will appear the way it did without having to refresh the screen. You can switch between consoles sessions by pressing the PrtSc key or by pressing the Ctrl key twice. Once a session is selected, press Enter to start the session.

To start a local graphical session, type `startx` at the AIX prompt. To quit the local graphical session again, press `Ctrl+Alt+Backspace`. Because `startx` is an AIX command, you can also use the command for remote graphical sessions. The difference between the local and the remote graphical session is mainly the location where the output is presented. Therefore, when you use the console, it is always a local output on the DR550 monitor. The remote output must be redirected to whatever system and display are able to show the output. It depends on the technical resources and configuration to redirect graphical output. See “Graphical output” on page 83 for more information about remote graphical output.

Using IBM Director

A remote session can be started from IBM Director console by selecting a DR550 SSAM server (p52A) and then selecting **Tasks** → **Remote Session** → **Remote Session** → **AIX host name**. You will be prompted to enter a user name and password. Refer to 9.3.2, “IBM Director ISS Extensions for DR550” on page 361 for details and illustrations. The limitations of remote login using superuser privileges also applies to the sessions initiated through IBM Director. You must use a user ID that is enabled for remote login, such as `dr550adm`. After you enter a user ID and password, you should see the system prompt of the operating system on the AIX server. To close the session, enter the `exit` command and press Enter or use the `Ctrl+D` key combination.

Using the local area network

Administering the DR550 SSAM Server(s) from a remote server (network workstation or administration workstation) is exactly the same for both the DR550 DR1 and the DR550 DR2.

Figure 3-4 shows the different ways to manage the DR550 DR1.

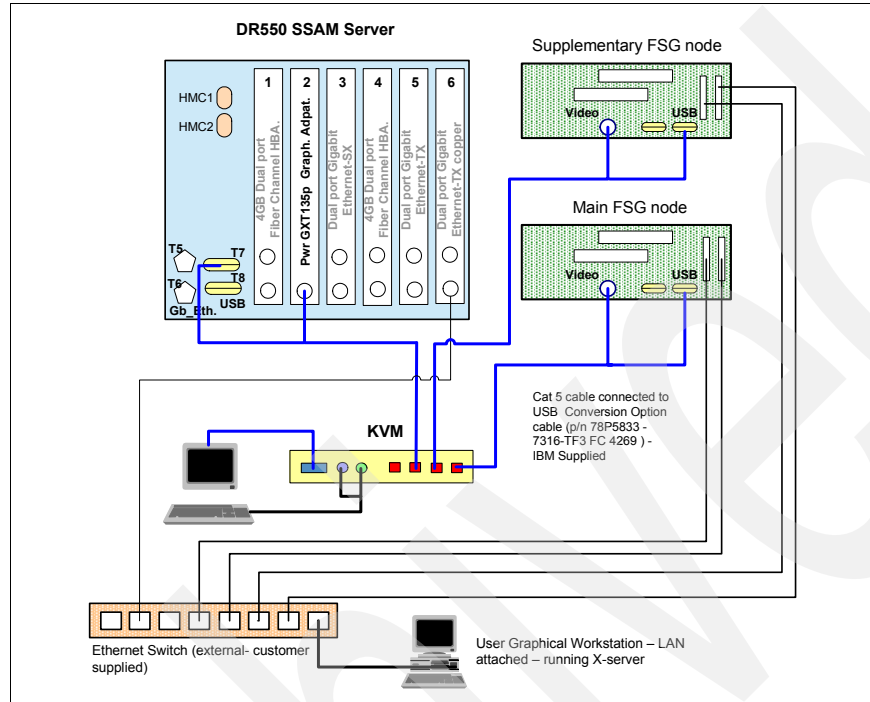


Figure 3-4 Managing the DR550 DR1 (with dual FSG) through KVM or LAN

Command-line output

To use the command line on the DR550 SSAM Servers through the local area network, you need a network workstation running a Secure Shell (SSH) client. For example, on UNIX-based and Linux-based network workstations, the SSH protocol and SSH clients are typically included in the operating system. On a Microsoft® Windows-based network workstation, you can find several SSH clients on the Internet, available either as freeware or for purchase.

A well-known example of an SSH client for Windows is the program named *PuTTY*. This is a free implementation of Telnet and SSH clients for Win32® and UNIX platforms, along with an xterm terminal emulator. For further information about PuTTY, visit:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/>

Graphical output

Note: Remotely displaying the AIX graphical output is technically feasible as described below. However, because of the DR550 restrictions for remote accounts, there is probably no real or practical need for this type of access. We only mention it for completeness.

To use the graphical output of the DR550 SSAM Server(s) through the local area network, you need a network workstation running an X-server. For example, on UNIX-based and Linux-based network workstations, the X-server program is included in the operating system. On a Windows-based network workstation, you can find several X-servers on the Internet.

Cygwin/X is an example of an X-server for Windows. You can find more information about *Cygwin/X* at:

<http://x.cygwin.com/>

After the installation of an X-server program on your Windows-based network workstation, or by using the X-server on a UNIX-based or Linux-based network workstation, you can use one of the following procedures to display graphical output from the DR550.

Procedure 1 (export DISPLAY variable)

1. Connect the network workstation to the same network that the DR550 is connected to.
2. Start the X-server in the network workstation. Depending on your operating system, the X-server might already be running (UNIX or Linux).
3. Allow the AIX operating system of the DR550 SSAM Server to access the X-server in the network workstation. For this to work, the DR550 SSAM Server must be in the list of accepted X-server hosts. Depending on the version of X-server, you do this by entering the appropriate command on your network workstation, or by using the X-server functionality. A common UNIX or Linux command for this task is **xhost**. For example, if you want to accept the DR550 SSAM Server with an address 192.168.1.22 (factory cluster service), you can type **xhost + 192.168.1.22** to accept this host on the X-server.
4. Establish a SSH session to the DR550 SSAM Server from where you need the graphical output. Use a SSH client of your choice to connect to the DR550 SSAM Server and log in to the AIX operating system.
5. Redirect the display of the DR550 SSAM Server to the network workstation. For that purpose, type the AIX command **export DISPLAY=192.168.1.99:0.0**. This command sets the AIX environment variable **DISPLAY** to the address of the administration workstation, in our example, to the address 192.168.1.99; this might be different in your environment.
6. Start the graphical application. For example, run **xclock** to validate your setup.

Procedure 2 (SSH tunnel)

You must change the configuration settings of the SSH daemon to enable the forwarding of the X11 protocol. The configuration steps 1 through 6 must be done only the first time you establish the connection:

1. Log on at the management console as **dr550** and switch to root authority with the **su -** command.
2. Change to the directory **/etc/ssh** with the **cd /etc/ssh** command.
3. Edit the configuration file of the SSH daemon with **vi sshd_config**.
4. Scroll to the line **X11Forwarding** and change the option to **yes**. Remove the **#** in the front of the line. Save the file and exit the editor. (See Example 3-1)
5. Restart the SSH daemon with the commands **stopsrc -s sshd** and **startsrc -s sshd**.

Example 3-1 sshd_config file (excerpt)

```
# Set this to 'yes' to enable PAM keyboard-interactive authentication
# Warning: enabling this may bypass the setting of 'PasswordAuthentication'
#PAMAuthenticationViaKbdInt no
```

X11Forwarding yes

```
#X11DisplayOffset 10
#X11UseLocalhost yes
#XAuthLocation /usr/bin/x11/xauth
```

```
#PrintMotd yes
```

- Now you can establish a SSH session to the DR550 SSAM Server from where you need the graphical output. Use a SSH client of your choice to connect to the DR550 SSAM Server and log in to the AIX operating system. Make sure that X11 forwarding is enabled in your SSH client configuration. If you use the PuTTY SSH client, you will find this at **Connection** → **SSH** → **Tunnels**. Check the **Enable X11 forwarding** box and save the settings in your session profile (see Figure 3-5).
- Log on from your remote workstation with the user id dr550adm to the DR550 and start the graphical application. For example, start the **xclock**.

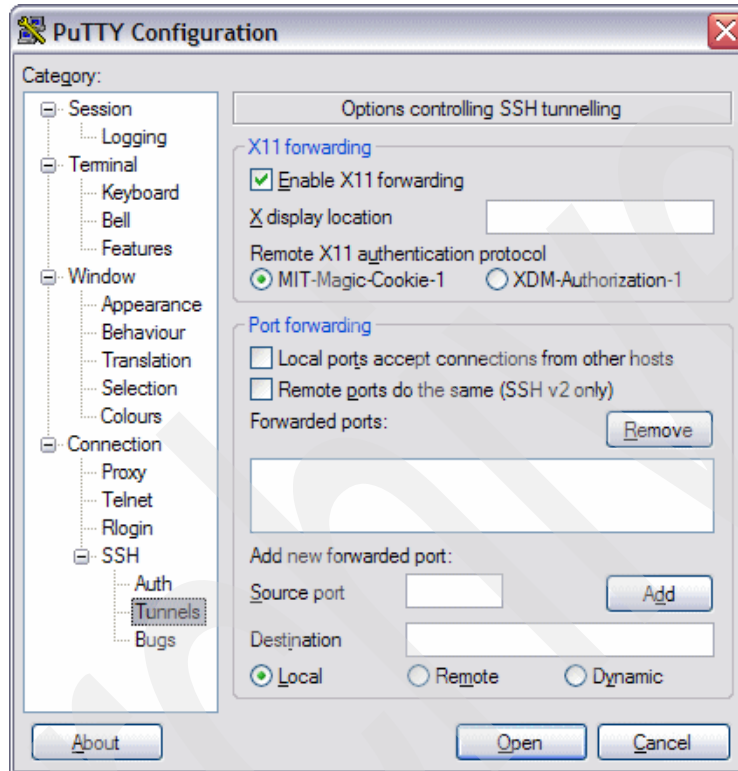


Figure 3-5 Configuration X11 forwarding in PuTTY

Figure 3-6 illustrates the different options available to manage the DR550.

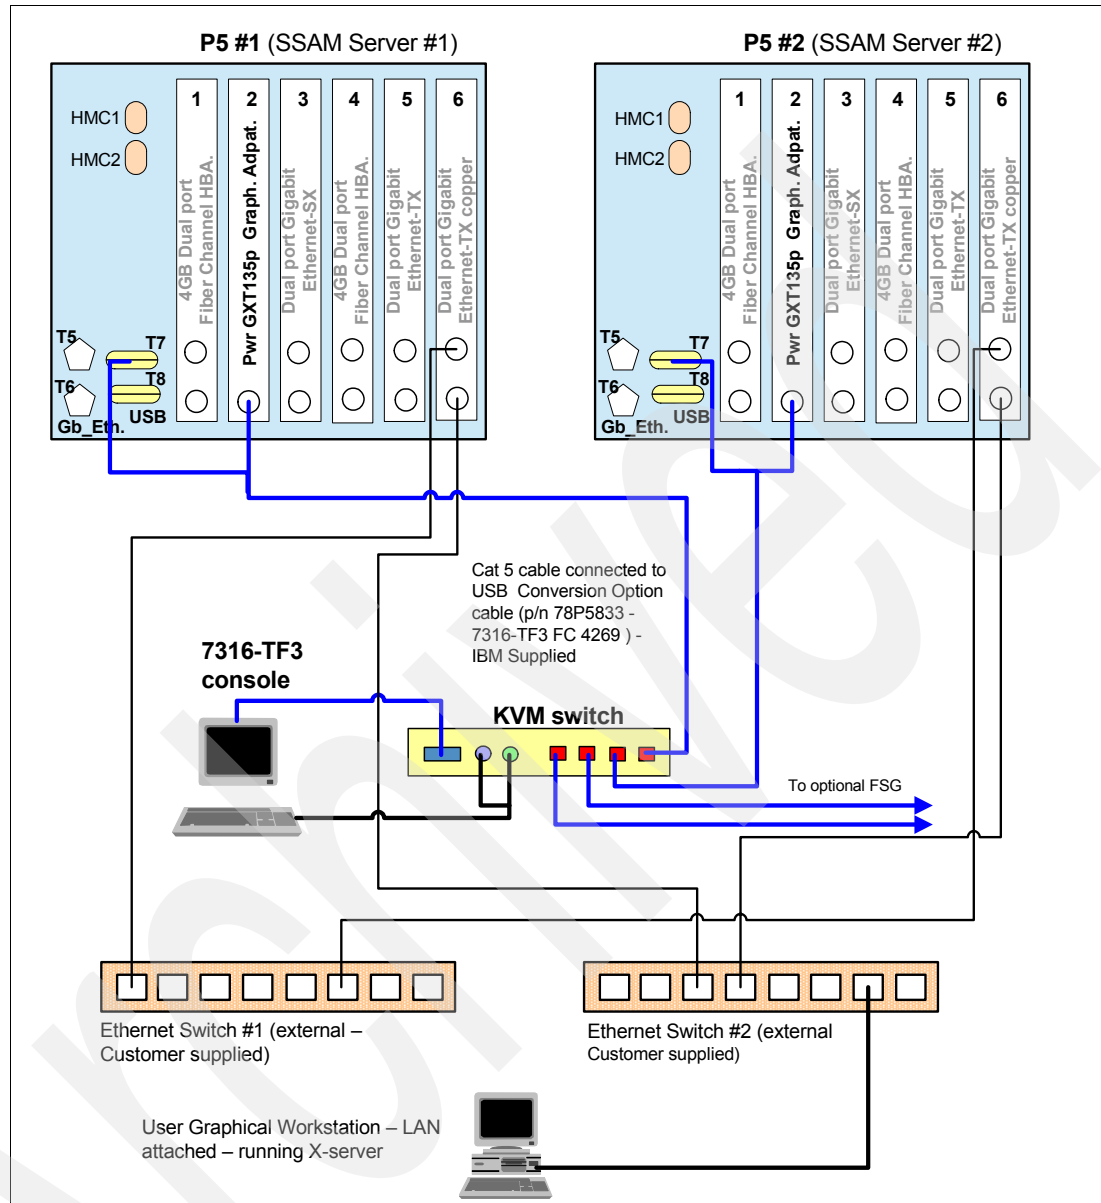


Figure 3-6 Managing the DR550 through KVM, or LAN

3.5 Accessing the FSG nodes

To administer the DR550 File System Gateways, you can connect (locally) directly to the FSG node through the KVM or access the node remotely over the LAN. Remote LAN access to FSG server is only permitted for non-root users. Switching the account to root is only allowed for local users (there is no remote login capability). To gain superuser (root) privileges on the FSG server, you have to use the system console to log in as a local user and then switch your account to root.

Using the keyboard-video-mouse (KVM) switch

From the console keyboard of the DR550, you can also start a local management console for the FSG nodes. You can start the management console with the **PrtSc** key or by pressing the **Ctrl** key twice, select the appropriate node, and press **Enter**. This gives you the graphical output of the FSG.

Using IBM Director

From an IBM Director console, Start a remote session by selecting the name of an FSG server and then selecting **Tasks** → **Remote Session** → **Remote Session: FSG hostname**. You will be prompted to enter a user name and password. Refer to 9.3.2, “IBM Director ISS Extensions for DR550” on page 361 for details and illustrations. The limitations on remote login using superuser privileges also apply to the sessions initiated through IBM Director. You must use a user ID that is enabled for remote login, such as **fsgadm**. After you enter user ID and password, you should see the system prompt of the operating system on the FSG server. To close the session, enter the **exit** command and press **Enter** or use the **Ctrl+D** key combination.

Using the Local Area Network

To administer the FSG node or nodes from a remote server or network workstation, you can use the Local Area Network (LAN) as the communication path. After you have set up the IP address or addresses for the FSG, and the FSG is connected to the network, you can use the connection for command-line or graphical output.

Command-line output

To use the command line on the FSG node or nodes through the local area network, you need a network workstation running a Secure Shell (SSH) client. For example, on UNIX-based and Linux-based network workstations, the SSH protocol and SSH clients are included in the operating system. On a Microsoft Windows-based network workstation, you can find several SSH clients on the Internet, available either as freeware or for purchase.

A well-known example of an SSH client for Windows is a program named *PuTTY*. This is a free implementation of Telnet and SSH for Win32 and UNIX platforms, along with an xterm terminal emulator. For further information about PuTTY, visit:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/>

3.6 Required configuration tasks for the DR550

The DR550 DR1 and DR550 DR2 are preconfigured in manufacturing. Nevertheless, some tasks are still be required before deployment:

- Basic setup

The first task is to set up the system time, time zone, and, if necessary, time zone variables, and then accept the time within the SSAM server (use the **accept time** command).

- Network address configuration

This integrates the solution into existing networks. Without a network configuration that matches your network settings, the DR550 will not be accessible by applications for archiving their data.

- ▶ HACMP network configuration (DR550 DR2 dual-node configuration only)

For the dual-node configuration, customizing IP address settings also involves reconfiguring HACMP due to the fact that the IP address information is included into the HACMP configuration.

- ▶ SSAM and Data retention policy setup

Another crucial task is to set up your SSAM environment and the data retention policies for the archived objects. Data retention policies can vary from very simple to very complex, and so can the setup. The data retention policies for the IBM System Storage DR550 will be configured within the IBM System Storage Archive Manager. Furthermore, you might want to configure additional processes in relation to the new data retention policies, for example, schedules to migrate data or create copies. This can be an ongoing activity as new applications are added or new retention rules are required.

A few additional setup steps are required if you plan to use the Enhanced Remote Mirroring (ERM) capability on the DR550 (see 13.7, “Enhanced Remote Mirroring on DR550: Step-by-step” on page 538).

In the sections that follow, we explain how to configure the network and reconfigure HACMP (dual-node configuration only). Configuring IBM Tivoli Storage Manager client and IBM System Storage Archive Manager server will be described separately in Chapter 5, “IBM System Storage Archive Manager” on page 157.

3.6.1 Basic setup

This basic setup is identical for all configurations (DR550 DR1 and DR550 DR2 single or dual-node).

Changing the system time

This section explains how to change the system time, if required. This has to be done for both DR550 SSAM Servers. Follow these steps:

1. Log in to AIX and switch your account to root.
2. To set the time zone, date, and time, issue the command:

```
smitty chtz
```

Fill in the information for the time zone, date, and time in the corresponding fields. Commit the change, and press F10 to exit from SMITTY when you are done.

Managing the Time Zone variable

The setup of the time zone has to be done on all DR550 SSAM Servers.

The DR550 ships with the time zone set to GMT and without Daylight Savings Time being enabled. If you need to enable Daylight Savings Time, refer to “About Daylight Savings Time (DST)”. The fixes required to enable new US laws regarding DST have been installed already.

About Daylight Savings Time (DST)

If the Daylight Savings Time option is enabled, the default in AIX is for the system time to move forward 1 hour (to DST) at 2:00 a.m. the first Sunday in April, and move back one hour (to Standard Time) at 2:00 a.m. on the last Sunday in October. Starting in 2007, this switch occurs on the second Sunday in March and ends the first Sunday in November. The default is hardcoded and is not stored in any user accessible file. However, the date and time at which the switch to DST and ST occurs can be altered by the root user. To see if DST is enabled, issue the command `echo $TZ`; if the time zone variable ends in DT, DST is enabled.

Follow these steps:

- ▶ To set Daylight Savings Time, run **smitty chtz** and answer “1 yes” to “Use Daylight Savings Time?”.
- ▶ If Daylight Savings Time does not apply to your location, answer “2 no” to “Use Daylight Savings Time?”.

Changing the effective date to switch to DST

To change the date or time at which the system switches to DST and back to ST, edit the TZ line in `/etc/environment`. Change the line to read like the following:

```
TZ=CST6CDT,M4.1.0/1:00:00,M10.1.0/1:00:00
```

This would effect a change to Daylight Savings Time at 1:00 a.m. on the first Sunday in April and change back to 1:00 a.m. on the first Sunday in October, and keep the US Central Time Zone time offset from GMT.

The breakdown of the string is:

- ▶ CST6CDT is the time zone you are in.
- ▶ M4 is the fourth month.
- ▶ .1 is the first occurrence of a day in the month.
- ▶ .0 is Sunday.
- ▶ /1:00:00 is the time.

In more detail, the format is TZ = local_timezone,date/time,date/time. Here date is in the form of Mm.n.d, day d(0-6) of week n (1-5, where week 5 means “the last d day in month m” and which might occur in either the fourth or the fifth week) of month m of the year. Week 1 is the first week in which day d occurs. Day zero is Sunday. This format is compliant with POSIX 1003.1 standards for Extensions to Time Functions.

A reboot is required to activate the new settings.

Accept time within SSAM on DR550 DR1 and DR550 DR2 single-node

To accept this time within SSAM, do the following:

First, kill the dsmserv process. Here is an example:

```
ps -ef | grep dsm
root 188502 590004 0 15:01:33 pts/2 0:00 grep dsm
root 266276 1 0 May 08 - 0:48 /usr/tivoli/tsm/server/bin/dsmserv quiet
```

End the process by entering the **kill** command followed by the process number. In the case of our example:

```
kill -9 761886
```

Start SSAM in foreground mode

The SSAM server starts as part of the HACMP cluster, so here is the manual process to bring up SSAM server to enable SSAM server sessions (HACMP should be running and the LVs should be mounted):

1. Log in as root at node 1.
2. Enter **cd /usr/bin** at the command line.
3. Copy the original SSAM startup script by entering **cp startserver startserver.foreground**

4. Open the copied file with the vi editor (**vi startserver.foreground**) and scroll to the last line (\$DSMSERV_DIR/dsmserv quiet &).
5. Remove the quiet & at the end of the line and change this line to the following:
\$DSMSERV_DIR/dsmserv
6. Save the file and exit the editor.
7. Enter **startserver.foreground**. (This will start the SSAM server in foreground mode.)

Accept the date in SSAM - for all configurations

Wait for the TSM> prompt and follow these steps:

1. tsm: TSM> **accept date** (to match the date with system date and time)
ANR0894I Current system has been accepted as valid.
2. TSM:TSM> **enable sessions all**
ANR2552I Server now enabled for ALL access.
3. TSM:TSM> **halt** (to halt the SSAM server)
ANR2017I Administrator SERVER_CONSOLE issued command: HALT
ANR0991I Server shutdown complete.
4. **shutdown -Fr** (to restart the server which will automatically restart SSAM server)

Accept time within SSAM on a DR550 DR2 dual-node

The SSAM server starts as part of HACMP cluster, so here is the manual process to bring up SSAM server to enable SSAM server sessions. When starting this procedure, HACMP must already be stopped on both cluster nodes.

On DR550 SSAM Server 1, varyon the following volume groups:

```
varyonvg TSMApps
varyonvg TSMDBLogs
varyonvg TSMDBBkup
varyonvg TSMStg
varyonvg TSMDBLogsMirr
```

When all volume groups are varied on, mount the following file systems:

```
mount /tsm
mount /tsmDb
mount /tsmLog
mount /tsmdbbkup
mount /tsmDbM
mount /tsmLogM
```

On DR550 SSAM Server 1, make a copy of startserver script:

```
cp -p /usr/bin/startserver /tmp
```

On DR550 SSAM Server 1, edit the startserver script:

1. Run **vi /tmp/startserver**.
2. Change the last line of the file from \$DSMSERV_DIR/dsmserv quiet & to \$DSMSERV_DIR/dsmserv.
3. Save and exit the startserver script.

Start SSAM server in the foreground mode with **/tmp/startserver**.

Wait for the TSM> prompt and follow these steps:

1. tsm: TSM> **accept date** (to match the date with system date and time)

ANR0894I Current system has been accepted as valid.

2. TSM:TSM> **enable sessions all**

ANR2552I Server now enabled for ALL access.

3. TSM:TSM> **halt** (to halt the SSAM server)

ANR2017I Administrator SERVER_CONSOLE issued command: HALT

ANR0991I Server shutdown complete.

Remove the modified startserver script:

```
rm /tmp/startserver
```

Start HACMP on both nodes. The SSAM server process will be started automatically as part of the HACMP cluster.

3.6.2 Attaching the DR550 DR1 to the network

First, to avoid conflicts in an existing network environment, check that the IP addresses set at the factory (see 2.5.3, “Network configuration” on page 33) are not already in use by another device on your network. Note that there should rarely be a conflict, because the factory set IP addresses are usually not used in real network environments.

Second, make sure to plug your cable into the correct DR550 SSAM Server network interface card port. On the DR550 DR1, your network cable will be connected to the Ethernet adapter in PCI-X slot 6 of the DR550 SSAM Server. The Ethernet adapter in slot 6 (you have a choice between fiber or copper) is used for attaching to the network.

Onboard Ethernet ports T5 and T6 are used for connecting DR550 DR1 to the optional DR550 Storage Controller (DS4200).

The RSM server connects to port 1 of the Ethernet adapter in slot 5 of the DR550 SSAM Server (see also 9.4.2, “Connecting the Remote Support Manager” on page 374).

To configure a correct IP address for an interface, you have to identify that interface first. The easiest way to identify an interface is by using the adapter location code (See Figure 3-7).

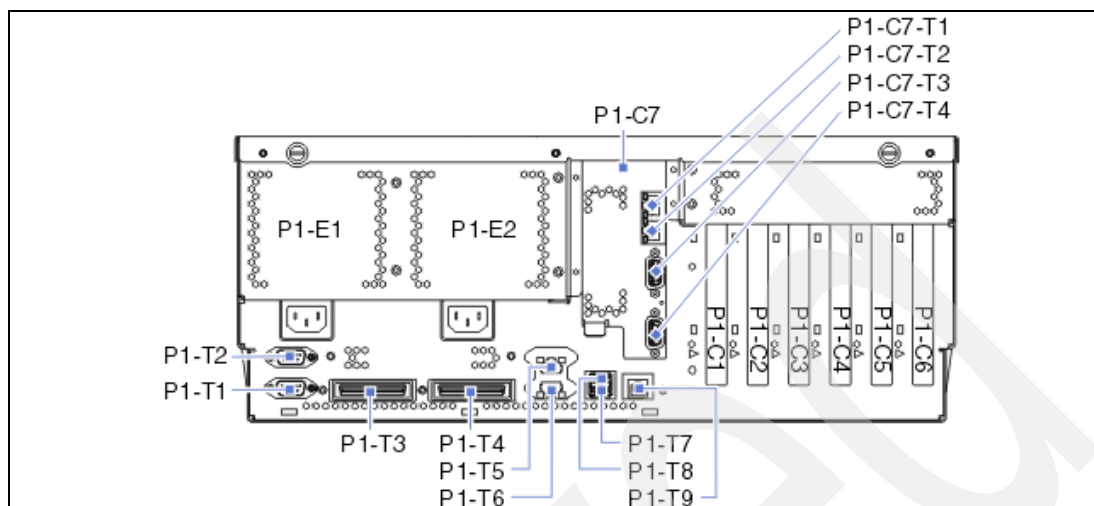


Figure 3-7 DR550 SSAM Server (p52A) rear view with I/O adapters and PCI slots

Use the following AIX command to display the adapter name and the corresponding location code:

```
lscfg -vp | grep ent[0-9] | awk '{ print $1, " ", $2 }' | sort
ent0 U787F.001.DPM2KTL-P1-C5-T1
ent1 U787F.001.DPM2KTL-P1-C5-T2
ent2 U787F.001.DPM2KTL-P1-C6-T1
ent3 U787F.001.DPM2KTL-P1-C6-T2
ent4 U787F.001.DPM2KTL-P1-T5
ent5 U787F.001.DPM2KTL-P1-T6
```

The location codes in the form PX-XX are highlighted in the output of the command. In this example, ent4 is the adapter in location P1-T4, the location of the top on-board Ethernet adapter. If a location code has two additional characters after the location code, it signifies the position of a port on a multi-port adapter. P1-C5-T1 is the Ethernet port at the top and the P1-C5-T2 is the port at the bottom on a physical PCI-X card installed in PCI slot #5 (P1-C5 location). At the AIX level, there is a logical interface name that corresponds to a physical interface name. In our example, “ent4” is the physical interface name and “en4” is the logical interface name. In a DR550 configuration, it is always a one-to-one mapping, that is, en4 corresponds to ent4, ent5 to ent5, and so on. This method is applicable to both copper and fiber Ethernet adapters.

For further information about DR550 SSAM Server (p52A) location codes, refer to:

<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/iphau/1occodes.htm>

Gigabit (fiber) Ethernet network

Fibre Channel was originally designed to support fiber optic cabling only. When copper support was added, the committee decided to keep the name in principle, but change its spelling from fiber to fibre. When referring to specific cabling, the correct American English spelling of fiber should be used. Therefore, we refer to fiber optic cabling only.

3.6.3 Attaching the DR550 DR2 to the IP network

In order to avoid conflicts in an existing network environment, check that the IP addresses set at the factory (see 2.5.3, “Network configuration” on page 33) are not already in use by another device on your network. Note that there should rarely be a conflict, because the factory set IP addresses are usually not used in real network environments.

Make sure to plug your cable or cables into the correct DR550 SSAM Server interface or interfaces. On the DR550, your network will always be connected to the Ethernet adapter in PCI-X slot 6 (location code P1-C6) of the DR550 SSAM Server (p52A).

Identify the interfaces within the AIX operating system so that you configure the correct IP addresses for the interfaces being used. For further assistance with the identification of Ethernet interfaces, see 3.6.2, “Attaching the DR550 DR1 to the network” on page 91.

In the following sections, we successively explain the 10/100/1000 Mbps network setup for the single-node and the dual-node configurations.

Important: Make sure to use the correct adapters and ports on the DR550 SSAM Servers when connecting to your network; check the interfaces during the configuration. Do not connect the DR550 if you already use the factory IP addresses for any other device within your network environment.

10/100/1000 Mbps (copper) Ethernet network

When you use the two-port Gigabit Ethernet-TX PCI-X adapter (copper, FC 1983) for connecting to a customer's network:

- ▶ For the DR550 DR2 single-node configuration, always plug your customer supplied Ethernet cable into the upper port of the adapter installed in PCI-X slot number 6 (location code P1-C6).
- ▶ For the DR550 DR2 dual-node configuration, plug your customer's supplied Ethernet cables into both ports of the adapter installed in PCI-X slot number 6 (location code P1-C6).

Connect the RSM server directly to port 15 of the two DR550 internal Ethernet switches. See 9.4.2, “Connecting the Remote Support Manager” on page 374 for more information.

Gigabit (fiber) Ethernet network

Ethernet connectivity for fiber is similar to Ethernet connectivity for copper.

When you use the two-port Gigabit Ethernet-SX PCI-X adapter (fiber, FC 3550) for connecting to a customer's network:

- ▶ For the DR550 DR2 single-node configuration, always plug your customer supplied LC-LC fiber cable into the upper port of the adapter installed in PCI-X slot number 6 (location code P1-C6).
- ▶ For the DR550 DR2 dual-node configuration, plug your customer's supplied LC-LC fiber cables into both ports of the adapter installed in PCI-X slot number 6 (location code P1-C6).

3.6.4 Configuring the IP network for a DR550 DR1 and DR2 single-node

To configure the DR550 DR1 for use within your network, you must change only one TCP/IP address. The steps are very similar to those for the DR550 DR2 single-node configuration.

DR550 DR1 and DR2 single-node steps summary

There are different ways to change the IP settings. There are various AIX commands that are available for these tasks. For convenience, we recommend that you use the System Management Interface Tool (SMIT) that is part of AIX. This is a quick and efficient tool, and, in addition, the SMIT session is logged into a file (smit.log). Then, if necessary, you can use the log file to analyze a situation in case of problems.

The following sequence of steps assumes the use of SMIT:

1. Obtain one IP address from your network administrator.
2. Connect to drs_engine through the management console.
3. Edit the /etc/hosts file on drs_engine.
4. Change the TCP/IP address for the interface en2 (en0 for fiber optic).
5. Verify the new IP address.
6. Edit the IBM System Storage Archive Manager client option file (dsm.sys).
7. Verify the connection from the administrative client to the SSAM server.
8. Restart the SMIS API with `/opt/IBM/ISS/resetSMIS.sh`.
9. Stop IBM Director Agent with `/opt/ibm/director/bin/twgstop`.
10. Start IBM Director Agent with `/opt/ibm/director/bin/twgstart`.
11. Log off all sessions.

DR550 DR1 and DR2 single-node steps

This section illustrates in detail how to set up the new IP address.

Most settings are done using the System Management Interface Tool (SMIT) of AIX. SMIT is an interactive interface application designed to simplify system management tasks. We start SMIT on the AIX command line prompt by typing `smitty`. The `smitty` command displays a hierarchy of menus that can lead to interactive dialogs. SMIT builds and runs commands as directed by the user.

The use of the name SMIT throughout this section refers to the tool and not to the command.

For the DR550 DR1 or DR2 single node, complete the following steps:

1. Obtain one IP address from your network administrator.

The first step is to set a new address in conformance with your network. You need to obtain one IP address from your network administrator. We suggest that you create a table where you write the factory IP address and the actual IP address side by side, as illustrated in Table 3-2.

Table 3-2 Translation table for IP addresses with DR550 Express

IP description	AIX interface	Factory IP address	Actual IP address
drs_engine	en0 or en2	192.168.1.5 255.255.255.0	100.100.51.121 255.255.255.0

2. Connect to drs_engine through the management console.
 - a. Log in to the DR550 SSAM Server (drs_engine) with user dr550 through the console.
 - b. After you are successfully logged in as dr550, issue the AIX command **su - root** to switch to root. Now you have the necessary AIX system rights to change the network settings.
3. Change the TCP/IP address for the interface en2 (en0 for fiber optic).
 - a. First, make a copy of the /etc/hosts file as it shipped from the factory (see Example 3-2) with the AIX command **cp /etc/hosts /etc/hosts.factory**.

Example 3-2 Preconfigured /etc/hosts file on drs_engine (excerpt)

```
127.0.0.1          loopback localhost      # loopback (lo0) name/address

192.168.1.5      drs_engine
```

- b. Edit /etc/hosts to change the factory default address 192.168.1.5 to the actual address provided by your network administrator (100.100.51.121 in our example in Example 3-3) and save /etc/hosts.

Example 3-3 Custom-made /etc/hosts file on drs_engine (excerpt)

```
127.0.0.1          loopback localhost      # loopback (lo0) name/address

100.100.51.121    drs_engine
```

- c. Save the file and exit the editor.
 - d. From the AIX command line, start the SMIT to change the IP characteristics of a network interface with the command **smitty chinet**.

Highlight the network interface you want to change (en2). Press Enter to change it. In the INTERNET ADDRESS (dotted decimal) field, type the new IP address 100.100.51.121. In the Network MASK (hexadecimal or dotted decimal) field, type the network mask 255.255.255.0 of the new network. You might also need to change the IP address and domain name for the Name Server and default gateway. When done, press Enter.

Check the result of the command at the COMMAND STATUS window of SMIT. The Command field shows OK while the text below says en0 changed. Exit SMIT by pressing F10.
4. Verify the new IP address and host name.

From the AIX command line, use the **ping 100.100.51.121** command to verify that the new address is working correctly. Use the **ping drs_engine** command to verify that the name is being correctly resolved and that the right IP address is being used by the **ping** command.
5. Edit the IBM System Storage Archive Manager client option file (dsm.sys).

You need to adjust the IBM System Storage Archive Manager client system options file so that the IBM System Storage Archive Manager API or any IBM System Storage Archive client (dsmadm) can find the IBM System Storage Archive Manager server.

On the DR550 SSAM Server, replace the value in the `dsm.sys` file for the `tcpserveraddress` with your new server name or IP address. We recommend that you do not use a dot address, for example, 100.100.51.111. Use the TCP/IP host name instead, which is `drs_engine`.

- a. Change to the directory where the file is located by using `cd /usr/tivoli/tsm/client/ba/bin`.
- b. Edit the file by entering `vi dsm.sys`. The file shown in Example 3-4 should appear.

Example 3-4 Example of dsm.sys file

SERvername	TSM
COMMmethod	TCPip
TCPPort	1500
TCPServeraddress	drs_engine

Tip: When using the TCP/IP name for the `tcpserveraddress` field in the IBM System Storage Archive Manager client system options file (`dsm.sys`), you do not need to change the value after a change of the IP address. This is not true when you use the dot address.

6. Connect to the System Storage Archive Manager server from `drs_engine1` with the command `dsmadmc`. Log in with the System Storage Archive Manager administrator `admin` and its password. Verify the connection with the System Storage Archive Manager command `query session`. Exit the administrative command line client with the System Storage Archive Manager command `quit`.
7. Restart the SMIS API with `/opt/IBM/ISS/resetSMIS.sh`.
8. Stop IBM Director Agent with `/opt/ibm/director/bin/twgstop`.
9. Start IBM Director Agent with `/opt/ibm/director/bin/twgstart`.
10. Log off all sessions. After a successful configuration, you should exit all shells with the AIX command `exit`. Depending on the actual shells, you must type the `exit` command more than once. For example, type it one time to close the shell of `root` and one time to close the shell of `dr550`. Repeat this until you see the AIX login prompt on the management console session.

Continue with the setup of SSAM and the retention policies by going to 3.6.6, “SSAM and data retention policy setup” on page 112.

3.6.5 Configuring the IP address and HACMP for a dual-node DR550 DR2

The dual-node configuration relies on High Availability Cluster Multi-Processing (HACMP) for AIX. HACMP provides the automated monitoring and recovery for all the resources on which the application (System Storage Archive Manager) depends.

As part of the IP network configuration, you also have to set up the HACMP network. This is described in detail in the following sections. First, we provide an introduction to HACMP networks in general, and then we provide the step-by-step guide for using this information for the actual setup.

To configure the DR550 dual-node configuration for use within your network, you must change the following pre-configured TCP/IP addresses:

- ▶ On the DR550 SSAM Server 1 (drs_engine1):
 - The boot address of the DR550 SSAM Server (drs_engine1_ext_boot)
 - The standby address of the DR550 SSAM Server (drs_engine1_ext_stdby)
 - The HACMP cluster service address (drs_cluster_ext_svc)
- ▶ On the DR550 SSAM Server 2 (drs_engine2):
 - The boot address of the DR550 SSAM Server (drs_engine2_ext_boot)
 - The standby address of the DR550 SSAM Server (drs_engine2_ext_stdby)

Figure 3-8 shows the IP addressing scheme.

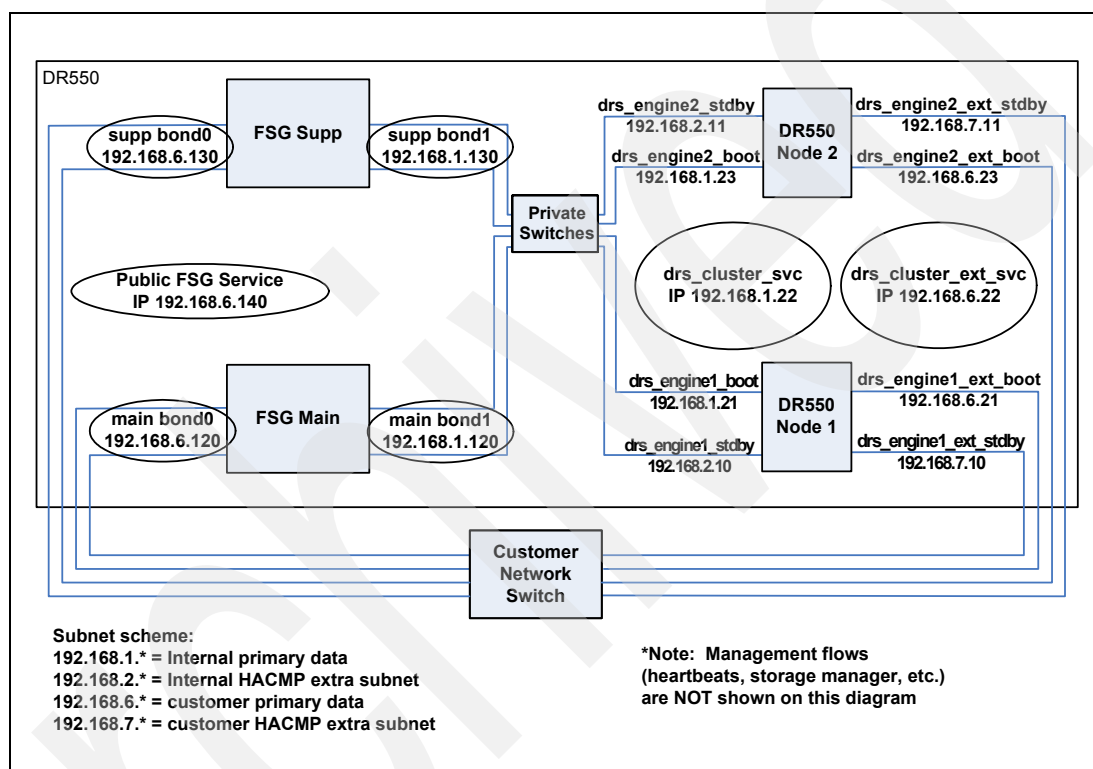


Figure 3-8 DR550 DR2 dual-node dual FSG data traffic* IP addressing scheme

Important: Do not change the factory IP configuration for the DS4200 controllers or any IP address on the DR550 SAN Switches (2005-B16).

Introduction to HACMP networks

An HACMP cluster is made up of the following physical components:

▶ Nodes

Nodes form the core of an HACMP cluster. A *node* is a processor that runs the AIX 5L software, the HACMP software, and the application software. Therefore, drs_engine1 and drs_engine2 are called *cluster nodes*.

- Shared external disk devices

Each node must have access to one or more shared external disk devices. A *shared external disk device* is a disk physically connected to multiple nodes. The shared disk stores mission-critical data, typically mirrored or RAID-configured for data redundancy. The shared external disk devices in the DR550 are the DS4000 disks.

A node in an HACMP cluster must also have internal disks that store the operating system and application binaries (System Storage Archive Manager), but these disks are not shared. Each DR550 SSAM Server in the DR550 dual-node configuration has two internal disks.

The shared external disk devices in the DR550 are configured with *non-concurrent access*. In non-concurrent access environments, only one connection is active at any given time, and the node with the active connection owns the disk. When a node fails, disk takeover occurs when the node that currently owns the disk leaves the cluster and a surviving node assumes ownership of the shared disk.

Note: The DR550 Dual-Node configuration uses disk heartbeat to determine the health of the active node at all times. The disk used for this is defined as *concurrent access within AIX LVM*. This is because both nodes need access to this volume in order to determine the health status.

- Networks

As an independent, layered component of AIX 5L, the HACMP software is designed to work with any TCP/IP-based network. Nodes in an HACMP cluster use the network to allow clients to access the cluster nodes, enable cluster nodes to exchange heartbeat messages, and in concurrent access environments, serialize access to data.

- Network interfaces

HACMP software defines two types of communication networks, characterized by whether these networks use communication interfaces based on the TCP/IP subsystem (TCP/IP-based) or communication devices based on non-TCP/IP subsystems (device-based). The DR550 uses the TCP/IP subsystem.

- Clients

A *client* is a processor that can access the nodes in a cluster over a local area network. Clients each run a “front end” or client application that queries the server application running on the cluster node. A client for the DR550 is a node, where a document workstation uses the System Storage Archive Manager API. The HACMP software provides a highly available environment for critical data and applications on cluster nodes. The HACMP software does not make the clients themselves highly available.

Cluster nodes communicate with each other over communication networks. If one of the physical network interface cards on a node within a network fails, HACMP preserves the communication to the node by transferring the traffic to another physical network interface card on the same node. If a “connection” to the node fails, HACMP transfers resources to another node to which it has access.

The Reliable Scalable Cluster Technology (RSCT) sends heartbeats between the nodes over the cluster networks to periodically check on the health of the cluster nodes themselves. If HACMP detects no heartbeats from a node, a node is considered to have failed, and resources are automatically transferred to another node.

In addition to this, HACMP uses disk heartbeat to detect failure; this feature is discussed in detail in “Disk heartbeat” on page 101.

This heartbeat mechanism is preconfigured, and no additional configuration is required.

HACMP physical and logical networks

A *physical network* connects two or more physical network interfaces. For example, HACMP is able to work with different physical networks, such as TCP/IP-based and non-IP-based networks. Because the DR550 is running within TCP/IP local area networks only, there are only two physical networks for HACMP in your environment (one uses internal DR550 Ethernet switches and the other external customer supplied Ethernet switches).

A *logical network* is a portion of a physical network that connects two or more logical network interfaces or devices. A logical network interface or device is the software entity that is known by an operating system. There is a one-to-one mapping between a physical network interface or device and a logical network interface or device. Each logical network interface can exchange packets with each other's logical network interface on the same logical network.

If subsets of logical network interfaces on the logical network need to communicate with each other (but with nothing else) while sharing the same physical network, subnets are used. A *subnet mask* defines the part of the IP address that determines whether one logical network interface can send packets to another logical network interface on the same logical network.

The IP address of the physical adapter will be replaced by HACMP if the adapter is changing its role in the HACMP cluster. For example, in the case of a network problem, the standby address of an adapter will be replaced by the HACMP service address. This process is known as HACMP adapter swap.

HACMP has its own, similar concept of a logical network. All logical network interfaces in an HACMP network can communicate HACMP packets with each other directly. HACMP generates a name for each HACMP logical network, such as `net_ether_01`.

An HACMP logical network can contain one or more subnets. RSCT takes care of routing packets between logical subnets.

Important: These are important requirements for HACMP networks:

- ▶ The netmask for all adapters in an HACMP network must be the same even though the service and standby adapters are on different logical subnets.
- ▶ The IP addresses of the Standby adapters must be on a separate subnet from the service adapters even though they are on the same physical network.

See the following sections for more information about HACMP networks, subnetworks, and communication.

HACMP communication interfaces

An HACMP communication interface is a grouping of a logical network interface, a service IP address, and a service IP label that you defined to HACMP. HACMP communication interfaces combine to create IP-based networks.

An HACMP communication interface is a combination of:

- ▶ A *logical network interface*, which is the name to which AIX 5L resolves a port (for example, `en0`) of a physical network interface card.
- ▶ A *service IP address*, which is an IP address (for example, `192.168.1.22`) over which services, such as an application (System Storage Archive Manager server), are provided and over which client nodes communicate.

- ▶ A *service IP label*, which is a logical label (for example, `drs_cluster_svc`) that maps to the service IP address.

Subnet routing requirements in HACMP

A *subnet route* defines a path, defined by a subnet, for sending packets through the logical network to an address on another logical network. Beginning with AIX 5L V5.1, you can add multiple routes for the same destination in the kernel routing table. If multiple matching routes have equal criteria, routing can be performed alternatively using one of the several subnet routes.

It is important to consider subnet routing in HACMP because of the following considerations:

- ▶ HACMP cannot distinguish between logical network interfaces that share the same subnet route. If a logical network interface shares a route with another interface, HACMP has no means to determine its health. For more information about network routes, see the AIX 5L man pages for the **route** command (AIX command **man route**).
- ▶ Various constraints are often imposed on the IP-based networks by a network administrator or by TCP/IP requirements. The subnets and routes are also constraints within which HACMP must be configured for operation.

Important: We recommend that each communication interface on a node belongs to a unique subnet so that HACMP can monitor each interface. This is not a strict requirement in all cases and depends on several factors. In such cases where it is a requirement, HACMP enforces it.

HACMP service IP address and label

A *service IP label* is a label that maps to the service IP address and is used to establish communication between client nodes and the server node. Services, such as the System Storage Archive Manager server, are provided using the connection made over the service IP label.

A service IP label can be placed in a resource group as a resource, which allows HACMP to monitor its health and keeps it highly available, either within a node or, if *IP address takeover* is configured, between the cluster nodes by transferring it to another node in the event of a failure.

HACMP heartbeating

In order for an HACMP cluster to recognize and respond to failures, it must continually check the health of the cluster. Some of these checks are provided by the heartbeat function.

There are two forms of HACMP heartbeating configured on the DR550, as depicted in Figure 3-9 on page 101.




Figure 3-9 IP heartbeat and disk heartbeat as used in the DR550

Heartbeat over IP

Each cluster node sends heartbeat messages at specific intervals to other cluster nodes and expects to receive heartbeat messages from the nodes at specific intervals. If messages stop being received, HACMP recognizes that a failure has occurred.

The heartbeat in the DR550 is established through the TCP/IP network. Therefore, RSCT Topology Services use the HACMP network topology to dynamically create a set of heartbeat paths that provides coverage for all TCP/IP interfaces and networks. These paths form heartbeat rings, so all components can be monitored without requiring excessive numbers of heartbeat packets.

In order for RSCT to reliably determine where a failure occurs, it must send and receive heartbeat packets over specific interfaces. This means that each NIC configured in HACMP must have an IP label on a separate subnet. For the DR550, this is configured by using heartbeating over IP interfaces (in contrast to heartbeating over IP aliases). With this method, you must configure all service and non-service IP labels on separate subnets; otherwise, HACMP will not work correctly.

Disk heartbeat

In addition to heartbeat over IP, the DR550 HACMP offering uses another method called *disk heartbeat*.

Disk heartbeat is a form of non-IP heartbeat that utilizes the existing shared disks of any disk type. This feature, which was introduced in HACMP V5.1, is quickly becoming the preferred method of non-IP heartbeat, because it eliminates the need for serial cables and 8-port asynchronous adapters. It also can easily accommodate greater distances between nodes when using a SAN environment.

This feature requires using enhanced concurrent volume groups to allow access to the disk by each node as needed. It utilizes a special reserved area on the disks to read and write the heartbeat data, as shown in Figure 3-10.

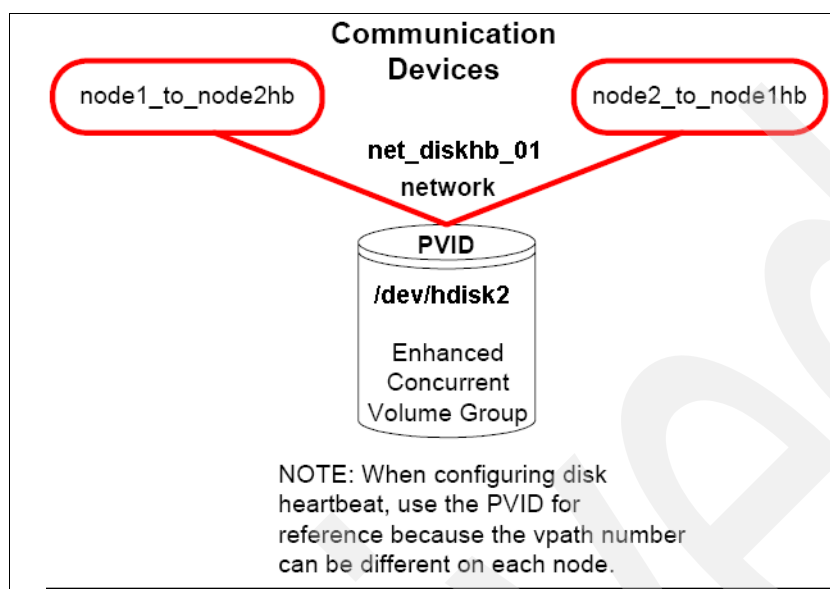


Figure 3-10 Example of how the disk heartbeat works

In addition to the service and non-service IP labels, the DR550 uses LUN0 (hdisk2) and LUN2 (hdisk4) from the DR550 Storage Controller (DS4200) for disk heartbeat.

Once the cluster is up and running, you can monitor the activity of the disk (actually all) heartbeats using `lssrc -ls topsvcs`. This command gives an output similar to the one shown in Example 3-5.

Example 3-5 Output of `lssrc`

```

NIM's PID: 721058
diskhb_0      [ 4] 2      2      S    255.255.10.3    255.255.10.3
diskhb_0      [ 4] rhdisk4    0x87c2ee52    0x87c2ee59
HB Interval = 2.000 secs. Sensitivity = 4 missed beats
Missed HBs: Total: 36 Current group: 36
Packets sent   : 51434 ICMP 0 Errors: 0 No mbuf: 0
Packets received: 48888 ICMP 0 Dropped: 0
NIM's PID: 704662
diskhb_1      [ 5] 2      2      S    255.255.10.2    255.255.10.2
diskhb_1      [ 5] rhdisk2    0x87c2ee51    0x87c2ee59
HB Interval = 2.000 secs. Sensitivity = 4 missed beats
Missed HBs: Total: 30 Current group: 30
Packets sent   : 51430 ICMP 0 Errors: 0 No mbuf: 0
Packets received: 48887 ICMP 0 Dropped: 0

```

Be aware that there is a grace period for heartbeats to start processing. This is normally around 60 seconds. So if you run this command quickly after starting the cluster, you may not see anything at all until heartbeat processing is started after the grace period time has elapsed.

For information about how to set up this feature, refer to the following:

<http://publib.boulder.ibm.com/infocenter/clresctr/vxrx/index.jsp?topic=/com.ibm.cluster.hacmp.doc/hacmpbooks.html>

Note: It is not necessary for the customer to set up disk heartbeating, because it is preconfigured with the DR550 dual-node.

For more information about disk heartbeating, refer to *Implementing High Availability Cluster Multi-Processing (HACMP) Cookbook*, SG24-6769.

Application Server Monitoring

Application Monitoring enables you to configure multiple monitors for an application server to monitor specific applications and processes; you can define actions to take upon detection of an unexpected termination of a process or other application failures.

HACMP can monitor applications that are defined to application servers. Process monitoring detects the termination of a process, using RSCT Resource Monitoring and Control (RMC) capability.

DR550 HACMP is configured to use a long-running monitoring mode. In this mode, the application monitor periodically checks that the application server is running. The monitor is run multiple times based on the monitoring interval that you specify. If the monitor returns a zero code, it means that the application is running successfully. A non-zero return code indicates that the application has failed. The checking starts after the specified stabilization interval has passed.

Action on application failure is configured to use the “failover” action. HACMP will recover the application server (SSAM server) on the secondary HACMP cluster node in the event of an application failure at the primary node. For more information, refer to:

<http://publib.boulder.ibm.com/infocenter/clresctr/vxrx/index.jsp?topic=/com.ibm.cluster.hacmp.doc/hacmpbooks.html>

Dual-node configuration steps summary

There are several ways to change the IP settings and reconfigure HACMP through AIX and HACMP commands. For convenience, we recommend that you use the System Management Interface Tool (SMIT) that is part of AIX. This is a quick and efficient tool, and in addition, the SMIT session is logged into a file (smit.log). You can use the log file later to analyze a situation in case of problems.

The following sequence of steps assumes the use of SMIT:

1. Obtain five applicable IP addresses from your network administrator.
2. Connect to drs_engine1 through the management console (Engine1).
3. Stop HACMP (if necessary) on both cluster nodes.
4. Edit the /etc/hosts file on drs_engine2.
5. Change the TCP/IP addresses for the interfaces en2 and en3.
6. Connect to drs_engine2 through the management console (Engine2).
7. Edit the /etc/hosts file on drs_engine2.
8. Change the TCP/IP addresses for the interfaces en2 and en3.
9. Remove the network from the HACMP cluster.
10. Define the new HACMP network.

11. Create the network resource cluster address within HACMP.
12. Add the new network resource to the HACMP resource group.
13. Verify and synchronize the HACMP cluster.
14. Start the HACMP cluster.
15. Edit the SSAM (Tivoli Storage Manager) client option file (dsm.sys).
16. Restart the SMIS API with `/opt/IBM/ISS/resetSMIS.sh`.
17. Stop IBM Director Agent with `/opt/ibm/director/bin/twgstop`.
18. Start IBM Director Agent with `/opt/ibm/director/bin/twgstart`.
19. Restart the SMIS API with `/opt/IBM/ISS/resetSMIS.sh`.

Dual-node configuration steps

This section illustrates in detail how to set up all five new IP addresses.

Most settings are done using the System Management Interface Tool (SMIT) of AIX. SMIT is an interactive interface application designed to simplify system management tasks. We start SMIT on the AIX command line prompt by typing `smitty`. The `smitty` command displays a hierarchy of menus that can lead to interactive dialogs. SMIT builds and runs commands as directed by the user.

Use of the name SMIT throughout this section refers to the tool and not to the command.

For the dual-node configuration, complete the following steps:

1. Obtain five applicable IP addresses from your network administrator.

The first step is to set new addresses in conformance with your network and HACMP requirements (see “Introduction to HACMP networks” on page 97). You need to obtain five IP addresses from your network administrator. We suggest that you create a table where you write the factory IP addresses and the actual IP addresses side by side, as illustrated in Table 3-3.

Table 3-3 Translation table for IP addresses with dual-node configuration

IP description	AIX interface	Factory IP address	Actual IP address
drs_engine1_ext_boot	en2	192.168.6.21 255.255.255.0	100.100.51.121 255.255.255.0
drs_engine1_ext_stdbby	en3	192.168.7.10 255.255.255.0	100.100.52.110 255.255.255.0
drs_cluster_ext_svc	n/a	192.168.6.22 255.255.255.0	100.100.51.122 255.255.255.0
drs_engine2_ext_boot	en2	192.168.6.23 255.255.255.0	100.100.51.123 255.255.255.0
drs_engine2_ext_stdbby	en3	192.168.7.11 255.255.255.0	100.100.52.111 255.255.255.0

2. Connect to drs_engine1 through the management console (Engine1).
 - a. Log in to the DR550 SSAM Server drs_engine1 with user dr550 through the management console. Use the PrtSc key on the keyboard-video-mouse (KVM) keyboard to reach the console. Press Enter to see the login prompt.
 - b. After you have successfully logged in as dr550, issue the AIX command `su - root` to switch to root. Now you have the necessary AIX system rights to change the network settings and to reconfigure HACMP.
3. Stop HACMP (if necessary).
 - a. Check to see if HACMP is running or stopped with the AIX command `lssrc -g cluster`.
 - b. If it is running, stop HACMP (see 4.1.3, “Stopping cluster services” on page 144). We assume there is no data traffic when stopping HACMP, because this is the implementation phase. It can take several minutes to stop HACMP. After HACMP is stopped, continue to the next step.
4. Change the TCP/IP addresses for the interfaces en2 and en3.
 - a. First, create a copy of the /etc/hosts file that is shipped from the factory (see Example 3-6) with the AIX command `cp /etc/hosts /etc/hosts.factory`.

Example 3-6 Preconfigured /etc/hosts file on drs_engine1 and drs_engine2 (excerpt)

```

127.0.0.1          loopback localhost      # loopback (100) name/address

#
#      Engine 1
#

192.168.6.22      drs_cluster_ext_svc
192.168.7.10      drs_engine1_ext_stdby
192.168.6.21      drs_engine1_ext_boot

#
#      Engine 2
#

192.168.6.23      drs_engine2_ext_boot drs_engine2
192.168.7.11      drs_engine2_ext_stdby drs_engine2
192.168.4.24      drs_engine2_FASTt1_ctr1A
192.168.5.26      drs_engine2_FASTt1_ctr1B

```

- b. From the AIX command line, start SMIT to change the IP characteristics of the en0 network interface with the command `smitty chinnet`.
 - c. Highlight the network interface you want to change (en2). Press Enter to change it. In the INTERNET ADDRESS (dotted decimal) field, type the new IP address 100.100.51.121. In the Network MASK (dotted decimal) field, type the network mask 255.255.255.0 of the new network. You might also need to change the IP address and domain name for the Name Server and default gateway. When done, press Enter.
- Check the result of the command at the COMMAND STATUS window of SMIT. The Command field shows OK while the text below says en0 changed. Exit SMIT by pressing F10.

- d. From the AIX command line, start SMIT to change the IP characteristics of the en1 network interface with the command **smitty chinet**.

Highlight the second network interface you want to change, for example, en3. Press Enter to change it. In the INTERNET ADDRESS (dotted decimal) field, type the new IP address 100.100.52.110. In the Network MASK (hexadecimal or dotted decimal) field, type the network mask 255.255.255.0 of the network. When done, press Enter.

Check the result of the command at the COMMAND STATUS window of SMIT. The Command field shows OK while the text below says en1 changed. Exit SMIT by pressing F10.

5. Edit the /etc/hosts file on drs_engine1.
 - a. Open the /etc/hosts file with an editor of your choice, for example, use vi (the AIX command to start vi and open the file is **vi /etc/hosts**). If you are unfamiliar with AIX editors, you can use xedit, an X Window System-based editor. For the latter, you first need to start an X-session. Start the X-session with the AIX command **startx** (see the **startx** man pages for details about issuing **man startx**). Then, open the /etc/hosts file with the AIX command **xedit /etc/hosts**. This will open a graphical editor session.

Change all of the preconfigured addresses according to the table you have prepared in step 1 on page 104. Only change these addresses; do not change addresses that are not listed in Table 3-3 on page 104 or any words in the /etc/hosts file. Change the addresses in all the stanzas of the /etc/hosts file, that is, you have to change five addresses. See Example 3-7.

Example 3-7 Custom-made /etc/hosts file on drs_engine1 and drs_engine2 (excerpt)

```
127.0.0.1          loopback localhost      # loopback (lo0) name/address

#
#      Engine 1
#

100.100.51.122    drs_cluster_ext_svc
100.100.52.110    drs_engine1_ext_stdby
100.100.51.121    drs_engine1_ext_boot

#
#      Engine 2
#

100.100.51.123    drs_engine2_ext_boot drs_engine2
100.100.52.111    drs_engine2_ext_stdby
192.168.4.24      drs_engine2_FASTt1_ctr1A
192.168.5.26      drs_engine2_FASTt1_ctr1B
```

- b. Save the file and exit the editor. If you are running an X-session for your editor, quit the X-session by pressing the Ctrl+Alt+Backspace key combination.
6. Connect to drs_engine2 through the management console (Engine2).
 - a. Log in to the DR550 SSAM Server drs_engine2 with user dr550 through the management console. Use the PrtSc key on the keyboard-video-mouse (KVM) keyboard to reach the console. Press Enter to get the login window.
 - b. After you have successfully logged in as dr550, issue the AIX command **su - root** to switch to root. Now you have the necessary AIX system rights to change the network settings and to reconfigure HACMP.

7. Change the TCP/IP addresses for the interfaces en2 and en3.
 - a. First, create a copy of the `/etc/hosts` file that is shipped from the factory (see Example 3-6 on page 105) with the AIX command `cp /etc/hosts /etc/hosts.factory`.
 - b. From the AIX command line, start SMIT to change the IP characteristics of the en0 network interface with the command `smitty chinet`.
 - c. Highlight the network interface you want to change (en2). Press Enter to change it. In the INTERNET ADDRESS (dotted decimal) field, type the new IP address 100.100.51.123. In the Network MASK (dotted decimal) field, type the network mask 255.255.255.0 of the new network. You might also need to change the IP address and domain name for the Name Server and default gateway. When done, press Enter.

Check the result of the command at the COMMAND STATUS window of SMIT. The Command field shows OK while the text below says en0 changed. Exit SMIT by pressing F10.
 - d. From the AIX command line, start SMIT to change the IP characteristics of the en1 network interface with the command `smitty chinet`.

Highlight the second network interface you want to change, for example, en3. Press Enter to change it. In the INTERNET ADDRESS (dotted decimal) field, type the new IP address 100.100.52.111. In the Network MASK (hexadecimal or dotted decimal) field, type the network mask 255.255.255.0 of the network. When done, press Enter.

Check the result of the command at the COMMAND STATUS window of SMIT. The Command field shows OK while the text below says en3 changed. Exit SMIT by pressing F10.
8. Edit the `/etc/hosts` file on `drs_engine2`.
 - a. Open the `/etc/hosts` file with an editor of your choice, for example, use vi (the AIX command to start vi and open the file is `vi /etc/hosts`). If you are unfamiliar with AIX editors, you can use `xedit`, an X Window System-based editor. For the latter, you first need to start an X-session. Start the X-session with the AIX command `startx` (see the `startx` man pages for details by issuing `man startx`). Then, open the `/etc/hosts` file with the AIX command `xedit /etc/hosts`. This will open a graphical editor session.

Change all of the preconfigured addresses according to the table you have prepared in step 1 on page 104. Only change these addresses; do not change addresses that are not listed in Table 3-3 on page 104 or any words in the `/etc/hosts` file. Change the addresses in all the stanzas of the `/etc/hosts` file, that is, you have to change five addresses. See Example 3-7 on page 106.
 - b. Save the file and exit the editor. If you are running an X-session for your editor, quit the X-session by pressing the Ctrl+Alt+Backspace key combination.

Important: The `/etc/hosts` files must be identical on both cluster nodes. Failure to comply can result in unpredictable HACMP cluster behavior

9. Verify that the new addresses are working.

- a. From the AIX command line, use the **ping** command for all the addresses (except the cluster service address) to verify that they are working correctly. Because you are logged in to `drs_engine2`, run **ping** on `drs_engine2`. See Example 3-8 for the **ping** command and its results. Stop the **ping** command with the Ctrl+C key combination each time a result is displayed.

Example 3-8 Ping command to verify that IP addresses are working

```
# ping 100.100.51.123
PING 100.100.51.123: (100.100.51.123): 56 data bytes
64 bytes from 100.100.51.123: icmp_seq=0 ttl=255 time=0 ms
64 bytes from 100.100.51.123: icmp_seq=1 ttl=255 time=0 ms
64 bytes from 100.100.51.123: icmp_seq=2 ttl=255 time=0 ms
^C
----100.100.51.123 PING Statistics----
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms
```

- b. Use the keyboard-video-mouse (KVM) console kit to go back to `drs_engine1` using the PrtSc key. Because you have never logged off from that engine, you do not need to identify yourself again. Use the **ping** command again to verify that all addresses (except the cluster service address) are working correctly. Stop the **ping** command each time a result is displayed with the Ctrl+C key combination.

Important: Network connectivity is important for HACMP functionality. Verify network connectivity any time you make changes to the network interfaces or network environment.

10. Remove the network from the HACMP cluster.

Regardless of which engine you are currently working with, start the SMIT HACMP menus by typing **smitty hacmp** on an AIX command line.

Select Extended Configuration → Extended Topology Configuration → Configure HACMP Networks → Remove a Network from the HACMP Cluster. In the Select a Network to Remove window, select `net_ether_02` (make sure not to select `net_ether_01`, because this is the HACMP internal network). Press Enter.

- a. In the ARE YOU SURE? window, press Enter to proceed. Check that the result is OK, and then press F3 twice.

11. Define the new HACMP network.

While still within the HACMP menus of SMIT, configure the new HACMP network for both nodes:

- a. Select Configure HACMP Nodes → Change/Show a Node in the HACMP Cluster and select node `drs_engine1`. Highlight the Communication Path to Node field and press F4 to select a new path for this node. From the list, select the new boot IP address `drs_engine1_ext_boot` (100.100.51.121). This is just an example; use the address you received from the network administrator) and press Enter. Press Enter again to change the path and verify that the result is OK.
- b. Press F3 twice. Go to Change/Show a Node in the HACMP Cluster again and select node `drs_engine2`. Highlight the Communication Path to Node field (where the old value for the path is displayed) and press F4 to select a new path for this node. From the list, select the new boot IP address `drs_engine2_ext_boot` (100.100.51.123) and press Enter. Press Enter again to change the path and verify that the result is OK.

- c. Go back by pressing F3 (four times) until you reach the Extended Configuration window. Select Discover HACMP-related Information from Configured Nodes and verify that the result is OK. Exit by pressing F3.
- d. Select Extended Topology Configuration → Configure HACMP Networks → Add a Network to the HACMP Cluster, and select “ether” below the discovered IP-based network types stanza. In the Network Name field, type `net_ether_02` and verify that the correct netmask for your network is presented in the Netmask field. Change the Enable IP Address Takeover through IP Aliases field to No (you can press Tab or F4). Press Enter, and verify that the result is OK. Press F3 three times.
- e. In the Extended Topology Configuration window, select Configure HACMP Communication Interfaces/Devices → Add Communication Interfaces/Devices, and select Add Discovered Communication Interface and Devices. Select Communication Interfaces. Select `net_ether_02`, and press Enter.

In the list, scroll down and select the boot (en2) and standby (en3) interfaces for *both* nodes by pressing F7. Make sure that you have selected four interfaces in all. Press Enter, and verify that the result is OK. Press F3 three times.

12. Create the network resource cluster address within HACMP.

- a. Go to Extended Resource Configuration → HACMP Extended Resources Configuration → Configure HACMP Service IP Labels/Addresses → Add a Service IP Label/Address → Configurable on Multiple Nodes, select `net_ether_02` and press Enter.
- b. In the IP Label/Address field, press F4. From the list, select `drs_cluster_ext_svc`, and press Enter. Press Enter again to start the SMIT process, and verify that the result is OK.

13. Add the new network resource to the HACMP resource group.

- a. Press F3 (normally five times) until you reach the Extended Configuration window. Select Extended Resource Configuration → HACMP Extended Resource Group Configuration → Change/Show Resources and Attributes for a Resource Group. Select `drs_550_rg` and press Enter.
- b. In the Service IP Label/Addresses field, press F4. Select `drs_cluster_svc`, and press Enter. Press Enter again to start the SMIT process, and verify that the result is OK.

14. Verify and synchronize the HACMP cluster.

- a. Press F3 (normally four times) to go back to the Extended Configuration window. Select Extended Verification and Synchronization, and press Enter.
- b. In the next window, use the default settings, and press Enter.
- c. Check the SMIT result window for an OK status, and quit SMIT by pressing F10.

15. Setting up the default gateway IP address.

- a. Obtain the IP address of a default gateway for the network in which the external services IP address (drs_cluster_ext_svc) is registered. In our example, the external services IP address is 100.100.51.122 and network address is 100.100.51. Before configuring the default gateway IP address, verify that the server can ping it and it is being accessed through an IP address on the 100.100.51 network (in our example):

```
drs_engine1 </>
# ping -R -c 2 100.100.51.10
PING 100.100.51.10: (100.100.51.10): 56 data bytes
64 bytes from 100.100.51.10: icmp_seq=0 ttl=64 time=0 ms
RR:      100.100.51.10
100.100.51.10
drs_engine1_ext_boot (100.100.51.121)
64 bytes from 100.100.51.10: icmp_seq=1 ttl=64 time=0 ms      (same route)
----100.100.51.10 PING Statistics----
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms
```

in this example default gateway IP address is **100.100.51.10**, local interface being used to access default gateway **100.100.51.121**. Both conditions are met and you can proceed with configuring the default gateway IP address on the system.

- b. Start **smitty route**.
- c. Go to Add static Route.
- d. In **DESTINATION Address** field, enter 0.0.0.0 (this is a predefined IP address for default gateway).
- e. In **Default GATEWAY Address**, enter the IP address of a default gateway provided by the network administrator, in our example, 100.100.51.10, and press Enter.

16. Start the HACMP cluster.

Now that the configuration of the HACMP cluster is completed, the cluster can be started. Whether you start it for the first time, which would normally be the case after the procedure above, or as part of the normal operations of the DR550, the procedure is the same. The procedure is covered in 4.1.2, "Starting cluster services" on page 140. Refer to that section to start your cluster now.

17. Edit the SSAM (Tivoli Storage Manager) client option file (dsm.sys).

You need to adjust the client system options file so that the API or any client (such as dsmdmcc) can find the System Storage Archive Manager server:

- a. On the first DR550 SSAM Server drs_engine1, replace the value in the /usr/tivoli/tsm/client/ba/bin/dsm.sys file for the TCPServeraddress with your new HACMP cluster service address. We recommend that you do not use a dot address, for example, 100.100.51.122. Use the TCP/IP domain name instead, which is drs_cluster_ext_svc.

To replace the value, use an editor of your choice, for example, use vi (the AIX command to start vi and open the file is **vi /usr/tivoli/tsm/client/ba/bin/dsm.sys**). If you are not familiar with AIX editors, you can use xedit, an X Window System-based editor. For the latter, you first need to start an X-session. Start the X-session with the AIX command **startx**. Then, open the /usr/tivoli/tsm/client/ba/bin/dsm.sys file with the AIX command **xedit /usr/tivoli/tsm/client/ba/bin/dsm.sys**. This will open a graphical editor session.

Example 3-9 on page 111 shows our example dsm.sys file.

Example 3-9 Example of dsm.sys file

SErvername	TSM
COMMmethod	TCPip
TCPPort	1500
TCPServeraddress	drs_cluster_ext_svc

- b. Save the file and exit the editor. If you are running an X-session for your editor, quit the X-session by pressing the Ctrl+Alt+Backspace key combination.
- c. Connect to the System Storage Archive Manager server from drs_engine1 with the command **dsmadm**. Log in with the System Storage Archive Manager administrator admin and its password. Verify the connection with the System Storage Archive Manager command **query session**. Exit the administrative command line client with the System Storage Archive Manager command **quit**.
- d. On the second DR550 SSAM Server drs_engine2, replace the value in the /usr/tivoli/tsm/client/ba/bin/dsm.sys file for the TCPServeraddress with your new HACMP cluster service address. We recommend that you do not use a dot address, for example, 100.100.51.122. Use the TCP/IP domain name instead, which is drs_cluster_ext_svc.

To replace the value, use an editor of your choice, for example, use vi (the AIX command to start vi and open the file is **vi /usr/tivoli/tsm/client/ba/bin/dsm.sys**). If you are not familiar with AIX editors, you can use xedit, an X Window System-based editor. For the latter, you first need to start an X-session. Start the X-session with the AIX command **startx**. Then, open the /usr/tivoli/tsm/client/ba/bin/dsm.sys file with the AIX command **xedit /usr/tivoli/tsm/client/ba/bin/dsm.sys**. This will open a graphical editor session.
- e. Save the file and exit the editor. If you are running an X-session for your editor, quit the X-session by pressing the Ctrl+Alt+Backspace key combination.
- f. Connect to the System Storage Archive Manager server from drs_engine2 with the command **dsmadm**. Log in with the System Storage Archive Manager administrator admin and its password. Verify the connection with the System Storage Archive Manager command **query session**. Exit the administrative command line client with the System Storage Archive Manager command **quit**.

Tip: When using the TCP/IP domain name for the TCPServeraddress field in the System Storage Archive Manager client system options file (dsm.sys), you do not need to change the value after a network reconfiguration. This is not true when you use the dotted decimal address.

18. Restart the SMIS API with **/opt/IBM/ISS/resetSMIS.sh**.
19. Stop IBM Director Agent with **/opt/ibm/director/bin/twgstop**.
20. Start IBM Director Agent with **/opt/ibm/director/bin/twgstart**.
21. Log off all sessions.

After successful configuration, you should exit all shells on both nodes with the AIX command **exit**. Depending on the actual shells, you must type the **exit** command more than once, for example, one time to close the shell of root and one time to close the shell of dr550. Repeat this until you see the AIX login prompt on both management console sessions.

Continue with the setup of SSAM and the retention policies in 3.6.6, "SSAM and data retention policy setup" on page 112.

3.6.6 SSAM and data retention policy setup

The last task is to set up your SSAM environment and the data retention policies for the archive objects. Data retention policies can vary from very simple to very complex, and so can the setup. The data retention policies for the IBM System Storage DR550 DR1 and DR550 DR2 will be configured within the IBM System Storage Archive Manager. Furthermore, you might want to configure additional processes in relation to the new data retention policies, for example, schedules to migrate data, create storage pool backups or create additional database backups.

Some SSAM definitions, such as a default storage pool with the corresponding retention policy, are already predefined by the factory with some small amount of data archived. This is done in order to activate archive retention protection and therefore prevent data deletion from SSAM controlled storage pools.

To set up your SSAM configuration:

1. Configure data shredding.

The storage pool ARCHIVEPOOL has a predefined SHRED parameter set to 3. Change this value and the shredding method to your requirements. You can find the procedures and more information about data shredding in 5.2.10, “Data shredding” on page 183.

Important: If you do not want to use data shredding, you must disable data shredding by changing the predefined SHRED parameter value of the ARCHIVEPOOL to 0. Otherwise, the space of expired or moved objects cannot be reclaimed by SSAM and is still occupied until a manual shredding is done! The procedure for disabling data shredding is described in “Disable data shredding” on page 188.

2. Set up data retention policies.

To adjust the preconfigured data retention period of 365 days to your requirements, you should change the values of the predefined policy set or add new policies in SSAM. You will find all necessary information about the policy concept and the configuration of the policy parameters in 5.2.1, “Archive copy group retention parameters” on page 170.

3. Configure encryption (if required).

If you need to encrypt the data transfer from the client to the SSAM server, or to encrypt data that is stored on tapes, you have to proceed with some configuration. Refer to 5.2.7, “Expiration processing” on page 178 for more information about the encryption of the client server connection, and to 5.2.11, “Device support for data retention” on page 188 for more information about the tape drive encryption.

4. Configure your archive clients.

You enter the server name or IP address of your SSAM server into the configuration files of your archive clients. Go to the Tivoli information center, Storage Manager Clients, for more information, at:

<http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp>

3.7 Required configuration tasks for FSG

The DR550 File System Gateway can be installed in two separate configurations: stand-alone or as a high availability cluster. We have explained the manufacturing settings for the FSG in 2.6, “What is preconfigured on the optional FSG” on page 56. In summary, the FSG is preconfigured with:

- ▶ Two disks used for the operating system and configured as RAID 1.
- ▶ Four data space and cache disks configured as RAID 6.
- ▶ Disks are partitioned.
- ▶ Operating system SLES10 is installed.
- ▶ FSG software is preinstalled but not configured.
- ▶ SLES10 firewall is disabled. (Integration with firewall is not supported in the first release.)
- ▶ Host name is set to main or supplementary (for the second node of an HA configuration) and must not be changed.

To customize the FSG for your environment, the following configuration tasks are required before deployment:

- ▶ Prepare SSAM for FSG attachment.
- ▶ Attach the FSG to the network.
- ▶ Network address configuration. This integrates the File System Gateway into existing networks.
- ▶ Configure the time zone.
- ▶ Configuration of Network Time protocol (NTP) connectivity; mandatory for cluster configurations.
- ▶ Apply the additional FSG file sets and kernel patches for performance tuning.
- ▶ Customize FSG software.

3.7.1 Prepare DR550 SSAM for FSG attachment

If you order a DR550 with software bundle V4.0, then all necessary configuration steps for the FSG attachment are already included in the initial SSAM configuration.

To prepare DR550 SSAM for FSG attachment:

1. Log on to the DR550 engine1 with the dr550adm user ID.
2. Run the command **dsmadm** to start the SSAM command-line administration interface.
3. You will be prompted for the SSAM user ID. Log on with the SSAM user admin.

4. Check for the preconfigured FSG definition by entering the SSAM command **q node**. You will find the node name DRG-NODE if the FSG is preconfigured (see Example 3-10). In this case, you can skip this section and exit the SSAM session with the **quit** command; otherwise, go to step 5 and enter the FSG definition into the SSAM configuration.

Example 3-10 Output of q node command

tsm: TSM>**q node**

Node Name	Platform	Policy Domain Name	Days Since Last Acce-ss	Days Since Password Set	Locked?
DR550_1	Sample-- API	STANDARD	614	736	No
DRG-NODE	Bycast DRG	DRG-DOMAIN	<1	12	No

5. Enter the FSG definition into SSAM by typing the following SSAM commands:

```
define domain DRG-DOMAIN
define policyset DRG-DOMAIN STANDARD
define mgmtclass DRG-DOMAIN STANDARD DRG-DEFAULTMC
define copygroup DRG-DOMAIN STANDARD DRG-DEFAULTMC type=archive
dest=ARCHIVEPOOL retinit=event retmin=1 retver=0
define mgmtclass DRG-DOMAIN STANDARD DRG-SYSMC
define copygroup DRG-DOMAIN STANDARD DRG-SYSMC type=archive dest=ARCHIVEPOOL
retinit=event retmin=0 retver=0
assign defmgmtclass DRG-DOMAIN STANDARD DRG-DEFAULTMC
activate policyset DRG-DOMAIN STANDARD
register node DRG-NODE abcd2357 passexp=0 domain=DRG-DOMAIN
```

Note: The values `retmin=1` and `retver=0` are only examples. Other values might be used; however, `retmin` should be equal to the protection period of the default FSG profile. Ignore the no backup copy group warning that occurs when entering the **activate** command.

Important: SSAM will not allow reducing the retention intervals or modifying the retention initialization method of a copygroup after the corresponding policyset is activated. Archiveretentionprotection also disallows deletion of policy domains, policy sets, management classes, and copygroups after a new definition is committed (policyset activation).

6. Exit the SSAM session with the **quit** command.

3.7.2 Attaching the DR550 File System Gateway to the network

First, to avoid conflicts in an existing network environment, check that the IP addresses set at the factory (see Table 3-4 on page 118 for stand-alone FSG configurations or Table 3-5 on page 120 for dual-node FSG configurations) are not already in use by another device on your network. Note that there should rarely be a conflict, because the factory set IP addresses are usually not used in real network environments.

In the DR550 DR1 configuration, communication between the SSAM server and FSG nodes goes through the external (customer) network. PCI adapters in slots 1 and 2 are only used if the external network uses fiber Ethernet switches. The 2-port Gigabit Ethernet-TX adapters in slots 3 and 4 (half height) are used if the external network uses copper Ethernet switches. Port number 1 (the top most) on the 2-port Gigabit Ethernet-TX adapters is used for connection to the external network connection (see Figure 3-11 on page 116).

Both Ethernet ports on the PCI Ethernet adapter cards located in slot 3 and 4 of the FSG must be connected to the external (customer) network. The interface, 10/100/1000 copper or gigabit fiber, depends on the DR550 configuration. You will always have the same Ethernet interface for the DR550 SSAM Servers and the FSGs.

In DR550 DR2 configurations, communication between the SSAM server and FSG nodes goes through the private (DR550 internal) network. Port number 2 (bottom) of each of the 2-port Gigabit Ethernet-TX adapters in slots 3 and 4 (half height) is used for communication with the SSAM server(s). Port number 1 (top) of each of the 2-port Gigabit Ethernet-TX adapters in slots 3 and 4 (half height) is used for communication with external network servers; the public network uses copper Ethernet switches. If the external (customer) network uses fiber Ethernet switches, the Gigabit Ethernet-SX adapters in FSG PCI slots 1 and 2 must be installed to provide the connection to the external network.

The two onboard Ethernet interfaces are only used with the high availability FSG cluster. The first onboard Ethernet interface of the main FSG is wired directly with a cross-over cable to the first onboard Ethernet interface of the supplementary FSG node. The same cabling is done between the second on-board Ethernet adapter of the main and supplementary FSG.

See the illustrated cable diagram in Figure 3-11.

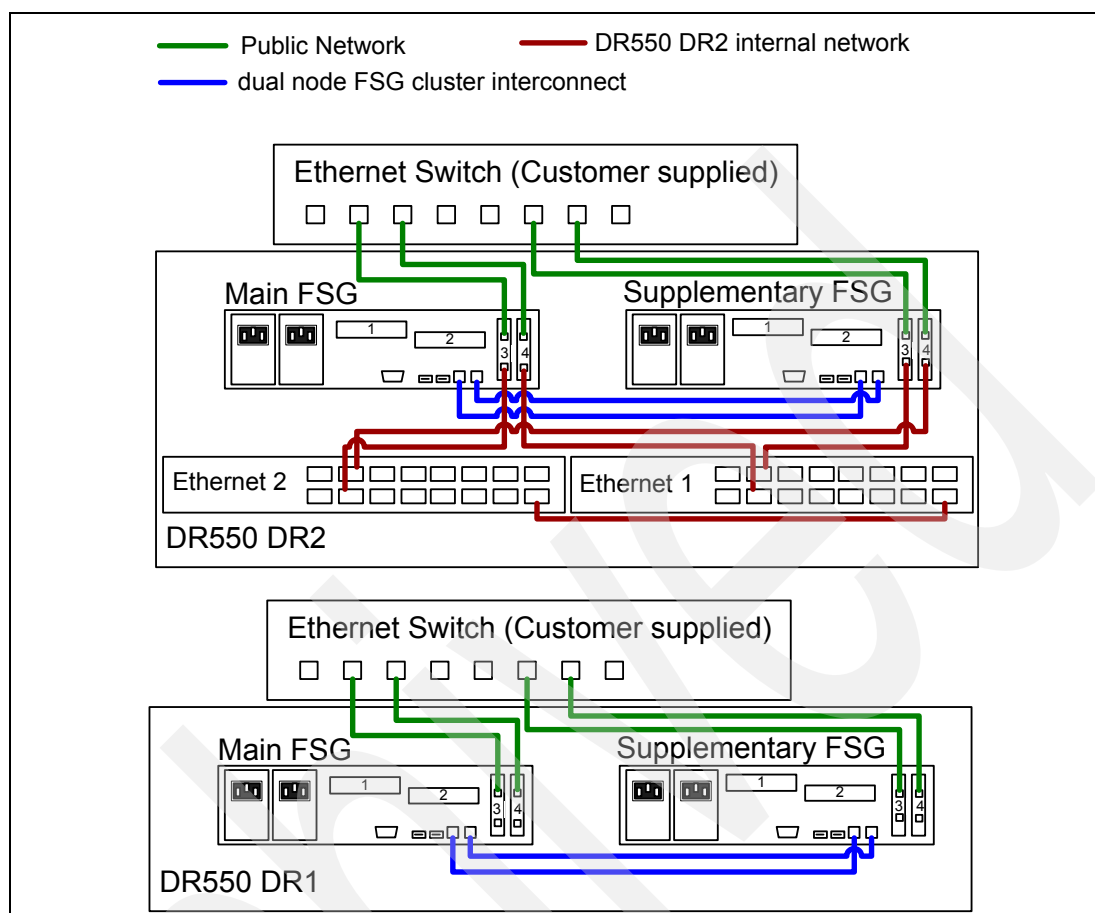


Figure 3-11 Network cabling for a dual-node FSG configuration

3.7.3 Configuring the network settings at the DR550 FSG

This section describes in detail how to set up the new IP addresses. The setup of the main FSG and the supplementary FSG are very similar.

Most settings are done using the graphical configuration utility *yast*, which is part of the SUSE Linux distribution.

Configuring the network settings for a stand-alone FSG node

For the main FSG, complete the following steps:

1. Make sure that the Ethernet cables (customer provided) are correctly connected in PCI port locations 0000:1c:00.0 and 0000:24:00.0 if you use copper interfaces and in PCI port locations 0000:08:03.0 and 0000:0f:04.0 if you use optical interfaces. See Figure 3-12 on page 117.

2. Obtain one IP address from your network administrator.

Set a new address in conformance with your network. The physical network adapter `eth0` and `eth1` are bonded to the logical network interface `bond0`.

Ethernet bonding refers to aggregating multiple Ethernet channels together to form a single channel. This is primarily used for redundancy of Ethernet paths.

The bonding of the IP address increases the reliability of the network attachment from the FSG. The application will communicate only with the logical interface bond0. This logical interface is connected to the two physical network interfaces eth0 and eth1 and the logical interface manages the failover of the network traffic from the primary network interface to the secondary physical network interface in case of an adapter, cable, or switch port failure.

Note: Do not change the preconfigured IP addresses for eth0 and eth1. These addresses are only placeholders and they are not used and not visible to the external (customer) network.

The FSG Linux operating system uses the udev facility for dynamic device mapping. udev maps a PCI slot location (see Figure 3-12) to a logical interface name such as eth0. The name mapping rules used by udev are located in the file /etc/udev/rules.d/30-net_persistent_names.rules on the FSG.




Figure 3-12 System x 2229 FSG rear view with I/O adapters and PCI slots locations and location codes

There are two sets of rules used on FSG. One is for a configuration with optical Ethernet adapters (Example 3-11) and one for a configuration with copper Ethernet adapters only (Example 3-12).

Example 3-11 Preconfigured 30-net_persistent_names.rules file for copper adapters only

```
SUBSYSTEM=="net", ACTION=="add", ID="0000:03:00.0", IMPORT="/lib/udev/rename_netiface %k eth2"
SUBSYSTEM=="net", ACTION=="add", ID="0000:1c:00.0", IMPORT="/lib/udev/rename_netiface %k eth0"
SUBSYSTEM=="net", ACTION=="add", ID="0000:24:00.0", IMPORT="/lib/udev/rename_netiface %k eth1"
SUBSYSTEM=="net", ACTION=="add", ID="0000:1c:00.1", IMPORT="/lib/udev/rename_netiface %k eth5"
SUBSYSTEM=="net", ACTION=="add", ID="0000:24:00.1", IMPORT="/lib/udev/rename_netiface %k eth4"
SUBSYSTEM=="net", ACTION=="add", ID="0000:06:00.0", IMPORT="/lib/udev/rename_netiface %k eth3"
```

Example 3-12 Preconfigured 30-net_persistent_names.rules file for optical and copper adapters

```
SUBSYSTEM=="net", ACTION=="add", ID="0000:03:00.0", IMPORT="/lib/udev/rename_netiface %k eth2"
SUBSYSTEM=="net", ACTION=="add", ID="0000:1c:00.0", IMPORT="/lib/udev/rename_netiface %k eth6"
SUBSYSTEM=="net", ACTION=="add", ID="0000:24:00.0", IMPORT="/lib/udev/rename_netiface %k eth7"
SUBSYSTEM=="net", ACTION=="add", ID="0000:08:03.0", IMPORT="/lib/udev/rename_netiface %k eth0"
SUBSYSTEM=="net", ACTION=="add", ID="0000:0f:04.0", IMPORT="/lib/udev/rename_netiface %k eth1"
SUBSYSTEM=="net", ACTION=="add", ID="0000:1c:00.1", IMPORT="/lib/udev/rename_netiface %k eth4"
SUBSYSTEM=="net", ACTION=="add", ID="0000:24:00.1", IMPORT="/lib/udev/rename_netiface %k eth5"
SUBSYSTEM=="net", ACTION=="add", ID="0000:06:00.0", IMPORT="/lib/udev/rename_netiface %k eth3"
```

Table 3-4 shows the translation table for IP address for a stand-alone FSG configuration with copper interfaces

Table 3-4 Translation table for IP address for a stand-alone FSG configuration with copper interfaces

Linux interface	Port Location	Factory IP address	Actual IP address
bond0	n/a	192.168.1.120	should be changed
eth0	0000:1c:00.0	192.168.1.121	do not change
eth1	0000:24:00.0	192.168.1.122	do not change
eth2	0000:03:00.0	169.254.1.1	do not change
eth3	0000:06:00.0	169.254.1.5	do not change
eth4	0000:24:00.1	n/a	do not change
eth5	0000:1c:00.1	n/a	do not change

3. Connect to the main FSG through the local console.

Log in to the main FSG with user fsgadm through the local console.

Switch the user to root by entering **su - root**, and enter root's password.

Note: Steps 4 and 5 are normally checked by the manufacturer.

4. Issue the **ifconfig -a** command and make note of what adapter MAC address corresponds to what interface device: eth0 and eth1 should correspond to the adapters in PCI slots 1 and 2 respectively, while eth2 and eth3 correspond to the onboard Ethernet adapters.
5. View the 30-net_persistent_names.rules file (in /etc/udev/rules.d) to validate the logical adapter name assignment based on the location of the adapter.
6. Change the TCP/IP address for the interface bond0:
 - a. Issue the **vi /etc/sysconfig/network/ifcfg-bond0** command to edit the configuration file of the bond0 interface.
 - b. Replace the factory IP address in the IPADDR field with the new IP address and modify the netmask if it is necessary (see Example 3-13). The netmask will be entered in a bitwise format; 24 represents a netmask of 255.255.255.0.
 - c. Save the file and exit the vi editor.
 - d. Enter **/etc/init.d/network restart** to restart the bond0 interface and activate the new IP address.

Example 3-13 Preconfigured /etc/sysconfig/network/ifcfg-bond0 file

```

BOOTPROTO='static'
IPADDR='192.168.1.120/24' <- Replace with your IP address/netmask
REMOTE_IPADDR=''
STARTMODE='onboot'
BONDING_MASTER='yes'
BONDING_MODULE_OPTS='mode=1 miimon=100 downdelay=200 updelay=200'
BONDING_SLAVE0='eth0'
BONDING_SLAVE1='eth1'

```

7. Check the status of the bonded network interfaces.

To ensure that both physical network interfaces are working, you can check the status with the command **cat /proc/net/bonding/bond0** (see Example 3-14). The status of all interfaces should be up.

Example 3-14 Output of cat /proc/net/bonding/bond0

```
main:/ # cat /proc/net/bonding/bond0
Ethernet Channel Bonding Driver: v3.0.1 (January 9, 2006)

Bonding Mode: fault-tolerance (active-backup)
Primary Slave: None
Currently Active Slave: eth0
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 200
Down Delay (ms): 200

Slave Interface: eth0
MII Status: up
Link Failure Count: 0
Permanent HW addr: 00:1a:64:23:64:7e

Slave Interface: eth1
MII Status: up
Link Failure Count: 0
Permanent HW addr: 00:1a:64:23:64:80
```

8. Enter the IP address of the default gateway in the network for the FSG.

- a. In the command line, enter the command **yast2 routing &**. The graphical configuration interface of SUSE Linux displays.
- b. Enter the IP address of your default gateway into the Default Gateway field and click **Finish**.
- c. Enter **/etc/init.d/network restart** at the command line to activate the changes.

9. Configure the DNS name server IP address (if required).

- a. In the command line, enter the command **yast2 dns &**.
- b. The graphical configuration interface of SUSE Linux displays (see Figure 3-13 on page 122). Do not change the local host name. Enter your domain name and the IP address of your name server in the appropriate field and click Finish.

Important: Do not change the local host name of the FSG. The FSG application is expecting main as the local host name for the Main FSG.

Configuring network settings for dual-node FSG configurations

Follow these steps:

1. Make sure that the Ethernet cables (customer provided) are correctly connected in PCI port locations 0000:1c:00.0 and 0000:24:00.0 if you use copper interfaces and in PCI port locations 0000:08:03.0 and 0000:0f:04.0 if you use optical interfaces. See Figure 3-12 on page 117.

2. Obtain three IP addresses from your network administrator.

Set a new address in conformance with your network. The physical network adapter eth0 and eth1 are bonded to the logical network interface bond0.

Note: Do not change the preconfigured IP addresses for eth0 and eth1. These addresses are only placeholders and they are not used and are not visible to the public network.

We suggest that you create a table where you write the factory IP address and the actual IP address side by side, as illustrated in Table 3-5.

Table 3-5 Translation table for IP address with dual-node FSG configurations

Server/location	Linux interface	Factory IP address	Actual IP address	FSG attributes
main	bond0	192.168.1.120	100.100.51.220.	MNCI
main 0000:1c:00.0	eth0	192.168.1.121	Do not change.	
main 0000:24:00.0	eth1	192.168.1.122	Do not change	
main 0000:03:00.0	eth2	169.254.1.1	Do not change.	MNGI
main 0000:06:00.0	eth3	169.254.1.5	Do not change.	
supplementary	bond0	192.168.1.130	100.100.51.230.	SPCI
supplementary 0000:1c:00.0	eth0	192.168.1.131	Do not change.	
supplementary 0000:24:00.0	eth1	192.168.1.132	Do not change.	
supplementary 0000:03:00.0	eth2	169.254.1.2	Do not change.	SPGI
supplementary 0000:06:00.0	eth3	169.254.1.6	Do not change.	
Virtual cluster IP address		192.168.1.140	100.100.51.240.	

3. Connect to the main FSG through the local console.

Log in to the main FSG with user root through the local console.

4. Change the TCP/IP address for the interface bond0

- a. Issue the `vi /etc/sysconfig/network/ifcfg-bond0` command to edit the configuration file of the bond0 interface.
- b. Replace the factory IP address in the IPADDR field with the new IP address and modify the netmask if it is necessary (see Example 3-13 on page 118). The netmask will be entered in a bitwise format. 24 represents a netmask of 255.255.255.0.
- c. Save the file and exit the vi editor.

Note: In a DR550 dual-node with dual FSG configuration, there is also a bond1 interface configured, which is used exclusively for communication between FSG and SSAM. The IP address assigned to the bond1 interface must not be changed.

- d. Enter `/etc/init.d/network restart` to restart the bond0 interface and activate the new IP address.

Example 3-15 shows the preconfigured `/etc/sysconfig/network/ifcfg-bond0` file.

Example 3-15 Preconfigured `/etc/sysconfig/network/ifcfg-bond0` file

```
BOOTPROTO='static'
IPADDR='192.168.1.120/24' <- Replace with your IP address/netmask
REMOTE_IPADDR=''
STARTMODE='onboot'
BONDING_MASTER='yes'
BONDING_MODULE_OPTS='mode=1 miimon=100 downdelay=200 updelay=200'
BONDING_SLAVE0='eth0'
BONDING_SLAVE1='eth1'
```

5. Configure the default gateway for the FSG.
 - a. From the command line, enter the command `yast2 routing &`. The graphical configuration interface of SUSE Linux displays.
 - b. Enter the IP address of your default gateway into the field Default Gateway and click **Finish**.
 - c. Enter `/etc/init.d/network restart` at the command line to activate the changes.
 6. Configure the DNS name server IP address (if required).
 - a. From the command line, enter the command `yast2 dns &`.
 - b. The graphical configuration interface of SUSE Linux displays (see Figure 3-13 on page 122). Do not change the local host name! Enter your domain name and the IP address of your name server in the appropriate field and click **Finish**.
- Important:** Do not change the local host name of the FSG. The FSG application is expecting `main` as the local host name for the main FSG and `supplementary` as the local host name for the supplementary FSG.
7. Connect to the supplementary FSG through the management console.

Log in to the supplementary FSG with user `root` through the management console.
 8. Repeat steps 3-8 from the procedure above for the supplementary FSG to set up the IP address, default gateway, and DNS server for the supplementary FSG.

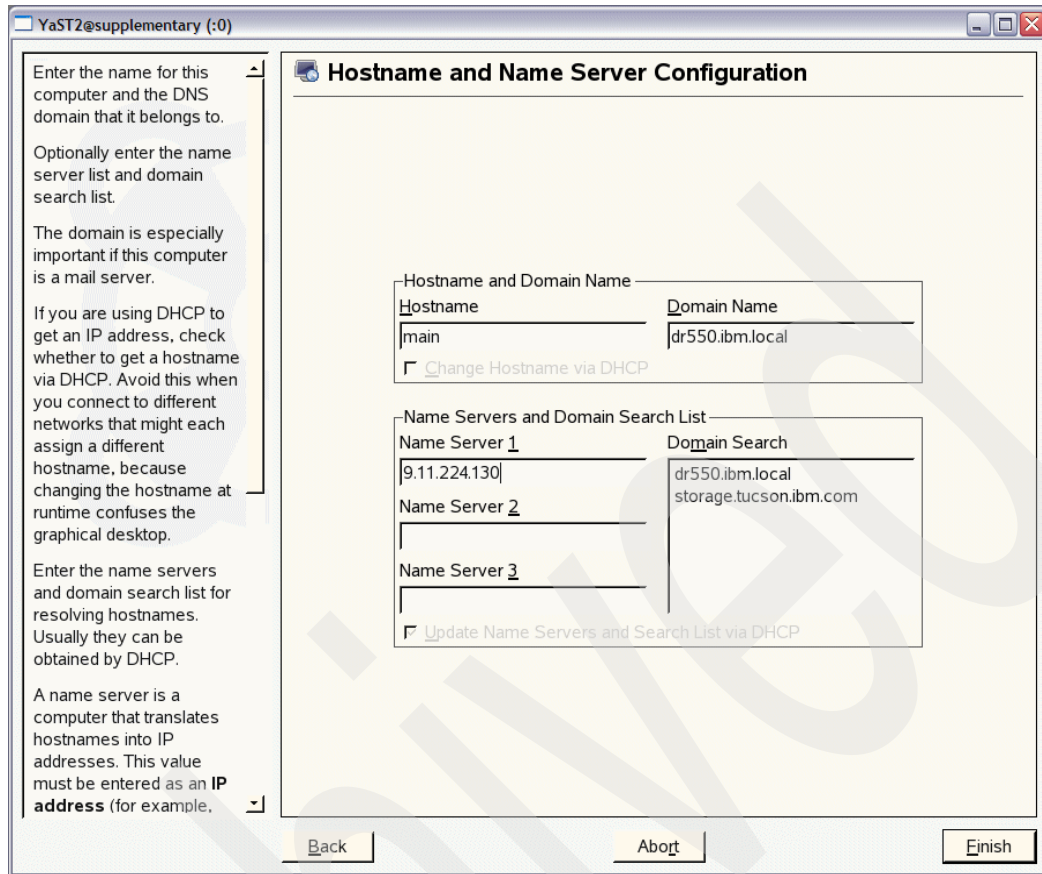


Figure 3-13 The yast2 dns window

3.7.4 Time zone configuration and NTP setup at the DR550 FSG node

The setup of the time zone and the setup of the NTP connectivity is identical for the main FSG node and the supplementary FSG node. If you have a dual-node FSG configuration, then you are required to configure the NTP connectivity on both nodes. The NTP protocol is for synchronizing the clocks of the FSGs. For a stand-alone FSG node, this configuration step is optional.

Follow these steps:

1. Log in as the root user to the main FSG node at the management console.
2. Set the correct time zone for your geographical region:
 - a. From the command line, enter the command **yast2 timezone &**. The graphical configuration interface of SUSE Linux displays.
 - b. Select your region first and then your time zone and click **Accept**. See Figure 3-14 on page 123.

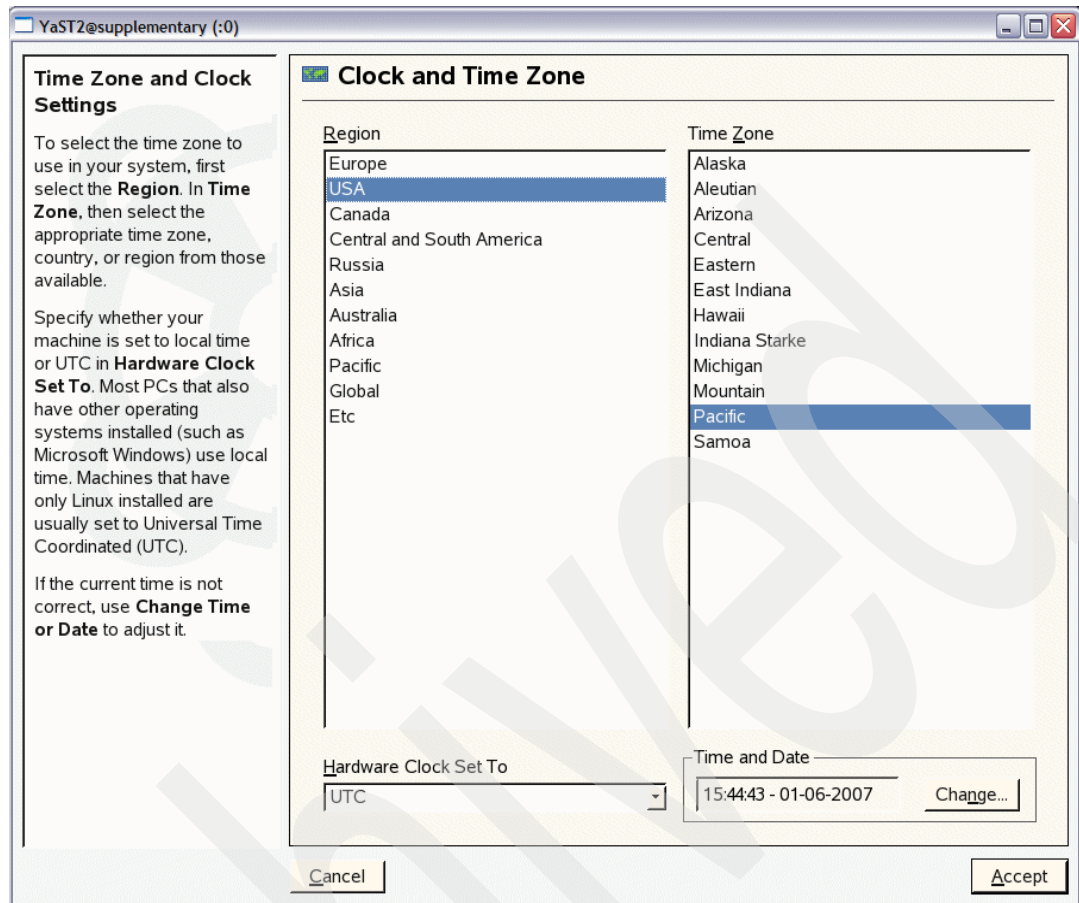


Figure 3-14 The yast2 time zone

3. Configure the NTP time server connection. (This is an optional step for stand-alone FSG node configurations, but is required for dual-node FSG configurations.)
 - a. From the command line, enter the command **yast2 ntp-client &**. The graphical configuration interface of SUSE Linux displays.
 - b. Select **During Boot** in the Automatically Start NTP Daemon field.
 - c. Enter the IP address or the host name of your NTP server in the Address field and click **Test** to verify the connectivity.
 - d. If the test was successful, you can leave the configuration window by clicking **Finish**.
4. Verify the NTP server's connectivity by opening the log file of the NTP daemon.
 - a. Enter **ps -ef |grep ntpd** at the command line. Here is an example:


```
main:/ # ps -ef |grep ntpd
ntp      11180      1  0 May15 ?        00:00:00 /usr/sbin/ntpd -p
/var/lib/ntp/var/run/ntp/ntpd.pid -u ntp -i /var/lib/ntp
root     8022  5288  0 11:01 pts/2    00:00:00 grep ntpd
```
 - b. Stop the running ntpd daemon by killing the process. Enter the **kill** command followed by the process number. In the case above, run **kill 11180**.

- c. From the command line, enter the command **ntpd -q**.

If the ntpd daemon can communicate to the NTP server, you will get a response about the time difference between the FSG and the NTP server. Here is an example where the communication to the NTP server was successful:

```
main:/ # ntpd -q
ntpd: time skew +0.005796s
```

- d. Restart the NTP daemon with the command **/etc/init.d/ntp start**.
- e. If the synchronization of the clock was successful, you will find a line at the bottom of the log file that is similar to the line in bold in Example 3-16.

Example 3-16 Output of the tail /var/log/ntp

```
14 May 10:34:48 ntpd[5109]: synchronized to LOCAL(0), stratum 10
14 May 10:34:48 ntpd[5109]: kernel time sync disabled 0041
14 May 10:35:54 ntpd[5109]: synchronized to 148.167.132.200, stratum 2
```

Repeat steps 1-4 for your supplementary FSG node if you have a dual-node FSG configuration.

3.7.5 Check network connectivity

At this stage, it is a good idea to verify proper network connectivity by issuing a **ping** command from the FSG (from each of the FSG nodes in a high availability cluster) to the IP address of your DNS according to the IP addressing schema example shown in Figure 3-15.

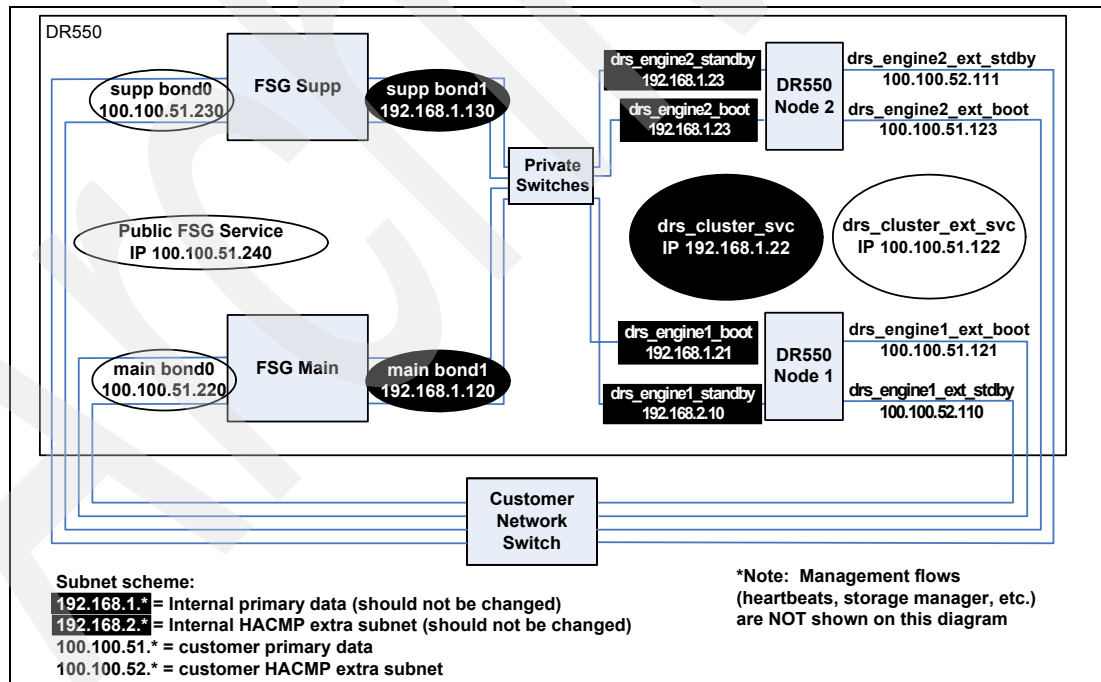


Figure 3-15 DR550 DR2 dual node with dual FSG data traffic - IP addressing scheme example

3.7.6 Customize the FSG software

In this section, we describe the customization process of the FSG node or nodes. Follow these steps:

1. Connect to the main FSG at the management console and log in with user root.

2. Copy the appropriate resources file:

- If using a stand-alone FSG, enter the command:

```
cp /usr/local/drg/resources.sngl /usr/local/drg/resources
```

- If using a high availability FSG, enter the command:

```
cp /usr/local/drg/resources.main /usr/local/drg/resources (on the main node)
```

```
cp /usr/local/drg/resources.suppl /usr/local/drg/resources (on the supplementary node)
```

3. Enter the command **enable_drg**.

This command will replace some kernel modules of the SUSE Linux and install some additional filesets at the FSG node. There are also some tuning parameters that will be modified by this command to optimize Linux for the usage of the FSG software.

4. A license agreement text is displayed. Read the license agreement and accept it by typing **I agree** into the input field and then pressing Enter twice.

5. An additional license agreement window displays. Accept it by clicking the **I Agree** button.

6. The FSG node will reboot automatically after the successful customization. When the node is rebooted, the FSG software will start automatically.

If you have a supplementary node, connect to the node, log in with root at the supplementary node, and repeat steps 2-5.

The next step is to configure your FSG software.

3.7.7 FSG post-installation and initial setup

Additional configuration within the FSG is still required to allow client machines to store and retrieve data from the FSG and allow the FSG to communicate with the SSAM storage on the DR550. This consists of configuring the DRG service, which is the main software component running at the FSG (the DRG service is explained in 6.2, “DR550 FSG architecture overview” on page 230).

Important: Make sure that your SSAM server is running and accessible. The FSG (stand-alone or HA configuration) will not be operational until it can effectively connect to the SSAM server.

Configuration of the DRG service is done by editing a profile file known as the DRGC bundle.

Note: The DRGC bundle is the main configuration file on a DR550 FSG and contains various parameters pertinent to the system, including networking information, SSAM parameters such as server name, and server password.

Table 3-6 lists some of the DRGC bundle attributes.

Table 3-6 FSG software (or DRG service) configuration attributes

Attribute Name	Preconfigured value	Description
BTIM	UI64	Bundle Time: The time (in microseconds since the epoch 00:00:00 UTC on January 1 1970) at which the bundle was updated.
BVER	UI64	Bundle Version: Indicates the bundle version. The version number is incremented when the format of fields is changed in a non-compatible way.
CRGC	CONT	Bundle container.
HTBT	CONT	Heartbeat Configuration Branch: This branch contains items used to generate Heartbeat configuration files.
MCLS	DRG-SYSMC	Default Management Class: The class will be used by SSAM API. The management class must exist on the SSAM server.
MNCI		Main node client IP: The IP address to public client network for the main FSG node. This is the IP address of the bond0 network interface.
MNGI	169.254.1.1	Main node Grid IP: IP address for intercluster communication between the main and supplementary nodes. This should be 169.254.1.1.
NDNM	DRG-NODE	SSAM node name: This is the node name of the FSG in the SSAM configuration.
PNGT	0.0.0.0	Pingtest IP: This is only used in dual-node FSG configurations. An IP address on the client network will be periodically pinged (ICMP echo) in order to monitor connectivity. We recommend that you use the default gateway IP address for this value. Make sure that your firewall rules allow a ping. In stand-alone configurations, this has a value of 0.0.0.0.
SPCI	0.0.0.0	Supplementary Node Client IP: This is only used in dual-node FSG configurations. It is an IP address to public client network for the supplementary FSG node. This is the IP address of the bond0 network interface. In stand-alone configurations, this has a value of 0.0.0.0.
SPGI	0.0.0.0	Supplementary Node Grid IP: This is only used in dual-node FSG configurations. This is the IP address for intercluster communication between the main and supplementary nodes. In stand-alone configurations, this has a value of 0.0.0.0. This should be changed to 169.254.1.2. in dual-node configurations.
SVFS	/drg-10	File Space Name: The file space name under which objects from DR550 FSG will be stored.

Attribute Name	Preconfigured value	Description
SVIP	192.168.1.11	SSAM Server IP: This is the IP address of your DR550 SSAM server. In DR550 dual-engine configurations, you must enter the HACMP service IP address here.
SVNM	server_a	SSAM Server Name: The SSAM server name as defined in /opt/tivoli/tsm/client/api/bin/dsm.sys.
SVPW	abcd2357	SSAM password.
SVUS	DRG-NODE	SSAM User Name: The user associated with the DR550 FSG or FSG cluster.
VINM	0	Virtual IP Netmask: This is only used in dual-node FSG configurations. Values from 0 to 31 are allowed. The value of 24 represents a netmask of 255.255.255.0.
VTIP	0.0.0.0	Virtual IP Address: The virtual IP address to be shared by the clustered FSG gateways. All client access to the FSG must occur through this address. For stand-alone configurations, the VTIP should have the same value as the value of MNCL.
WNBD	FC32	Winbind Support: Enables support for managing the winbind daemon (required for Active Directory group support). Samba must be correctly configured and joined to a domain before winbind support is enabled. The values of this parameter are: <ul style="list-style-type: none"> ► ENBL: Enable winbind. ► DABL: Disable winbind.

The DRGC bundle is configured as an XML file that can be updated with any text editor. After updating the DRGC file, save it to the /bundle-import/import directory. XML files are automatically imported into the system and converted back to a bundle. To facilitate this, the DR550 File System Gateway is configured with the following directories:

- **Current:** Bundles currently in use by the DR550 File System Gateway.
- **Import:** Place XML files here so that they can be used by the DR550 File System Gateway.
- **Imported:** Files successfully imported into the DR550 File System Gateway.
- **Invalid:** Files not imported into the DR550 File System Gateway.

Important: In a high availability FSG cluster, the DRGC bundle is automatically copied from the Main to the Supplementary. Both bundles must be identical. However, if you make any changes to the IP configuration, you will have to manually import the DRGC file into the Supplementary FSG.

Configuring the DRGC bundle for a stand-alone FSG

In this section we describe the configuration settings to connect a stand-alone FSG to the DR550. If you have a dual-node FSG configuration, skip this section and proceed directly to “Configuring the DRGC bundle for a high availability FSG cluster” on page 129.

Follow these steps:

1. Connect to FSG and log in as fsgadm, then **su** to the root user.
2. Copy the original cluster bundle configuration file into a temporary directory. At the command line, enter the command:

```
cp /var/local/bundle-import/current/DRGC /tmp/install
```
3. Edit the file with the command **vi /tmp/install/DRGC** and customize all the settings for your specific environment. Save the file and exit the editor when done.

You must change the following attributes in a stand-alone FSG configuration:

- SVIP (SSAM Server IP)
- MNCI (Main Node Client IP)
- VTIP (Virtual IP Address, which is the same value as MNCI in stand-alone configurations)
- VINM (Virtual IP Address Netmask; possible values 0-31, value 24 = 255.255.255.0)

All other attributes can stay at their default values.

Refer to Table 3-6 on page 126 to see other possible attributes that could be changed.

4. Copy the modified configuration file to the import directory:

```
cp /tmp/install/DRGC /var/local/bundle-import/import
```

The DR550 FSG software automatically imports the updated file and moves it to the /var/local/bundle-import/imported directory. If the file is rejected by the system, the file is moved to the /var/local/bundle-import/invalid directory. In this case, you have to recheck your edited DRGC file.

5. Enter **cat /var/local/bundle-import/current/DRGC** and verify that this file reflects the new settings.
6. Enter the following two commands:

```
ln -s /opt/tivoli/tsm/client/api/bin/dsm.sys /opt/tivoli/tsm/client/ba/bin/dsm.sys  
ln -s /opt/tivoli/tsm/client/api/bin/dsm.opt /opt/tivoli/tsm/client/ba/bin/dsm.opt
```

7. Reboot the main node by entering the command **reboot**.
8. To verify the status of the FSG node, you should log in to the main FSG as root and enter the command **crm_mon -i 15**. The main node should be in an online state and all services should have the status started main (Refer to Example 3-17 on page 129.) Exit the command with **Ctrl-C**.

Note: Make sure that your SSAM server is running and accessible. The FSG will not be operational until it can effectively connect to the SSAM server.

Example 3-17 Output of the `crm_mon -i 15` command for stand-alone configurations

```
=====
Last updated: Tue May  8 18:10:48 2007
Current DC: main (4da2fba3-ea94-5f03-9aaf-d1f14b65b0f0)
1 Nodes configured.
1 Resources configured.
=====

Node: main (4da2fba3-ea94-5f03-9aaf-d1f14b65b0f0): online

Resource Group: res_group
  drgbase      (Bycast::ocf:FSGBase): Started main
  ip_resource  (heartbeat::ocf:IPAddr2): Started main
  pingtest     (Bycast::ocf:PingTest): Started main
  drgactive    (Bycast::ocf:FSGActive): Started main
  Filesharing-NFS (Bycast::ocf:NFSServer): Started main
  Filesharing-CIFS (Bycast::ocf:Samba): Started main
```

Tip: If you need to stop and restart the cluster, do so by issuing the command `/etc/init.d/drg restart`.

Continue the configuration by referring to 3.7.8, “Testing connectivity and initial backup” on page 131.

Configuring the DRGC bundle for a high availability FSG cluster

In a clustered high availability FSG configuration, we have one active node that is usually the main node and one standby node, which is usually the supplementary node. If the active node is failing due to a hardware defect, the standby node will automatically take over and become the active FSG node. The application is always reachable through the virtual cluster IP address. This IP address points to the active FSG node even if it is the main or the supplementary FSG node.

All client applications must use the virtual cluster IP address to access the FSG.

There is no automatic fallback functionality. If you want to fail back from the supplementary FSG node to the main FSG node, you have to initiate the fail back manually by stopping and restarting the FSG software at the active node. This procedure is described in detail in 4.3, “Starting/Stopping the FSG” on page 151.

Configuration steps 1-6 have to be run at the main node only; as stated before, the DRG service will automatically replicate the DRGC bundle configuration at the supplementary node:

1. Connect to the main FSG and log in as the root user.
2. Copy the original cluster bundle configuration file into a temporary directory. At the command line, enter:

```
cp /var/local/bundle-import/current/DRGC /tmp/install
```

3. Edit the file with the command:

```
vi /tmp/install/DRGC
```

and customize all settings to your specific environment. Save the file and exit the editor when done.

Example 3-18 shows the original contents of the DRGC file.

Example 3-18 Output of the original DRGC file

```
<?xml version="1.0" encoding="UTF-8"?>
<containerxml version="1"
xmlns="http://www.bycast.com/schemas/XML-container-1.0.0">
  <container name="DRGC">
    <atom name="BVER" type="UI32" value="1"/>
    <atom name="BTIM" type="UI64" value="1177545891154500"/>
    <container name="TVSM">
      <atom name="SVNM" type="CSTR" value="server_a"/>
      <atom name="SVIP" type="IP32" value="192.168.1.11"/>
      <atom name="NDNM" type="CSTR" value="DRG-NODE"/>
      <atom name="SVUS" type="CSTR" value="DRG-NODE"/>
      <atom name="SVPW" type="CSTR" value="abcd2357"/>
      <atom name="MCLS" type="CSTR" value="DRG-SYSMC"/>
      <atom name="SVFS" type="CSTR" value="/drg-10"/>
    </container>
    <container name="HTBT">
      <atom name="MNNI" type="UI32" value="2310001"/>
      <atom name="SPNI" type="UI32" value="2310002"/>
      <atom name="MNGI" type="IP32" value="169.254.1.1"/>
      <atom name="SPGI" type="IP32" value="0.0.0.0"/>
      <atom name="MNCI" type="IP32" value="192.168.1.120"/>
      <atom name="SPCI" type="IP32" value="0.0.0.0"/>
      <atom name="VTIP" type="IP32" value="0.0.0.0"/>
      <atom name="VINM" type="UI32" value="0"/>
      <atom name="PNGT" type="IP32" value="0.0.0.0"/>
      <atom name="WNBD" type="FC32" value="DABL"/>
    </container>
  </container>
</containerxml>
```

You must change the following attributes in a dual-node FSG configuration:

- SVIP (SSAM Server IP)
- VTIP (Virtual IP Address)
- VINM (Virtual IP Address Netmask, possible values 0-31, value 24 = 255.255.255.0)
- MNCI (Main Node Client IP)
- SPCI (Supplementary Node Client IP)
- PNGT (Pingtest IP)
- MNGI (Main Node Grid IP, value should be 169.254.1.1)
- SPGI (Supplementary Node Grid IP, value should be 169.254.1.2)

All other attributes can keep their default values.

Refer to Table 3-6 on page 126 to see other attributes that can possibly be changed, and see Example 3-18 for the original contents of the DRGC file.

4. Copy the modified configuration file to the import directory with the command **cp /tmp/install/DRGC /var/local/bundle-import/import**. The DR550 FSG automatically imports the updated file and moves it to the /var/local/bundle-import/current directory. If the import of the file was successful, the configuration file will automatically be copied to the supplementary node. If the file is rejected by the system, the file is moved to the

/var/local/bundle-import/invalid directory. In this case, you have to recheck your edited DRGC file.

5. Enter **cat /var/local/bundle-import/current/DRGC** and verify that this file reflects the new settings.
6. Reboot the main node by entering the command **reboot**.
7. Start a reboot of the supplementary node when the reboot of the main node is completed. Connect to the supplementary node and log in with the root user. Enter **reboot** at the command line.
8. Both nodes will join the cluster and start the FSG application. The first node that comes up will be the active node and the second will be the standby node. To verify the status of the FSG cluster, you should log in to the main FSG with the root user and enter the command **crm_mon -i 15**. The main and the supplementary nodes should be in an online state and all services should run at the main FSG node (refer to Example 3-19). Exit the command with CTRL-C.

Note: Make sure that your SSAM server is running and accessible. The FSG will not be operational until it can effectively connect to the SSAM server.

*Example 3-19 Output of the **crm_mon -i 15** command for dual-node FSG configurations*

```
=====
Last updated: Tue May 15 13:47:59 2007
Current DC: main (4da2fba3-ea94-5f03-9aaf-d1f14b65b0f0)
2 Nodes configured.
1 Resources configured.
=====

Node: main (4da2fba3-ea94-5f03-9aaf-d1f14b65b0f0): online
Node: supplementary (430e222e-c5bf-84b8-22e0-5fb725f1753b): online

Resource Group: res_group
  drgbase      (Bycast::ocf:FSGBase): Started main
  ip_resource  (heartbeat::ocf:IPAddr2): Started main
  pingtest    (Bycast::ocf:PingTest): Started main
  drgactive    (Bycast::ocf:FSGActive): Started main
  Filesharing-NFS (Bycast::ocf:NFSserver): Started main
  Filesharing-CIFS (Bycast::ocf:Samba): Started main
```

Note: If you are just starting the cluster, you must wait long enough (usually more than one minute) for the final status of the cluster to be displayed.

3.7.8 Testing connectivity and initial backup

We recommend that you do an initial backup of the FSG metadata to the DR550 SSAM. The goal at this stage is mainly to test the connectivity between FSG and DR550 SSAM to ensure that the archiving of data works before deploying the system for production.

Backup for stand-alone FSG node

Follow these steps:

1. Log in to the main FSG node with the drg user and switch the user to root.
2. At the command line, enter **/usr/local/drg/forcebackup**. This command will back up the metadata for the files at the FSG and the file system structure at the FSG to the DR550 SSAM server storage pool.
3. To check the log file (for the **forcebackup** command), enter:

```
cat /var/local/attributes/drg/PBKI.xml
```

If the value of the field CSTR contains any value other than N/A or 0, then the backup was successful and the connection between FSG and SSAM is working. (See Example 3-20 for a sample output of a successful **forcebackup**.)

Backup for high availability FSG configuration

1. Log in to the active FSG node, using the virtual cluster IP address, with:
- ```
ssh virt.clusterIP_address
```
2. At the command line, enter **/usr/local/drg/forcebackup**. This command will back up the metadata for the files at the FSG and the file system structure at the FSG to the DR550 SSAM server storage pool.
  3. To check the log file (for the **forcebackup** command), enter:

```
cat /var/local/attributes/drg/PBKI.xml
```

If the value of the field CSTR contains any value other than N/A or 0, then the backup was successful and the connection between FSG and SSAM is working. (See Example 3-20 for a sample output of a successful **forcebackup**.)

#### *Example 3-20 Output of the PBKI.xml file after a successful forcebackup*

---

```
main:/var/local/bundle-import/invalid # cat /var/local/attributes/drg/PBKI.xml
<?xml version="1.0" encoding="UTF-8"?>
<containerxml version="1"
xmlns="http://www.bycast.com/schemas/XML-container-1.0.0">
 <container name="PBKI">
 <atom name="AVER" type="UI32" value="2"/>
 <atom name="AVTP" type="FC32" value="CSTR"/>
 <atom name="APER" type="FC32" value="READ"/>
 <atom name="ATIM" type="UI64" value="1179265643418089"/>
 <container name="AVAL">
 <atom name="0x00000000" type="UI32" value="2"/>
 <atom name="0x00000001" type="CSTR"
value="594EC50C-FFA5-4A3D-B893-F060F2178B9D"/>
 </container>
 </container>
</containerxml>
```

---

The File System Gateway is now configured for operation. The next step is to customize the retention protection settings (see 6.3, “File System Gateway administration and operations” on page 238) and to integrate the gateway with the application (see 6.4, “Integrating client applications with DR550 FSG” on page 245).

## 3.8 Changing passwords on base DR550

In this section, we describe how to change the passwords for the DR550 DR1 and DR2. Although not technically, absolutely required, we strongly recommend that you change the default passwords that were set at the factory. It is also a good practice to change the passwords on a regular basis.

Here we discuss changing the passwords for the AIX operating system and Tivoli Storage Manager users.

### AIX passwords

On the AIX operating system, you should change the passwords for all four preconfigured system accounts (dr550, dr550adm, ibmce, and root). Because of AIX security restrictions with the DR550 and DR550 Express, the most convenient way to change passwords is to log in with the user whose password you want to change, and then change the password; root is allowed to change passwords for other users, but then these users are forced to change their password again the next time they log in.

Therefore, the procedure for every distinct user account is to use the management console to access the (first) p5 (on the DR550, press the **PrtSc** button, select **Engine1** on **Port 01** from the panel, and then press **Enter**):

1. Log in to AIX using the dr550 user with the appropriate password.
2. Change the password of dr550 with the AIX command **passwd**. Follow the instructions on the screen.
3. Switch to the root user ID and its environment with the AIX command **su - root** and enter root's password.
4. Change the password of root with the AIX command **passwd**. Follow the instructions on the screen.
5. Switch to the dr550adm user ID and its environment with the AIX command **su - dr550adm**.
6. Change the password of dr550adm with the AIX command **passwd**. Follow the instructions on the screen.
7. Exit the shell back to the root environment with the AIX command **exit**.
8. Switch to the ibmce user ID and its environment with the AIX command **su - ibmce**.
9. Change the password of ibmce with the AIX command **passwd**. Follow the instructions on the screen.
10. Exit the shell back to the root environment with the AIX command **exit**.
11. Exit the shell back to the dr550 environment with the AIX command **exit**.
12. Exit the shell back to the AIX login screen with the AIX command **exit**.

For the DR550, use the management console to access the second p5 520 and repeat the above procedure for all four user IDs. Once changed, you should test the new passwords by logging in again.

## System Storage Archive Manager passwords

The System Storage Archive Manager accounts will be changed within the System Storage Archive Manager server. Use the System Storage Archive Manager administrative command line to type a command or use the Administration Center interface to change the passwords.

- ▶ For the account `admin`, you can use the command `update admin admin new_password` within an administrative command line. Or, set the new password within the Administration Center interface by selecting **Object view** → **Administrators** → **ADMIN** → **Operations: Update an Administrator**.
- ▶ For the account `hacmpadm` on the DR550, you can use the command `update admin hacmpadm new_password` within an administrative command line. Or, set the new password within the Administration Center interface by selecting **Object view** → **Administrators** → **CLIENT** → **Operations: Update an Administrator**.

Once changed, you should test the new passwords. You can test these by logging in the user to the Tivoli Storage Manager server again.

## DR550 Storage Controller (DS4200) Storage Manager password

The DS4000 Storage Manager allows you to set a password to prevent unauthorized users to perform certain actions. If the password protection is enabled, you have to enter the password for each Storage Manager session where a configuration change or other maintenance activities are attempted on the DR550 Storage Controller (DS4200) or the DR550 Expansion Drawers (EXP420). The monitoring is still possible without entering a password.

The password protection can be enabled and changed in the SMclient.

To set or change the password:

1. Connect to engine2 through the management console.
2. Log in to AIX with the `dr550` user.
3. Switch to root authority using the `su -` command.
4. Start the graphical desktop with the command `startx`.
5. Start the SMclient at the command line with the `SMclient` command.
6. Go to the menu bar and change the password by selecting **Storage Subsystem** → **Set Password**. (See Figure 3-16 on page 135.)

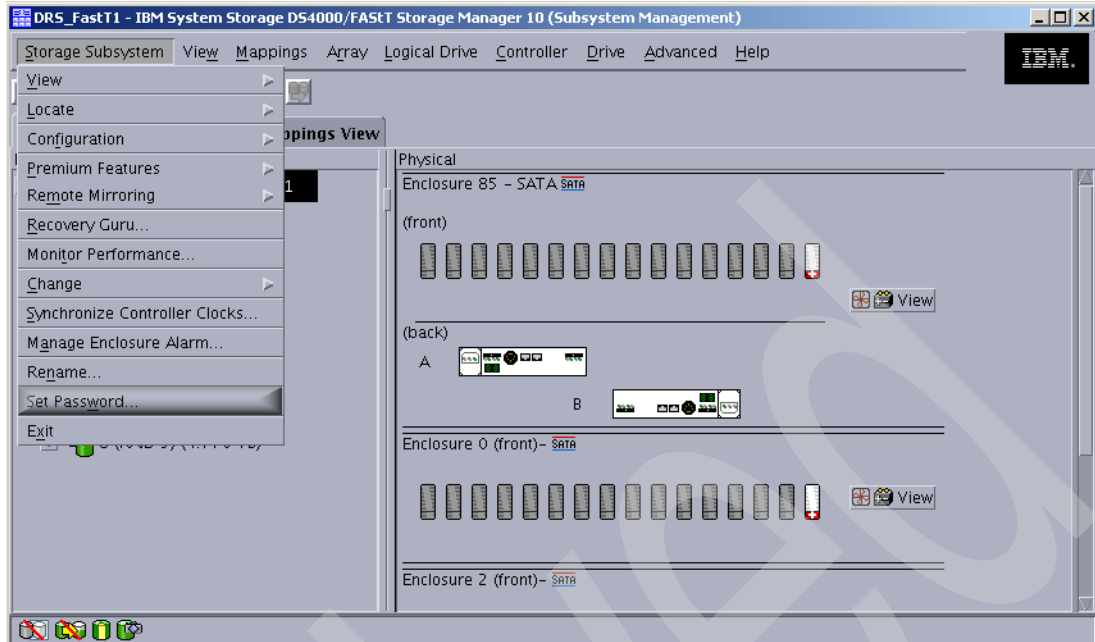


Figure 3-16 Change of the SMclient password

### 3.9 Changing passwords on the FSG

In this section, we describe how to change the passwords in the optional DR550 File System Gateway. We strongly recommend that you change all the default passwords that were set at the factory. On the SUSE Linux operating system, you should change all the passwords of the preconfigured system accounts (root, fsgadm, and drg).

To change the password on the FSG:

1. Connect to the main FSG and log in with fsgadm, then use the **su** - command to change to the root user.
2. Change the password of root with the command **passwd**. Follow the instructions on the screen.
3. Switch to the drg user ID and its environment with the command **su - drg**.
4. Change the password of drg with the **passwd** command. Follow the instructions on the screen.
5. Exit the shell with the **exit** command.
6. Switch to the fsgadm user ID and its environment with the command **su - fsgadm**.
7. Change the password of fsgadm with the **passwd** command. Follow the instructions on the screen.
8. Exit the shell with the **exit** command.

If you have a dual-node FSG configuration, you should proceed with the change of the passwords at the supplementary FSG node by following these steps:

1. Connect to the supplementary FSG and log in with the root user.
2. Repeat steps 2-5 from the procedure above to change the passwords at the supplementary FSG node.

Archived

## Operating the DR550 and FSG

This chapter provides information about operating the DR550 and FSG components. The chapter describes:

- ▶ The procedure and commands to start and stop the cluster services on the dual-node DR550 (2233-DR2) and to stop and start the IBM System Storage Archive Manager for all configurations of DR550 DR1 and DR2 models.
- ▶ The procedure and commands to start and stop the File System Gateway, including a description of the failover and failback procedures.
- ▶ A summary of post-installation checkout procedures for the DR550 and the FSG.

## 4.1 Starting and stopping HACMP cluster services

This section applies to the DR550 dual-node only.

Before starting or stopping cluster services on the DR550 dual-node configuration, you must have a thorough understanding of the node interactions and the impact on the system availability that starting or stopping cluster services causes. This section defines the HACMP cluster services and briefly describes the process for starting up or shutting down these services. We also provide a detailed step-by-step procedure. The processes described here for starting or stopping cluster services use the C-SPOC utility.

This information is reproduced in part from *HACMP for AIX: Administration Guide*, SC23-4862 and *HACMP for AIX: Troubleshooting Guide*, SC23-5177. You can obtain more details in those documents. Copies can be found on the Web at:

<http://publib.boulder.ibm.com/epubs/pdf/c2348629.pdf>

<http://publib.boulder.ibm.com/epubs/pdf/c2351773.pdf>

**Important:** The DR550 is not configured for HACMP to be automatically started with the AIX boot and has to be started manually after AIX is booted. HACMP will, however, be stopped automatically when AIX is shutting down.

### 4.1.1 Cluster services

The AIX System Resource Controller (SRC) controls the HACMP daemons (except for clclockd, which is a kernel extension). SRC provides a consistent interface for starting, stopping, and monitoring processes by grouping sets of related programs into subsystems and groups. In addition, it provides facilities for logging abnormal terminations of subsystems or groups and for tracing one or more subsystems.

The HACMP daemons are collected into the SRC subsystems and groups shown in Table 4-1.

Table 4-1 HACMP daemons

| Daemon                                 | Subsystem  | Group   | Function                          |
|----------------------------------------|------------|---------|-----------------------------------|
| /usr/sbin/rsct/bin/topsvcs             | topsvcs    | topsvcs | Cluster Topology Services         |
| /usr/sbin/rsct/bin/hagsglsm            | grpplsm    | grpsvcs | Cluster Group Services            |
| /usr/sbin/rsct/bin/haemd_HACMP         | emsvcs     | emsvcs  | Cluster Event Management          |
| /usr/sbin/rsct/bin/haemRM/harmad_HACMP | emaixos    | emsvcs  | Event Management AIX              |
| /usr/sbin/rsct/bin/rmcd                | ctrmc      | rsct    | Resource Monitor & Control Daemon |
| /usr/es/sbin/cluster/clstrmgr          | clstrmgrES | cluster | Cluster Manager                   |
| /usr/es/sbin/cluster/clinfo            | clinfoES   | cluster | Cluster Information Program       |
| /usr/es/sbin/cluster/clcomd            | clcomdES   | clcomdS | Cluster Communication             |



When using the SRC commands, you can control the clstrmgrES and clinfo daemons by specifying the SRC cluster group. Notice that if you list the daemons in the AIX System Resource Controller (SRC), you will see ES appended to their names. The actual executables do not have the ES appended; the process table shows the executable by path (/usr/es/sbin/cluster...).

## Understanding how to start cluster services

Figure 4-1 shows the main commands and scripts executed on both cluster nodes when cluster services are started in the cluster using the C-SPOC utility.

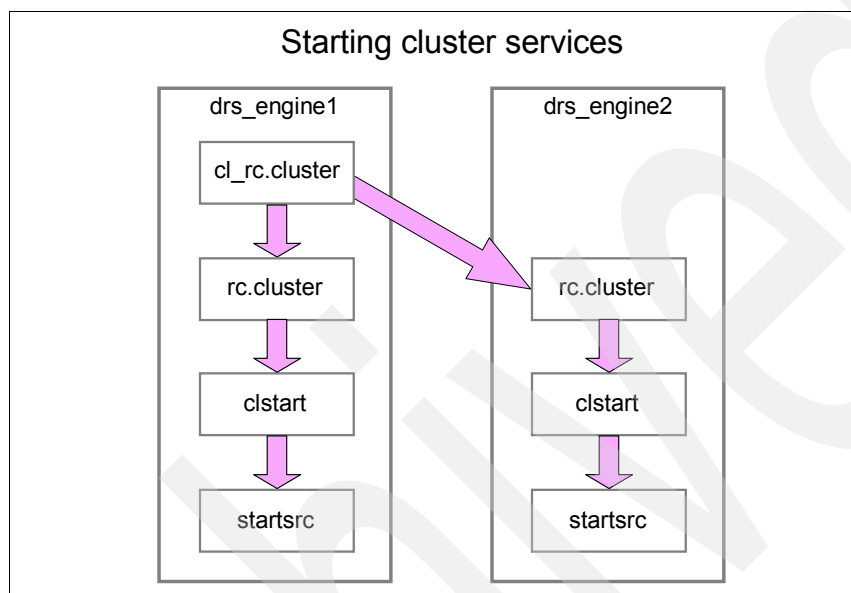


Figure 4-1 Flow of commands used at cluster startup by C-SPOC utility

Using the C-SPOC utility, you can start cluster services on one node or on both nodes in the cluster by executing the C-SPOC `/usr/es/sbin/cluster/sbin/cl_rc.cluster` command on a single cluster node. The C-SPOC `cl_rc.cluster` command calls the `rc.cluster` command to start the cluster services on the nodes specified from the one node. The `rc.cluster` script initializes the environment required for HACMP by setting environment variables, and then calls the `/usr/es/sbin/cluster/utilities/clstart` script to start the HACMP daemons. The `clstart` script is the HACMP script that starts all the cluster services. The `clstart` script calls the SRC `startsrc` command to start the specified subsystem or group.

The nodes are started in sequential order, not in parallel. The output of the command run on the remote node is returned to the originating node. Because the command is executed remotely, there can be a delay before the command output is returned.

The HACMP daemons are started in the following order:

1. RSCT daemons (Topology Services, Group Services, and then Event Management)
2. Cluster Manager and Cluster Communication daemon
3. Cluster Information Program daemon (optional)

## Understanding how to stop cluster services

Figure 4-2 provides an illustration of how commands and scripts are executed when cluster services are stopped on a single node in the cluster using the C-SPOC utility.

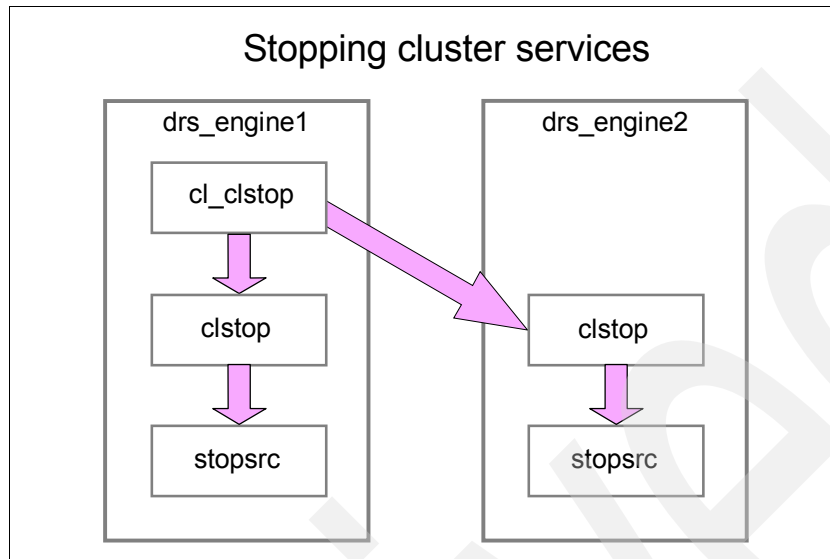


Figure 4-2 Flow of commands used at cluster shutdown by C-SPOC utility

Using the C-SPOC utility, you can stop cluster services on a single node or on both nodes in the cluster. The C-SPOC **c1\_stop** command performs some cluster-wide verification and then calls the **clstop** script to stop cluster services on the specified nodes. The **clstop** script stops an HACMP daemon or daemons. The **clstop** script stops all the cluster services or individual cluster services by calling the SRC command **stopsrc**.

The nodes are stopped in sequential order, not in parallel. The output of the command run on the remote node is returned to the originating node. Because the command is executed remotely, there can be a delay before the command output is returned.

You typically stop cluster services:

- ▶ Before making any hardware or software changes, or other scheduled node shutdowns or reboots. Failing to do so might cause unintended cluster events to be triggered on other nodes.
- ▶ Before certain reconfiguration activities. Some changes to the cluster information stored in the AIX Object Data Manager (ODM) require stopping and restarting the cluster services on all nodes for the changes to become active. For example, if you want to change the name of the cluster, the name of a node, or the name of a network interface, you must stop and restart the cluster.

When the AIX operating system is shut down (with the **shutdown** command) while HACMP is up, the processing in `/usr/es/sbin/cluster/etc/rc.shutdown` performs a forced shutdown. If you prefer to have the resources taken over prior to issuing the AIX **shutdown** command, then you should stop HACMP with a graceful with takeover option. If you do not want to have the resources taken over, take no special action.

### 4.1.2 Starting cluster services

The following steps describe the procedure for starting the cluster services on a single cluster node or both nodes in the cluster by executing the C-SPOC `/usr/es/sbin/cluster/sbin/c1_rc.cluster` command on one of the cluster nodes.

**Tip:** Because the DR550 is running a two-node HACMP cluster, the normal, typical starting procedure for HACMP in the DR550 always includes starting *both* of the cluster nodes. Starting only one of the two cluster nodes is abnormal and atypical. The only time you need to start one cluster node is for maintenance tasks.

To start the cluster services:

1. Log in to any of the DR550 SSAM Servers, such as `drs_engine1`. Use the `dr550` user and switch to root. You can switch to root with the AIX command `su - root`.
2. On the AIX command line, start SMIT by issuing the command `smitty cl_admin`.
3. Select Manage HACMP Services → Start Cluster Services, and press Enter.
4. Use the default settings (see Figure 4-3) when you want to start only one of the two nodes (recommended for maintenance tasks only). Press Enter.

Start Cluster Services

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

|                                                          | [Entry Fields] |   |
|----------------------------------------------------------|----------------|---|
| * Start now, on system restart or both                   | now            | + |
| Start Cluster Services on these nodes                    | [drs_engine1]  | + |
| * Manage Resource Groups                                 | Automatically  | + |
| BROADCAST message at startup?                            | true           | + |
| Startup Cluster Information Daemon?                      | true           | + |
| Ignore verification errors?                              | false          | + |
| Automatically correct errors found during cluster start? | Interactively  | + |

F1=Help  
F5=Reset  
F9=Shell

F2=Refresh  
F6=Command  
F10=Exit

F3=Cancel  
F7=Edit  
Enter=Do

F4=List  
F8=Image

Figure 4-3 Start Cluster Services on one node

Press F4 in the Start Cluster Services on these nodes field when you want to start the cluster services on both nodes (recommended). A window (Start Cluster Nodes on these nodes) opens. Select both nodes with F7. Both nodes are marked in front of their line (see Figure 4-4). Press Enter.

Start Cluster Services on these nodes  
Move cursor to desired item and press F7.  
ONE OR MORE items can be selected.  
Press Enter AFTER making all selections.

```

> drs_engine1
> drs_engine2

```

F1=Help  
F7=Select  
Enter=Do

F2=Refresh  
F8=Image  
/=Find

F3=Cancel  
F10=Exit  
n=Find Next

Figure 4-4 Start Cluster Services on both nodes

Back in the Start Cluster Services screen, check that both nodes are shown in the appropriate line and press Enter.

5. Starting the cluster services might take a few minutes. Do not interrupt this process, and wait for the OK message at the end of the process. See Figure 4-5.

```
COMMAND STATUS

Command: OK stdout: yes stderr: no

Before command completion, additional instructions may appear below.

Verifying Cluster Configuration Prior to Starting Cluster Services.

There are no active cluster nodes to verify against.
Verifying node(s): drs_engine1 drs_engine2 requested to start.....

Successfully verified node(s): drs_engine1 drs_engine2

Starting Cluster Services on node: drs_engine1
This may take a few minutes. Please wait...
drs_engine1: start_cluster: Starting HACMP
drs_engine1: replay files found:
drs_engine1: /usr/es/sbin/cluster/etc/vg/TSMApps.replay
drs_engine1: 0513-029 The portmap Subsystem is already active.
drs_engine1: Multiple instances are not supported.
drs_engine1: 0513-029 The inetd Subsystem is already active.

F1=Help F2=Refresh F3=Cancel F6=Command
F8=Image F9=Shell F10=Exit /=Find
n=Find Next
```

Figure 4-5 Starting the HACMP Cluster services

6. When command execution completes, and HACMP Cluster Services are started on all the specified nodes, SMIT displays a command status screen. It should display an OK message in the upper-left corner. If not, you have to analyze the given messages in the same screen, fix the problem, and start the process again. See Figure 4-6 on page 143.
- If cluster services fail to start on any cluster node, check the C-SPOC utility log file named /tmp/cspoc.log for error messages. This file contains the command execution status of the C-SPOC command executed on each cluster node. /tmp/hacmp.out file should also be inspected in the event of the stop process failure.

```

Command: OK stdout: yes stderr: no

Before command completion, additional instructions may appear below.

Starting Cluster Services on node: drs_engine1
This may take a few minutes. Please wait...
drs_engine1: start_cluster: Starting HACMP
drs_engine1: replay files found:
drs_engine1: /usr/es/sbin/cluster/etc/vg/TSMApps.replay
drs_engine1: 0513-029 The portmap Subsystem is already active.
drs_engine1: Multiple instances are not supported.
drs_engine1: 0513-029 The inetd Subsystem is already active.
drs_engine1: Multiple instances are not supported.
drs_engine1: 438290 - 0:00 syslogd
drs_engine1: Setting routerevalidate to 1
drs_engine1: INFORMATION: must wait at least 2 minutes before cluster restart
drs_engine1: Sleeping 1 minutes.
drs_engine1: 0513-059 The topsvcs Subsystem has been started. Subsystem PID is 192678.
drs_engine1: 0513-059 The grpsvcs Subsystem has been started. Subsystem PID is 618604.
drs_engine1: 0513-059 The emsvcs Subsystem has been started. Subsystem PID is 90194.
drs_engine1: 0513-059 The emaixos Subsystem has been started. Subsystem PID is 594128.
drs_engine1: 0513-059 The gscclvmd Subsystem has been started. Subsystem PID is 278758.
drs_engine1: 0513-059 The clinfoES Subsystem has been started. Subsystem PID is 540882.
drs_engine1: Feb 28 2008 10:56:29 Starting execution of
/usr/es/sbin/cluster/etc/rc.cluster

F1=Help F2=Refresh F3=Cancel F6=Command
F8=Image F9=Shell F10=Exit /=Find
n=Find Next

```

*Figure 4-6 Result output from Start the HACMP Cluster Services*

7. Select System Management (C-SPOC) → Manage HACMP Services → Show Cluster Services, and press Enter.

- Check the SMIT result screen. It should display an OK message in the upper-left corner, three running cluster subsystems, and their AIX process IDs (PID). See Figure 4-7.

| COMMAND STATUS                                                       |             |            |             |
|----------------------------------------------------------------------|-------------|------------|-------------|
| Command: OK                                                          | stdout: yes | stderr: no |             |
| Before command completion, additional instructions may appear below. |             |            |             |
| Status of the RSCT subsystems used by HACMP:                         |             |            |             |
| Subsystem                                                            | Group       | PID        | Status      |
| topsvcs                                                              | topsvcs     | 192678     | active      |
| grpsvcs                                                              | grpsvcs     | 618604     | active      |
| grpglsm                                                              | grpsvcs     |            | inoperative |
| emsvcs                                                               | emsvcs      | 90194      | active      |
| emaixos                                                              | emsvcs      | 594128     | active      |
| ctrmc                                                                | rsct        | 172128     | active      |
| Status of the HACMP subsystems:                                      |             |            |             |
| Subsystem                                                            | Group       | PID        | Status      |
| clcomdES                                                             | clcomdES    | 442394     | active      |
| clstrmgrES                                                           | cluster     | 585914     | active      |
| Status of the optional HACMP subsystems:                             |             |            |             |
| Subsystem                                                            | Group       | PID        | Status      |
| clinfoES                                                             | cluster     | 540882     | active      |
| F1=Help                                                              | F2=Refresh  | F3=Cancel  | F6=Command  |
| F8=Image                                                             | F9=Shell    | F10=Exit   | /=Find      |
| n=Find Next                                                          |             |            |             |

Figure 4-7 Result window: Show Cluster Services

- Quit the SMIT session. Press F10 or ESC+0.

**Tip:** The HACMP cluster will automatically start the System Storage Archive Manager server. Depending on your DR550 configuration, it might take several minutes to start the System Storage Archive Manager server. Therefore, if a System Storage Archive Manager login fails directly after HACMP start, try again later.

## Reintegrating nodes

A node reintegrating into the cluster is just like a node starting up, except that the Cluster Manager is already running on the other cluster node. The cluster's fallback configuration determines whether the node joining the cluster takes back shared resources the other node took over during failover. The DR550 configuration is defined as cascading with fallback. Therefore, the node joining the cluster (the original owner) will take back the shared resources.

### 4.1.3 Stopping cluster services

The following steps describe the procedure for stopping cluster services on a single node or on both nodes in the cluster by executing the C-SPOC `/usr/es/sbin/cluster/sbin/cl_c1stop` command on one of the cluster nodes. When stopping multiple nodes, C-SPOC stops them sequentially, not in parallel. If any node specified to be stopped is inactive, the shutdown operation aborts.

**Important:** When stopping cluster services, minimize activity on the system. If the node you are stopping is currently providing highly available services, notify users of your intentions if their applications will be unavailable. Let them know when services will be restored.

**Note:** With HACMP V5.4, the Cluster Manager (clstrmgrES) will remain active while HACMP is stopped.

To stop cluster services:

1. Log in to any of the DR550 SSAM Servers; use the dr550 user and switch to root. You can switch to root with the AIX command `su - root`.
2. On the AIX command line, start SMIT by issuing the command `smitty clstop`.
3. Use the default settings (see Figure 4-8) and change nothing when you want to stop only one of the two nodes (the one you are logged in to) with a graceful shutdown mode.

**Tip:** Because the DR550 is operating as a two-node HACMP cluster, stopping only one of the two cluster nodes is not a usual task. Stopping only one cluster node is needed for maintenance tasks only. Therefore, the normal stopping procedure for HACMP with the DR550 always includes stopping *both* cluster nodes.

| Stop Cluster Services                                                                   |            |                               |          |
|-----------------------------------------------------------------------------------------|------------|-------------------------------|----------|
| Type or select values in entry fields.<br>Press Enter AFTER making all desired changes. |            |                               |          |
|                                                                                         |            | [Entry Fields]                |          |
| * Stop now, on system restart or both                                                   |            | now                           | +        |
| Stop Cluster Services on these nodes                                                    |            | [drs_engine1]                 | +        |
| BROADCAST cluster shutdown?                                                             |            | true                          | +        |
| * Select an Action on Resource Groups                                                   |            | Bring Resource Groups Offline | +        |
| F1=Help                                                                                 | F2=Refresh | F3=Cancel                     | F4=List  |
| F5=Reset                                                                                | F6=Command | F7=Edit                       | F8=Image |
| F9=Shell                                                                                | F10=Exit   | Enter=Do                      |          |

Figure 4-8 Stop Cluster Services on a single-node

Press F4 in the Stop Cluster Services on these nodes field when you want to stop the cluster services on both nodes. A screen (Stop Cluster Nodes on these nodes) opens. Select both nodes with F7. Both nodes are marked in front of the line. Press Enter. See Figure 4-9.

```

 Stop Cluster Services on these nodes

Move cursor to desired item and press F7.
ONE OR MORE items can be selected.
Press Enter AFTER making all selections.

> drs_engine1
> drs_engine2

F1=Help F2=Refresh F3=Cancel
F7=Select F8=Image F10=Exit
Enter=Do /=Find n=Find Next

```

Figure 4-9 Stop Cluster Services on both nodes

If you want to change the Action on Resource Groups, press F4 in the Select an Action on Resource Groups line. A screen (Shutdown Mode) opens. Select the shutdown mode, and press Enter.

```

 Stop Cluster Services

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

* Stop now, on system restart or both [Entry Fields]
Stop Cluster Services on these nodes now +
BROADCAST cluster shutdown? [drs_engine1,drs_engine2] +
* Select an Action on Resource Groups true +
 Bring Resource Groups Offline +

 Select an Action on Resource Groups

Move cursor to desired item and press Enter.

Bring Resource Groups Offline
Move Resource Groups
Unmanage Resource Groups

F1=Help F2=Refresh F3=Cancel
F1=Help F8=Image F10=Exit
F5=Reset /=Find n=Find Next
F9=Shell

```

Figure 4-10 HACMP Action on resource group

The selected shutdown mode refers to the following types of shutdowns:

- Bring Resource Groups Offline (Graceful): With this option, the local node shuts itself down gracefully. Remote node(s) interpret this as a graceful shutdown of the cluster and do *not* take over the resources. HACMP will stop monitoring the applications on the selected node(s) and will also release cluster resources (such as applications) from HACMP control. This means that the application will be offline on that node.



- **Move Resource Groups (Takeover):** With this option, the local node shuts itself down gracefully. Remote node(s) interpret this as a non-graceful shutdown and take over the resources. The mode is useful for system maintenance. If you select this option, HACMP will stop monitoring the applications on the selected node(s) and will attempt to recover the groups and applications to other active nodes in the cluster.
- **Unmanaged Resource Groups (Forced):** The resources on the local node remain active on the node. HACMP marks the state of such resource groups as UNMANAGED. You can use this option to bring down a node while you perform maintenance or make a change to the cluster configuration such as adding a network interface. The node's applications remain available, except for those that access enhanced concurrent mode volume groups (but without the services of HACMP for AIX daemons).

**Important:** It is very important that you only use Unmanaged Resource Groups (Forced) mode on one cluster node at a time and for as short a time period as possible. When you need to stop cluster services for an extended period and on more than one cluster node, you should use one of the other options (either graceful or graceful with takeover).

4. Back on the Stop Cluster Services screen, check that the node or nodes are shown in the appropriate line, make sure you specified the correct shutdown mode, and press Enter.
5. The system stops the cluster services (excepted clstmgrES) on the nodes specified. SMIT displays a command status screen. It should display an OK message in the upper-left corner. If not, you have to analyze the given messages in the same screen, fix the problem, and start the process again.

If the stop operation fails, check the C-SPOC utility log file named /tmp/cspoc.log for error messages. This file contains the command execution status of the C-SPOC command executed on each cluster node. The /tmp/hacmp.out file should also be inspected in the event of the stop process failure.

6. Quit the SMIT session. Press F10 or ESC+0.
7. Verify on the AIX command line that all the HACMP services are stopped successfully. Use the command `lssrc -g cluster`. With HACMP V5.4, the Cluster Manager (clstrmgrES) will remain active while HACMP is stopped; the `lssrc -g cluster` command should only return the clstmgrES service as running.

### AIX shutdown and cluster services

When the AIX operating system is shut down (with the **shutdown** command) while HACMP is up, the processing in the /usr/sbin/cluster/etc/rc.shutdown script performs a forced shutdown. If you prefer to have resources taken over prior to issuing the AIX **shutdown** command, stop HACMP with a graceful with takeover option. If you do not want to have resources taken over, take no special action. Both methods, shutting down AIX and stopping HACMP, stop the Tivoli Storage Manager server in the DR550.

## 4.2 Starting and stopping IBM System Storage Archive Manager

In this section, we describe how to start and stop the System Storage Archive Manager server. This is applicable to both the DR550 DR1 and DR2.

### 4.2.1 Starting the System Storage Archive Manager server

Depending on your configuration, you can start the System Storage Archive Manager server with the following procedures:

- DR550 DR2 single-node configuration and DR550 DR1:
  - a. Connect to the DR550 SSAM Server through the management console.
  - b. Log in as the dr550 user and switch to root using the AIX command **su - root** to switch the user root.
  - c. On the AIX command line, start the server with the command **/usr/tivoli/tsm/server/bin/rc.admserv &**.
  - d. Verify that the server is running using the AIX command **ps -ef | grep dsm**.

The output should look like this:

```
root 520320 593920 0 10:08:58 pts/0 0:00 grep dsm
root 614418 1 0 12:53:01 - 0:30
/usr/tivoli/tsm/server/bin/dsmserv quiet
```

- DR550 Dual-node configuration:

When HACMP is running on DR550 cluster nodes, the SSAM application is automatically started and is being monitored continuously by the HACMP application monitor process. Under normal circumstances, there should be no need to start the SSAM process manually.

You might still need to start the SSAM process manually in a situation where some maintenance is required (accepting the date in SSAM, for example), which cannot be done with SSAM managed by HACMP and running in background mode. Another example of a situation that might require having the SSAM server started in the foreground is when the SSAM server is unresponsive even after the process is successfully started by HACMP (you cannot log in to the server using the ISC or dsmadm administrative client). In this case, HACMP has to be shut down first before you can start the SSAM process manually. See 4.1.3, “Stopping cluster services” on page 144 for more information about how to stop HACMP services. After HACMP is stopped on both cluster nodes, you can use the following procedure to start the SSAM server:

- a. Connect to the DR550 SSAM Server where the cluster service address is available. Use the management console to connect.
- b. Log in as the dr550 user and switch to root using the AIX command **su - root** to switch the user.
- c. Make sure that HACMP is not running by running **smitty hacmp** and selecting System Management (C-SPOC) → Manage HACMP Services → Show Cluster Services, and press Enter.
- d. Varyon the following volume groups:

```
varyonvg TSMApps
varyonvg TSMDBLogs
varyonvg TSMDBBkup
varyonvg TSMStg
varyonvg TSMDBLogsMirr
```

- e. After all volume groups are varied on, mount the file systems:

```
mount /tsm
mount /tsmDb
mount /tsmLog
mount /tsmdbbkup
mount /tsmDbM
mount /tsmLogM
```

- f. Make a copy of the startserver script:

```
cp -p /usr/bin/startserver /tmp
```

- g. Edit the startserver script:

```
vi /tmp/startserver
```

- h. Change the last line of the file from `$DSMSERV_DIR/dsmserv quiet &` to `$DSMSERV_DIR/dsmserv`.

- i. Save and exit the startserver script.

- j. Start SSAM server in the foreground mode by running `/tmp/startserver` and wait for the TSM> prompt to appear.

**Note:** It might take a few minutes for the SSAM startup process to complete (that is, for the TSM> prompt to appear). During the startup process, AIX CPU utilization is on average 50% in “id” (idle) and 50% in “wa” (waiting on disk I/O). The `vmstat 1` command can be used to monitor CPU utilization. An indication that the SSAM startup process is complete is when the “wa” (waiting on disk I/O) column is reduced to 0%.

**Important:** When the SSAM application is started manually on an HACMP cluster with HACMP cluster services stopped, it will remain unavailable for all external clients. FSG nodes and all external SSAM clients are configured to use IP addresses associated with the `drs_cluster_svc` and `drs_cluster_ext_svc` labels, respectively. These IP addresses are only available when HACMP is started and fully operational.

After required SSAM repairs or maintenance tasks have been done, you should stop the SSAM server by entering the `halt` command at the TSM> prompt. Please note that if inadvertently you enter this command at the AIX prompt that the operating system will be halted immediately.

## 4.2.2 Stopping the System Storage Archive Manager server

Depending on your configuration, you can stop the System Storage Archive Manager server with the following procedures:

- DR550 DR2 single-node configuration and DR550 DR1:

To avoid unnecessary interruption of administrative and client node sessions, stop the server only after the current sessions have completed or have been canceled. To stop the server, do the following:

- a. Connect to the System Storage Archive Manager server through an administrative command-line client or the Administration Center interface (see 11.4.1, “Using the IBM SSAM Administration Center” on page 469).

- b. Verify that you can stop the server without impacting running processes and sessions. To shut down the server without severely impacting administrative and client node sessions, perform the following steps:
  - i. Use the **disable session** command to prevent new client node sessions from being started.
  - ii. Use the **query sessions** command to identify currently running administrative and client node sessions.
  - iii. SSAM clients should be notified about the fact that SSAM server will shut down (you must do this outside of System Storage Archive Manager).
  - iv. Use the **cancel sessions** command to cancel any remaining administrative or client node sessions.
  - v. Use the **query process** command to display active background process information.
  - vi. Use the **cancel process process\_number** command to cancel any existing background process.
  - vii. In the case of tape attachment, use the **query mount** command to display information about the status of one or more sequential access volumes that are mounted.
  - viii. Use the **dismount volume volume\_name** command to dismount an idle volume by its volume name. If a drive cannot dismount the volume, manual intervention is required.
- c. Issue the **halt** command at the TSM> prompt to shut down the server.
 

The **halt** command forces an abrupt shutdown, which cancels all administrative and client node sessions even if they have not completed processing. Any transactions in progress that are interrupted by the **halt** command are rolled back when you restart the server.

  - If you issue the **halt** command from a command line within the Web administrative client, System Storage Archive Manager will shut down immediately without any further inquiry. The Web session will freeze, because the connection is shut down.
  - If you issue the **halt** command from a command line within the administrative command-line client, System Storage Archive Manager will ask “Do you wish to proceed?”. Type yes or Y to shut down the server.

**Important:** Use the System Storage Archive Manager **halt** command only after the administrative and client node sessions have completed or been canceled. Do not use the command on an AIX command line.

► DR550 DR2 dual-node configuration:

To avoid unnecessary interruption of administrative and client node sessions, stop the server only after the current sessions have completed or have been canceled. To stop the server, do the following:

- a. Connect to the System Storage Archive Manager server through an administrative command-line client or the Administration Center interface (see 11.4.1, “Using the IBM SSAM Administration Center” on page 469).
- b. Verify that you can stop the server without impacting the running processes and sessions. To shut down the server without severely impacting administrative and client node sessions, perform the following steps:
  - i. Use the **disable session** command to prevent new client node sessions from being started.

- ii. Use the **query sessions** command to identify currently running administrative and client node sessions.
- iii. SSAM clients should be notified about the fact that the SSAM server will shut down (you must do this outside of System Storage Archive Manager).
- iv. Use the **cancel sessions** command to cancel any remaining administrative or client node sessions.
- v. Use the **query process** command to display active background process information.
- vi. Use the **cancel process *process\_number*** command to cancel any existing background process.
- vii. If you have tape attachment, use the **query mount** command to display information about the status of one or more sequential access volumes that are mounted.
- viii. Use the **dismount volume *volume\_name*** command to dismount an idle volume by its volume name. If a drive cannot dismount the volume, manual intervention is required.

**Important:** When HACMP is running on DR550 cluster nodes, the SSAM application is being continuously monitored by the HACMP application monitor. If an attempt is made to shut down the SSAM server using the **halt** command while it is still running under HACMP control, the application monitor will restart the SSAM server almost immediately.

- c. To stop SSAM, HACMP cluster services have to be stopped. For further information about how to stop HACMP services, see 4.1.3, “Stopping cluster services” on page 144.

## 4.3 Starting/Stopping the FSG

In this section, we describe how to start and stop the FSG.

### 4.3.1 Clustered FSG configuration

Under normal operations, for a high availability File System Gateway system, the Main DR550 FSG is in “active” mode while the Supplementary DR550 FSG is in “standby” mode. If the active gateway fails, the cluster service automatically transitions the Supplementary DR550 FSG to the “active” state. The Supplementary DR550 FSG immediately takes over from the failed Main DR550 FSG.

In situations where you need to perform maintenance on the “active” Main DR550 FSG, or if you need to restore the Main DR550 FSG to the “active” state after required repairs or maintenance tasks have been done, the following procedures can be used to initiate a manual failover:

1. Log on to the “active” server as root.
2. Determine the status of the clusters with the following command:

```
crm_mon -i 15
```

The **-i** interval option is the time (in seconds) between status updates. **crm\_mon** defaults to 15 seconds. This is a continuously running command that updates the output every 15 seconds. The output shows the current state of the cluster. Refer to 4.5, “Check FSG status” on page 154 for a detailed output from this command.

**Note:** If you are just starting the cluster, you must wait long enough (usually more than one minute) for the final status of the cluster to be displayed.

3. Press Ctrl-C to exit and return to the command line.
4. Stop the “active” DR550 File System Gateway. At the command line, enter:  

```
/etc/init.d/drg stop
```

The “standby” DR550 File System Gateway transitions to “active” status. Failover is complete. Client operations that are in progress when the failover occurs are interrupted. Once the failover is complete, Windows CIFS clients resume normal operation without remapping their connections. NFS clients, however, must be remounted.
5. Restart the DR550 File System Gateway stopped in step 4. At the command line, enter:  

```
/etc/init.d/drg start
```
6. Monitor the DR550 File System Gateway with `crm_mon`. At the command line, enter:  

```
crm_mon -i 15
```

The output will be similar to the one shown in Figure 4-11 on page 155.

## 4.4 DR550 check procedures summary

This section outlines the checkout procedure to be performed after a successful installation and configuration.

### 4.4.1 Check software levels

The actual software levels recommended for DR550 are listed in the section “DR550 Single-Node Components and DR550 Dual-Node Components” in *IBM System Storage DR550 Version 4.5 Problem Determination and Service Guide*, GA32-0576. It is imperative to install the recommended levels only. It is also important to respect the required code levels to keep the whole software stack consistent and compatible.

**Note:** The DR550 uses common, off-the-shelf software for most components; updates for a specific component are in accordance with the standard procedure for that component or product. For example, to update the System Storage Archive Manager, refer to the official Tivoli Storage Manager procedure. Read and understand the installation instructions for the software to be updated. In addition, consider the recommended DR550 code levels and corequisite levels.

Use the commands listed to check the software levels in AIX. Using these commands requires a logon to AIX:

- ▶ AIX level: `oslevel -r` and `oslevel -s`
- ▶ System Storage Archive Manager (SSAM) level: `lspp -l | grep tivoli.tsm.server`
- ▶ HACMP level: `lspp -l | grep cluster`
- ▶ To check the level of the DS4000 components, start the DS4000 Storage Manager Client (see “DS4000 Storage Manager for DR550” on page 20), select **Manage Devices**, and then select **Subsystem** → **View Profile**.

## 4.4.2 AIX basic check

The commands listed in Table 4-2 can be used to perform a basic checkout of the DR550 at the AIX level. Most commands can be executed with dr550 authority.

Table 4-2 AIX commands

| Component checked                                 | Command                                                                                     |
|---------------------------------------------------|---------------------------------------------------------------------------------------------|
| Check if the Tivoli Storage Manager server is up. | <code>ps -ef   grep dsm</code>                                                              |
| Check if HACMP is started.                        | <code>lssrc -g cluster</code><br><code>/usr/es/sbin/cluster/clstat -a -o<sup>a</sup></code> |
| Check the status of the volume groups.            | <code>lsvg -o, lsvg -p &lt;vg-name&gt;, lsvg -l &lt;vg-name&gt;</code>                      |
| Check the status of file systems.                 | <code>df, mount, lsfs</code>                                                                |
| Check the status of the disk systems.             | <code>lsdev -Cc disk</code>                                                                 |
| Check status of the array controllers.            | <code>fget_config -v1 dar0<sup>a</sup></code>                                               |

a. These require root authority.

## 4.4.3 Tivoli Storage Manager basic check

Use the commands listed in Table 4-3.

Table 4-3 Tivoli Storage Manager commands

| Component checked                                 | Command                                    |
|---------------------------------------------------|--------------------------------------------|
| Check if the Tivoli Storage Manager server is up. | <code>ps -ef   grep dsm</code>             |
| Log on to Tivoli Storage Manager Server.          | <code>dsmadm</code>                        |
| Check DB volumes.                                 | <code>q dbvol</code>                       |
| Check DB status.                                  | <code>q db</code>                          |
| Check LOG volumes.                                | <code>q dblog</code>                       |
| Check LOG status.                                 | <code>q log</code>                         |
| Check Storage pools.                              | <code>q stg</code>                         |
| Check storage pool volumes.                       | <code>q vol</code>                         |
| Check activity log.                               | <code>q actlog</code>                      |
| Check tape drives.                                | <code>q drives</code>                      |
| Check tape library.                               | <code>q libr</code>                        |
| Check paths.                                      | <code>q path</code>                        |
| Check the number of files stored.                 | <code>q occ &lt;node-name&gt;</code>       |
| Check file space utilization.                     | <code>q filespace &lt;node-name&gt;</code> |

#### 4.4.4 HACMP basic check

This section only applies to dual-node DR550 systems. Table 4-4 details some commands used to check the HACMP status.

Table 4-4 HACMP check

| Component checked    | Command                                        |
|----------------------|------------------------------------------------|
| AIX error report     | <code>errpt -a</code>                          |
| List file systems    | <code>df, mount, and lsfs</code>               |
| HACMP process status | <code>lssrc -g cluster</code>                  |
| HACMP cluster status | <code>/usr/es/sbin/cluster/clstat -a -o</code> |

To verify the HACMP configuration, run `smitty hacmp` and select Extended Configuration → Synchronization and Verification. Note, upon changes of the HACMP resource group, a synchronization has to be performed as well.

#### HACMP failover (manual)

For the dual-node system to only check the HACMP operability, perform a HACMP failover. This requires root authority. Do the following steps:

1. Enter the following command from the AIX command line:  
`smit hacmp`
2. Select System Management → HACMP Resource Group and Application Management → Move a Resource Group to another Node. Select Resource Group (dr550\_rg) and then select the Target Node to activate. Press Enter twice to start the failover.
3. The failover might take a few minutes. Check the HACMP status periodically with the command `smit hacmp` and select Problem Determination → View current State. Repeat until the cluster state is stable.

#### 4.4.5 Tivoli Storage Manager API checkout

In order to verify the functionality, we recommend you use the Tivoli Storage Manager API test program. The Tivoli Storage Manager API client includes an API test program. This test program is located in the directory `<tsm-client-path>\api\samprun` for Windows and is called *dapismp*. This test program allows you to archive, query, and retrieve objects in the DR550.

Refer to 5.3.2, “SSAM/Tivoli Storage Manager API (using the sample application dapismp)” on page 214 for more details.

### 4.5 Check FSG status

For the high availability configuration, the most commonly used command to find the status of the nodes in the cluster is the following command:

```
crm_mon
```

The output from this command is as shown in Figure 4-11 on page 155 (if the `-i` option is not used, a default refresh interval of 15 seconds is used).



```

Refresh in 10s...
=====
Last updated: Fri Feb 29 09:53:30 2008
Current DC: main (4da2fba3-ea94-5f03-9aaf-d1f14b65b0f0)
2 Nodes configured.
1 Resources configured.
=====

Node: main (4da2fba3-ea94-5f03-9aaf-d1f14b65b0f0): online
Node: supplementary (430e222e-c5bf-84b8-22e0-5fb725f1753b): online

Resource Group: res_group
 drgbase (Bycast::ocf:FSGBase): Started main
 ip_resource (heartbeat::ocf:IPAddr2): Started main
 pingtest (Bycast::ocf:PingTest): Started main
 drgactive (Bycast::ocf:FSGActive): Started main
 Filesharing-NFS (Bycast::ocf:NFS_Server): Started main
 Filesharing-CIFS (Bycast::ocf:Samba): Started main

```

Figure 4-11 The `crm_mon` command output

**Note:** If you are just starting the cluster, you must wait long enough (usually more than one minute) for the final status of the cluster, as shown in Figure 4-11, to be displayed.

Archived

# IBM System Storage Archive Manager

Starting with Version 5.2.2, IBM Tivoli Storage Manager has been enhanced to help meet additional requirements for data retention and data protection (regulatory compliance). The enhanced product initially referred to as the IBM Tivoli Storage Manager for Data Retention is now called the IBM System Storage Archive Manager (SSAM).

In this chapter, we describe the architecture and features of the base product (IBM Tivoli Storage Manager).

This chapter is intended primarily for readers who are new to Tivoli Storage Manager. However, we also indicate which features of Tivoli Storage Manager, or SSAM in this case, are relevant to the IBM System Storage DR550.

Furthermore, we describe the installation and configuration of IBM Tivoli Storage Manager Operational Reporting. This is additional tool can be used to generate IBM SSAM application reports from a remote service Windows workstation.

## 5.1 IBM System Storage Archive Manager overview

IBM System Storage Archive Manager (SSAM) provides a comprehensive solution focused on the key data protection activities of backup, archive, recovery, space management, and disaster recovery.

When its data retention and protection features are enabled, the IBM Tivoli Storage Manager is called the IBM System Storage Archive Manager. The key software component of the IBM System Storage DR550 Models DR1 and DR2 is the IBM System Storage Archive Manager V5.5.

The IBM System Storage DR550 is an archive solution focused on providing a data retention repository for applications that use the Archive Manager Application Program Interface (API) or the DR550 File System Gateway to archive data to the DR550 System Storage Archive Manager server.

It is beyond the scope of this book to explain SSAM in detail. This book focuses on the System Storage Archive Manager fundamentals necessary to understand the IBM System Storage DR550 solution and explores what customizing has already been done to the System Storage Archive Manager server provided in the IBM System Storage DR550.

**Tip:** For a detailed overview of SSAM/Tivoli Storage Manager and its complementary products, refer to the IBM Tivoli software Information Center at the following location:

<http://publib.boulder.ibm.com/infocenter/tivihelp>

### 5.1.1 IBM System Storage Archive Manager architecture overview

SSAM is implemented as a client/server software application. As depicted in Figure 5-1, the core product is the System Storage Archive Manager server.

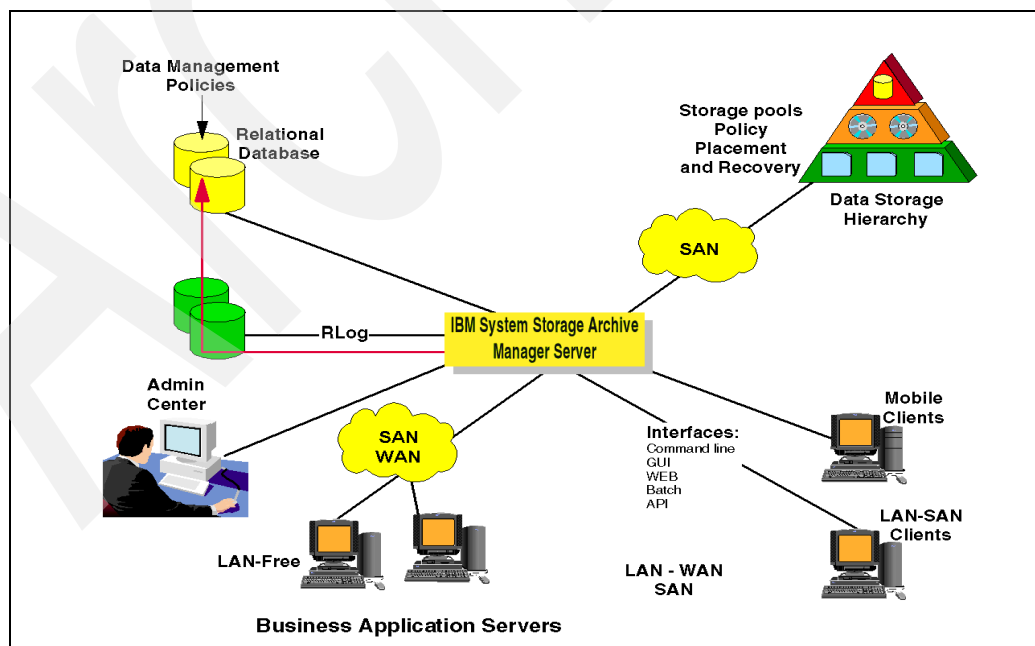


Figure 5-1 IBM System Storage Archive Manager architectural overview

Clients using the Web client as well as applications or the File Systems Gateway using the API can transfer data to the server through a LAN, WAN, or SAN (see Figure 5-1 on page 158).

**Note:** Client applications can only communicate over the LAN or WAN when archiving to a DR550 System Storage Archive Manager server.

## **System Storage Archive Manager server**

As part of the IBM System Storage Open Software Family, IBM SSAM protects data from hardware failures, errors, and unforeseen disasters by storing backup and archive copies on offline and off-site storage. Scaling to protect hundreds to thousands of computers running more than a dozen operating systems, ranging from mobile computers to mainframes and connected together through the Internet, WANs, LANs, or SANs, Storage Manager Extended Edition's centralized Web-based management, intelligent data move and store techniques, and comprehensive policy-based automation all work together to minimize administration costs and the impact to both computers and networks.

The SSAM server consists of a runtime environment and a relational database. The proprietary database with its recovery log stores all information about the running environment and the managed data. In the case of the IBM System Storage DR550, the IBM System Storage Archive Manager Server has been preconfigured and enabled for data retention protection and runs within AIX.

## **Server storage**

A SSAM server can write data to more than 400 types of devices, including hard disk drives, disk arrays and subsystems, stand-alone tape drives, tape libraries, and other forms of random and sequential-access storage. The media that the server uses are grouped into storage pools. The storage devices can be connected directly to the server through SCSI, through Fibre Channel directly attached, or over a storage area network (SAN). For the DR550-DR2, the storage pools are implemented through SAN attachment to the DR550 Storage Server (DS4200) and its EXP420 expansion units. For the DR550-DR1, storage pools are defined on the internal SCSI drives of the DR550 SSAM Server. However, this can also differ if the option to add disks to the DR550-DR1 is available; in that case, the storage pools will be defined on the DR550 Storage Server.

## **Client nodes**

A client node, in the context of the IBM System Storage DR550, is an application that communicates and transfer data objects for archiving to the System Storage Archive Manager server. A client node is registered in a policy domain on the server.

There are three types of nodes that can now be used with the DR550 SSAM Server:

- ▶ IBM Tivoli Storage Manager API
- ▶ IBM Tivoli Storage Manager Backup/Archive Client (Web-based)
- ▶ File System Gateway

## **Application program interface (API)**

Tivoli Storage Manager provides a data management application program interface (API) that can be used to implement application clients to integrate popular business applications, such as databases or groupware applications. The API also adheres to an open standard and is published to enable customers and vendors to implement specialized or custom clients for particular data management needs or nonstandard computing environments. The API enables an application client to use the Tivoli Storage Manager storage management functions. The API includes function calls that you can use in an application to perform the

following operations: start or end a session, assign management classes to objects before they are stored on a server, archive objects to a server, and signal retention events for retention such as activate, hold, or release.

Alternatively, some vendor applications exploit the Tivoli Storage Manager data management API by integrating it into their software product to implement new data management functions or to provide archival functionality on additional system platforms. Some examples are IBM DB2 Content Manager, IBM DB2 Content Manager On Demand, IBM CommonStore for SAP® R/3, Lotus® Domino®, and Microsoft Exchange data archival.

The API is published to enable customers or vendors to implement their own solutions following their special needs, including full documentation available on the Internet. For more information, see *IBM Tivoli Storage Manager: Using the Application Programming Interface*, SC32-0147, available at:

[http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/topic/com.ibm.itsmfdt.doc/b\\_api.pdf](http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/topic/com.ibm.itsmfdt.doc/b_api.pdf)

Refer to “Testing the new features with dapismp” on page 214 for examples about how to exercise the API features for archival and data retention.

## IBM Tivoli Storage Manager Backup Archive Client

The Backup Archive (BA) Client provides an easy and effective way to archive and retrieve data from a workstation. The process is easy and menu driven. The BA client can be accessed either directly as an application installed on the client node or can be accessed remotely through a Web browser.

**Important:** With a System Storage Archive Manager server (and consequently with a DR550), archiving and setting retention attributes on archived data is only feasible through the Web-based BA client. See “BA client remote access (Web Client)” on page 206.

We discuss the BA Client in more detail, including examples of how to archive data using the different methods available. Refer to 5.3.1, “SSAM/Tivoli Storage Manager BA Client V5.5” on page 204.

## DR550 File System Gateway

The DR550 File System Gateway (FSG) hardware and software enable a broad range of content-management applications to access the DR550 SSAM Server by adding NFS and CIFS network file access. The gateway takes files that are sent using network file protocols and associates these files with an IBM System Storage Archive Manager management policy. This is done by using customer-configurable path and file naming pattern matching rules. The DR550 FSG sends these files, with their associated policies, to the DR550 SSAM Server using the SSAM application program interface.

We discuss the DR550 File System Gateway (FSG) in more detail in Chapter 6, “IBM DR550 File System Gateway” on page 229.

## Administrative interfaces

To configure and manage your DR550 SSAM server, you need an administrative interface. Two different interfaces are available: the command-line administrative client and the Web browser based Administration Center.

### **Command-line administrative client (dsmadm)**

The command-line administrative client is preinstalled and preconfigured at your DR550 SSAM Server. You can start it with any user authority. We recommend that you log on to AIX with the dr550adm user ID, because this user has no remote login restriction:

1. Log on to AIX at IBM System Storage DR550 #1 and issue the command **dsmadm**.
2. Enter your DR550 SSAM Server user ID and password. The default administrative user is admin with the password admin.
3. You will get a shell with the TSM> prompt where you can enter Tivoli Storage Manager/SSAM commands. (See Example 5-1.)
4. The **help** command gives you a help for all possible commands and their syntax.
5. To exit the shell, enter the command **quit**.

#### *Example 5-1 Tivoli Storage Manager/SSAM shell with dsmadm command*

---

```
$ dsmadm
IBM Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 5, Level 0.0
(c) Copyright by IBM Corporation and other(s) 1990, 2007. All Rights Reserved.

Enter your user id: admin

Enter your password:

Session established with server TSM: AIX-RS/6000
Server Version 5, Release 5, Level 0.0
Server date/time: 02/27/08 11:23:01 Last access: 02/27/08 10:40:41

tsm: TSM>
```

---

### **Administration Center**

For the central administration of one or more System Storage Archive Manager instances, as well as the whole data management environment, SSAM provides a Java-based graphical administration interface called the *Administration Center* (it is installed as an *Integrated Solution Console* or ISC component). The Administration Center and the ISC are preinstalled and started automatically at the IBM System Storage DR550. If you want to install the Administration Center and the ISC on a separate server in your network, you can find the software on your SSAM / Tivoli Storage Manager server CDs.

The administrative interface enables administrators to control and monitor server activities, define management policies for clients, and set up schedules to provide services to clients at regular intervals (see Figure 5-2).

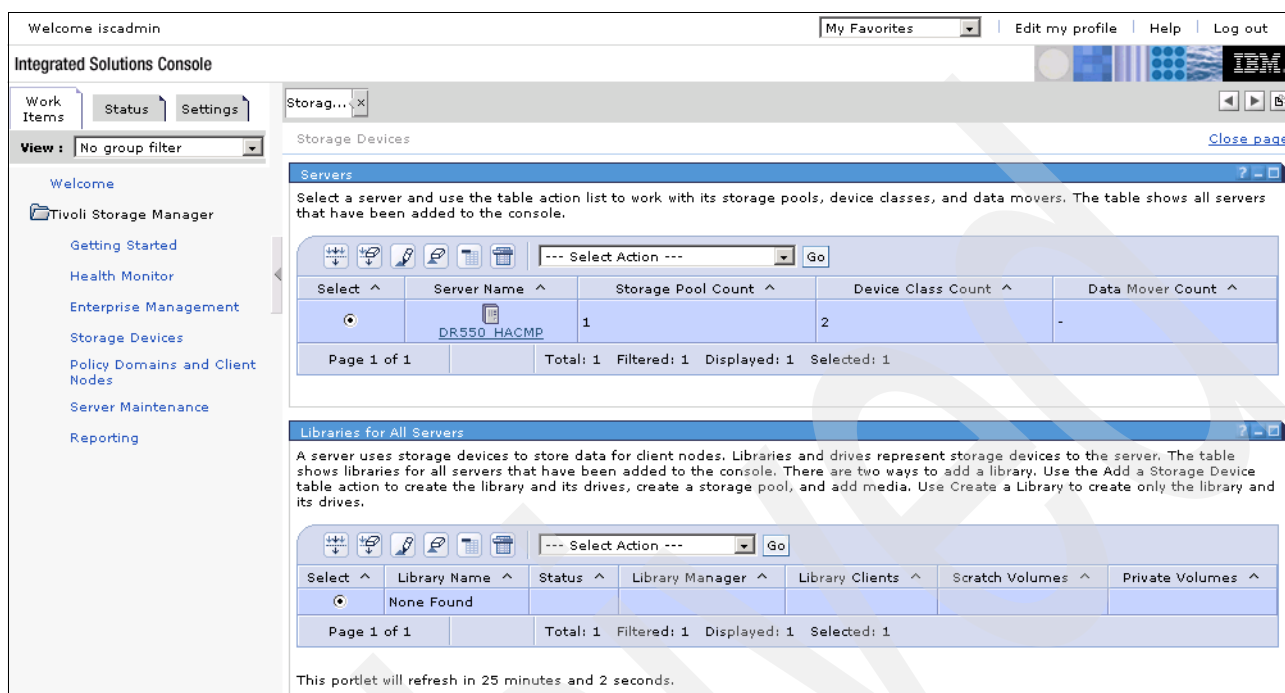


Figure 5-2 Administration Center Web interface: Storage devices main window

### Work with the Administration Center / ISC

This section gives you a short introduction on how to start and configure the Administration Center:

1. To connect to the Administration Center Web interface, start a Web browser and start an http session to the IP address of the node or workstation where the Administration Center and the ISC are installed, using the port number specified when installing the ISC:  
`http://ip_of_management_station:8421/ibm/console`
2. Log in at ISC with user ID `iscadmin`. The password for the preinstalled ISC at the IBM System Storage DR550 is `iscadmin`. If you have installed the ISC and Administration Center at a separate server you have configured this password during the installation process.
3. To configure your DR550 SSAM Server at the Administration Center and select **Tivoli Storage Manager** → **Server Maintenance** → **select Action** → **add Server Connection**. Enter a DR550 SSAM Server administration user ID, usually `admin`, the password of the DR550 SSAM Server user, and the IP address of the DR550 SSAM Server, and click **OK**.

Alternatively, if you install the ISS extensions for IBM Director (see 9.3.2, “IBM Director ISS Extensions for DR550” on page 361), you can use the IBM Director Console as a launch point for the SSAM Administration Center.

Now you can select various functions to administrate your DR550 SSAM Server.

**Tip:** If you do not want to work with the ISC at the IBM System Storage DR550, you should disable the startup of the ISC during boot of the AIX to free up the used resources. The next section will show you the appropriate procedure.



### **Disable ISC startup at the IBM System Storage DR550**

To disable ISC startup at the IBM System Storage DR550:

1. Log in to AIX with the dr550 user ID at the management console and switch to root authority by using `su -`.
2. Enter `rmitab isc6`.
3. Reboot the system to activate the change with the `shutdown -Fr` command.
4. Repeat steps 1-3 for the DR550 SSAM Server # 2 if you have a dual-node IBM System Storage DR550.

*Example 5-2 /etc/inittab (excerpt)*

---

```
#Begin AC Solution Install block
#Start the Cloudscape database server
si:23456789:wait:/usr/ibm/common/acsi/bin/acsisrv.sh -start
#End AC Solution Install block
conserver:2:once:/opt/conserver/bin/conserver -d -i -m 64
autosrvr:2:once:/usr/tivoli/tsm/server/bin/rc.admserv >/dev/console 2>&1 #Start
the Tivoli Storage Manager server
smmonitor:2:wait:/usr/sbin/SMmonitor start > /dev/console 2>&1 # start SMmonitor
daemon
isc6:23:once:/opt/IBM/ISC601/PortalServer/bin/startISC.sh ISC_Portal
sbchfscheck:2:wait:/usr/sbin/sbchfscheck
```

---

### **Automation**

DR550 SSAM Server includes a central scheduler that runs on the DR550 SSAM Server and provides services for use by the server and clients. You can schedule administrative commands to tune server operations and to start functions that require significant server or system resources during times of low usage. You can also schedule a client action, but that would be unusual for a data retention-enabled client. Each scheduled command (administrative or client) is called an *event*. The server tracks and records each scheduled event in the database.

## **5.1.2 IBM System Storage Archive Manager basic concepts**

DR550 SSAM Server manages client *data objects* based on information provided in administrator-defined *policies*.

Data objects can be subfile components, files, directories, or raw logical volumes that are archived from client systems; they can be objects, such as tables, logs, or records from database applications, or simply a block of data that an application system archives to the server. The DR550 SSAM Server stores these objects on disk volumes and tape media that it groups into *storage pools*.

We explain these concepts with regard to the IBM System Storage DR550 and the DR550 System Storage Archive Manager server.

### **SSAM storage pools and storage hierarchy**

SSAM manages data as objects stored in SSAM storage pools (see Figure 5-3 on page 164). Each object has an associated management policy to which it is “bound.” The policy defines how long to keep that object and where the object enters the storage hierarchy.

The physical location of an object within the storage pool hierarchy has no effect on its retention policies. An object can be migrated or moved to another storage pool within a SSAM storage hierarchy. This can be useful when freeing up storage space on higher performance devices, such as disk, or when migrating to new technology. Objects can and should also be copied to copy storage pools for disaster recovery protection. To store these data objects on storage devices and to implement storage management functions, SSAM uses logical definitions to classify the available physical storage resources. Most important is the logical entity called a *storage pool*, which describes a storage resource for a single type of media such as disk volumes, which are files on a file system, or tape volumes, which are cartridges in a library.

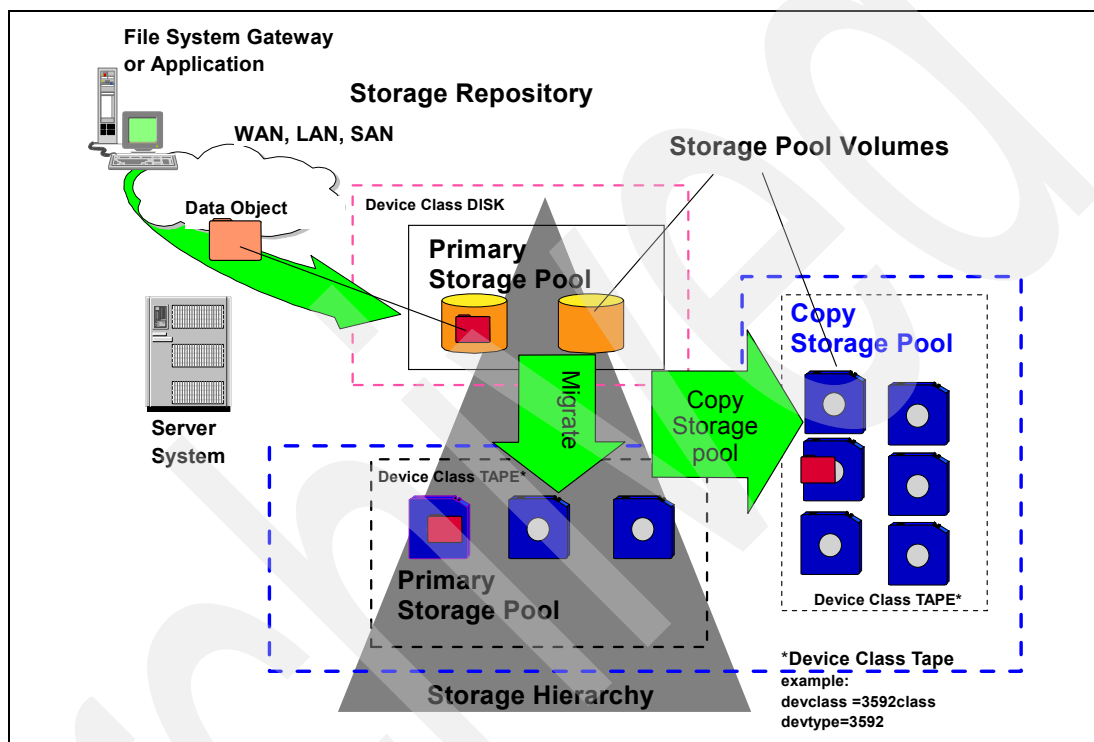


Figure 5-3 IBM System Storage Archive Manager storage hierarchy

## Device classes

A storage pool is built up from one or more physical storage pool volumes. For example, the archivepool on the DR550 SSAM Server consists of several raw logical volumes (see Figure 5-16 on page 193).

A logical entity called a *device class* is used to describe how SSAM can access those physical volumes to place the data objects on them. Each storage pool is bound to a single device class.

The storage devices used with SSAM vary mainly in their technology and total cost.

To reflect this, you can imagine the storage as a pyramid (or triangle), with high-performance storage in the top (typically disk), normal performance storage in the middle (typically optical disk or cheaper disk), and low-performance, but high-capacity, storage at the bottom (typically tape). Figure 5-3 illustrates this.

Disk storage devices are random access media, making them better candidates for storing frequently accessed data.

Tape, however, is a high-capacity sequential access media, which can easily be transported off-site for disaster recovery purposes. Access time is much slower for tape due to the amount of time needed to load a tape in a tape drive and locate the data. However, for many applications, that access time is still acceptable.

Disk storage is referred to as online storage, while tape storage has often been referred to as *off-line* and *near-line* with regard to Hierarchical Storage Management (HSM) in the past. With SSAM/Tivoli Storage Manager HSM, tape volumes, located in a tape library, are accessed by the application that is retrieving data from them (near-line) transparently. Tapes no longer in the library are off-line, requiring manual intervention. The introduction of lower cost mass storage devices, such as Serial Advanced Technology Attachment (SATA) disk systems, offers an alternative to tape for near-line storage. Figure 5-4 illustrates the use of a SATA disk as near-line storage. The IBM System Storage DR550 uses SATA disks.




Figure 5-4 Online, near-line, and off-line storage

### Device types

Each device defined to SSAM is associated with one device class. Each device class specifies a *device type*.

A device type identifies a device as a member of a group of devices that share similar media characteristics. For example, the 3592 device type applies to IBM System Storage Enterprise Tape Drive 3592 or IBM System Storage TS1120.

The device type also specifies management information, such as how the server gains access to the physical volumes, recording format, estimated capacity, and labeling prefixes.

Device types include DISK, FILE, and a variety of removable media types.

Note that a device class for a tape or optical drive must also specify a *library*.

### Device access strategy

The access strategy of a device is either random or sequential. Primary storage pools can use random devices (such as disk) or sequential devices (such as tape). Copy storage pools use sequential access devices. Certain SSAM/Tivoli Storage Manager processes use only sequential access strategy device types:

- ▶ Copy storage pools
- ▶ SSAM database backups
- ▶ Export
- ▶ Import (disabled on DR550 System Storage Archive Manager server)

### Predefined device classes on the DR550 SSAM Server

To illustrate the application of the concepts we just defined, we show in this section how they are used for the DR550 System Storage Archive Manager server.

### Device class DISK

Magnetic disk devices are the only random access devices. All disk devices share the same predefined device class: DISK. The DISK device class supports creating disk storage pool volumes on file systems using raw logical volumes. The predefined primary storage pool “archivepool” on the DR550 SSAM Server uses the device class DISK. Each storage pool volume is a raw logical volume (see Figure 5-5).

### Device class DBBKUP

Device type FILE creates files on a file system as sequential access volumes. SSAM manages these files the same way it manages tape volumes. To the server, these files have the characteristics of a tape volume. A FILE device class named DBBKUP is predefined on the DR550 SSAM Server for the purpose of SSAM full database backups. Copy storage pools can also use the device type FILE because it is a sequential access device.

| Device Class Name | Device Access Strategy | Storage Pool Count | Device Type | Format |
|-------------------|------------------------|--------------------|-------------|--------|
| DISK              | Random                 | 1                  |             |        |
| DBBKUP            | Sequential             | 0                  | FILE        | DRIVE  |

Figure 5-5 Predefined device classes on the DR550

## Tape devices

System Storage Archive Manager supports a wide variety of enterprise class tape drives and libraries. The following link connects you to the product support Web site where you will find a link to the currently supported devices list:

[http://www.ibm.com/software/sysmgmt/products/support/IBM\\_TSM\\_Supported\\_Devices\\_for\\_AIXHPSUNWIN.html](http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_AIXHPSUNWIN.html)

We recommend that you use tape devices for the purpose of backing up your primary storage pools to copy storage pools and backing up the database. Tape devices are well-suited for this, because the media can be transported off-site for disaster recovery purposes. A tape drive or tape library is not included in the IBM System Storage DR550 Models DR1 and DR2; however, any system is tape-ready and you can attach tape devices that are supported by SSAM/Tivoli Storage Manager on the AIX platform and that best suit your data retention requirements. We recommend that you use the IBM System Storage TS1120 Tape Drive in combination with rewritable and WORM media or the IBM Ultrium 3 LTO drives in combination with 3589 rewritable and WORM media. We discuss attaching tape in Chapter 11, “Tape attachment” on page 447.

## SSAM policy concepts

A data storage management environment consists of three basic types of resources: *client systems* (for example, applications using the API/BA clients to archive data), *policy*, and *data*.

The client systems run the applications that create or collect data to be managed.

The policies are the rules to specify how to manage the archived objects, for example, how long to retain an archive object in storage, whether chronological or event-based archive retention is used, in which storage pool to place an object, or, in the case of backup, how many versions to keep, where they should be stored, and what SSAM does to the archive object once the data is no longer on the client file system.

Client systems, or *nodes*, in SSAM/Tivoli Storage Manager terminology, are grouped together with other nodes with common storage management requirements into a *policy domain*. The policy domain links the nodes to a *policy set*, which is a collection of storage management rules for different storage management activities.

**Note:** The term *client node* refers to the application sending data to the DR550 SSAM Server.

A policy set consists of one or more *management classes*. A management class contains the rule descriptions called *copy groups* and links these to the data objects to be managed.

A copy group is the place where all the storage management parameters are defined, such as the number of stored copies, retention period, and storage media. When the data is linked to particular rules, it is said to be bound to the management class that contains those rules. There are two types of copy groups: backup and archive. Another way to look at the components that make up a policy is to consider them in the hierarchical fashion in which they are defined, that is, consider the policy domain containing the policy set, the policy set containing the management classes, and the management classes containing the copy groups and the storage management parameters (see Figure 5-6).

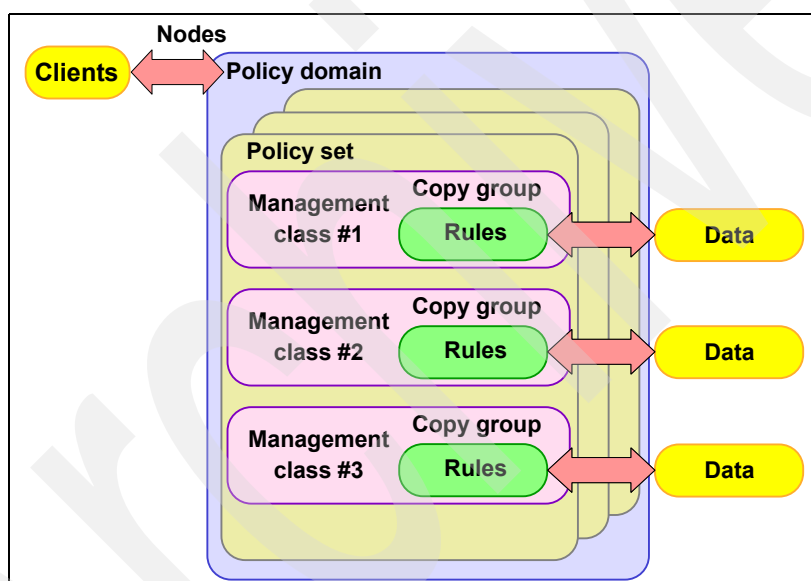


Figure 5-6 Policy relationships and resources

We explain the relationships in the following sections, leaving out the backup copy group that does not apply to the DR550 SSAM Server.

### **Archive copy group**

This group controls the archive processing of files associated with the management class. An archive copy group determines the following:

- ▶ How to handle files that are in use during archive
- ▶ Where the server stores archived copies of files
- ▶ How long the server keeps archived copies of files

### **Management class**

The management class associates client files with archive copy groups. A management class can contain one backup or archive copy group, both a backup and an archive copy group, or no copy groups. Users can bind (that is, associate) their files to a management class through the include-exclude list.

**Note:** With the IBM System Storage DR550, management classes can only contain archive copy groups, since backups are not possible on a System Storage Archive Manager server.

### **Policy set**

The policy set specifies the management classes that are available to groups of users. Policy sets contain one or more management classes. You must identify one management class as the default management class. Only one policy set, the ACTIVE policy set, controls policies in a policy domain.

### **Policy domain**

This enables an administrator to group client nodes by the policies that govern their files and by the administrators who manage their policies. A policy domain contains one or more policy sets, but only one policy set (named ACTIVE) can be active at a time. The server uses only the ACTIVE policy set to manage files for client nodes assigned to a policy domain.

You can use policy domains to:

- ▶ Group client nodes with similar file management requirements
- ▶ Provide different default policies for different groups of clients
- ▶ Direct files from different groups of clients to different storage hierarchies based on need (different file destinations with different storage characteristics)
- ▶ Restrict the number of management classes to which clients have access

Figure 5-7 on page 169 summarizes the relationships among the physical device environment, SSAM storage and policy objects, and clients. The numbers in the following list correspond to the numbers in the figure.

**Note:** Figure 5-7 on page 169 shows an outline of the policy structure. We show all the steps necessary to create a valid policy in this chapter.

These are the steps:

1. When clients are registered, they are associated with a policy domain. Within the policy domain are the policy set, management class, and copy groups.
2. When a client (application) archives an object, the object is bound to a management class. A management class and the archive copy group within it specify where files are stored first (destination), and how they are managed when they are archived.
3. Storage pools are the destinations for all stored data. An archive copy group specifies a destination storage pool for archived files. Storage pools are mapped to device classes, which represent devices. The storage pool contains volumes of the type indicated by the associated device class. For example, the storage pool *archivepool* on the DR550, which uses the device class DISK, contains only disk volumes.

Data stored in disk storage pools can be migrated to tape or optical disk storage pools and can be backed up to copy storage pools.




Figure 5-7 Example of the policy structure for archive

## 5.2 IBM System Storage Archive Manager

In this section, we discuss the System Storage Archive Manager (sometimes referred to as SSAM) features. We concentrate our explanation on additional archive retention features, such as event-based archive retention, deletion hold, and data retention protection.

IBM System Storage Archive Manager makes the deletion of data before its scheduled expiration extremely difficult. Short of physical destruction of the storage media or server, or deliberate corruption of data or deletion of the Archive Manager database, SSAM will not allow data to be deleted before its scheduled expiration date. Content management and archive applications can apply business policy management for the ultimate expiration of archived data at the appropriate time.

SSAM was introduced as a separately licensed product in Version 5.2.2 and was designed to help meet data retention and disposition compliance regulations and policies. SSAM uses the IBM Tivoli Storage Manager Extended Edition code.

Archive Manager offers rich functionality and features giving you a powerful and comprehensive archive retention solution:

- ▶ Archive Manager runs on vendor neutral storage technology, giving you the ability to utilize hundreds of different types of disk, tape, optical, and DVD media on which to retain your data.
- ▶ Hierarchical storage capabilities allow you to create policies so data is stored on the type of media that best meets data longevity, access speed, and cost needs.
- ▶ Migration automates moving data from one type of media to another as media needs change, and as new types of media become available in the market. This not only helps ensure data longevity, but also allows data to be stored on the type of media that best meets its speed of access and cost needs.

- ▶ Archive Manager's expiration policies expire the data when it is no longer needed, thus freeing up the non-WORM storage media and saving you money.
- ▶ Off-site protection of the data is standard in Archive Manager. Off-site copies can be created onto any of the hundreds of types of media supported, and like the primary copy, is policy-managed to allow for expiration.

## Activation of System Storage Archive Manager (SSAM)

After installation of a new server instance, activate the data protection features by issuing the administrative command **set archiveretentionprotection on**. Once activated, you *cannot* store data as backup objects or initiate backup sessions using a normal backup-archive client. SSAM becomes an archive-only retention compliance solution.

The System Storage Archive Manager in the IBM System Storage DR550 is already activated. In addition, the IBM System Storage DR550 ships from manufacturing with a small amount of test data in the standard archive copy group. The purpose is to provide protection from disabling the archive retention protection on the DR550 SSAM Server (this can only be done on a DR550 SSAM Server that does not contain *any* data objects).

The following steps can be used to verify that the DR550 System Storage Archive Manager server is configured correctly.

From the administrative command line client (**dsmadm**) or from the Administration Center interface, enter the following commands to verify the status of data retention protection and licenses:

### 1. Type **query status**.

The second to last output line shows the SSAM/Tivoli Storage Manager server status:

Archive Retention Protection: On

### 2. Type **query license**.

You should have the following set:

Is IBM System Storage Archive Manager in use ?: Yes

Is IBM System Storage Archive Manager licensed ?: Yes

**Tip:** Define a test policy domain for test data. We highly recommend that you define a test policy domain and policy set for any pre-production testing. Remember, all of the test data that you archive to a DR550 SSAM server cannot be prematurely deleted. For chronological archive retention, use realistic retention periods for test data, such as **RETVER=*n* days**. For event-based archive retention, remember that for each and every object you archive, there must also be a retention event signaled to activate the retention period (**RETVER=*n* days**). Without it, the data remains in SSAM storage forever. The **RETMIN** parameter should also be set to something realistic for test purposes (**RETMIN=*n* days**). We discuss the different archive retention methods in 5.2.1, "Archive copy group retention parameters" on page 170.

## 5.2.1 Archive copy group retention parameters

In order to use the archive function of SSAM, you must define valid policies that include defining a policy domain, policy set, management class or classes, and an archive copy group, as well as setting archive retention parameters in the archive copy group and associating your application clients with the SSAM policies.



The archive copy group parameters that govern retention are RETVER, RETINIT, and RETMIN. The RETINIT and RETMIN parameters were introduced in ISSAM/Tivoli Storage Manager V 5.2.2 to make it possible for applications using the API or the BA Client 5.3.2 to further control the retention period (RETVER) for archive objects. Chronological archive retention has always been possible with SSAM and was controlled solely by the RETVER parameter. With SSAM/Tivoli Storage Manager V5.2.2, event-based archive retention and two new archive copy group parameters have been introduced.

## Two methods of archive retention

There are two methods of archive retention, which are defined by the parameters of the archive copy group:

- ▶ Chronological archive retention
- ▶ Event-based archive retention

We now look at the parameters of the archive copy group and their possible values for the two archive retention methods.

### ***The existing archive retention parameter***

The existing archive retention parameter is RETVER (retain version). Possible values are RETVER=0 to 30,000 days or NOLIMIT.

**Important:** Selecting the NOLIMIT value on DR550 System Storage Archive Manager server, as is the case on a IBM System Storage DR550, means that you will never be able to delete the data.

The retain version parameter (RETVER) within the archive copy group specifies the number of days to retain each archive object. Possible values are 0 to 30,000 days or NOLIMIT, which means that an archive copy is maintained indefinitely.

### ***New archive retention parameters***

The two new archive retention parameters are:

- ▶ RETINIT (retention initiation)

The possible values are RETINIT=creation or event.

The retention initiation (RETINIT) parameter specifies when the time specified by the retain version (RETVER=*n* days) attribute is initiated. The possible values for this parameter are creation or event. The default value is creation. In the following list, we explain both values:

- RETINIT=creation (chronological archive retention)

By setting this parameter to creation (RETINIT=creation) in the archive copy group, you specify that the retention time specified by the RETVER attribute (RETVER=*n* days) is initiated right at the time an archive copy is stored on the server. This is referred to as *chronological archive retention*.

- RETINIT=event (event-based archive retention)

By setting this parameter to event (RETINIT=event) in the archive copy group, you specify that the retention time (RETVER=*n* days) for the archived data is initiated by an application that used API function calls or the Web Client. If the application never initiates the retention, the data is retained indefinitely. This method of archive retention is referred to as *event-based archive retention*.

Possible events to signal through the API or the BA Web Client to the DR550 System Storage Archive Manager (SSAM) server are:

- Activate: Activates the countdown of the RETVER value for the given event-based object.
- Hold: Prevents the DR550 SSAM server from deleting the object, even if the RETVER period has ended. Signaling a “hold” does not extend the retention period, but a hold object will only expire after a release event is sent.
- Release: Removes the hold status of an object. The DR550 SSAM server will then treat the object again according to the RETVER and RETMIN values.

► RETMIN (retain minimum)

Possible values are RETMIN=0 to 30,000 days.

The retain minimum (RETMIN) parameter applies only to event-based archive retention policy and specifies the minimum number of days to retain an archive object regardless of the value of RETVER. The default value is 365. Possible values are 0 to 30,000 days.

We provide the following examples to give you insight into archive copy groups and defining policy.

***Defining and updating an archive copy group***

Before you can define a copy group, you must first define a policy domain, a policy set, a management class, and finally, a storage pool as the destination for the archived data.

The following administrative command lines are examples of defining an archive copy group and are based on having a policy domain named *testdom*, a policy set named *testset*, two management classes, one named *testclass\_chrono*, which is the assigned default management class, and the other named *testclass\_even*.

The management class *testclass\_chrono* is set up as a chronological class with a retention period of one year. The management class *testclass\_event* is set up as an event-based class with a minimum retention period of seven years and an activate-event with a retention period of two years. In addition, we have a storage pool named *archivepool*. The copy groups themselves can only be named *standard* in SSAM. To define and update an archive copy group:

► Define an archive copy group for chronological archive retention:

```
define copygroup testdom testset testclass_chrono standard type=archive
retver=365 retinit=creation destination=archivepool
```

The policy set must now be validated and activated before these policy changes go into effect:

```
validate policyset testdom testset
activate policyset testdom testset
```

► Define an archive copy group for event-based archive retention:

```
define copygroup testdom testset testclass_event standard type=archive
retver=730 retinit=event retmin=2555 destination=archivepool
```

The policy set must now be validated and activated before these policy changes go into effect:

```
validate policyset testdom testset
activate policyset testdom testset
```

## 5.2.2 Chronological archive retention

Figure 5-8 shows a simplified view of a chronological retention policy. With RETINIT=creation and RETVER=365 days, a file that is archived on day 0 is retained for 365 days and becomes eligible for expiration. In this case, after 365 days from the time the data was created, all references to that data are deleted from the database, making the data irretrievable from SSAM storage volumes. This kind of archive retention is called *chronological retention*. By default, the RETINIT value is set to creation.

**Note:** Choose chronological archive retention when the application that is doing the archiving is not able to send retention events such as activate, hold, and release. Chronological archive retention is also intended for use when you are archiving to a regular Tivoli Storage Manager server (not enabled for data retention protection) through the normal backup/archive client.

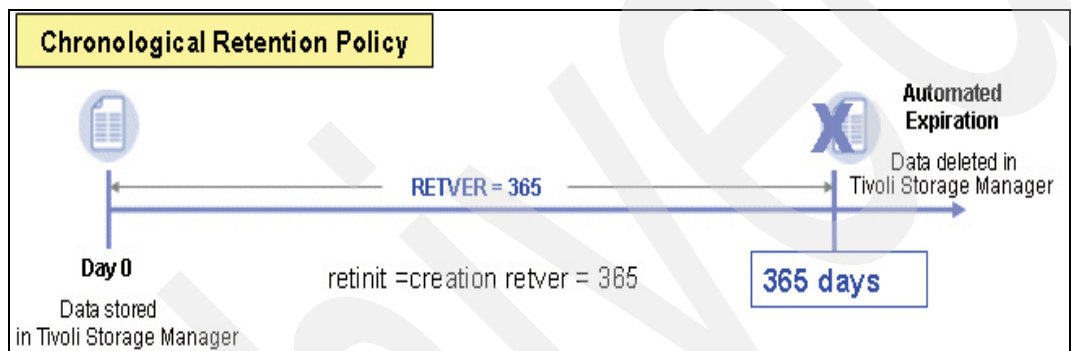


Figure 5-8 Chronological retention policy

Archive copy groups using the chronological retention policy satisfy many archive retention requirements.

## 5.2.3 Event-based retention policy

In certain situations, data retention periods cannot be easily defined, or they depend on events taking place long after the data is archived. Event-based archive retention is designed to meet these requirements. Event-based retention policy is designed for applications that use the API function calls to trigger events also known as *retention events*. You can also use the Web-based BA client (Version 5.3.2 or higher) to archive client objects (data) using event-based policies and trigger retention events against those objects.

Figure 5-9 on page 174 shows a time line depicting an event-based policy. In this example, an application using the API archives data using the retention values shown. The archived data is retained for a minimum of 2,555 days (RETMIN=2555). If the retention time (RETV) is activated through an API retention event, SSAM assigns an expiration date for this object. The expiration date that SSAM assigns is whichever comes later, either:

- ▶ The date the object was archived, plus the number of days specified in the RETMIN parameter.
- ▶ The date the event was signaled, plus the number of days specified in the RETVER parameter.

After reaching this expiration date, the data is eligible for expiration. When the time for expiration occurs, all references to that data are deleted from the SSAM database, making the data irretrievable from SSAM storage volumes. This kind of archive retention is referred to as *event-based retention*. In 5.2.13, “IBM System Storage Archive Manager policy examples” on page 201, we show examples of defining policies that use event-based retention.

**Note:** Use event-based archive retention if the archive application you are using (such as DB2 Content Manager together with Record Manager) uses the API function calls to activate the retention period of the archived data objects.




Figure 5-9 Event-based retention policy

Table 5-1 shows the information gathered from two archive queries that run after archiving a file, one using creation-based archive policy and one using event-based archive policy.

**Note:** When an object is archived using event-based retention, SSAM manages that object as though the RETVER parameter were set to NOLIMIT until an event initiates the retention period (see Table 5-1).

Table 5-1 Status of files archived with creation-based and event-based retention

| Object attributes in SSAM/Tivoli Storage Manager database | RETINIT=CREATION   | RETINIT=EVENT                |
|-----------------------------------------------------------|--------------------|------------------------------|
| Insert date                                               | 2006/2/28 12:16:30 | 2006/2/29 1:23:56            |
| Expiration date                                           | 2016/3/9 12:16:30  | 65535/0/0 0:0:0 (= no limit) |
| Mgmt class                                                | CREATION           | EVENT                        |
| Retention initiated                                       | STARTED            | PENDING                      |
| Object Held                                               | FALSE              | FALSE                        |

Notice that the status of the Retention-Initiated attribute is STARTED for the management class CREATION, and PENDING for the management class EVENT. Also, compare the expiration dates.

## 5.2.4 Deletion hold and release

Some regulations require that the data is retained longer than the minimum retention period in certain cases. This might be due to any litigation, a legally-required or a company-required audit, or a criminal investigation requiring the data as evidence. The API (and Web-based BA Client 5.3.2) supports new function calls used to place a deletion hold on an archive object. These functions are also called *retention events*. A deletion hold can be applied at any point in time during the retention period for an archive object. The object will then be retained until a deletion release is applied. If a deletion release is not applied, the object is retained indefinitely. Although deletion hold and release are events, they can be applied to objects archived not only using the event-based policies, but also the chronological, creation-based policies.

Figure 5-10 shows a time line depicting deletion hold and release. In “Sending retention events using dapismp” on page 221, we demonstrate setting deletion activate, hold, and release on archived objects.

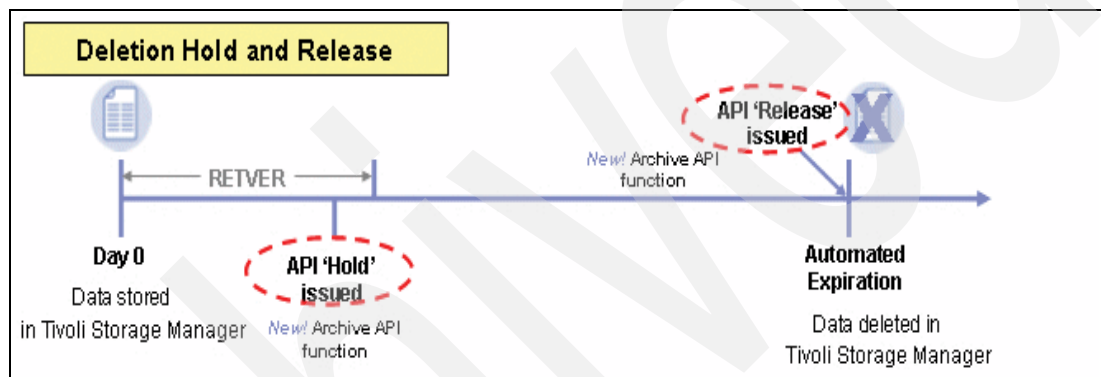


Figure 5-10 Deletion hold and release

## 5.2.5 IBM System Storage Archive Manager prerequisites

System Storage Archive Manager has the following key prerequisites and general requirements:

- ▶ A newly installed Tivoli Storage Manager Version 5.2.2 (or later) server, that has the data retention option (set archiveretentionprotection on) turned on during server setup, prior to storing any data.
- ▶ License package consisting of the Tivoli Storage Manager Extended Edition license.
- ▶ Already defined valid SSAM policies.
- ▶ Tivoli Storage Manager BA Client (Web-based) or API on the client node is enabled for communication with a SSAM server by specifying the following option in the client system options file (dsm.opt in Windows or dsm.sys in UNIX):  
enablearchiveretentionprotection yes
- ▶ Adequate storage device is attached to the DR550 SSAM Server.

### Available features and restrictions

A DR550 SSAM server is much like any other Tivoli Storage Manager server. All features to administer the server and manage data objects and the storage repository are still available. All devices that are supported with Tivoli Storage Manager server are available for a DR550 SSAM Server.

The available features of System Storage Archive Manager are:

- ▶ Access control
- ▶ Data retention protection
- ▶ Creation-based retention
- ▶ Event-based retention
- ▶ Deletion hold and release
- ▶ Data encryption (available in Version 5.3 and later)
- ▶ Data shredding (available in Version 5.4 and later)
- ▶ Tape Drive encryption (available in Version 5.4 and later)

Table 5-2 summarizes the differences between System Storage Archive Manager and Tivoli Storage Manager Extended Edition.

*Table 5-2 Tivoli Storage Manager Extended Edition and SSAM*

| Function                          | IBM Tivoli Storage Manager Extended Edition    | IBM System Storage Archive Manager                                                                                |
|-----------------------------------|------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| Install                           | IBM Tivoli Storage Manager Extended Edition CD | IBM Tivoli Storage Manager Extended Edition CD <i>and</i> set archiveretentionprotection (preconfigured on DR550) |
| Devices supported                 | More than 400                                  | More than 400                                                                                                     |
| Server-to-server backup           | Yes                                            | No                                                                                                                |
| Library sharing                   | Yes                                            | Yes                                                                                                               |
| Client data                       | Backup, archive, HSM                           | Archive                                                                                                           |
| API data                          | Backup, archive                                | Archive                                                                                                           |
| Import data                       | Yes                                            | No                                                                                                                |
| Export data                       | Yes                                            | Yes                                                                                                               |
| Delete data, node, file space     | Yes                                            | No                                                                                                                |
| Lower archive retention criterion | Yes                                            | No                                                                                                                |
| Archive hold/release              | No                                             | Yes                                                                                                               |
| Chronological archive             | Yes                                            | Yes                                                                                                               |
| Event-based archive               | No                                             | Yes                                                                                                               |

### System Storage Archive Manager safety features

To ensure that objects stored under data retention policies remain compliant to those policies, there are some restrictions with the use of SSAM features.

The following restrictions apply:

- ▶ A registered node cannot be reassigned to a different policy domain.
- ▶ You cannot define a device class with device type SERVER.
- ▶ You cannot import data to an SSAM server.
- ▶ You cannot activate a policy set that contains weaker retention parameters than the ones in place in the active policy set.
- ▶ You cannot remove data retention protection on an DR550 SSAM server before the retention requirements for all data have been satisfied and all data has expired.

## 5.2.6 Data retention protection

Data retention protection ensures that archive objects will not be deleted from the DR550 SSAM server until the policy-based retention requirements for that object have been satisfied. After an archive object is stored on a DR550 SSAM server, retention protection cannot be removed.

Retention protection is based on the retention criterion for each object, which is determined by the RETVER and RETMIN parameters of the archive copy group of the management class to which the object is bound. If an object uses event-based retention (RETINIT=EVENT), the object will not expire until whatever comes later: either the date the object was archived plus the number of days in the RETMIN parameter, or the date the event was signaled plus the number of days specified in the RETVER parameter. When using the chronological retention (RETINIT=CREATION), the archive object will expire after the time that is set with the RETVER parameter has elapsed.

**Note:** You always need an event to start the expiration period for event-based retention!

Table 5-3 shows the relationship between the different parameters and their use within certain retention policies.

Table 5-3 Archive copy group parameters

| Archive copy group parameters                                                                                      | Chronological retention                                                                                  | Event-based retention                                                                                                  |
|--------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------|
| <b>RETINIT</b><br>Defines when to initiate the retention period defined in the RETVER attribute.                   | <b>RETINIT=CREATION</b><br>The expiration date is based on the date the object was archived plus RETVER. | <b>RETINIT=EVENT.</b><br>The expiration date is based on the date of the retention initiation event plus RETVER.       |
| <b>RETVER</b><br>Number of days to retain the archive object after retention is initiated.                         | RETVER=0 to 30,000 days or NOLIMIT.                                                                      | RETVER=0 to 30,000 days.                                                                                               |
| <b>RETMIN</b><br>Minimum number of days to retain archive object.                                                  | Not applicable.                                                                                          | RETMIN=days.<br>Based on date object was archived.                                                                     |
| What is the earliest date when the object could become eligible for expiration after retention has been initiated? | (date object was archived) + RETVER.                                                                     | (date retention was initiated through Event) + RETVER or (date object archived) + RETMIN, whichever is <i>longer</i> . |

### Deletion protection

The following operations cannot delete archived data on a DR550 System Storage Archive Manager server:

- ▶ Requests from the application client to delete an archive object prematurely
- ▶ DELETE FILESPACE (from either a client or administrative command)
- ▶ DELETE VOLUME DISCARDATA=YES
- ▶ AUDIT VOLUME FIX=YES

## 5.2.7 Expiration processing

The expiration processing deletes expired client files from storage pools. Expiration processing also removes from the database any restartable restore sessions that exceed the time limit for saving such sessions. You can run expiration processing either automatically or by command. You should ensure that expiration processing runs periodically to allow the server to reuse storage pool space that is occupied by expired client files.

**Note:** An archive file is not eligible for expiration if there is a deletion hold on it. If a file is not held, it will be handled according to the existing expiration processing.

### Running expiration processing automatically

You control automatic expiration processing by using the expiration interval option (EXPINTERVAL) in the server options file (dsmserve.opt). You can set the options by editing the dsmserve.opt file. If you use the server options file to control automatic expiration, the server runs expiration processing each time you start the server. After that, the server runs expiration processing at the interval you specified with the option, measured from the start time of the server.

### Setup the automatic expiration processing

Follow these steps:

1. Edit the server options file with the AIX command:  

```
vi /usr/tivoli/tsm/server/bin/dsmserve.opt
```
2. Go to the line EXPInterval 24. You can find this line by pressing ESC and entering /EXPInterval. Then press N until the cursor is at the line \* EXPInterval 24. Remove the \* at the beginning of the line and save the file. The SSAM server will now run the expiration process every 24 hours.
3. Restart the DR550 SSAM server by entering **stopserver** at the AIX command line. This will stop the DR550 SSAM server. Then start the DR550 SSAM server with the **startserver** command.

#### Example 5-3 The dsmserve.opt file (excerpt)

---

```
* =====
* EXPINTERVAL
* *****
* EXPInterval
* Specifies the number of hours between automatic inventory expiration
* runs.
* Syntax
* +-----+-----+
* | EXPInterval | value |
* +-----+-----+
* Parameters
* value Specifies the number of hours between automatic
* inventory expiration runs. The minimum value
* is 0, where automatic expiration will not
* execute and must be started with the
* EXPIRE INVENTORY command. The maximum value
* is 336 (14 days).
* Examples
* EXPInterval 24
```

---



## Using commands and scheduling to control expiration processing

You can manually start expiration processing by issuing the **expire inventory** command. Expiration processing then deletes expired files from the database.

You can schedule this command by using the **DEFINE SCHEDULE** command. If you schedule the **EXPIRE INVENTORY** command, set the expiration interval to 0 (zero) in the server options so that the server does not run expiration processing when you start the server.

See the Tivoli Information Center for more information about scheduling commands:

<http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp?toc=/com.ibm.itsm.fdt.doc/toc.xml>

### ***Start manual expiration processing***

Follow these steps:

1. At the SSAM admin client, enter the command **expire inventory**.
2. You can check the status of the expiration process with the **query process** command.

## 5.2.8 Encryption

In order to make the archived data more secure, the SSAM/Tivoli Storage Manager client implements an encryption function, which allows you to encrypt data before it is sent to the DR550 System Storage Archive Manager server. This helps secure archived-data during transmission, and it means that the data stored on the DR550 SSAM server is encrypted and thus is unreadable even by the administrator.

The encryption processing is the last task performed on the client system before the data is sent to the server; other client operations such as compression happen before encryption is done.

### ***API encryption***

You can use either a 56-bit DES or 128 AES (Advanced Encryption Standard); the 128 AES was introduced in SSAM/Tivoli Storage Manager V5.3. The default, 56-bit DES, can be overridden by setting the parameter **ENCRYPTIONTYPE AES128** in the **dsm.opt** (Windows) or **dsm.sys** (UNIX).

The encryption function enables you to choose which files are subject to encryption using an include/exclude list. Set the **include.encrypt** parameter in the option file (**dsm.opt** or **dsm.sys**) for the objects to encrypt (the default is NO encryption) and the **exclude.encrypt** for the objects that you do not want to encrypt.

For example, to encrypt all data, set:

```
include.encrypt /.../* (AIX)
```

and

```
include.encrypt *\...* (Windows)
```

To encrypt the object **/FS1/DB2/FULL**, set:

```
include.encrypt /FS1/DB2/FULL
```

For client applications using the API, there are two methods to handle encryption:

- ▶ Application-managed encryption
- ▶ Transparent encryption

The two methods are exclusive. In other words, only one method should be chosen for any given application client node. For both methods, an encryption password is used to generate the real encryption key. The encryption password can be up to 63 characters in length, but the key generated from it is always 8 bytes for 56 DES and 16 bytes for 128 AES.

Application-managed encryption means that the client application (your document management system in the case of a IBM System Storage DR550) is responsible for managing the keys (actually encryption passwords used by SSAM to generate the encryption keys). In addition, the client application code might have to be changed to communicate the password to the API on each archive or retrieve operation.

On the other hand, transparent encryption provides encryption of application data without requiring any changes to the client application and delegates all key management operations (generation, storage, and retrieval) to the DR550 SSAM server.

For those reasons, we recommend using transparent encryption. This method requires a Version 5.3 Tivoli Storage Manager server (or System Storage Archive Manager V5.3 server in the case of the DR550) and can only be used through the API.

### ***Transparent encryption***

This is the simplest and safest method to implement data encryption. One random encryption key is generated per session (every time a client initiates a session with the DR550 SSAM Server for archiving). The key is generated with a random number generator on the client side.

For each archived object, the generated encryption key is sent to and stored in the DR550 SSAM server database. However, before it is sent to the DR550 SSAM server along with the encrypted archived object, the key is encrypted using DES 56 encryption. Once the server receives the structure containing the encrypted encryption key, it decrypts the key, re-encrypts the key using a specific server-based encryption mechanism, and stores it in the database along with the corresponding object\_ID.

**Important:** If the encryption key is not available, data cannot be retrieved under any circumstances. Be sure that you back up your server database frequently to prevent data loss.

During a retrieval, the server uses the server-based mechanism to decrypt the key, re-encrypts, and sends the re-encrypted key to the client along with the encrypted object. In turn, the client (API) extracts the key and decrypts it. Finally, the decrypted key is used to decrypt the data.

To enable transparent encryption, specify - ENABLECLIENTENCRYPTKEY YES in the system option file dsm.opt (Windows) or dsm.sys (AIX).

## **5.2.9 Tape drive encryption**

This is a new function available with SSAM Version 5.4 and the appropriate tape drive devices. In this section, we explain the different tape drive encryption methods and we describe the setup of the application based tape encryption for the IBM System Storage DR550.

## Overview

It is often critical to secure client data, especially when that data might be of a sensitive nature. To ensure that data for offsite volumes is protected, IBM Tape encryption technology is available. This technology utilizes a stronger level of encryption by requiring 256-bit Advanced Encryption Standard (AES) encryption keys. Keys are passed to the drive by a key manager in order to encrypt and decrypt data.

## Tape drive encryption methods

IBM tape drives supports three methods of drive encryption.

### *Application managed encryption*

Encryption keys are managed by the application, in this case, the SSAM. SSAM generates and stores the keys in the server database. Data is encrypted during WRITE operations when the encryption key is passed from the server to the drive. Data is decrypted on READ operations. The application encryption method is only supported for storage pool volumes. To use application encryption, set the DRIVEENCRYPTION parameter to ON in the associated DEVCLASS.

### *Library managed encryption*

Encryption keys are managed by the tape library. Keys are stored in an encryption key manager and provided to the drive transparent to SSAM. If the hardware is set up to use library encryption, SSAM can allow this method to be utilized by setting the DRIVEENCRYPTION parameter to ALLOW in the associated DEVCLASS.

### *System managed encryption*

System encryption is currently available on AIX and MVS™. Encryption keys are managed by the device driver or operating system and stored in an encryption key manager. They are provided to the drive transparent to SSAM. If the hardware is set up to use system encryption, SSAM can allow this method to be utilized by setting the DRIVEENCRYPTION parameter to ALLOW.

**Note:** Hardware based tape encryption will only encrypt data written to the tape. No encryption will be done for data stored on disk.

In this book, we only discuss application managed encryption. Additional information about key management and how to initiate tape encryption is available at:

<http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp?toc=/com.ibm.itstorage.doc/toc.xml>

## Hardware requirements for drive encryption

At the time of the writing of this book, there were two different IBM tape drives supporting drive encryption:

- ▶ IBM TS1120 Enterprise Tape Drives, which are encryption capable (Type: 3592-E05). Application managed tape encryption using the IBM TS1120 Tape Drives is supported in the following libraries:
  - IBM System Storage TS3400 Tape Library
  - IBM System Storage TS3500 Tape Library
  - IBM TotalStorage 3494 Tape Library

**Note:** Only TS1120 with the ENC sticker at the rear side of the drive is encryption capable. Earlier TS1120 might have no encryption support. If you do not know how to determine the encryption capability of your tape drive, call IBM support.

- ▶ IBM TS1040 LTO4 Tape Drive (Type: 3588-F4A). Application managed tape encryption using IBM TS1040 Tape Drives is supported in the following IBM libraries:
  - IBM System Storage TS3100 Tape Library
  - IBM System Storage TS3200 Tape Library
  - IBM System Storage TS3310 Tape Library
  - IBM System Storage TS3500 Tape Library

**Note:** You might need to update the IBM tape device driver and the SSAM server version to have the full LTO4 encryption support in the DR550. Refer to the Tivoli support Web site for more information.

The Tivoli support web site is located at:

[http://www.ibm.com/software/sysmgmt/products/support/IBM\\_TSM\\_Supported\\_Devices\\_for\\_AIXHPSUNWIN.html](http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_AIXHPSUNWIN.html)

### ***Setup of encryption method at the tape drive***

Before you can use the drive encryption in SSAM, you must set up the encryption method at the tape drive or at the tape library that contains the tape drive. The configured encryption method in the drive or library must match your operating system or application settings. The setup depends on the library type that you are using:

- ▶ TS3500 Tape Library

For all tape drives that are installed within a TS3500 or 3584 library, you can set up the encryption method through the library Web interface (select **Manage Library** → **By Logical** → **Select Library** → **choose** → **Modify Encryption Method** → **GO**). Refer to the TS3500 Users Guide for more information:

[http://www-1.ibm.com/support/docview.wss?rs=1159&context=STCMML8&dc=DA400&uid=sg1S7000149&loc=en\\_US&cs=utf-8&lang=en](http://www-1.ibm.com/support/docview.wss?rs=1159&context=STCMML8&dc=DA400&uid=sg1S7000149&loc=en_US&cs=utf-8&lang=en)

- ▶ 3494 Tape Library

For all TS1120 tape drives that are installed in a 3494 library, you have to change the encryption method at the operator panel of the tape drive. Contact your IBM Hardware Support or refer to the TS1120 Maintenance Information (section Drive Encryption) documentation, which comes with your machine, for more information.

### **Setup of drive encryption in SSAM**

SSAM can manage the encryption keys associated with hardware based encryption in the IBM TS1120 tape drives. Because the encryption keys are stored within the SSAM database, you cannot read the content of an encrypted tape with a tape drive that is not assigned to that SSAM server. You should not use drive encryption for SSAM database backups to tape!

The drive encryption will be turned on or off in the device class definition of the tape drives with the `driveencryption` option. Three options are available: ALLOW (the default), ON, or OFF:

- ▶ **ON:** Specifies that SSAM is the key manager for drive encryption and will permit drive encryption for *empty* volumes only if the application method is enabled. If you specify ON and you enable either the library or system method of encryption, drive encryption will not be permitted and backup operations will fail.
- ▶ **ALLOW:** Specifies that SSAM does not manage the keys for drive encryption. However, drive encryption for empty volumes is permitted if either the library or system method of encryption is enabled.
- ▶ **OFF:** Specifies that drive encryption will not be permitted. If you enable either the library or system method of encryption, backups will fail. If you enable the application method, SSAM will disable encryption, and backups will be attempted.

The following simplified example shows the steps you would take to permit the encryption of data for empty volumes in a storage pool:

1. Log in at the active DR550 SSAM server to AIX with the `dr550adm` user ID.
2. Use the `dsmdmc` command at the DR550 SSAM server to start an administrative SSAM session. Log in with the SSAM user ID `admin`.
3. Define a library with the SSAM command:

```
define library library_name libtype=SCSI
```

Example:

```
define library 3584 libtype=SCSI
```

4. Define a device class with the SSAM command:

```
define devclass device_class_name library=library_name driveencryption=on
```

Example:

```
define devclass 3592_encrypt library=3584 driveencryption=on
```

5. Define a storage pool with the SSAM command:

```
define stgpool pool_name device_class_name maxscratch=maximum_scratch_tapes
```

Example:

```
define stgpool 3592_encrypt_pool 3592_encrypt maxscratch=10
```

## 5.2.10 Data shredding

Data shredding was introduced with SSAM Version 5.4. In this section, we describe the two different shredding methods and their setup.

After client data has expired, it might still be possible to recover it. For sensitive data, this condition is a potential security exposure. The destruction of deleted data, also known as *shredding*, lets you store sensitive data so that it is overwritten one or more times after it has expired. This process increases the difficulty of discovering and reconstructing the data later. System Storage Archive Manager performs shredding only on data in random access disk storage pools. Shredding occurs only after a data deletion commits, but it is not necessarily completed immediately after the deletion (this is controlled by the parameters `automatic` or `manual`). The space occupied by the data to be shredded remains occupied while the shredding takes place and is not available as free space for new data until the shredding is complete.

**Important:** The default shred value for the predefined ARCHIVEPOOL storage pool is 3. If you do not want to shred your data, you have to disable data shredding, otherwise the SSAM server cannot reclaim the disk space from expired or moved objects until a manual shredding is started. The procedure for disabling data shredding is described in “Disable data shredding” on page 188.

There are two different shredding methods available: automatic and manual.

**Tip:** If there is a need to use shredding, we recommend the manual shredding method for the DR550. The advantage of the manual method is that it can be performed when it will not interfere with other server operations.

## Configure the shredding method in SSAM

You have to configure which shredding method will be used by SSAM. Only one method, automatic or manual, can be activated. This section shows you how to query and change the current shredding method.

### Determine the current shredding method

1. Use the **dsmdmc** command at the DR550 SSAM server to start an administrative SSAM session. Log in with the SSAM user ID admin.
2. Enter the **query option shredding** command. The output gives you the information about the active shredding method (see Example 5-4).

*Example 5-4 Output of the query option shredding command*

---

tsm: TSM>**query option shredding**

| Server Option | Option Setting |
|---------------|----------------|
| -----         | -----          |
| SHREDDing     | Manual         |

---

### Change the current shredding method

To change the shredding method used by SSAM:

1. Enter the **setopt shredding automatic** or the **setopt shredding manual** command. This will change the default shredding method to the given parameter.
2. Answer the Do you wish to proceed? question with y (see Example 5-5).

*Example 5-5 Output of the setopt shredding command*

---

tsm: TSM>**setopt shredding automatic**

Do you wish to proceed? (Yes (Y)/No (N)) y

ANR2119I The SHREDDING option has been changed in the options file.

---

## Automatic shredding

The advantage of automatic shredding is that it is performed without administrator intervention whenever deletion of data occurs. This limits the time that sensitive data might be compromised. Automatic shredding also limits the time that the space used by deleted data is occupied. Shredding performance is affected by the amount of data to be shredded, the number of times that data is to be overwritten, and the speed of the disk and server hardware.

You can specify that the data is to be overwritten up to 10 times. The greater the number of times, the greater the security, but also the greater the impact on server performance.

**Note:** If you want to shred your data more than one time, we strongly recommend that write caching be disabled for any disk devices used to store sensitive data. If write caching is enabled, the overwrite operations are adversely affected.

### ***Set up automatic shredding***

To set up automatic shredding:

1. Disable the write cache of the DR550 Storage Controller (DS4200). This step is only necessary if your shred value is greater than 1.
  - a. Connect to DR550 SSAM Server #1 for Model DR1 and Model DR2 single-node or connect to DR550 SSAM Server #2 for Model DR2 dual-node systems.
  - b. Log in to AIX at the console with the dr550 user ID. Switch to root authority by entering `su -`.
  - c. Enter the command `SMcli 192.168.4.101 192.168.5.102 -c "set alllogicaldrives writecacheenabled=false;"`. This will disable the write cache for all logical volumes at the DR550 Storage Controller.
2. Use the `dsmdmc` command at the DR550 SSAM server to start an administrative SSAM session. Log in with the SSAM user ID admin.
3. Set the shredding method to automatic with the `setopt shredding automatic` command. See "Configure the shredding method in SSAM" on page 184 for details.
4. Change the shredding counter with the SSAM command `update stgpool ARCHIVEPOOL shred=5`. This will change the shredding counter for the predefined ARCHIVEPOOL to 5. (see Example 5-6). With the shredding counter parameter, you can enter the amount of overwrites that complies with your security policy. Possible values are 0-10. A value of 0 will turn off the shredding.

Each object will be automatically shredded immediately after deletion or expiration.

#### *Example 5-6 Output of update stgp command*

---

```
tsm: TSM>update stgpool ARCHIVEPOOL shred=5
ANR2202I Storage pool ARCHIVEPOOL updated.
```

---

Refer to the *Tivoli Storage Manager Administration's Guide* within the IBM Tivoli software Information Center for more information about data shredding:

<http://publib.boulder.ibm.com/infocenter/tivihelp>

**Note:** Disabling the DR550 Storage Controller write cache might affect the write performance of your system.

### **Manual shredding**

The advantage of manual shredding is that it can be performed when it will not interfere with other server operations. Manual shredding is possible only if automatic shredding is disabled.

Note that to guarantee that all shreds are written to the disk, disk caching needs to be disabled while the shred is being run. Therefore, shredding should be done when archiving of data is at a minimum. If you do most of the archiving during the day, shredding could be scheduled to run during the night.

## Set up manual shredding

To set up manual shredding:

1. Disable the DR550 Storage Controller write cache with the **SMcli** command (requires root authority).
  - a. Connect to DR550 SSAM Server #1 for Model DR1 and Model DR2 single-node or connect to DR550 SSAM Server #2 for Model DR2 dual-node systems.
  - b. Log in to AIX at the console with the dr550 user ID. Switch to root authority by entering **su -**.
  - c. Enter the command **SMcli 192.168.4.101 192.168.5.102 -c "set alllogicaldrives writecacheenabled=false;"**. This will disable the write cache for all logical volumes at the DS4200 disk subsystem.
2. Use the **dsmadm** command at the DR550 SSAM server to start an administrative SSAM session. Log in with the SSAM user ID admin.
3. Set the shredding method to automatic with the **setopt shredding manual** command. See "Configure the shredding method in SSAM" on page 184 for more details.
4. Change the shredding counter with the SSAM command **update stgpool ARCHIVEPOOL shred=5**. This will change the shredding counter for the predefined ARCHIVEPOOL to 5. (see Example 5-6 on page 185). With the shredding counter parameter, you can enter the amount of overwrites that comply with your security policy. Possible values are 0-10. A value of 0 will turn off the shredding.
5. Start the shredding process manually with the SSAM command **shred data**.

Refer to the *Tivoli Storage Manager Administration's Guide* within the IBM Tivoli software Information Center for more information about data shredding:

<http://publib.boulder.ibm.com/infocenter/tivihelp>

**Note:** Disabling the DR550 Storage Controller write cache might affect the write performance of your system.

## Using the *datashred.sh* script

A script is provided by IBM to enable multiple manual shreds. It is located in /usr/bin and is called *datashred.sh*. This script can be used as needed, or it can be set up as a regular task that runs at a specified time.

The advantage of this script is that the write cache of the DR550 Storage Controller will only be disabled during the shredding process. This avoids the performance impact of a permanently disabled DS4200 write cache.

The following actions will be done by the *datashred.sh* script (see Example 5-7 on page 187):

- ▶ Disable (temporarily) the write cache for both DS4200 controllers
- ▶ Log in to SSAM and run the **shred data du=120** command (shredding data for two hours)
- ▶ Enable write cache for both DS4200 controllers at



#### Example 5-7 The *datashred.sh* script

---

```
#!/usr/bin/ksh
#
datashred.sh: execute the data shredding process for 2 hrs daliy
#
#
SMcli 192.168.4.101 192.168.5.102 -c "set alllogicaldrives
writecacheenabled=false;"
/usr/tivoli/tsm/client/ba/bin/dsmadm -id=admin -password=admin shred data du=120
sleep 7300
SMcli 192.168.5.102 192.168.4.101 -c "set alllogicaldrives
writecacheenabled=true;"
```

---

The modifications of the following parameters might be necessary in this script:

- ▶ SSAM admin password: Change *-password=admin*
- ▶ Duration time (in minutes) before the shredding process will be cancelled: *du=120*
- ▶ IP address of DR550 Storage Controllers: *192.168.5.102 192.168.4.101*

To modify the script, start the editor with **vi /usr/bin/datashred.sh**.

#### **Run the script *datashred.sh***

This section describes the two ways to run the *datashred.sh* script: manually or scheduled within the AIX crontab.

To run this script manually at the command line, you need root authority!

Follow these steps:

1. Log in to AIX with the dr550 user ID and switch to root authority with **su-**.
2. To run the script, enter **/usr/bin/datashred.sh** at the command line.

You can insert this script into your AIX crontab to run it on a regular basis. For example, if you want to run the data shredding each day at 11 p.m., do the following:

1. Log in to AIX with the dr550 user ID and switch to root authority with **su-**.
2. Open the crontab with the AIX command **crontab -e**.
3. Scroll to the bottom of the file and append the following line:

```
0 23 * * * /usr/bin/datashred.sh
```

4. Save the file and exit the editor.

#### **Query the shredding status**

If you want to determine the number of objects waiting for shredding, you can use the SSAM command **query shredstatus**:

1. Use the **dsmadm** command at the DR550 SSAM server to start an administrative SSAM session. Log in with the user ID **admin**.

2. Issue the SSAM command **q shred f=d** (see Example 5-8). The output gives you information about the amount and the occupied space of objects awaiting the shredding.

*Example 5-8 query shredstatus command*

---

```
tsm: TSM>q shred f=d
```

```
Shredding Active: No
Objects Awaiting Shred: 0
Occupied Space (MB): 0
Data Left To Shred (MB): 0
```

---

### Disable data shredding

If you do not want to use the data shredding functionality, you have to change the shredding value in the storage pool definition of the ARCHIVEPOOL.

Follow these steps:

1. Use the **dsmdmc** command at the DR550 SSAM server to start an administrative SSAM session. Log in with the user ID **admin**.
2. Set the shredding value to 0 with the command **update stgpool1 ARCHIVEPOOL shred=0**.
3. Shred all expired data with the **shred data** command.

## 5.2.11 Device support for data retention

SSAM supports more than 400 storage devices. These are the same devices that are supported by Tivoli Storage Manager Extended Edition. Depending on the regulatory requirement that customers are trying to meet, there might or might not be specific types of media required.

Most regulations allow the stored data to be on any type of device as long as the content management application establishes a retention policy. This ability for the content management application to establish a retention policy is changing, such as in the case of the old paradigm of having regulatory data stored on optical media, and has opened up the ability to store the data on other types of media such as disk and tape.

**Tip:** We recommend the IBM System Storage TS1120 Enterprise Tape Drive in combination with the IBM System Storage 3592 WORM media, or the new generation of IBM Ultrium 4LTO drives in combination with the 3589 WORM media to extend the IBM System Storage DR550 characteristics for non-erasable and non-rewritable data to the tape storage pool.

For more information about WORM media support, see Chapter 8, “Managing Removable Media Operations”, and the section titled “Special Considerations for WORM Tape Media” in the *IBM Tivoli Storage Manager for AIX Administrator's Guide Version 5.5*, SC32-0117.

## 5.2.12 IBM System Storage Archive Manager and IBM System Storage DR550

System Storage Archive Manager is installed on each delivered IBM System Storage DR550.

Therefore, the DR550 SSAM server executable files are installed locally on each of the DR550 SSAM Servers (p5 520 servers) in a single-node or dual-node system. This is the standard configuration for a DR550 SSAM server running within an HACMP cluster in the case of a dual-node configuration.

For the DR550 Model DR1, all configuration files, the SSAM database, the recovery log, and the storage pool volumes are configured on specific file systems within the internal drives in the DR550 SSAM Server. This will differ from time to time when the option to purchase an external disk is available, and then the storage pool volumes will be configured on the external disk, not the internal disk.

For the DR550 Model DR2, all configuration files, the SSAM database, the recovery log, and the storage pool volumes are configured on specific logical volumes residing on the DR550 Storage Controller. With the dual-node configuration, all those files or resources are shared in the HACMP cluster.

## **What is predefined on the System Storage Archive Manager server**

The definitions in the following sections apply.

### ***System Storage Archive Manager database and database volumes***

For the DR550 Model DR1, the SSAM database volumes are defined on the internal drives in the DR550 SSAM Server. For the DR550 Model DR2 single-node or dual-node, database volumes reside on an AIX JFS2 file system created on DR550 Storage logical drives. The DR550 Storage arrays are configured as RAID 5.

**Note:** Starting with DR550 V4.5, it is possible to use/order the DR550 Storage with RAID 6 (refer to Chapter 2, “DR550 and File System Gateway overview” on page 9 for details on RAID 6).

Other important characteristics are:

- ▶ **SSAM DB size:** The size of the preconfigured database volumes is determined based on the total size of storage ordered (see Figure 5-11 on page 190 for an example of the SSAM DB configuration).
- ▶ There is additional space available for emergency expansion of the database files if needed. This space is in the same volume group as the existing database and log files. 10 GB is reserved for the primary database, 10 GB reserved for the secondary copy, and 20 GB is reserved for the backup space. If you expand the primary database, be sure to expand the secondary copy and the backup space as well.
- ▶ System Storage Archive Manager database volumes created on a JFS2 file system. These volumes are mirrored by SSAM (see Figure 5-12 on page 190).
- ▶ DBBackuptrigger: 75%.
- ▶ System Storage Archive Manager database mirrored: Yes.
- ▶ There is a predefined administrative schedule DBB, which is set for daily full database backups (at 6:00) to device class DBBKUP. This device class is a FILE device class. There is also a database backup trigger defined, which automatically starts a full database backup when the recovery log utilization reaches 75%. The files that are created by this backup process are placed in the /tsmdbbkup directory. The system will maintain two backup volumes. Older backup volumes will be automatically deleted. This is scheduled by another administrative schedule called DELVOLH.

| Available<br>Space<br>(MB) | Assigned<br>Capacity<br>(MB) | Maximum<br>Extension<br>(MB) | Maximum<br>Reduction<br>(MB) | Page<br>Size<br>(bytes) | Total<br>Usable<br>Pages | Used<br>Pages | Pct<br>Util | Max.<br>Pct<br>Util |
|----------------------------|------------------------------|------------------------------|------------------------------|-------------------------|--------------------------|---------------|-------------|---------------------|
| 300,000                    | 300,000                      | 0                            | 299,932                      | 4,096                   | 76,800,00                | 80,759        | 0.1         | 0.1                 |

Figure 5-11 SSAM DB information for DR550 Model DR2: 48 TB configuration (query db)

| Volume Name<br>(Copy 1) | Copy<br>Status | Volume Name<br>(Copy 2) | Copy<br>Status | Volume Name<br>(Copy 3) | Copy<br>Status |
|-------------------------|----------------|-------------------------|----------------|-------------------------|----------------|
| /tsmDb/dbvo10           | Sync'd         | /tsmDbM/dbvo10          | Sync'd         |                         | Undefined      |
| /tsmDb/dbvo11           | Sync'd         | /tsmDbM/dbvo11          | Sync'd         |                         | Undefined      |
| /tsmDb/dbvo12           | Sync'd         | /tsmDbM/dbvo12          | Sync'd         |                         | Undefined      |
| /tsmDb/dbvo13           | Sync'd         | /tsmDbM/dbvo13          | Sync'd         |                         | Undefined      |
| /tsmDb/dbvo14           | Sync'd         | /tsmDbM/dbvo14          | Sync'd         |                         | Undefined      |
| /tsmDb/dbvo15           | Sync'd         | /tsmDbM/dbvo15          | Sync'd         |                         | Undefined      |
| /tsmDb/dbvo16           | Sync'd         | /tsmDbM/dbvo16          | Sync'd         |                         | Undefined      |

Figure 5-12 SSAM DB volumes (query dbvol)

### System Storage Archive Manager recovery log and recovery log volumes

For the DR550 Model DR1, the SSAM recovery log volumes are defined on the internal drives in the DR550 SSAM Server. For the DR550 Model DR2 single-node or dual-node, the SSAM recovery log volume resides on an AIX JFS2 file system created on DR550 Storage logical drives. The DR550 Storage Controller arrays are configured as RAID 5.

**Note:** Since DR550 V4.5, there is the possibility to use the DR550 Storage with RAID 6.

Other important characteristics are:

- ▶ SSAM recovery log available space: The size of the preconfigured log volumes is determined based on the total size of storage ordered (see Figure 5-13).
- ▶ SSAM recovery log volumes created on a JFS2 file system (see Figure 5-14).
- ▶ Recovery log mirrored: Yes.
- ▶ Recovery log mode: Roll forward (to support point-in-time restores of the database).

| Available<br>Space<br>(MB) | Assigned<br>Capacity<br>(MB) | Maximum<br>Extension<br>(MB) | Maximum<br>Reduction<br>(MB) | Page<br>Size<br>(bytes) | Total<br>Usable<br>Pages | Used<br>Pages | Pct<br>Util | Max.<br>Pct<br>Util |
|----------------------------|------------------------------|------------------------------|------------------------------|-------------------------|--------------------------|---------------|-------------|---------------------|
| 13,000                     | 12,500                       | 500                          | 12,488                       | 4,096                   | 3,199,488                | 1,857         | 0.1         | 0.8                 |

Figure 5-13 SSAM recovery log for DR550 Model DR2: 48 TB configuration (query log)

| Volume Name<br>(Copy 1) | Copy<br>Status | Volume Name<br>(Copy 2) | Copy<br>Status | Volume Name<br>(Copy 3) | Copy<br>Status |
|-------------------------|----------------|-------------------------|----------------|-------------------------|----------------|
| /tsmLogM/logvo10        | Sync'd         | /tsmLog/logvo10         | Sync'd         |                         | Undefined      |
| /tsmLogM/logvo11        | Sync'd         | /tsmLog/logvo11         | Sync'd         |                         | Undefined      |

Figure 5-14 SSAM recovery log volumes (query logvol)

Table 5-4 lists the volume naming conventions used for the DR550 Storage logical drives that are mapped to the AIX servers.

*Table 5-4 Naming convention for the logical volumes*

| AIX volume group | AIX logical volume | AIX file system    | IBM System Storage Archive Manager usage |
|------------------|--------------------|--------------------|------------------------------------------|
| TSMApps          | tsmappslv          | /tsm               | Configuration files                      |
| TSMDBLogs        | tsmdblv0           | /tsmDb             | Database                                 |
|                  | tsmloglv0          | /tsmLog            | Recovery log                             |
| TSMDBLogsMirr    | tsmdblv0           | /tsmDbM            | Database mirror                          |
|                  | tsmloglv0          | /tsmLogM           | Recovery log mirror                      |
| TSMDBBkup        | tsmdbbkuplv        | /tsmdbbkup         | Database backups                         |
| TSMStg           | tsmstglv1          | Raw logical volume | Storage pool volume                      |
|                  | tsmstglv2          | Raw logical volume | Storage pool volume                      |
|                  | ...                | Raw logical volume | Storage pool volume                      |
|                  | tsmstglvn          | Raw logical volume | Storage pool volume                      |

### **DR550 Storage configuration**

It is possible to order the DR550 Model DR1 with additional storage; in that case, this section will apply. If you did not order the DR550 Model DR1 with additional storage, this section will not apply to the DR550 Model DR1.

In the SMclient graphical interface view for the logical and physical configuration of the DR550 Storage, you can see the volume sizes and the positioning of the logical drive TSMDBBkup, which is used for SSAM database backups.

**Tip:** The SSAM database volumes and the SSAM recovery log volumes are mirrored across different arrays of the DR550 Storage Controller.

### **DR550 Storage Controller and Expansion Drawer configuration summary**

The following is a summary of the DR550 Storage configuration:

- ▶ DR550 Storage Controller has two RAID 5 arrays configured. The preferred path for array 1 goes to controller A and the preferred path of array 2 goes to controller B.
- ▶ One hot spare disk is in the DR550 Storage Controller to recover single disk failures for the RAID 5 arrays in the DR550 Storage Controller.
- ▶ Each DR550 Expansion Drawer (EXP420) will contain two RAID 5 disk arrays. The hot spare coverage is provided by one hot spare disk per DR550 Expansion Drawer (EXP420).

Figure 5-15 illustrates the array and the volume configuration of a DR550 Storage with 750 GB SATA disk drives.

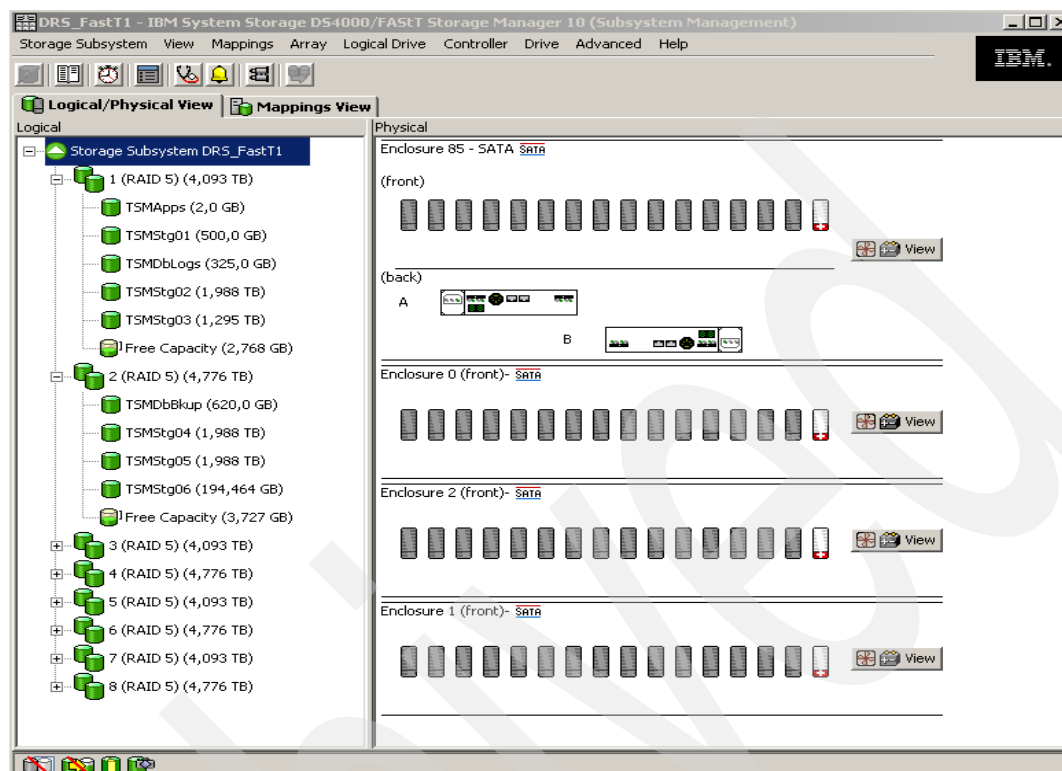


Figure 5-15 Physical configuration of the DR550 Storage in a 48 TB configuration

### **AIX Logical Volume Manager (LVM) configuration**

DR550 Model DR1 ships with internal disks in two RAID configurations. The first two disks are on a RAID 1 array and are used for the operating system. Drives 3 – 8 are in a RAID 5 array (5+P). All data (SSAM database, logs, database backup space, and archive spare) resides on the RAID 5 array within the internal disks.

For the DR550 Model DR1 ordered with a DR550 Storage Controller (DS4200), and for DR550 Model DR2, LUNs are configured based on the total storage capacity ordered. The whole disk capacity is allocated for use by SSAM. No disk capacity is set aside for other applications or servers.

For detailed information about the LUN configurations set as the defaults, refer to Appendix C, “DS4000 Logical Volume configuration”, of the *IBM System Storage DR550 Version 4.5 Problem Determination and Service Guide*, GA32-0576.

Figure 5-16 on page 193 shows the complete AIX LVM layout and volume group design for the DR550 Model DR2 single and dual nodes, including the DR550 Storage logical drives. AIX logical volumes can be formatted with a file system, as is the case for the volume groups rootvg, TSMApps, TSMdbLogs, and TSMdbBkup (JFS2). System Storage Archive Manager volumes defined in these file systems are basically AIX files residing in the directory structure of the file system. Contrary to that, the logical volumes within the TSMStg volume group do not carry a file system, but are accessed as raw logical volumes by System Storage Archive Manager and make up the primary pool ARCHIVEPOOL.

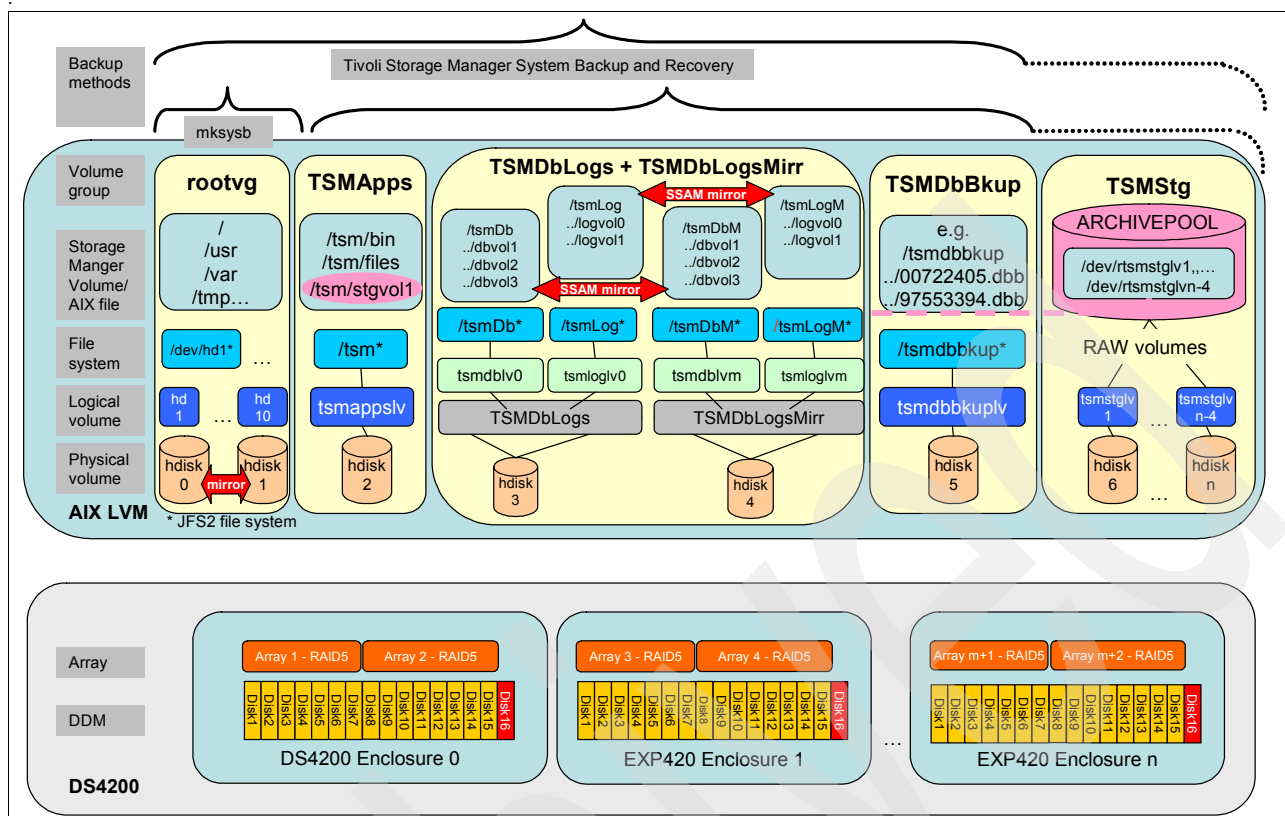


Figure 5-16 AIX LVM design and DR550 Storage configuration for DR550 Model DR2

### Predefined primary storage pools and storage pool volumes

**Note:** A comparably small file (5 MBs) called `/tsm/stgvol1` that is in the volume group TSMApps is also allocated to the Tivoli Storage Manager ARCHIVEPOOL. This is the only System Storage Archive Manager volume in that pool not residing in the volume group TSMStg.

All remaining space on the DR550 Storage arrays is divided into 2,000 GB LUNs (the last LUN of each array is always smaller because of the remaining space of the array) and mapped to the AIX servers. Once mapped, they are configured as AIX physical volumes (*hdisks*). On top of these *hdisks*, logical volumes are configured with the AIX Logical Volume Manager (LVM). Each of these logical volumes has been defined to the primary storage pool archivepool as a raw logical volume. These volumes together make up the total capacity of the primary storage pool.

SSAM is configured with one storage pool volume (500 GB) set at the factory. The remaining volumes are predefined in AIX volume group TSMStg, but not assigned as storage pool volumes to the System Storage Archive Manager.

To add the predefined volumes to the SSAM storage pool, stop HACMP (if applicable) and SSAM (on both DR550 SSAM Servers in dual-node systems). From the AIX command line, issue the command `/usr/bin/addstg`. Note that this must be run from DR550 SSAM Server #1 on dual-node systems. The script creates `/tmp/bldstg`. This SSAM script lists all storage pool volumes that can be added to the SSAM Storage Pool. If you do not want to add the whole capacity, you can edit the `bldstg` SSAM script and remove the logical volume lines.

To execute the bldstg script, log in to SSAM and run the macro /tmp/bldstg. Now the storage pool volumes will be added to the ARCHIVEPOOL within SSAM. This may take a few minutes to complete. After the script completes, you will need to restart HACMP if applicable and SSAM (see 4.1, “Starting and stopping HACMP cluster services” on page 138).

To check the storage pool, issue the commands **q vol** and **q stg** in SSAM.

### Device classes

Figure 5-17 shows predefined device classes on the DR550 SSAM Server:

- ▶ Device class DISK  
Used by *primary storage pool archivepool*.
- ▶ Device class DBBKUP, device type FILE  
DBBKUP uses the device type FILE. It is a sequential access device class that has been predefined and is used for full database backups that are run daily as specified in an administrative schedule on the DR550 SSAM Server. The sequential access files (volumes) created by this process are located in the /tsmdbbkup file system.

| Device Class Name | Device Access Strategy | Storage Pool Count | Device Type | Format |
|-------------------|------------------------|--------------------|-------------|--------|
| DBBKUP            | Sequential             | 0                  | FILE        | DRIVE  |
| DISK              | Random                 | 1                  |             |        |

Figure 5-17 Predefined device classes (SSAM command query devclass)

### Predefined administrative schedules

There are three predefined administrative schedules that are executed daily on the Tivoli Storage Manager server, as listed in Table 5-5. You can check the schedules with the SSAM command **query schedules type=admin f=d**.

Table 5-5 Predefined administrative schedules

| Schedule name | Start time | Runs  | Command                                                             |
|---------------|------------|-------|---------------------------------------------------------------------|
| DBB           | 6:00       | Daily | <b>backup db type=full devclass=dbbkup</b>                          |
| BAVOLH        | 10:00      | Daily | <b>backup volhist file=/usr/tivoli/tsm/server/bin/backupvolhist</b> |
| DELVOLH       | 23:59      | Daily | <b>delete volhist type=dbb todate=-2</b>                            |

The predefined administrative schedules are:

1. Schedule DBB starts a full database backup of the SSAM database using the device class DBBKUP. The resulting backup volumes are created in the /tsmdbbkup directory. This schedule runs daily at 6:00. Figure 5-18 on page 195 shows the output of a **query volhist** command showing a sequential access volume created by the backup database process.



```

Date/Time: 02/27/08 06:07:58
Volume Type: BACKUPFULL
Backup Series: 133
Backup Operation: 0
Volume Seq: 1
Device Class: DBBKUP
Volume Name: /tsmdbbkup/04117678.DBB
Volume Location:
Command:

```

Figure 5-18 Output of query volhist t=dbb

2. Schedule BAVOLH will back up the volume history file to the /usr/tivoli/tsm/server/bin/backupvolhist file.
3. Schedule DELVOLH deletes records in the volume history for database backups that are older than two days.

These schedules are preconfigured to provide the appropriate protection. However, they can be customized to fit the customer's environment. For example, if the full database backup at 6:00 interferes with a heavy archiving workload, the start time can be modified to another time that fits better into the workload profile. For information about how to update schedules, refer to the *IBM Tivoli Storage Manager for AIX Administrator's Guide Version 5.5*, SC32-0117.

**Important:** These administrative schedules protect the DR550 System Storage Archive Manager server database with a daily backup. We recommend that you check this process on a daily basis.

### Performance settings in dsmserv.opt

The performance tuning parameters shown in Table 5-6 have been set to optimize the mainstream DR550 SSAM Server environment. However, the workloads and configurations of specific customer environments might benefit from additional performance tuning and optimization. Use the *IBM Tivoli Storage Manager Performance and Tuning Guide V5.5*, SC32-0141, which can be found at:

<http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp>

Table 5-6 IBM Tivoli Storage Manager server option file settings

| Option      | Tivoli Storage Manager V5.5 Standard | DR550 SSAM Server factory setting |
|-------------|--------------------------------------|-----------------------------------|
| BUFPOOLSIZE | 32768 KB                             | 512000 KB                         |
| MAXSESSIONS | 25                                   | 100                               |
| TXNGROUPMAX | 256                                  | 512                               |
| COMMTIMEOUT | 60                                   | 300                               |

The following factors can affect the System Storage Archive Manager (SSAM) server's performance significantly:

- ▶ Average client file size
- ▶ Client hardware (CPUs, RAM, disk drives, and network adapters)
- ▶ Client activity (non-SSAM workload)
- ▶ Server hardware (CPUs, RAM, disk drives, and network adapters)
- ▶ Server storage pool devices (disk, tape, and optical)

- ▶ Server operating system
- ▶ Server activity (non-SSAM workload)
- ▶ Network hardware and configuration
- ▶ Network utilization
- ▶ Network reliability
- ▶ Communication protocol
- ▶ Communication protocol tuning
- ▶ Final output repository type (disk, tape, and optical)

### **Installed default policy settings**

The policies that are automatically defined during the DR500 SSAM server installation are the installed default policies. These are intended to be updated or replaced to suit your data retention requirements. Before you begin setting up your policies, you must have a thorough understanding of your data retention requirements and the policy settings that govern the retention and expiration of the data you intend to store on the DR550 SSAM Server.

### **Installed default policy domains**

Two policy domains are preconfigured: STANDARD and DRG-DOMAIN.

#### ▶ STANDARD

The default policy domain STANDARD is predefined on the DR550 SSAM server. This is the default policy for archiving through the API or the Web-based archive client. You can either edit this policy domain to suit your data retention requirements, or you can create new policy domains using your own naming conventions. Figure 5-19 shows the default settings in the policy domain STANDARD.

```
tsm: TSM>q domain standard f=d

 Policy Domain Name: STANDARD
 Activated Policy Set: STANDARD
 Activation Date/Time: 05/10/05 14:40:13
 Days Since Activation: 741
 Activated Default Mgmt Class: ARCHIVE_EVENT
 Number of Registered Nodes: 1
 Description:
Backup Retention (Grace Period): 60
Archive Retention (Grace Period): 90
Last Update by (administrator): ADMIN
 Last Update Date/Time: 05/10/05 14:40:13
 Managing profile:
 Changes Pending: No
 Active Data Pool List:
```

*Figure 5-19 Default settings in the policy domain STANDARD*

#### ▶ DRG-DOMAIN

The default policy domain DRG-DOMAIN is predefined on the DR550 SSAM server. This is the default domain for the optional DR550 File System Gateway (FSG). Figure 5-20 on page 197 shows the default settings for the policy domain DRG-DOMAIN.

```

tsm: TSM>q domain drg-domain f=d

 Policy Domain Name: DRG-DOMAIN
 Activated Policy Set: STANDARD
 Activation Date/Time: 05/04/07 23:25:38
 Days Since Activation: 17
 Activated Default Mgmt Class: DRG-DEFAULTMC
 Number of Registered Nodes: 1
 Description:
Backup Retention (Grace Period): 30
Archive Retention (Grace Period): 365
Last Update by (administrator): SERVER_CONSOLE
 Last Update Date/Time: 05/04/07 23:25:38
 Managing profile:
 Changes Pending: No
 Active Data Pool List:

```

*Figure 5-20 Default settings for the policy domain DRG-DOMAIN*

### **Installed default policy sets**

There are two default policy sets predefined, one for each domain: STANDARD and DRG-DOMAIN.

Figure 5-21 shows the active policy set in the STANDARD domain and Figure 5-22 shows the active policy set in the DRG-DOMAIN.

```

tsm: TSM>query policyset standard active f=d

 Policy Domain Name: STANDARD
 Policy Set Name: ACTIVE
 Default Mgmt Class Name: ARCHIVE_EVENT
 Description:
Last Update by (administrator): ADMIN
 Last Update Date/Time: 05/10/05 14:39:38
 Managing profile:
 Changes Pending: No

```

*Figure 5-21 Default settings for the active policy set STANDARD*

```

tsm: TSM>q policyset drg-domain active f=d

 Policy Domain Name: DRG-DOMAIN
 Policy Set Name: ACTIVE
 Default Mgmt Class Name: DRG-DEFAULTMC
 Description:
Last Update by (administrator): SERVER_CONSOLE
 Last Update Date/Time: 02/20/08 01:15:44
 Managing profile:
 Changes Pending: No

```

*Figure 5-22 Default settings for the active policy set DRG-DOMAIN*

### Installed default management classes

The management classes STANDARD and DRG-DEFAULTMC are predefined as default management classes for their policy domains (see Figure 5-23). Defining additional management classes pointing to appropriate archive copy groups with different retention rules is the recommended way to separate objects with different retention requirements within the DR550 SSAM Server. The management class is the distinguishing attribute used by a document management application to feed objects into the DR550 SSAM Server. If the document management system does not specify a management class at the delivery of an object to the DR550 SSAM Server, the default management class STANDARD will be used to store the object.

```
tsm: TSM>q mgmtclass
```

| Policy Domain Name | Policy Set Name | Mgmt Class Name | Default Mgmt Class ? | Description |
|--------------------|-----------------|-----------------|----------------------|-------------|
| DRG-DOMAIN         | ACTIVE          | DRG-DEFAULTMC   | Yes                  |             |
| DRG-DOMAIN         | ACTIVE          | DRG-SYSMC       | No                   |             |
| DRG-DOMAIN         | STANDARD        | DRG-DEFAULTMC   | Yes                  |             |
| DRG-DOMAIN         | STANDARD        | DRG-SYSMC       | No                   |             |
| STANDARD           | ACTIVE          | ARCHIVE_EVENT   | Yes                  |             |
| STANDARD           | STANDARD        | ARCHIVE_EVENT   | Yes                  |             |

Figure 5-23 Default management classes

### Preconfigured management classes

The classes are:

#### ► ARCHIVE\_EVENT

One management class, ARCHIVE\_EVENT, is preconfigured in the factory (see Figure 5-24). This class is used to write a small amount of test data into the DR550 SSAM Server with a retention period of one year. With active archive data in the system, the DR550 SSAM server archive retention protection cannot be disabled, which adds another level of protection from corruption to the offering. This management class is not meant to be used by the customer.

```
tsm: TSM>q mgmtclass standard
```

| Policy Domain Name | Policy Set Name | Mgmt Class Name | Default Mgmt Class ? | Description |
|--------------------|-----------------|-----------------|----------------------|-------------|
| STANDARD           | ACTIVE          | ARCHIVE_EVENT   | Yes                  |             |
| STANDARD           | STANDARD        | ARCHIVE_EVENT   | Yes                  |             |

Figure 5-24 Preconfigured management class ARCHIVE\_EVENT for test data

► DRG-DEFAULTMC and DRG-SYSMC

Two management classes, DRG-DEFAULTMC and DRG-SYSMC, are preconfigured for the optional DR550 FSG in the factory (see Figure 5-25). These classes are used exclusively by the DRG software (the FSF application code) to archive data and back up the DR550 FSG metadata.

```
tsm: TSM>q mgmtclass drg-domain
```

| Policy Domain Name | Policy Set Name | Mgmt Class Name | Default Mgmt Class ? | Description |
|--------------------|-----------------|-----------------|----------------------|-------------|
| DRG-DOMA-IN        | ACTIVE          | DRG-DEFA-ULTMC  | Yes                  |             |
| DRG-DOMA-IN        | ACTIVE          | DRG-SYSMC       | No                   |             |
| DRG-DOMA-IN        | STANDARD        | DRG-DEFA-ULTMC  | Yes                  |             |
| DRG-DOMA-IN        | STANDARD        | DRG-SYSMC       | No                   |             |

Figure 5-25 Default settings for management class DRG-DOMAIN

**Preconfigured default archive copy groups**

There are three archive copy groups predefined, one for each predefined management class. The default archive copy groups are shown in Figure 5-26. If additional archive copy groups are created, they will always be named STANDARD. Because they are always tied to a specific management class, this is the criteria to differentiate them.

```
tsm: TSM>q copyg * active * type=archive
```

| Policy Domain Name | Policy Set Name | Mgmt Class Name | Copy Group Name | Retain Version |
|--------------------|-----------------|-----------------|-----------------|----------------|
| DRG-DOMA-IN        | ACTIVE          | DRG-DEFA-ULTMC  | STANDARD        | 0              |
| DRG-DOMA-IN        | ACTIVE          | DRG-SYSMC       | STANDARD        | 0              |
| STANDARD           | ACTIVE          | ARCHIVE_-EVENT  | STANDARD        | 365            |

Figure 5-26 Settings for default archive copy group STANDARD

► Default copy class for DR550 SSAM Server

In the default archive copy group for the management class ARCHIVE\_EVENT, the retention initiation (RETINIT) is set to event, the retain minimum (RETMIN) parameter is set to 365 days, and the retain version (RETVR) parameter is set to 365 days.

Files archived using this copy group are held until they get the event signal for a minimum of 365 days. If an event signal was sent, the data will be retained 365 days from the date the signal was sent; otherwise, the data expires at the end of the minimum retention period and is deleted from System Storage Archive Manager storage.

**Note:** A deletion hold can still be applied during the 365 day period, which prevents that object from being deleted from storage until a deletion release is applied for that same object.

If you want to change the parameters for the default copy class, you can use the **update copyclass** command.

Example 5-9 shows how to change the minimum retention period to 200 days.

*Example 5-9 Change the minimum retention period*

```
tsm: TSM>update copygroup standard standard archive_event type=archive retmin=200
```

ANR1618W UPDATE COPYGROUP: the new RETMIN value (200) is less than the value previously stored (365) for STANDARD STANDARD ARCHIVE\_EVENT.

ANR1537I Archive copy group STANDARD updated in policy domain STANDARD, set STANDARD, management class ARCHIVE\_EVENT.

► **Default copy classes for the optional DR550 FSG**

There are two copy classes predefined for the optional DR550 FSG. For the copy group for the DRG-DEFAULTMC management class, the retention initiation (RETINIT) is set to event, the retain minimum (RETMIN) parameter is set to 1 day, and the retain version (RETVER) parameter is set to 0 days. For the copy group for the DRG-SYSMC management class, the retention initiation (RETINIT) is set to event, the retain minimum (RETMIN) parameter is set to 0 days, and the retain version (RETVER) parameter is set to 0 days. Files archived from the DR550 FSG to DR550 SSAM Server will be held until they get an ACTIVATE event from the DR550 FSG and then will expire. The DR550 FSG is managing the retention period for data that is archived through the DR550 FSG. You must modify the DRG profile to configure the retention period to your requirements for all objects that are archived through the FSG. Usually there is no need to modify the predefined values. (See the *IBM System Storage DR550 File System Gateway Software Version 1.1.1 Integration Guide*, GC27-2124 for more information about modifying DRG profiles.)

**Preconfigured client nodes**

Two client nodes are preconfigured (see Figure 5-27). The node DR550\_1 is for the purpose of generating the test data, as previously mentioned, to protect the DR550 SSAM server from being disabled for retention protection. This registered node is not intended to be used by the customer. The node DRG\_NODE is used by the optional DR550 File System Gateway to archive data, transferred through CIFS or NFS protocol to the DR550 FSG, to the DR550 SSAM server.

```
tsm: TSM>q node
```

| Node Name | Platform     | Policy Name | Domain | Days Since Last Access | Days Since Password Set | Locked? |
|-----------|--------------|-------------|--------|------------------------|-------------------------|---------|
| DR550_1   | Sample-- API | STANDARD    |        | 619                    | 741                     | No      |
| DRG-NODE  | Bycast DRG   | DRG-DOMAIN  |        | <1                     | 17                      | No      |

Figure 5-27 Preconfigured client nodes

### 5.2.13 IBM System Storage Archive Manager policy examples

In this section, we illustrate with two examples possible customer requirements and show how to create the corresponding rules within the DR550 SSAM Server. The commands to create a node, a domain, a policy set, a management class, and an archive copy group, and then how they are validated and activated, are shown.

#### Example one: event-based retention policy

A company is required to keep all records and communications pertaining to customer transactions and correspondence for at least seven years after a transaction is final. The records must be stored in such a way that they cannot be altered or prematurely deleted (records must be non-erasable and non-rewritable). In the case of pending litigation, these records must not be deleted until all litigation has ended. This means that the company might have to keep this data much longer than seven years. This data must also be readily available for review and proof of compliance.

Data pertaining to a transaction is being created long before the transaction is final. The correspondence and consulting can start in January, but the deal might not be final until December. The data is being archived as it is created, meaning that data created in January must be retained longer than data created in June or August. We probably would not choose to chronologically archive here (RETINIT=creation). We must be able to get the data on the same time line as when the transaction is deemed final so that all of the data will expire together.

#### Retention initiation method

We use event-based retention for this solution, meaning that the RETINIT parameter in the archive copy group is set to *event* for the management class to which this data is bound. Data archived with event-based retention is stored in Tivoli Storage Manager initially as though the retain version were set to *no limit* (RETVER=nolimit). This means that the data is kept indefinitely. When the application storing the data initiates the retention by sending a retention initiation event called *activate* through the API to the DR550 SSAM server (in our case, when the transaction becomes final), the value of the retain version (RETVER=2555) goes into effect (Retention Initiated: STARTED), and all documents activated expire 2555 days or seven years later.

What if the files must be retained even longer? For example, they are needed as evidence in a court case that starts six years after the transaction, but could take two or three years to resolve?

In this case, the application could send a retention event *hold* for the necessary data. SSAM now keeps this data even after the expiration date has passed. When the litigation is over, and the data is no longer needed, the application can release the data by sending a *release* retention event for these objects. These objects then expire according to RETVER=2555 (from the date the retention was initiated earlier with *activate*).

The following list shows the commands used to define, assign, validate, and activate the policy that fulfills the requirements:

- ▶ **define domain compliance archretention=2555**
- ▶ **define policyset compliance compliance\_set**
- ▶ **define mgmtclass compliance compliance\_set compliance\_class\_event**
- ▶ **define copygroup compliance compliance\_set compliance\_class\_event type=archive destination=archivepool retinit=event retver=2555 retmin=2555**
- ▶ **assign defmgmt compliance compliance\_set compliance\_class\_event**

- ▶ **validate policyset compliance compliance\_set** (Ignore warnings about the missing backup copygroup!)
- ▶ **activate policyset compliance compliance\_set** (Ignore warnings about the missing backup copygroup!)
- ▶ **register node node\_1 password domain=compliance**

Figure 5-28 shows the policy defined for our example.

```
query copygroup type=archive format=detailed

Policy Domain Name: COMPLIANCE
Policy Set Name: ACTIVE
Mgmt Class Name: COMPLIANCE_CLASS_EVENT
Copy Group Name: STANDARD
Copy Group Type: Archive
Retain Version: 2,555
Retention Initiation: Event
Retain Minimum Days: 2,555
Copy Serialization: Shared Static
Copy Frequency: CMD
Copy Mode: Absolute
Copy Destination: ARCHIVEPOOL
Last Update by (administrator): ADMIN
Last Update Date/Time: 03/23/06 22:19:43
Managing profile:
```

*Figure 5-28 Policy defined for example one*

The same information is depicted in Figure 5-29 on page 203 using the policy structure template that we introduced in 5.1.2, “IBM System Storage Archive Manager basic concepts” on page 163.



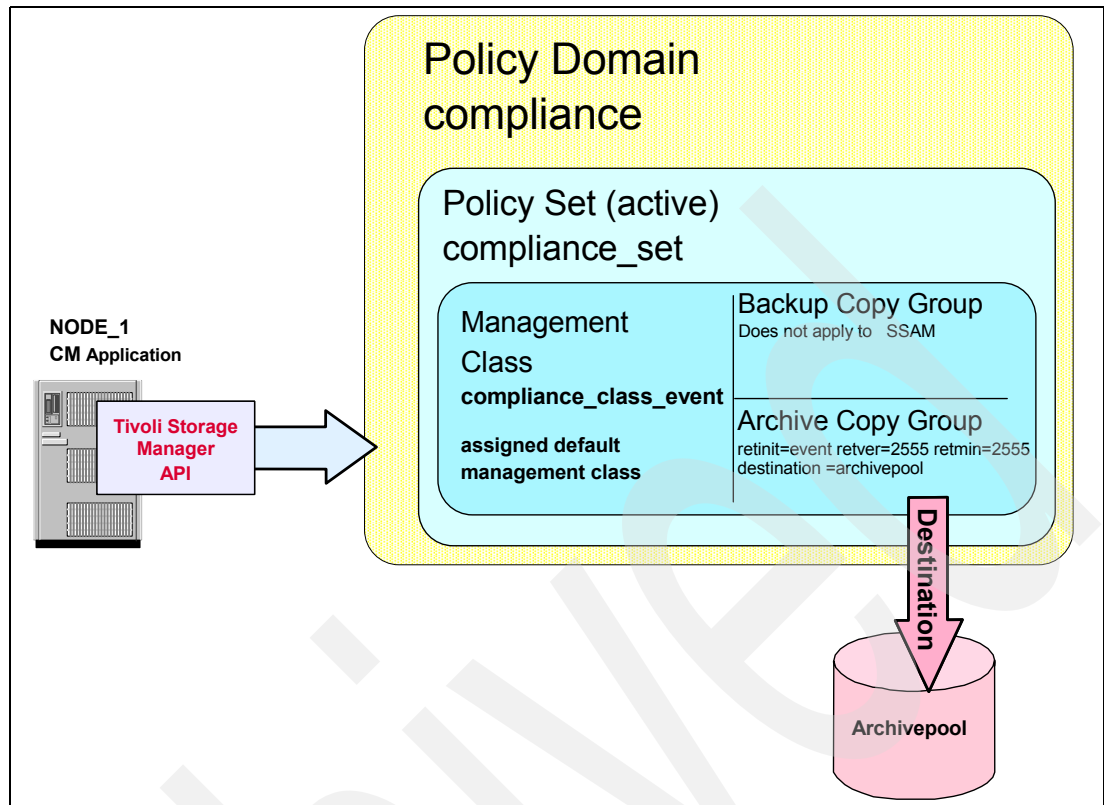


Figure 5-29 Policy structure template for example one

### Example two: event-based retention policy

The finance department of a business is required to keep payroll information for employees for three years after they retire, but at least 30 years in total. When an employee retires, a retention event activate is sent to the Tivoli Storage Manager for Data Retention server by the application through the API that manages the employee record's retention. This event starts the counter of RETVER=1095, which ensures that the records are kept for the required period of three years after retirement. If this period has passed, and there is still time remaining for RETMIN=10950, the Tivoli Storage Manager server will keep the record for the remaining time. Otherwise, the object becomes eligible for deletion.

Here we list the commands used to define, assign, validate, and activate the policy that fulfills the requirements:

- ▶ **define domain financial archretention=10950**
- ▶ **define policyset financial standard**
- ▶ **define mgmtclass financial standard payroll**
- ▶ **define copygroup financial standard payroll standard type=archive destination=archivepool retinit=event retver=1095 retmin=10950**
- ▶ **assign defmgmt financial standard payroll**
- ▶ **validate policyset financial standard** (Ignore warnings about the missing backup copygroup!)
- ▶ **activate policyset financial standard** (Ignore warnings about the missing backup copygroup!)
- ▶ **register node "node name" "password" domain=financial**

Figure 5-30 shows the defined policy for our example.

```
query copygroup type=archive format=detailed

 Policy Domain Name: COMPLIANCE
 Policy Set Name: ACTIVE
 Mgmt Class Name: COMPLIANCE_CLASS_EVENT
 Copy Group Name: STANDARD
 Copy Group Type: Archive
 Retain Version: 2,555
 Retention Initiation: Event
 Retain Minimum Days: 2,555
 Copy Serialization: Shared Static
 Copy Frequency: CMD
 Copy Mode: Absolute
 Copy Destination: ARCHIVEPOOL
 Last Update by (administrator): ADMIN
 Last Update Date/Time: 03/23/06 22:19:43
 Managing profile:
```

Figure 5-30 Policy defined for example two

## 5.3 Explore archive retention features

There are two methods to explore the archive retention features. The first is by using the Tivoli Storage Manager BA Client V5.3.2 (Web Client), which can be launched remotely or directly from the client machine. Alternatively, you can use the Tivoli Storage Manager API that comes with a sample application called *dapismp*. We discuss and illustrate both methods in the remainder of this section.

### 5.3.1 SSAM/Tivoli Storage Manager BA Client V5.5

The SSAM/Tivoli Storage Manager BA client component sends data to, and retrieves data from, a DR550 SSAM server. The SSAM/Tivoli Storage Manager client must be installed on every machine that will transfer data to server-managed storage. The DR550 SSAM server uses a unique node name to identify each client instance. A password can be used to authenticate communications between the SSAM/Tivoli Storage Manager client and server. Data can be recovered from the same client machine that initially transferred it, or to another client with a compatible file system format.

The BA client basically consists of the software component and a customization file. This customization file, called the client options file (*dsm.opt*), specifies client/server communications parameters and other SSAM/Tivoli Storage Manager client settings. Client communications parameters must agree with those specified in the server options file. The client options file is located in the client directory and can be modified using a text editor.

The BA Client allows archiving data to a System Storage Archive Manager. This will only be possible if you have enabled the client for archive retention protection in the *dsm.opt* file. Remember that once enabled for data retention, you can no longer use the BA client to do backups (see Figure 5-31 on page 205). You can only archive data (no backups) when connecting to an IBM System Storage Archive Manager (SSAM) server installed on a IBM System Storage DR550.



Figure 5-31 Trying to back up data with archive protection enabled

**Note:** To set retention events for archived data, like hold or release a file, you must use the Web-based client. The regular BA client does not provide these functions!

In the sections that follow, we explain how to install, configure, and use the BA client for archive retention and protection.

### Install and configure BA Client V5.5 for data retention

The steps below apply to a Windows environment:

1. Download the SSAM/Tivoli Storage Manager Client V5.5 or later version. You can find the current maintenance levels of the software at:  
<ftp://ftp.software.ibm.com/storage/tivoli-storage-management/maintenance/client/v5r5/Windows/x32/v550>  
 Within the download directory, download the self-extracting executable client code. Refer to the readme (5.5.0.0-TIV-TSMBAC-WinX32-README.FTP) file in the same directory; for example, the code to download should be a file named 5.5.0.0-TIV-TSMBAC-WinX32.exe.
2. Start the installation by extracting the client code in 5.5.0.0-TIV-TSMBAC-WinX32.exe.
3. In the first window (Location to Save Files), choose a folder where the client software can be unpacked. In our case, it is done in c:\tsm\_images\TSM\_BA\_CLIENT. Click **Next**.  
 The install wizard extracts all the files into the specified directory.
4. Once the install wizard has completed the extraction, the setup wizard starts executing. In the Choose Setup Language window, choose your language, such as **English (United States)**, and click **OK**.
5. In the Welcome to the Install Wizard window, click **Next**.
6. In the Destination Folder window, select the installation folder, such as C:\Program Files\Tivoli\tsm\, and then click **Next**.
7. In the Setup Type window, leave the default setting as Typical and click **Next**.
8. In the Ready to Install the Program window, click **Install**. The InstallShield Wizard starts installing the software.
9. When the InstallShield Wizard Completed window opens, check that the installation is successful, and click **Finish**. If the install failed, correct the problem and repeat the installation.
10. If there is no dsm.opt file in the BA client installation folder, copy the dsm.smp file from the C:\Program Files\Tivoli\tsm\config directory to the BA client installation folder and rename the sample option file from dsm.smp to dsm.opt.

11. Edit the dsm.opt file within the BA client installation folder (see Figure 5-32). Set the following parameters:

- IP Address of the DR550 SSAM Server (TCPServeraddress)
- commethod tcpip
- tcpport 1500
- enablearchiveretentionprotection yes
- passwordaccess generate

Save the file.




Figure 5-32 Example of a dsm.opt file

### BA client remote access (Web Client)

You need to set up the remote access function in order to access the Web client. The Web client only (or the CLI), but not the regular BA client, can (and must) be used to set retention events for archive data to a System Storage Archive Manager (as is the case with the IBM System Storage DR550).

**Note:** To set retention events for archived data, like hold or release a file, you must use the Web-based client. The regular BA client does not provide these functions!

You can use the GUI Setup Wizard of the BA client or command line to install and configure the Web client. We will describe both ways in the following section.

#### ***Installation of the Web Client through the GUI of the common BA client***

You must have installed and configured the native BA client before you can start this procedure, as shown in "Install and configure BA Client V5.5 for data retention" on page 205.

Perform the following steps:

1. Start your native BA client.
2. From the native client GUI main window, open the Utilities menu and select **Setup Wizard**.
3. Select the **Help me configure the TSM Web Client** check box and click **Next**.
4. Select **Install a new Web Client Agent** and click **Next** and then follow the instructions on the window.

### ***Installation of the Web Client at the command line***

To install and configure the Web client from the command line, perform the following steps:

1. Ensure that you specify `passwordaccess generate` in the client options file (`dsm.opt`).
2. Install the Client Acceptor Service by entering the following command:

```
dsmcutil install cad /name:"TSMBA_web" /node:nodename /password:password
/autostart:yes
```

where `nodename` and `password` are your Storage Manager node name and password. `TSMBA_web` is an example. You can use any name you want. The default name is Tivoli Storage Manager Client Acceptor.

3. Install the Remote Client Agent Service by entering the following command:

```
dsmcutil install remoteagent /name:"TSM AGENT" /node:nodename
/password:password /partnername:"TSMBA_web"
```

where `nodename` and `password` are your Storage Manager node name and password. `TSM AGENT` is an example. You can use any name as long as it is different from the Client Acceptor Daemon (CAD) name. The default name is TSM Remote Client Agent. The `/partnername` option value must match the name of the CAD service. The default name is TSM Client Acceptor.

4. Start the Client Acceptor Service by entering `net start "TSM CAD"` on the command line, or do the following:
  - a. Open the Windows Start menu and select **Settings** → **Control Panel**.
  - b. Double-click **Administrative Tools** and then double-click **Services**.
  - c. In the Services window, right-click **TSMBA\_web**, and select **Start** from the pop-up menu. The window shown in Figure 5-33 should appear.

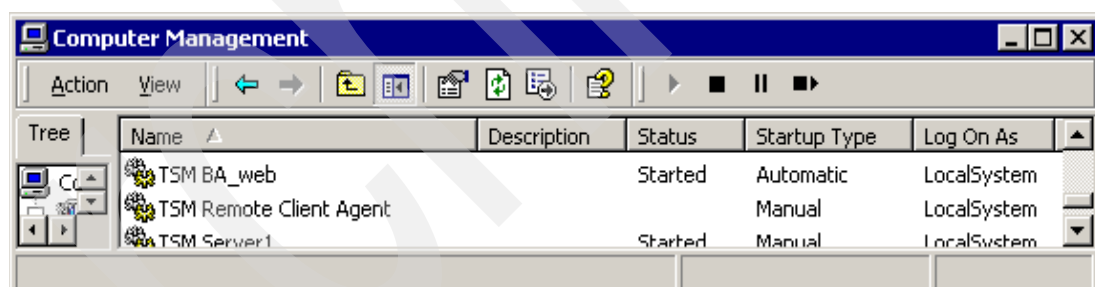


Figure 5-33 Services show Tivoli Storage Manager components

To access the Web client, enter the following URL from any supported Web browser:

`http://your_machine_name:1581`

where `your_machine_name` is the host name of the machine running the regular BA client.

The IBM Tivoli Storage Manager Web client interface for client machines requires a Java Swing-capable Web browser:

- ▶ Microsoft Internet Explorer® 6.0 or later with Java Plug-in 1.4.1
- ▶ Firefox 1.6 or later

For more information about how to set up the Web Client, refer to:

<http://www.redbooks.ibm.com/redbooks/pdfs/sg246762.pdf>

## Testing the archive function using the Web Client

To use the Web client for archiving, we define a new policy domain, policy set, and management classes in our DR550 SSAM Server using the administrative command line:

1. To create a policy domain named CLITEST\_PD, we use the following command:

```
define domain CLITEST_PD
```

2. Within the policy domain CLITEST\_PD, we create one policy set named CLITEST\_PS:

```
define policyset CLITEST_PD CLITEST_PS
```

3. We create two separate management classes for the purpose of testing creation-based retention and event-based retention:

```
define mgmtclass CLITEST_PD CLITEST_PS CLITEST_MG_CR
define mgmtclass CLITEST_PD CLITEST_PS CLITEST_MG_EV
```

4. We assign the first management class as the default:

```
assign defmgmtclass CLITEST_PD CLITEST_PS CLITEST_MG_CR
```

5. Next, we define archive copy groups (type=archive) for each of the management classes. The archive copy groups must be defined along with the appropriate parameters to differentiate between creation-based retention and event-based retention:

- Archive Copy Group (Chronological Archive):

```
define copygroup CLITEST_PD CLITEST_PS CLITEST_MG_CR type=archive
destination=archivepool retver=1825 retinit=creation
```

- Archive Copy Group (Event-Based Archive):

```
define copygroup CLITEST_PD CLITEST_PS CLITEST_MG_EV type=archive
destination=archivepool retver=365 retinit=event
```

6. We validate the Policy Set using the following command:

```
validate policyset CLITEST_PD CLITEST_PS
```

The command returns the information that the default management class does not have a backup copy group, and that files will not be backed up by default if policyset is activated. This message is normal and expected in our case because the DR550 is an archive-only solution.

7. We now activate the Policy Set:

```
activate policyset CLITEST_PD CLITEST_PS
```

8. Finally, we register the client node (CLITEST) that we use for the test:

```
register node CLITEST password domain=CLITEST_PD
```

We can now archive data using the creation-based management class (Example 1) or the event-based management class (Example 2). For each example, we show how to trigger retention events.

### **Example one**

Archive data using the Creation-based management class (Chronological Archive):

1. Launch the Web Client from a Web browser by entering the URL `http://BAclient_IP:1581`, where BAclient\_IP represents the address of the BA client and select some files you would like to archive, as shown in Figure 5-34 on page 209.

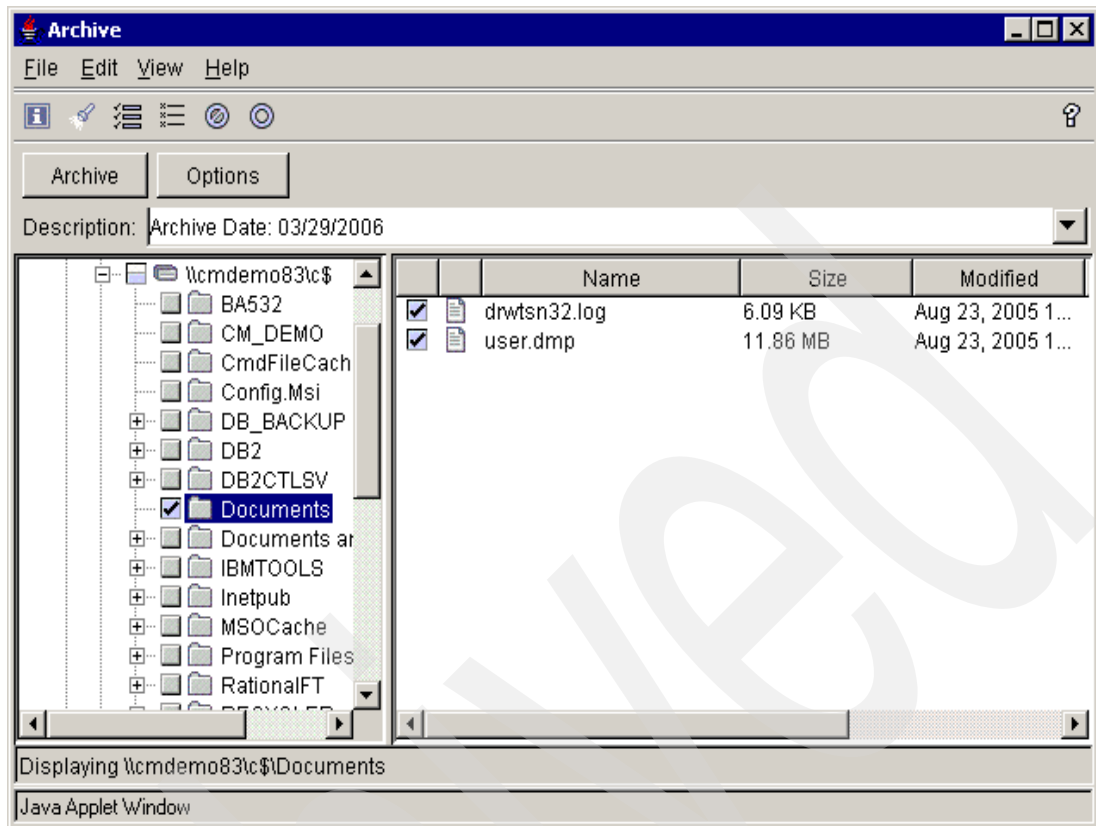


Figure 5-34 Test files archived to test Chronological Archive

2. Click the **Archive** tab to archive these files using the default (creation-based management class). Once the Archive is complete, the message box shown in Figure 5-35 displays.




Figure 5-35 Archive complete

You can verify that the data that has been archived and that it has adopted the correct management class as well as the correct retention period. You can see an example in Figure 5-36. Notice that the status of Retention Initiation is *Started*. This is correct, because with chronological-based retention, the retention period starts counting down as soon as the data has been archived.

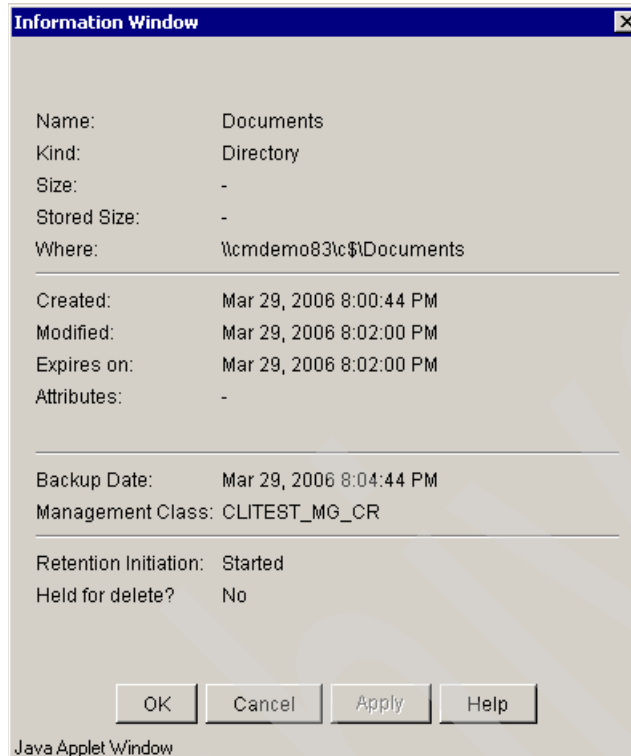


Figure 5-36 Example of Chronological Archive

It is possible to put a hold on the archived data by first selecting the data that is required to be held, then selecting **Hold** from the drop-down menu for **Select Event Type**, and clicking **SetEvent**. See Figure 5-37.

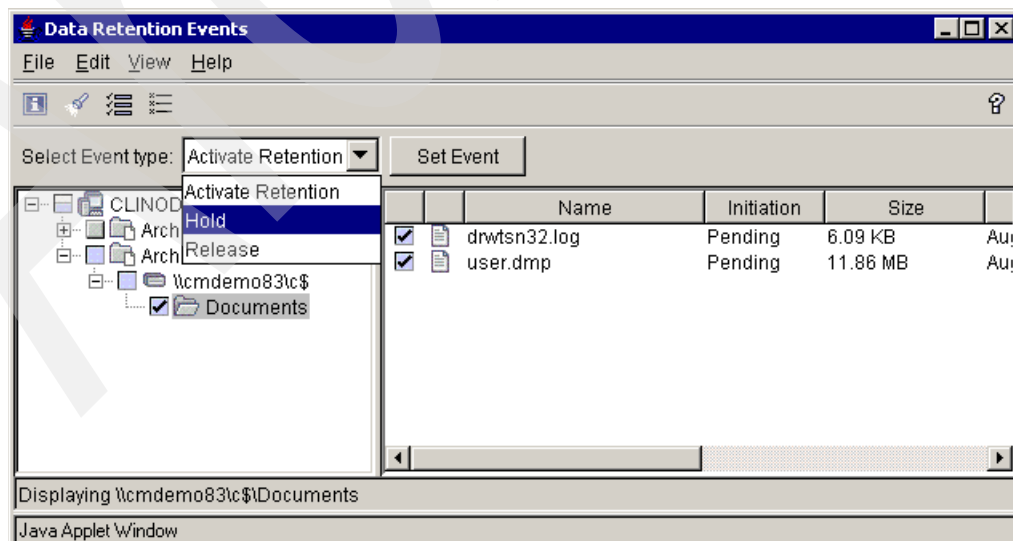


Figure 5-37 Example of how to set a Hold event



You can see in Figure 5-38 that items on hold are indicated by a lock.

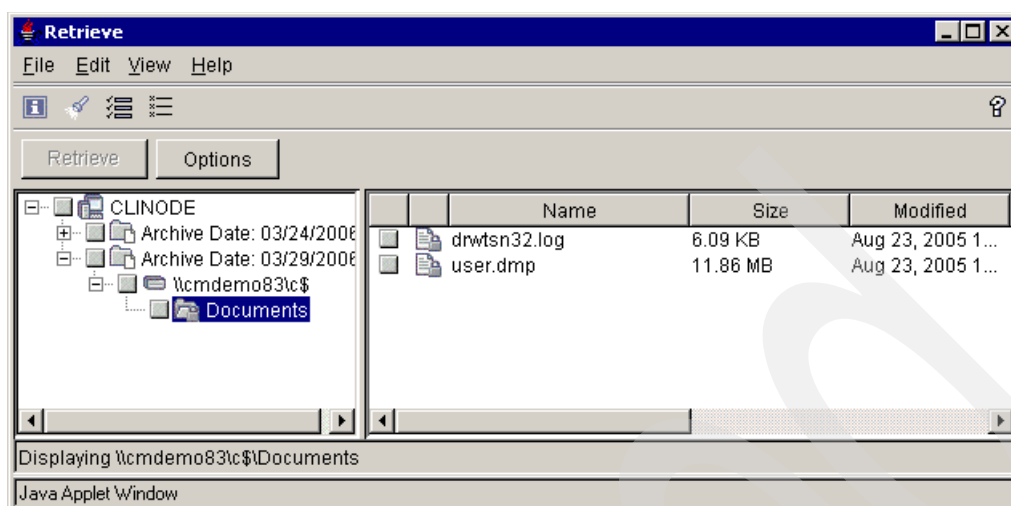


Figure 5-38 Hold event set

The selected data will now be held indefinitely, until a release event is triggered by the user.

To release the hold, select **Release** from the drop-down menu for **Select Event Type**, and click **SetEvent** (see Figure 5-37 on page 210). The countdown towards expiration resumes as though it was never put on hold.

### Example two

Example two shows Archive data using the Event-based management class (Event Archive).

We follow the same process as in example one:

1. Invoke the Web client, and select files to archive.
2. Select **Options** → **Override Include Exclude List** and choose the desired Management Class. We select **CLITEST\_MG\_EV** (the management class we created for Event-Based Retention). See Figure 5-39.




Figure 5-39 Changing the Management Class from the BA Client before archiving

You can now verify the characteristics of the archived data by selecting one of the files you just archived and clicking **View** → **File Details**. The result is shown in Figure 5-40. Notice that in this case that the Retention Initiation shows as Pending, which is to be expected since we used Event-Based retention and no Activate event has been sent yet.



Figure 5-40 Example of event-based retention

The countdown to expiration starts once an Activate Retention event is sent for that file. Figure 5-41 on page 213 shows how to activate the retention: Select the file, then choose **Activate Retention** from the menu for the Select Event Type, and click **Event**.

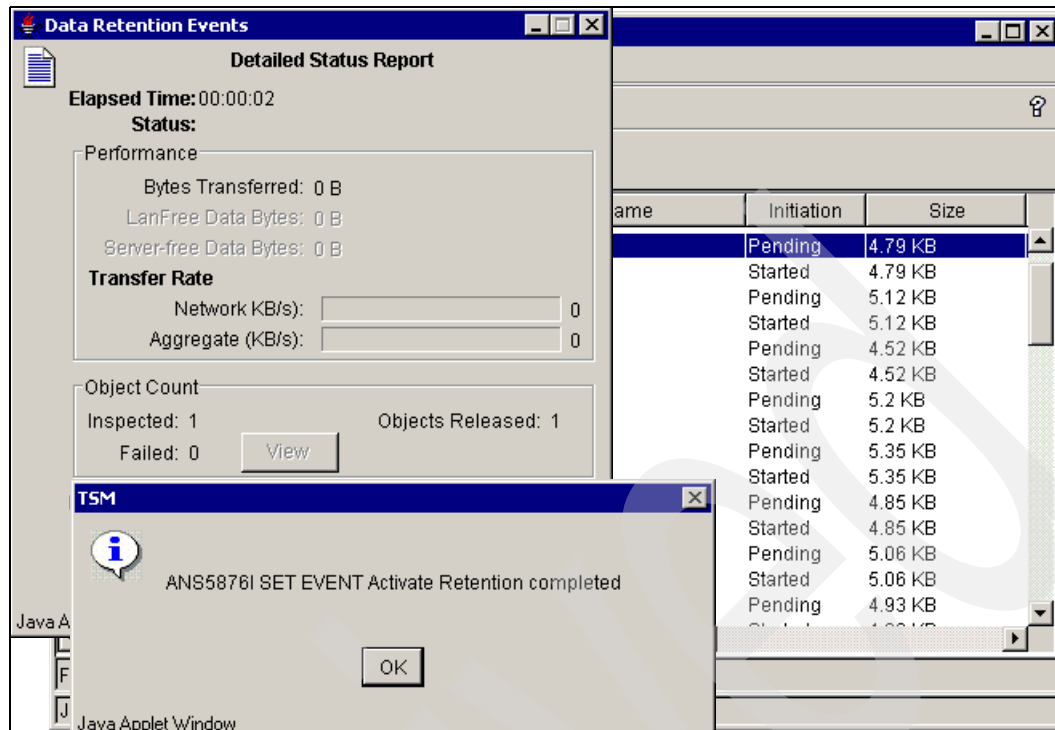


Figure 5-41 Set Activate Event

As seen in Figure 5-42, the file characteristics of this file have now changed from Retention Initiation pending to Retention Initiation Started.

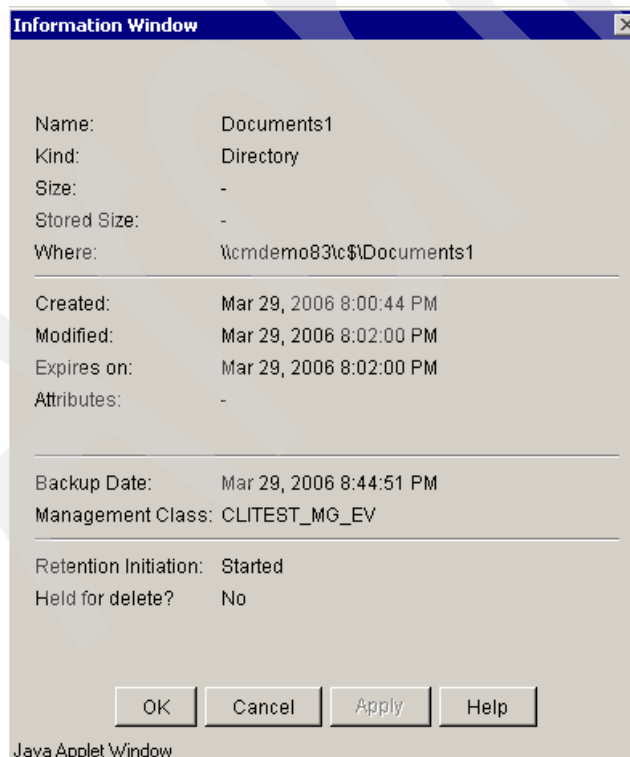


Figure 5-42 Activate Retention on file

The server will reject any attempt to delete the archived data, as shown in Figure 5-43.




Figure 5-43 Example of the data that, once archived, cannot be deleted from SSAM

### 5.3.2 SSAM/Tivoli Storage Manager API (using the sample application dapismp)

SSAM/Tivoli Storage Manager API comes with a sample application called *dapismp*. You can use this sample program to explore and better understand the new data retention and compliance-enhanced features.

The sample API program *dapismp* creates objects and feeds them to the retention policies of a previously defined management class. You can then use this program to query the DR550 SSAM server for information about the objects that were created and trigger retention events for these objects. We use *dapismp* throughout this section of the book as we explore the new features of SSAM/Tivoli Storage Manager. Furthermore, we use *dapismp* on a Microsoft Windows client system; in this environment, you can use the sample API program right after the installation and configuration of the API (on UNIX-based systems, you will need to compile the sample API program before you can run it). The executable file *dapismp.exe* can typically be found in the directory `C:\Program Files\Tivoli\TSM\api\SAMPRUN`, or an equivalent location, depending on where the SSAM/Tivoli Storage Manager client files have been installed. The *dapismp* sample API program requires a *dsm.opt* file in the same directory that must contain at least one of the following statements:

```
TCPSERVERADDRESS <IP address of TSM DR server>
ENABLEARCHIVERETENTIONPROTECTION yes
```

#### Testing the new features with dapismp

We demonstrate the following new features:

- Creation-based retention initiation (chronological archive)  
RETINIT=CREATION

Eligible retention events:

- Hold
- Release

► Event-based retention initiation

RETINIT=EVENT

Eligible retention events:

- Activate
- Hold
- Release

## Policy used during testing

For our tests, we set up a new policy domain named APITEST and defined two management classes. The assigned default management class is named CREATION and uses the creation-based retention initiation. The second management class is named EVENT and uses the event-based retention initiation. Figure 5-44 and Figure 5-45 show detailed information about the retention settings in each management class. Our test node is named apitest1 and is registered in the policy domain APITEST1.

```
Policy Domain Name: APITEST1
Policy Set Name: ACTIVE
Mgmt Class Name: CREATION
Copy Group Name: STANDARD
Copy Group Type: Archive
Retain Version: 1825
Retention Initiation: Creation
Retain Minimum Days:
Copy Serialization: Shared Static
Copy Frequency: CMD
Copy Mode: Absolute
Copy Destination: ARCHIVEPOOL
Last Update by (administrator): ADMIN
Last Update Date/Time: 03/23/2006
Managing profile:
```

Figure 5-44 Archive copy group settings for management class CREATION

```
Policy Domain Name: APITEST1
Policy Set Name: ACTIVE
Mgmt Class Name: EVENT
Copy Group Name: STANDARD
Copy Group Type: Archive
Retain Version: 365
Retention Initiation: Event
Retain Minimum Days: 730
Copy Serialization: Shared Static
Copy Frequency: CMD
Copy Mode: Absolute
Copy Destination: ARCHIVEPOOL
Last Update by (administrator): ADMIN
Last Update Date/Time: 03/23/2006 10:26:33
Managing profile:
```

Figure 5-45 Archive copy group settings for management class EVENT

The management class CREATION has been updated to be the default management class (see Figure 5-46). This means that objects delivered (by dapismp or a document management system) through the API to the DR550 SSAM server without a specific management class assigned will be stored in the DR550 SSAM Server with the policies of the standard management class, in this case, CREATION.

```
tsm: TSM>query mgmtclass apitest standard
```

| Policy Domain Name | Policy Set Name | Mgmt Class Name | Default Mgmt Class ? | Description |
|--------------------|-----------------|-----------------|----------------------|-------------|
| APITEST            | STANDARD        | CREATION        | Yes                  |             |
| APITEST            | STANDARD        | EVENT           | No                   |             |

Figure 5-46 Default management class CREATION

Using the sample API program dapismp

To use the sample API program dapismp, complete the following steps:

- 1. Start dapismp and sign in (connect to the DR550 SSAM Server).  
To start dapismp on a Microsoft Windows client system:
  - a. Start a command window and change to the C:\Progra~1\tivoli\TSM\api\SAMPRUN directory (or the appropriate installation directory).
  - b. At the command prompt, type **dapismp** and press Enter, which starts the dapismp executable and brings you to the first screen, as shown in Figure 5-47.

**Note:** The actual screens contain more options than those shown here. In the interest of saving space, we show only the minimum input needed to attain the desired results. We edited out the options not used in this example.

```

* Welcome to the sample application for the Tivoli Storage Manager API. *
* API Library Version = 5.3.0.8 (unicode) *

Choose one of the following actions to test:

0. Signon
1. Backup
2. Restore
3. Archive
4. Retrieve
5. Queries
6. Change Password
7. Utilities : Deletes, Updates, Logevent, SetAccess, RetentionEvent
8. Set preferences, envSetUp
9. Exit to system

Enter selection ==>
```

Figure 5-47 First window of sample API program dapismp after startup

2. Start a session with the DR550 System Storage Archive Manager Server.
  - a. Select the option 0. Signon to attempt a session with the DR550 SSAM server. The only information that must be provided here is your node name and password, as shown in Figure 5-48. The other fields can be skipped. An example of a successful signon is shown in Figure 5-48.

**Note:** If you have problems at this stage, check to see that the API environment variables DSMI\_DIR, DSMI\_CONFIG, and DSMI\_LOG have been set.

```

Enter selection ==>0
Node name:apitest1
Owner name:
Password:password
API Config file:
Session options:
User Name:
User pswd:
Are the above responses correct (y/n/q)?
y
Doing signon for node apitest1, owner , with password password
Handle on return = 1

```

Figure 5-48 Example of successful signon

- b. Submit the **query session** command on the DR550 SSAM Server to verify that a session was started. Figure 5-49 shows that the attempt was successful. Now that you have successfully signed on to the server, proceed to step 3 on page 218.

| Sess<br>Number | Comm.<br>Method | Sess<br>State | Wait<br>Time | Bytes<br>Sent | Bytes<br>Recvd | Sess<br>Type | Platform        | Client Name |
|----------------|-----------------|---------------|--------------|---------------|----------------|--------------|-----------------|-------------|
| 15             | Tcp/Ip          | IdleW         | 13 S         | 468           | 299            | Node         | Sample--<br>API | APITEST1    |

Figure 5-49 Output of query session command verifying the session

3. Create archive objects using dapism.

Use dapism to create two objects and archive them to the DR550 SSAM Server. Then look at their retention policies. Repeat this step and override the default management.

- a. From screen 1 of dapism, select option 3. Archive, as shown in Figure 5-50. You are prompted for information about the file that dapism creates and sends to the server. You are also prompted to enter the name of another management class, which would override the assigned default management class. Figure 5-50 shows the minimum input required to create the first object in the chronological management class. Repeat this step with different file name qualifiers and override the default management class with the event-based management class EVENT, as shown in Figure 5-51. Continue to the next step.

```
3. Archive
Enter selection ==>3
 Filespace:apitest1
 Highlevel:\
 Lowlevel:\test1
 Object Type(D/F):f
 Object Owner Name:
Object already compressed?(Y/N):
 Wait for mount?(Y/N):
 File size:1000000 (in bytes)
 Number of files:1
 Seed string:1
 Archive description:apitest1
 Mgmt class override:
Are the above responses correct (y/n/q)?
y
Creating 1 object(s) called apitest1\\test1(nnn) each of size 1,000,000.
Creating object 1 of 1 Size=1,000,000 Name=apitest1\\test1
```

Figure 5-50 Output of the archive function of dapism into a standard management class

```
Enter selection ==>3
 Filespace:apievent1
 Highlevel:\apievent1
 Lowlevel:\eventtest
 Object Type(D/F):f
 Object Owner Name:
Object already compressed?(Y/N):
 Wait for mount?(Y/N):
 File size:1000000
 Number of files:1
 Seed string:1
 Archive description:"test event based"
 Mgmt class override:event
Are the above responses correct (y/n/q)?
y
Creating 1 object(s) called apievent1\apievent1\eventtest(nnn) each of size 1,000,000.
Creating object 1 of 1 Size=1,000,000 Name=apievent1\apievent1\eventtest
 Object: 1 Buffer: 1 Bytes sent: 1,000,000 Bytes left: 0
```

Figure 5-51 Event-based retention overrides the management class



**Note:** As a reminder, the actual screens contain more options than those shown here. In the interest of saving space, we show only the minimum input needed to attain the desired results. We have edited out the options not used in this example.

#### 4. Query the DR550 SSAM server.

We now query the server and compare the policy information for both objects. Querying the DR550 SSAM server for archives can be done with the `dapismp` sample program or with select statements from the SSAM administrative command line. We show both methods here:

- From the first screen of the `dapismp` sample program, select option 5. Queries and then option 2. Archive Query in the following screen. Enter the name of the file space you want to query, which is required. In addition, the high-level and low-level qualifiers are required, as specified when the object has been created. In the low-level qualifier, a wildcard (\*) can be used. For detailed output, answer yes when prompted, as shown in Figure 5-52, and continue. Figure 5-54 on page 220 shows the output of the query.

```
Enter selection ==>2
 Filespace:apitest1
 Highlevel:\
 Lowlevel:*
 Object Type(D/F/A):f
Show detailed output? (Y/N):y
Are the above responses correct (y/n/q)?
y
```

Figure 5-52 Minimum input required for archive query using `dapismp`

- From a SSAM administrative command-line prompt, enter the following select statement:

```
select * from archives where node_name='APITEST1'
```

The output in Figure 5-53 shows that an object was archived to the server by node `APITEST1`; the object is bound to the default management class.

```

NODE_NAME: APITEST1
FILESPACE_NAME: apitest1
FILESPACE_ID: 1
TYPE: FILE
HL_NAME: \
LL_NAME: test1
OBJECT_ID: 3074
ARCHIVE_DATE: 2006-03-23 12:16:30.000000
OWNER:
DESCRIPTION: apitest1
CLASS_NAME: DEFAULT
```

Figure 5-53 Output of the select statement

5. Compare the results.

Examine the information that SSAM has associated with the objects. Figure 5-54 and Figure 5-55 on page 221 show the output of the archive query issued in the previous step. You can identify the parameters RETINIT and RETVER, which we discussed earlier in this book:

– RETINIT=creation

- The file that the dapismp program created was bound to the assigned default management class, in our case, management class CREATION (see Figure 5-54), which uses creation-based retention initiation (RETINIT=creation).
- Retention Initiated is STARTED (RETVER=*n* days is initiated).
- The Expiration date for this object is 2006/3/23 12:16:30.
- The Object Held is FALSE (deletion hold is not set).
- The high-level Object ID is 0-3074. This is important. You will need this information later.

When expiration processing runs on the server any time after 2011/3/2 12:16:30, this file will be deleted from the database, unless a “deletion hold” retention event is triggered for this object. We demonstrate this in “Sending retention events using dapismp” on page 221.

```
Item 1: apitest1\\test1
Object type: File
Desc: apitest1
Insert date: 2006/3/23 12:16:30
Expiration date: 2011/3/2 12:16:30
Owner:
Restore order: 4-0-35-0-0
Object id: 0-3074
Copy group: 1
Media class: Library
Mgmt class: DEFAULT
Object info is :Tivoli Storage Manager API Verify Data
Object info length is :60
Estimated size : 0 1000000
Retention Initiated: STARTED
Object Held : FALSE
```

Figure 5-54 Creation-based retention initiation: Output of select statement

– RETINIT=event

- The file that the dapismp program created was bound to the EVENT management class. (You chose to override the default and use the event management class; see Figure 5-51 on page 218.)
- The status of Retention Initiated is PENDING, because no retention “activate” event has been issued yet.
- The expiration date for this object is 65535/0/0 0:0:0 (the same is true when RETVER=nolimit).
- The status of Object Held is FALSE (the deletion hold is not set).
- The high-level Object ID is 0-3076. This is important. You will need this information later.

```

Item 1: apievent1\apievent1\eventtest
 Object type: File
 Desc: 1"test event based"
 Insert date: 2006/3/23 1:23:56
 Expiration date: 65535/0/0 0:0:0
 Owner:
 Restore order: 4-0-37-0-0
 Object id: 0-3076
 Copy group: 1
 Media class: Library
 Mgmt class: EVENT
 Object info is :Tivoli Storage Manager API Verify Data
 Object info length is :60
 Estimated size : 0 1000000
 Retention Initiated: PENDING
 Object Held : FALSE
Press any key to continue

```

Figure 5-55 Event-based retention initiation: Output from dapismp archive query

## Sending retention events using dapismp

To send retention initiation events using the dapismp sample program:

1. Starting from the first screen of dapismp, select option 7. Utilities, which brings you to the Utilities screen.
2. Select option 12. Retention Event. You are then prompted for the high-level object ID of the file for which you will trigger a retention event. In this case, the object ID is 0-3074.
3. Next, you are prompted for the low-level object ID, 3074 in this case.
4. Finally, you are prompted for the type of event you want to trigger. There are two possibilities for creation-based retention initiation: Hold (deletion hold) and Release (release the hold).
5. Select h for Hold and press Enter twice. Figure 5-56 and Figure 5-57 on page 222 show the output resulting from these actions.

```

7. Utilities : Deletes, Updates, Logevent, SetAccess, RetentionEvent
8. Set preferences, envSetUp
9. Exit to system

Choose one of the following actions:
.....
12. Retention Event

Enter selection ==>12
 Object ID (HI) to signal:0-3074
 Object ID (LOW) to signal:3074
Activate (A) Hold (H) Release (R):h
Are the above responses correct (y/n/q)?
y
Finished Retention Event successfully

```

Figure 5-56 Triggering retention events

This action triggers a deletion hold event for an archive object. Figure 5-57 shows that the object has a “deletion hold” status.

```
Item 1: apitest1\\test1
 Object type: File
 Desc: apitest1
 Insert date: 2006/3/23 12:16:30
 Expiration date: 2006/4/9 12:16:30
 Owner:
 Restore order: 4-0-35-0-0
 Object id: 0-3074
 Copy group: 1
 Media class: Library
 Mgmt class: DEFAULT
 Object info is :Tivoli Storage Manager API Verify Data
 Object info length is :60
 Estimated size : 0 1000000
 Retention Initiated: STARTED
 Object Held : TRUE
Press any key to continue
```

Figure 5-57 Deletion hold is set

The object is held, but there is no change to the expiration date.

### **Delete archive from application**

Because of data retention protection, attempting to delete the object (as we show in Figure 5-58) results in a failed status (which is the expected result).

```
10. Object Rename
11. Object Delete
12. Retention Event

Enter selection ==>11
 Object ID (HI) to DELETE:0-3074
 Object ID (LOW) to DELETE:3074
 Backup or Archive(B/A):a
Are the above responses correct (y/n/q)?
y

*** dsmDeleteObj failed:
ANS0266I (RC2302) The dsmEndTxn vote is ABORT, so check the reason field.

Choose one of the following actions:
```

Figure 5-58 Attempt to delete an archive object in hold status

Table 5-7 on page 223 illustrates the initial status of “Retention initiated” and “Object Held” after the creation of an object in the two different management classes. While “Retention initiated” is already STARTED in the chronological (CREATION) management class, it is still PENDING in the event-based class. This will change to STARTED as soon as the retention event activation has been issued through the API.

“Object Held” is FALSE for both, because no retention event hold has been issued. This will show TRUE once a hold is received and FALSE again after a release event.

Table 5-7 Initial status of files archived with creation-based and event-based retention

| Field                | RETINIT=creation   | RETINIT=event     |
|----------------------|--------------------|-------------------|
| Insert date:         | 2006/3/23 12:16:30 | 2006/3/23 1:23:56 |
| Expiration date      | 2011/3/2 12:16:30  | 65535/0/0 0:0:0   |
| Mgmt class:          | CREATION           | EVENT             |
| Retention Initiated: | STARTED            | PENDING           |
| Object Held:         | FALSE              | FALSE             |

For more information about the IBM SSAM/Tivoli Storage Manager API, consult *Tivoli Storage Manager Using the Application Program Interface V5.5*, SC32-0147.

## 5.4 IBM Tivoli Storage Manager operational reporting

The IBM Tivoli Storage Manager Operational Reporting feature automates some of the monitoring tasks you typically perform manually. By generating reports and monitors, operational reporting notifies you if the Tivoli Storage Manager server in the DR550 requires attention.

Operational reports can be scheduled to run daily and are generated even if there are no problems. Operational monitors are special types of reports and can be scheduled to run hourly. The monitors send you a notification only if there are issues. Operational reporting does not maintain a separate database of information and is not a trending tool.

Operational reporting is included as part of the IBM Tivoli Storage Manager for Windows server and is also available as a stand-alone package on a Windows server. Since the DR550 runs a Tivoli Storage Manager server on AIX, operational reporting is not part of the DR550. You have to use the stand-alone package, which is free of charge.

Operational reporting is administered through the Microsoft Management Console on a Windows machine. All platforms for IBM Tivoli Storage Manager servers, Version 5.1.8 or Version 5.2.2 and later, are supported. Operational reporting runs as a service and supports multiple Tivoli Storage Manager servers. The latter is useful if you need to administer not only the Tivoli Storage Manager server within the DR550, but other servers as well.

An operational report consists of the following: a standard report, customized summary, and optional extensions that you create. You can select which sections to include in the report. The operational reporting installation package contains two default custom summary templates: one for a report and one for a monitor.

Default e-mail messages notify you if the server is running smoothly, or if there are issues such as failed or missed schedules. You can also link to a Web summary page to check operational reports about your server. An operational monitor will notify you either through e-mail or by sending an instant message to your Windows desktop. Operational reporting can write data to a file that can be read by a Tivoli Enterprise Console® log file adapter. The log file adapter reads the information and forwards it to the Tivoli Enterprise Console.

For more information about operational reporting, see *IBM Tivoli Storage Manager for Windows Quick Start*, GC32-0784 or *IBM Tivoli Storage Manager for Windows Administrator's Guide*, GC32-0782.

Another detailed source of information about operational reporting is an IBM Redpaper titled *Integrating IBM Tivoli Storage Manager Operational Reporting with Event Management*, REDP-3850. This IBM Redpaper is available at:

<http://www.redbooks.ibm.com/redpapers/pdfs/redp3850.pdf>

### **Installing and configuring operational reporting**

We list the steps you must take to install and configure the stand-alone package of operational reporting to access the IBM Tivoli Storage Manager server of the DR550. You must perform all of the following steps on your designated management station:

1. Install operational reporting on the designated management station.
  - b. Download the most current IBM Tivoli Storage Manager management console for Windows from the Internet. You can find the most current maintenance levels of the software at:  
  
<ftp://ftp.software.ibm.com/storage/tivoli-storage-management/maintenance/server/v5r3/WIN/LATEST/>  
  
Within the download folder, download the self-extracting executable management console code. Refer to the readme file within the same folder to learn what the code is named, for example, a file named TSMCON5230\_WIN.exe.
  - c. Start the installation by starting the self-extracting executable client code, such as TSMCON5310\_WIN.exe.
  - d. Within the first window (InstallShield Wizard welcome page), click **Next** to continue.
  - e. In the next window (Location to Save Files) that displays, choose the folder where the software can be unpacked, such as c:\tsm\_images\TSMCON5310\_WIN, and click **Next**.  
  
The install wizard extracts the files and prepares to install.
  - f. In the upcoming window (InstallShield for the console welcome page), click **Next**.
  - g. You must read the license agreement carefully on the next page (License Agreement) and accept it by selecting the **I accept the terms in the license agreement** field. Click **Next** to proceed.
  - h. In the Customer Information window, complete the User Name and Organization fields with the appropriate information. In the lower area of the window, select **Anyone who uses this computer (all users)** to install the application for the use of all users. Click **Next** to continue.
  - i. Select **Complete** in the Setup Type window, and then **Next**.
  - j. If you want to review or change any of your installation settings, click **Back** in the Ready to Install the Program window. Click **Install** to begin the installation.
  - k. The installation ends in the last window (InstallShield Wizard Completed), and you click **Finish** to finish the installation and to close the window.

2. Configure operational reporting on the designated management station.

After a successful installation, you must configure the DR550 within the operational reporting tool.

It is impossible here to discuss every parameter. The only purpose of this section is to show how to establish the communication between IBM Tivoli Storage Manager in the DR550 and the operational reporting tool. Also, we show how to use one preconfigured report as well as one preconfigured monitor.

Do the following:

- a. In Windows, select **Start** → **Programs** → **Tivoli Storage Manager** and start the Management Console.
- b. In the management console (tsmw2k), right-click **Tivoli Storage Manager** in the left pane (Tree tab). Click **Add TSM Computer** to open the Tivoli Storage Manager Computer Resource Properties window.

Complete the fields with all the information you need to connect to the Tivoli Storage Manager server of your DR550. See Figure 5-59 for an example.

**TSM Computer Resource Properties**

Computer Information

Machine Name:  ...

Operating System:

☐ Access remote machine using current Windows account information

☒ Include Report Features

☒ TSM Web Client

☒ TSM Web Administrator

HTTP Port:

| Server Name | TCP/IP Address | TCP  | HTTP |
|-------------|----------------|------|------|
| TSM         | 192.168.1.22   | 1500 | 1580 |

Buttons: Edit, Delete, Account..., Add

Bottom buttons: OK, Cancel, Apply

Figure 5-59 Tivoli Storage Manager Computer Resource Properties

Click **Add**. In Figure 5-60 on page 226 (Logon Information), use the IBM Tivoli Storage Manager administration account admin of your DR550 to log on and click **OK** to close this window again. The server is added and you can click **OK** to close the Tivoli Storage Manager Computer Resource Properties window.

- c. In the management console (tsmw2k), right-click **Tivoli Storage Manager** in the left tree (Tree tab). Click **TSM Operational Reporting** to open the Properties window.

At the Operational Reports tab, Operational Monitors tab, Report Service tab, and Summary Information tab, there is no immediate configuration needed.

At the E-mail Account tab, fill the fields with your e-mail parameters and try the communication with your e-mail server by using the **Test** button.

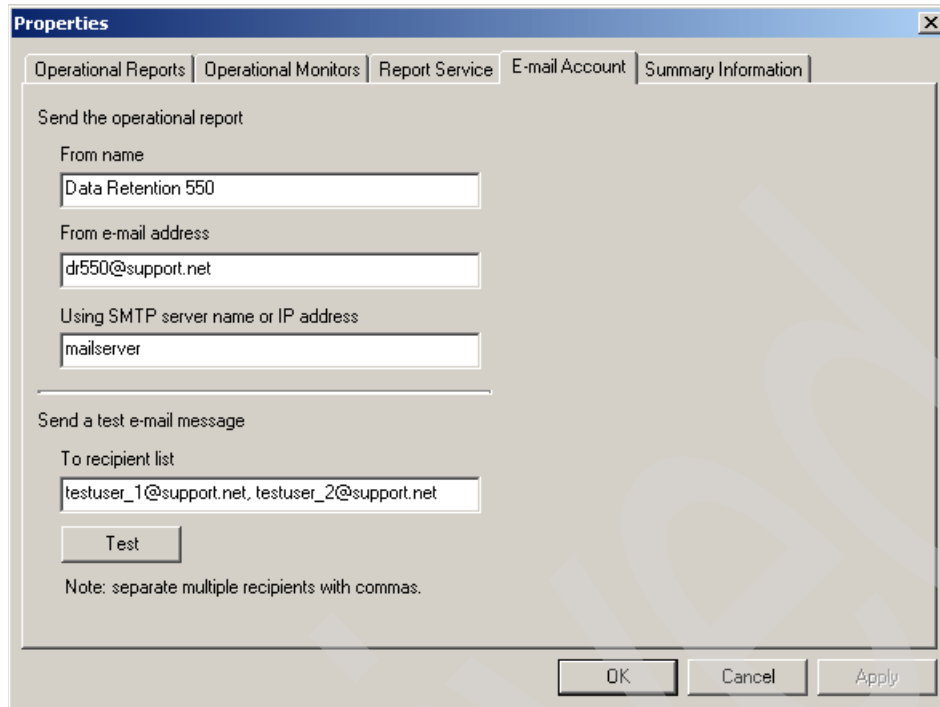


Figure 5-60 Operational Reporting Properties window: E-mail Account tab

- d. Back in the management console (tsmw2k) window, in the left pane, expand the tree for the above configured IBM Tivoli Storage Manager server, for example, DRS\_ENGINE (UNIX). Expand the tree named TSM and the tree named Reports.

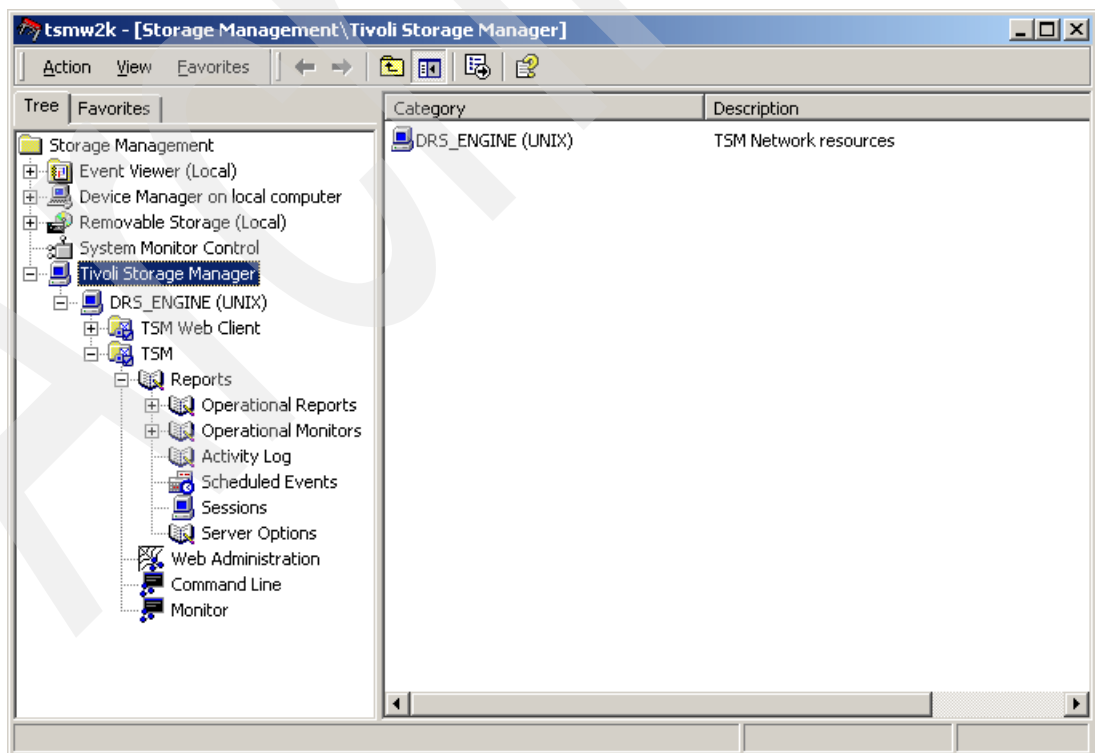


Figure 5-61 IBM Tivoli Storage Manager management console



- e. In the expanded tree structure, right-click **Operational Reports**, and click **New**. In the next window, click **OK** to create an operational report named Daily Report.

At the Report Details tab, customize the settings so they fit your demands. For example, disable the Client Schedules Status, because the DR550 will not operate any client schedules. See Figure 5-62 for an example.

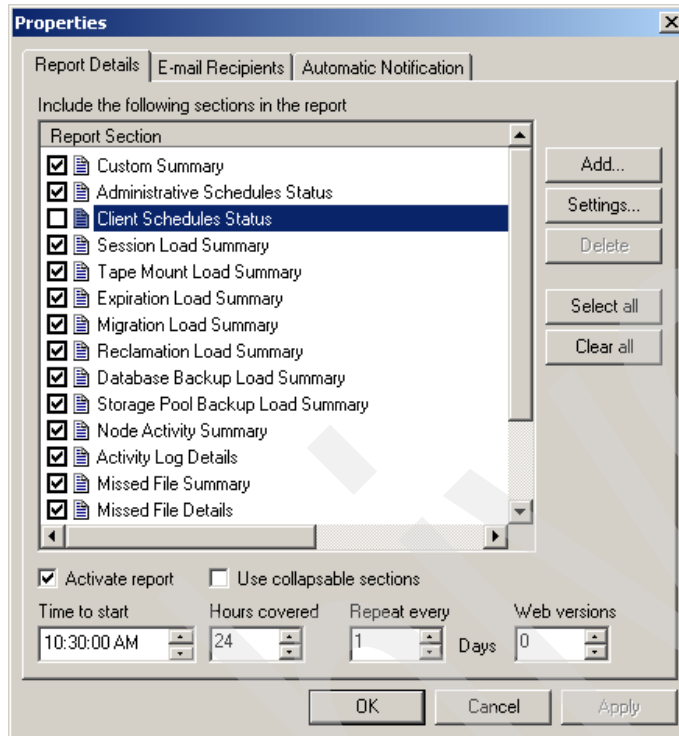


Figure 5-62 Operational Reports: Report Details

At the E-mail Recipients tab, complete the Recipient name field, the E-mail address field, and switch the format to HTML if your e-mail client supports that format. Click **Add** to configure the recipient. Repeat the steps for other e-mail recipients.

At the Automatic Notification tab, do not do any configuration and click **OK**.

- f. Within the expanded trees, right-click **Operational Monitors**, and click **New**. In the next window, click **OK** to create an operational report titled Hourly Report.

At the Monitor Details tab, customize the settings so that they fit your demands.

At the E-mail Recipients tab, complete the Recipient name field, the E-mail address field, and switch the format to HTML if your e-mail client supports that format. Click **Add** to configure the recipient. Repeat the steps for other e-mail recipients.

At the Net Send Recipients tab, do not do any configuration and click **OK**.

Archived



## IBM DR550 File System Gateway

This chapter provides an overview of the DR550 File System Gateway (DR550 Feature Code 4101, MTM 2229-FSG) architecture, describes its functionality, discusses concepts of deletion protection policies, and describes some basic administration tasks, including preparing the gateway for client access and recovery and replacement procedures.

## 6.1 File System Gateway overview

The main function of the DR550 File System Gateway (2229-FSG) server allows client applications to store and retrieve data by means of the Common Internet File System (CIFS) or Network File System (NFS) protocols.

On the front end, the FSG provides NFS and CIFS interfaces to client applications. At the back end, the FSG interacts with the IBM System Storage DR550 through the System Storage Archive Manager (SSAM) API. In other words, the DR550 FSG takes files that are sent using network file protocols (CIFS or NFS) and associates these files with an IBM System Storage Archive Manager management class (note that the management class must use event based retention). Different files or set of files can be associated with different management classes. This is done at the DR550 FSG by configuring profiles that define path and file naming patterns and associating them with corresponding management classes. The files maintained by the DR550 FSG can be protected against deletion by setting specific parameters at the DR550 FSG. In addition, the DR550 FSG can assign a minimum retention period to files. After this period, files are maintained until deleted by explicit action by the content management application. Deletion of files at the DR550 FSG triggers an activation of the event based policy that controls the corresponding objects in the SSAM Server.

The files received at the DR550 FSG are stored locally (cached) in addition to being forwarded to SSAM Server for archiving. The DR550 FSG maintains the file system directory tree locally, keeping pointers to the archived copy of the files in the SSAM Server. When a client application requests a file, the file is sent directly by the DR550 FSG if it is still in the cache; otherwise, the DR550 FSG requests the file from the SSAM Server and forwards it to the client application.

## 6.2 DR550 FSG architecture overview

In this section, we review the physical and logical DR550 FSG architecture.

### 6.2.1 Physical architecture

We have already discussed the physical components (hardware and software) that make up the DR550 FSG in 2.3, “DR550 File System Gateway overview” on page 25.

The DR550 FSG is connected to and communicates with the SSAM Server engines through IP. The data movement is accomplished by the DR550 FSG using SSAM API to communicate with the SSAM Server.

**Note:** With a DR550 DR2 that contains an internal Ethernet network, the IP communication between the FSG and the DR550 engines takes place over that internal network.

The DR550 FSG is available as either a stand-alone configuration or a high availability cluster configuration. Figure 6-1 on page 231 and Figure 6-2 on page 231 give a logical view of the data flow in DR550 and the File System Gateway solution.

### Stand-alone configuration

The DR550 FSG stand-alone configuration consists of a single server connected to a single-node DR550 Model DR1 and DR2. There is no failover protection provided. If the DR550 File System Gateway fails, client applications are unable to access the DR550 SSAM Server until recovery is complete.

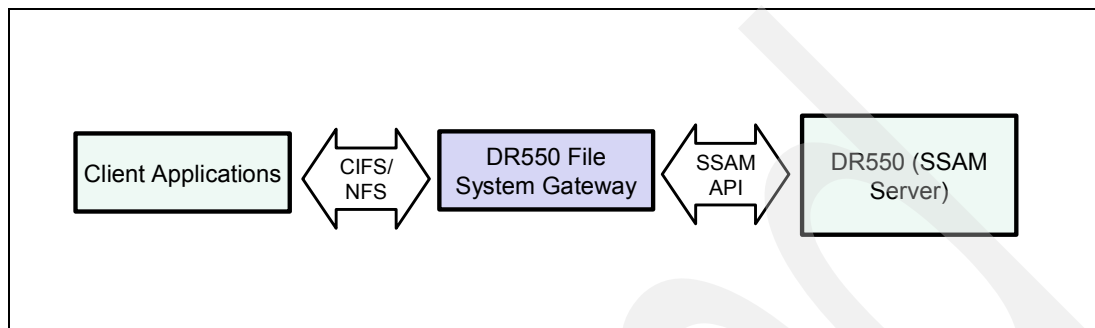


Figure 6-1 DR550 File System Gateway stand-alone configuration

### High availability cluster

The DR550 File System Gateway can be configured to operate in a high availability cluster. In this configuration, a Main DR550 File System Gateway server and a Supplementary DR550 File System Gateway server together form a cluster providing failover functionality, as shown in Figure 6-2.

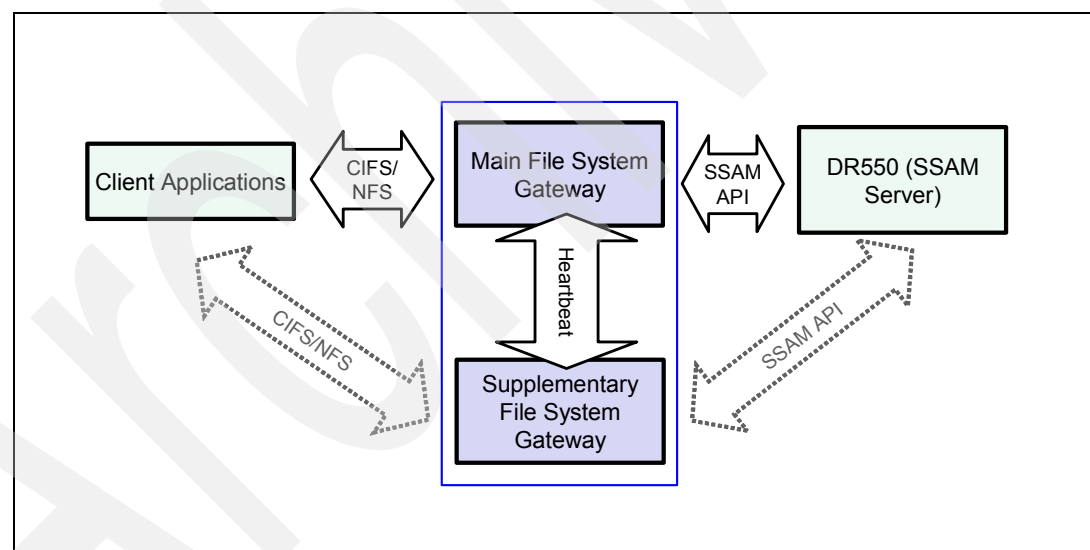


Figure 6-2 DR550 File System Gateway high availability cluster configuration

By default, the Main DR550 File System Gateway acts as the “active” gateway, while the Supplementary DR550 File System Gateway acts as the “standby” gateway. The condition of the two nodes is monitored by a heartbeat service.

Any time the heartbeat service detects a failure in the Main “active” DR550 File System Gateway, it triggers an automatic failover to the Supplementary “standby” DR550 File System Gateway, whereby the Supplementary now becomes the active server. As a result, there is minimal loss of access to the DR550 File System Gateway and there is no loss of data from the DR550.

Once the Supplementary has become active due to a failover, even when the disabling condition that caused the failover is no longer current, and the Main DR550 FSG is back online, the supplementary continues to be the “active” cluster node until a failback to the default configuration is done manually.

## Replication group

Another defining feature of the DR550 FSG architecture is the replication group, especially with regard to the high availability configuration. The DR550 FSG is configured with a replication group that has only one member: the primary File System Gateway (DRG) service (the DRG service is the main software module of the File System Gateway). If the DR550 File System Gateway is configured in a high availability cluster, the replication group on the supplementary DR550 File System Gateway is configured with a secondary DRG service. The Primary DRG service in the replication group provides read and write access to the DR550 SSAM Server for a set of clients. As files are saved to the primary DRG service, and file pointers and file system references are replicated to the secondary DRG service in the Supplementary DR550 File System Gateway. This secondary DRG service provides a “mirror” of the file system. In the event that the primary DRG service or the DR550 File System Gateway server hosting it fails, clients can continue to access files through the “mirrored” file system on the Supplementary DR550 File System Gateway. Access to the DR550 SSAM Server is interrupted until failover is completed, and the redundancy of file system information in the DR550 SSAM Server is reduced while repairs are in progress.

## Conditions for automatic failover

The Main DR550 FSG checks for the following failures that will result in a failover from the “active” DR550 FSG to the “standby” DR550 FSG:

- ▶ System or Cache disk failure on the “active” DR550 FSG
- ▶ Loss of network connectivity from the “active” DR550 FSG
- ▶ Catastrophic CPU failure on the “active” DR550 FSG
- ▶ Software crash

## 6.2.2 Logical architecture

From a logical architecture standpoint, the DR550 FSG is made up of the following layered elements, from top to bottom:

### ▶ Services

A service is a software module providing a set of capabilities to the DR550 File System Gateway system. Each service consists of components with attributes that deliver a particular capability. Services are identified by a three letter acronym, such as the DRG (DR550 File System Gateway) service. The DR550 File System Gateway (DRG) service is the main software module of the File System Gateway. It streams data to the DR550 SSAM Server and caches a copy of the data to speed its retrieval. The DRG Service and its components are discussed in more detail in “DRG service components” on page 233.

### ▶ Components

Services have components associated with them. A component within a service delivers a particular capability. Each component has a set of configurable attributes that can be monitored automatically. Components are the first layer under services. Components have names defined by the software.

### ▶ Attributes

The most granular level of the system is an attribute. This is a single value or property of a component in a service. Attributes are the elements that are monitored for the status and state of the DR550 File System Gateway.

Every service and component of the DR550 File System Gateway contains a set of attributes associated with the capabilities that the service or component provides.

Each attribute has its own XML file that contains current configuration or status information for that attribute.

Evaluating the data of each XML file assists in determining the current status of the DR550 FSG.

### **Attribute XML files**

Attribute XML files are composed of a valid subset of XML. Within the file, information is structured as *containers* that might contain other containers or *atoms*. Each atom has a value of a specified data type. The values of the atom define the attribute and its values.

All attributes are used for reporting purposes only. Updating Attribute XML files has no impact on the behavior of the DR550 FSG.

Attribute XML file have two distinct parts:

- The XML wrapper that identifies the content as XML and defines the content of the XML to be a container. An example is:

```
<?xml version="1.0" encoding="UTF-8"?>
<containerxml version="1" xmlns="http://www.ibm.com/
schemas/XML-container-1.0.0">
<container name="FCST">
...
</containerxml>
```

- One or more attribute containers, each of which include atoms that define the attribute. An example is:

```
<container name="FCST">
<atom name="AVER" type="UI32" value="2"/>
<atom name="AVTP" type="FC32" value="ENUM"/>
<atom name="APER" type="FC32" value="READ"/>
<atom name="ATIM" type="UI64" value="1170896136591928"/>
<container name="AVAL">
<atom name="0x00000000" type="UI32" value="2"/>
<atom name="0x00000001" type="ENUM" value="4"/>
</container>
```

- Each atom is only included once per container. For information about the atom elements that can be contained within an attribute XML file, see “Atom Elements” in the *IBM System Storage DR550 File System Gateway Software Version 1.1.1 User Guide*, GC27-2125.

### **DRG service components**

The DRG service is the main software module of the DR550 File System Gateway. It streams data to the SSAM Server and caches a copy of the data to speed its retrieval.

- The DRG service supports the following components: (for a description of attributes relevant to each of the components, refer to the *IBM System Storage DR550 File System Gateway Software Version 1.1.1 User Guide*, GC27-2125).

## ***Storage***

The Storage component provides information regarding the storage space used by the DRG service, and the objects cached and processed.

The attributes for this component describe:

- ▶ Cache configuration (total space, space available, and so on)
- ▶ File system operations (number of files and directories created, number of reads and writes, and rate of reads and writes)
- ▶ Storage to SSAM Server (number of files stored, number of files retrieved, removal of files, and so on)

## ***Replication***

The Replication component reports the status of the high availability cluster.

- ▶ If a DR550 File System Gateway is not part of a high availability cluster, its cluster status is reported as “N/A” (not applicable).
- ▶ The cluster status is “Normal” if both DR550 File System Gateways are healthy: the Main Primary is “active,” and the Supplementary Primary is in “standby.”
- ▶ If only one DR550 File System Gateway server in the cluster is available and “active,” the cluster status is “Vulnerable”.
- ▶ If the cluster is not providing service, but the clustering mechanism is expected to automatically restore service, the status is “Transitional”.
- ▶ If no DR550 File System Gateway in a cluster can provide an active service and the failure cannot be recovered from automatically, the cluster status is “Failed”.

## ***Backup***

The Backup component provides information about backups of the managed file system. The replication group session file (a description of changes to the file system) is backed up to the SSAM Server as soon as it is finalized. By default these backups occur every thirty minutes. Should the shared file system become corrupted, a backup stored in the SSAM Server can be used to restore the system.

## ***Client Services***

The Client Services component provides information about the support services used to manage the file system share (NFS or CIFS), client integrations, or the high availability cluster (heartbeat). If the relevant support service is not running, you do not have access to the SSAM Server managed file system, client integration, or high availability clustering services.

## ***Events***

The Events component reports on the audit messages and attribute values generated by the DR550 File System Gateway, and their progress to relay services that forward these messages or values to their final destination.

## ***Resources***

The Resources component provides information about the computational, storage, and network resources available, as well as low level information about the operation of the DR550 File System Gateway.



### ***Timing***

This component indicates the current condition of the DR550 File System Gateway time (the time kept by the DR550 File System Gateway software with respect to the hardware clock). It is either synchronized or synchronizing. It also has another attribute that gives the offset between the DR550 File System Gateway time and the operating system clock.

## **6.2.3 Deletion protection**

The DR550 File System Gateway offers retention protection by using three flags that control the deletion of files and the option to set a retention period:

- ▶ Write Once Read Many (WORM)
- ▶ Content Handle Release Inhibit (CHRI)
- ▶ No-Delete

These flags can be set alone or in combination with each other to enable different levels of deletion protection.

**Important:** The deletion protection flags (WORM, CHRI, and No-Delete) have a systemwide scope and the settings are effective on all files and directories in the FSG.

### **Write Once Read Many (WORM)**

This is a strong deletion protection policy that protects files from deletion or modifications. For files enabled with WORM protection, after they have been committed to the file system, the DR550 File System Gateway will not allow any alterations or modifications, such as moving them between directories, renaming them, appending to them, modifying them, replacing their content, or deleting them. Attempts to alter or delete a file receive a “permission denied” response. The WORM flag can be used in combination with a Data Protection period that allows you to perform any of the above operations after a predefined period of time. Empty directories might be removed after their protection period has expired. They are, however, not automatically removed by the FSG.

**Note:** Do not use Windows Explorer to create folders on a WORM enabled share. The folder gets created with a name of “New Folder” and cannot be renamed. If you end up creating such folders, then the only option is to delete them and the procedure for this is described in “New Folder problem for CIFS shares” on page 244

The following directory operations are permitted even with WORM protection active (the DRG service must be running):

- ▶ Create directory
- ▶ Change ownership
- ▶ Modify permissions
- ▶ Change Timestamps
- ▶ Read directory

After a directory's protection period has elapsed, the following file and directory operations are permitted:

- ▶ Remove directory (if empty)
- ▶ Create file
- ▶ Change ownership of file
- ▶ Change Permissions on a file
- ▶ Change Timestamps on a file

- Read file

**Note:** A newly created file, until it is committed to the SSAM Server, is not protected from deletion or modification. There is usually a ten second interval between the moment the DR550 FSG detects that the file is no longer in use and the time it becomes protected.

### **Content Handle Release Inhibit**

When the Content Handle Release Inhibit (CHRI) flag is enabled, the files, even after they are committed to the DR550 FSG, can be altered or deleted by the application. However, the change only affects the DR550 FSG view of the data; prior versions and deleted content remain undeleted in the SSAM Server. They cannot be retrieved by the application using the DR550 FSG. Special tools and a services engagement will be required for retrieving these items.

The CHRI protection might be a more appropriate choice than the more stringent WORM for applications that require the ability to delete or modify temporary files as they write content to the archive.

### **No-delete**

The no-delete flag is the most stringent form of protection and, used in combination with the WORM flag, can impose global protection on all files. This will prevent any file from deletion from the SSAM Server, regardless of whether the data protection period for that file has expired. Extraordinary circumstances, such as fraud investigation or litigation, are good examples of where this flag might be used.

**Tip:** Any of the three flags described above may be used, individually or in combination, to implement a customized deletion protection policy. When all three flags are enabled, the no-delete flag has predominance and nothing can be deleted until this flag is disabled.

### ***NFS and CIFS in WORM Mode***

There are subtle differences between NFS or CIFS objects in how the deletion protection is handled and the sequence of events and are tied to the NFS and CIFS protocols. The key difference in NFS and CIFS file handling is that a CIFS client informs the server that it has closed the file as soon as it has finished saving it, whereas an NFS client does not do so.

This means that the DR550 FSG is aware when a CIFS client is done with a file and a short time (about ten seconds) after that awareness occurs, it begins the process of ingesting and forbids deletion or alteration after that. In the case of NFS, the DR550 FSG waits a short time after the last block of the file has been saved, and then forbids further alterations and deletions as it starts to ingest the file.

The wait period mentioned above, that is, the time between when a file has been written and the time when the DR550 FSG begins ingesting the file, is configurable (it is set to a default of ten seconds).

### ***Enabling deletion protection***

Deletion protection is effective for all new files stored in the FSG after you enable or disable specific parameters on the DR550 FSG. This is done by means of editing a configuration profile file (FOPT) using the ADE Console Interface. This procedure is described in 6.3.2, “Enabling deletion protection” on page 239.

Setting the retention period is done by editing a DRG profile file (FPRF) using the ADE Console Interface. This procedure is described in 6.3.3, “Editing DRG profiles” on page 241.

## 6.2.4 DRG profiles

A DRG profile consists of a list of parameters that define the behavior of the DR550 File System Gateway with respect to the affected directories. The directory or directories where these parameters apply are defined by means of a regular expression.

A DR550 File System Gateway has a “default” profile that displays first in the list of profiles and applies to all directories not otherwise specified.

When more than one profile is defined for a DR550 File System Gateway, the behavior of files in that directory is governed by the last match made to the directory name. With the caching priority and the preloading parameter, you can change the performance of these directories.

A DRG Profile permits you to customize the behavior of a DR550 File System Gateway on a directory-by-directory basis to meet any particular needs for data access and data protection. You can configure the items shown in the following sections (refer to 6.3.3, “Editing DRG profiles” on page 241 and to Table 6-1 on page 242 for instructions on how to set the different parameters).

### File protection period

File protection periods enforce WORM-level file deletion protection for a defined period of time. During its file protection period, a file cannot be modified or deleted. After the file protection period for a file expires, clients are permitted to delete it through the DR550 File System Gateway. File protection periods can be useful in enforcing data retention periods that are required by policy or regulation, while permitting storage space to be freed for other uses as the protection period expires for each piece of content.

**Attention:** The default protection period is forever, which means files that are archived to the DR550 FSG will never expire. Therefore, it is important to define the protection period for specific files before you archive data to the FSG.

### Parallel loading

Parallel loading is the ability to enable a DR550 File System Gateway to preload the remaining items in a directory when the first item is requested. This improves access times in many cases, because, statistically, when one file is accessed in a directory, there is a great probability other files in that directory might also be accessed.

### Secondary preloading

In a high availability cluster, a Supplementary DR550 File System Gateway can provide access to files stored in the SSAM Server. These features speed access to new data (Secondary Preload) and to older data (Secondary Preload Access) through the Supplementary DR550 File System Gateway.

### Caching priority

When a file is saved to the SSAM Server, the File System Gateway caches a copy of the file. This copy is used to speed subsequent retrievals of that file. Cached copies are treated as transient content, and the least-recently-accessed copies may be purged to free up space on the DR550 File System Gateway as needed.

**Tip:** The DRG profile is mirrored automatically to the standby DR550 FSG in a high availability environment. You must configure the DRG profile only at the active DR550 FSG.

### **Recommendations**

We recommend that you:

- ▶ Do not delete or change the expiration settings of the default profile to ensure that the root directories of your CIFS or NFS shares will not expire or will become erasable.
- ▶ Define different management classes in SSAM for each protection period you want to define. If you have three different protection periods, you should define three different SSAM management classes and copy groups. The RETMIN value of the copy group should cover the time frame as the protection period in the DRG profile.
- ▶ Create a DRG profile for /New Folder if you have CIFS shares. See 6.3.3, “Editing DRG profiles” on page 241 for examples of how this can be accomplished.

## **6.3 File System Gateway administration and operations**

In this section, we discuss the DR550 FSG general administration and operations tasks.

### **6.3.1 Integrating the DR550 FSG with SSAM Server**

Once the network setup is done and the DR550 FSG has reliable IP connectivity with the SSAM Server, you have to update your SSAM Server environment to integrate the DR550 FSG.

We have already explained what to do in 3.7.1, “Prepare DR550 SSAM for FSG attachment” on page 113. In this case the configuration was done by using the SSAM administration command line (`dsmadm`) at the DR550 SSAM Server. Alternatively, you could invoke the `dsmadm` from the DR550 FSG. However, before you can run the `dsmadm` command, you must first log in as root on the File System Gateway, and enter the following two commands:

```
ln -s /opt/tivoli/tsm/client/api/bin/dsm.sys /opt/tivoli/tsm/client/ba/bin/dsm.sys
ln -s /opt/tivoli/tsm/client/api/bin/dsm.opt /opt/tivoli/tsm/client/ba/bin/dsm.opt
```

The result of running these two commands is that `dsmadm` will use the same configuration file as the API client on the DR550 FSG.

**Note:** To log in as root on the DR550 FSG, do the following:

- ▶ Log in as `fsgadmin`.
- ▶ Switch to root by running `su -`.

The following setup in the SSAM Server needs to be done (This is just an example.):

1. Log in to either the DR550 SSAM Server or DR550 FSG and obtain root privileges.

2. Run the **dsmadm** command:

```
dsmadm
IBM Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 3, Level 0.0
(c) Copyright by IBM Corporation and other(s) 1990, 2004. All Rights Reserved.
```

Enter your user id: **admin**

Enter your password:

Session established with server TSM: AIX-RS/6000

Server Version 5, Release 5, Level 0.0

Server date/time: 03/06/08 10:58:38 Last access: 03/06/08 09:43:53

tsm: TSM>

3. Run the following commands from the TSM prompt:

```
tsm: TSM>define domain DRG-DOMAIN
tsm: TSM>define policyset DRG-DOMAIN STANDARD
tsm: TSM>define mgmtclass DRG-DOMAIN STANDARD DRG-DEFAULTMC
tsm: TSM>define copygroup DRG-DOMAIN STANDARD DRG-DEFAULTMC
type=archive dest=ARCHIVEPOOL retinit=event retmin=1 retver=0
tsm: TSM>define mgmtclass DRG-DOMAIN STANDARD DRG-SYSMC
tsm: TSM>define copygroup DRG-DOMAIN STANDARD DRG-SYSMC type=archive
dest=ARCHIVEPOOL retinit=event retmin=0 retver=0
tsm: TSM>assign defmgmtclass DRG-DOMAIN STANDARD DRG-DEFAULTMC
tsm: TSM>activate policyset DRG-DOMAIN STANDARD
tsm: TSM>register node DRG-NODE abcd2357 passexp=0 domain=DRG-DOMAIN
```

### 6.3.2 Enabling deletion protection

We defined the concepts of deletion protection in 6.2.3, “Deletion protection” on page 235.

Deletion protection is enabled and disabled by editing the FOPT configuration file. Editing must be done through the ADE console interface.

**Important:** The ADE Console should be used with caution. It permits low-level operations and can possibly interrupt operations and corrupt data, if used improperly.

Proceed as follows:

1. Log on to the DR550 FSG as fsgadm.
2. At the command line, telnet to the ADE console:  

```
telnet localhost 1413
```
3. Start the Bundle Editor module of the console:  

```
cd /proc/BDED
```
4. Open the FOPT configuration file with this command:  

```
edit FGRP/10/FOPT
```

  - To enable WORM Mode:
    - i. Enter:  

```
option -e #1
```

- ii. Save the changes and close the file with this command:

submit

WORM Mode is now enabled and files can no longer be deleted or altered through the DRG file system on the DR550 File System Gateway.

See Example 6-1 for an illustration.

*Example 6-1 Enabling WORM Mode*

---

```
ade nodeid: /proc/BDED > edit FGRP/10/FOPT
```

```
Opening bundle 'FGRP/10/FOPT'...
```

```
Bundle 'FGRP/10/FOPT' opened
```

#	FONM	FOVL
1	WORM Mode	Disabled
2	Content Handle Release Inhibit	Disabled
3	No-Delete Flag	Disabled

```
ade nodeid: /proc/BDED > option -e #1
```

#	FONM	FOVL
1	WORM Mode	Enabled
2	Content Handle Release Inhibit	Disabled
3	No-Delete Flag	Disabled

```
ade nodeid: /proc/BDED > submit
```

```
Submitting bundle 'FGRP/10/FOPT'...
```

```
ade nodeid: /proc/BDED > Bundle 'FGRP/10/FOPT' submitted
```

```
ade nodeid: /proc/BDED >
```

---

– To enable CHRI Mode:

- i. Open the bundle for editing:

```
edit FGRP/10/FOPT
```

- ii. Enable CHRI Mode:

```
option -e #2
```

- iii. Save the changes and close the bundle editor:

```
submit
```

The CHRI flag is now enabled. Files can now be modified or deleted on the DRG file system of the DR550 File System Gateway, but they are not deleted from the SSAM Server.

– To disable CHRI:

- i. Open the bundle for editing:

```
edit FGRP/10/FOPT
```

- ii. Disable CHIR Mode:

```
option -d #2
```

- iii. Save the changes and close the bundle editor:

```
submit
```

The CHRI flag is now disabled.

- To enable Global Protection of data, enable both the WORM and the no-delete flags as follows:

- i. Open the bundle for editing with this command:

```
edit FGRP/10/FOPT
```

- ii. Enable the WORM flag:

```
option -e #1
```

- iii. Enable the “no-delete” flag:

```
option -e #3
```

- iv. Save the changes and close the bundle editor:

```
submit
```

Global protection is now in effect for all files.

- 5. Type **exit** to log out of the ADE console or proceed with editing the DRG profiles if required (see 6.3.3, “Editing DRG profiles” on page 241).

### 6.3.3 Editing DRG profiles

DRG profiles were defined in 6.2.4, “DRG profiles” on page 237. This section details how you can edit DRG profiles to set your required parameters:

- 1. Log on as fsgadm to the main DR550 FSG.

- 2. Enter the DRG profile editor:

```
telnet localhost 1413
```

- 3. Enter **cd /proc/BDED**.

- 4. Enter **edit /FGRP/10/FPRF**.

You will get a list off all defined DRG profiles. Line 1 is the default profile.

#### Example 6-2 Default DRG profile example

```
ade nodeid: /proc/BDED > edit /FGRP/10/FPRF
```

```
Opening bundle 'FGRP/10/FPRF'...
```

```
Bundle 'FGRP/10/FPRF' opened
```

```
FSG Profiles Bundle: FGRP/10/FPRF
```

#	Pattern	Ingest	Cache	Protection	Preload (D R S)	Mgmt Class
1			1	forever	DABL DABL DABL	DRG-DEFAULTMC

```
ade nodeid: /proc/BDED >
```

### DRG profile editor commands

We include here an overview of the DRG profile editor commands. Refer to the *IBM System Storage DR550 File System Gateway Software Version 1.1.1 Integration Guide*, GC27-2124 for a complete listing of the DRG profile editor commands and their options.

► **editprofile -a|#N [options]**

Table 6-1 shows the command's parameters.

Table 6-1 *editprofile* command parameters

Parameter	Description
N	The number identifying the profile # in the list to modify.
-a	Add a new profile with the characteristics granted by the options included in this command.
-o [+n -n N]	Change the order of this profile in the profiles list. The +n or -n flags move the profile up or down n positions in the list. The number flag N places the profile at that position in the list.
-p [regex]	Sets the directory or directories where this profile applies, where regex is a regular expression specifying a directory name. Note that /fsg (which is the directory under which all shares must be defined) must not be included in the path name.
-c [priority]	Sets the caching priority for the directory, where the value of priority must be 1, 2, 3, 4, or 5. Directories with a caching priority of 1 are swapped out before directories with a larger priority value.
-t time<unit>	Sets the file deletion protection period for files protected by WORM Mode (see "Enabling deletion protection" on page 236). If the value of time = forever, the file deletion protection period is disabled and WORM Mode never expires. If the value of time = none, the file is immediately available for modification. If a unit is not specified, seconds are used. Use the following abbreviations with time to specify a different unit. Do not leave a space between time and the unit. For example, indicate 1 year as 1y.  <unit> is one of the following: s Seconds (default), i Minutes, h Hours, d Days, w Weeks, m Months, or y Years
-d [ENBLIDABL]	Disable/Enable Parallel Loading (Preload D).
-r [ENBLIDABL]	Disable/Enable preloading of new files on the secondary. (Preload R).
-s [ENBLIDABL]	Disable/Enable preloading files to the secondary on each file access. (Preload A).
--mgmt-class [CLASS]	Assigns a management class for files ingested into the DR550. CLASS Tivoli Storage Manager or SSAM Management Class.

► **profiles**

Displays the entire list of profiles in the configuration bundle that you are editing.

► **revert**

Closes the current configuration bundle without saving changes. (The bundle name can optionally be included with the command.)

► **rmprofile #N**

Removes an existing profile. #N indicates which profile in the list to remove.



► **submit** <-r>

Save and close the bundle currently being edited. If the -r option is used, the bundle is saved without closing.

## Adding, modifying, and deleting DRG profiles

This section illustrates DRG profile editing. It is assumed that the appropriate management classes (DRG-one-year-class and DRG-one-month-class) have already been defined in the SSAM Server.

Examples of DRG profile editing:

► To add a new DRG profiles, enter:

```
editprofile -a -p "/cifs-shares/x-ray" -t 1y --mgmt-class DRG-one-year-class
editprofile -a -p "/cifs-shares/lab" -t 1m --mgmt-class DRG-one-month-class
```

**Important:** Note that /fsg (which is the directory under which all shares must be defined) must *not* be included in the path name.

This will configure a retention period of one year for all objects in the folder /cifs-shares/x-ray and a retention period of one month for all objects in the folder /cifs-shares/lab.

**Note:** Before specifying a new management class in the profile, it must be defined to SSAM Server.

After executing these commands, the DRG profile bundle appears as shown in Example 6-3.

*Example 6-3 DRG profiles example*

#	Pattern	Ingest	Cache	Protection	Preload (D R S)	Mgmt Class
1		1	forever		DABL DABL DABL	DRG-DEFAULTMC
2	/cifs-shares/x-ray	1	1.0 y		DABL DABL DABL	DRG-one-year-class
3	/cifs-shares/lab	1	1.0 m		DABL DABL DABL	DRG-one-month-class

**Important:** When changing the retention period for an existing directory pattern entry in the profile, the new retention period will be applied to data stored after the change only. Data stored before the change is not affected.

**Tip:** When more than one profile is defined for the DR550 FSG, the behavior of files in that directory is governed by “the last match” made to the directory name. This means you must check and if necessary change the order of your profiles to ensure that the correct retention policy will be assigned to the archived objects.

You can change the order of the profiles with the **editprofile -o +/-nN** command:

► To change an existing DRG profile, enter:

```
editprofile #3 -t 1y --mgmt-class DRG-one-year-class
```

This will change the retention period for all new objects in the /cifs-shares/lab directory to one year, as shown in Example 6-4.

*Example 6-4 Changing an existing profile*

#	Pattern	Ingest	Cache	Protection	Preload (D R S)	Mgmt Class
1			1	forever	DABL DABL DABL	DRG-DEFAULTMC
2	/cifs-shares/x-ray		1	1.0 y	DABL DABL DABL	DRG-one-year-class
3	/cifs-shares/lab		1	1.0 y	DABL DABL DABL	DRG-one-year-class

- To delete an existing DRG profile, enter:

```
rmprofile #2
```

This will delete the profile with the line number 2.

## New Folder problem for CIFS shares

You should not use the Windows Explorer to create new folders in DR550 FSG shared drives. Windows Explorer always creates a new directory as “New Folder”. Because of the deletion protection enforced by the DR550 FSG, the directory cannot be deleted or renamed. To avoid a lot of “New Folders” in your directory structure, we recommend that you define a separate DRG profile that will associate a retention period of only one second to the folder “New Folder”. The following procedure can be used for that purpose:

1. Enter the DRG profile editor:

```
telnet localhost 1413
cd /proc/BDED
edit /FGRP/10/FPRF
```

2. Add a profile:

```
editprofile -a -p "/New Folder/" -t 1s
```

The result is shown in Example 6-5.

*Example 6-5 DRG profile entry for New Folder*

#	Pattern	Ingest	Cache	Protection	Preload (D R S)	Mgmt Class
1			1	forever	DABL DABL DABL	DRG-DEFAULTMC
2	/cifs-shares/x-ray		1	3.0 dy	DABL DABL DABL	DRG-DEFAULTMC
3	/cifs-shares/lab/test		1	3.0 dy	DABL DABL DABL	DRG-DEFAULTMC
4	/New Folder/		1	1.0 s	DABL DABL DABL	DRG-DEFAULTMC

Each New Folder will expire after one second and can be deleted again. The directory pattern depends on the language setting of the client system. If you have a German language Windows, for example, then Windows will create a folder named “Neuer Ordner”. In this case, you should modify the line of the profile to your language requirements.

3. Save and activate changes.

To save and activate your DRG profile changes, enter the **submit** command.

4. Exit the DRG profile editor.

Enter the command **exit** to leave the DRG profile editor.

**Note:** You can only delete but not rename the “New Folder” that is created after the definition of the profile. This works for each subfolder.

**Tip:** If the value of time = forever, the file deletion protection period is disabled and WORM MODE never expires. If the value of time = none, the file is immediately available for modification.

## 6.4 Integrating client applications with DR550 FSG

The primary function of the DR550 FSG is to provide shares for both NFS and CIFS clients and enable deletion protection on the file systems that the clients access. In a high availability cluster, share properties must be set on both the Main and the Supplementary DR550 FSG and both must be identical. However, the share directories are created only in the Main DR550 FSG. This section describes the procedure to set up directories that can be exported for NFS mounts by client applications or attached by Windows clients as network drives using the CIFS protocol.

**Important:** Before using the DR550 File System Gateway, make sure that your application works when using file shares provided by the DR550 FSG. This is because the DR550 FSG modifies the file share behavior by disallowing deletions or modifications to any file stored on the DR550 FSG file share.

### 6.4.1 Create local users and groups on the DR550 FSG

Before we set up the properties of directories to be exported or shared, we need to create the directories and assign the right permissions. We also need to create user IDs at the DR550 FSG for NFS or CIFS clients to authenticate and map the network drives. (Note that instead of local user management, a more efficient approach is to use a centralized directory service, such as one based on LDAP or Microsoft Active Directory. This is discussed in Chapter 7, “Centralized user management and FSG scenarios” on page 275.)

**Important:** If you work with a dual-node cluster DR550 FSG and local user management, you need to be sure that you have exactly the same user and groups on both nodes, including uid, gid, and password. The DR550 FSG will *not* automatically replicate that information.

To start the User and Group Administration in SLES10, log on as root and run **yast2 users &**

**Tip:** When you create users and groups for NFS, they should match the users and groups set up on the remote machine, where the NFS share is imported. Otherwise, you must use some squashing and allow nobody or guest to access the share, which is not ideal.

From the User and Group Administration menu shown in Figure 6-3, you can administer local users and groups. To define a new user, make sure the **Users** button is selected and click the **Add** button. The New Local User window is displayed, as shown in Figure 6-4.

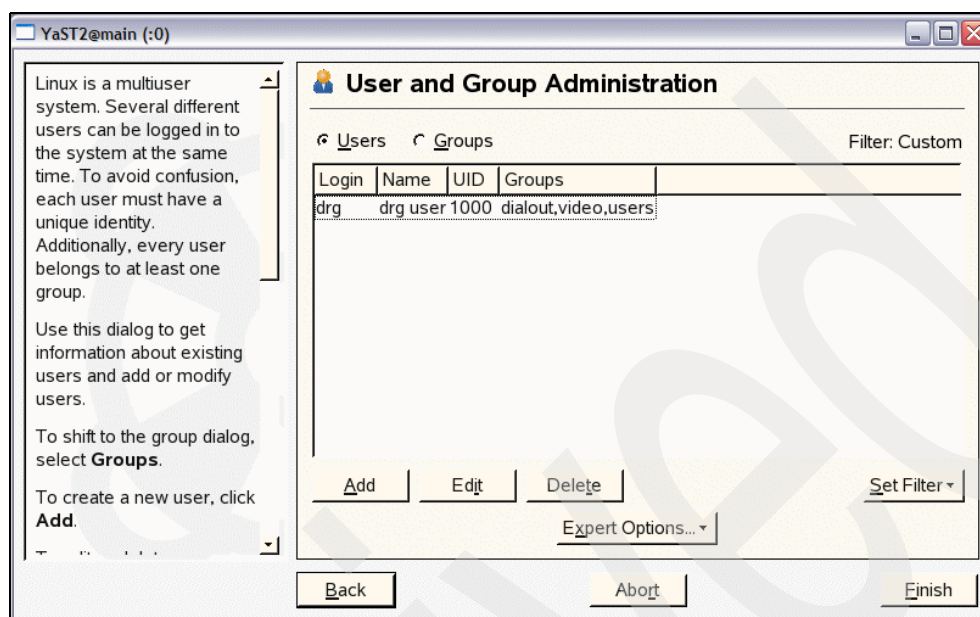


Figure 6-3 User and Group Administration window

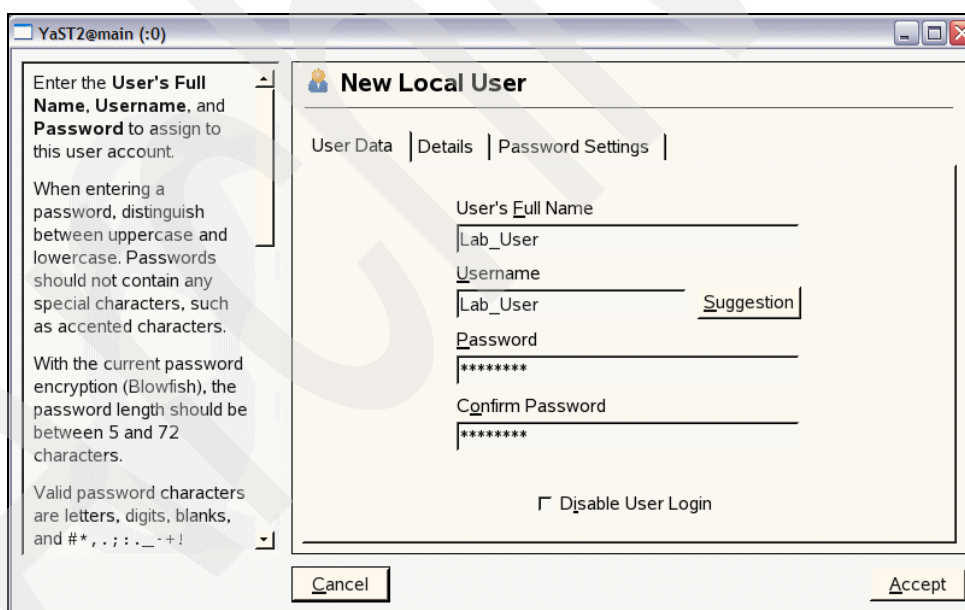


Figure 6-4 Add new local user

Type in a new user name (say lab\_user) and a password. Click on the **Details** tab and type in the required information for the user, as shown in Figure 6-5 on page 247.

You get a menu with default values for the user, including the ID, home directory, login shell, and group. You can change any of these or accept the defaults. Click **Accept**.

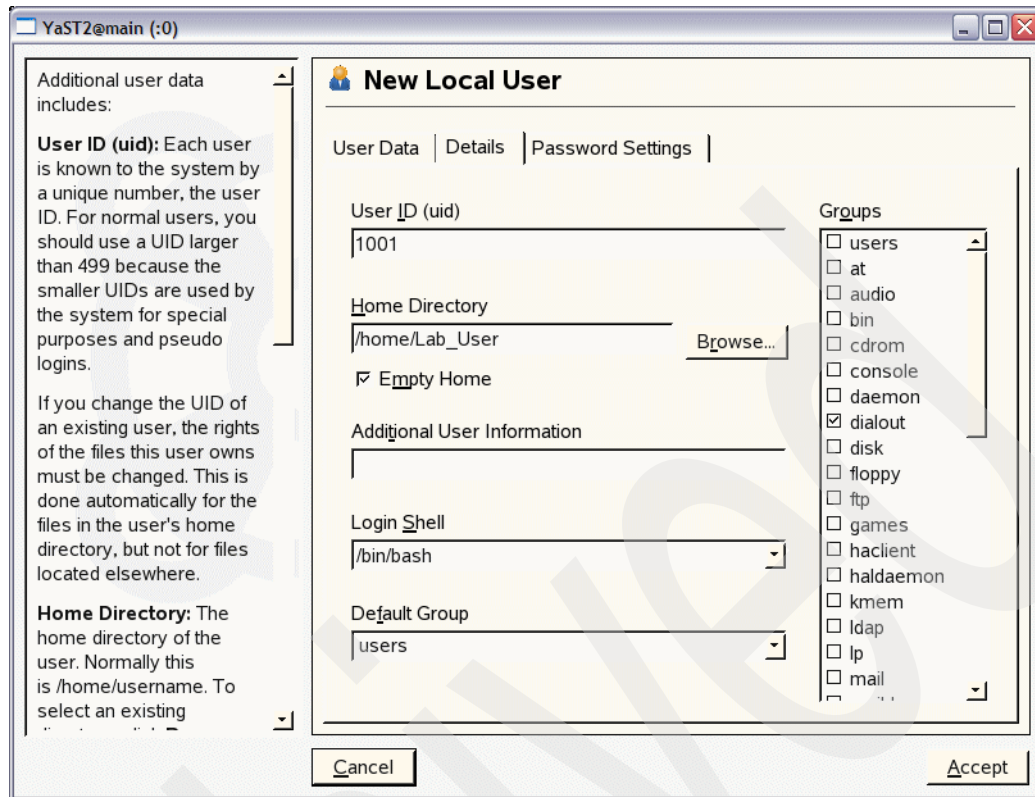


Figure 6-5 User details

Back on the Users and Groups administration window, you can start defining and configuring groups. Each user is assigned one or more groups and each group can have one or more users as members. It is a many to many relationship. To define groups, select **Groups** and you obtain the list of default (preconfigured) groups, as shown in Figure 6-6.

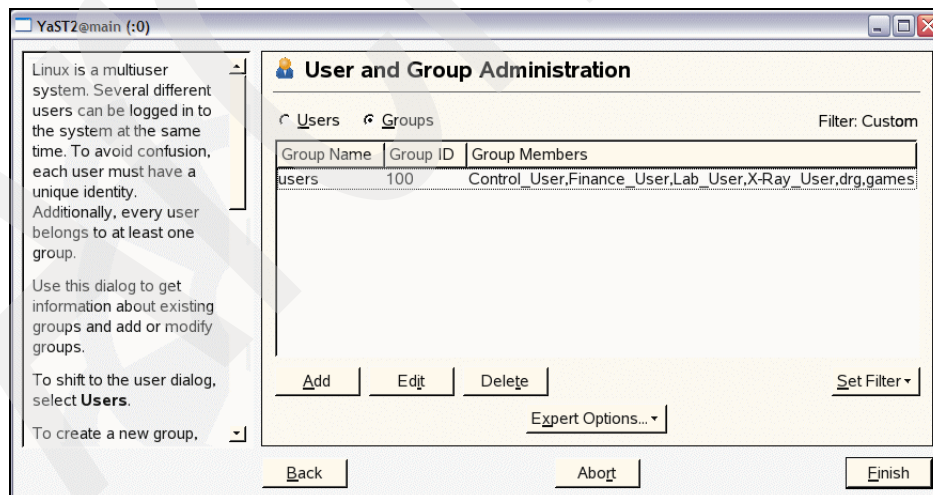


Figure 6-6 Default groups on DR550 FSG

You can now click the **Add** button. The New Local Group window is displayed, as shown in Figure 6-7.

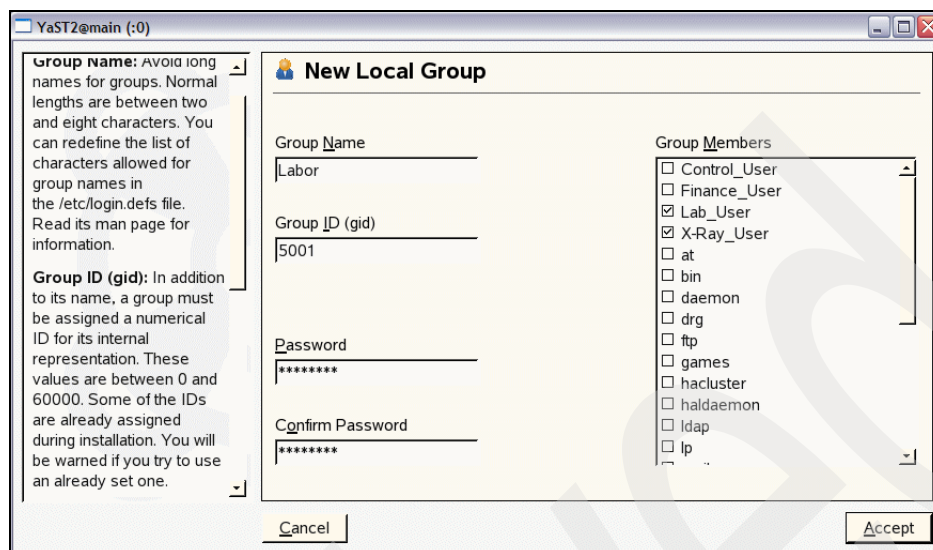


Figure 6-7 Adding a group with the name Labor

After you have entered the required information, including the group name and password, click the **Accept** button. Once you are done adding users and groups, click the **Finish** tab, and you are now ready to set up shared directories for usage by CIFS or NFS clients. The CIFS or NFS clients will use the user IDs and passwords you just defined to authenticate to the DR550 FSG and get access to the CIFS or NFS shares.

## 6.4.2 Create and configure shares for NFS exports

The first step is to create the directories (or file tree structures) locally at the DR550 FSG and then configure them for export.

**Note:** In a high availability cluster, shared properties must be set on both the Main and the Supplementary DR550 FSG and both configurations must be identical. However, the creation of shared directories is only performed at the active (usually Main) DR550 FSG.

### Creating directories for NFS exports

The directories to be exported must be created and be assigned the appropriate permissions for access by clients.

**Important:** The directories must be created under the /fsg directory, and the drg service must be active. In an HA configuration, directories must only be created at the active DR550 FSG.

Example 6-6 on page 249 shows the sequence of commands to create a new directory tree nfs-shares/finance. You must be logged in as fsgadm on the DR550 FSG.

*Example 6-6 Create directory tree /nfs-shares/finance*

```
cd /fsg
ls nfs-shares (check that the directory does not already exist)
mkdir nfs-shares
cd nfs-shares
mkdir finance
```

Assign the required permissions, as illustrated in Example 6-7.

*Example 6-7 Assign permissions*

```
chmod 755 /fsg/nfs-shares/finance <- sets permissions to read,write,execute (7)
 to owner (root), read, execute (5) to group,
 and read, execute (5) to user.

chgrp finance /fsg/nfs-shares/finance <- sets the group to finance
 (finance must exist as a group)

chown finance_user /fsg/nfs-shares/finance <- this sets the user to finance_user
```

**Configure Shares for NFS exports**

To set NFS share properties, follow this procedure:

1. Enter **yast2 nfs-server**. This opens the NFS Server Configuration window, as shown in Figure 6-8.

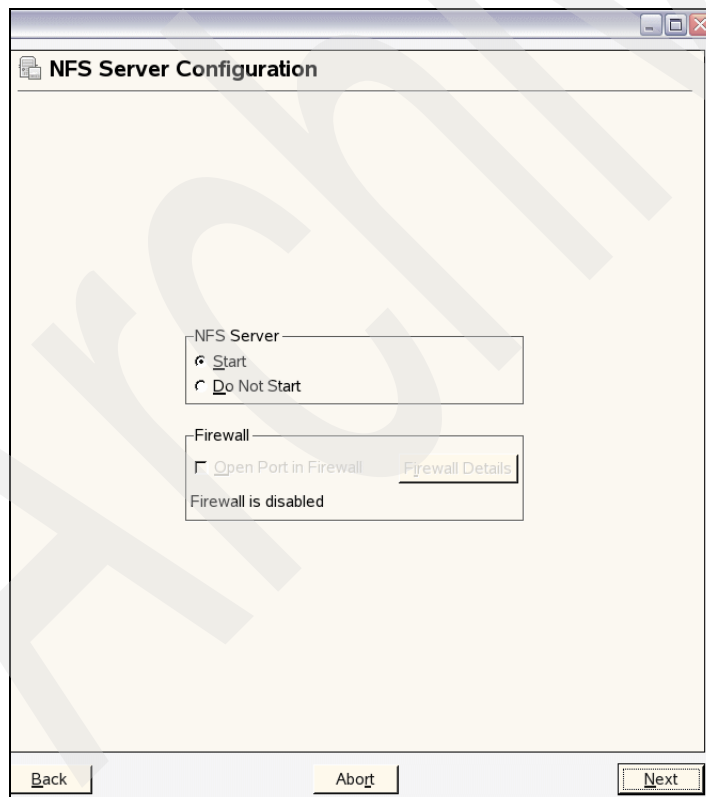


Figure 6-8 NFS Server Configuration



2. Select **Start** and then click **Next** to display the Directories to Export window (Figure 6-9). At first, there will be no entries because the DR550 FSG has no shares preconfigured for NFS export.

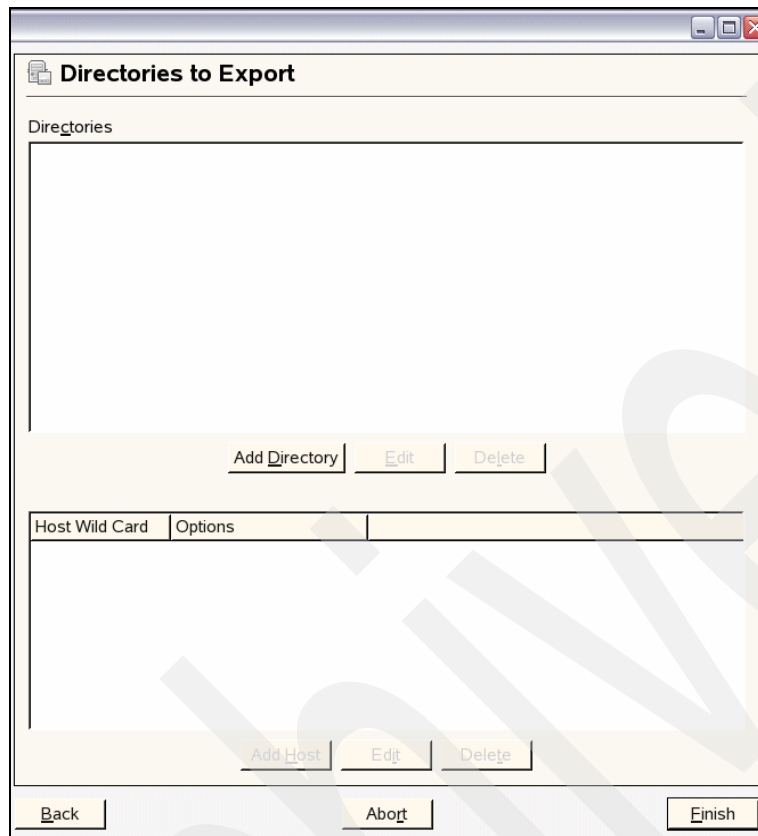


Figure 6-9 NFS Configuration - Directories to Export

3. Click the **Add Directory** button. A window displays where you can specify the information for the exported shares (Figure 6-10).




Figure 6-10 Add Directory /fsg/nfs-shares/finance for exporting

4. Enter the name of the directory to be exported and then click **OK**.
5. Next, specify which hosts are allowed to access that particular share. Here we selected a wild card for host (meaning any host can mount the selected directories), as shown in Figure 6-11 on page 251. (Refer to the Linux documentation for the meaning of the different parameters.)






Figure 6-11 NFS configuration - permissions on exported directory

You can also set other permissions and options, such as a file system ID (fsid). In Figure 6-12, we specified fsids of 1 and 2 for the two shares we had created (/fsg/nfs-shares/finance and /fsg/nfs-shares/management). The fsid number can be selected by the user but it must be unique. The concept and significance of fsids is explained in “High availability considerations for NFS” on page 252.

Once done with all the configurations, click **Finish**.

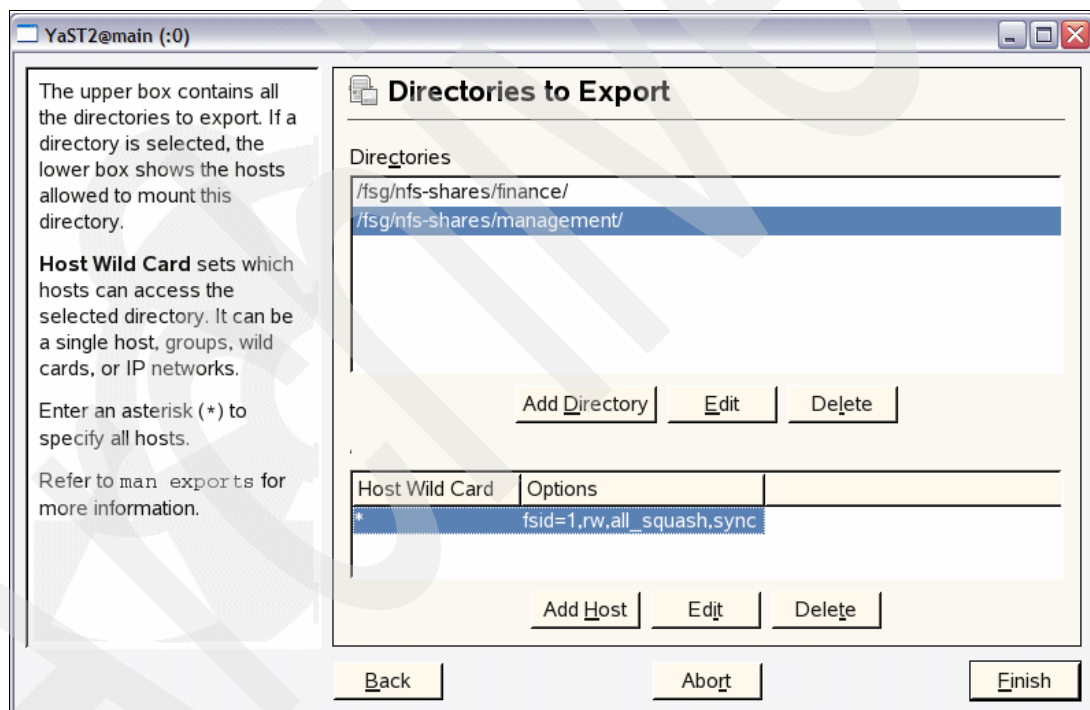


Figure 6-12 NFS Configuration - Add directory to export

6. Repeat steps 3 -5 to add as many directories and hosts as required.

To verify your settings, use the following command:

```
cat /etc/exports
```

With our example, you would see the following output:

```
/fsg/nfs-shares/finance/ *(fsid=2,rw,all_squash,sync)
/fsg/nfs-shares/management/ *(fsid=1,rw,all_squash,sync)
main:/home #
```

## High availability considerations for NFS

If you are adding or creating NFS shares in a high availability DR550 FSG cluster, you must specify a unique fsid for each share on both the Main and Supplementary. The fsid can be any 32-bit number except 0, but it must be unique among all the exported file systems.

For example:

The Main DR550 File System Gateway might have:

- fsid=1 for /fsg/share1
- fsid=2 for /fsg/share2

The Supplementary DR550 File System Gateway might have:

- fsid=3 for /fsg/share1
- fsid=4 for /fsg/share2

The fsid option forces the file system identification portion of the file handle and file attributes used on the wire to be num instead of a number, which is derived from the major and minor number of the block device on which the file system is mounted. It is important for NFS that both servers of the failover pair use the same NFS file handles for the shared file system, thus avoiding stale file handles after failover.

You can enter the fsid with the syntax fsid=1 in the Option field if you are using **yast2** to configure the NSF exports or you can enter the fsid into the /etc/exports file, using the **vi** editor.

Example of /etc/exports with a fsid:

```
/fsg/nfs-shares/finance/ *(fsid=2,rw,all_squash,sync)
/fsg/nfs-shares/management/ *(fsid=1,rw,all_squash,sync)
```

## Verify NFS is running

After configuring an NFS share, verify that the NFS server is running on the “active” DR550 FSG. If the NFS server is not running, you must start the NFS server.

To verify that the NFS server is running, at the command line, enter, as fsgadm:

```
sudo /etc/init.d/nfsserver status
```

To start the NFS server, at the command line, enter, as fsgadm:

```
sudo /etc/init.d/nfsserver start
```

**Note:** When using YaST on a high availability configuration, NFS services are inadvertently started on the current “standby” DR550 File System Gateway. NFS services should only be running on the “active” DR550 File System Gateway after configuration is complete. You must manually stop services on the “standby” DR550 File System Gateway with the following command at the command line as fsgadm:

```
sudo /etc/init.d/nfsserver stop
```

### 6.4.3 Create and configure shares for CIFS

Similar to what we did for NFS exports, the first step is to create the directories (or file tree structures) locally at the DR550 FSG and then configure them for export.

#### Set Samba stat cache size

Samba is an open source software that can be run on a platform other than Microsoft Windows, for example, UNIX, Linux, IBM System 390, OpenVMS, and other operating systems, allowing that host to interact with a Microsoft Windows client or server as though it is a Windows file and print server. The DR550 FSG comes preloaded with the Samba software and uses it to provide CIFS file access to Windows clients.

The Samba services are enabled by running the `smbd` daemon. The `smbd` is configured using the `smb.conf` file.

The `smbd` daemon maintains a cache (stat cache) of recent file names for speeding up file name mappings (between Windows and Linux). This stat cache can grow unbounded. You must, therefore, place a limit on Samba's stat cache size.

When a file is open by a client application, the operating system will normally lock the file to prevent access by other applications. The `smbd` daemon should normally check the status of the lock before attempting any operation with the file and this can slow down the system. By default, Samba only makes locking calls when asked to by a client (assuming that it is up to the client to manage the locks). This is called *opportunistic locking* (oplocks). However, with the DR550 FSG, we cannot assume that client applications will necessarily manage the locks. This would compromise data integrity. The oplocks must be set to no for the DR550 FSG.

To make the required changes to the `smbd.conf` file, proceed as follows:

1. Edit the `/etc/samba/smb.conf` file and add the following line in the global section:  
`max stat cache size = 51200`
2. Update oplocks to:  
`oplocks = no`
3. Save the `smb.conf` file.

#### Create directories and set permissions for CIFS shares

**Important:** The directories must be created under the `/fsg` directory, and the `drg` service must be active. In an HA configuration, the directories must only be created at the active (usually main) FSG node.

Example 6-8 shows the sequence of commands to create a new directory tree `cifs-shares/lab` (the names `cifs-shares` and `lab` are just examples). You must be logged in as `fsgadm` on the DR550 FSG.

*Example 6-8 Create directory tree /cifs-shares/lab*

---

```
cd /fsg
ls cifs-shares (check that the directory does not already exist)
mkdir cifs-shares
cd cifs-shares
mkdir lab
```

---

Assign the required permissions, as illustrated in Example 6-9.

*Example 6-9 Assign permissions*

```
chmod 755 /fsg/cifs-shares/lab <- sets permissions to read,write, execute (7)
 to owner (root), read, execute (5) to group,
 and read, execute (5) to user.

chgrp Labor /fsg/cifs-shares/lab <- sets the group to Labor
 (Labor must exist as a group)

chown lab_user /fsg/cifs-shares/lab <- this sets the user to lab_user
```

## Configure CIFS shares

**Important:** In an HA configuration, the CIFS shares must be configured on both the main and supplementary DR550 FSG (steps 1-6 below).

To set CIFS share properties, do the following:

1. Start the graphical Desktop with **startx**. Enter **yast2 samba-server** to display the Samba Configuration window, as shown in Figure 6-13 (note that the system will first display an Initializing window).

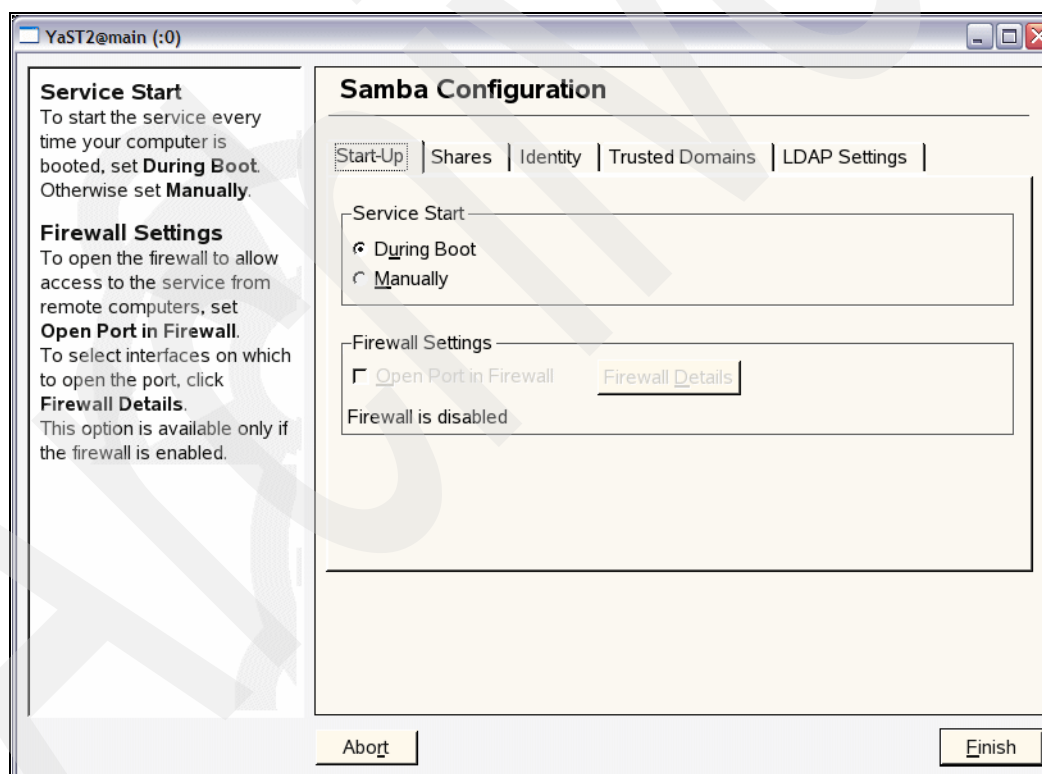


Figure 6-13 Samba Configuration Start-Up tab

2. Select to start service **During Boot**, then click the **Shares** tab.
3. On the Shares tab, you will see a list of already configured shares (if any). See Figure 6-14 on page 255.

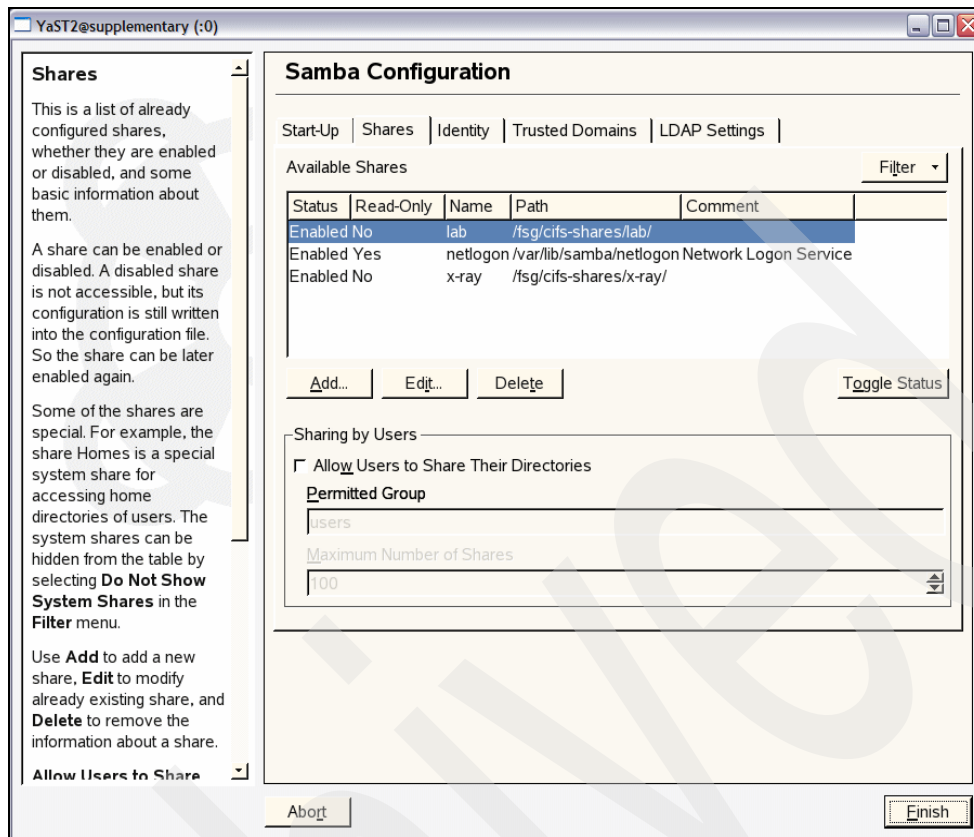


Figure 6-14 Samba Configuration Shares tab

- You can add new shares by clicking the **Add** button and then entering the required information in the New Share window (Figure 6-15).

**New Share**

Identification

Share Name

lab

Share Description

CIFS share Lab

Share Type

☐ Printer

☒ Directory

Share Path

/fsg/cifs-shares/lab/ Browse...

☐ Read-Only ☒ Inherit ACLs

Figure 6-15 CIFS shares configuration - Add new share "lab"

Enter the share name, the share path, and the permissions (Read Only or Inherit ACLs), then click **OK**. Add as many shares as you need by repeating this step.

5. Click the **Identity** tab to set up the server identity and its role in the network, as shown in Figure 6-16.

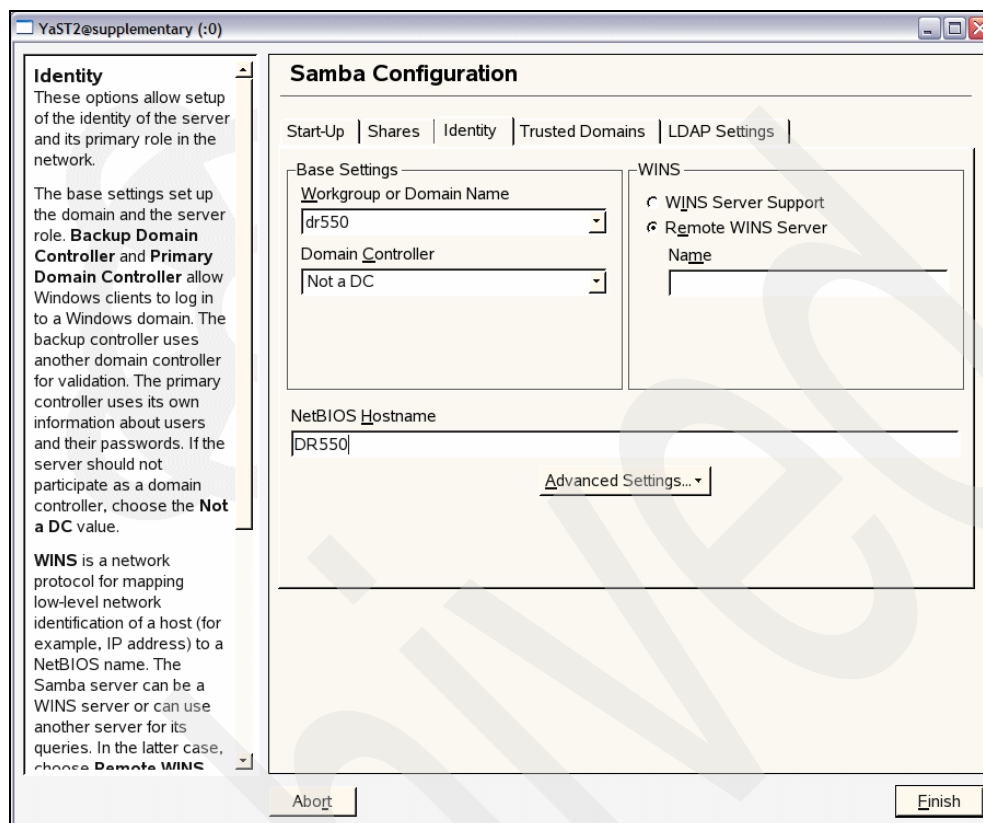


Figure 6-16 Samba Configuration - Server identity

6. When you are done, click **Finish**.

## Define the Samba password for the CIFS clients

**Important:** In an HA configuration, the Samba passwords must be defined on both the main and supplementary DR550 FSGs.

Samba maintains a separate list of passwords for Samba users, such as CIFS clients. To define the Samba password, enter the command **smbpasswd -a labuser**. (This command is used to set the password for user labuser.)

This will prompt you for a password and then display another prompt for you to confirm the password entered.

Repeat this command for all of the CIFS clients.

**Tip:** Samba maintains a separate list of passwords and they could potentially be different than the Linux password set for the same user ID (see 6.4.1, “Create local users and groups on the DR550 FSG” on page 245). However, to simplify maintenance, we highly recommend that you make the Samba password the same as the Linux password. Ideally, this should also match the Windows client user ID and password.

Refer to Chapter 7, “Centralized user management and FSG scenarios” on page 275 to learn how a centralized user management approach could greatly simplify the administration and maintenance of the various user IDs and passwords required by the DR550 FSG.

#### 6.4.4 Client access to the DR550 FSG

We recommend that a fixed port for the NFS mount D daemon be configured in `/etc/sysconfig/nfs`. For example:

```
MOUNTD_PORT="903"
```

The specific port number should be consistent with your defined firewall rules.

##### Mount clients

For NFS clients, enter the following command (at the client):

```
mount -t nfs -o hard, intr, udp, nfsvers=3, wsize=32768, rsize=32768 <ip_address of DR550 FSG>:/share1 <local_mountpoint>
```

For CIFS clients, on the Windows client system, enter:

```
net use x: \\<ip_address>\share1 /user: drg
```

You can also open Windows Explorer and select **Tools** → **Map Network Drive**.

### 6.5 DR550 FSG maintenance activities

In this section, we review some of the tasks that you might have to occasionally do to administer the DR550 FSG in your environment. The tasks include:

- ▶ Updating network configuration
- ▶ Changing and resetting passwords
- ▶ Changing file share customizations

#### 6.5.1 Updating the network configuration

The communication (data movement) between DR550 FSG and the SSAM Server, and between the DR550 FSG and the client systems, is based on TCP/IP over Ethernet. In a high availability configuration, the health of the clusters is monitored by heartbeat connections between the Ethernet ports of the two nodes connected together through a cross-over cable. The IP configuration is maintained in the DRGC bundles, which can be edited when modifications are required. (The DRGC was introduced in 3.7.7, “FSG post-installation and initial setup” on page 125.)

##### Changing the IP Address

To change the IP address of the main DR550 FSG:

1. Log on as `fsgadm` and `su` to root.
2. Copy the DRGC bundle (located at `/var/local/bundle-import/current/DRGC`) to a temporary working directory, for example, `/tmp`.
3. With a text editor (such as `vi`), open the DRGC bundle:

```
vi /tmp/DRGC
```

4. Make the required changes, such as MNCI (Main Client IP) or SPCI (Supplementary Client IP), and save the file.

**Note:** When you update the IP address of the DR550 File System Gateway, client services (NFS and CIFS) are restarted; any open client connections are dropped. Re-establish these connections or remount the share.

5. Copy this saved file to the import directory with this command:

```
cp /tmp/DRGC /var/local/bundle-import/import
```

The DR550 FSG automatically imports the updated DRGC file and moves it to the imported directory. The DRGC bundle is updated.

6. Restart the DG550 FSG:

```
sudo /etc/init.d/drg stop
sudo /etc/init.d/drg start
```

If using IBM Director, you must also stop and restart the agent after an IP address change.

**Important:** In a high availability cluster, the DRGC bundle is automatically copied to the Supplementary DG550 FSG. If cross-over cable IP addresses are updated, the DRGC bundle must be manually copied to the Supplementary DR550 FSG.

## 6.5.2 Changing and resetting passwords

Typically there are two categories of passwords that might need changing or resetting on the DR550 FSG:

- ▶ Client access for CIFS shares
- ▶ DR550 FSG access to the SSAM Server

### Client access for CIFS shares

1. Log in as fsgadm.
2. At the command line, enter **sudo smbpasswd <username>**. You will then be prompted for a new password.
3. Enter a new password.

A new password will be created and will replace the old one after you retype the new password correctly.

Example 6-10 gives an overview of this process.

#### *Example 6-10 Resetting the Samba password*

---

```
fsgadm@main:~> sudo smbpasswd <username>
New SMB password:<new password>
Retype new SMB password:<new password>
fsgadm@main:~>
```

---

As noted before, Samba maintains a separate list of passwords and they could potentially be different than the Linux password set for the same user ID. However, to simplify maintenance, we highly recommend that you make the Samba password the same as the Linux password.



Refer to Chapter 7, “Centralized user management and FSG scenarios” on page 275 to learn how a centralized user management approach could greatly simplify the administration and maintenance of the various user IDs and passwords required by the DR550 FSG.

### Access to the SSAM

To set or reset passwords for access to the DR550, the configuration has to be done in two places: the DR550 FSG and the SSAM Server. On the DR550 FSG, this is done by configuring the DRGC bundle (see 3.7.1, “Prepare DR550 SSAM for FSG attachment” on page 113 and 3.7.7, “FSG post-installation and initial setup” on page 125). On the SSAM Server, this is done using the **dsmdmc** command. Keep in mind that the passwords have to be identical on both sides.

To reset passwords giving the DR550 File System Gateway access to the SSAM Server:

1. Log on as fsgadm.
2. Copy the DRGC bundle to your temporary working directory (such as /tmp).  
The DRGC bundle is located in /var/local/bundle-import/current.
3. With a text editor, such as vi, open the DRGC bundle and edit the file.  
Update the SVPW attribute with a new password.
4. Save the DRGC bundle.
5. Copy the DRGC bundle to the import directory /var/local/bundle-import/import.
6. The DR550 FSG will automatically import the updated bundle and move it to the imported directory. The DRGC bundle is then updated. In a high availability cluster, the bundle is automatically copied to the Supplementary DR550 FSG.
7. From the DR550 FSG command prompt, connect to the SSAM Server by entering **sudo dsmdmc**.
8. At the command line, enter this Tivoli Storage Manager command:  

```
update node <node_name> <password>
```

(This password must be the same as the one stored in the DRGC bundle on the DR550 FSG.)

### 6.5.3 Changing file share customizations

Customizations of the DR550 FSG’s behavior with regard to the file shares is done by modifying the following features:

- ▶ Deletion protection on content
- ▶ File retention periods
- ▶ Parallel loading of content to retrieve all remaining items in a directory when the first item is requested
- ▶ Preloading of content to the Supplementary DR550 FSG
- ▶ Caching priority of directories

Each of these features can be changed at any time, but the deletion protection period and file type detection is only effective on content that has not yet been ingested into the DR550 FSG. In 6.3.3, “Editing DRG profiles” on page 241, you can find information about changing the features.

## 6.6 DR550 FSG data archiving and expiration scenarios

File System Gateway data archiving and expiration scenarios are provided here to illustrate the life cycle of an archived object from the time it is created on an FSG managed file system until its expiration in SSAM.

The DR550 environment used for running the test scenario consists of a single SSAM server, dual-node FSG, and an NFS client configured on a SLES 10 Linux server. An NTP client is used to synchronize clocks on all four servers with an external time source. This allows us to make time sensitive test result interpretation more reliable. Both FSG nodes are online during the test with “res\_group” resource group components active on the “main” FSG node.

We start our testing with the definition of an SSAM management class and copy group. For illustration purposes, we define a copygroup with the “Retain Minimum Days” (RETMIN) attribute set to 0 in order to be able to demonstrate the object expiration process in SSAM. An object will eventually expire even if this attribute is set to a value greater than 0, but in this case it will be retained in SSAM for a longer period of time.

---

### *Example 6-11 SSAM management class and copy group definitions*

---

```
dsmadm
IBM Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 5, Level 0.0
(c) Copyright by IBM Corporation and other(s) 1990, 2007. All Rights Reserved.
```

```
Enter your user id: admin
```

```
Enter your password:
```

```
Session established with server TSM: AIX-RS/6000
Server Version 5, Release 5, Level 0.0
Server date/time: 03/19/08 15:18:54 Last access: 03/19/08 15:09:40
```

```
tsm: TSM>define mgmtclass DRG-DOMAIN STANDARD DRG-10-MINUTES-MC
ANR1520I Management class DRG-10-MINUTES-MC defined in policy domain DRG-DOMAIN,
set STANDARD.
```

```
tsm: TSM>define copygroup DRG-DOMAIN STANDARD DRG-10-MINUTES-MC type=archive
destination=archivepool retver=0 retinit=event retmin=0
ANR1535I Archive copy group STANDARD defined in policy domain DRG-DOMAIN, set
STANDARD, management class DRG-10-MINUTES-MC.
```

```
tsm: TSM>validate policyset DRG-DOMAIN STANDARD
ANR1557W The space management migration destination in management class
DRG-10-MINUTES-MC does not refer to a defined storage pool: SPACEMGPOOL. If this
pool does not exist when policy set STANDARD is activated, clients will fail when
using this management class to migrate space-managed files to the server.
ANR1515I Policy set STANDARD validated in domain DRG-DOMAIN (ready for
activation).
```

```
tsm: TSM>activate policyset DRG-DOMAIN STANDARD
ANR1557W The space management migration destination in management class
DRG-10-MINUTES-MC does not refer to a defined storage pool: SPACEMGPOOL. If this
pool does not exist when policy set STANDARD is activated, clients will fail when
using this management class to migrate space-managed files to the server.
```

```

Do you wish to proceed? (Yes (Y)/No (N)) y
ANR1557W The space management migration destination in management class
DRG-10-MINUTES-MC does not refer to a defined storage pool: SPACEMGP00L. If this
pool does not exist when policy set STANDARD is activated, clients will fail when
using this management class to migrate space-managed files to the server.
ANR1514I Policy set STANDARD activated in policy domain DRG-DOMAIN.

```

---

Before data can be archived (using the new management class definition through the file system gateway), we need to define a new FSG profile that will link a file name with an SSAM management class. We also need to assign FSG retention rules (WORM protect period). In our example, we create a profile with a WORM protection period set to 10 minutes. Example 6-12 illustrates the profile definition (you must be logged on as root):

*Example 6-12 FSG profile definition*

```

main:~ # telnet localhost 1413
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^]'.
ADE 4.1 - Console Ready.

ade 2310001: / > cd /proc/BDED

ade 2310001: /proc/BDED > edit FGRP/10/FPRF

Now editing 'FGRP/10/FPRF'
FSG Profiles Bundle: FGRP/10/FPRF

Pattern Ingest Cache Protection Preload (D R S) Mgmt Class

1 1 forever DABL DABL DABL DRG-DEFAULTMC

ade 2310001: /proc/BDED > editprofile -a -p "/10minutes" -t 10i --mgmt-class
DRG-10-MINUTES-MC

FSG Profiles Bundle: FGRP/10/FPRF

Pattern Ingest Cache Protection Preload (D R S) Mgmt Class

1 1 forever DABL DABL DABL DRG-DEFAULTMC
2 /10minutes 1 10.0 mi DABL DABL DABL DRG-10-MINUTES-MC

ade 2310001: /proc/BDED > submit

Submitting bundle 'FGRP/10/FPRF'...
ade 2310001: /proc/BDED > Bundle 'FGRP/10/FPRF' submitted
ade 2310001: /proc/BDED > exit

Connection closed by foreign host.
main:~ #

```

---

Once the new FSG profile is defined, the new set of archive retention rules can be used on a client. In our example, we use an NFS client configured on a Linux SLES 10 host (refer to Example 6-13).

*Example 6-13 Exporting NFS share on main node and mounting filesystem on a client*

---

```
#On main FSG node
main# echo '/fsg *(rw,no_root_squash,sync)'> /etc/exports
main# cat /etc/exports
/fsg *(rw,no_root_squash,sync)
main # exportfs -a
main # exportfs
/fsg <world>

#On NFS client
nfsclient # mkdir /archivetest
nfsclient # mount 100.100.51.240:/fsg /archivetest
```

---

In this example, 100.100.51.240 is the public FSG service IP address. The next step in our scenario is to create the directory that matches the FSG profile pattern definition /10minutes. In our case, the fully qualified directory name on nfsclient is /archivetest/10minutes.

The following command on nfsclient can be used to create the directory: **mkdir /archivetest/10minutes**. The test environment is now ready for data archiving using the newly defined retention rules.

In the following scenario, a single file is created in the /archivetest/10minutes directory on nfsclient. After the file is created, we verify that the file cannot be removed immediately from the client (WORM protected), the file metadata is replicated to the supplementary FSG node, and the file data is stored in the SSAM server. This is illustrated in Example 6-14.

*Example 6-14 Creating file on nfsclient, verifying supplementary FSG node and SSAM application*

---

```
#On nfsclient:
nfsclient # date ; echo 'testing'> /archivetest/10minutes/test_file.txt
Wed Mar 19 16:51:38 MST 2008
nfsclient # ls -l /archivetest/10minutes/test_file.txt
-rw-r--r-- 1 root root 8 Mar 19 16:51 /archivetest/10minutes/test_file.txt
nfsclient # rm /archivetest/10minutes/test_file.txt
rm: cannot remove '/archivetest/10minutes/test_file.txt': Permission denied

#On supplementary FSG node:
supplementary # ls -l /fsg/10minutes/test_file.txt
-rw-r--r-- 1 root root 8 Mar 19 17:51 /fsg/10minutes/test_file.txt

#On SSAM server:
dsmadm
IBM Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 5, Level 0.0
(c) Copyright by IBM Corporation and other(s) 1990, 2007. All Rights Reserved.

Enter your user id: admin

Enter your password:

Session established with server TSM: AIX-RS/6000
Server Version 5, Release 5, Level 0.0
```

Server date/time: 03/19/08 17:18:50 Last access: 03/19/08 16:01:36

```
tsm: TSM>select * from archives where description='/10minutes/test_file.txt'
```

ANR2963W This SQL query may produce a very large result table, or may require a significant amount of time to compute.

Do you wish to proceed? (Yes (Y)/No (N)) y

```
 NODE_NAME: DRG-NODE
FILESPACE_NAME: /drg-10
 FILESPACE_ID: 1
 TYPE: FILE
 HL_NAME: /uuid/D1/1F/05/
 LL_NAME: D11F0520-1617-4254-A991-F0C1E520EF42
 OBJECT_ID: 27485
ARCHIVE_DATE: 2008-03-19 17:51:48.000000
 OWNER: DRG-USER
DESCRIPTION: /10minutes/test_file.txt
 CLASS_NAME: DRG-10-MINUTES-MC
```

tsm: TSM>

---

The archived object is bound to the DRG-10-MINUTES-MC SSAM management class according to the SSAM `select` command output, as shown in Example 6-14 on page 262. More details about the chronology of events can be found in `/var/local/log/bycast.log` log file on both FSG nodes. To narrow down the search and review only the log file entries relevant to the sample file, use the following command on the main FSG nodes:

```
grep test_file.txt /var/local/log/bycast.log
```

*Example 6-15 Main FSG “bycast.log” file entries related to “test\_file.txt”*

---

```
main# grep test_file.txt /var/local/log/bycast.log
Mar 19 17:51:38 main ADE: |2310001 26934 000046 FSGC ACTN
2008-03-19T23:51:38.655203| NOTICE 4271 FSGC: Creating new map entry for
ObjectPath '/fsg/10minutes/test_file.txt', opCode 'LCRF'
Mar 19 17:51:38 main ADE: |2310001 28245 000046 FSGC ACTN
2008-03-19T23:51:38.655837| NOTICE 3638 FSGC: Submitting CHEK
'/fsg/10minutes/test_file.txt' file to NewfileWatcher
Mar 19 17:51:38 main ADE: |2310001 29349 000056 NEWF NEWF
2008-03-19T23:51:38.655978| NOTICE 0862 NEWF: EVENT_CHECKNEWFILE:
/fsg/10minutes/test_file.txt
Mar 19 17:51:38 main ADE: |2310001 28245 000046 FSGC EPOA
2008-03-19T23:51:38.659546| NOTICE 3522 FSGC: Operation for
'/fsg/10minutes/test_file.txt' is done
Mar 19 17:51:38 main ADE: |2310001 26934 000056 NEWF SCNF
2008-03-19T23:51:38.668455| NOTICE 0763 NEWF: File no longer in use:
/fsg/10minutes/test_file.txt
Mar 19 17:51:48 main ADE: |2310001 28572 000056 NEWF INNF
2008-03-19T23:51:48.838711| NOTICE 0890 NEWF: Starting ingest of file
'/fsg/10minutes/test_file.txt'
Mar 19 17:51:48 main ADE: |2310001 29696 000122 INGS %CEA
2008-03-19T23:51:48.838917| NOTICE 0180 INGS: Ingest starting for
/fsg/10minutes/test_file.txt; filesize 8
```

```

Mar 19 17:51:49 main ADE: |2310001 30024 000056 NEWF %DED
2008-03-19T23:51:49.006196| NOTICE 0382 NEWF: New UUID
D11F0520-1617-4254-A991-F0C1E520EF42 for /fsg/10minutes/test_file.txt
Mar 19 17:51:49 main ADE: |2310001 27917 000046 FSGC NFCP
2008-03-19T23:51:49.006929| NOTICE 1685 FSGC: Newfile completed for
'/fsg/10minutes/test_file.txt', notifying Swapout
Mar 19 17:57:21 main ADE: |2310001 27917 000046 FSGC ACTN
2008-03-19T23:57:21.410673| WARNING 8478 FSGC: Operation not permitted for key
'/fsg/10minutes/test_file.txt', errno 13

```

---

According to the log records in the bycast.log file on the main FSG node, the file was created at 23:51:38.655203 (GMT time), which corresponds to 17:51:38.655203 time in a local time zone. An SSAM object with new UUID **D11F0520-1617-4254-A991-F0C1E520EF42** was sent to SSAM at 23:51:49.006196 (GMT time). The UUID identifier from the bycast.log record corresponds to the LL\_NAME attribute from SSAM **select** command output seen in Example 6-14 on page 262.

The WARNING message Operation not permitted for key **'/fsg/10minutes/test\_file.txt'** recorded at 23:57:21.410673 (GMT time) represents our attempt to remove the file while it is still WORM protected.

For more information about the log message format and DRG Service Attributes, see *IBM System Storage DR550 File System Gateway Software Version 1.1.1 User Guide*, GC27-2125.

In the described test scenario, the FSG profile was defined to provide WORM protection for a 10 minute period. Ten minutes after the file was created, it can be removed from the file system mounted on nfsclient.

**Note:** After the WORM protection period has expired for a file on the FSG file system, the file will not be removed automatically. The file can be removed explicitly by any user with read-write authority. If the file is not explicitly removed, it will remain indefinitely on the FSG file system.

---

*Example 6-16 Removing file after WORM protection period expiration*

---

```

#On nfsclient:
nfsclient # date ; rm /archivetest/10minutes/test_file.txt
Wed Mar 19 17:48:42 MST 2008

#On main FSG node:
main # ls /fsg/10minutes/test_file.txt
/bin/ls: /fsg/10minutes/test_file.txt: No such file or directory

#On supplementary FSG node:
supplementary # ls /fsg/10minutes/test_file.txt
/bin/ls: /fsg/10minutes/test_file.txt: No such file or directory

#On SSAM server:
dsmdmc
IBM Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 5, Level 0.0
(c) Copyright by IBM Corporation and other(s) 1990, 2007. All Rights Reserved.

Enter your user id: admin

```

Enter your password:

Session established with server TSM: AIX-RS/6000

Server Version 5, Release 5, Level 0.0

Server date/time: 03/19/08 17:18:50 Last access: 03/19/08 16:01:36

tsm: TSM>**select \* from archives where description='/10minutes/test\_file.txt'**

Session established with server TSM: AIX-RS/6000

Server Version 5, Release 5, Level 0.0

Server date/time: 03/19/08 18:49:20 Last access: 03/19/08 17:52:40

ANR2963W This SQL query may produce a very large result table, or may require a significant amount of time to compute.

Do you wish to proceed? (Yes (Y)/No (N)) y

      NODE\_NAME: DRG-NODE  
      FILESPEC\_NAME: /drg-10  
      FILESPEC\_ID: 1  
      TYPE: FILE  
      HL\_NAME: /uuid/D1/1F/05/  
      LL\_NAME: D11F0520-1617-4254-A991-FOC1E520EF42  
      OBJECT\_ID: 27485  
      ARCHIVE\_DATE: 2008-03-19 17:51:48.000000  
      OWNER: DRG-USER  
      DESCRIPTION: /10minutes/test\_file.txt  
      CLASS\_NAME: DRG-10-MINUTES-MC

tsm: TSM>**expire inventory wait=yes quiet=no**

ANR0984I Process 7 for EXPIRE INVENTORY started in the FOREGROUND at 18:49:32.

ANR0811I Inventory client file expiration started as process 7.

ANR2369I Database backup volume and recovery plan file expiration starting under process 7.

ANR0812I Inventory file expiration process 7 completed: examined 25 objects, deleting 0 backup objects, 19 archive objects, 0 DB backup volumes, and 0 recovery plan files. 0 errors were encountered.

ANR0987I Process 7 for EXPIRE INVENTORY running in the FOREGROUND processed 19 items with a completion state of SUCCESS at 18:49:32.

tsm: TSM>**select \* from archives where description='/10minutes/test\_file.txt'**

ANR2963W This SQL query may produce a very large result table, or may require a significant amount of time to compute.

Do you wish to proceed? (Yes (Y)/No (N)) y

ANR2034E SELECT: No match found using this criteria.

ANS8001I Return code 11.

---

The removal of the file on the nfsclient triggers a series of events:

1. The file is removed directly from the /fsg file system on the main FSG node.
2. The FSG replication process removes file metadata from the supplementary FSG node.
3. The Main FSG node sends an API event to the SSAM application specifying that the object is expired.

For more information about event-based retention policies, refer to 5.2.1, “Archive copy group retention parameters” on page 170. The SSAM object becomes eligible for expiration but not yet expired. To illustrate the object expiration process in SSAM, an **expire inventory** command had to be issued manually. After the inventory expiration process completes, the test object no longer exists in SSAM (see the ANR2034E error message returned by the second **select** command in Example 6-16 on page 264).

## 6.7 DR550 File System Gateway support

In this section, we review actions that need to be taken for several situations that might arise in case of technical problems, such as a failover between the DR550 FSG nodes, loss of connectivity between the DR550 FSG and SSAM Server, and other types of software and hardware failures.

### 6.7.1 Failover and data access

In general, failover should be seamless for clients of the High Availability DR550 FSG cluster. However, the data access can be impacted in some ways. When a failover occurs, any client operations that are in progress are interrupted. The client operations initiated while the standby DR550 FSG makes the transition to the “active” state are also interrupted. The effect of this is different on CIFS and NFS clients. Windows CIFS clients might need to remap their drive connections to the DR550 FSG once Samba starts on the newly “active” Supplementary DR550 FSG. NFS clients, however, will need to reissue a mount to their file systems on the FSG.

You should also be aware of the following behavior with a high availability cluster:

- ▶ If there are disruptions in the connection between the DR550 FSG and the SSAM Server, and they do not communicate with each other, this will not trigger a failover and the currently “active” and “standby” DR550 FSGs will continue their roles respectively.
- ▶ Data saved to the Main DR550 FSG, but not yet archived in SSAM Server at the time of a failover, will be archived only after the Main DR550 FSG is restored.
- ▶ If data has not been replicated, it must be manually replicated to the other DR550 FSG to ensure correct behavior in case of a failover later.
- ▶ If failover occurs, NFS clients must be remounted to allow for the storage and retrieval of data to resume. CIFS shares, however, generally, do not require a remapping of their drives.



## 6.7.2 Manual failover in a high availability cluster

In case of a high availability configuration, if there is a failure in the Main DR550 FSG, while it is in the “active” mode, more than likely that would trigger an automatic failover. In this case, the “standby” Supplementary node becomes active and all data traffic from the client to the DR550 FSG and the DR550 FSG to the SSAM Server will be through the Supplementary DR550 FSG. However, there might be instances when there was no automatic failover (for example, maintenance). In that case, a manual failover has to be initiated. The following steps will do the task:

1. Log on to the active DR550 FSG as fsgadm.
2. Enter the following command to monitor the current status of the clusters:

```
sudo crm_mon -i 15
```

This command will display the DR550 FSG cluster status and refresh every 15 seconds.

Find out if the failed DR550 FSG is the active one. Press <Ctrl>-C to exit and return to the command line.

3. If the failed DR550 FSG (for example, Main) is active, then run the following command from the active DR550 FSG:

```
sudo /etc/init.d/drg stop
```

4. Monitor the status of the clusters again with the command:

```
sudo crm_mon -i 15
```

Eventually the “standby” DR550 FSG transitions to “active” status and at this time failover is complete. Windows CIFS clients resume normal operations without having to be remapped, though operations in progress when the failover occurred are interrupted. NFS clients will have to remount the directories.

5. Restart the DR550 FSG stopped in step 3 with the following command:

```
sudo /etc/init.d/drg start
```

6. Monitor the status of the clusters again with the command:

```
sudo crm_mon -i 15
```

## 6.7.3 Clearing replication errors after failover

During a failover, there is a risk that a small set of data, which are normally files that have been written to the now failed “active” node, have not been archived into the SSAM Server or replicated to the “standby” DR550 FSG. Replication errors are logged to the bycast-err.log. These files need to be identified. The following steps will do that task:

1. Log in as fsgadm.
2. Switch to root with the command `su -`.
3. Run the command `cat /var/local/attributes/drg/RPER.xml`.
4. Check for replication errors. These are noted in the AVAL element container's 0x00000001 atom:

```
<atom name="0x00000001" type="UI64" value="0"
```

A value of 0 indicates there are no replication errors. Go to 6.7.4, “Recovery from a failed file system” on page 269 or 6.7.8, “DR550 File System Gateway replacement” on page 272. If the value is greater than 0, then continue with the next step.

5. Locate replication errors in the /var/local/log/bycast-err.log file.

Each event message includes the path to a file that is at risk of loss. Search for the keyword `FRPA_REPLICATION_ERROR`. The path and name of the replication error file is highlighted in bold in the sample event message shown in Example 6-17.

*Example 6-17 Sample Output of bycast-err.log*

---

```
|20130041 30062 000044 FSGC FCRA 2007-02-23T18:53:43.558259| ERROR 4896 FSGC:
FRPA_REPLICATION_ERROR for object '/fsg/directory/lostfile.txt',error 'ENOE',
opcode 'FACB', SC 262
```

---

6. Make a list of the complete file path to all directories that include files with replication errors.

7. Now manually replicate these files using the fsgtool utility.

- a. Verify the first file found in the replication error list is missing in the “active” DR550 FSG.

- i. Log on to the newly active Supplementary DR550 FSG as fsgadm.

- ii. Enter command `su -`.

- iii. Use the `ls` command to list the first file in the replication error list. For example:

```
ls /fsg/cifs-shares/itso/images/image1.jpg
```

(If the file is missing, you will get the message no such file or directory.)

- b. Verify that the file existed on the failed DR550 FSG.

- i. From the active DR550 FSG, enter the following:

```
ssh <ip-address of failed FSG>
```

- ii. Log in as fsgadm and enter the command `su -`.

- iii. Verify that the file exists with the following command:

```
ls /fsg/cifs-shares/itso/images/image1.jpg
```

- c. Using the fsgtool, dump the metadata information for the file with this command:

```
/usr/local/drg/fsgtool -getstat-x /fsg/cifs-shares/itso/images/image1.jpg
> /tmp/image1.jpg_info.txt
```

- d. Create the lost file in the “active” DR550 FSG.

- i. Return to the “active” DR550 FSG with this command:

```
exit
```

- ii. Stop the “active” DR550 FSG with this command:

```
/etc/init.d/drg stop
```

- iii. Copy the file containing the metadata dump (from the failed DR550 FSG):

```
scp fsgadm@<ip-address of failed FSG>:/tmp/image1.jpg_info.txt /tmp
```

- iv. Now create the file using the metadata dump created in step c:

```
/usr/local/drg/fsgtool -create /tmp/image1.jpg_info.txt -i
```

- v. Start the DR550 FSG again, which is stopped in step ii.

```
/etc/init.d/drg start
```

- vi. Verify that the lost file has been restored in the “active” node:

```
ls /fsg/cifs-shares/itso/images/image1.jpg
```

- e. Enter **exit** to log out of the DR550 FSG.

## 6.7.4 Recovery from a failed file system

A failure of the managed file system is very rare. If a failure does occur, it might be one of two issues:

- ▶ **Hardware**

At least two of the DR550 File System Gateway's data array drives have failed.

- ▶ **Software**

Failure of the operating system or other software leading to a corrupted file system

The procedure you select to restore the DR550 FSG to a normal functioning state depends on which of these above issues have caused the failure. There is also the option of a completely failed DR550 FSG system, which will need replacement (see 6.7.8, "DR550 File System Gateway replacement" on page 272).

**Note:** This recovery procedure cannot restore files that have been saved to the DR550 FSG share, but have not yet been archived into the SSAM Server. In fact, those files are lost as a result of the recovery procedure.

### Replace hard drives from the data array

This is required only in the case of multiple disk drive failures in a data array. The drives must be replaced and re-initialized before restoring the system. The server with the failed drives is first shut down, and the failed drives are replaced with new drives of the same size, type, and characteristics. After that the server is started and using utilities from the drive manufacturer, the proper RAID type, stripe size, and so on are set. The new drive array must then be partitioned and formatted.

### Reformat the Data Array

Log in as **fsgadm** on the DR550 FSG and enter the following commands:

- ▶ Switch to root with **su -**.
- ▶ Clear the directory with the command **rm -rf /var/local/fsg/\***.
- ▶ Run **umount /fsg**.
- ▶ Use standard YaST tool procedures to reinitialize the data array by running **yast2**.

### ***Close client connections (only for a stand-alone configuration)***

- ▶ Disconnect all CIFS shares from the Windows clients.
- ▶ Unmount all NFS shares from the NFS client machines.

### ***Start the DR550 FSG***

- ▶ Log in as **fsgadm** to the DR550 FSG.
- ▶ Enter the command **sudo /etc/init.d/drg start**.

The DRG service gets started and proceeds to automatically restore the file system from the last backup and the entire set of session files. This process might take a long time to complete, depending on the size of the file systems.

## 6.7.5 Identify orphaned files

Sometimes after a failure, files might have been saved to the DR550 FSG but not to the SSAM Server. Restoration of the last saved session file will result in “orphaned” files, files for which no metadata exists.

When a DR550 File System Gateway is restored, a restore-orphaned file is created that lists potentially orphaned objects. The file is appended with a time stamp. The restore-orphaned file is located at `/var/local/fsg/restore-orphaned.<time stamp>`.

Orphaned files can be retrieved, deleted, or ignored. It is up to the user to decide on how to administer orphaned files. If recovery is decided upon, the method of recovery is also up to the user.

## 6.7.6 Wait for pending files to complete replication (HA cluster only)

If content is actively being ingested to the SSAM Server through the “active” Supplementary DR550 File System Gateway, you must ensure that all content saved during failover has been saved to the SSAM Server before restoring the DR550 File System Gateways to their original roles.

1. Log in as fsgadm.
2. Switch to root by running `su -`.
3. Monitor the DRG services attribute FSGP.xml:

```
cat /var/local/attributes/drg/FSGP.xml
```

The contents of the FSGP.xml file are displayed, as shown in Example 6-18.

*Example 6-18 Output of FSGP.xml*

---

```
main:~ # cat /var/local/attributes/drg/FSGP.xml
<?xml version="1.0" encoding="UTF-8"?>
<containerxml version="1" xmlns="http://www.ibm.com/schemas/XML-container-1.0.0">
 <container name="FSGP">
 <atom name="AVER" type="UI32" value="2"/>
 <atom name="AVTP" type="FC32" value="UI64"/>
 <atom name="APER" type="FC32" value="READ"/>
 <atom name="ATIM" type="UI64" value="1205179567514644"/>
 <container name="AVAL">
 <atom name="0x00000000" type="UI32" value="2"/>
 <atom name="0x00000001" type="UI64" value="0"/>
 </container>
 </container>
</containerxml>
main:~ #
```

---

The highlighted line contains the “Files Stored to SSAM Server – Pending” value. Example 6-18 indicates that there are no pending files. An alarm notification is sent when this value passes its threshold. Monitor this attribute until the “Files Stored to SSAM Server – Pending” value decreases to zero.

The “Files Stored to SSAM Server – Pending” is larger than zero when files are saved to the DR550 File System Gateway more quickly than they can be written to the SSAM Server. If the value is greater than zero, ensure that the DR550 File System Gateway is operating normally.

You might also need to temporarily reduce ingests to the SSAM Server before you can proceed.

### 6.7.7 Restore DR550 FSGs to their original roles (HA cluster only)

In a high availability cluster, if the Main DR550 File System Gateway's state remains "standby" after the file restoration is complete, manually force the Main DR550 File System Gateway to return to its default role as "active." Recovery is complete once the Main DR550 File System Gateway returns to its default "active" role.

To restore the DR550 File System Gateways to their default roles, do the following steps:

1. Log in as fsgadm.
2. At the command line for the Supplementary DR550 File System Gateway, enter:  
`sudo /etc/init.d/drg stop`
3. Monitor the status of the DR550 File System Gateway. At the command line, enter:  
`sudo crm_mon -i 15`

The output of this command is shown in Figure 6-17.

```
Refresh in 15s...

=====
Last updated: Mon Mar 10 14:42:59 2008
Current DC: main (430e222e-c5bf-84b8-22e0-5fb725f1753b)
2 Nodes configured.
1 Resources configured.
=====

Node: main (4da2fba3-ea94-5f03-9aaf-d1f14b65b0f0): online
Node: supplementary (430e222e-c5bf-84b8-22e0-5fb725f1753b): online

Resource Group: res_group
 drgbase (Bycast::ocf:FSGBase): Started main
 ip_resource (heartbeat::ocf:IPAddr2): Started main
 pingtest (Bycast::ocf:PingTest): Started main
 drgactive (Bycast::ocf:FSGActive): Started main
 Filesharing-NFS (Bycast::ocf:NFS_Server): Started main
 Filesharing-CIFS (Bycast::ocf:Samba): Started main
```

Figure 6-17 Output of `crm_mon`

4. When the Main DR550 File System Gateway assumes the "active" mode, start the Supplementary DR550 File System Gateway.

At the command line for the Supplementary DR550 File System Gateway, enter:

```
sudo /etc/init.d/drg start
```

The Supplementary DR550 File System Gateway starts and assumes the "standby" state.

### ***Restore client connections***

The last stage of the recovery process is to restore client access to the restored DR550 File System Gateway. This procedure is not necessary if the DR550 File System Gateway is part of a high availability cluster.

To restore client connections, mount the NFS file systems or connect the CIFS network drives from the client.

### ***Verify connectivity***

Verify that the DR550 File System Gateway is connected to the network. To verify connectivity at the command line, enter:

```
ping <default gateway ip address>
```

or

```
ping <NTP server ip address>
```

## **6.7.8 DR550 File System Gateway replacement**

If it is determined that the system drives or the software on the DR550 File System Gateway server have failed to such an extent that the old DR550 File System Gateway's system drives cannot be used in a new server to recover the DR550 File System Gateway, you will need to follow the procedures listed below to replace the DR550 File System Gateway server.

### ***Process overview***

The replacement process for a DR550 File System Gateway uses many of the same procedures as those used to restore a failed file system or data array. There are some variations in the process as you work through it. An overview of the recovery process for stand-alone and high availability cluster configurations follows. As you execute the procedures, read each section carefully and perform only those steps needed to correct the failure you are repairing.

### ***Stand-alone configuration***

If the failed DR550 File System Gateway is operating in a stand-alone configuration, read-write access or read-only access to the file system is interrupted.

To return the system to full functionality, perform the following steps, each of which is described in detail below:

1. Install a new DR550 File System Gateway.
2. Close client connections to the DR550 File System Gateway.
3. Start the DR550 File System Gateway.
4. Re-apply any DR550 File System Gateway customizations.
5. Create new share directories exactly the same as the original share directories.
6. Restore client access to the recovered DR550 File System Gateway.

### ***High availability cluster***

In a high availability cluster, if the “active” Main DR550 File System Gateway fails, file system service to the DR550 is only briefly interrupted as the clustering service transitions the “standby” Supplementary DR550 File System Gateway to “active” status.

To return the system to full functionality, perform the following steps, each of which is described in detail below:

1. Replace the failed DR550 File System Gateway server, reinstall the operating system, and reinstall the DR550 File System Gateway software (as for a new server).
2. Start the DR550 File System Gateway. The repaired DR550 File System Gateway assumes the “standby” role.
3. Re-apply any DR550 File System Gateway customizations.
4. Create new share directories exactly the same as the original share directories.
5. Wait for any pending files stored to the “active” Supplementary DR550 File System Gateway to complete replication.
6. Restore DR550 File System Gateways to their original roles.

### **Preparation**

Replacement of a DR550 File System Gateway requires reloading the software, which requires:

- ▶ SUSE Linux Enterprise Server (SLES) 10 installation CDs
- ▶ Enablement Layer for IBM System Storage DR550 File System Gateway CD
- ▶ IBM System Storage DR550 File System Gateway Software CD

To restore a failed DR550 File System Gateway, you must also consult the following:

- ▶ *IBM DR550 File System Gateway Software Installation Guide*, GC27-2123
- ▶ *IBM DR550 File System Gateway Software Integration Guide*, GC27-2124

### **Replacing a DR550 File System Gateway Server**

Follow the instructions in the *IBM DR550 File System Gateway Software Installation Guide*, GC27-2123 and the *IBM DR550 File System Gateway Software Integration Guide*, GC27-2124 to prepare the DR550 File System Gateway on the replacement server.

Preparing the server installs the server's operating system, sets up the server by customizing the operating system and configuring drives, and loads the DR550 File System Gateway software. Do not start the DR550 File System Gateway until directed to do so by the procedure in the *IBM DR550 File System Gateway Software Installation Guide*, GC27-2123.

### **Close client connections**

If the failed DR550 File System Gateway is not part of a cluster, you must prevent clients from accessing the file system of the DR550 File System Gateway before restarting services by disconnecting the network drives.

### **Start the reinstalled server**

Start the reinstalled DR550 File System Gateway server and exit the command shell session:

1. Log on to the server as fsgadm.
2. At the command line, enter:  

```
sudo /etc/init.d/drg start
```

When started, the DRG service performs a startup scan and discovers that the file system requires restoration. The DR550 File System Gateway proceeds to restore the file system automatically from the last backup and the entire set of session files. This process might take a few hours to complete.

### ***Reapply DR550 File System Gateway customizations***

There are a number of customizations that might have been applied to a DR550 File System Gateway at the time that it was first integrated.

To reapply any required customizations, update the IP address configured in the DRGC bundle (see 6.5.1, “Updating the network configuration” on page 257).

### ***Recreate share directories***

Recreate new CIFS and NFS share directories exactly the same as the original share directories. For more information, see the *IBM DR550 File System Gateway Software Integration Guide*, GC27-2124. Share directories (on /fsg) will be restored automatically.

### ***Wait for pending files to complete replication (HA cluster only)***

Follow the instructions in 6.7.6, “Wait for pending files to complete replication (HA cluster only)” on page 270 to monitor the number of pending files, and ensure that they have been ingested before continuing.

### ***Restore DR550 File System Gateways to their original roles (HA cluster only)***

In a high availability cluster, if the DR550 File System Gateway is the Main DR550 File System Gateway, the DR550 File System Gateway status remains in “standby” after the file restoration is complete. Manually force the Main DR550 File System Gateway to return to its configured status as the “active” Main DR550 File System Gateway, as described in 6.7.2, “Manual failover in a high availability cluster” on page 267.

### ***Restore client connectivity***

The last stage of the recovery process is to restore client access to the restored DR550 File System Gateway. This procedure is not necessary if the DR550 File System Gateway is part of a high availability cluster.

To restore client connections, mount the NFS file systems or connect the CIFS network drives from the client.



## Centralized user management and FSG scenarios

We have seen in 6.4, “Integrating client applications with DR550 FSG” on page 245 that groups, user IDs, and passwords must be known to the SLES 10 Linux operating system on the FSG to grant access to the CIFS or NFS shares hosted by the FSG. The method detailed in 6.4.1, “Create local users and groups on the DR550 FSG” on page 245 does not require any additional software on the FSG, but has the major disadvantage of requiring user administration at the FSG level. Worse, if you have a high availability FSG cluster, you must repeat the same administration effort for the Supplementary FSG (there is no option to automatically copy the local user IDs and group definitions between Main and Supplementary).

In other words, local user management is probably acceptable and sustainable only for organizations with a very small number of users. Medium size and larger organization will want to use a centralized approach to FSG user management. In fact, they might already have such an infrastructure in place and will want to use it with their FSG.

The goal of this chapter is to illustrate, through practical scenarios, how to integrate two popular centralized user management technologies with the DR550 FSG.

**Important:** You always have to configure the file system rights and the NFS or CIFS shares, as explained in 6.4.2, “Create and configure shares for NFS exports” on page 248 or 6.4.3, “Create and configure shares for CIFS” on page 253, independently of the user management technique used (local or centralized).

## 7.1 Introduction to directories and LDAP

To improve functionality and ease-of-use, and to enable cost-effective administration of distributed applications, information describing the various users, applications, files, printers, and other resources accessible from a network is often collected into a special database that is called a directory.

### 7.1.1 Directory components

A directory contains a collection of objects organized in a tree structure. The LDAP naming model defines how entries are identified and organized. Entries are organized in a tree-like structure called the Directory Information Tree (DIT).

Entries are arranged within the DIT based on their distinguished name (DN). An DN is a unique name that unambiguously identifies a single entry. DNs are made up of a sequence of relative distinguished names (RDNs). Each RDN™ in a DN corresponds to a branch in the DIT leading from the root of the DIT to the directory entry. A Distinguished Name is composed of a sequence of RDNs separated by commas, such as `cn=thomas,ou=itso,o=ibm`.

You can organize entries, for example, after organizations and within a single organization; you can further split the tree into organizational units, and so on.

### 7.1.2 Directory and directory services

Sun's Network Information System (NIS) was one of the first attempts at centralizing user administration.

The Lightweight Directory Access Protocol (LDAP) and Microsoft Active Directory are popular examples of technologies that support centralized user management based on directories.

#### LDAP

Directories in LDAP are accessed using the client/server model. An application that wants to read or write information in a directory does not access the directory directly, but uses a set of programs or APIs that cause a message to be sent from one process to another. The second process retrieves the information on behalf of the first (client) application and returns the requested information if the client has permission to see the information. The format and contents of the messages exchanged between client and server must adhere to an agreed-upon protocol (LDAP conforms to RFC2307).

#### Microsoft Active Directory

Active Directory is an implementation of LDAP directory services by Microsoft for use primarily in Windows environments.

## 7.2 FSG configuration for open LDAP

This section illustrates how to set up and integrate the File System Gateway (FSG) for authentication of NFS clients using LDAP.

### 7.2.1 Environment for our NFS scenario

For our scenario, we used the following environment:

- ▶ SUSE Linux Enterprise Server (SLES) 10 as the LDAP server with the following properties:
  - IP address 9.11.218.254
  - DN DR550.IBM.COM
- ▶ We used the following users:
  - Control\_User
  - Finance\_User
- ▶ High availability FSG cluster:
  - FSG main with the following properties:
    - IP address 9.11.218.251
    - Server name main.DR550.IBM.COM
  - FSG supplementary with the following properties:
    - IP address 9.11.218.252
    - Server name supplementary.DR550.IBM.COM

### 7.2.2 Installing required LDAP packages

To implement a LDAP server on a SLES 10 system, you need to install the following SLES10 packages:

- ▶ openldap2
- ▶ openldap2-clients
- ▶ openldap2-devel
- ▶ nss\_ldap
- ▶ pam\_ldap

You can search for the LDAP modules by entering the command **yast2 sw\_single &**.

This opens the window shown in Figure 7-1 on page 278. Enter the word **ldap** in the Search field and click **Search** to obtain the list of packages that contain the word **ldap** in their name.

Select the required packages (using the check boxes) and click **Accept** (note that the screen capture in Figure 7-1 on page 278 shows more selected packages than necessary).

The installation of the different packages is now taking place. Wait until the whole process completes.

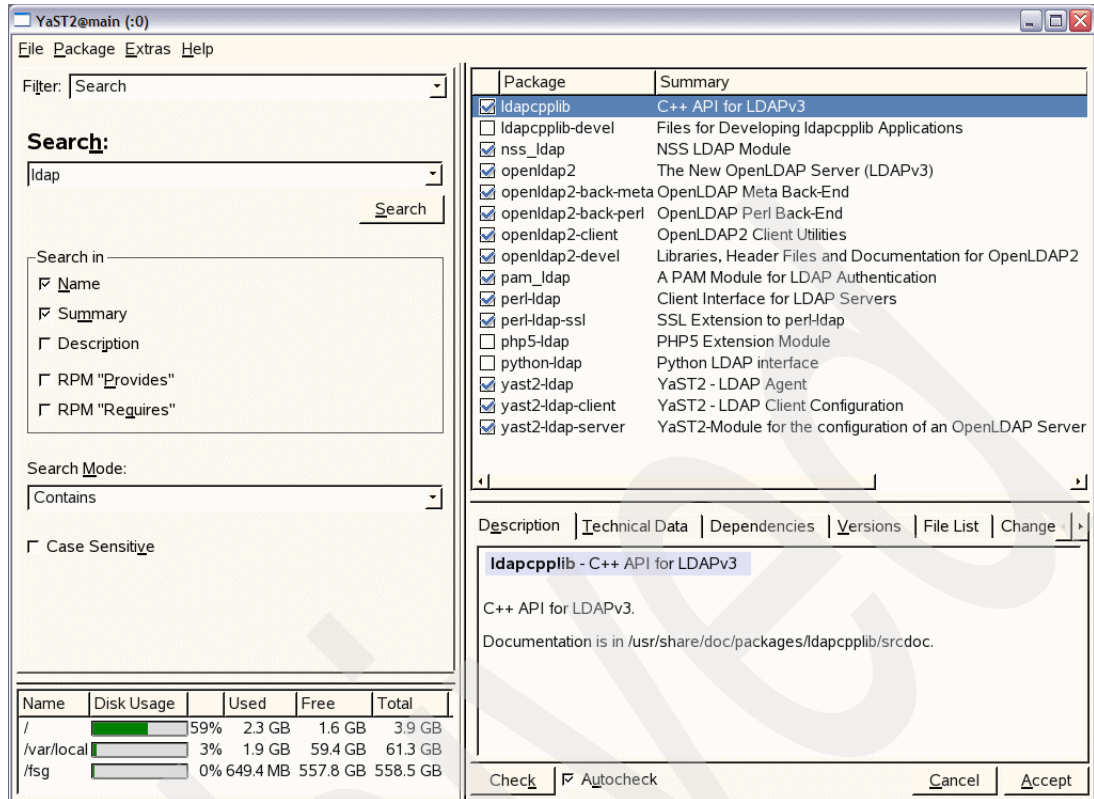


Figure 7-1 YaST software installation for LDAP

### 7.2.3 Defining users and groups in LDAP

The actual registration of user and group data when using LDAP is very similar to the procedure for creating local users in SLES10 and what we have described in 6.4.1, "Create local users and groups on the DR550 FSG" on page 245.

1. Enter the **yast2 &** command. The window shown in Figure 7-2 should appear. Select **Security & Users → User Management**.

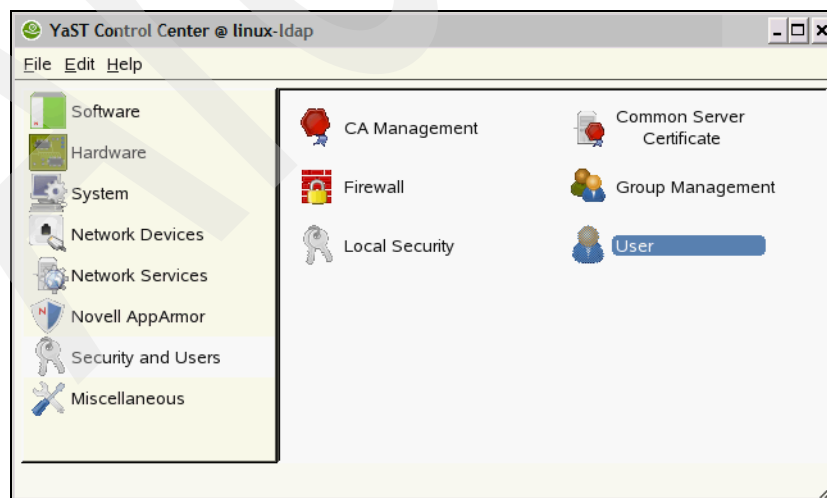


Figure 7-2 YaST - User management

2. Because we have installed LDAP on this server, you are prompted to enter the LDAP Administrator password, as shown in Figure 7-3.




Figure 7-3 Enter LDAP administrator password

3. Enter the password and click **OK**. A new window opens, as shown in Figure 7-4. You can use Set Filter to limit the view of users to the LDAP users.

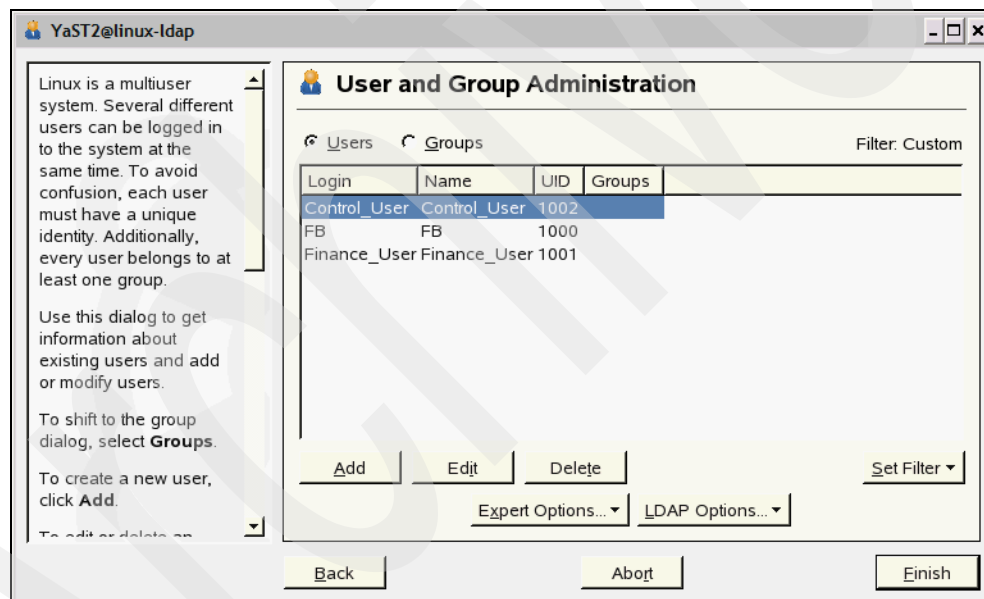


Figure 7-4 User administration

4. Click **Add** and enter the configuration of a new user. A dialog with four tabs opens (Figure 7-5 on page 280):
  - a. Specify user name, login, and password in the User Data tab.
  - b. Check the Details tab for the group membership, login shell, and home directory of the new user. If necessary, change the default to values that better suit your needs. The default values, as well as those of the password settings, can be defined with the procedure described in 6.4.1, “Create local users and groups on the DR550 FSG” on page 245.
  - c. Modify or accept the default Password Settings.
  - d. Go to the Plug-Ins tab, select the LDAP plug-in, and click **Launch** to configure additional LDAP attributes assigned to the new user.

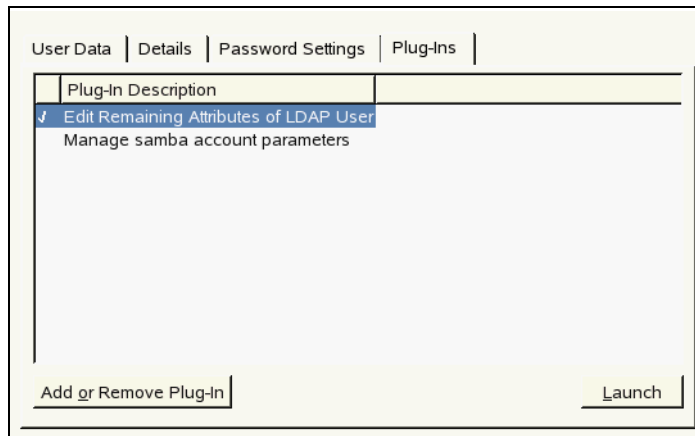


Figure 7-5 Additional LDAP attributes

5. Click **Accept** to apply your settings and leave the user configuration.

## 7.2.4 Configuring the LDAP client

To have the FSG (main and supplementary nodes) use the LDAP server we just configured for user authentication, the LDAP client code must be installed and configured on each node.

**Important:** In a high availability cluster, you must perform the configuration explained in this section both on the supplementary and main FSG nodes.

Enter the command `yast2 ldap &`. This displays the window shown in Figure 7-6 on page 281.

1. Select the **Use LDAP** button in the User Authentication box.
2. Enter the LDAP server IP address in the Address of LDAP Servers field.
3. Enter the LDAP Distinguished name in the LDAP base DN field. Alternatively, you can click **Fetch DN** after you have entered the LDAP server name (and assuming the service is started). In this case, a window is displayed and you can select the DN.
4. Click **Advanced Configuration**. This takes you to the window shown in Figure 7-7 on page 281.

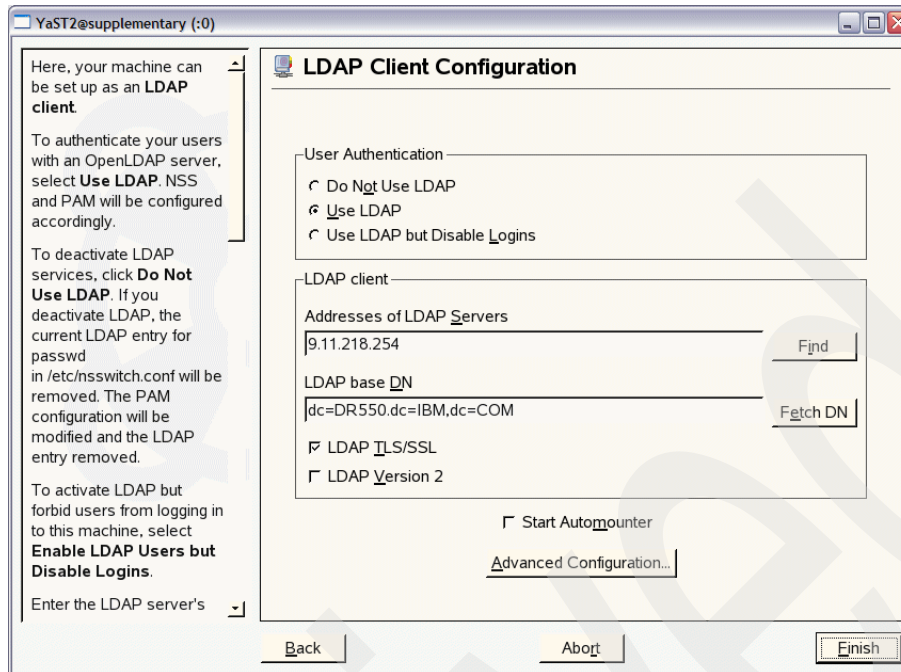


Figure 7-6 LDAP Client Configuration - main window

In the Advanced Configuration window, select the **Client Settings** tab and enter the various Naming Contexts. They should normally match the base DN specified in the previous window.

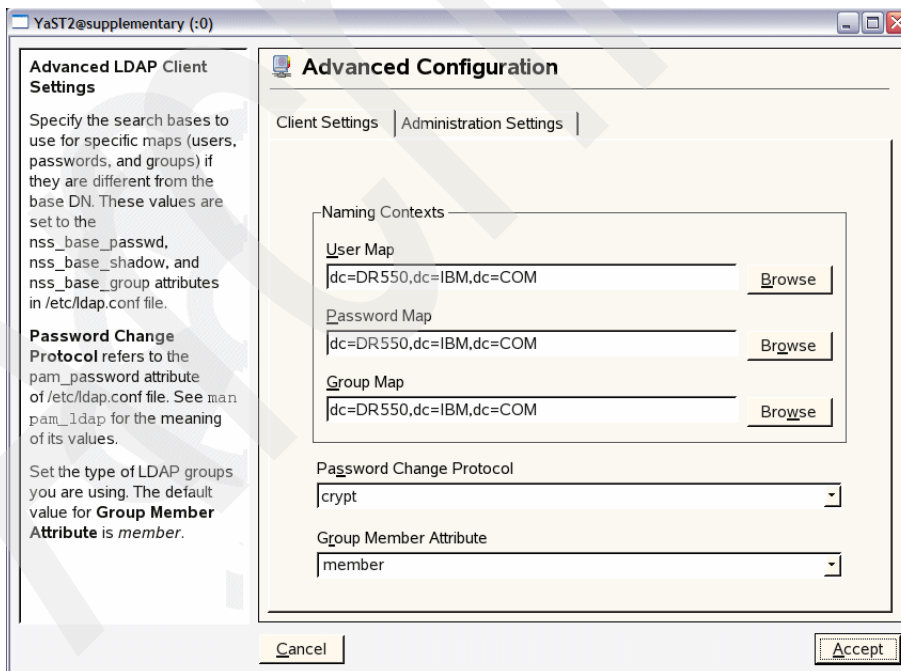


Figure 7-7 LDAP Client Configuration - Advanced Configuration

Select the **Administration Settings** tab, as shown in Figure 7-8, and click **Accept** to complete the LDAP client configuration.

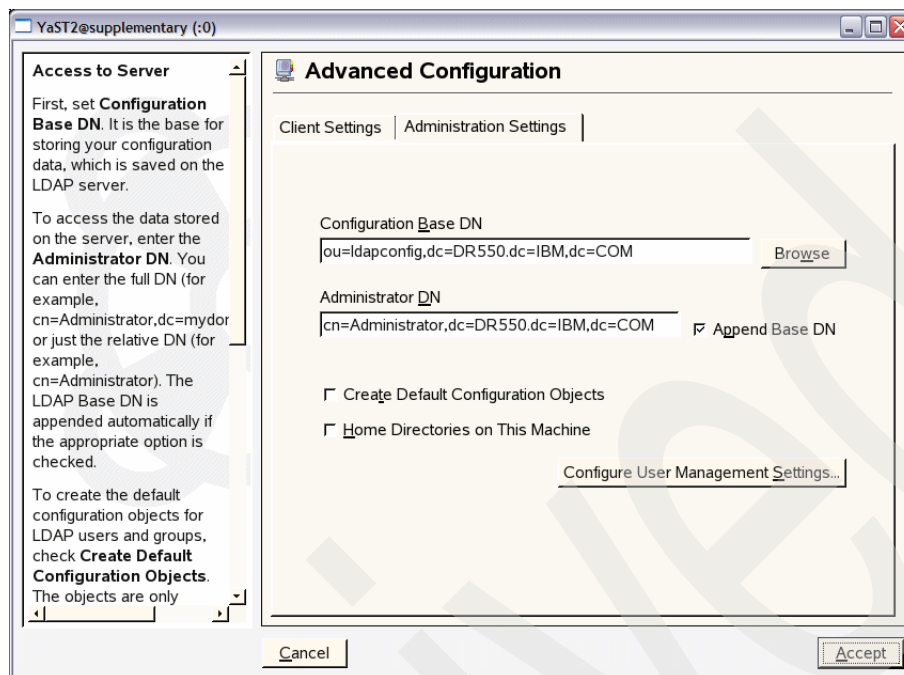


Figure 7-8 LDAP Client Configuration - Advanced Configuration - Admin Settings

At this point, you must restart the network services either by entering the command **/etc/init.d/network restart** or by rebooting the system.

Next, you can verify that your LDAP client is properly configured and working by entering the command **getent passwd**.

You should see a listing showing all the configured Linux accounts, including those defined on the LDAP server (shown in bold in Figure 7-9).

```
supplementary:~ # getent passwd
at:x:25:25:Batch jobs daemon:/var/spool/atjobs:/bin/bash
bin:x:1:1:bin:/bin:/bin/bash
daemon:x:2:2:Daemon:/sbin:/bin/bash
ftp:x:40:49:FTP account:/srv/ftp:/bin/bash
....
....
User:/var/lib/YaST2/suse-ncc-fakehome:/bin/bash
uucp:x:10:14:UNIX-to-UNIX CoPy system:/etc/uucp:/bin/bash
wwwrun:x:30:8:WWW daemon apache:/var/lib/wwwrun:/bin/false
drp:x:1000:100:drp:/home/drp:/bin/bash
hacluster:x:90:90:heartbeat
processes:/var/lib/heartbeat/cores/hacluster:/bin/false
ldap:x:76:70>User for OpenLDAP:/var/lib/ldap:/bin/bash
Control_User:x:1002:100:Control_User:/home/Control_User:/bin/bash
Finance_User:x:1001:100:Finance_User:/home/Finance_User:/bin/bash
supplementary:~ #
```

Figure 7-9 Output of getent passwd command



## Setting permissions and ownerships

Now you are able to use your LDAP user and group accounts when setting file system rights for the shares you want to use with NFS, as described in 6.4.2, “Create and configure shares for NFS exports” on page 248.

**Important:** Setting the permissions and ownership at the file system level must always be done, regardless of the authentication method (local, LDAP, or AD) that you use at your installation.

## 7.3 FSG configuration for Active Directory

This section illustrates how to set up and integrate the File System Gateway (FSG) for authentication of CIFS Windows clients using Active Directory.

### 7.3.1 Environment for our CIFS scenario

Our CIFS scenario is based on the following environment.

- ▶ Windows Server® 2003 R2 Domain Controller
  - IP address 9.11.218.253
  - Server name NDXTEAM54.DR550.IBM.LOCAL
  - Domain name DR550.IBM.LOCAL
- ▶ We used the following OU structure:
  - DR550.IBM.LOCAL
  - DR550
  - CIFS\_FSG
  - FSG\_Server
  - Lab
  - X-Ray
  - DR550\_Admins
  - Service\_Accounts

An OU is the smallest container that can be defined. A domain can contain many OUs, which can therefore be described as subunits or subcontainers of a domain. Each OU may contain objects such as users, groups, computers, security groups, printers, applications, security policies, and file shares. Other groups may be defined through the schema. Moreover, administrators can delegate administrative tasks for the OU to specific users, thus creating a kind of subadministrator. OUs enable you to define different administrative units while keeping a central instance of superusers, which are the domain administrators.

Table 7-1 on page 284 shows you the accounts we used. You do not have to create the computer accounts (main and supplementary, because they will show up automatically after joining the domain, but you should move them to a separate OU).

**Note:** Using different OUs is necessary only if you want to specify different group policies. In our scenario, we did not create different group policies, but it is possible with SLES10.

Table 7-1 Accounts created in AD

OU	Account	Description
FSG_Server	Main	Computer account FSG main
FSG_Server	Supplementary	Computer account FSG supplementary
Lab	Lab_User	Test user
X-Ray	X-Ray_User	Test user
DR550_Admins	fsg	Group for fsg users
Service_Accounts	SLES10Bind	LDAP Bind account

- ▶ FSG main
  - IP address 9.11.218.251
  - Server name main.DR550.IBM.LOCAL
- ▶ FSG supplementary
  - IP address 9.11.218.252
  - Server name supplementary.DR550.IBM.LOCAL

### 7.3.2 Preparing Active Directory

In this section, we review the settings required on the AD server to support the FSG.

#### Enabling Identity Management for UNIX

Identity Management for UNIX makes it easy to integrate computers running Windows into your existing UNIX enterprise. Active Directory network administrators can use Server for NIS to manage Network Information Service (NIS) domains, and Password Synchronization automatically synchronizes passwords between Windows and UNIX operating systems.

With minor differences, Identity Management for UNIX is compliant with the Internet Engineering Task Force (IETF) standard Request for Comments (RFC) 2307, meaning that your network's password and NIS attributes can be resolved by the Lightweight Directory Access Protocol (LDAP).

Because FSG, a UNIX based system, will use Active Directory for user authentication, you need to make sure that the Identity Management for UNIX component is installed as part of the Active Directory Services. On your Windows Server 2003 R2 Domain Controller, enable Identity Management for UNIX by going to the Control Panel, select **Add/Remove Programs** → **Add Windows Components** → **Active Directory Services**, and check **Identity Management for UNIX**, as shown in Figure 7-10 on page 285. Note that this will require a reboot and does require Schema Admin privileges. This will add a UNIX Properties tab to each user account in AD Users and Computers that will allow you to control the user UID, primary group GID, NIS Server setting, and user shell settings (such as /bin/bash).

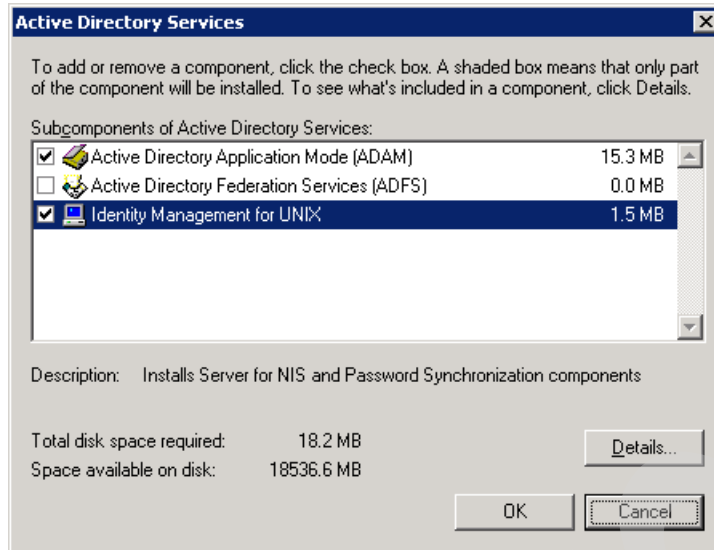


Figure 7-10 Install Identity Management for UNIX

## Create an LDAP Bind Account

Active Directory is an implementation of LDAP directory services for use primarily in Windows environments. You need to create an account in Active Directory that will be used to bind to Active Directory for LDAP queries. This account does *not* need any special privileges; in fact, making the account a member of Domain Guests and *not* a member of Domain Users is fine. This helps minimize any potential security risks as a result of this account.




Figure 7-11 Properties of the bind user SLES10Bind

## Prepare Active Directory accounts

Each Active Directory account that will authenticate from Linux (as would be the case with FSG users) must be configured with a UID and other UNIX attributes. This is accomplished through the UNIX Attributes tab on the properties dialog box of a user account. (Installing the Identity Management for UNIX component enables this, as mentioned in “Enabling Identity Management for UNIX” on page 284) Be sure to set login shell, home directory, UID, and primary UNIX group ID.

After all the user accounts have been configured, you are ready to configure the Linux server(s) for authentication against Active Directory. Figure 7-12 shows you the OU structure used for our scenario.

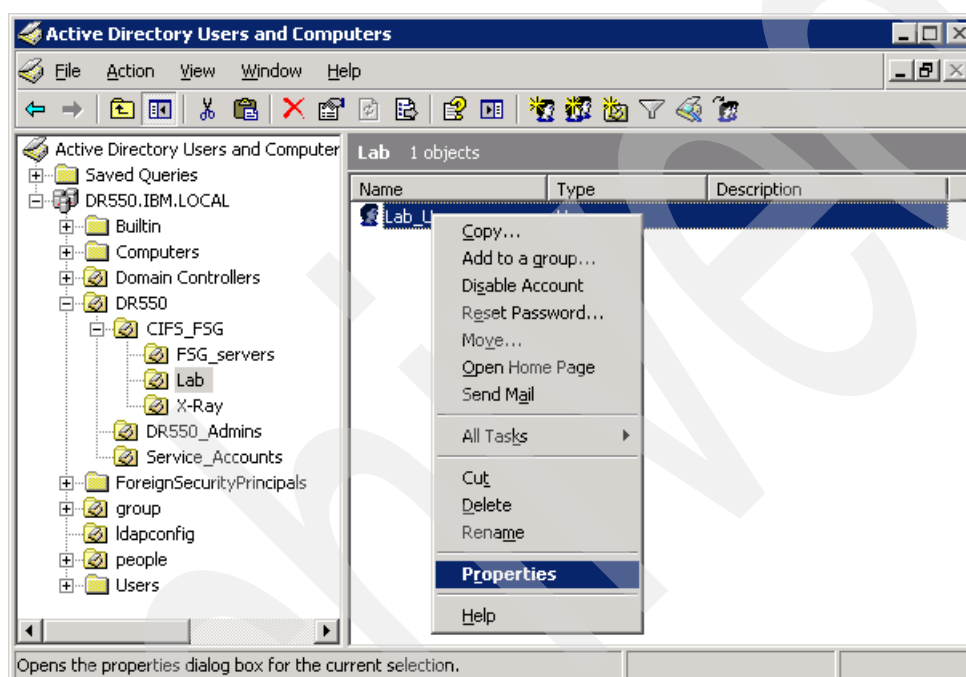
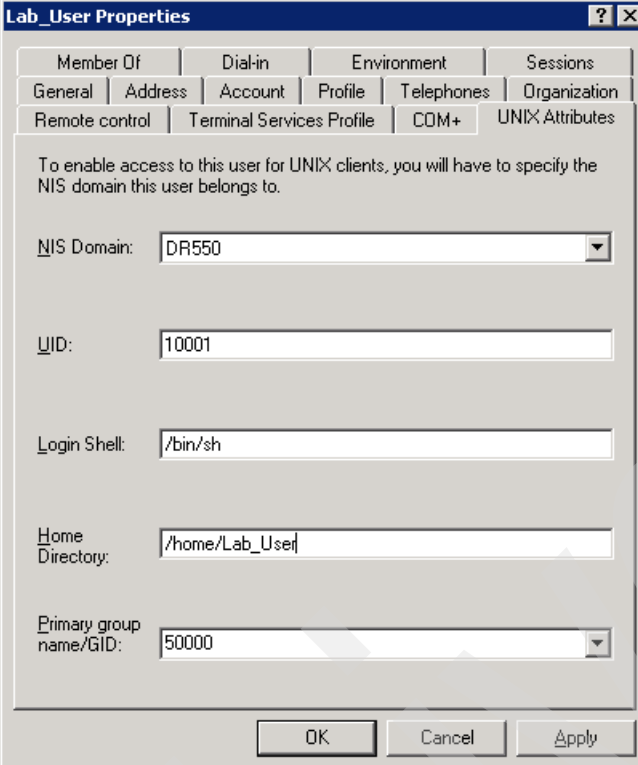


Figure 7-12 OU structure

You can verify or change the properties of the user account, but be sure you give a unique UID to the account. For the GID, we created the group fsg with the GID 50000, as shown in Figure 7-13 on page 287.



**Lab\_User Properties** [?] [X]

Member Of	Dial-in	Environment	Sessions
General	Address	Account	Profile
Telephones	Organization	Remote control	Terminal Services Profile
COM+	UNIX Attributes		

To enable access to this user for UNIX clients, you will have to specify the NIS domain this user belongs to.

NIS Domain:

UID:

Login Shell:

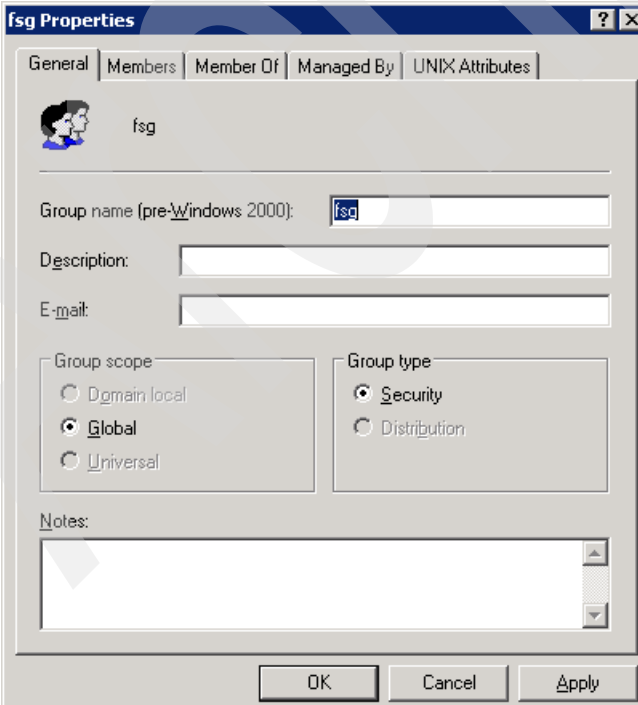
Home Directory:

Primary group name/GID:

OK Cancel Apply


Figure 7-13 Lab\_User Properties

Figure 7-14 and Figure 7-15 on page 288 show the configuration setting for the group fsg. We need this to create a GID, which we can use for the user accounts.



**fsg Properties** [?] [X]

General	Members	Member Of	Managed By	UNIX Attributes
---------	---------	-----------	------------	-----------------

 fsg

Group name (pre-Windows 2000):

Description:

E-mail:

Group scope

☐ Domain local

☒ Global

☐ Universal

Group type

☒ Security

☐ Distribution

Notes:

OK Cancel Apply

Figure 7-14 Properties of the group fsg

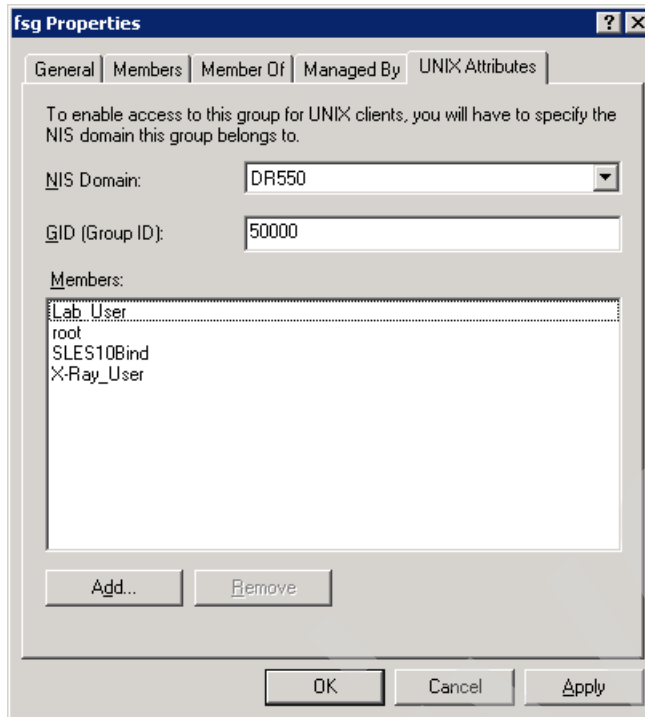


Figure 7-15 Members of the group fsg

### 7.3.3 Preparing the FSG for Active Directory

You need to install the following filesets on the FSG (both main and supplementary FSG on a high availability FSG cluster):

- ▶ krb5-client
- ▶ pam-krb5
- ▶ pam\_ldap
- ▶ pam\_smb
- ▶ nss\_ldap
- ▶ openldap2
- ▶ openldap2-client

You can use `yast2` to search for the required filesets by entering the command `yast2 sw_single &`. You will get the window shown in Figure 7-17 on page 290.

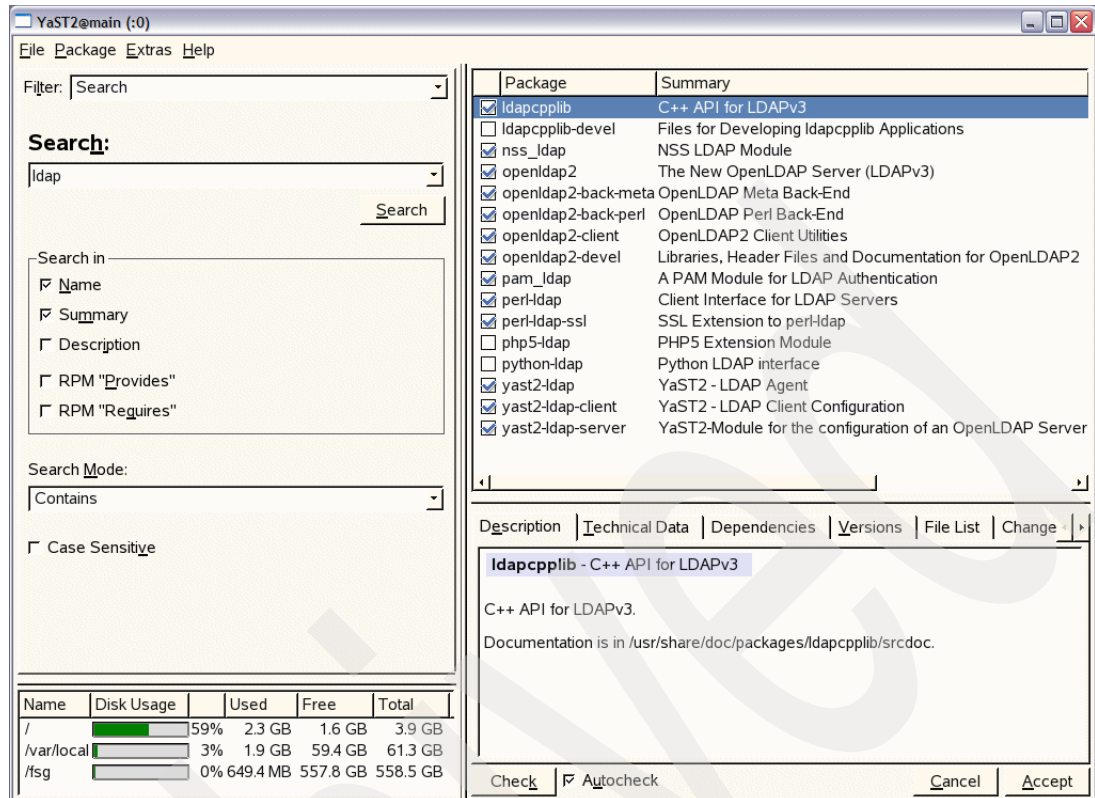


Figure 7-16 Installing LDAP modules with yast2

Select **Autocheck** to verify that all prerequisites will be installed and click **Accept** after selecting all the necessary files.

## Join the Active Directory with the FSG

The FSG CIFS share must now be configured to use Active Directory for authenticating CIFS clients (Samba clients). There are different ways to join the Active Directory. You can do it with the yast command:

```
yast2 samba-client &
```

This opens the window shown in Figure 7-17 on page 290. In this window, select the domain to join by directly entering its name. Alternatively, you can use the Browse button.

Select **Use SMB Information for Linux Authentication**.

We do not need home directories on the FSG because we only want to use CIFS share.

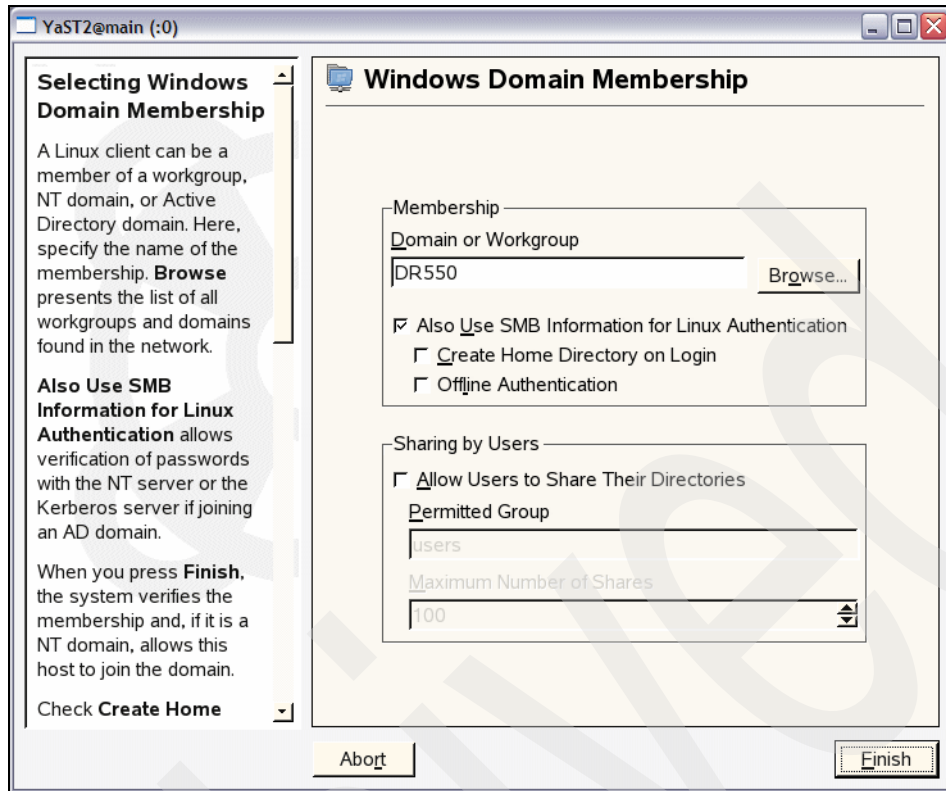


Figure 7-17 YaST - Windows Domain Membership

Click **Finish**. A window displays, as shown in Figure 7-18. Click **Yes** to join the domain.




Figure 7-18 Join Active Directory

You are now prompted for a user name and password, as shown in Figure 7-19 on page 291. You need to authenticate with a user name and password having the appropriate rights on your Active Directory server to add workstations to the AD domain.





Figure 7-19 Join Active Directory - Authentication

Enter the user name and password and click **OK**. If your settings were properly specified, you should get a box informing you that the domain was successfully joined (Figure 7-20).




Figure 7-20 Join Active Directory - success

Click **OK** to confirm.

Once the FSGs (main and supplementary) have successfully joined the Active Directory domain, you can move the computer objects from the OU Computer to a specific OU you have set in your domain for these type of objects. In our example, we move these computer objects to the FSG\_Servers OU. This is shown in Figure 7-21 on page 292.

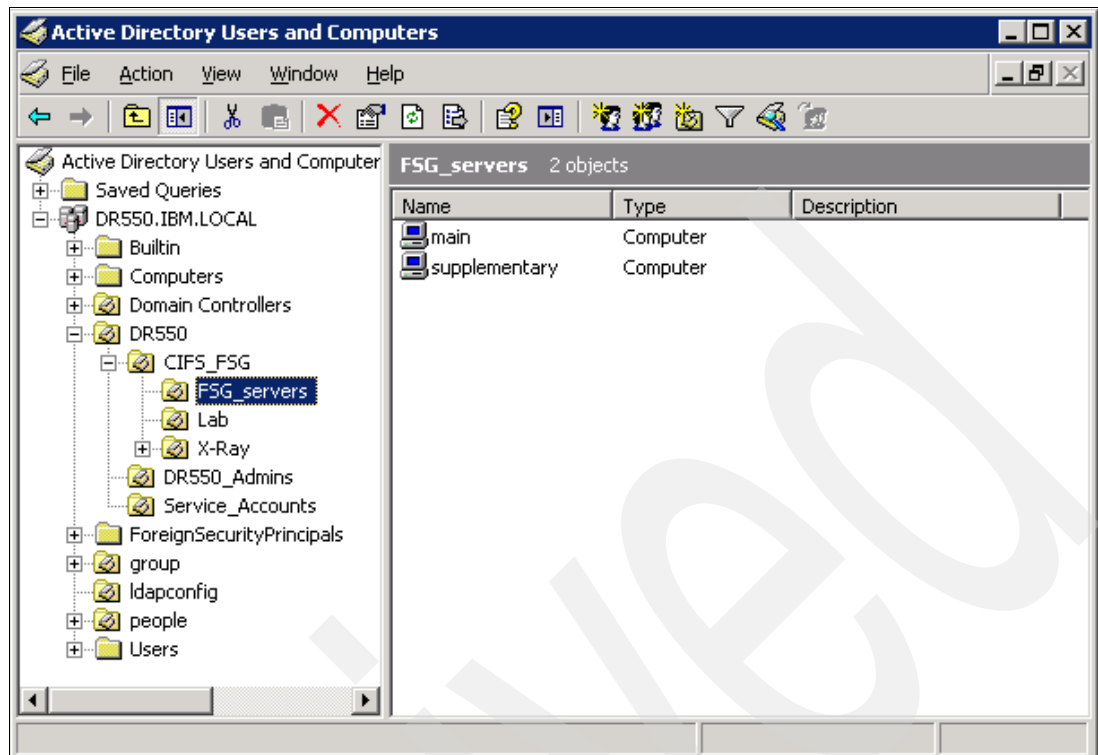


Figure 7-21 OU for the FSGs

You need to change the properties of these computer objects, as shown in Figure 7-22 and Figure 7-23 on page 293.

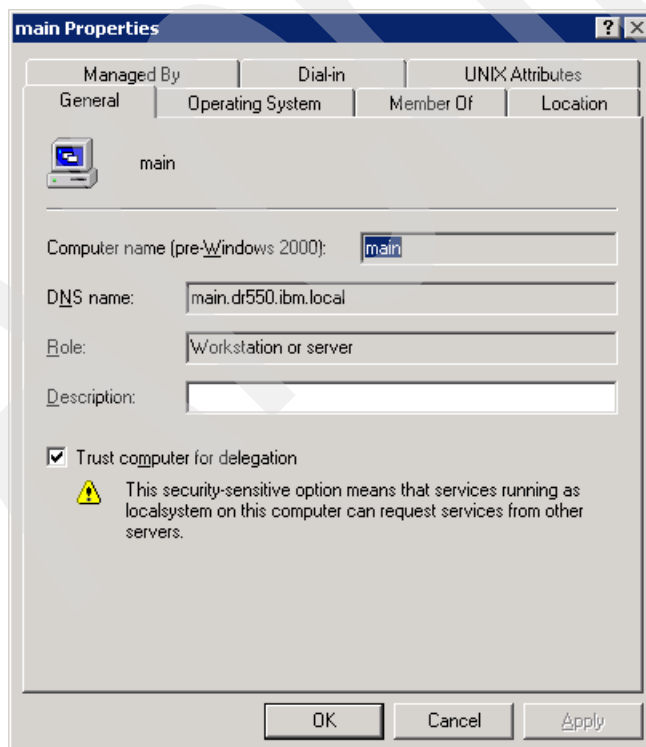


Figure 7-22 General Tab for FSG main

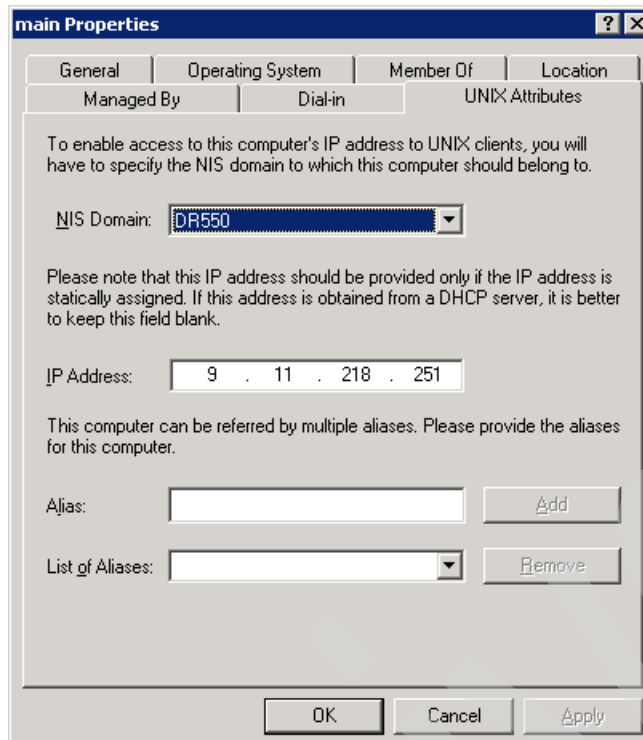


Figure 7-23 UNIX attributes for FSG main

This completes the setup work for Windows Active Directory. Now you have to configure a few files on each FSG manually. The following sections show you how the files look after applying these changes.

## Configure the FSG

On the FSG, set up your config files as illustrated in the following configuration files. See the file comment headers for the file names and locations (replace items such as domain names with settings specific to your environment). Some of the entries are already there from the FSG initial configuration, such as DNS names, for example. You need to add or change the other lines manually (using vi, for example):

1. We start first with the `/etc/hosts` file:

```
vi /etc/hosts
#####
/etc/hosts
#####
This file describes a number of host name-to-address
mappings for the TCP/IP subsystem. It is mostly
used at boot time, when no name servers are running.
On small systems, this file can be used instead of a
"named" name server.
Syntax:
#
IP-Address Full-Qualified-Hostname Short-Hostname
#
127.0.0.1 localhost
9.11.218.253 NDXTEAM54.dr550.ibm.local NDXTEAM54
special IPv6 addresses
::1 localhost ipv6-localhost ipv6-loopback
```

```
fe00::0 ipv6-localnet
```

```
ff00::0 ipv6-mcastprefix
```

```
ff02::1 ipv6-allnodes
```

```
ff02::2 ipv6-allrouters
```

```
ff02::3 ipv6-allhosts
```

2. Configure the /etc/krb5.conf file:

```
#####
/etc/krb5.conf for connecting with Windows Server 2003#
#####
[logging]
kdc = FILE:/var/log/krb5/krb5kdc.log
admin_server = FILE:/var/log/krb5/kadmind.log
default = SYSLOG:NOTICE:DAEMON

[libdefaults]
ticket_lifetime = 24000
default_realm = DR550.IBM.LOCAL
default_tkt_enctypes = des3-hmac-sha1 des-cbc-crc des-cbc-md5 aes256-cts
arcfour-hmac-md5
default_tgs_enctypes = des3-hmac-sha1 des-cbc-crc des-cbc-md5 aes256-cts
arcfour-hmac-md5
[realms]
DR550.IBM.LOCAL = {
 kdc = 9.11.218.253
 default_domain = DR550.IBM.LOCAL
 admin_server = 9.11.218.253
}
[domain_realm]
.dr550.ibm.local = DR550.IBM.LOCAL
dr550.ibm.local = DR550.IBM.LOCAL
```

3. Once Kerberos is configured, configure LDAP:

```
#####
custom /etc/ldap.conf for connecting with Server 2003 R2
#####
Your LDAP server. Must be resolvable without using LDAP.
host 9.11.218.253
The distinguished name of the search base.
base DC=DR550,DC=IBM,DC=LOCAL
url ldap://NDXTEAM54.DR550.IBM.LOCAL/
binddn cn=SLES10Bind,cn=Service_Accounts,cn=DR550,dc=DR550,dc=IBM,dc=LOCAL
bindpw password
scope sub
bind_timelimit 15
timelimit 15

OpenLDAP SSL mechanism
start_tls mechanism uses the normal LDAP port, LDAPS typically 636
ssl no
nss_map_attribute uniqueMember member
pam_filter objectclass=posixAccount
nss_base_passwd DC=DR550,DC=IBM,DC=LOCAL
nss_base_shadow DC=DR550,DC=IBM,DC=LOCAL
```

```
nss_base_group DC=DR550,DC=IBM,DC=LOCAL
nss_map_objectclass posixAccount user
nss_map_objectclass shadowAccount user
nss_map_objectclass posixGroup group
nss_map_attribute gecos cn
nss_map_attribute homeDirectory unixHomeDirectory
nss_map_attribute uniqueMember member
nss_initgroups_ignoreusers root,ldap
```

#### 4. Configure the Name Switch Service to use LDAP:

```
#####
/etc/nsswitch.conf
#####
An example Name Service Switch config file. This file should be
sorted with the most-used services at the beginning.
#
The entry '[NOTFOUND=return]' means that the search for an
entry should stop if the search in the previous entry turned
up nothing. Note that if the search failed due to some other reason
(like no NIS server responding) then the search continues with the
next entry.
#
Legal entries are:
#
compat Use compatibility setup
nisplus Use NIS+ (NIS version 3)
nis Use NIS (NIS version 2), also called YP
dns Use DNS (Domain Name Service)
files Use the local files
[NOTFOUND=return] Stop searching if not found so far
#
For more information, please read the nsswitch.conf.5 manual page.
#
passwd: files nis
shadow: files nis
group: files nis

passwd: files ldap
shadow: files ldap
group: files ldap

hosts: files dns wins
networks: files dns

services: files
protocols: files
rpc: files
ethers: files
netmasks: files
netgroup: files nis
publickey: files

bootparams: files
automount: files nis
aliases: files
```

5. Next, verify NTP. We only need to verify this file because in our case we configured NTP already for the FSG basic functions:

```

/etc/ntp.conf file

Sample NTP configuration file.
See package 'ntp-doc' for documentation, Mini-HOWTO and FAQ.
Copyright © 1998 S.u.S.E. GmbH Fuerth, Germany.

Author: Michael Andres,

Radio and modem clocks by convention have addresses in the
form 127.127.t.u, where t is the clock type and u is a unit
number in the range 0-3.

Most of these clocks require support in the form of a
serial port or special bus peripheral. The particular
device is normally specified by adding a soft link
/dev/device-u to the particular hardware device involved,
where u correspond to the unit number above.

Generic DCF77 clock on serial port (Conrad DCF77)
Address: 127.127.8.u
Serial Port: /dev/refclock-u

(create soft link /dev/refclock-0 to the particular ttyS?)

server 127.127.8.0 mode 5 prefer

Undisciplined Local Clock. This is a fake driver intended
for backup and when no outside source of synchronized time
is available.

server 127.127.1.0 # local clock (LCL)
fudge 127.127.1.0 stratum 10 # LCL is unsynchronized

Outside source of synchronized time

server xx.xx.xx.xx # IP address of server
server 9.11.218.253

Miscellaneous stuff

driftfile /var/lib/ntp/drift/ntp.drift # path for drift file

logfile /var/log/ntp # alternate log file
logconfig =syncstatus + sysevents
logconfig =all

statsdir /tmp/ # directory for statistics files
filegen peerstats file peerstats type day enable
```

```
filegen loopstats file loopstats type day enable
filegen clockstats file clockstats type day enable

#
Authentication stuff
#
keys /etc/ntp.keys # path for keys file
trustedkey 1 2 3 4 5 6 14 15 # define trusted keys
requestkey 15 # key (7) for accessing server variables
controlkey 15 # key (6) for accessing server variables
```

## 6. Configure Samba

At this point, we have Kerberos authentication configured, LDAP configured, NSS configured to use LDAP, and time synchronization configured and running. Now we need to get Samba configured to help automate the process of integrating into Active Directory:

```
#####
/etc/samba/smb.conf file
#####
smb.conf is the main Samba configuration file. You find a full
commented version at /usr/share/doc/packages/samba/examples/.
[global]
workgroup = DR550
realm = DR550.IBM.LOCAL
security = ads
encrypt passwords = yes
use kerberos keytab = true
password server = NDXTEAM54.DR550.IBM.LOCAL
netbios name = main
winbind use default domain = yes
winbind separator = +
idmap uid = 1000-59999
idmap gid = 1000-59999
winbind enum users = yes
winbind enum groups = yes
deadtime = 10
winbind cache time = 10
winbind nested groups = yes
template homedir = /home/%U
template shell = /bin/bash
client use spnego = yes
socket options = TCP_NODELAY SO_RCVBUF=16384 SO_SNDBUF=16384
idmap backend = ad
ldap idmap suffix = dc=DR550,dc=IBM,dc=LOCAL
ldap admin dn = cn=Administrator,cn=Users,dc=DR550,dc=IBM,dc=LOCAL
ldap suffix = dc=DR550,dc=IBM,dc=LOCAL
dns proxy = no
domain master = no
preferred master = no
max log size = 100
log file = /var/log/samba/%m.log
printing = cups
printcap name = cups
printcap cache time = 750
cups options = raw
map to guest = Bad User
```

```

include = /etc/samba/dhcp.conf
logon path = \%Lprofiles.msprofile
logon home = \%L%U.9xprofile
logon drive = P:
usershare allow guests = no
[admin]
comment = Windows Admin Access
path = /
valid users = "@Domain_Admins"
admin users = "@Domain_Admins"
read only = No
create mask = 0664
browseable = No
inherit permissions = Yes
[printers]
comment = All Printers
path = /var/tmp
printable = Yes
create mask = 0600
browseable = No
[print$]
comment = Printer Drivers
path = /var/lib/samba/drivers
write list = @ntadmin root
force group = ntadmin
create mask = 0664
directory mask = 0775

```

## 7. Configure Pluggable Authentication Mechanism (PAM)

SLES uses Pluggable Authentication Mechanism (PAM) to control authentication and authorization, so we next need to configure PAM to use Kerberos and LDAP, where necessary. There are a number of files that need to be configured to make this happen:

```

#####
/etc/pam.d/common-account - authorization settings common to all services
#####
This file is included from other service-specific PAM config
files, and should contain a list of the authorization modules
that define the central access policy for use on the system.
The default is to only deny service to users whose accounts
are expired.
#
account sufficient pam_krb5.so
account required pam_unix2.so

#####
/etc/pam.d/common-auth - authentication settings common to all services
#####
This file is included from other service-specific PAM config
files, and should contain a list of the authentication modules
that define the central authentication scheme for use on the
system (for example, /etc/shadow, LDAP, Kerberos, and so on). The default
is to use the traditional UNIX authentication mechanisms.
#
auth required pam_env.so
auth sufficient pam_krb5.so

```



```
auth required pam_unix2.so
```

```

/etc/pam.d/common-password - password-related modules common to all services

This file is included from other service-specific PAM config
files, and should contain a list of modules that define the
services to be used to change user passwords. The default is
pam_unix2 in combination with pam_pwcheck.
The "nullok" option allows users to change an empty password, else
empty passwords are treated as locked accounts.

To enable Blowfish or MD5 passwords, you should edit
/etc/default/passwd.

Alternate strength checking for passwords should be configured
in /etc/security/pam_pwcheck.conf.

pam_make can be used to rebuild NIS maps after password change.

password required pam_pwcheck.so nullok
password required pam_unix2.so nullok use_first_pass use_authtok
#password required pam_make.so /var/yp

/etc/pam.d/common-session - session-related modules common to all services

This file is included from other service-specific PAM config
files, and should contain a list of modules that define tasks
to be performed at the start and end of sessions of *any*
kind (both interactive and non-interactive). The default is
pam_unix2.

session required pam_limits.so
session required pam_unix2.so
session required pam_mkhomedir.so umask=0077 skel=/etc/skel
```

8. All these files are referenced by a master PAM configuration file:

```
##PAM-1.0
#####line above is part of this file#####
#/etc/pam.d/su config file

#auth sufficient pam_rootok.so
auth include common-auth
account include common-account
password include common-password
session include common-session
session optional pam_xauth.so
```

## Additional configuration steps

1. Run **getent passwd**. You should only see SLES10 local users in the command's output.
2. Run **kdestroy** to destroy any cached Kerberos tickets you might currently have.
3. Run **kinit Administrator@DR550.IBM.LOCAL** to create a new Kerberos ticket for the machine with Domain Administrator credentials.

4. Run **klist** to verify the Kerberos ticket, as shown in Figure 7-24.

```
main:/etc # klist
Ticket cache: FILE:/tmp/krb5cc_0
Default principal: Administrator@DR550.IBM.LOCAL

Valid starting Expires Service principal
06/04/07 16:31:27 06/04/07 23:11:27 krbtgt/DR550.IBM.LOCAL@DR550.IBM.LOCAL

Kerberos 4 ticket cache: /tmp/tkt0
klist: You have no tickets cached
```

*Figure 7-24 Output of kinit*

5. Run **net ads join -U Administrator@DR550.IBM.LOCAL** to join the machine to the domain using the Kerberos ticket of the domain administrative user, as shown in Figure 7-25.

```
main:/etc # net ads join -U Administrator@DR550.IBM.LOCAL
Administrator@DR550.IBM.LOCAL's password:
[2007/06/04 16:34:55, 0] libads/ldap.c:ads_add_machine_acct(1479)
 ads_add_machine_acct: Host account for main already exists - modifying old
 account
Using short domain name -- DR550
Joined 'MAIN' to realm 'DR550.IBM.LOCAL'
```

*Figure 7-25 Joining the Active Directory Domain*

6. Restart the applicable services and daemons:

```
/etc/init.d/smb stop
/etc/init.d/winbind stop
/etc/init.d/smb start
/etc/init.d/winbind start
```

7. Now you can run **wbinfo -u** to see all the domain users and **wbinfo -g** to see all domain groups, as shown in Figure 7-26 on page 301.

```

main:/etc # wbinfo -u
Administrator
Guest
SUPPORT_388945a0
ndxadmin
krbtgt
X-Ray_User
Lab_User
root
IUSR_NDXTEAM54
IWAM_NDXTEAM54
SLES10Bind

main:/etc # wbinfo -g
Domain Computers
Domain Controllers
Schema Admins
Enterprise Admins
Domain Admins
Domain Users
Domain Guests
Group Policy Creator Owners
DnsUpdateProxy

```

Figure 7-26 List of domain users and groups

8. Enter the command **su Administrator**. This should prompt you for the Administrator's password, as shown in Figure 7-27. Create a home dir for that user if necessary, and then switch to the user.

```

main:~ # su Administrator
administrator@main:/root>

```

Figure 7-27 Run su to an Active Directory account

- You should add the LDAP bind account (used to bind to LDAP for queries such as SLES10Bind10 in our example) to the list of SMB users with the command **smbpasswd -w**.
- Finally, using YaST (select **System** → **RunLevel Editor** to access it), enable the NTP, SMB, and Winbind daemons. Also, disable the nscd daemon to avoid caching problems and unwanted interaction with Winbind.

At this point, if you were successful in using **su** to switch to a Windows user, you should be able to reboot the system and log back into the system as a Windows user (be sure to use a Windows server that has UNIX attributes assigned in Active Directory).

**Note:** If you happen to get yourself locked out of the system, it will likely be an `/etc/nsswitch.conf` file problem. Simply boot with the SLES 10 installation disk using the “Recover System” option, then issue these commands to change the `/etc/nsswitch.conf` file:

**mount -w /dev/hda1 /mnt** (Where “`/dev/hda1`” is your system partition.)

**vi /mnt/etc/nsswitch.conf** (Remove the word “`ldap`” from the lines `passwd`, `group`, and `shadow`. They should only show “`files`” or “`compat`”.)

You can now reboot and log in as root to troubleshoot the problem.

## Setting permissions and ownerships

Now you are able to use your Active Directory user and group accounts to set file system rights for the shares you want to use with CIFS, as described in 6.4.3, “Create and configure shares for CIFS” on page 253 (the last section, “Define the Samba password for the CIFS clients” on page 256, does not apply when using Active Directory).

**Important:** Setting the permissions and ownership at the file system level must always be done, regardless of the authentication method (local, LDAP, or AD) that you use at your installation.

## Using IBM System Storage DR550

In this chapter, we discuss enterprise content management and data retention solutions, including brief descriptions of some of the solutions that are available today.

We cover the following topics in the chapter:

- ▶ Enterprise content management systems and IBM System Storage DR550
- ▶ IBM Enterprise Content Management (ECM) portfolio
- ▶ Using the IBM Tivoli Storage Manager Application Program Interface (API) with content management applications
- ▶ Integrating IBM Content Manager with IBM System Storage Archive Manager Server
- ▶ Integrating IBM Content Manager OnDemand with IBM System Storage Archive Manager Server
- ▶ IBM Optim overview, and integration of IBM Optim Archive with the DR550

## 8.1 Enterprise content management systems and DR550

Any enterprise content management systems that can use the archive and retention functions offered through the IBM Tivoli Storage Manager APIs should be able to communicate and use the DR550 SSAM Server.

For information regarding IBM Tivoli Storage Manager APIs, refer to:

<http://publib.boulder.ibm.com/infocenter/tivihelp>

Consult with your software application vendors to determine if their applications support the IBM Tivoli Storage Manager APIs for data archiving in an IBM System Storage Archive Manager environment. Several vendors have qualified (or have committed to qualify) one or more of their applications with IBM System Storage DR550. These vendors include:

- ▶ AXS-One
- ▶ BrainTribe (Formerly Compendium)
- ▶ Caminosoft
- ▶ Ceyoniq
- ▶ Easy Software
- ▶ Hummingbird
- ▶ Hyland Software (OnBase)
- ▶ Hyperwave
- ▶ IRIS Software (Documentum Connector)
- ▶ MBS Technologies (iSeries Connector for IBM CM V5)
- ▶ OpenText (formerly IXOS)
- ▶ Saperion
- ▶ SER Solutions
- ▶ TRIADE (NFS/CIFS/FTP Gateway)
- ▶ Veritas Enterprise Vault (formerly KVS)
- ▶ Waters (Creon Labs, NuGenesis)
- ▶ Windream
- ▶ Zantaz

Additional information about qualified independent software vendors (ISVs) is available at:

[http://www.ibm.com/servers/storage/disk/dr/pdf/DR550\\_interop\\_matrix.pdf](http://www.ibm.com/servers/storage/disk/dr/pdf/DR550_interop_matrix.pdf)

You can also find information about ISV support for the IBM Tivoli Storage Manager API at the Tivoli Knowledge Center, available at:

<http://www-306.ibm.com/software/tivoli/partners/public.jsp?tab=connect>

## 8.2 IBM Enterprise Content Management portfolio

Products offered through the IBM Enterprise Content Management (ECM) portfolio can be configured or integrated with IBM System Storage DR550. The IBM ECM suite of products manage content and core business process, and help ensure compliance while integrating with existing applications and infrastructure. They integrate and deliver critical business information when and where it is needed, in context, and under control.

Key products offered within IBM Enterprise Content Management portfolio are:

- ▶ *OmniFind™ Enterprise Edition*: Provides secure enterprise search among multiple repositories.
- ▶ *CommonStore for Lotus Domino* and *CommonStore for Exchange Server*: Provide e-mail management, including archive, search, and retrieval. E-mail management also includes e-mail attachments management.
- ▶ *IBM Content Manager* and *FileNet Content Manager*: Provide a comprehensive, scalable, and secure content management system that supports multiple platforms.
- ▶ *Content Manager OnDemand*: Provides efficient enterprise report management, including archive, search, and retrieve. Enterprise reports are usually the computer-generated reports, such as daily, monthly, quarterly, and annual accounting reports, and they are usually produced in high volume.
- ▶ *Document Manager*: Manages the complete life cycle of business documents, including check-in, check-out, and version control. They are usually used by engineering firms with complex design documents that go through multiple review and revision cycles.
- ▶ *IBM Records Manager* and *FileNet Records Manager*: Enable organizations to securely capture, declare, classify, store, and dispose of both electronic and physical records, to help ensure legal, regulatory, and industry compliance. IBM Records Manager provides the records management engine that can be embedded in the existing business applications.
- ▶ *FileNet Business Process Manager*: Automates, streamlines, and optimizes critical business processes by managing the flow of work between people and systems.
- ▶ *FileNet Image Manager Active Edition*: Provides comprehensive image management that include high volume capturing of paper documents as images, and search and retrieval of the images.

IBM ECM solutions provide the repository back-end services necessary to address an enterprise content management. It is common to use several products together in an enterprise-wide solution. For example, IBM Records Manager might be used together with IBM Content Manager or Content Manager OnDemand to provide the records management capability to the Content Manager or Content Manager OnDemand solutions.

Because this chapter discusses the usage of DR550 SSAM Server, in this section, we introduce the following IBM ECM products, which provide the core enterprise content repositories that interface with DR550 SSAM Server:

- ▶ IBM Content Manager
- ▶ IBM Content Manager OnDemand
- ▶ IBM FileNet Content Manager
- ▶ IBM FileNet Image Manager Active Edition

In addition, we also introduce the IBM FileNet P8 family of products.

More information about the IBM Enterprise Content Management portfolio of products is available at:

<http://www.ibm.com/software/data/cm/>

## 8.2.1 IBM Content Manager

IBM Content Manager is one of the core products in IBM ECM portfolio. Content Manager manages all types of digitized content across multiple platforms, databases, and applications. Built on a multi-tier, distributed architecture, it provides the scalability to grow from a single business unit to a geographically dispersed enterprise.

Content that Content Manager supports include HTML and XML Web content, document images, electronic office documents, printed output, audio, and video. Content Manager provides the content infrastructure (acting as the back-end content repository) for solutions such as compliance in a regulated life sciences environment, records management, document life cycle management, Lotus Notes® e-mail management, Exchange Server e-mail management, and digital media and Web content management.

Unlike simple file systems, Content Manager uses a powerful relational database to provide indexed search, security, and granular access control at the individual content item level. Content Manager provides check-in and check-out capabilities, version control, object-level access control, a flexible data model that enables compound document management, and advanced searching based on user-defined attributes. It also includes workflow functions, automatically routing and tracking content through a business process according to predefined rules.

Content Manager also supports replication to store and manage content objects in multiple locations, as well as application-transparent local caching (LAN cache). It enables content to be stored close to its point of use while remaining under central management control, reducing bandwidth requirements and increasing disaster protection.

A Content Manager solution consists of one Library Server, and one to many Resource Managers. The Library Server responds to user queries, while the Resource Managers maintain collections of content. Figure 8-1 shows how Content Manager system components interface with IBM System Storage DR550.

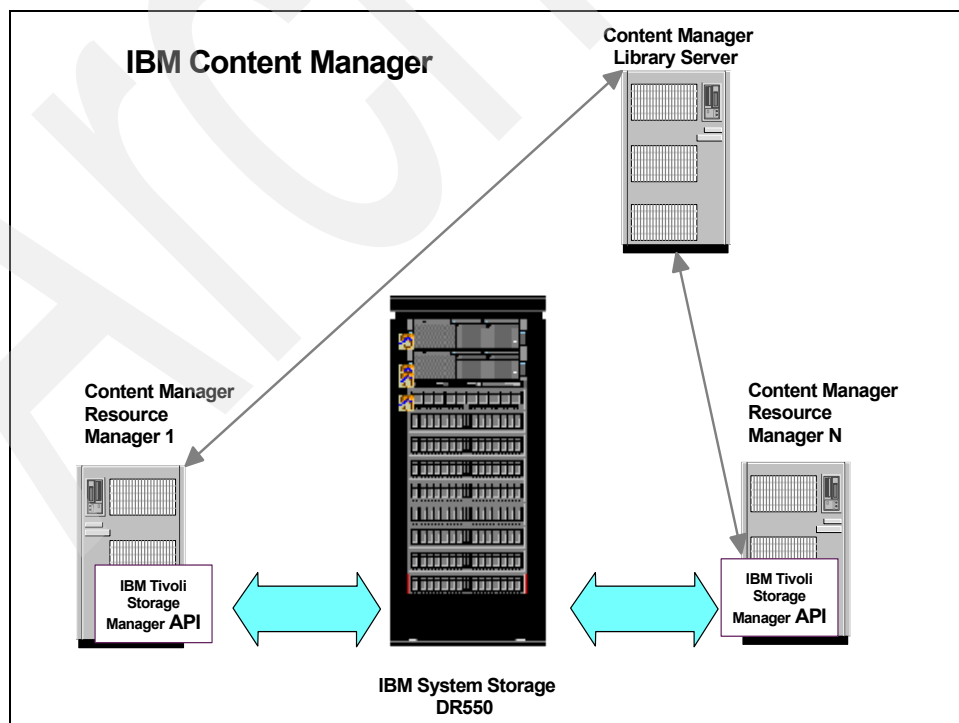


Figure 8-1 IBM Content Manager and IBM System Storage DR550



For more information, refer to the following IBM Redbooks publications:

- ▶ *Content Manager Implementation and Migration Cookbook*, SG24-7051, found at:  
<http://www.redbooks.ibm.com/abstracts/SG247051.html>
- ▶ *DB2 Content Manager for z/OS V8.3 Installation, Implementation, and Migration Guide*, SG24-6476, found at:  
<http://www.redbooks.ibm.com/abstracts/SG246476.html>
- ▶ *Content Manager Backup/Recovery and High Availability: Strategies, Options, and Procedures*, SG24-7063, found at:  
<http://www.redbooks.ibm.com/abstracts/SG247063.html>
- ▶ *Performance Tuning for Content Manager*, SG24-6949, found at:  
<http://www.redbooks.ibm.com/abstracts/SG246949.html>

## 8.2.2 IBM Content Manager OnDemand

IBM Content Manager OnDemand is a high-performance repository optimized for managing computer output. Content Manager OnDemand provides a highly reliable and flexible system to meet data archive and retrieval requirements. It can store and index about *two million* pages per hour, which is the performance demanded by high-volume billing or statement processing applications. OnDemand transforms any type of print output format, such as invoices, customer statements, bills, reports, and check images, into searchable, Web-integrated, electronic content that can be deployed in a variety of ways to meet customers' requirements and resolve their problems. OnDemand allows computer-printed output to be bundled, redirected over the network, and automatically distributed based on business rules.

One of the key strengths of OnDemand is its ability to directly archive computer print data streams. OnDemand is optimized to capture, search, present, and manage large collections of small objects, such as statements or bills.

Document capture is fast and unattended with multiple data types, such as line data, Xerox meta code, and optionally, PDFs. Index values are automatically extracted and stored in the database. Application report templates are predefined for easier administration. Access to the documents is provided by means of convenient and numerous search and presentation options.

An OnDemand solution consists of one Library Server and one or more Object Servers. The Library Server stores data indexes and the Object Servers store data objects. Object Servers can be local or remote. Each Object Server can have Tivoli Storage Manager connected to manage long-term archival to other magnetic, optical, and storage. The OnDemand Object Server communicates with the Tivoli Storage Manager server through the Tivoli Storage Manager API.

Figure 8-2 shows how Content Manager OnDemand components interface with IBM System Storage DR550.

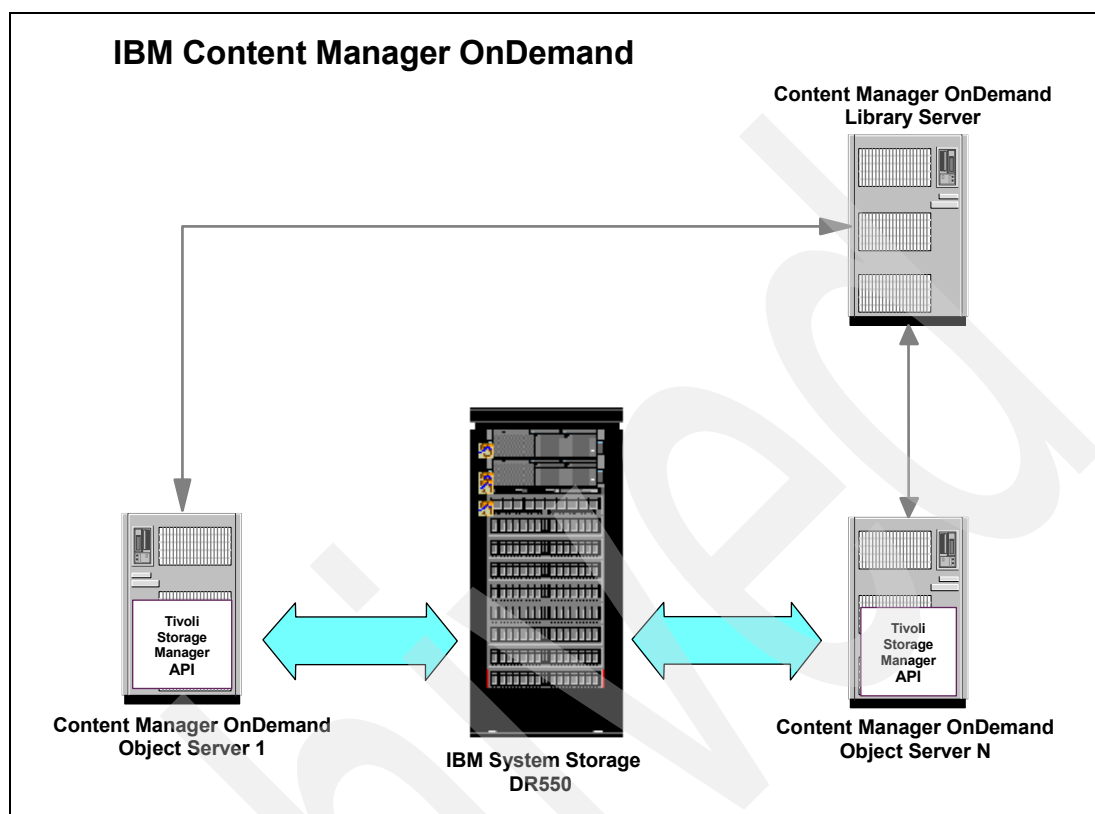


Figure 8-2 Content Manager OnDemand Object Servers interfacing with Tivoli Storage Manager storage

For more information, refer to the following IBM Redbooks publications:

- ▶ *Content Manager OnDemand Guide*, SG24-6915, found at:  
<http://www.redbooks.ibm.com/abstracts/SG246915.html>
- ▶ *Content Manager OnDemand Backup, Recovery, and High Availability*, SG24-6444, found at:  
<http://www.redbooks.ibm.com/abstracts/SG246444.html>

### 8.2.3 IBM FileNet Content Manager

Also part of the IBM Enterprise Content Management portfolio is the FileNet P8 family of products, which includes back-end services, development tools, and applications that address enterprise content and process management requirements. IBM FileNet Content Manager is one of the core products in the FileNet P8 family.

IBM FileNet Content Manager provides full content life cycle and extensive document management capabilities for digital content. It combines document management with workflow and process capabilities to automate and drive content-related tasks and activities.

FileNet Content Manager streamlines document management tasks by providing content versioning and parent-child capabilities, approval workflows, and integrated publishing support. It delivers the ability to actively manage content across the enterprise regardless of what repository it resides in, using FileNet Content Federation Services.

FileNet Content Manager delivers the unique advantage of Active Content: content and documents that actively drive process automation to completion. With Active Content, documents and other forms of content can drive task resolution and reduce time, cost, and risk. Businesses can respond immediately to the business and transaction events that set critical processes in motion and drive them quickly to completion, increasing the overall responsiveness and agility of the operations.

FileNet Content Manager consists of a Content Engine and one to many object stores (among other components). At the core of the Content Engine are repository services for capturing, managing, and storing business related digital assets. Multiple repositories, called object stores, can be created and managed within a single system to serve the business requirements. Object stores can be configured to store content in a database, a file system, a fixed content device, or a combination of these options. An object store is capable of storing a variety of business-related data, for example, an insurance claim, a customer loan account, or information about Business Partners. It can also store any type of structured or unstructured content such as XML documents, Web pages, photos, voice data, images, process definitions, and templates. Figure 8-3 shows how object stores interface with the IBM System Storage DR550.

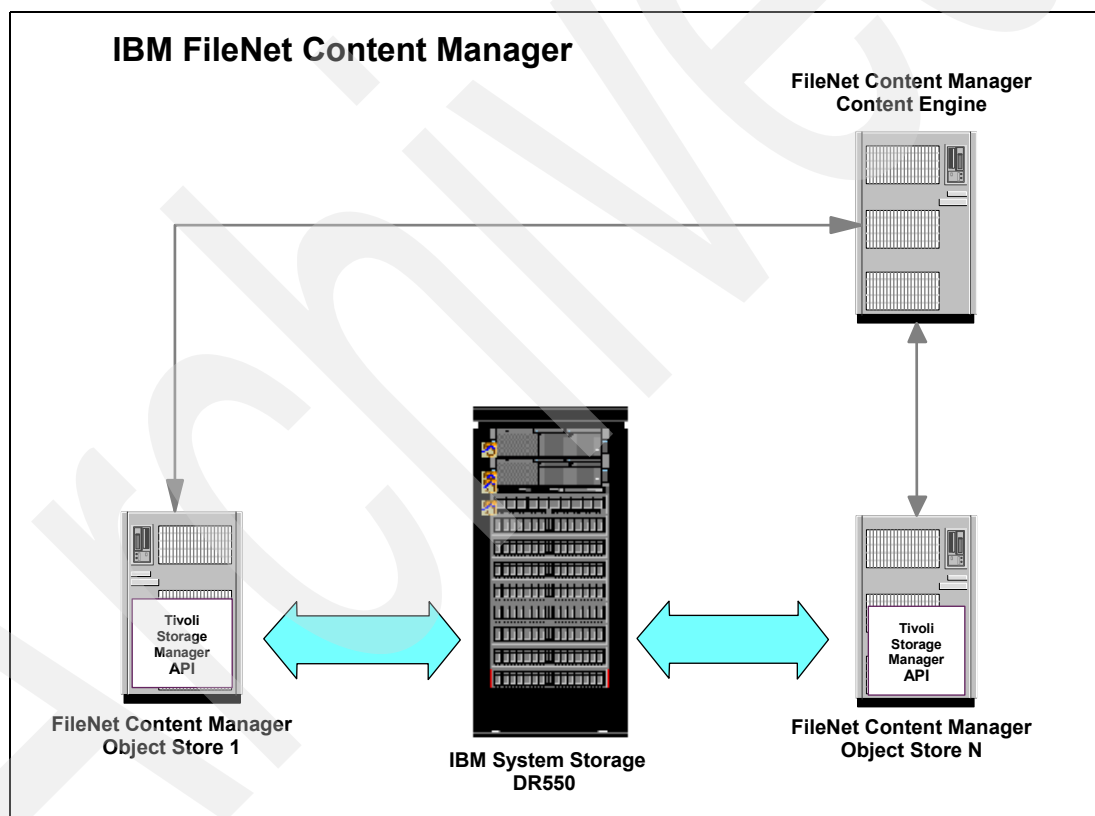


Figure 8-3 FileNet Content Manager object stores interfacing with Tivoli Storage Manager storage

## 8.3 Integrating Content Manager with DR550 SSAM Server

A Content Manager system contains a library server and one or more resource managers. The Content Manager resource manager relies on Tivoli Storage Manager for accessing secondary storage devices other than local file systems.

Starting with Content Manager Version 8.2 Fix Pack 3, the Content Manager resource manager provides support for the IBM System Storage DR550 in compliance mode, using the Tivoli Storage Manager application program interface (API). In this mode, an active retention protection ensures availability of objects, such as files, for a period of time, which can be determined by the administrator.

Also with Content Manager V8.2 Fix Pack 3, a single Content Manager resource manager can now manage volumes from different Tivoli Storage Manager servers. This function enables existing customers to have Tivoli Storage Manager volumes with and without retention protection on the same Content Manager resource manager.

**Important:** To use Content Manager in compliance mode in conjunction with the IBM System Storage DR550 offering, the software level of Content Manager must be Version 8.2 Fix Pack 3 or later and the IBM Tivoli Storage Manager API must be Version 5.2.2 or later. The Tivoli Storage Manager archive copy group must use event-based retention.

The following rules apply to a Content Manager environment set up for System Storage Archive Manager:

- ▶ You cannot migrate data out of Content Manager volumes that are Tivoli Storage Manager volumes under retention control.
- ▶ You cannot have more than one local Content Manager storage class in a Content Manager policy where the primary storage class contains a Tivoli Storage Manager volume that is under retention control.
- ▶ If the first Content Manager storage class in the Content Manager policy does not have a Tivoli Storage Manager volume under retention control, you can:
  - Have other storage classes. In that case, if you also have a storage class with a Tivoli Storage Manager volume under retention control, it must be the last storage class.
  - Have a remote storage class that contains a Tivoli Storage Manager volume under retention control.

There are no restrictions on Content Manager replication, because the source or target collections can have migration policies with a Tivoli Storage Manager volume under retention control.

### Configuring Content Manager for DR550 SSAM Server

You need to configure different entities within Content Manager before data will be archived in the Tivoli Storage Manager storage repository (see Figure 8-4 on page 311):

- ▶ You must have a Tivoli Storage Manager server, and the policies must include archive copy groups with retention values matching the retention requirements of the item types in Content Manager that will use Tivoli Storage Manager. The archive copy group must use event-based retention, because this is the only configuration Content Manager supports for Tivoli Storage Manager when the archive retention protection mode is on.
- ▶ You must register a node in that Tivoli Storage Manager policy domain.

- ▶ The Tivoli Storage Manager API software (Version 5.2.2 or later) must be installed and configured on the Content Manager resource server.
- ▶ Several options must be set in Content Manager to allow the system to use the System Storage Archive Manager server. To enable the compliance mode, you need to define a unique Content Manager storage group for Content Manager volumes that correspond to Tivoli Storage Manager volumes under retention control.

Depending on your retention requirements, the configuration of the different entities within the Content Manager and Tivoli Storage Manager can be very complex. We now describe and discuss some of the Content Manager concepts and constructs. We also provide an example to help illustrate these concepts.

In a Content Manager for Microsoft Windows system, the Content Manager System Administrator GUI is used to set parameters.

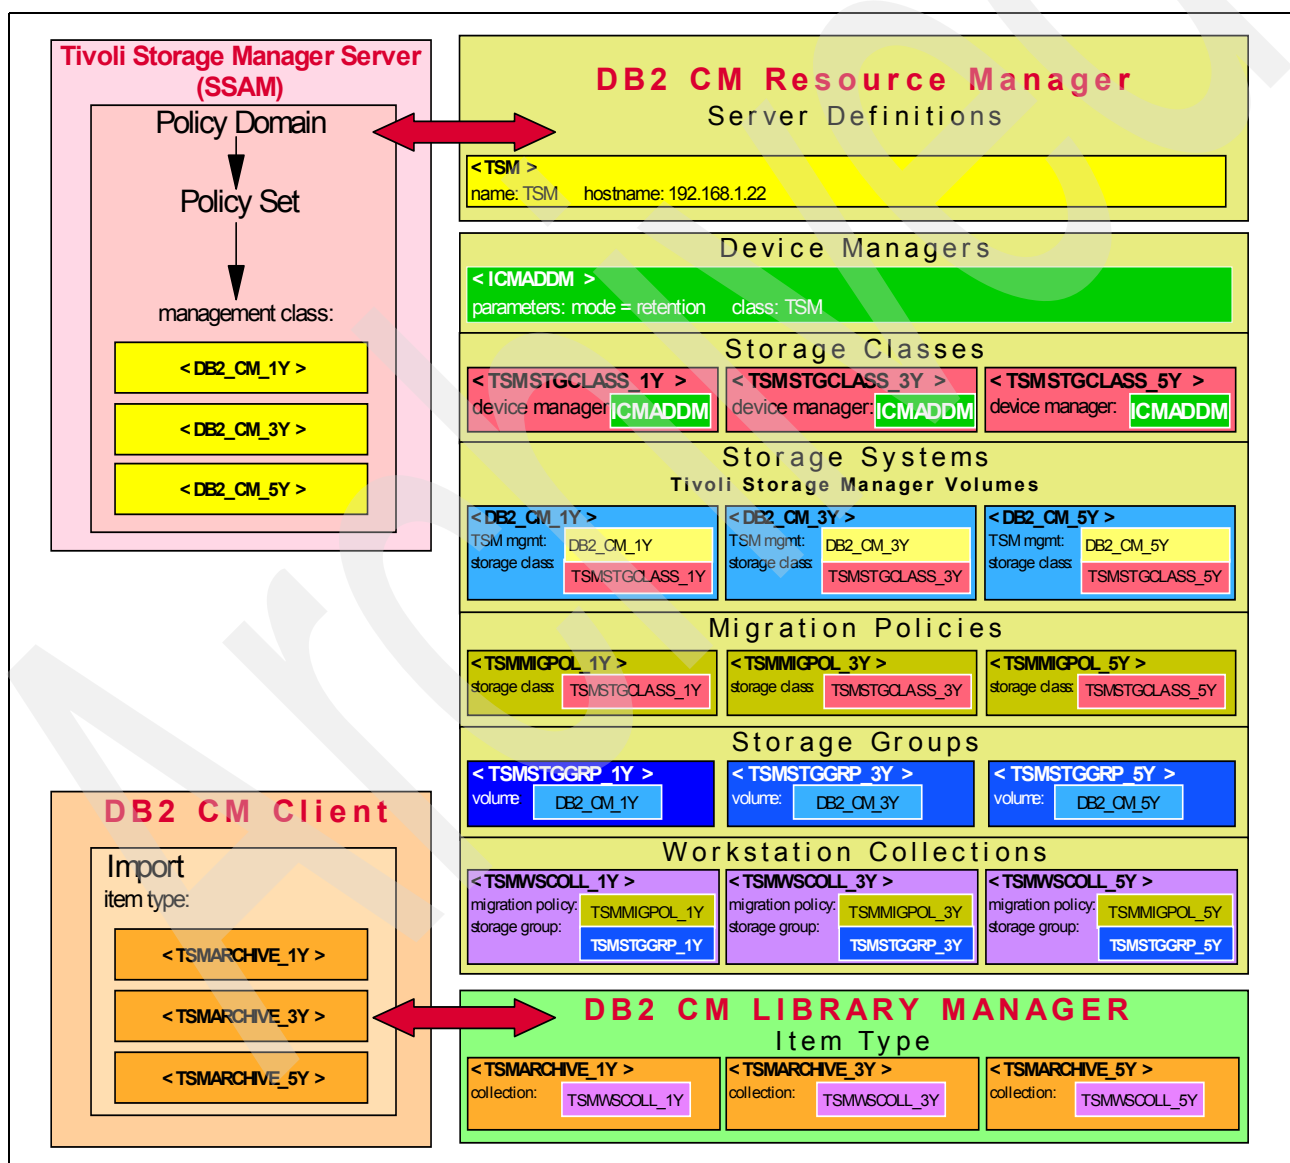


Figure 8-4 Overview: Content Manager for Tivoli Storage Manager archive management

Some important terms include:

<b>Device manager</b>	A software artifact that acts as an intermediary between your resource manager and physical storage. It is the interface between the resource manager and the storage system defined with it in a migration policy. It communicates the tasks that you define for the resource manager to the storage system where you store your objects. You assign device managers to a storage class so that the storage class can communicate with the storage systems.
<b>Storage class</b>	A logical grouping of similar storage types that identifies the type of media on which an object is stored. It is not directly associated with a physical location; however, it is directly associated with the device manager, which is the interface between the resource manager and the actual physical location. You can assign only one device manager to each storage class. Types of storage classes include <i>fixed disk</i> , <i>VideoCharger™</i> , <i>media archive</i> , and <i>Tivoli Storage Manager</i> .
<b>Storage system</b>	An actual physical device or unit where the objects are stored. There are different types of storage systems, such as volumes on Windows, file systems on UNIX, Content Manager VideoCharger, media archive, and Tivoli Storage Manager. Storage systems are also known as <i>volumes</i> . A storage system is associated with a storage class.
<b>Migration policy</b>	A user-defined schedule for moving objects from one storage class to the next. It describes the retention and class transition characteristics for a group of objects in a storage hierarchy. Creating a migration policy and defining the migrator schedule automates the migration of objects so that you do not have to manually monitor migration.

**Note:** Tivoli Storage Manager calls its migration policies *management classes*.

**Storage group** A storage group contains one or more storage systems and storage classes. It associates each storage system to a storage class.

## Content Manager for Windows and Tivoli Storage Manager configuration

We assume that the Content Manager V8.2 Fix Pack 3 (or later) software is installed and configured and that a Content Manager client is available for testing archive functions. The Tivoli Storage Manager server is located in the IBM System Storage DR550; therefore, the archive retention protection is set on, which makes it a System Storage Archive Manager server.

To enable Content Manager for Windows to access the SSAM server for archive management, complete the following steps on the Content Manager resource server, and then the Tivoli Storage Manager administrative command-line client, and finally, the Content Manager System Administration Client GUI, as outlined in the following sections.

### Content Manager resource server

First, on the Content Manager resource server, complete the following steps:

1. Download the current Tivoli Storage Manager backup-archive client, API, and the Tivoli Storage Manager administrative client command-line files (all packaged as one download file). You can find the current maintenance levels of the software at:

<ftp://ftp.software.ibm.com/storage/tivoli-storage-management/maintenance/client/v5r5/Windows/x32/v550>

Within the download folder, download the self-extracting executable client code. Refer to the readme.ftp file within the same folder as the code is named, for example, a file named TSMBAC-WinX32.exe.

2. Start the installation by starting the self-extracting executable client code, such as TSMBAC-WinX32.exe.
3. In the first window (Location to Save Files), choose a folder where the software can be unpacked, such as c:\tsm\_images\TSM\_BA\_Client, and click **Next**.  
The install wizard extracts the files.
4. In the Choose Setup Language window, choose your language, such as **English (United States)**, and click **OK**.  
The install wizard prepares the installation.
5. In the Welcome to the InstallShield Wizard window, click **Next**.
6. In the Destination Folder window, select the installation folder, such as c:\Program Files\Tivoli\TSM\, and then click **Next**.
7. In the Setup Type window, change the default setting from Typical to **Custom**, and then click **Next**.
8. In the Custom Setup window, select the **Administrative Client Command Line Files** and **Client API SDK Files** additional features (three are already selected), and then click **Next**. (See Figure 8-28 on page 332.) Although the administrative client command line is not necessary for the Content Manager, we use this interface to set up the Tivoli Storage Manager server and check the results of archive sessions. This step is optional and you do not need to install this product if you prefer to use the administrative Web client. The Client API SDK Files contain the dapismp command line for testing purposes. It might be useful to test the server connection to the DR550 SSAM Server with this tool.

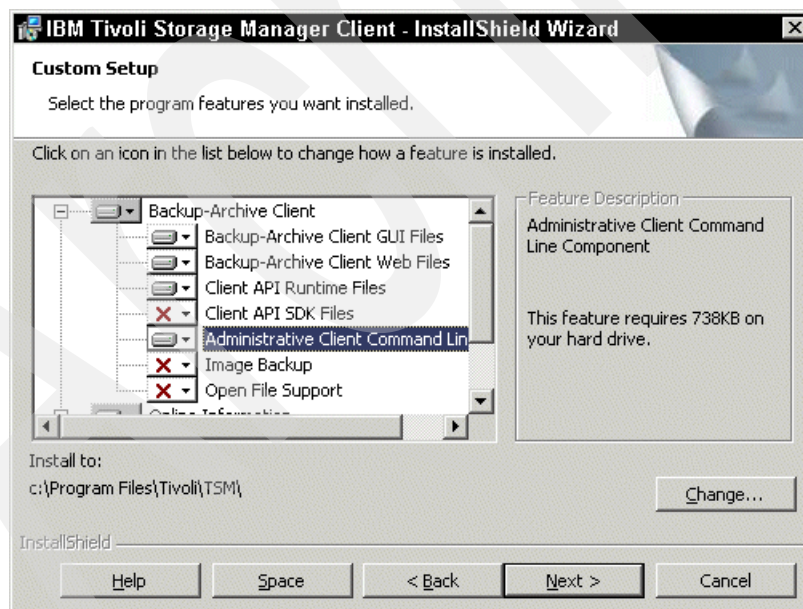


Figure 8-5 Custom Setup window

9. In the Ready to Install the Program window, click **Install**. The InstallShield Wizard starts installing the software.
10. When the InstallShield Wizard Completed window opens, check that the installation is successful and click **Finish**. If it is not successful, correct the problem and repeat the installation.



11. The API uses unique environment variables to locate files. Set up the API environment variables DSMI\_CONFIG, DSMI\_DIR, and DSMI\_LOG in Microsoft Windows (select **System Properties** → **Environment Variables**). It is a best practice to configure the variables for the entire system (system variables) rather than for a single user (user variables). See Figure 8-29 on page 332 for details.

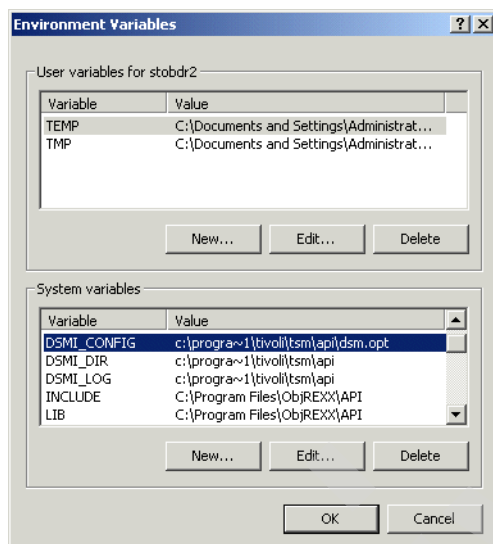


Figure 8-6 Set API Environment Variables window

12. Copy the dsm.opt file from the backup-archive client installation folder to the API installation folder. If there is no dsm.opt file, copy the dsm.smp sample option file from the Tivoli Storage Manager configuration directory (C:\Program Files\Tivoli\TSM\config) to the backup-archive client installation folder and to the API installation folder. Rename the sample option file from dsm.smp to dsm.opt in both folders.
13. Edit the dsm.opt file within the backup-archive client installation folder. Set the IP address of your System Storage Archive Manager server to (TCPServeraddress), commmethod tcpip, tcpport 1500, enablearchiveretentionprotection on, and passwordaccess generate. Save the changes. This step is optional, and you do not need to configure this file if you do not use the administrative command-line client.
14. Edit the dsm.opt file within the API client installation folder. Set the IP address of your System Storage Archive Manager server to (TCPServeraddress), commmethod tcpip, tcpport 1500, enablearchiveretentionprotection on, and passwordaccess prompt. Save the changes.

The Tivoli Storage Manager API access method “generate” is supported by Content Manager, but the resource manager first attempts to access Tivoli Storage Manager with “prompt”. If using prompt is not successful, it retries using generate. If you use generate, you need to use the Tivoli Storage Manager API sample program dapismp to change the password, which in turn, enables this feature.

**Tip:** You can configure Content Manager to signal Tivoli Storage Manager to use the retention mode instead of using the Tivoli Storage Manager parameter enablearchiveretentionprotection. To do this, in the Device Manager Properties window, configure your Tivoli Storage Manager device manager, ICMADDMM, and set Parameters to mode=retention. By using this configuration, you do not have to configure the Tivoli Storage Manager API options file with enablearchiveretentionprotection on.



### ***Tivoli Storage Manager administrative command-line client***

Next, use the Tivoli Storage Manager administrative command-line client to perform these steps:

1. With the administrative command-line client, first create a new Tivoli Storage Manager policy domain exclusively for Content Manager systems. The policy domain is named DB2\_CM\_PD, where the letters PD stand for policy domain. Create the new policy domain with the following Tivoli Storage Manager command:

```
define domain db2_cm_pd archretention=3650
```

This command creates the policy domain and sets the *archive retention grace period* to 3650 days, which is 10 times longer than the default. The grace period specifies the number of days to retain an archive copy when the management class for the file no longer contains an archive copy group and the default management class does not contain an archive copy group. The retention grace period protects archive copies from being immediately expired.

2. Within the policy domain DB2\_CM\_PD, we create one policy set named DB2\_CM\_PS, where the letters PS stand for policy set. Create the policy set by issuing the following Tivoli Storage Manager command:

```
define policyset db2_cm_pd db2_cm_ps
```

3. Create three different Tivoli Storage Manager management classes within the Tivoli Storage Manager policy set so that you can configure different retention policies. Because the plan is to archive some of your data for one year, some data for three years, and other data for five years, make sure to reflect that in your Tivoli Storage Manager management classes. The Tivoli Storage Manager management classes are named DB2\_CM\_1Y, DB2\_CM\_3Y, and DB2\_CM\_5Y. Use the following commands to create the three management classes:

```
define mgmtclass db2_cm_pd db2_cm_ps db2_cm_1y
define mgmtclass db2_cm_pd db2_cm_ps db2_cm_3y
define mgmtclass db2_cm_pd db2_cm_ps db2_cm_5y
```

Assign the first management class as the default by issuing the following command:

```
assign defmgmtclass db2_cm_pd db2_cm_ps db2_cm_1y
```

4. The next step is to define archive copy groups (type=archive) for each of the three management classes. The archive copy groups must be defined along with the correct parameters. First, they need to work with the event-based retention (RETINIT=event) and specify the retention values (RETMIN, RETVER) to reflect the different policies. In our example, the following Tivoli Storage Manager commands apply:

```
define copygroup db2_cm_pd db2_cm_ps db2_cm_1y type=archive
destination=archivepool retver=0 retinit=event retmin=365
define copygroup db2_cm_pd db2_cm_ps db2_cm_3y type=archive
destination=archivepool retver=0 retinit=event retmin=1095
define copygroup db2_cm_pd db2_cm_ps db2_cm_5y type=archive
destination=archivepool retver=0 retinit=event retmin=1825
```

See 5.2.3, “Event-based retention policy” on page 173 for a detailed description of the parameters.

5. Validate the policyset by issuing this command:

```
validate policyset db2_cm_pd db2_cm_ps
```

The command will return the information that the default management class does not have a backup copy group, and that files will not be backed up by default if this set is activated. Because the IBM System Storage DR550 is an archive-only solution, and indeed we want to archive Content Manager objects, you can ignore the message.

6. Activate the policyset with the following command:

```
activate policyset db2_cm_pd db2_cm_ps
```

7. After the successful definition of all policies, you can register a node in the newly created policy domain. Name the Content Manager resource manager cmarchive and register it in the DB2\_CM\_PD domain:

```
register node cmarchive password domain=db2_cm_pd archdelete=yes
```

### **Content Manager System Administration Client GUI**

Finally, in the Content Manager System Administration Client GUI, complete the following steps:

1. Start the Content Manager System Administration Client GUI for Windows and select **Content Manager** as the server type and select the instance you want to enable for Tivoli Storage Manager use, for example, **ICMNLSDDB**.
2. Log in with your user ID and password. A window similar to the one shown in Figure 8-7 opens.

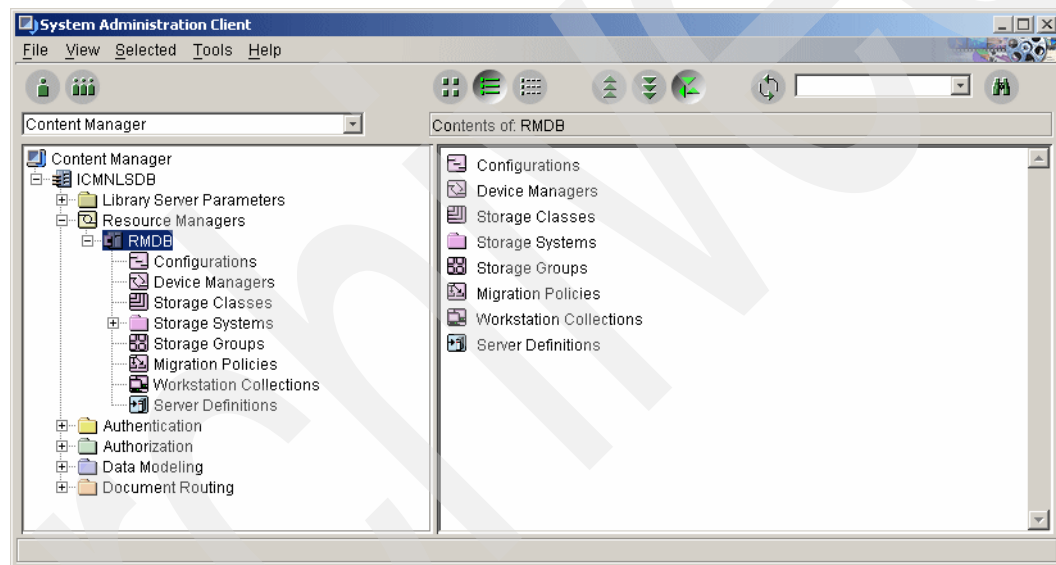


Figure 8-7 Content Manager System Administration Client

3. Click **Resource Managers** and then click **RMDb**. This will expand the tree of the resource manager database (RMDb), as shown in Figure 8-7. If your resource manager is not running, or there are problems in the communication between the library server and the resource manager, the message shown in Figure 8-8 will be displayed instead of an enlarged tree. Start the resource manager or correct the problems and click **RMDb** again.




Figure 8-8 Problems connecting to the resource manager database (RMDb)

4. Configure the resource manager as follows:

- a. Right-click **Server Definitions** in the left pane of the window and click **New**, as shown in Figure 8-9, to open the New Server Definition window (see Figure 8-10 on page 318). This is the general way to create new entries for all of the entities within the resource manager; therefore, we do not show this in detail again.

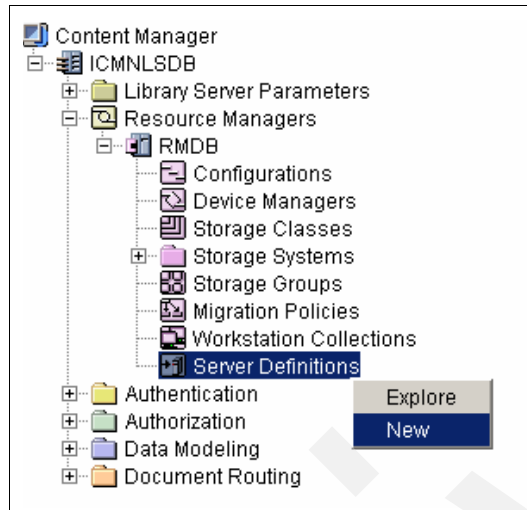
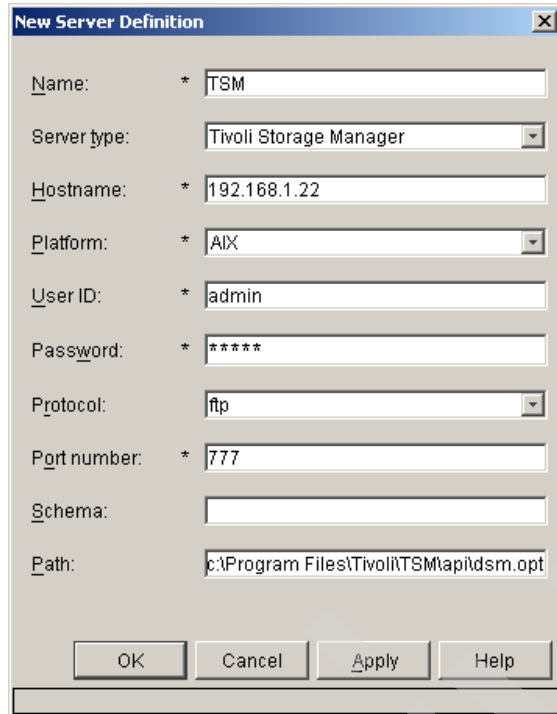


Figure 8-9 Create new Server Definitions for the resource manager

- b. In the Server Definition Properties window, specify the parameters that pertain to the DR550 System Storage Manager Server. With the exception of the values in the Hostname and the Password fields, the values shown in Figure 8-10 on page 318 can normally be used. Enter the host name of your DR550 SSAM Server (for example, the HACMP cluster service address for a dual-node configuration) and the password of the administrative user ID (admin) of your DR550 System Storage Archive Manager server. Select **ftp** from the Protocol drop-down list. Choose an arbitrary port number for the Port number field. Any port number will work for Tivoli Storage Manager. Leave the Schema field blank, but enter a fully-qualified path to the Tivoli Storage Manager API option file in the Path field (this is optional if you only use one Tivoli Storage Manager server).



The 'New Server Definition' window contains the following fields and values:

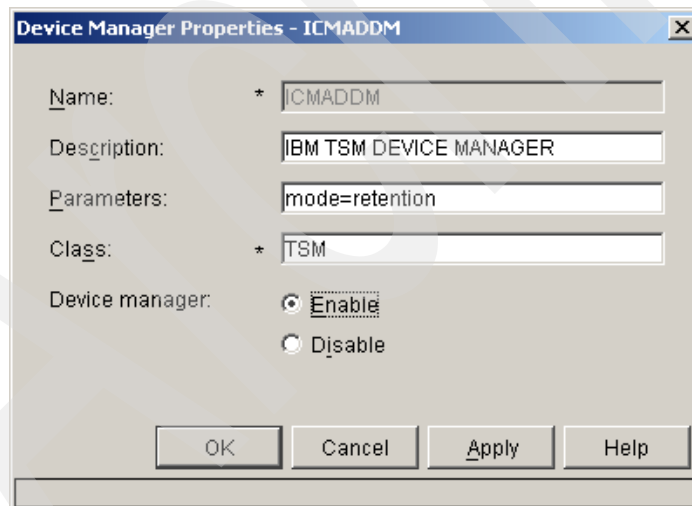
Field	Value
Name:	TSM
Server type:	Tivoli Storage Manager
Hostname:	192.168.1.22
Platform:	AIX
User ID:	admin
Password:	*****
Protocol:	ftp
Port number:	777
Schema:	
Path:	c:\Program Files\Tivoli\TSMapi\dsm.opt

Buttons at the bottom: OK, Cancel, Apply, Help.

Figure 8-10 New Server Definition window

Click **OK** to save the server information.

- c. Click **Device Managers** and then double-click **ICMADDM** in the right pane. This opens the Device Manager Properties window for ICMADDM, as shown in Figure 8-11.



The 'Device Manager Properties - ICMADDM' window contains the following fields and values:

Field	Value
Name:	ICMADDM
Description:	IBM TSM DEVICE MANAGER
Parameters:	mode=retention
Class:	TSM
Device manager:	<input checked="" type="radio"/> Enable <input type="radio"/> Disable

Buttons at the bottom: OK, Cancel, Apply, Help.

Figure 8-11 Device Manager Properties: ICMADDM window

In the Parameters field, type mode=retention and enable the device manager by selecting **Enable**. Click **OK** to save the information.

**Important:** Within Content Manager, you can configure the Tivoli Storage Manager device manager ICMADDMM to signal to Tivoli Storage Manager that archive protection is in use. Therefore, the Parameters field must contain mode=retention. If this is not set, you must enable the archive protection in the Tivoli Storage Manager API option file dsm.opt with ENABLEARCHIVERETENTIONPROTECTION ON. We recommend that you always set both parameters in your environment.

- d. Right-click **Storage Classes** and click **New** to open the New Storage Class window. In the Name field, type a meaningful name for your new storage class. Select **Local destination**, and select **ICMADDMM** as the Device manager. Click **OK** to save the storage class.




Figure 8-12 New Storage Class window

In our example, we created three storage classes named TSMSTGCLASS\_1Y, TSMSTGCLASS\_3Y, and TSMSTGCLASS\_5Y.

- e. Double-click **Storage Systems** to expand its contents. Right-click **Tivoli Storage Manager Volumes** and click **New** to open the New Tivoli Storage Manager Volume window.

Define your new Tivoli Storage Manager volume, but do not assign it at this time (Assignment: **Unassigned**). In the Tivoli Storage Manager management class field, type the Tivoli Storage Manager management class you want to use with this Content Manager storage system. Select the Server name and Storage class that you created before and that belong to the volume. See Figure 8-14 on page 320. Click **OK** to save the configuration.

When defining Tivoli Storage Manager volumes, Content Manager connects to the configured System Storage Archive Manager server. Therefore, the System Storage Archive Manager server should be available and configured for Content Manager at this time; otherwise, Content Manager will display an error message, as shown in Figure 8-13.



Figure 8-13 Tivoli Storage Manager configuration error message

**Important:** Always enter your Tivoli Storage Manager management class in uppercase. Refer only to Tivoli Storage Manager management classes that use the event-based archive retention.

Figure 8-14 shows an example of how to configure the first of three Tivoli Storage Manager volumes. Associate this volume with the Tivoli Storage Manager management class of one year retention (DB2\_CM\_1Y); this is the name resource manager gives to the volume. The storage class you created for this configuration is named TSMSTGCLASS\_1Y and it is referenced in the third line.

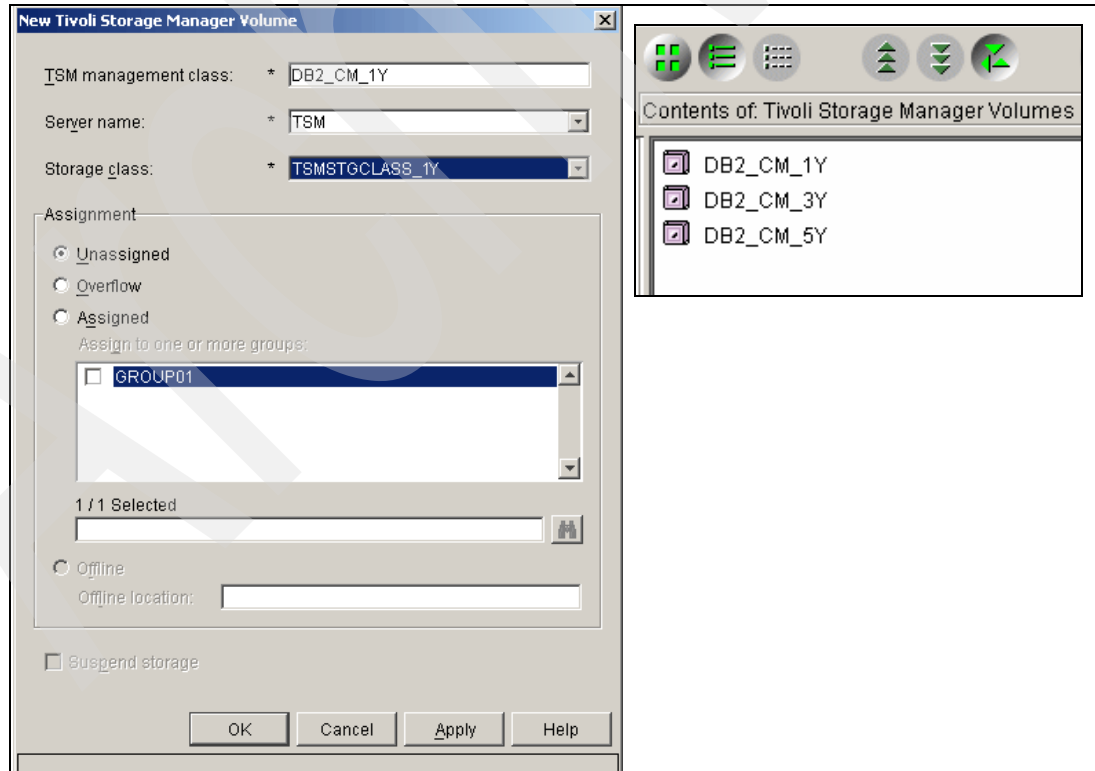


Figure 8-14 New Tivoli Storage Manager Volume window

Create three volumes in total (DB2\_CM\_1Y, DB2\_CM\_3Y, and DB2\_CM\_5Y) and assign the same Server name. Choose the appropriate Storage class each time.

The result should show three Tivoli Storage Manager volumes with names belonging to the Tivoli Storage Manager management classes, as shown on the right side of Figure 8-14 on page 320.

- f. Right-click **Storage Groups** and click **New** to open the Storage Group Properties window, as shown in Figure 8-15.

In the Name field, type the name you want to give to the new storage group, for example, TSMSTGGRP\_1Y.

The Storage systems list identifies the available storage systems. From this list, choose the storage system that you want to associate with this storage group. For example, choose the volume DB2\_CM\_1Y for the storage group TSMSTGGRP\_1Y.

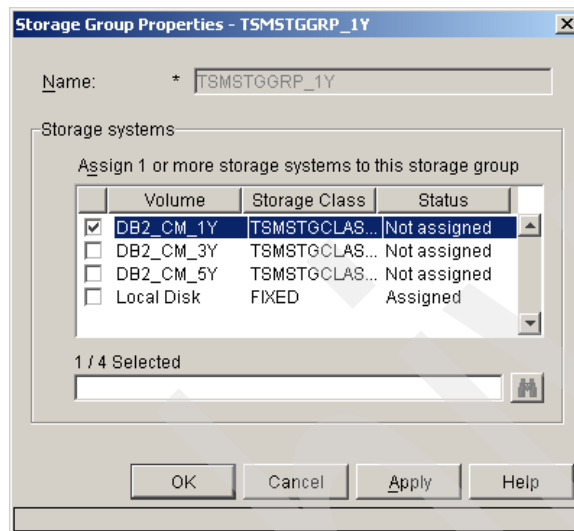


Figure 8-15 Storage Group Properties window

Click **OK** to save the configuration.

Create three storage groups (TSMSTGGRP\_1Y, TSMSTGGRP\_3Y, and TSMSTGGRP\_5Y) and assign the appropriate Tivoli Storage Manager volume each time. Only assign one volume to one storage group.

- g. Right-click **Migration Policies** and click **New** to open the New Migration Policy window, as shown in Figure 8-16.

In the Name field, type the name of the migration policy and click **Add**. The New Migration Policy Entry window opens. Select the correct Storage Class and the Retention period. Always select **Forever** as the Retention period.

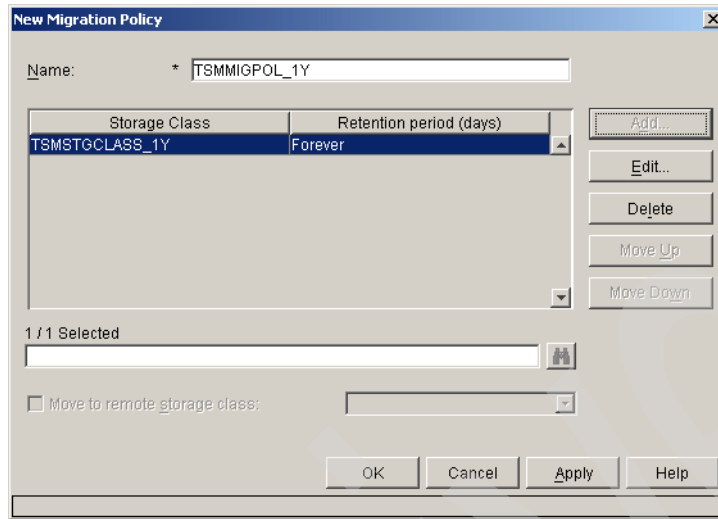


Figure 8-16 New Migration Policy window

Click **OK** to save the configuration.

Create three migration policies (TSMMIGPOL\_1Y, TSMMIGPOL\_3Y, and TSMMIGPOL\_5Y) and assign the appropriate Storage Class each time.

- h. Right-click **Workstation Collections** and click **New** to open the New Workstation Collection window, as shown in Figure 8-17 on page 323.

In the Name field, type a unique name for your workstation collection, for example, TSMWSCOLL\_1Y. In the Migration policy field, select the dedicated migration policy you want to use, for example, **TSMMIGPOL\_1Y**, and the Resource Manager will automatically fill in the Storage group field, in this case, with TSMSTGGRP\_1Y.

You can replicate objects in this collection to several other collections that are on different resource managers. Because we only have one resource manager in our environment, we do *not* use the **Add** button, but save the configuration instead.






Figure 8-17 New Workstation Collection window

Click **OK** to save the configuration.

Create three workstation collections (TSMWSCOLL\_1Y, TSMWSCOLL\_3Y, and TSMWSCOLL\_5Y) and assign the appropriate Migration policy and Storage group each time.

##### 5. Configure the library server.

The Content Manager library server can be used for different operations and therefore has a variety of entities to configure. We concentrate on the item type only, because this is the only entity we need in our environment. This might be different in your production environment.

An *item type* is a template that consists of a root component, zero or more child components, and a classification. By classifying the item type, you make a judgement about the purpose of the items created using this item type. The classifications are item, resource item, document, and document part.

The following example shows you how to create document item types. The Content Manager client applications require that each document item type has a base part. Typically, document item types have ICMBASE (base part), ICMANNOTATION (graphical annotations that overlay the base part), and ICMNOTELOG (separate textual comments). There are additional parts (ICMBASETEXT and ICMBASESTREAM) available:

**ICMANNOTATION** Contains additions to, or commentary about, the main data; following the document metaphor, annotations include sticky notes, color highlights, stamps, and other graphical annotations in the text of a document. These are the typical annotation parts from previous releases of Content Manager. Using the Client for Windows or the eClient, users can create graphical annotations, which are viewed on top of the file or document being displayed. Most client applications can show or hide these annotations.

**ICMBASE** Contains the fundamental content of a document item type that stores any non-textual type of content, including image and audio. Requirement: To be viewable in the eClient, all document item types must include at least one base document part.

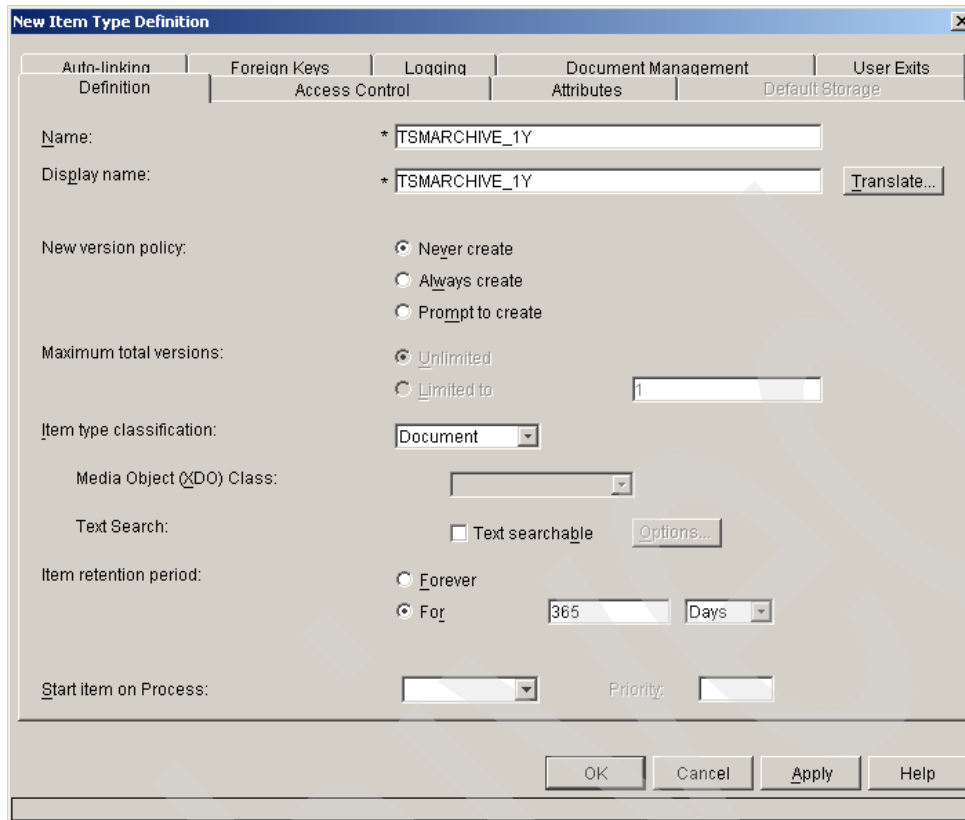
- ICMBASETEXT** Contains the fundamental content of a document item type that stores text content. If you plan to index a text part of your document, you should store the part in this part item type. Indexing a text part enables a text search to be performed on the content of the part.
- ICMNOTELOG** Contains a log of information entered by users, for example, indicating the reason that the insurance application was denied or instructions to the next reviewer of the document. These are the typical notelog parts from previous releases of Content Manager. Using the Client for Windows or eClient, your users can create, view, and edit notelog parts. Notelog parts contain the user ID, time stamp, and text comments as entered by client users.
- ICMBASESTREAM** Contains streamed data, such as video.

To configure the library server:

- a. Expand **Data Modeling** in the system administration tree.
- b. Right-click **Item Types** and click **New** to open the New Item Type Definition window, as shown in Figure 8-18 on page 325:
  - i. On the Definition page, in the Name field, type a meaningful name. Item type names are case-sensitive and must be unique. Use names that are easy to remember and that reflect the folders and documents are included in item type.

**Note:** The item type names in our example reflect the use of Tivoli Storage Manager and the retention period. This might not be useful in your environment, and you might prefer names that reflect the folders and documents that are included.

- ii. Click **Translate** to open the Translate Display Name window. All of the available languages defined in the system are listed. In the Translated Name column, type the translated display name for the other languages.  
Click **OK** to save the information.
  - iii. In the New version policy field, select **Never create**. In the Item type classification list, specify the new item type as **Document**. In the Item retention period field, select the retention period for the item. This number is the expiration date calculated by the library server when an item is created. See Figure 8-18 on page 325 for other settings.



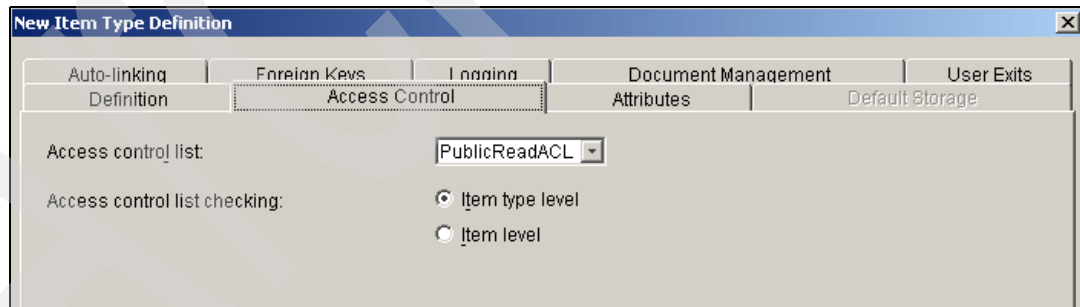
The 'New Item Type Definition' window is shown with the 'Definition' tab selected. The window has a title bar with a close button. Below the title bar are five tabs: 'Auto-linking', 'Foreign Keys', 'Logging', 'Document Management', and 'User Exits'. The 'Definition' tab is active, showing the following fields and options:

- Name:** \*TSMARCHIVE\_1Y
- Display name:** \*TSMARCHIVE\_1Y (with a 'Translate...' button)
- New version policy:**
  - ☒ Never create
  - ☐ Always create
  - ☐ Prompt to create
- Maximum total versions:**
  - ☒ Unlimited
  - ☐ Limited to: 1
- Item type classification:** Document (dropdown)
- Media Object (XDO) Class:** (empty dropdown)
- Text Search:**
  - ☐ Text searchable (with an 'Options...' button)
- Item retention period:**
  - ☐ Forever
  - ☒ For: 365 Days (dropdown)
- Start item on Process:** (empty dropdown)
- Priority:** (empty dropdown)

At the bottom of the window are four buttons: 'OK', 'Cancel', 'Apply', and 'Help'.

Figure 8-18 New Item Type Definition window: Definition tab

- iv. Click the **Access Control** tab. On the Access Control page, in the Access control list field, select **PublicReadACL**. In the Access control list checking field, specify whether the access control list applies to the item type level or item level. For example, choose **Item type level**. See Figure 8-19.



The 'New Item Type Definition' window is shown with the 'Access Control' tab selected. The window has a title bar with a close button. Below the title bar are five tabs: 'Auto-linking', 'Foreign Keys', 'Logging', 'Document Management', and 'User Exits'. The 'Access Control' tab is active, showing the following fields and options:

- Access control list:** PublicReadACL (dropdown)
- Access control list checking:**
  - ☒ Item type level
  - ☐ Item level

Figure 8-19 New Item Type Definition window: Access Control tab

- v. Click the **Attributes** tab. On the Attributes page, select the attributes or attribute groups that you want to add into the item type from the Available attributes or groups list. Click **Add** to add them to the Selected attributes and components list. See Figure 8-20 for an example.

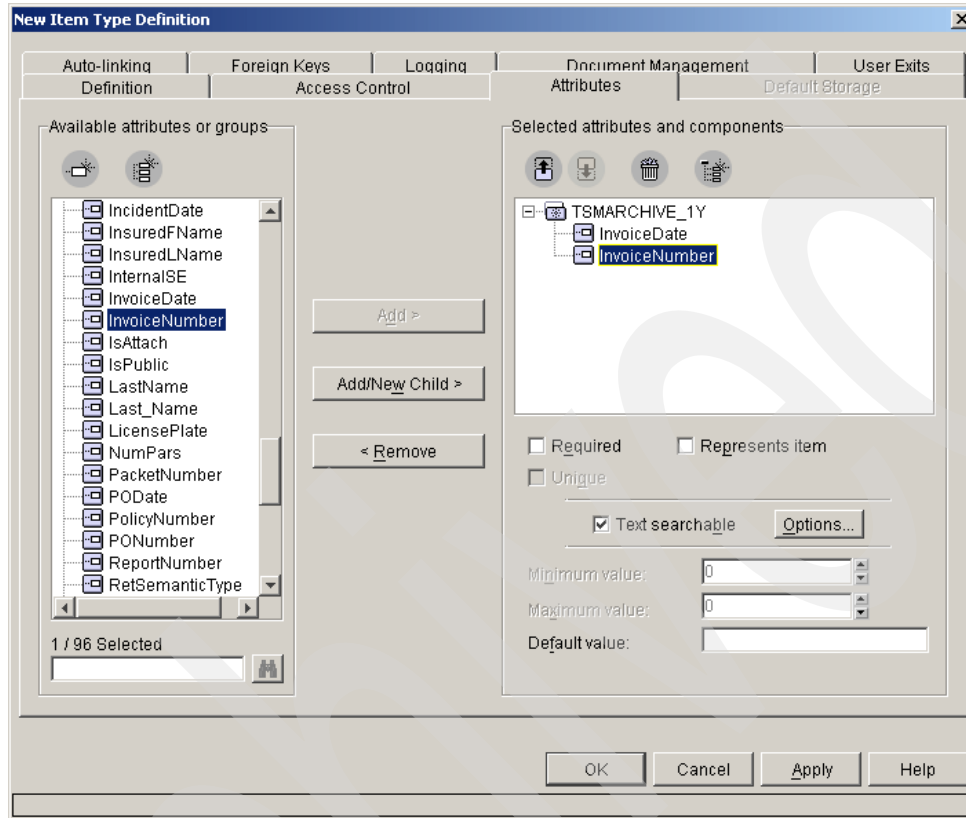


Figure 8-20 New Item Type Definition window: Attributes tab

In our example, the use of the Auto-linking, Foreign Keys, Logging, and User Exits tabs is optional. You should check if this is also true for your environment.

- vi. Click the **Document Management** tab. On the Document Management page, click **Add** to open the Define Document Management Relations window, as shown in Figure 8-21 on page 327. In the Part type field, select a first part (**ICMANNOTATION**) to associate with the document item type. From the Access control list drop-down list, select an access control list (**PublicReadACL**) to associate with the part type. In the Resource manager field, select the resource manager (**RMDB**) on which the part type is stored. In the Collection field, select the collection (**TSMWSCOLL\_1Y**) on which the part is stored. In the New version policy field, specify a version policy (**Never create**) for the part type. Click **Apply** to apply the first document management relation.




Figure 8-21 Define Document Management Relations window

- vii. In the Part type field, select a second part (**ICMBASE**) to associate with the document item type. From the Access control list, select an access control list (**PublicReadACL**) to associate with the part type. In the Resource manager field, select the resource manager (**RMDB**) on which the part type is stored. In the Collection field, select the collection (**TSMWSCOLL\_1Y**) on which the part is stored. In the New version policy field, specify a version policy (**Never create**) for the part type. Click **Apply** to apply the second document management relation.
- viii. In the Part type field, select a third part (**ICMNOTELOG**) to associate with the document item type. From the Access control list, select an access control list (**PublicReadACL**) to associate with the part type. In the Resource manager field, select the resource manager (**RMDB**) on which the part type is stored. In the Collection field, select the collection (**TSMWSCOLL\_1Y**) on which the part is stored. In the New version policy field, specify a version policy (**Never create**) for the part type. Click **OK** to apply the third document management relation and to close the window. See Figure 8-22 for the results.

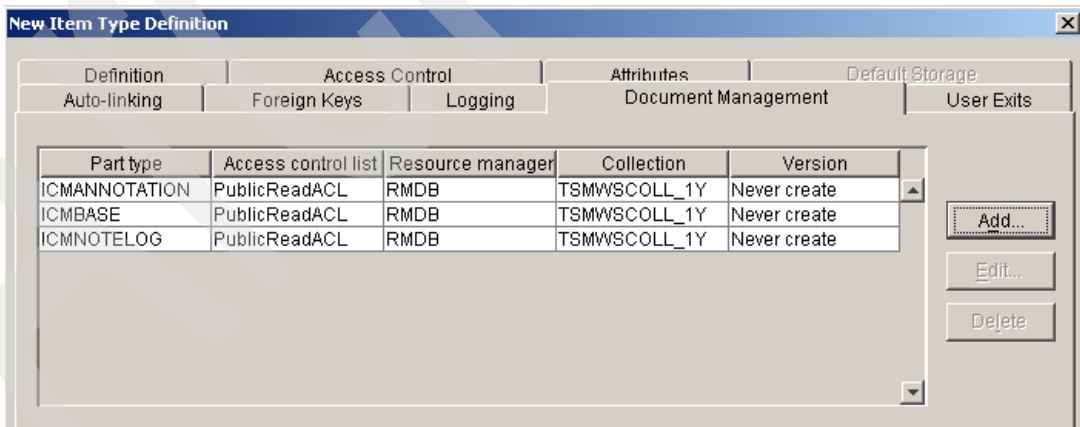


Figure 8-22 New Item Type Definition window: Document Management tab

Click **OK** at the bottom of the New Item Type Definition window. This saves the configuration of the new item type.

- c. Repeat the procedure described before to create two more item types (TSMARCHIVE\_3Y, TSMARCHIVE\_5Y) with the appropriate settings. The library server now contains three item types created for archive purposes, as shown in Figure 8-23. The three item types are associated with the DR550 SSAM Server as a storage unit, and they provide archive retentions of one year, three years, and five years.




Figure 8-23 Data Modeling: Item Types

6. Use the Content Manager Client for Windows for testing:
  - a. Start the Content Manager Client for Windows.
  - b. In the Welcome window, in the Server field, select the library server to which you want to connect, for example, **ICMNLSDDB**. In the User ID and Password fields, you must provide a user with the authority to import and search data on the library server. For example, use icmadmin as a user.
  - c. After successful login, close the Welcome window.
  - d. Go to **File** → **Import** to open the Import window. In the Import window, click **Add Files to Import** and select the files from the list. Use the buttons in the upper part of the window to navigate to the folder where the data can be found and click one of the files you want to import. The file will be displayed in the File name field. Click **Open**. The window closes and returns you to the Import window. The Import window now displays the selected file in the Files to be imported field. In the File Type window, select the type of file you selected before. Select the item type of the file at the Item Type window, for example, select **TSMARCHIVE\_1Y** for the first import test. Fill out the other fields with meaningful information. See Figure 8-24 on page 329.

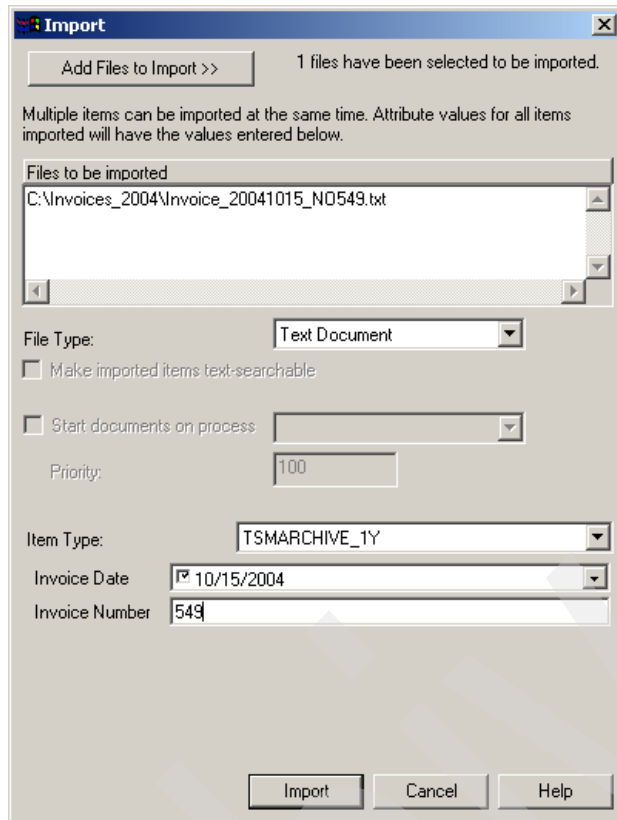


Figure 8-24 Content Manager Client: Import window

Click **Import** to import the selected file. The Content Manager Client starts importing the file and shows the progress in an import progress window, as shown in Figure 8-25.




Figure 8-25 Content Manager Client: Import progress window

- e. Repeat the above procedure twice to import two more files. Select different files each time. For the first file, in the Item Type field, select **TSMARCHIVE\_3Y**. For the second file, in the Item Type field, select **TSMARCHIVE\_5Y**.

- f. Go to **Search** → **Basic** to open the Basic Search window. In the Item Type field, select the item type **TSMARCHIVE\_1Y** and use search parameters for a general search. See Figure 8-26 for an example of search parameters.

Figure 8-26 Content Manager Client: Basic Search window

Click **OK** to start the search.

The Search Results window displays the search results, as shown in Figure 8-27.

Invoice Date	Invoice Number
10/15/2004	549
10/19/2004	548
10/19/2004	546

Figure 8-27 Content Manager Client: Search Results window

You can double-click the results to verify that the retrieve from storage is correct.

## 8.4 Integrating Content Manager OnDemand with DR550

A Content Manager OnDemand system contains a library server and one or more object servers. The object server stores data objects in its cache file systems, which can be defined on locally attached or SAN-attached storage. The object server also supports archive storage systems. The UNIX and Windows platforms' OnDemand object server supports Tivoli Storage Manager as their archive repository and uses the Tivoli Storage Manager API to communicate with and transfer data objects to archive storage. When data is loaded into the OnDemand system, OnDemand creates objects, which hold the compressed data and store it in its cache file systems. These objects can also be archived to Tivoli Storage Manager at the time the data is loaded into OnDemand, or after the objects have been stored in the OnDemand cache storage for a predetermined amount of time. This hierarchical use of storage is useful for storing data on fast access devices such as disk (online) during the time of the highest likelihood of access to the data and then migrating to archive storage.



## Configuring OnDemand for Tivoli Storage Manager archive management

There are several steps that you need to complete to enable Content Manager OnDemand to use the Tivoli Storage Manager server as an archive manager:

1. You must have a Tivoli Storage Manager server, and the policies must include archive copy groups with retention values coinciding with the retention requirements of the application groups in OnDemand that will use Tivoli Storage Manager.
2. You must register a node in that Tivoli Storage Manager policy domain.
3. The Tivoli Storage Manager API software must be installed and configured on the OnDemand object server.
4. Several options must be set in OnDemand to allow the system to use Tivoli Storage Manager.

In an OnDemand for Windows system, the OnDemand configurator is used to set this parameter so that you can use the Tivoli Storage Manager server as an archive manager. In an OnDemand UNIX-based system, the `ars.cfg` configuration file is updated to specify that Tivoli Storage Manager is to be used.

## OnDemand for Windows and Tivoli Storage Manager configuration

To enable OnDemand to access a Tivoli Storage Manager server for archive management, you must complete the following steps:

1. Perform these steps on the object server:
  - a. Download the most current Tivoli Storage Manager backup-archive client, API, and the Tivoli Storage Manager administrative client command-line files (all within one download file) from the Internet. You can find the most current maintenance levels of the software at:  
<ftp://ftp.software.ibm.com/storage/tivoli-storage-management/maintenance/client/v5r5/Windows/x32/v550>  
Within the download folder, download the self-extracting executable client code. Refer to the `readme.ftp` file within the same folder to learn what the code is named, for example, a file named `TSMBAC-WinX32.exe`.
  - b. Start the installation by starting the self-extracting executable client code, such as `TSMBAC-WinX32.exe`.
  - c. Within the first window (Location to Save Files), choose the folder where the software can be unpacked, such as `c:\tsm_images\TSM_BA_Client`, and click **Next**.  
The install wizard extracts the files.
  - d. In the Choose Setup Language window, choose the language to be used during installation, such as **English (United States)**, and click **OK**.  
The install wizard prepares the installation.
  - e. In the Welcome to the InstallShield Wizard window, click **Next**.
  - f. In the Destination Folder window, select the installation folder, such as `c:\Program Files\Tivoli\TSM\`, and then click **Next**.
  - g. In the Setup Type window, change the default setting from Typical to **Custom**. Click **Next**.

- h. In the Custom Setup window, select the **Administrative Client Command Line Files** additional feature (three are already selected), and then click **Next**. (See Figure 8-28.)

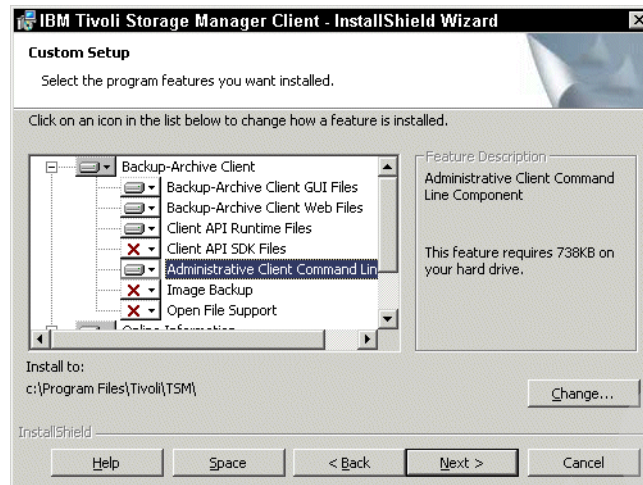


Figure 8-28 Custom Setup window

- i. When you are ready to install, click **Install** in the Ready to Install the Program window. The InstallShield Wizard begins to install the software.
- j. In the InstallShield Wizard Completed window, check that the installation is successful and click **Finish**. If it is not successful, correct the problem and repeat the installation.
- k. The API uses unique environment variables to locate files. Set up the API environment variables DSMI\_CONFIG, DSMI\_DIR, and DSMI\_LOG within Microsoft Windows after selecting **System Properties** → **Environment Variables**. See Figure 8-29 for details.

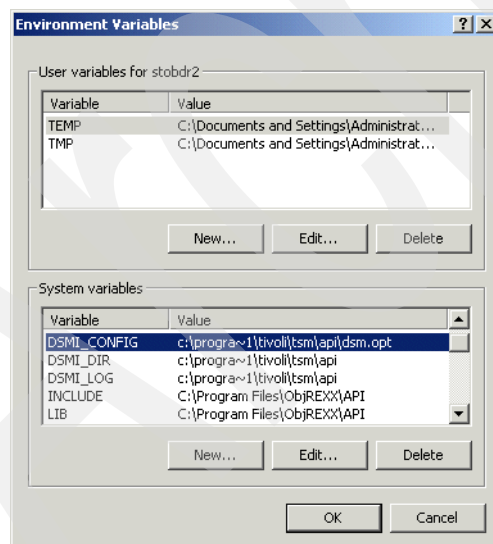


Figure 8-29 Set API Environment Variables window

- l. Copy the dsm.opt file from the backup-archive client installation folder to the API installation folder. If there is no dsm.opt file, copy the dsm.smp sample option file from the Tivoli Storage Manager configuration directory (C:\Program Files\Tivoli\TSM\config) to the backup-archive client installation folder and to the API installation folder. Rename the sample option file from dsm.smp to dsm.opt in both folders.

- m. Edit the `dsm.opt` file in the backup-archive client installation folder. Set the IP address of your Tivoli Storage Manager server with `(TCPServeraddress)`, `commethod tcpip`, `tcpport 1500`, and `passwordaccess generate`. Save the changes.
  - n. Edit the `dsm.opt` file in the API client installation folder. Set the IP address of your Tivoli Storage Manager server with `(TCPServeraddress)`, `commethod tcpip`, `tcpport 1500`, and `enablearchiveretentionprotection yes`. Save the changes.
2. In the Storage Manager administrative command-line client, we named the OnDemand storage node `odarchive` and registered it in the standard domain:
 

```
register node odarchive password domain=standard
```
  3. Use the OnDemand Configurator for these steps:
    - a. Start the OnDemand for Windows configurator and then select **Instances**. Click the `instance_name` of the instance you want to enable for Tivoli Storage Manager use.
    - b. Select the **Storage** tab.
    - c. In the Configuration area at the top of the Storage tab, select the **TSM** option.
    - d. After selecting TSM, click **TSM Options**. Enter the path to the Tivoli Storage Manager program files directory of the Tivoli Storage Manager API and the path to the Tivoli Storage Manager options `dsm.opt` file, as shown in Figure 8-30. Click **OK**. On the Storage tab, click **Apply**.

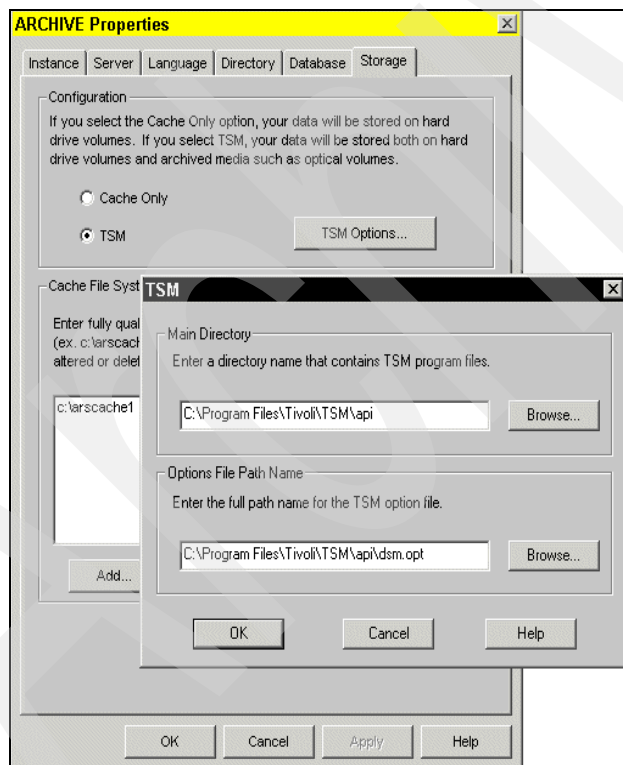


Figure 8-30 OnDemand for Windows configuration

- e. You will see a warning stating that the OnDemand services must be restarted for the changes to take effect, as shown in Figure 8-31.




Figure 8-31 Updating the OnDemand instance

4. Use the OnDemand Administrator for these steps:
  - a. Start the OnDemand Administrator client by selecting **Start** → **Programs** → **IBM OnDemand32** → **OnDemand Administrator**. Log on to the OnDemand server.
  - b. Navigate to the Storage Sets icon and select the storage set that you want to update. In our case, we chose the storage set **Library Server**.
  - c. Right-click and select **Update** storage set.

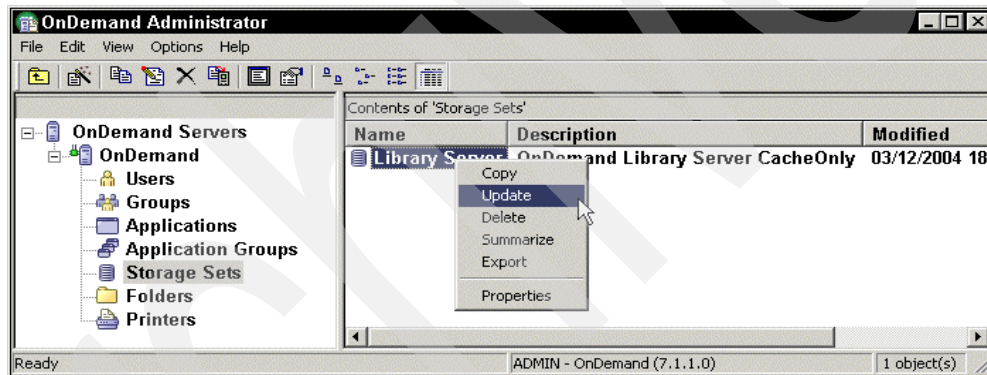


Figure 8-32 Update the storage set

- d. On the next window, choose the primary object server **\*ONDEMAND**, and click **Update** to update the primary object server named Library Server, as shown in Figure 8-33 on page 335. This brings you to the Update a Primary Node window.

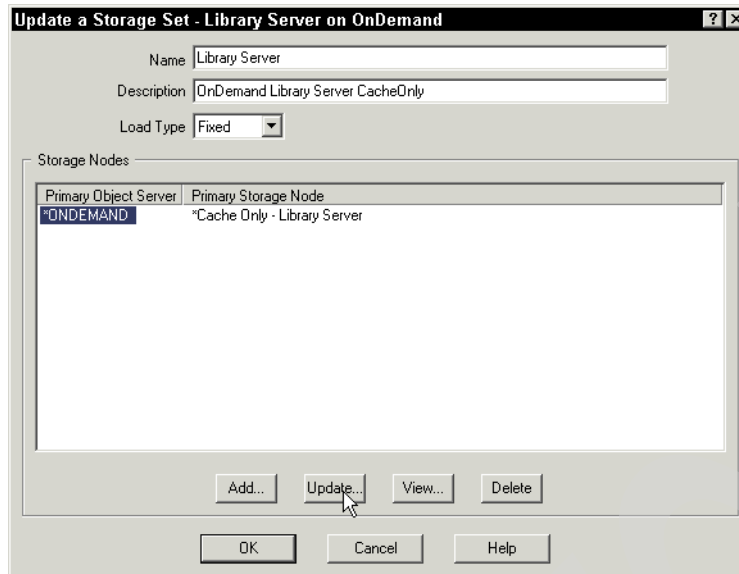


Figure 8-33 Update a Storage Set window

- e. From the Update a Primary Node window (see Figure 8-34):
  - i. Clear the **Cache Only** check box.
  - ii. In the Logon field, enter the Tivoli Storage Manager node name that you registered with the SSAM server; see 8.3, “Integrating Content Manager with DR550 SSAM Server” on page 310.
  - iii. In the Password field, enter the password you entered when registering the node to Tivoli Storage Manager and verify the password.
  - iv. You can update the Description field to reflect that this is no longer a cache-only primary storage node.
  - v. Select **OK** in the Update a Primary Node window.
  - vi. Now, you can update the description of the storage to reflect that this is no longer a cache-only storage set. Then, select **OK** in the Update a Storage Set window (see Figure 8-33).

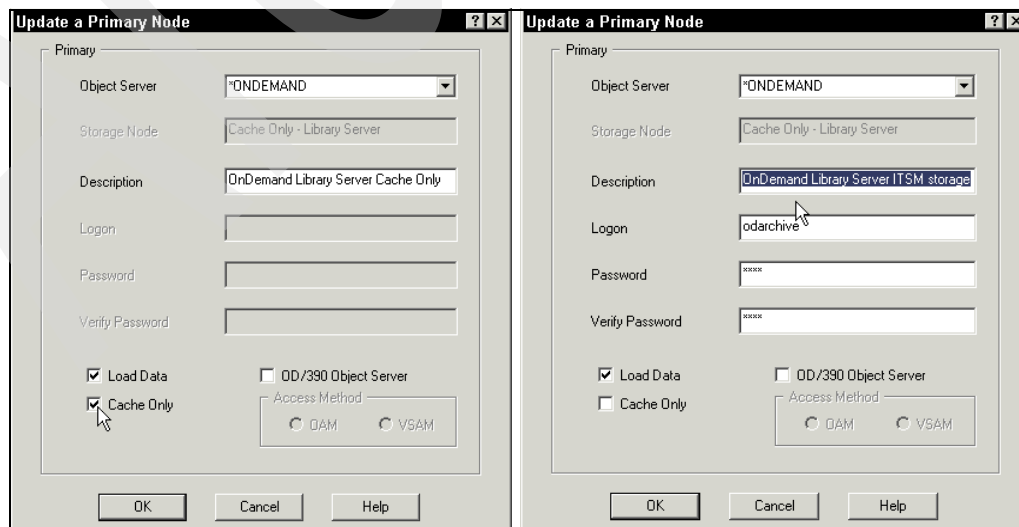


Figure 8-34 Update Primary Node windows

- f. This storage set is now able to store objects to the SSAM server. You now need to create or update an application group to use the new settings.
5. Use the OnDemand Administrator for these steps:
  - a. Navigate to the Application Groups icon and select the application group that you want to update. In our case, we chose the application group **jpeg1**.
  - b. Right-click and select **Update**, as shown in Figure 8-35.

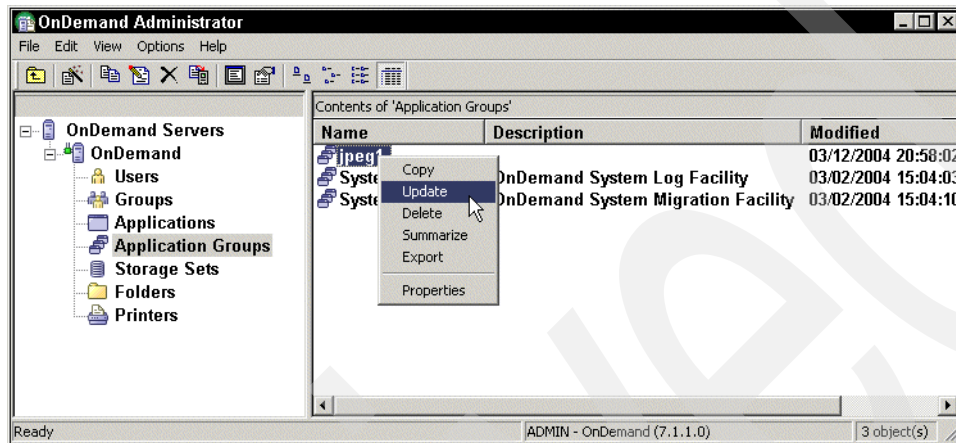


Figure 8-35 Update an Application Group

- c. Select the **Storage Management** tab from the Update an Application Group window. From the Storage Set Name list, choose the name of the storage set you updated in the previous steps (Figure 8-36 on page 337).
- d. Set the Cache Data values. The cache data setting determines if the report data is stored in the DASD cache, and if so, how long it is kept in cache before it expires. You can also choose to have the cache searched or not searched when retrieving documents for viewing. If you choose not to store reports in cache, a storage set that supports archive storage must be selected.
- e. The Life of Data and Indexes values determine when OnDemand can delete reports, resources, and index data from the application group. Choose from:
  - Never expires: OnDemand maintains application group data indefinitely.
  - Expires in \_\_\_ Days: After reaching this threshold, OnDemand can delete data from the application group. The default value is 2555 (seven years). The maximum value that you can type is 99999 (273 years).

**Important:** If you plan to maintain application group data in archive storage, the length of time that the archive storage manager maintains the data must be equal to or exceed the value that you specify for the Life of Data and Indexes fields. Consult the *IBM Content Manager OnDemand for Multiplatforms: Administration Guide*, SC18-9237 for more information.

- f. Do not select the Cache Data option. Click the **Advanced** button.



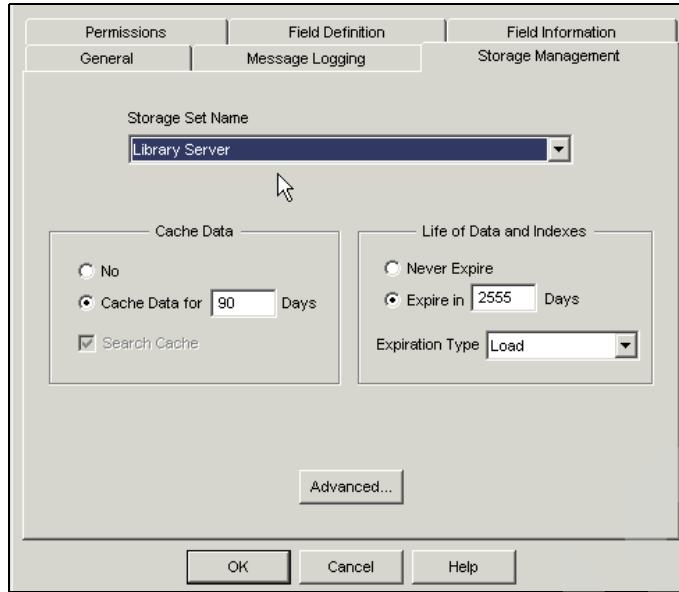


Figure 8-36 Update an application group storage management

- g. In the Advanced Storage Management window, choose when you want to have data objects migrated from the OnDemand *cache file system* to the System Storage Archive Manager server. If you leave the **When Data is Loaded** option selected, each time data is loaded by the OnDemand applications into OnDemand, the objects are stored in the cache file system and to Tivoli Storage Manager archive storage at the same time. This configuration setting has the advantage that if the cache file system of this OnDemand object server is damaged (disk failure), the objects are still accessible from the Tivoli Storage Manager storage.

6. Migrate data from cache.

This determines when documents and resources are migrated to archive storage:

- a. A storage set associated with a Tivoli Storage Manager client node must be selected to enable migration to archive storage. See Figure 8-37 on page 338.

The possible values are:

- No: Data is never migrated from cache. This option is unavailable when a storage set associated with a Tivoli Storage Manager client node is selected for the application group.
- When Data is Loaded: Data is migrated to archive storage when the data is loaded into the application group.
- Next Cache Migration: Data is migrated to archive storage the next time that ARSMANT is run with the -m option. The -m option indicates that data and resources are to be copied from cache to archive storage.
- After \_\_\_ Days in Cache: Specifies the number of days that data is to remain in cache-only storage. After reaching the prescribed number of days in cache storage, the data is copied to archive storage the next time that ARSMANT is run with the -m option for data migration.

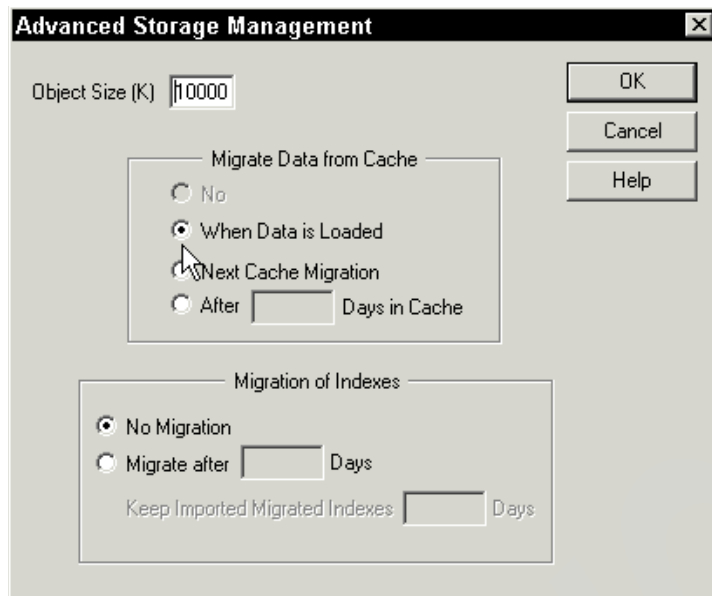


Figure 8-37 Advanced Storage Management window

- b. Click **OK** in the Advanced Storage Management window, and **OK** in the **Storage Management** tab of the application group.

You should now be able to load data using an application in the application group that we updated. This data should be migrated to the SSAM server and stored in the OnDemand cache file system.

Figure 8-38 and Figure 8-39 on page 339 show the **load** command used from the OnDemand command window to successfully load data with the generic indexer and the output of the select statement used to query the Tivoli Storage Manager database after the load that shows the object was archived to Tivoli Storage Manager. In this case, it was a SSAM server.

```
C:\Program Files\IBM\OnDemand for WinNT\bin>arsadmin load -g jpeg1 -u admin -p o
ndemand -i c:\arsload\gen.txt -d c:\arsload -h ondemand
OnDemand Load Id = >5014-1-0-4FAA-0-0<
Loaded 1 rows into the database
Document compression type used - 0D77. Bytes Stored = >9929<

C:\Program Files\IBM\OnDemand for WinNT\bin>
```

Figure 8-38 Load data to OnDemand with generic indexer, migrate to Tivoli Storage Manager



```

NODE_NAME: ODARCHIVE
FILESPACE_NAME: \CAA
FILESPACE_ID: 1
TYPE: FILE
HL_NAME: \DOC\
LL_NAME: 2FAAA
OBJECT_ID: 1043
ARCHIVE_DATE: 2004-03-12 20:57:51.000000
OWNER:
DESCRIPTION: IBM OnDemand
CLASS_NAME: STANDARD
select * from archives where node_name='ODARCHIVE'

```

Figure 8-39 Select statement output to Tivoli Storage Manager after OnDemand migration

Figure 8-40 illustrates how storage management works in OnDemand.

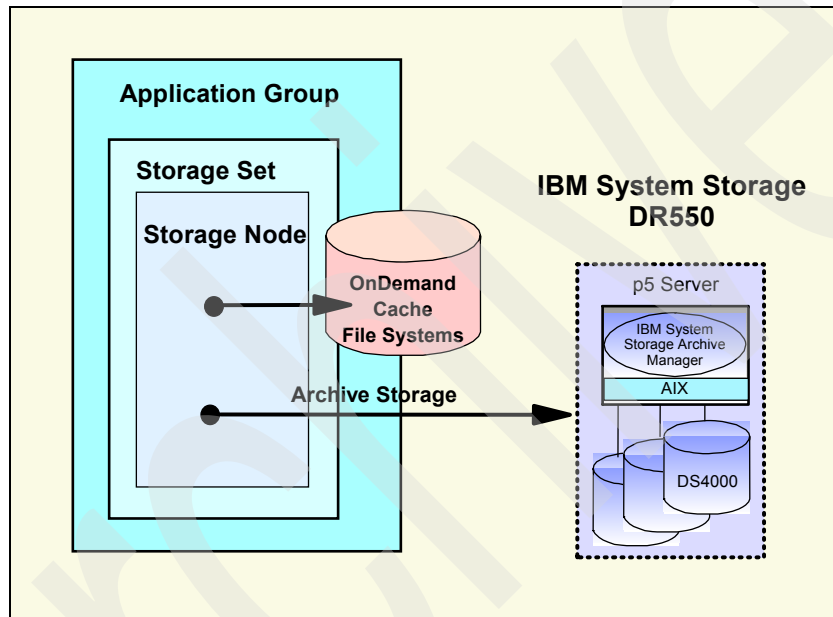


Figure 8-40 Storage management in OnDemand

If you are configuring an OnDemand for UNIX system to use Tivoli Storage Manager for archive storage, you need to be sure that the `ars.cfg` file has been updated to reflect that Tivoli Storage Manager (SSAM) is to be used as the storage manager. The file also needs to include valid paths for Tivoli Storage Manager options files and all of the Tivoli Storage Manager components that will be used.

## 8.5 IBM Optim overview

IBM Optim can manage enterprise data throughout every stage of the information life cycle. IBM Optim enables your company to assess, classify, subset, archive, store, and access enterprise application data.




Figure 8-41 IBM Optim solution

The IBM Optim solution offers the following functions:

**Archive:** Tier business transactions by age and status. Segregate historical information from current activity and safely remove it to a secure archive (such as the IBM System Storage DR550). Maintain a production database in a manageable size to simplify maintenance and speed disaster recovery.

**Classify:** Apply business rules to govern active, inactive, and reference data. Define and implement policies; when data needs to be available, where it should be stored, how long it should be retained and who can have access.

**Assess:** Determine where application data is growing fastest and assess the impact of data tiering strategies. Identify and address potential problems before they affect business results.

**Test Data Management:** Speed application deployment by streamlining the way you create and manage test environments. Subset and migrate data to build realistic and right sized test databases. Eliminate the expense and effort of maintaining multiple database clones.

De-identify confidential information to protect privacy.

**Access:** Enable decision makers to access the right data at the right time. Query and browse active, inactive, and reference data. Utilize familiar forms, screens, and panels. Generate standard and custom reports. Restore archived transactions if business processing becomes a requirement.

**Store:** Store application data according to its evolving business value. Maintain active transactions in high speed storage tiers. Relocate reporting data on secure “WORM” devices. Reclaim underutilized capacity and maximize the value of your existing storage infrastructure.

**Note:** We focus here on the Archive component and explain how it integrates with the IBM System Storage DR550.

Using the archiving features from IBM Optim, you can:

- ▶ Isolate historical data from current activity and safely remove it to a secure archive.
- ▶ Access archived historical data easily, using familiar tools and interfaces.
- ▶ Restore archived data to its original business context when it requires additional processing.

### 8.5.1 IBM Optim Archive

Archive provides everything you need to create and manage archives of relationally intact data from databases with any number of tables, interconnected with any number of DBMS and application-managed relationships, regardless of their complexity.

After creating an Archive File, Archive selectively removes data from the production database, according to your instructions, to maximize database performance and response time. An indexing feature allows you to quickly search Archive Files for needed information and, if necessary, restore all or a precisely selected, referentially intact, portion of the data.

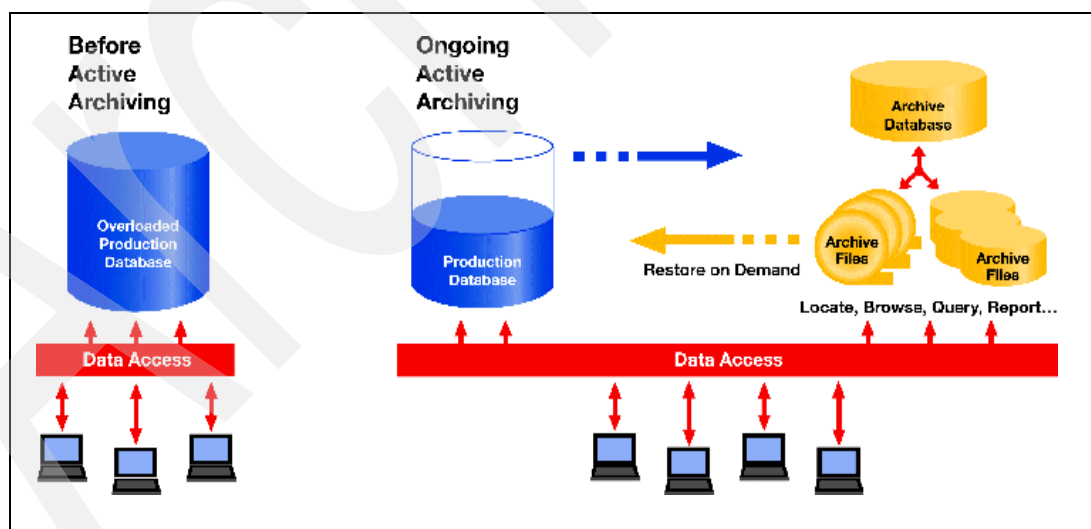


Figure 8-42 Implementing IBM Optim Archive

Archive allows you to introduce application-specific logic into your archiving operation, thereby integrating archiving with your applications. The Archive actions allow you to define and execute custom SQL statements or calls to a stored procedure at predefined points in the Archive and Restore processes.

For example, you might instruct Archive to call a routine you have written that updates a column with a value each time it deletes a row from your production database. Your application can then use this value to determine if a customer's data has been archived.

## How Archive works

Archiving data with Archive is a simple two-step process. It is depicted in Figure 8-43.

1. Create an Access Definition to specify the tables and relationships that define the set of referentially intact data to be archived. In the Access Definition, you also indicate any data to be deleted from the production database after archiving, set up indexing parameters, and define Archive Actions.
2. Archive copies the data described in the Access Definition to an Archive File, executes appropriate actions, creates indexes used subsequently to find archived data, and deletes the selected data from the database.

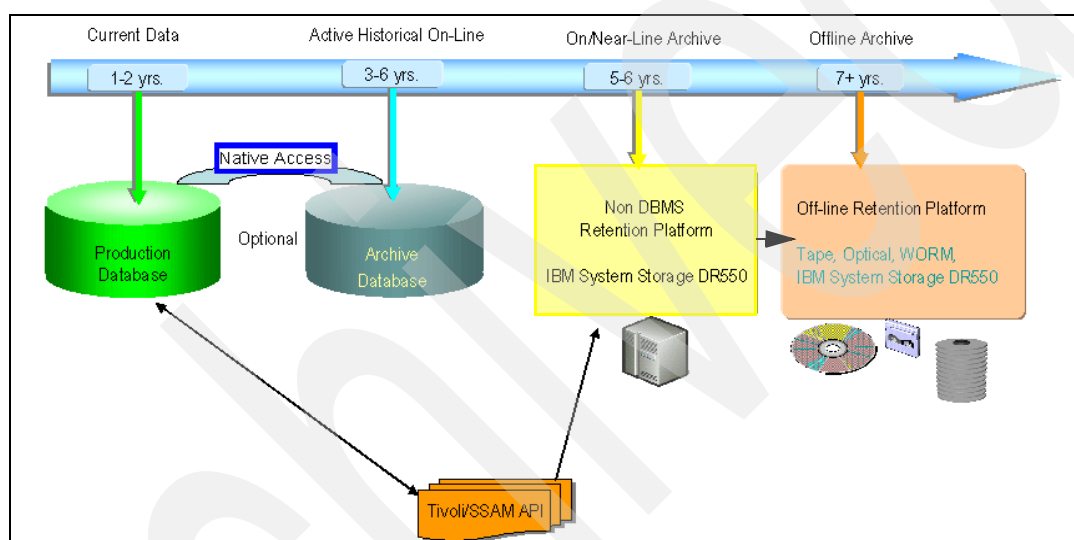


Figure 8-43 Data Retention Strategies with IBM System Storage DR550

## Delete feature

The powerful, yet safe, delete feature resolves the problem of deleting from the production database. Using standard facilities for all operations, Archive quickly and accurately deletes all or a portion of the archived data. For example, you might want to archive data for customers that have been inactive for the past year. You can create an Archive File of all data pertaining to the inactive customers and delete only the order and payment history from the production database, leaving the master account information, such as name and address, intact.

## Search Archive Files

The search facilities allow you to search Archive Files, specifying criteria to narrow the search. Search results can be presented in an interactive display, allowing you to browse archived data without having to restore it to your production system, a method useful in many situations, for example, to answer customer inquiries.

## Restore Archive Files

The Restore Process allows you to restore archived data to the production database or to a separate database and accommodates data model changes during the restoration. In addition, you can find or restore data using criteria that differ completely from that used to create the archive.

For organizations that have developed a comprehensive archiving strategy, Archive and Restore Processes can be automated, with archiving occurring on a regularly scheduled basis and restoration triggered by applications that provide the criteria for data to be restored.

### **Summary**

Archive addresses a critical operational need for organizations with large, complex databases. Old data can be archived in a precise manner and production databases optimized for peak performance. Archived data can be browsed or selectively restored as needed.

## **8.6 IBM Optim integration with IBM System Storage DR550**

In this section, we describe the IBM Optim software configuration required to connect to the DR550 System Storage Archive Manager Server.

### **8.6.1 Storage Profile**

The Storage Profile is needed to create a duplicate Archive File, to copy an Archive File to the DR550 SSAM Server, or to implement a retention policy for the primary Archive File. Without a Storage Profile, all management of the Archive File (such as copying, deleting, or saving to DR550 SSAM Server) would have to be done manually. You can automatically manage Archive Files and integrate with the DR550 by defining a Storage Profile.

### **8.6.2 Tivoli tab**

In the Storage Profile definition dialog, use the Tivoli tab to identify the file recall path and other parameters for the SSAM API (as used by the DR550). To access the SSAM API, you must specify a node name and password in Product Options. This is shown in Figure 8-44 on page 344.

**Note:** This tab is displayed only if you select the **Copy to Backup Device** check box and choose **Tivoli** from the corresponding drop-down list on the Primary Copy tab.

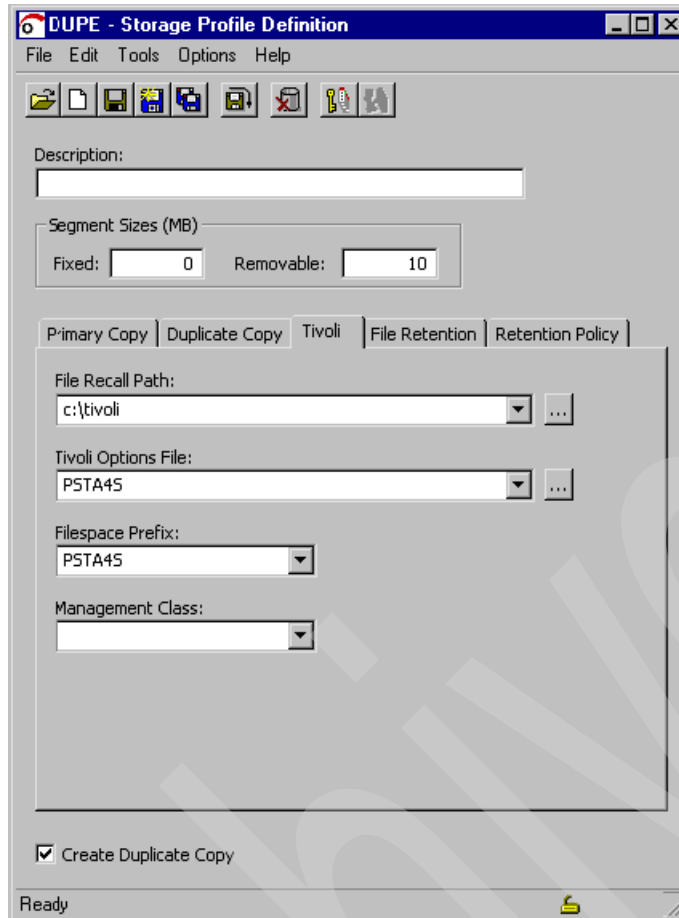


Figure 8-44 Tivoli Configuration tab

To integrate Optim Archive with the SSAM Server on the DR550, specify the following:

► File Recall Path

The complete path to the default directory for Archive Files recalled from SSAM API. Click the arrow to select it from the drop-down list, or click the browse button to select from your system directories. If you do not specify a recall path, the Archive File is recalled to the original path.

**Note:** Use the File Retention tab to specify the period of time for which a recalled Archive File is retained on disk. (Refer to 8.6.3, “File Retention tab” on page 345.)

► Tivoli Options File

The complete path to the Tivoli options file (dsm.opt) that contains a set of processing options. Click the arrow to select from the drop-down list, or click the **Browse** button to select from your system directories. If you do not specify a path, Tivoli looks in the directory from which Archive is executing.

**Important:** You must include ENABLEARCHIVERETENTIONPROTECTION Yes in dsm.opt file.

- Filespace Prefix

Enter the name of the file space prefix that identifies a group of Archive Files in Tivoli, or click the arrow to select it from the drop-down list. If you do not specify a file space prefix, Archive uses the default prefix, PSTA4S.

- Management Class

Enter the name of a valid SSAM management class, which is already defined in the DR550 SSAM Server. It defines the policy for the DR550 SSAM Server operations. Click the arrow to select it from the drop-down list. If you do not specify a management class, SSAM API uses the default management class.

### 8.6.3 File Retention tab

Use the File Retention tab to provide parameters for retaining an Archive File on disk after copying to or recalling from the DR550 SSAM Server (note that this is different from the retention period set through the SSAM policy in the associated management class), as shown in Figure 8-45.

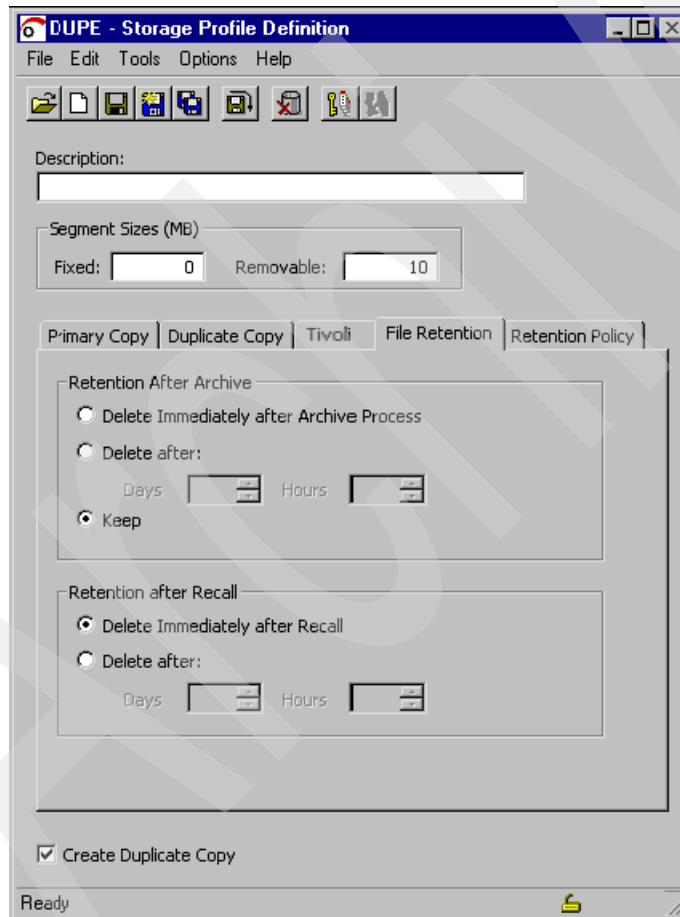


Figure 8-45 File Retention tab

**Note:** This tab is displayed only if you select the **Copy to Backup Device** check box and choose a backup device on the Primary Copy tab.

## Retention after archive

The Archive Process creates an Archive File on disk and, if required by the Storage Profile, copies it to the DR550 SSAM Server. You can select an option to delete the Archive File from disk immediately after the Archive Process is complete or after a specified period, measured from the time the Archive File was last accessed (whether by IBM Optim or other software). Once deleted, the Archive File is no longer in the disk file system.

You can also choose to retain the Archive File on disk as well as on the DR550 SSAM Server.

**Note:** Retention after Archive settings do not apply to a primary Archive File that has also been copied to a WORM device.

### *Delete immediately after Archive Process*

Delete the Archive File from disk immediately after the Archive Process is complete. A deleted Archive File is recalled by a deferred Delete Process.

**Note:** This option is not recommended if you plan to access the Archive File frequently.

### *Delete after*

Delete the Archive File from disk after the specified period, measured from the time the Archive File was last accessed. You can specify a number of days, hours, or a combination of both. (Storage space limitations should be considered when specifying how long to retain Archive Files.)

- Days:

Number of days (0 to 999) to keep the Archive File on disk (since the file was last accessed).

- Hours:

Number of hours (0 to 23) to keep the Archive File on disk (since the file was last accessed).

**Note:** You can specify 0 for either Days or Hours, but not for both.

- Keep:

Do not delete the Archive File from disk.

**Attention:** Use caution when creating and deleting Archive Files from disk after copying to the DR550 SSAM Server. If you delete an Archive File from disk and later reuse the Archive File name, you will not be prompted to confirm the overwrite. At run time, the new Archive File will overwrite the original Archive File copied to the DR550 SSAM Server.

## Retention after recall

When you process or browse an Archive File that was copied to the DR550 SSAM Server and removed from disk, IBM Optim references the Storage Profile that was used to create the Archive File to copy the Archive File from the DR550 SSAM Server to the recall path. You can choose to delete the Archive File from disk immediately after the recall or after a specified period of time.



To specify an Archive File that was copied to the DR550 SSAM Server and removed from disk, you can:

- ▶ Enter the Archive File name, if known, for an action request.
- ▶ Click the **Retrieve** button, if available, to select the last created Archive File.
- ▶ Select the Archive File from the Archive Directory, for browsing or a Restore Request.

Specify the period of time for which the recalled Archive File is retained on disk after the recall is complete:

- ▶ Delete immediately after Recall

Delete the Archive File from disk immediately after the recall is complete.

- ▶ Delete after

Delete the Archive File from disk after the specified period, measured from the time the recall is complete. You can specify a number of days, hours, or a combination of both. (Storage space limitations should be considered when specifying how long to retain Archive Files.)

- Days:

Number of days (0 to 999) to keep the recalled Archive File on disk after the recall is complete.

- Hours:

Number of hours (0 to 23) to keep the recalled Archive File on disk after the recall is complete.

**Note:** You can specify 0 for either Days or Hours, but not for both.

Archived



## DR550 call home features

This chapter describes the three separate call home features that can be used with the DR550.

All three features are required if you need call home support simultaneously for the DR550 SSAM Servers, the DR550 File System Gateway nodes, and for the DR550 Storage Controller.

We recommend that you implement the call home features in your environment.

## 9.1 Call home functions

There are three methods that should be used with the DR550 to manage alerts and provide call home support:

1. Electronic Service Agent (eSA) for System p (on AIX) provides call home support for the DR550 SSAM Servers (p52A systems). We describe this functionality in 9.2, “Electronic Service Agent for AIX” on page 351.
2. Electronic Service Agent (eSA) for System x, an extension of IBM Director, provides call home support for the optional File System Gateway (FSG) nodes. We provide an overview of the IBM Director and details about the Electronic Service Agent for System x in 9.3.1, “IBM Director overview” on page 359.
3. IBM Remote System Manager (RSM) provides call home support for the DR550 Storage Controller (DS4000). For details, refer to 9.4, “RSM for DR550” on page 372. Note that there is a specific version of RSM for the DR550.

Figure 9-1 depicts the three call home features and how the information flows from the managed systems to IBM support.

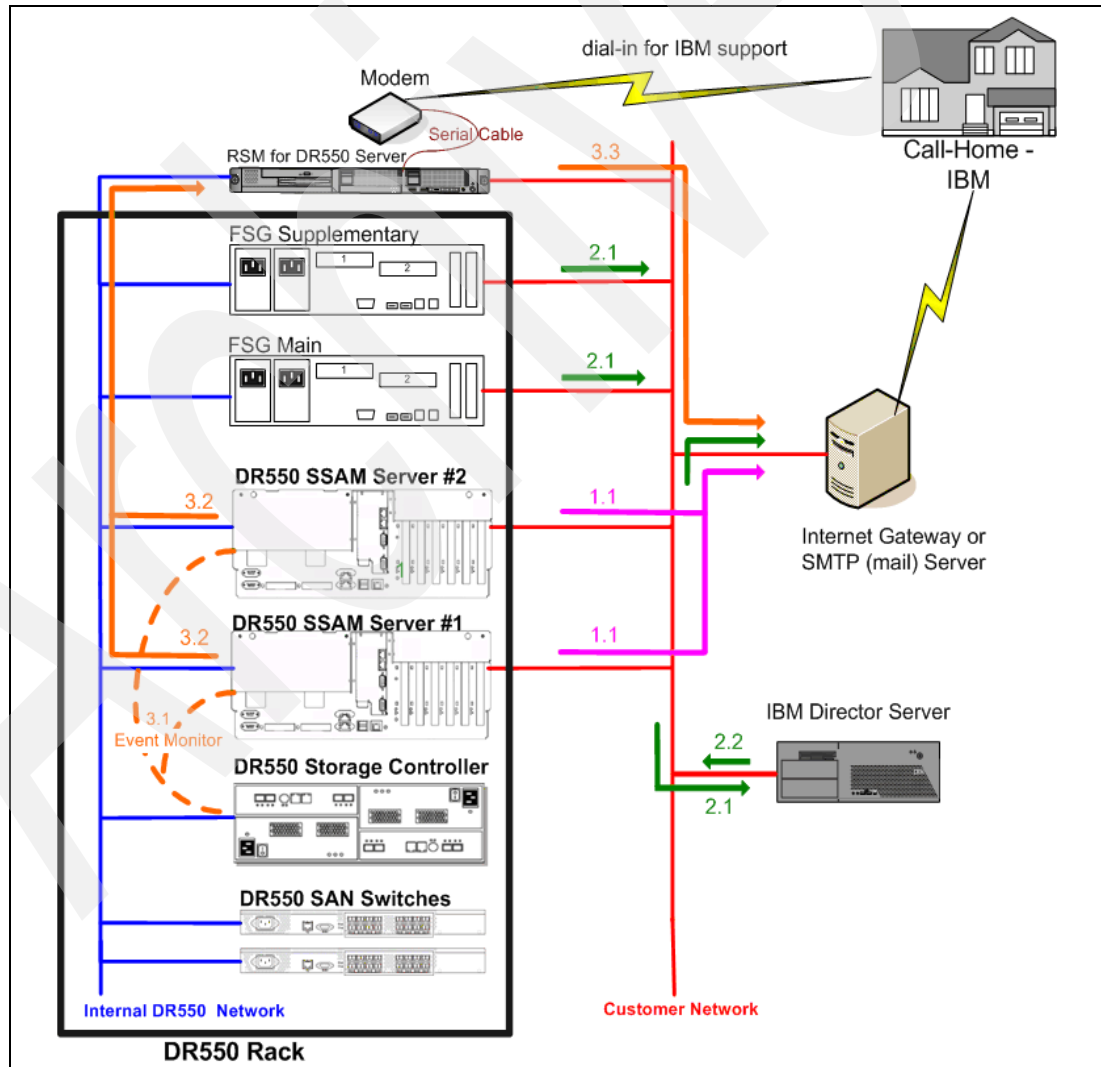


Figure 9-1 DR550 Call home features

For the DR550 SSAM servers that use eSA for AIX on System p, the call home information flow is indicated by the pink arrows (labelled 1.x). In this case, the eSA for System p sends the information directly from the DR550 SSAM server to IBM Support (Internet or dial-up).

For FSG nodes (System x) that rely on eSA for System x, the call home information flow is indicated by the green arrows (labelled 2.x). In this case, the events go first to the IBM Director Server, which in turn sends the information to IBM support (Internet or dial-up).

The monitoring for the DR550 Storage Controller (DS4000) is a bit different. Here the Event Monitor component of the DS4000 Storage Manage runs on the DR550 SSAM Servers, on AIX. The Event Monitor catches error conditions on the DS4000 (as indicated by the orange dotted lines) and sends the messages to the RSM server (indicated by the orange arrows, labeled 3.x). The RSM sends e-mails to IBM support through the SMTP server. When the message is received, IBM Support is able (during a specified time window, and if authorized by the customer) to dial in to the RSM server and collect additional diagnostic data or run some maintenance tasks.

In the following sections, we describe the configuration of the different call home functions in detail.

**Important:** Check with your local service representative that your DR550 2233-DR1 or 2233-DR2 is correctly recorded in the IBM support systems. This is prerequisite for the call home functions to be able to open calls in the IBM support systems.

## 9.2 Electronic Service Agent for AIX

The Electronic Service Agent for AIX is a “no-charge” software tool that resides on the customer’s System p to monitor events and transmit system inventory information to IBM. Information collected through this Electronic Service Agent is available to IBM support representatives to assist in diagnosing problems. With the early knowledge about potential problems provided by the Electronic Service Agent, IBM can proactively respond to customers and assist in maintaining higher availability and performance.

**Note:** There is a specific version of Electronic Service Agent (eSA) for AIX already installed on the DR550 SSAM Servers. Do not update this preinstalled version with a version available for download from the eSA for System p Web site.

### 9.2.1 How eSA for AIX works

The process works as follows:

1. An event is generated when a hardware error exceeds a preset threshold.
2. The Electronic Service Agent sends a service request to IBM.
3. The service request is forwarded to the call management system for the machine’s country. Details sent with the service request include unique identifiers, system machine type, serial number, machine name (as displayed in IBM Director), company, contact person, location details, and other useful inventory and diagnostic information.
4. IBM responds to the Electronic Service Agent by returning a service request number, branch number, and country code.

## 9.2.2 Configuring eSA for AIX

As already described, eSA for AIX is preinstalled on the DR550 SSAM servers. You need to be logged on as root to perform the configuration steps below.

**Note:** Make sure that port 5024 is enabled on the firewalls between your administrative computer and the DR550 SSAM servers.

1. First, verify that `/bin/false` is in the list of shells included in the file `/etc/security/login.cfg`. You can use the vi editor to edit the file. If `/bin/false` is not already in the list of shells, you must add it. Refer to Figure 9-2.

```
shells =
/bin/sh,/bin/bsh,/bin/csh,/bin/ksh,/bin/tsh,/bin/ksh93,/usr/bin/sh,/usr/bin/bsh
,/usr/bin/csh,/usr/bin/ksh,/usr/bin/tsh,/usr/bin/ksh93,/usr/sbin/uucp/uucico,
/usr/sbin/sliplogin,/usr/sbin/snappd,/bin/false
```

Figure 9-2 Part of `/etc/security/login.cfg`

2. Verify that the user account `esausuer` exists on the DR550 SSAM server. If not, you have to create it.

To create the `esausuer` account, enter **smitty users** at the AIX command line. In the smitty window, add the `esausuer` account with no other settings. Then use the Change/Show Characteristics of a User option to change the `esausuer`.

Make the following changes, as shown in Figure 9-3 on page 353:

- Set Group Set to system.
- Set Home directory to `/dev/null`.
- Set Initial Program to `/bin/false`.
- Set User can Login Remotely to false.
- Set Valid TTYs to `/dev/tty`.

Change / Show Characteristics of a User

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

[MORE...3]	[Entry Fields]	
Primary GROUP	[staff]	+
Group SET	[system]	+
ADMINISTRATIVE GROUPS	[]	+
ROLES	[]	+
Another user can SU TO USER?	true	+
SU GROUPS	[ALL]	+
HOME directory	[/dev/null]	
Initial PROGRAM	[/bin/false]	
User INFORMATION	[]	
EXPIRATION date (MMDDhhmmyy)	[0]	
Is this user ACCOUNT LOCKED?	false	+
User can LOGIN?	true	+
User can LOGIN REMOTELY(rsh,tn,rlogin)?	false	+
[MORE...35]		

F1=Help	F2=Refresh	F3=Cancel	F4=List
F5=Reset	F6=Command	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

Figure 9-3 Change characteristics of esauser

Change the password for the esauser by entering **passwd esauser** at the command line. Finally, run **pwdadm -c esauser** so that the esauser does not force you to change the password during the first login.

### 3. Run **smitty esa\_main**.

In the window that displays, select Configure Service Connectivity (see Figure 9-4), then Create/Change Service Connectivity (Figure 9-5 on page 354), and finally select Create/Change Primary Service configuration.

Electronic Service Agent

Move cursor to desired item and press Enter.

Configure Electronic Service Agent

**Configure Service Connectivity**

Start Electronic Service Agent

Stop Electronic Service Agent

F1=Help	F2=Refresh	F3=Cancel	F8=Image
F9=Shell	F10=Exit	Enter=Do	

Figure 9-4 eSA - Configure service connectivity

Configure Service Connectivity

Move cursor to desired item and press Enter.

**Create/Change Service Configuration**

- Restore Default Configuration
- Manage the Service and Support Proxy
- Change Connectivity Settings
- Configure PPP for Dial

F1=Help                      F2=Refresh                      F3=Cancel                      F8=Image  
F9=Shell                      F10=Exit                      Enter=Do

Figure 9-5 eSA - Create/Change Service configuration

Create/Change Service Configuration

Move cursor to desired item and press Enter.

**Create/Change Primary Service Configuration**

- Create/Change Secondary Service Configuration
- Create/Change Tertiary Service Configuration

F1=Help                      F2=Refresh                      F3=Cancel                      F8=Image  
F9=Shell                      F10=Exit                      Enter=Do

Figure 9-6 eSA - Create/change Primary Service Configuration

In the window shown in Figure 9-7 on page 355, enter the appropriate connection type, then select Yes for the Test service configuration field.



```

Create/Change Primary Service Configuration

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[TOP] [Entry Fields]
Connection type [Direct Internet] +
Test service configuration [Yes] +

If type is DIRECT_INTERNET, no entry required.

If type is DIAL,
TTY port number of modem [] +
Modem type [] +
Primary telephone number [] +
Alternate telephone number [] +
Dial prefix [] +

If type is HTTP_PROXY,
[MORE...5]

F1=Help F2=Refresh F3=Cancel F4=List
F5=Reset F6=Command F7=Edit F8=Image
F9=Shell F10=Exit Enter=Do

```

Figure 9-7 eSA - Create Primary Service Configuration

Press Enter to test your configuration (see Figure 9-8).

```

COMMAND STATUS

Command: OK stdout: yes stderr: no

Before command completion, additional instructions may appear below.

#####

Testing Direct Internet Service Configuration

Performing Basic Connectivity Test ... SUCCESS

F1=Help F2=Refresh F3=Cancel F6=Command
F8=Image F9=Shell F10=Exit /=Find
n=Find Next

```

Figure 9-8 eSA - Test configuration

- Once the test is completed, press F3 several times until you are returned to the Electronic service agent main menu (Figure 9-4 on page 353).

- From the Electronic Service Agent menu, select Configure Electronic Service Agent. Here you have to provide your (customer) information. See Figure 9-9. Upon completion of this step, the DR550 SSAM server will attempt the connection to IBM support.

Configuring Electronic Service Agent

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

[Entry Fields]

[IBM Redbook]

[Somebody]

[111111]

[somebody@de.ibm.com]

GERMANY

[222222]

UNITED STATES

Operational setting

\* Port number on which to receive connections

[5024]

#

F1=Help

F2=Refresh

F3=Cancel

F4=List

F5=Reset

F6=Command

F7=Edit

F8=Image

F9=Shell

F10=Exit

Enter=Do

Figure 9-9 Configuring Electronic Service Agent

Note the port number (5024) as the default on which to receive connections. That is why you must make sure that this port is open on your AIX system.

At this point, the eSA for AIX configuration is complete for providing the call home functionality

You should now be able to start a Web browser on your management workstation and point to <https://<ip-address>:5024> to start using the eSA for AIX agent user interface.

You have to log on with the esauser account.

The eSA for AIX main window displays, as shown in Figure 9-10 on page 357.

From the main window, you can run various tasks and change any additional settings for the Electronic Service Agent.

For example, as shown in Figure 9-11 on page 357, you can configure eSA for AIX to send traps to an SNMP destination (this is optional and not required for the call home functionality). We discuss that SNMP destination option in more detail in Chapter 10, “DR550 SNMP monitoring” on page 389.

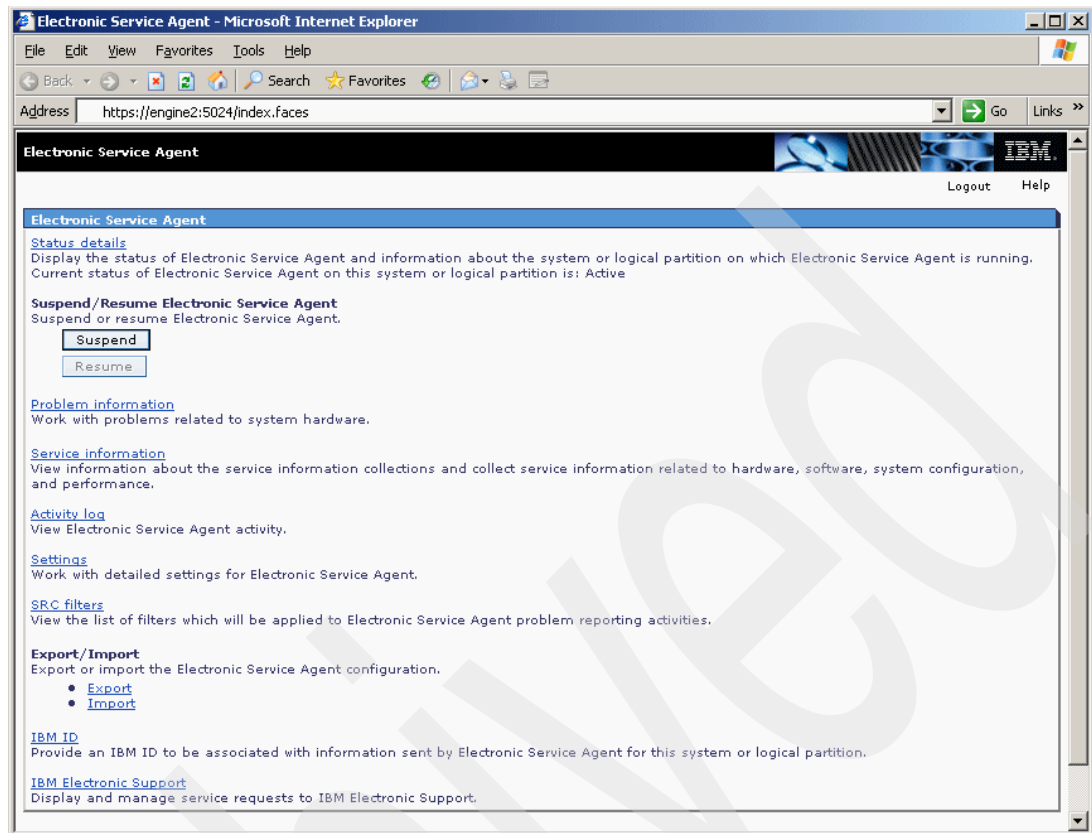


Figure 9-10 eSA for AIX - Main window

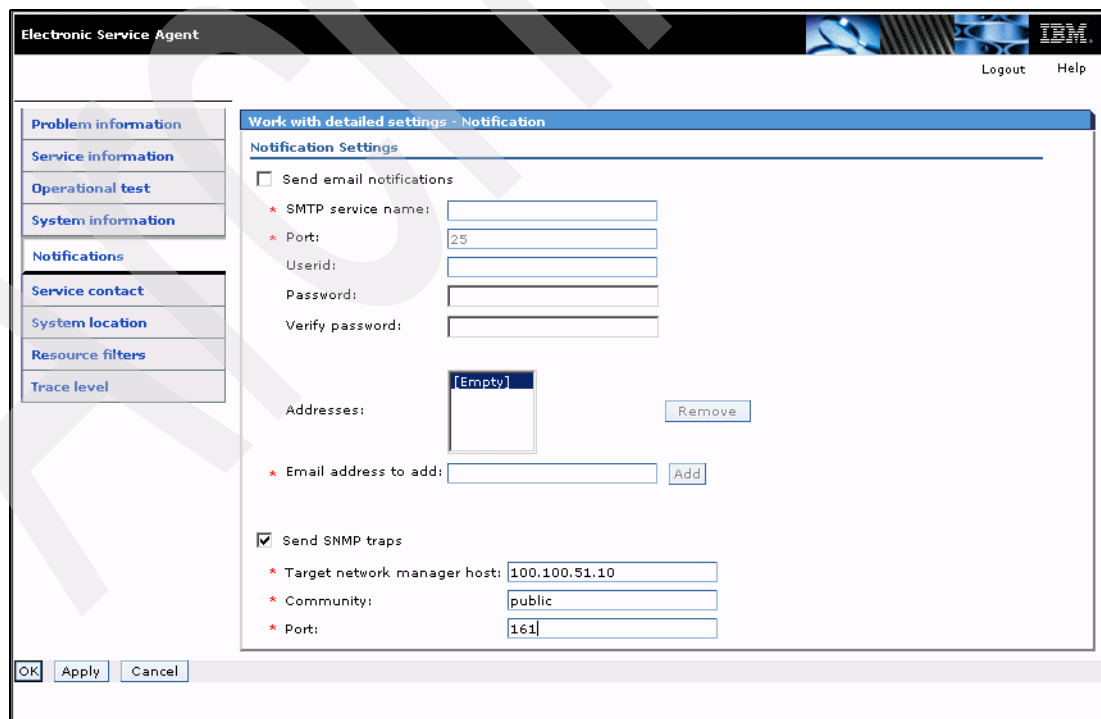


Figure 9-11 Configure SNMP alerting

To make sure that your configuration works, you should create a test problem. To do so, go under Problem Information and click **Send test problem**, as shown in Figure 9-12.

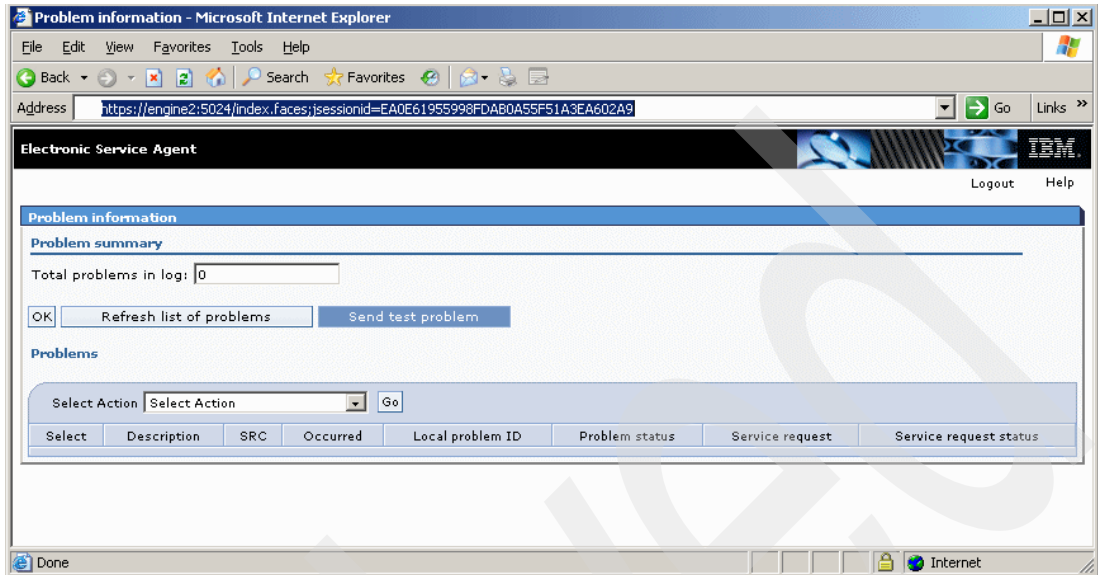


Figure 9-12 eSA for AIX - Problem Information window

If the test problem was successfully submitted, you should get a Success message window, as shown in Figure 9-13.

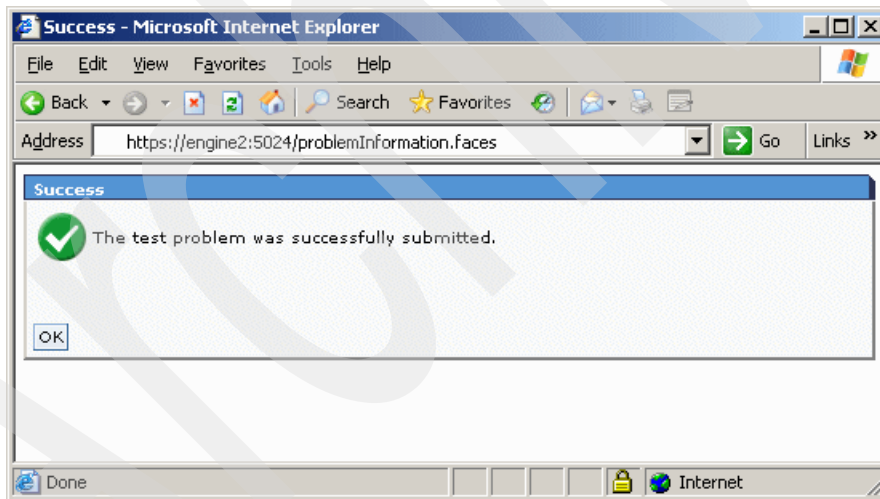


Figure 9-13 eSA for AIX - test problem

**Note:** Call your IBM service representative to verify that your systems are properly recorded within IBM support for successful processing.

In Figure 9-14 on page 359, you can see (indicated by a red circle) the service request number that is created by the IBM call home server. When you call IBM for that specific problem, reference this service request number.

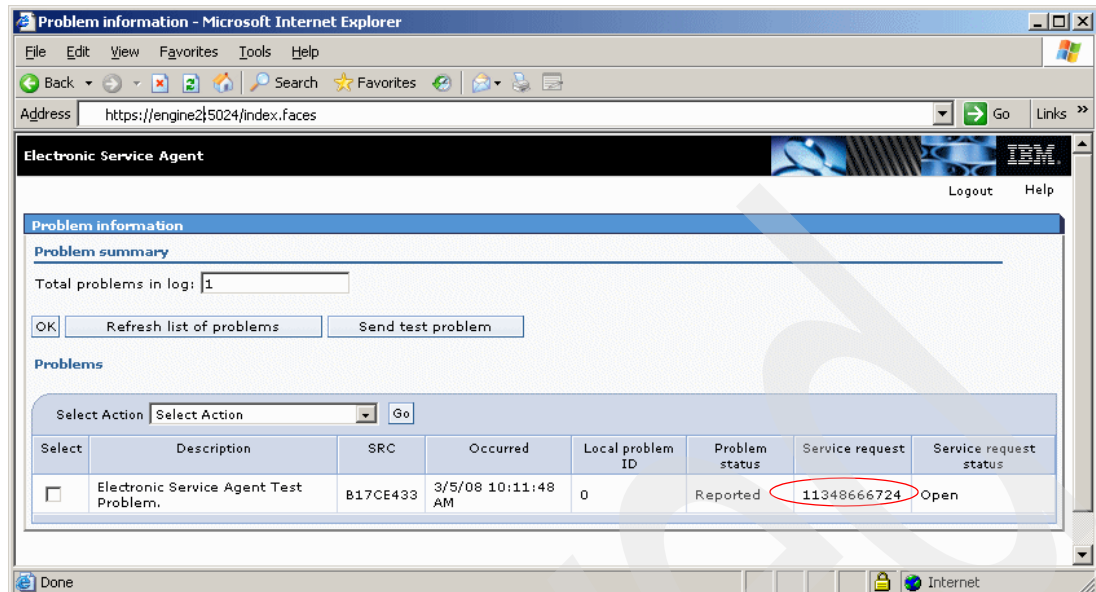


Figure 9-14 eSA for AIX - Problem Information window after the test

For further details, refer to the *Electronic Service Agent on AIX*, found at [http://www-304.ibm.com/jct03004c/support/electronic/portal/!ut/p/\\_s.7\\_0\\_A/7\\_0\\_CI?category=5&locale=en\\_US](http://www-304.ibm.com/jct03004c/support/electronic/portal/!ut/p/_s.7_0_A/7_0_CI?category=5&locale=en_US) and *IBM System Storage DR550 Version 4.5 Installation Roadmap*, GC27-2175.

## 9.3 IBM Director and eSA for System x

With the DR550, the Electronic Service Agent for System x is used to provide call home support functionality for the optional FSG nodes.

eSA for System x is implemented as an extension of IBM Director Server, which is thus a prerequisite. For this reason, we include here a brief overview of the IBM Director.

### 9.3.1 IBM Director overview

IBM Director provides an integrated suite of software tools for a consistent, single point of management and automation. With IBM Director, IT administrators can view and track the hardware configuration of remote systems in detail and monitor the usage and performance of critical components, such as processors, disks, and memory.

IBM Director enables monitoring and event management across a heterogeneous IT environment, including Intel and POWER systems that support Windows, Linux, NetWare, ESX Server, AIX 5L, and i5/OS® from a single Java-based user interface. From one access point, users can monitor system resources, inventory, events, task management, core corrective actions, distributed commands, and hardware control for servers, clients, and storage.

The hardware in an IBM Director environment can be divided into the following groups:

- ▶ *Management servers*: One or more servers on which IBM Director Server is installed.
- ▶ *Managed systems*: Servers, workstations, or any computer managed by IBM Director.
- ▶ *Management consoles*: Servers and workstations, from which you communicate with one or more IBM Director Servers.
- ▶ *SNMP devices*: Network devices, disk systems, or computers that have SNMP agents installed or embedded (in the case of the DR550, these would be the DR550 SSAM Servers, the FSG nodes, and the DR550 Storage System).

SNMP is described in Chapter 10, “DR550 SNMP monitoring” on page 389.

Figure 9-15 depicts a typical IBM Director management environment.

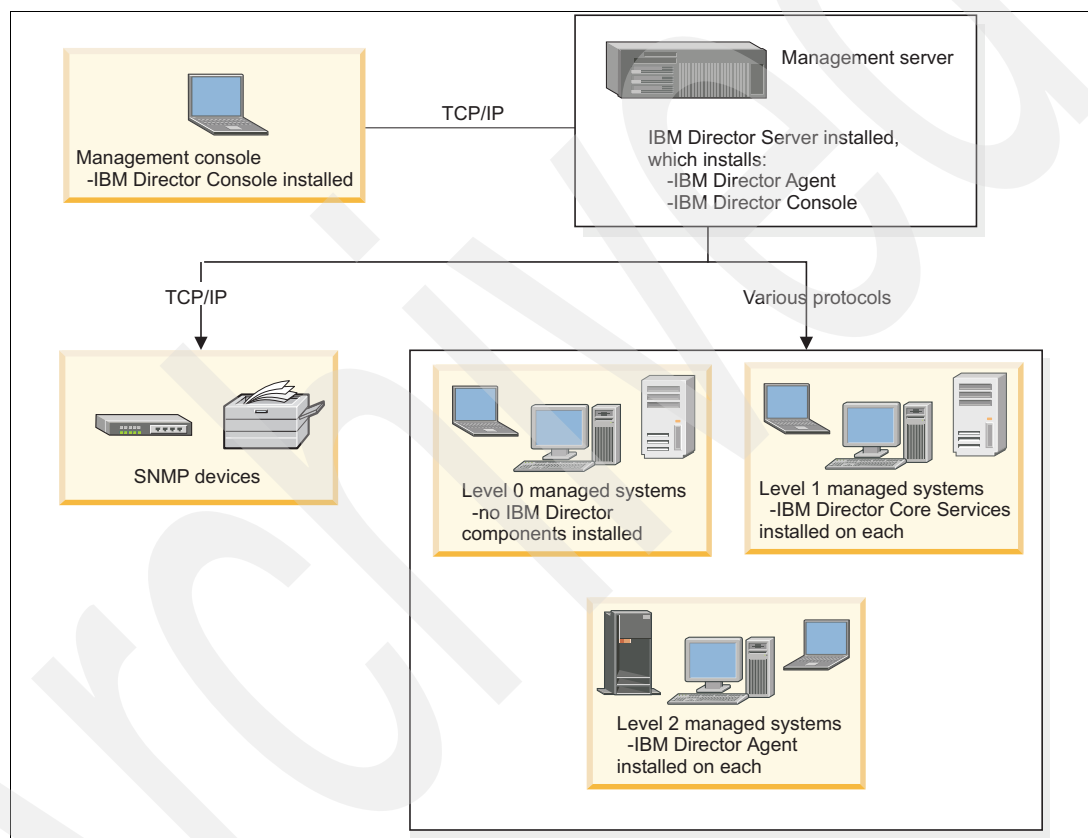


Figure 9-15 Typical IBM Director management environment

IBM Director software has four main components:

- ▶ **IBM Director Server**

IBM Director Server is the main component of IBM Director. IBM Director Server contains the management data, the server engine, and the application logic. It provides basic functions such as discovery of the managed systems, persistent storage of inventory data, SQL database support, presence checking, security and authentication, management console support, and administrative tasks.

In the default installation under Windows, Linux, and AIX, IBM Director Server stores management information in an embedded Apache Derby database. You can access information that is stored in this integrated, centralized, relational database even when the managed systems are not available. For large-scale IBM Director solutions, you can use a

stand-alone database application, such as IBM DB2 Universal Database™, Oracle®, or Microsoft SQL Server®.

- ▶ IBM Director Agent (also known as Level-2 Agent)

IBM Director Agent is installed on a managed system to provide enhanced functionality for IBM Director to communicate with and administer the managed system. IBM Director Agent provides management data to the management server through various network protocols.

- ▶ IBM Director Core Services (also known as Level-1 Agent)

IBM Director Core Services provides a subset of IBM Director Agent functionality that is used to communicate with and administer a managed system. Systems that have IBM Director Core Services (but not IBM Director Agent) installed on them are referred to as *Level-1* managed systems.

IBM Director Core Services provides management entirely through standard protocols. This includes discovery, authentication, and management. The IBM Director Core Services package installs an SLP service agent, an SSL-enabled CIMOM (on Linux), or CIM mapping libraries to WMI (on Windows), an optional ssh server, and platform-specific instrumentation.

- ▶ IBM Director Console

IBM Director Console is the graphical user interface (GUI) for IBM Director Server. Using IBM Director Console, system administrators can conduct comprehensive hardware management using either a drag-and-drop action or a single click. You can install IBM Director Console on as many systems as you need. The license is available at no charge.

All IBM customers can download the latest version of IBM Director code from the IBM Director Software Download Matrix page:

<http://www.ibm.com/systems/management/director/downloads.html>

### 9.3.2 IBM Director ISS Extensions for DR550

Starting with Version 4.5 of the DR550, specific DR550 extensions for IBM Director, known as the DR550 Integrated Storage Solution (ISS) extensions, are available. You can download these extensions by selecting the Product Support link from the following IBM Web site:

<http://www.ibm.com/systems/storage/disk/dr/index.html>

Download the file `issinstall.zip` (this is for a Microsoft Windows environment).

After the file is unzipped, run `ISSinstall` on a system where you have previously installed IBM Director Server V5.20.2. When running the `ISSinstall` program, the IBM Director Server service will be restarted.

**Important:** When running `ISSinstall`, the recommended procedure allows a reset of your IBM Director Server configuration. Understand that this will reset the IBM Director to its initial state and any existing configuration information will be lost. This should not be an issue if you install a Director Server specifically for the DR550.

The ISS extensions must be installed on the Director Server or any Director Console. The install procedure prompts you to indicate whether the ISS Extensions are being installed on the IBM Director server itself or an IBM Director Console. Type `Server` or `Console` as appropriate and press `Enter` (see Figure 9-16 on page 362 and Figure 9-17 on page 362).

```

C:\WINDOWS\system32\cmd.exe
Would you like to install the Archive System server extension (including
console) or only the Archive System console extension?
Respond with Server or Console:Server

Stopping IBM Director...
The IBM Director Support Program service is stopping...
The IBM Director Support Program service was stopped successfully.

Director stop completed

Installing Archive System extension.
The process cannot access the file because it is being used by another process.
The system cannot find the file specified.
The process cannot access the file because it is being used by another process.
1 file(s) copied.
1 file(s) copied.
1 file(s) copied.
1 file(s) copied.
1 file(s) copied.
1 file(s) copied.
1 file(s) copied.
1 file(s) copied.
Installing files completed.

!!
!In order to complete the installation of the Archive System
!extension, a reset of the IBM Director Server is recommended.
!This will result in the IBM Director Server being restored
!to its initial install state. The reset may overwrite any
!changes to IBM Director properties files you may have made.
!
!You may want to save your changes to a personal directory
!before selecting OK. Selecting cancel will restart IBM
!Director Server without clearing the properties files and
!may result in unexpected behavior.
!!

To continue with the reset type OK, otherwise Cancel:

```

Figure 9-16 Installation of the DR550 Extensions for IBM Director - Part 1

Next, you need to acknowledge and confirm that you accept to reset all IBM Director settings (understand that information about all systems in your environment that you had previously discovered, or added and configured, will be lost). You can type `Cancel` if you need to keep your old settings.

```

C:\WINDOWS\system32\cmd.exe

To continue with the reset type OK, otherwise Cancel: ^C
NOT PERFORMING RESET OF IBM DIRECTOR!!!

Starting IBM Director
The IBM Director Support Program service is starting.
The IBM Director Support Program service was started successfully.

IBM Director started

IBM Director console will need to be closed before changes can take effect.
Press any key to continue . . .

```

Figure 9-17 Installation of the DR550 Extensions for IBM Director - Part 2

Once the ISS Extensions are installed and the Director Server has been restarted, launch the Director Console to define your environment. Use the Discovery Options to find your DR550 systems (and other systems in the network that you want to manage from this Director Server console). For each system you discover or add to the console, right-click the system and select **Request Access** from the pop-up menu. Here you have to authorize access using root authority.

If you have requested access to your SSAM Storage Servers and to your FSGs, then your Director Console window should be similar to the one shown in Figure 9-18 on page 363. All systems within the DR550 should be listed under the DR550 group and prefixed with the 2233 Machine Type and Model (DR2 in our example).



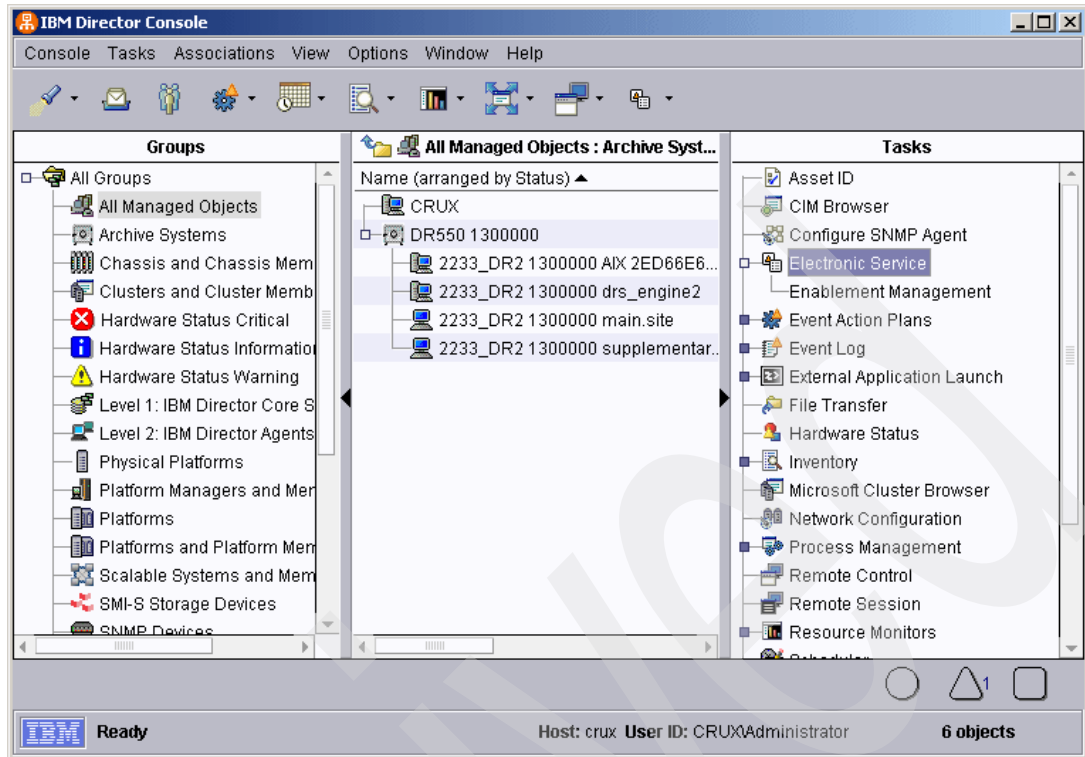


Figure 9-18 IBM Director Console with ISS Extensions for DR550

If you do not see the complete list of tasks (on the right), select **Options** → **User Administration** and check under Task Access that the user your use is not limited. Refer to Figure 9-19.

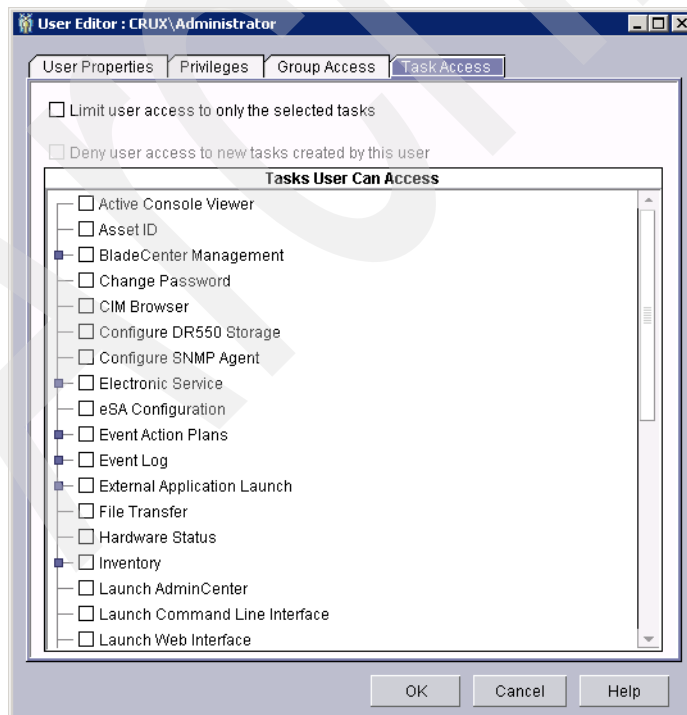


Figure 9-19 User Administration options for IBM Director

### 9.3.3 Electronic Service Agent for System x

The Electronic Service Agent (eSA) for System x is an extension of the IBM Director Server. Code installation is not required on the monitored System x servers. This extension is available for Windows and Linux based IBM Director servers.

The eSA monitors System x and BladeCenter® servers for hardware errors. Hardware errors that meet certain criteria for criticality are automatically reported to IBM. Problems typically reported include Predictive Failure Analysis® (PFA) based events, power failures, system overheating (as detected by the service processor), and RAID drive failures.

Before you can start to configure and use the eSA for System x, you have to make sure that the Electronic Service Agent extension for System x is installed on your IBM Director (it comes as an add-on of the IBM Director Server). If eSA for System x is installed, you will find it under the Tasks section in the IBM Director Console. In Figure 9-20, it is not installed.

Note also that if you are launching the Director Console from a different server than where the Director Server is installed, and if the eSA for System x is not installed at the console, it will show in the task pane but will fail to execute.

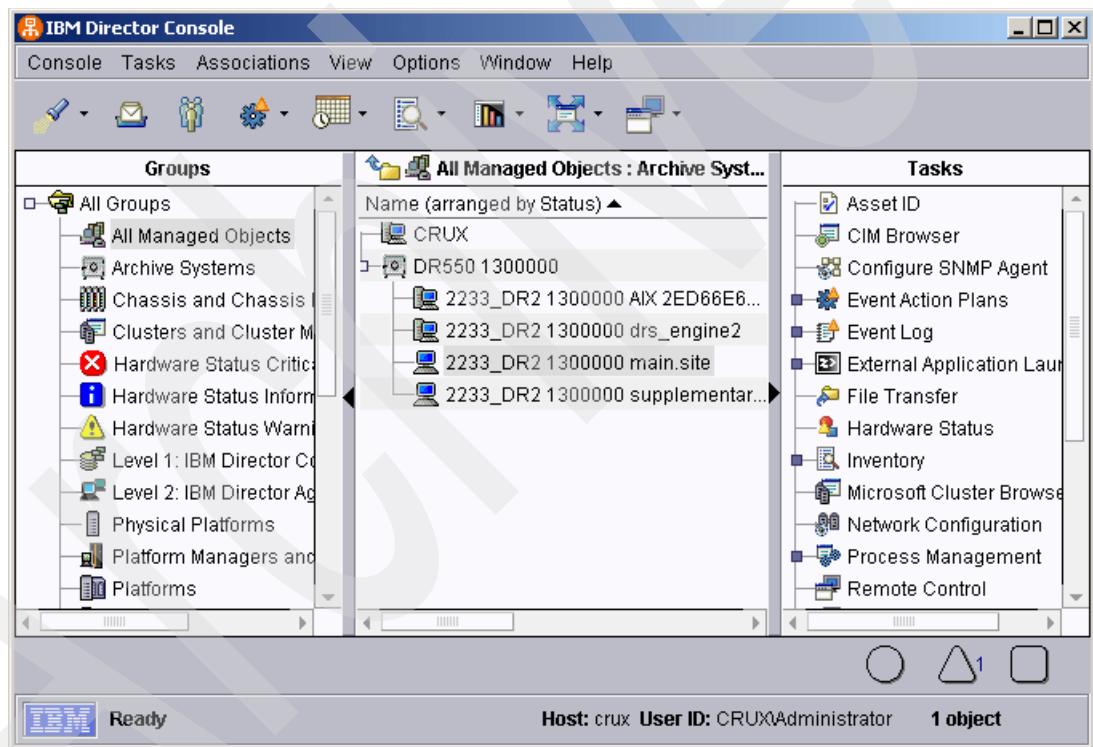


Figure 9-20 IBM Director Console main window

**Note:** Before starting the installation of the Electronic Service Agent for System x, make sure that the managed system (FSG nodes) appears in the IBM Director and that IBM Director has collected an inventory of those systems.

When the System x (that is, FSG nodes for the DR550) you want to enable for Electronic support shows up in the IBM Director Console window, and once you have requested access to it, you can collect its inventory (collecting the inventory is necessary for call home support), as shown in Figure 9-21 on page 365.

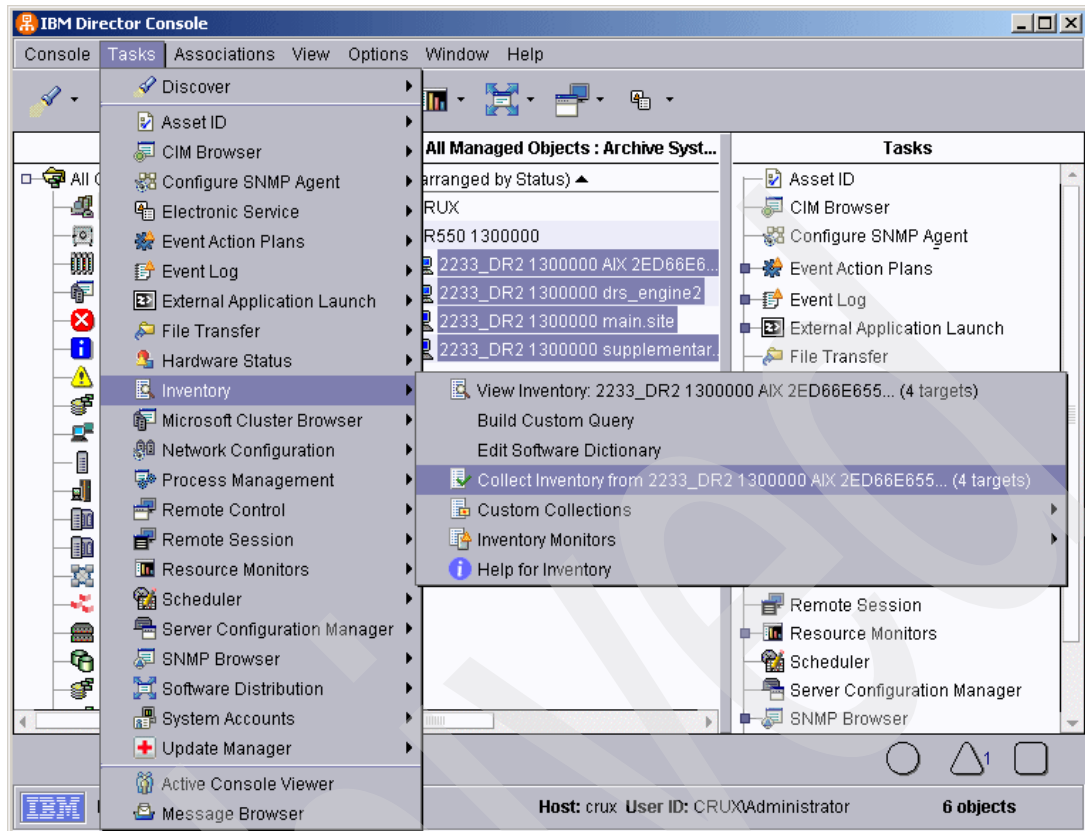


Figure 9-21 Collect Inventory

You will get a window, as shown in Figure 9-22, where you can check if the inventory was collected successfully.

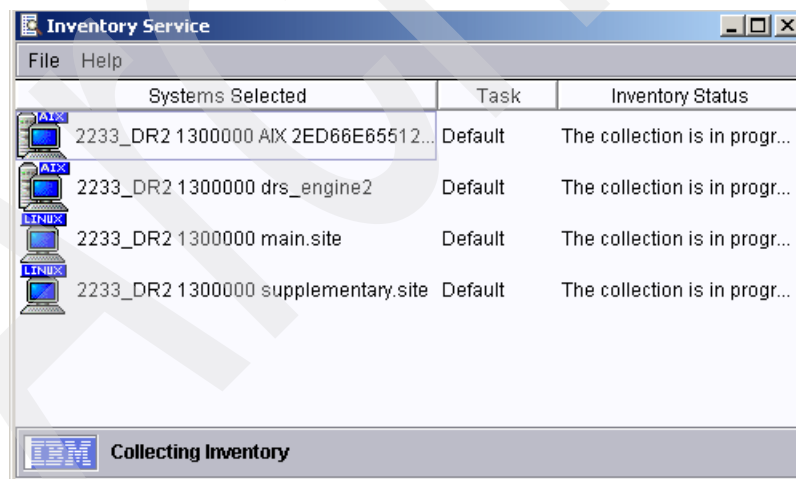


Figure 9-22 Status window for collecting the Inventory

After this is done, you can download, install, and configure the Electronic Service Agent.

## Install Electronic Service Agent for System x

The Electronic Service Agent for System x can be downloaded from the IBM Electronic support Web site. Follow the links from the IBM DR550 support Web site. Use the following link, then click **IBM Electronic Service Agent**, which should take you to Download. Then select **xSeries®**:

<http://www.ibm.com/support/electronic>

We illustrate here the installation of the Electronic Service Agent for xSeries (System x) on Windows.

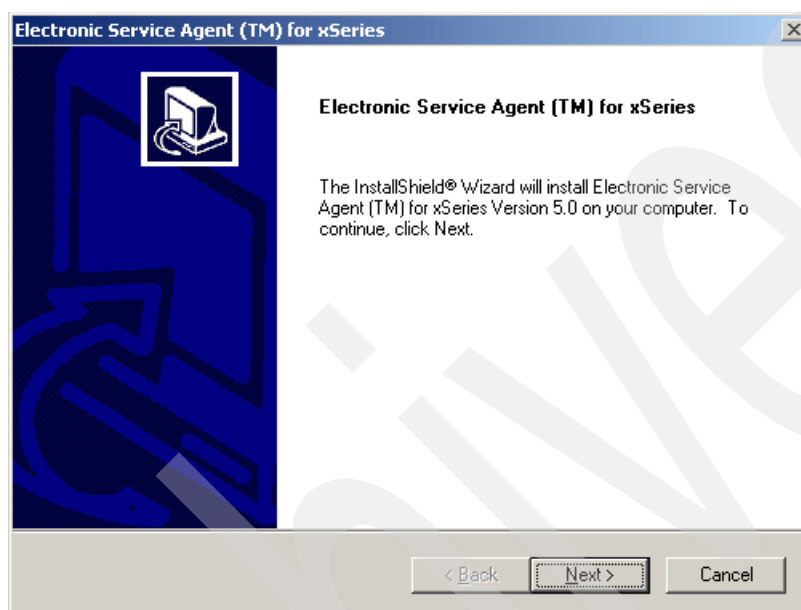


Figure 9-23 Electronic Service Agent for System x - Install

You are reminded, as we have explained above, that you must make sure that all System x servers are already added in the IBM Director Console and that you collected their inventory.




Figure 9-24 Electronic Service Agent for System x - Warning

Next you must agree the license agreement and to the fact that you are responsible for the communications charges to IBM.

Once the installation has completed, you have to restart the IBM Director Server.

Now you should see the Electronic Service showing as an additional task in the IBM Director Console window. This is visible in Figure 9-25 on page 367 (red circle).

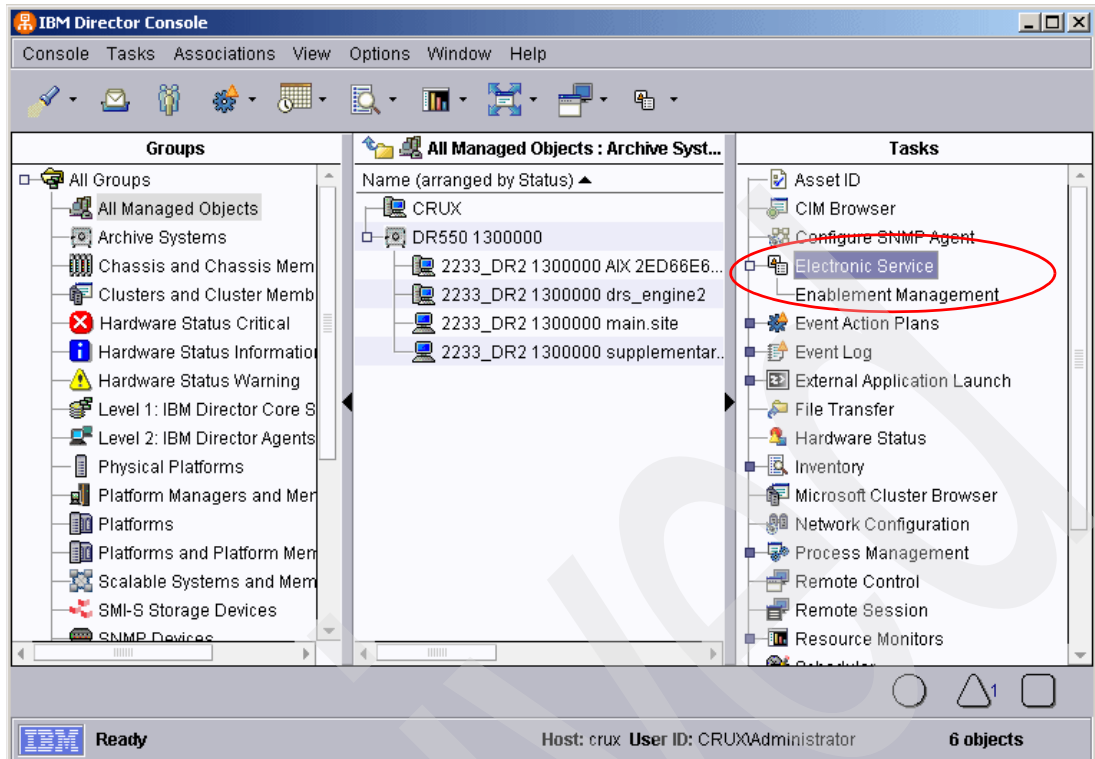


Figure 9-25 Electronic Service Task

## Configuring the Electronic Service Agent

There is a wizard available for configuring the Electronic Service Agent. You start the wizard by double clicking **Electronic Service** in the Director Console task pane.

The Electronic Service Agent Setup Wizard main window displays, as shown in Figure 9-26. Click **Next**.



Figure 9-26 Configuration Wizard for the Electronic Service

Fill out all the panels and then test the connection by clicking **Test Connection**, as shown in Figure 9-27.

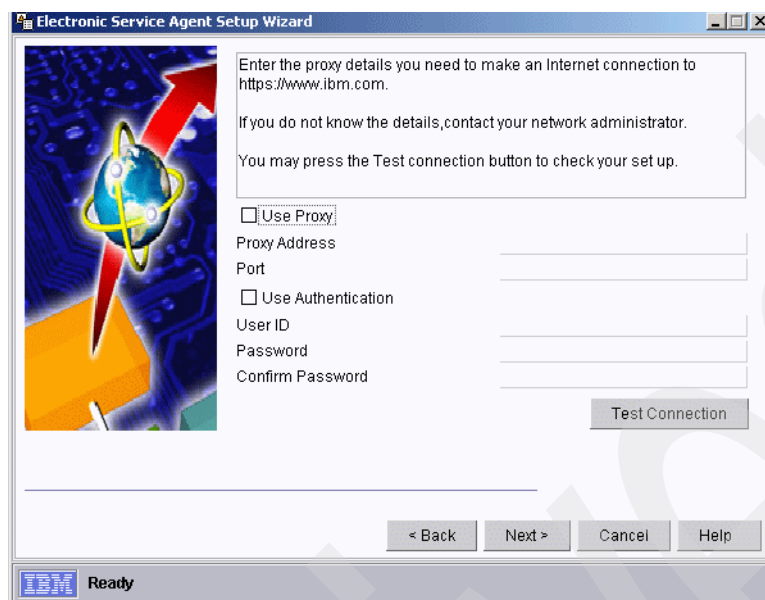


Figure 9-27 Electronic Service Agent Setup Wizard

Verify that the Connection Test was successful (as in Figure 9-28) before you proceed further with the configuration.



Figure 9-28 Electronic Service Agent Setup Wizard - Internet Connection Test

If the Internet Connection Test was successful, click **OK** and enter the required information in the next few windows displayed by the wizard to complete the agent configuration.

Click **Finish** in the last window shown in Figure 9-29 on page 369.

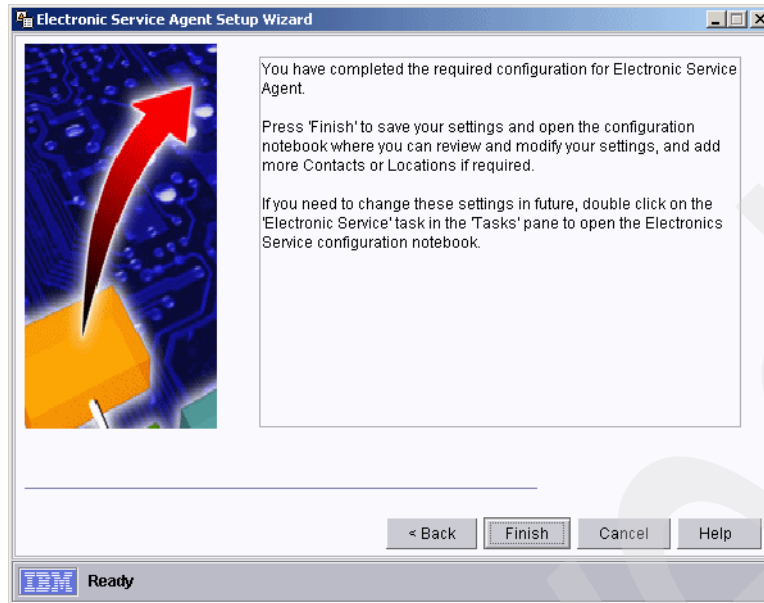


Figure 9-29 Electronic Service Agent Setup Wizard - Finish the Configuration

## Enablement Management

After configuring the agent, you must now enable the System x servers (that is, the FSG nodes in the case of the DR550) for which you want the call home functionality. To start the enablement procedure, double click **Enablement Management** in the Tasks pane of the IBM Director console (refer to Figure 9-25 on page 367).

The Enablement Management dialog displays, as shown in Figure 9-30.

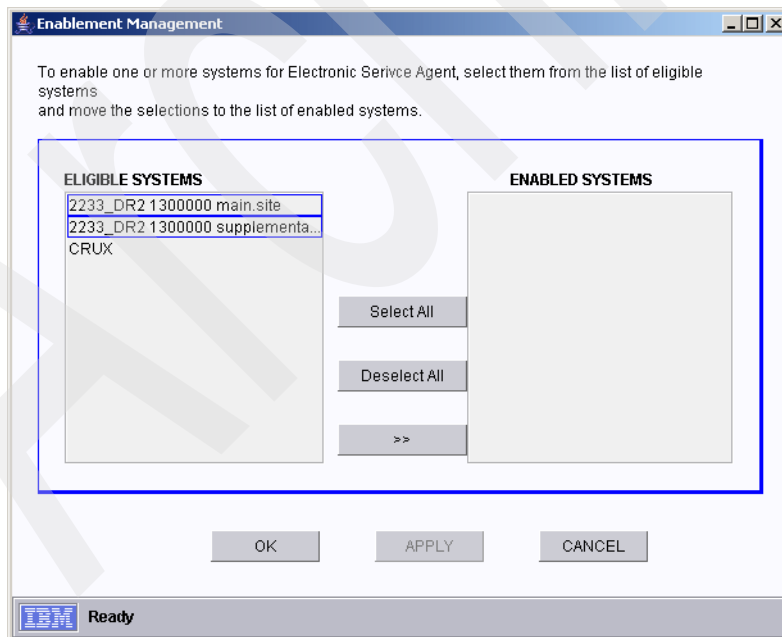


Figure 9-30 Select servers to enable

Select the FSG nodes and move the selected systems into the ENABLED SYSTEMS column and click **APPLY**.

When you click **APPLY**, a connection is attempted to IBM to check that the serial numbers of the systems (FSG nodes) that were added are indeed covered by warranty or service contracts, as shown in Figure 9-31.

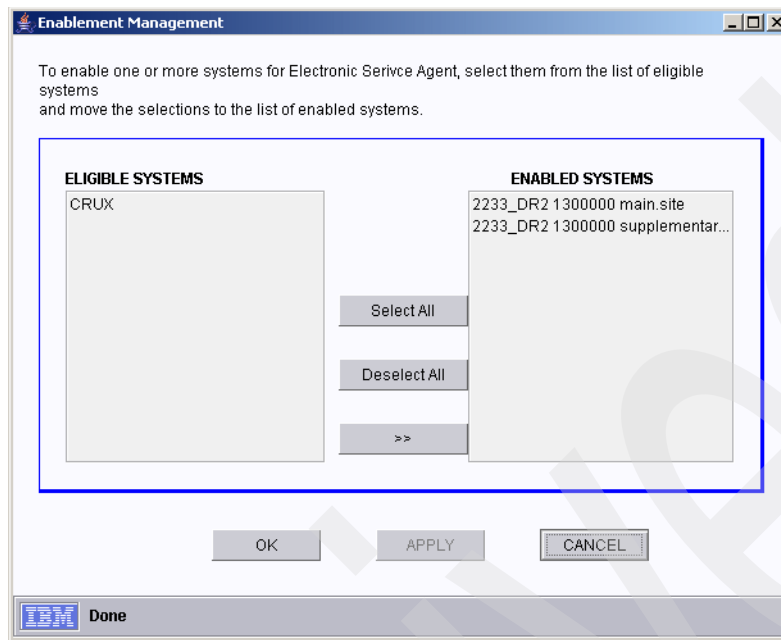


Figure 9-31 Systems are enabled to use eSA

**Note:** If you get a negative acknowledge during the enablement process, contact your IBM Service representative.

### Testing eSA for System x

When your System x servers are enabled, you can create a test call (recommended). In the IBM Director Console window, select one of the FSG servers, then right-click and select **Electronic Service** from the pop-up menu.

This should give you the Electronic Service Agent dialog window shown in Figure 9-32 on page 371.



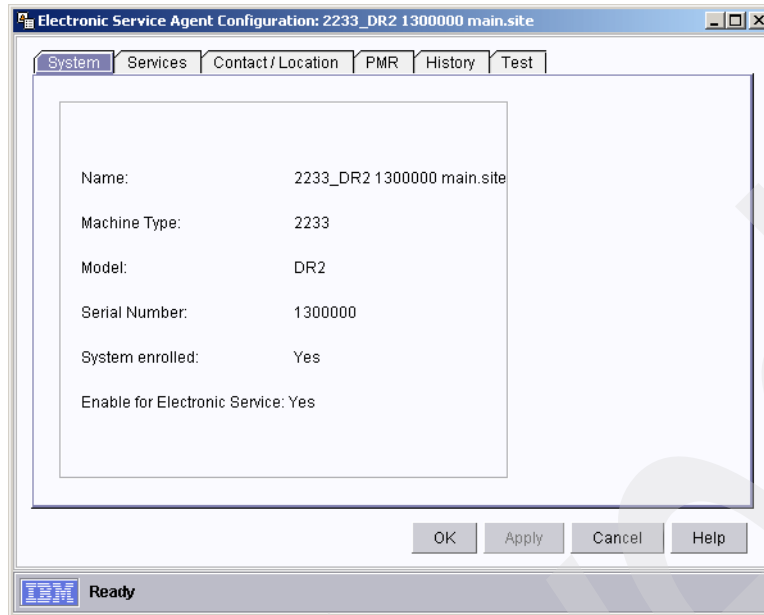


Figure 9-32 Electronic Service Agent - configure a System x server

Check the different tabs and set specific, appropriate options for this server. Then, go to the Test tab shown in Figure 9-33.

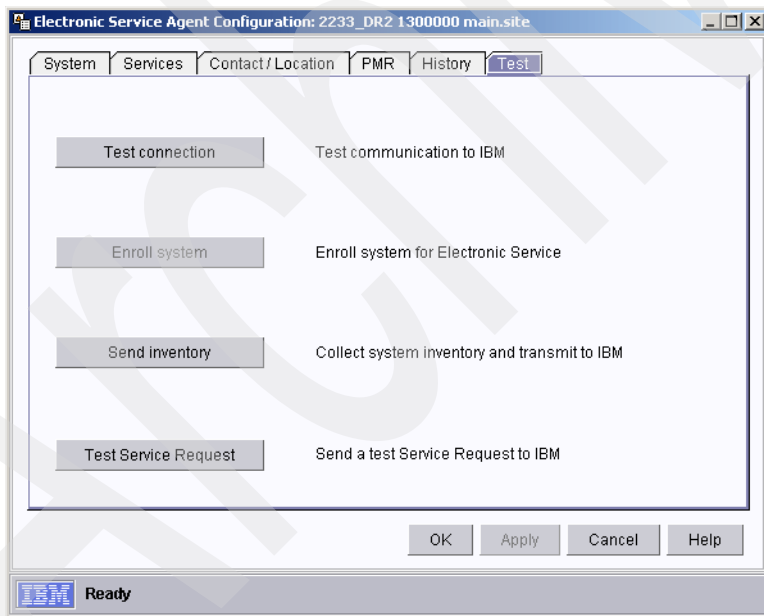


Figure 9-33 Electronic Service Agent - configure a System x server - Test Service Request

Here you can first do a connection test, or click directly on the Test Service Request, as shown in Figure 9-34.

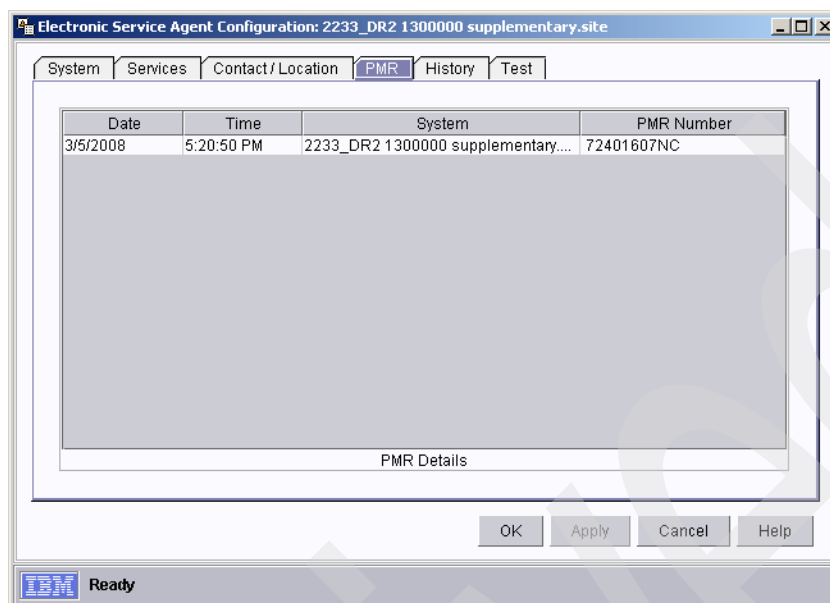


Figure 9-34 Electronic Service Agent - configure a System x server - PMR tab

Once you create a successful service request, you can see the call created and the associated number under the PMR tab.

Your IBM Director is now able to send service requests (call home) for this System x server.

## 9.4 RSM for DR550

This section explains the Remote Support Manager for DR550 Installation and Configuration. RSM is the recommended solution for alerting and call home support for the DR550 Storage controller (DS4200) subsystem. For completeness, this chapter includes a brief overview of RSM.

**Note:** The DS4000 Service Alert (Call Home) has been discontinued. The functionality is now provided by Remote Support Manager. There is a specific version of RSM for the DR550.

**Consideration:** To maintain the physical security of the RSM server, you must follow one of the procedures below:

- ▶ Order a second lockable rack, such as a 7014-S25, 7014-T00, or 7014-T42 with the locking feature. Install the RSM server in the second lockable rack and run unbroken cables from the RSM server rack to the DR550 rack. There must be no routers, hubs, or switches between the RSM server and the DR550.
- ▶ Alternatively, open an RPQ to have the RSM server placed in the DR550 rack.

## 9.4.1 Introduction to RSM for DR550

The IBM Remote Support Manager for DR550 (RSM for DR550) is an application that installs on an IBM System x server running Novell SUSE Linux Enterprise Server 9, SUSE Linux Enterprise Server 10, or Red Hat Enterprise Linux 4 Advanced Server, and provides problem reporting and remote access for IBM Service for the IBM DR550 (machine type 2233).

The RSM for DR550 is an optional application and is not sold as a feature of the DR550. Customers are responsible for RSM hardware and software installation.

The problem reporting provided by the RSM for DR550 application automatically creates an entry in the IBM call management system for the 2233-DR1 or 2233-DR2 that reports a problem. This is the equivalent of placing a voice call to IBM Service for a problem. Once in the system, problems are responded to with the same priority as specified by the maintenance agreement in place for the product.

RSM for DR550 controls security for remote access (IBM Support) by managing hardware and software components of the server it is installed on. Once installed, the server should be considered to be a single purpose appliance for problem reporting and remote access support for your storage subsystem. Only applications approved by IBM should be installed. (Management of the internal firewall and other configuration changes made by the software might prevent other applications from working.)

Remote access is provided by an external modem attached to the server. This connection provides IBM Service with a command-line interface to the server. All bulk data transfers for logs and other problem determination files are sent to IBM through e-mail or by FTP using the server's Ethernet interface. Isolation of remote and local users of the system from other devices on your intranet is performed by an internal firewall that is managed by the RSM for DR550 application. Remote users do not have the ability to change any security features of the application. The modem is configured to only answer when Remote Access is enabled by the RSM for Storage application. You can manually enable and disable remote access, or you can choose to have remote access automatically enabled when a storage subsystem reports a problem. When remote access is enabled, a timer is started that will automatically disable remote access when it expires. The person identified as the primary contact for the RSM for Storage system is notified by e-mail whenever a change in the remote access settings occurs and all state changes are also written to the Security Log.

Monitoring of storage subsystems is performed by your existing SMclient running on the System p nodes of the DR550, which is configured to send SNMP traps to the Remote Support Manager when critical events are detected. Configuration of the management application is addressed later in this book.

The Remote Support Manager's user interface allows you to control and view the status of four management areas:

- ▶ System configuration
- ▶ Reporting
- ▶ Remote access
- ▶ Internal firewall

## 9.4.2 Connecting the Remote Support Manager

For the DR550 Model DR1, the following connections are required, as shown in Figure 9-35:

- ▶ Connect the RSM server to the DR550 Storage Controller using the E2 ports of both Ctrl A and Ctrl B of the DR550 Storage Controller.
- ▶ Connect the RSM server to the DR550 SSAM Server using port 1 of the Ethernet adapter in slot 5 of the DR550 SSAM Server.
- ▶ Connect the RSM server to the customer network.

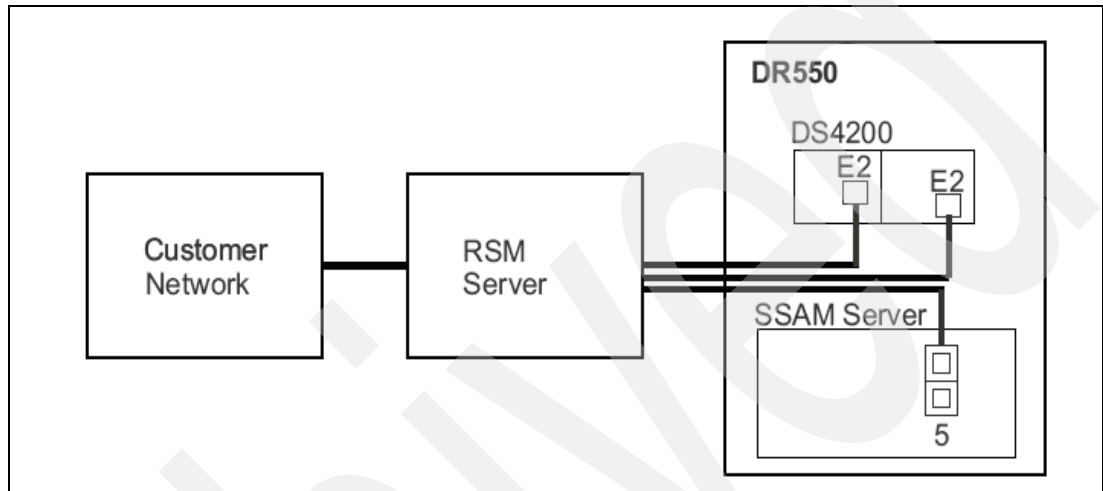


Figure 9-35 DR1 Ethernet connection

For DR550 DR2, the Remote Support Manager (RSM) server requires the following connections, as shown in Figure 9-36:

- ▶ Connect the RSM server directly to port 15 of the two DR550 internal Ethernet switches.
- ▶ Connect the RSM server to the customer network.

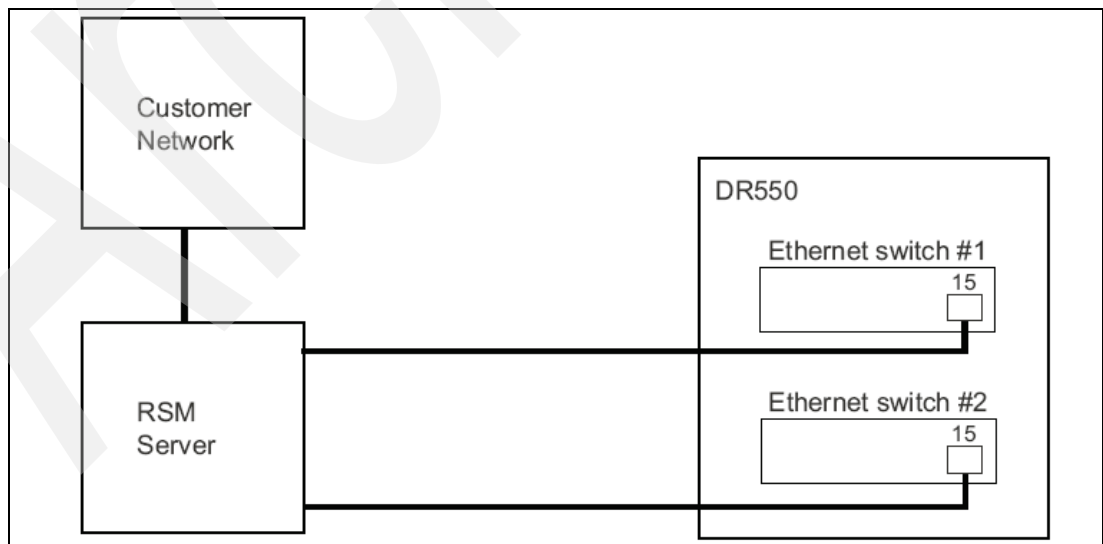


Figure 9-36 DR2 Ethernet connections

**Note:** For compliance, RSM only supports one DR550 at a time. If you have multiple DR550s or other DS4000 storage systems, you need to have multiple RSM servers. Only when the ERM feature is installed (Remote Mirroring between two DR550s) can you use the RSM for both DR550 Storage Controllers.

### 9.4.3 Server installation and configuration

This section describes the installation of RSM hardware and software.

#### Hardware

The RSM was designed to run on a System x (xSeries) server, and most System x servers can be used. Refer to the *RSM for Storage Compatibility Guide* found at [ftp://ftp.software.ibm.com/systems/support/system\\_x\\_pdf/rsm\\_for\\_storage-compatibility\\_guide-v3.5.pdf](ftp://ftp.software.ibm.com/systems/support/system_x_pdf/rsm_for_storage-compatibility_guide-v3.5.pdf) for the minimum server requirements, and a list of the specific servers that have been tested with the RSM software. If you choose to use a server that has not been specifically tested with the RSM software, the setup and configuration of the server's BIOS, and installation of the Linux OS, may be different than what is included in the documentation for the RSM software and you will need to contact IBM System x support for questions related to these other servers.

The following servers have been tested with the RSM for Storage software:

- ▶ IBM 4364 x3250
- ▶ IBM 8849 x306m

The RSM for Storage application is designed to work with an external modem attached to the first serial port. The functional requirements for the modem are minimal and most "Hayes-compatible" external modems can be used.

Be sure to use a modem and power cord appropriate for your country or region. The RSM for Storage application has been tested with the following modems:

- ▶ Multitech Multimodem II MT5600BA
- ▶ Multitech MultiModem ZBA MT5634ZBA

Refer to the *RSM for Storage Compatibility Guide* for updated information about which modems are supported. You will need to contact your IT support staff for installation and problem resolution related to the modem.

#### Software

The RSM for Storage application is supported on SUSE Linux Enterprise Server 9 Service Pack 3 from Novell, SUSE Linux Enterprise Server 10, or RHEL 4 Advanced Server from Red Hat. Refer to the *RSM for Storage Compatibility Guide* for current information or ordering one of the products. The RSM for Storage application receives SNMP traps from the Event Monitor included with IBM Storage Manager. The RSM for Storage application cannot be installed on the same system used to manage your storage network. Your DR550 has the IBM Storage Manager installed.

## 9.4.4 RSM installation and configuration

**Note:** DR550 requires a special version of RSM that is available only on the DR550 support Web site.

### Installing RSM

1. Go to the DR550 support Web site and download the DR550 Version 4.5 version of the RSM code:

<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/supportresources?brandind=5000028&familyind=5329490&taskind=1>

2. Go to the RSM Web site at and download the *IBM Remote Support Manager for Storage Planning, Installation and User's Guide*, GC26-7933:

<http://www.ibm.com/systems/storage/disk/rsm/index.html>

3. Follow the installation and configuration instructions in the *IBM Remote Support Manager for Storage Planning, Installation and User's Guide*, GC26-7933.

**Note:** During the activation step of the RSM installation, you will call the IBM Support Center. You must provide the DR550 machine type (2233) and model (DR1 or DR2) to get to the DR550 Support Center. They will help you to properly configure RSM for DR550.

### Additional RSM configuration for DR550

If the RSM interface is already running, click **Logout**. Otherwise, open the RSM interface by clicking the Manage icon. From the main RSM menu, select either **View System Configuration** or **Manage System Configuration**. The option available depends on the user ID that you used to log in. When prompted, log in as the lservice user. Select the storage subsystem name to bring up the window shown in Figure 9-37 on page 377. The values shown in the figures are examples only. You will see the information you created when the storage subsystems were configured during RSM installation.

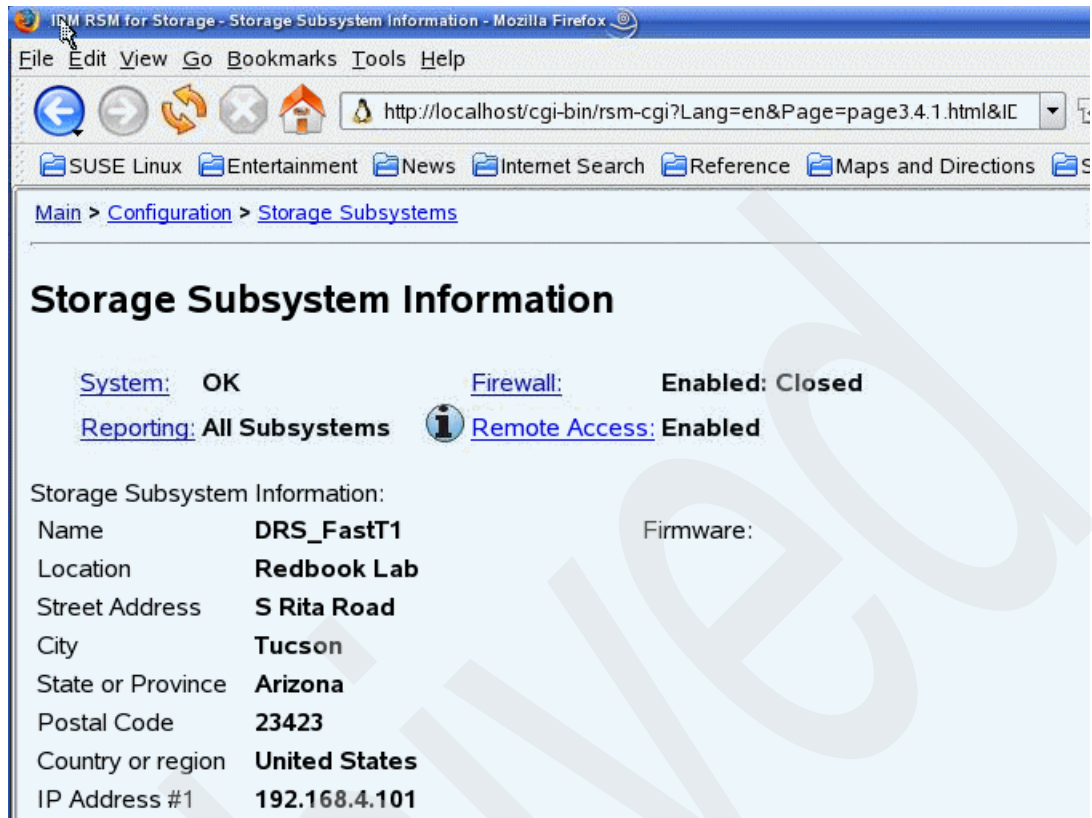


Figure 9-37 Storage Subsystem Information window

Select the **Update** link in the Part of Solution setting. The window is shown in Figure 9-38.

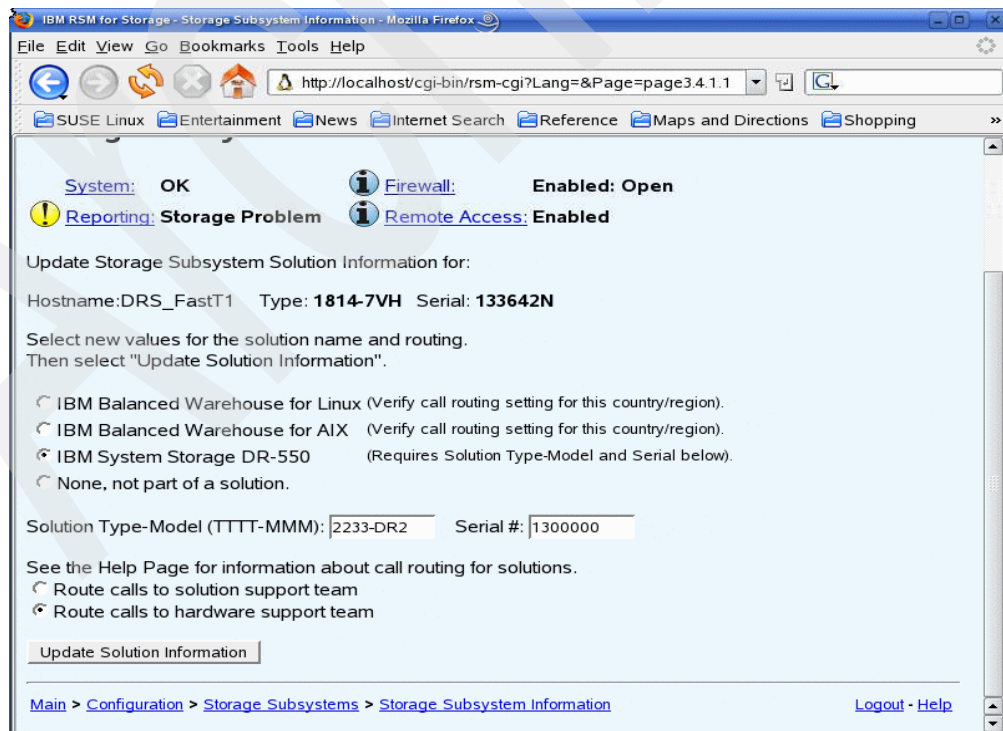


Figure 9-38 RSM 2233 Solution Type window

Select **IBM System Storage DR550**, and enter 2233 in the Solution Type-Model field. Enter the DR550 serial number in the Serial field.

Click the **Update Solution Information** button and click **Logout** to exit RSM. The resulting window is shown in Figure 9-38 on page 377.

## Configure SNMP traps to be sent to the RSM

The RSM for Storage system sends alerts to IBM Service when SNMP traps with event information are received from the subsystems being monitored. For DS4000 subsystems, SNMP traps are sent by the IBM Storage Manager client or the IBM Storage Manager's Event Monitor service. As the Storage Manager Client may not always be running, it is recommended that the Event Monitor be installed (It is already installed with the DR550).

To send traps to RSM, configure RSM as a destination for SNMP traps in the Storage Manager client. This configuration will also be applied to the Event Monitor on that same system.

- ▶ In the Enterprise Management window, right-click the Storage Subsystem and select **Configure Alerts®**, as shown in Figure 9-39.
- ▶ Click the **Configure Alerts** button.
- ▶ Select **All storage subsystems**, and then click **OK**.
- ▶ Select the **SNMP** tab.
- ▶ Enter the IP address of the RSM for Storage system in the Trap Destination window, as shown in Figure 9-40 on page 379
- ▶ Click **Add**.
- ▶ Click **Validate** to send a test trap the RSM for Storage system, as shown in Figure 9-41 on page 379.

The RSM for Storage application ignores test alerts such as these, but you can verify that it was received by viewing the RSM for Storage Activity Log. Click the magnifying glass icon on the desktop labeled Activity Log. The reception of the test alert and the fact that it was ignored should be logged.

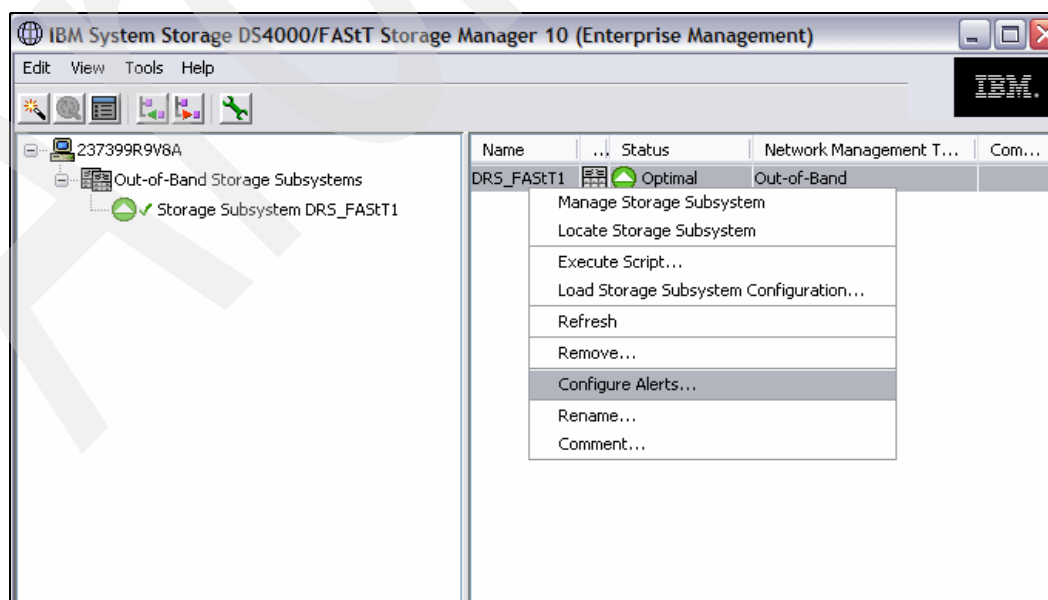


Figure 9-39 Enterprise Management window



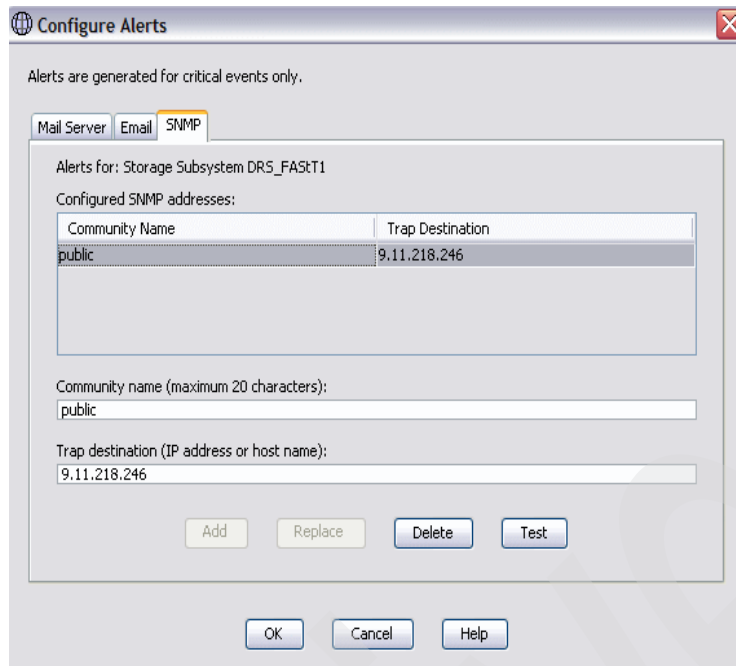


Figure 9-40 Configure Alerts window




Figure 9-41 Alert Validation window

### 9.4.5 RSM users and terminology

During installation of the RSM for Storage application, desktops for four users are configured: root, admin, rservice, and lservice

► admin

This is the customer account for managing the operation of the RSM application. The admin user desktop has a yellow background and contains icons for:

- Master Help browser page
- Management browser page
- Activity Log

This is the normal login for RSM management. The password for this user is set by the root user of the system. This account cannot be used or accessed by IBM support. Only customer admin person can use this account.

- ▶ **lservice**

This is the IBM Service login for on site service. The lservice user desktop is blue and contains the same icons as for the admin user. The lservice user has some restrictions on the directories and files it can access. This is to prevent this user from changing any of the RSM system security settings.

- ▶ **root**

It is the administrative user for the system. It can perform any task and access any file. The root user desktop has a red background. In addition to the icons placed on the desktop for other users, a System Log icon will display the most recent entries in the Linux system log.

- ▶ **rservice**

This is the Remote Service user that is used exclusively for remote access to the system. The user ID reserved for IBM Support remote access (rservice) is only valid when Remote Access is enabled. Attempts to log in using the root, admin, or lservice user IDs are rejected.

## Permissions and privileges

The Switch User (**su**) command is disabled to prevent a normal user from attempting to become root and have unrestricted access to the system. The RSM for DR550 Storage application makes other changes in program and directory permissions to limit what programs and files these users can access.

The remote connection made by IBM into the RSM for Storage system is a console interface and programs that can initiate an IP connection on this interface are removed from the system during installation of the RSM for Storage application. The following programs are removed from the system during installation of the RSM for Storage application:

- ▶ rcp
- ▶ rsh
- ▶ rexec
- ▶ tftpd
- ▶ vsftpd
- ▶ telnetd
- ▶ rlogin

The only daemons (or services) running on the RSM for Storage system are snmptrapd, sshd, and httpd, which listen for SNMP traps, secure shell session requests, and HTTP(S) requests, respectively.

## RSM terminology

The following terms are used to describe RSM operations:

**Event**

Problem reported to RSM by the DR550 Storage Manager application.

**Alert**

Almost all critical events will be sent to IBM. The SMclient or the Event Monitor take note from the critical events and sent an SNMP trap to the RSM Server. Afterwards, the alert will be sent as an e-mail to IBM and generates an Problem Record on the DR550 Support Center.

The following events are not reported to IBM Service. However, you will receive an e-mail from the RSM for Storage system when these events occur:

- 6200: FlashCopy® repository logical drive capacity threshold.
- 6202: FlashCopy logical drive failed.

- 4011L: Logical Drive not on a preferred path due to ADT/RDAC.
- None: The persistent monitor running on Host xxxxxxxx cannot reach the indicated Storage Subsystem.
- None: The persistent monitor running on Host xxxxxxxx can now reach the indicated Storage Subsystem.

#### Heartbeat

Customers receive a noon (local time) daily status e-mail telling them that the RSM is operational.

#### Problem Data

The DS4000 Support Bundle (or zipped file containing RLS counters, Profile, and MEL files). Collected when an alert occurs. The zipped file can be used by IBM Support for problem determination.

Figure 9-42 gives an overview of how the RSM process works.

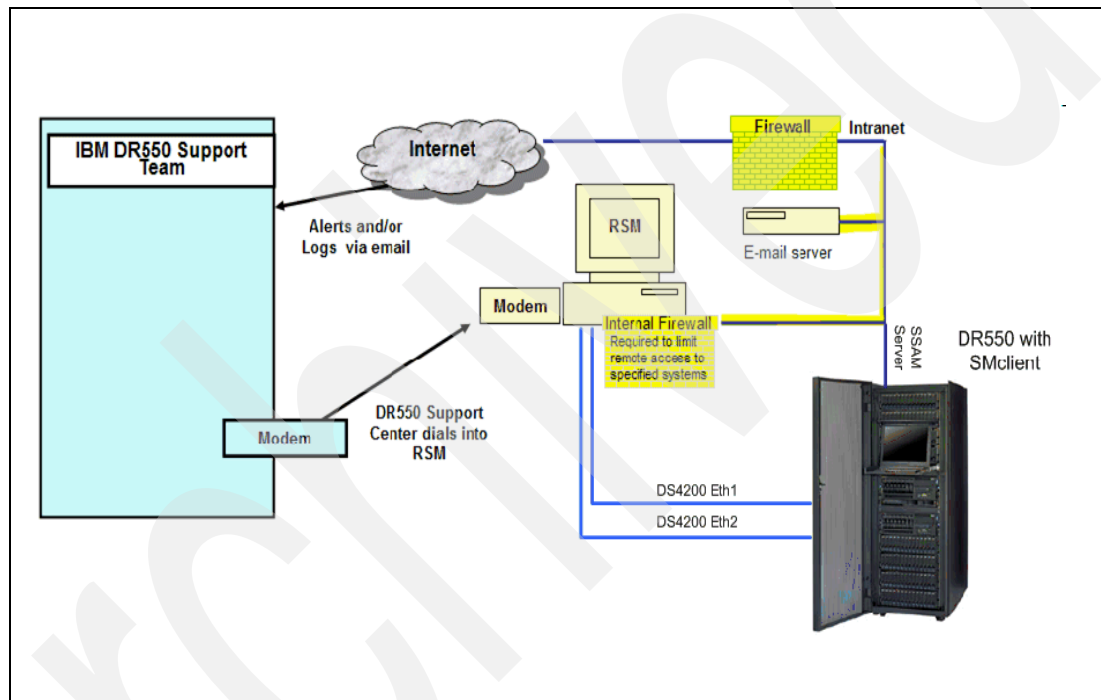


Figure 9-42 RSM process

### 9.4.6 RSM modem connectivity

During the RSM configuration, you normally perform a connection test. Part of this test is to verify that the RSM for Storage software can communicate with the modem. If the configuration test completed without errors, then the modem is connected to the correct serial port and is communicating with the RSM for Storage setup.

On the RSM Access page, click **Enable Remote Access**, shown in Figure 9-43 on page 382. This will enable the modem to answer when called. Verify modem connectivity by calling the modem phone number from a voice phone.

**IBM Remote Support Manager for Storage**

[Main](#) [Logout](#) - [Refresh status](#) - [Help](#)

**Remote Access**

[System:](#) **Problem**      [Firewall:](#)      **Enabled: Closed**  
[Reporting:](#) **Suspended**      [Remote Access:](#) **Disabled**

**Manage Remote Access**

Remote Access is: **Disabled** [Enable Remote Access](#)

The Remote Access Timeout is: **N/A**

Select one of the following to change the current (and default) Remote Access Timeout:

☐ 12 hours    ☐ 24 hours    ☒ 36 hours    ☐ 48 hours    ☐ 72 hours    ☐ 96 hours    [Update Timeout Value](#)

Figure 9-43 Enable Remote Access window

Before the Remote Support Manager for Storage will send alerts to IBM, it must be activated by contacting IBM Service. This is also the last step in verifying correct operation of the RSM for Storage setup. On the RSM for Storage user interface, click **Remote Access** and enable remote access.

## Reporting and alerts

Normally all configured storage subsystems will be enabled for reporting. This means that the RSM will accept and process any events related to the subsystem. If you are relocating a subsystem or performing any task that might generate events (that IBM does not need to respond to), you can disable reporting until the subsystem is fully operational again. While you are making configuration changes to the RSM for Storage application, the Reporting Status may be “Suspended”. This is a reminder that no events will be processed by the RSM for Storage system while any configuration problems exist. This page will show a summary of all alerts being tracked by the RSM for Storage application and allow you to view details about alerts that are active for each subsystem. When a subsystem first reports a problem, an alert is sent to IBM Service. Once IBM has been alerted to a problem, additional alerts for that subsystem are held at the RSM for Storage system. IBM will respond to the alert either by connecting to the RSM for Storage system or calling you. If IBM Service connects to the RSM for Storage system, they will either acknowledge or close all of the alerts for the subsystem. Alerts are acknowledged to indicate they have been seen by IBM Service but work on the problem has not been completed. Closing all of the alerts for a subsystem indicates that service is complete. When all alerts for a subsystem are closed, the RSM for Storage application will consider the next event from that subsystem to be a new problem and an alert will be sent to IBM Service. The Reporting and Alerts page will show the number of alerts sent, acknowledged, and pending for each subsystem that has active alerts. Pending alerts are ones that are candidates to be sent to IBM Service, but are being held at the RSM for Storage system for one of three reasons:

- Holding

Another alert has already been sent to IBM Service for the subsystem.

► Queued

The RSM for Storage application attempted to send the alert, but received an error. The most likely cause is a network problem that prevents the RSM for Storage application from reaching the SMTP server. The RSM for Storage application will attempt to re-send the alert every few minutes.

► Waiting

IBM Service was remotely connected to the RSM for Storage system when the alert occurred. If all other alerts have been closed and the remote user disconnects without acknowledging this alert, it will then be sent to IBM Service as a new problem.

In Figure 9-44, you can see an example for the Reporting and Alerts window.

**IBM Remote Support Manager for Storage**

[Main](#) [Logout](#) - [Refresh status](#) - [Help](#)

---

### Reporting and Alerts

[System](#): **Problem**      [Firewall](#): **Enabled: Closed**  
 [Reporting](#): **Suspended**      [Remote Access](#): **Disabled**

Products with Active Alerts:	<b>0</b>	Alerts sent to IBM:	<b>0</b>
Products with Reporting Enabled:	<b>1</b>	Alerts Pending:	<b>0</b>
Products with Reporting Disabled:	<b>0</b>	Alerts Acknowledged:	<b>0</b>

NOTE: Reporting has been suspended because of a configuration problem.

[View/Change](#) reporting state for each subsystem

Figure 9-44 Reporting and Alert window

On the Remote Access management page, you have the option of allowing remote access to be automatically enabled when an alert is sent to IBM. This allows IBM to connect to the RSM for Storage system without needing to first speak with the contact person for the system. If you choose to disable this function, IBM Service will contact you to have remote access enabled manually when an alert is received.

The acknowledge state of an alert is an indication that IBM Service has seen or is aware of the alert. When IBM Service dials into the RSM for Storage system, they will view the existing alerts and acknowledge them. In some situations, IBM Service may not require remote access to the RSM or subsystem in order to determine the cause of a problem and may call you to discuss the problem resolution. You can also acknowledge (or close) alerts.

The RSM for Storage application sends a single alert to IBM Service for each subsystem. Additional alerts that occur after the first alert is sent to IBM are available for examination in the RSM for Storage application. When all of these active alerts for a subsystem are closed, the next alert will again be sent to IBM Service. When intermittent problems are occurring, it may be necessary for all alerts for the subsystem to be closed, so the RSM for Storage application can be used to automatically notify IBM when the problem recurs. IBM Service will normally close all alerts for a subsystem then dialed into the RSM for Storage system, but resolution of the problem does not require dialing into the RSM for Storage system. You might be asked to close the existing alerts so that the RSM can resume sending alerts to IBM. Closing an alert in the RSM for Storage application does not close the problem report with IBM Service.

## 9.5 IBM Director agent for AIX

To provide an enhanced level of functionality, IBM Director Agent (Level-2) for AIX is now pre-installed on the DR550 SSAM servers, starting with Version 4.5. Table 9-1 shows the list of tasks supported with the Level 2 agent.

Table 9-1 Tasks available by IBM Director Agent Level

IBM Director base tasks	Level 0	Level 1	Level 2
Asset ID™	No	No	Yes
CIM Browser	No	No	Yes
Configure SNMP Agent	No	No	Yes
Event Action Plans	Yes	Yes	Yes
Event log	Yes	Yes	Yes
External application launch <sup>a</sup>	Yes	Yes	Yes
File transfer <sup>b</sup>	No	No	Yes
Hardware status	No	Yes	Yes
Inventory	Yes	Yes	Yes
Microsoft Cluster Browser	No	No	Yes
Network configuration	No	No	Yes
Power Management™	Yes	Yes	Yes
Process Management <sup>c</sup>	No	No	Yes
Remote Control	No	No	Yes
Remote Session	Yes	Yes	Yes
Resource Monitors	No	No	Yes
Scheduler	Yes	Yes	Yes
Server Configuration Manager	No	Yes	Yes
SNMP Browser <sup>d</sup>	Yes	Yes	Yes
Software Distribution	Yes	Yes	Yes
System Accounts	No	No	Yes
Update Manager	Yes	Yes	Yes

a. Varies with application.

b. Disabled with the DR550 if compliance script is applied.

c. Disabled with the DR550 if compliance script is applied.

d. SNMP service/daemon must be installed and running on managed system.

## 9.6 CIM agent for DR5550 Storage Controller

This chapter provides Information about the SMI-S provider (CIM Agent) for the DR550 Storage Controller.

### 9.6.1 Introduction

The Storage Management Initiative - Specification (SMI-S), driven by the Storage Networking Industry Association (SNIA), is an industry standard to access and manage storage devices. SMI-S expands on the CIM and WBEM standards, using XML over HTTP to communicate between storage management applications and the devices they manage.

The DR550 Storage Controller is not by default a CIM-ready device. Therefore, a DS4000 Device provider (acting as the CIM Agent) must be installed (on a host system) to bridge communications between the DR550 Storage Controller and the IBM Director. This is necessary because the storage subsystem and IBM Director do not speak the same language. The device provider (agent) for the DR550 Storage Controller is provided by Engenio and is referred to as the SMI-S provider.

### 9.6.2 Storage Management Initiative - Specification

The Storage Networking Industry Association (SNIA) has fully adopted and enhanced the CIM for Storage Management in its Storage Management Initiative - Specification (SMI-S). SMI-S was launched in mid-2002 to create and develop a universal open interface for managing storage devices, including storage networks. The idea behind SMI-S is to standardize the management interfaces so that management applications can use these and provide cross-device management. This means that a newly introduced device can be immediately managed if it conforms to the standards.

If you are interested in finding out what is new in SMI-S or just general information about it, you can visit the following Web site:

[http://www.snia.org/tech\\_activities/standards/curr\\_standards/smi/](http://www.snia.org/tech_activities/standards/curr_standards/smi/)

### 9.6.3 Prerequisites for using SMI-S

You must have the CIMOM supported firmware level on the storage devices (DR550 Version 4.5). If you have an incorrect version of the firmware, you might not be able to discover and manage the storage devices that have the incorrect level of software or firmware installed. CIMOMs by default should use either port 5988 (HTTP) or port 5989 (HTTPS) for communication.

### 9.6.4 Installing CIM agent for IBM DS4000

The CIM Agent for the DS4000 family is provided by Engenio and it is called the Engenio SANtricity SMI-S Provider. In DR550 Version 4.5, the CIM agent is already installed and ready to use on each DR550 SSAM Server. When needed, the Engenio SMI-S Provider version can be downloaded from the following Web site:

[http://www.lsi.com/storage\\_home/products\\_home/external RAID/management\\_software/smis\\_provider/index.html](http://www.lsi.com/storage_home/products_home/external RAID/management_software/smis_provider/index.html)

## 9.6.5 Engenio SMI-S Provider service availability

You can verify that the Engenio SMI-S Provider service has started or stopped with the following commands (see Figure 9-45):

```
cd /opt/engenio/SMI_SProvider/bin
./launch start - to start
./launch stop - to stop
./launch restart - to restart
./launch status - for status
```

```
drs_engine1 </opt/engenio/SMI_SProvider/bin>
./launch status
The Provider is running, PID is 397510
```

Figure 9-45 Provider Status Check

In the default installation directory, the installer has created a file named arrayhost.txt in the path /opt/engenio/SMI\_SProvider/bin/arrayhosts.txt (see Figure 9-46):.

```
drs_engine1 </opt/engenio/SMI_SProvider/bin>
cat arrayhosts.txt
192.168.4.101
```

Figure 9-46 arrayhosts.txt file

In this file, the IP addresses of the DR550 Storage Controller unit can be reviewed, added, or edited. After editing this file, you must restart the service. The configuration is stored in a file named opt/engenio/SMI\_SProvider/bin/providerStore.

Every time you change anything with the registered DS4000/FaStT Controllers and restart the Engenio CIMOM, and when you make a new discovery, the providerStore and arrayhosts.txt are updated with a new time stamp.

**Important:** After failover to another node, it is necessary to run the following script to restart the provider service:

```
/opt/IBM/ISS/resetSMIS.sh
```

The script will be executed automatically by HACMP.

Otherwise, the IBM Director cannot reach the DR550 SSAM Server. This is most important for HACMP systems, because it is configured for IP replacement.

## 9.6.6 IBM Director Storage implementation

If the DR550 Storage Controller is not visible in IBM Director, you may need to create the object in order for IBM Director Server to discover it. To do this:

1. Select **Console** → **New** → **Managed Object** → **SMI-S Storage Devices**.
2. In the pop-up window, enter the IP address of the SSAM Server in the DR550. Leave the other fields blank. IBM Director searches for the device and creates an object representing the DR550 storage system.



### Viewing components status

IBM Director provides status for all components, called Managed Objects, that it has discovered. In the IBM Director Console, you can view the hardware status and events that have been received by the IBM Director Server. To refresh the status of a Managed Object, right-click the Managed Object and select **Presence Check**.

The small lock icon to the left of an object name indicates that access to the system is blocked because authentication information has not been provided. To unlock a system, right-click the object and select **Request Access**, as shown in Figure 9-47. IBM Director requires root access to work correctly. Once authentication has completed successfully, the lock icon disappears. Each system requires authentication only once per IBM Director Server.

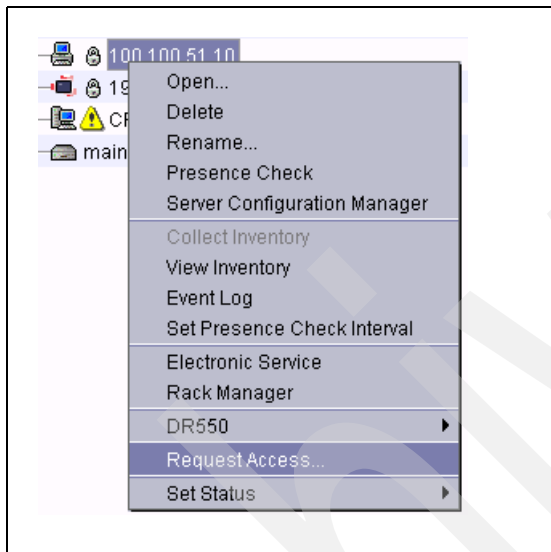
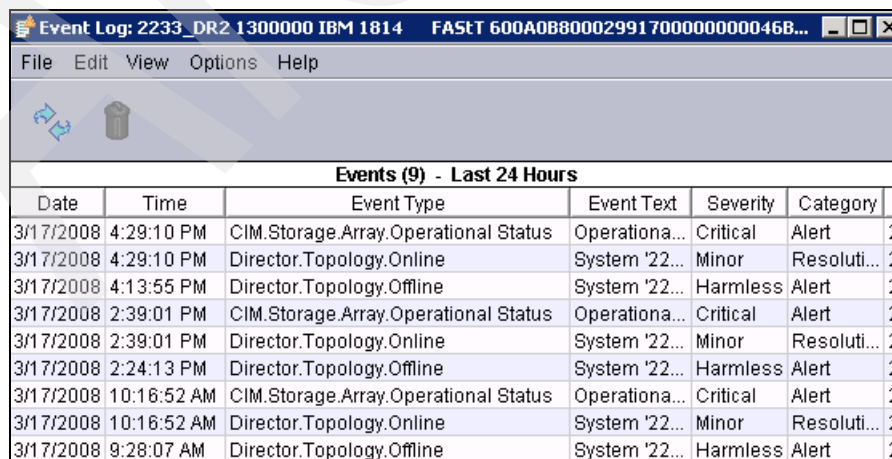


Figure 9-47 Request Access window

### Viewing event logs

You can view event logs as follows: Right-click the target object and select **Event Log**. Drag the Event Log icon from the Tasks pane and drop it on the target object. Select **Tasks** → **Event Log** and then select the desired target object. Click the red 'X' or yellow '!' icon in the lower left corner of the IBM Director Console window. The Event Log window is shown in Figure 9-48.



Events (9) - Last 24 Hours						
Date	Time	Event Type	Event Text	Severity	Category	
3/17/2008	4:29:10 PM	CIM.Storage.Array.Operational Status	Operationa...	Critical	Alert	2
3/17/2008	4:29:10 PM	Director.Topology.Online	System '22...	Minor	Resoluti...	2
3/17/2008	4:13:55 PM	Director.Topology.Offline	System '22...	Harmless	Alert	2
3/17/2008	2:39:01 PM	CIM.Storage.Array.Operational Status	Operationa...	Critical	Alert	2
3/17/2008	2:39:01 PM	Director.Topology.Online	System '22...	Minor	Resoluti...	2
3/17/2008	2:24:13 PM	Director.Topology.Offline	System '22...	Harmless	Alert	2
3/17/2008	10:16:52 AM	CIM.Storage.Array.Operational Status	Operationa...	Critical	Alert	2
3/17/2008	10:16:52 AM	Director.Topology.Online	System '22...	Minor	Resoluti...	2
3/17/2008	9:28:07 AM	Director.Topology.Offline	System '22...	Harmless	Alert	2

Figure 9-48 IBM Director Event Log window

## Hardware monitoring

You can monitor the hardware as follows: Right-click the target object and select **Hardware Status**. Drag the Hardware Status icon from the Tasks pane and drop it on the target object. Select **Tasks** → **Hardware Status**. If no target object is selected, all hardware with an error condition is displayed. The Hardware Status window, shown in Figure 9-49, displays error event messages that contain additional information for determining the cause of the problem.

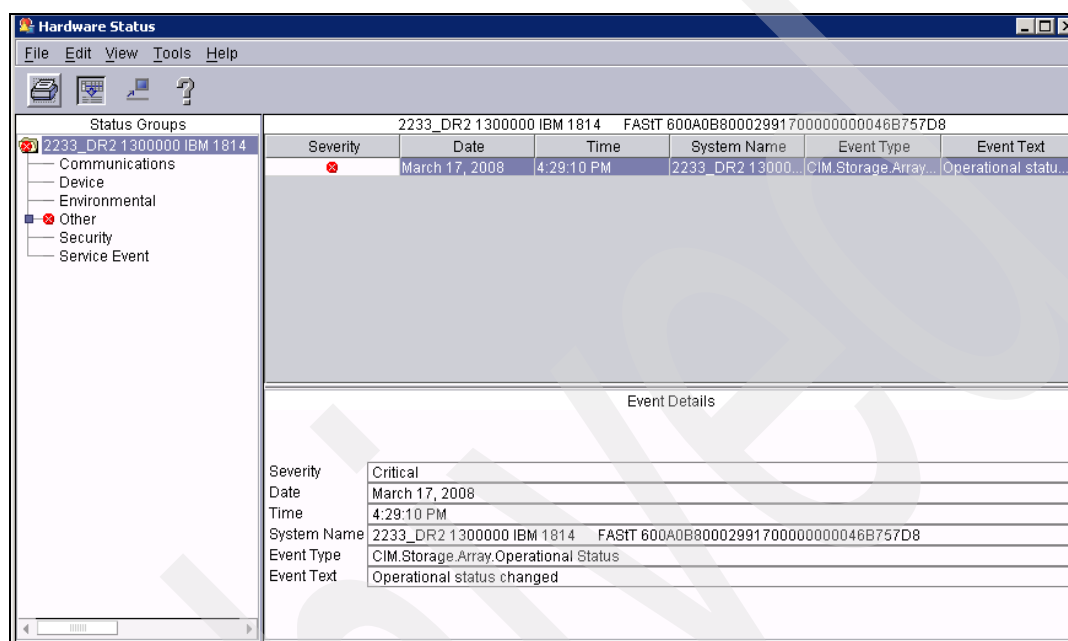


Figure 9-49 IBM Director Hardware Status window

To display the Status of the DR550 Storage Controller, double-click the manage object in the IBM Director Main menu. The Storage System Information window will be displayed as shown in Figure 9-50.

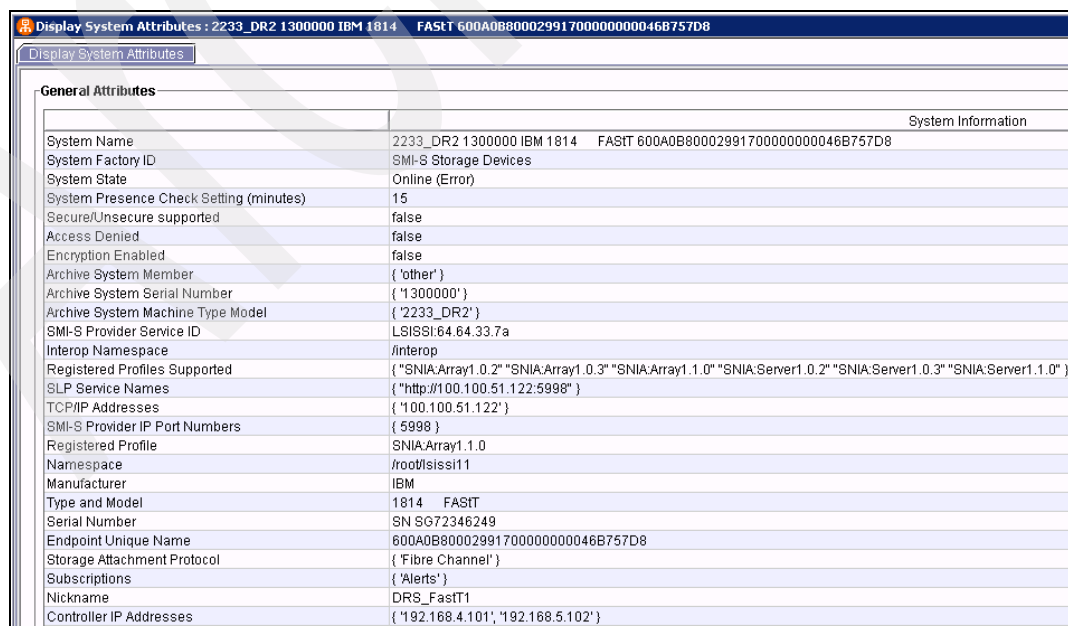


Figure 9-50 System Information window

## DR550 SNMP monitoring

This chapter discusses monitoring with Simple Network Management Protocol (SNMP) because this protocol is generally supported by all components of the DR550. You will find a high level introduction to SNMP as well as step-by-step configuration guides for SNMP for individual components of the system. The DR550 core software application, IBM System Storage Archive Manager (or SSAM), is also SNMP-enabled.

This chapter discusses the process of generating and processing SNMP events (traps). We have documented an implementation in which we use a two-tier SNMP management approach. One SNMP manager (IBM Director) is used to only present important and critical events to system administrators. The other SNMP manager (Net-SNMP) receives all the events from the managed devices and filters out duplicates and non-critical events before forwarding to the level one SNMP manager.

**Important:** The procedures and scripts presented and documented in this chapter are not part of the DR550 supported offering. They are provided as examples only.

## 10.1 Simple Network Management Protocol (SNMP) overview

This section explains the general SNMP functionality, the SNMP components, and its mode of operation.

SNMP is an industry-standard set of functions for monitoring and managing TCP/IP-based networks. SNMP includes a protocol, a database specification, and a set of data objects. A set of data objects forms a Management Information Base (MIB). SNMP provides a standard MIB that includes information such as IP addresses and the number of active TCP connections. The actual MIB definitions are encoded into the SNMP agent running on a system.

MIB-2 is the Internet standard MIB that defines over 100 TCP/IP specific objects, including configuration and statistical information, such as:

- ▶ Information about interfaces
- ▶ Address translation
- ▶ IP, ICMP (Internet-control message protocol), TCP, and UDP

SNMP can be extended through the use of the SNMP Multiplexing protocol (the SMUX protocol) to include enterprise-specific MIBs that contain information related to a specific environment or application. A management agent (a SMUX peer daemon) retrieves and maintains information about the objects defined in its MIB, and passes this information onto a specialized network monitor or network management station (NMS) such as Director Server.

The SNMP protocol defines two terms, agent and manager, instead of the client and server used in many other TCP/IP protocols.

### SNMP agent

An SNMP agent is a daemon process that provides access to the MIB objects on IP hosts on which the agent is running. The agent can receive SNMP GET requests or send SNMP SET requests from SNMP managers and can send SNMP trap requests to SNMP managers. On AIX, the SNMP agent is implemented as /usr/sbin/snmpd (see Example 10-30 on page 418).

Agents send traps to the SNMP manager to notify that a particular condition exists on the agent system, such as the occurrence of an error. In addition, the SNMP manager generates traps when it detects status changes or other unusual conditions while polling network objects.

### SNMP manager

An SNMP manager can be implemented in two ways:

- ▶ An SNMP manager can be implemented as a simple command tool that can collect information from SNMP agents.
- ▶ An SNMP manager also can be composed of multiple daemon processes and database applications. This type of complex SNMP manager provides you with monitoring functions using SNMP. It typically has a graphical user interface for operators.

The SNMP manager gathers information from SNMP agents and accepts trap requests sent by SNMP agents.

IBM Director is an example of an SNMP manager with a GUI interface and we recommend the use of IBM Director as the SNMP manager for your DR550 devices as an enterprise class Tier 1 SNMP manager.

NET-SNMP is another example of an SNMP management application. NET-SNMP is highly customizable and expandable command-line based tool with powerful trap processing capabilities. This tool comes pre-installed on SLES 10 servers and can be used to complement IBM Director SNMP manager functionality.

## SNMP trap

A trap is a message sent from an SNMP agent to an SNMP manager without a specific request from the SNMP manager.

SNMP defines six generic types of traps and allows the definition of enterprise-specific traps. The trap structure conveys the following information to the SNMP manager:

- ▶ Agent's object that was affected
- ▶ IP address of the agent that sent the trap
- ▶ Event description (either a generic trap or enterprise-specific trap, including trap number)
- ▶ Time stamp
- ▶ Optional enterprise-specific trap identification
- ▶ List of variables describing the trap

## SNMP communication

The SNMP manager sends SNMP get, get-next, or set requests to SNMP agents, which listen on UDP port 161, and the agents send back a reply to the manager. The SNMP agent can be implemented on any kind of IP host, such as UNIX workstations, routers, and network devices. In the case of the DR550, an SNMP agent can be configured on all hardware components and some SNMP-enabled applications.

You can gather various information about the specific IP hosts by sending the SNMP get and get-next requests, and you can update the configuration of IP hosts by sending the SNMP set request.

The SNMP agent can send SNMP trap requests to SNMP managers, which listen on UDP port 162. The SNMP trap1 requests sent from SNMP agents can be used to send warning, alert, or error notification messages to SNMP managers.

Note that you can configure an SNMP agent to send SNMP trap requests to multiple SNMP managers.

**Tip:** As depicted in Figure 10-1, we use Net-SNMP as a tier 2 SNMP manager. We use the RSM for DR550 server to run NET-SNMP, because its the only external server that is allowed to be attached to the DS4000 controllers and to the DR550 internal networks. It receives all events (traps) from the managed devices but forwards only enterprise-significant traps (events) to the tier 1 management application (IBM Director). We recommend this two-tier approach because IBM Director cannot filter out events or eliminate duplicates and the amount of messages dispatched or displayed at the Director Console would be overwhelming for the operators.




Figure 10-1 SNMP architecture and communication

## Generic SNMP security

The SNMP protocol uses the community name for authorization. Most SNMP implementations use the default community name public for read-only community and private for a read-write community. In most cases, a community name is sent in a plain-text format between the SNMP agent and manager. Some SNMP implementations have additional security features, such as restriction of the accessible IP addresses.

Therefore, you should be careful about the SNMP security. At the very least:

- ▶ Do not use the default community names (public and private).
- ▶ Do not allow access to hosts that are running the SNMP agent from networks or IP hosts that do not necessarily require access.

You might want to physically secure the network to which you would send SNMP packets by using a firewall, because community strings are included as plain text in SNMP packets.

If the Net-SNMP Tier 2 SNMP manager is implemented on the RSM server for DR550 as documented later in this chapter, the solution will provide a sufficient level of network security because the RSM server is connected to the internal DR550 network. The internal network is physically isolated from any devices outside of the DR550 frame.

## Message Information Base (MIB)

The objects, which you can get or set by sending SNMP get or set requests, are defined as a set of databases called Message Information Base (MIB). The structure of MIB is defined as an Internet standard in RFC 1155; the MIB forms a tree structure.

Most hardware and software vendors provide you with extended MIB objects to support their own requirements. The SNMP standards allow this extension by using the private sub-tree, called *enterprise specific MIB*. Because each vendor has a unique MIB sub-tree under the private sub-tree, there is no conflict among vendor original MIB extensions.

## SNMP trap request

An SNMP agent can send SNMP trap requests to SNMP managers (such as Net-SNMP or IBM Director) to inform them of the change of values or status on the IP host where the agent is running. There are seven predefined types of SNMP trap requests, as shown in Table 10-1.

Table 10-1 SNMP trap request types

Trap type	Value	Description
coldStart	0	Restart after a crash.
warmStart	1	Planned restart.
linkDown	2	Communication link is down.
linkUp	3	Communication link is up.
authenticationFailure	4	Invalid SNMP community string was used.
egpNeighborLoss	5	EGP neighbor is down.
enterpriseSpecific	6	Vendor specific event happened.

A trap message contains pairs of an OID and a value shown in Table 10-1 to notify you of the cause of the trap message. You can also use type 6, the *enterpriseSpecific* trap type, when you have to send messages that are not appropriate for other predefined trap types, for example, DISK I/O error and application down. You can also set an integer value field called *Specific Trap* on your trap message.

## 10.2 Implementation scenario example

This section and the following ones present a sample implementation scenario that we used for illustration purposes while writing this book.

As previously stated, we used a two tier SNMP manager approach for our implementation scenario. For the tier 2 SNMP manager, we use the Net-SNMP service, which is installed on the RSM server for DR550. IBM Director Server is the tier 1 SNMP manager.

**Consideration:** The scenarios, scripts and procedures presented here are not part of the DR550 product. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or functioning of these examples. They are provided on an *as-is* basis.

All sample scenarios described in this chapter were executed in the following DR550 environment:

- Hardware
  - DR550 Model DR2 with dual DR550 SSAM Servers
  - RSM for DR550 server used as Tier 2 SNMP Manager (Net-SNMP server)
  - IBM Director server used as Tier 1 SNMP Manager

- ▶ Software levels
  - RSM for DR550 server
    - OS Linux SLES 10
    - RSM for DR550 1.7.12
    - net-snmp-5.3.0.1-25.2
    - net-snmp-devel-5.3.0.1-25.2
  - IBM Director server
    - Microsoft Windows Server 2003 Enterprise Edition Service Pack 2
    - IBM Director Version 5.20.2
- ▶ IP addresses information:
  - RSM for DR550 server:
    - eth0 192.168.4.100 255.255.255.0 internal switch 1
    - eth1 192.168.5.100 255.255.255.0 internal switch 2
    - eth2 9.11.218.246 255.255.255.0 customer network to external
    - eth3 100.100.51.10 255.255.255.0 customer network
  - IBM Director Server
    - Ethernet internal: 100.100.51.11 255.255.255.0 customer network
    - Ethernet test: 9.11.218.6 255.255.255.0 customer network to external
- ▶ Firewall settings:
  - RSM for DR550 server
    - Open port 161 inbound/outbound UDP
    - Open port 162 inbound/outbound UDP

**Note:** Ask your IBM Support how to enable these ports on the RSM server for DR550.

Figure 10-2 on page 395 shows the environment and the flow of SNMP messages from the DR550 components to the tier-2 SNMP manager (Net-SNMP, installed on the RSM server for DR550), and from there to the tier-1 SNMP manager (IBM Director server).

As you can see, SNMP traps from the DR550 SAN switches are sent to the Net-SNMP server over the DR550 internal Ethernet network. This is possible because RSM is connected to the internal network. (In other words, using the Net-SNMP on the RSM server for DR550 as tier-2 SNMP manager allows you to manage the internal SAN switches, and this is another reason why we recommend the 2-tier approach.)

The other DR550 components (FSG nodes and DR550 SSAM Servers) send SNMP traps to the Net-SNMP server on RSM over the external (customer) Ethernet. Note that the DR550 SSAM Servers also run the DS4000 Event monitor (dotted lines), and as such they can also forward the DS4000 SNMP events to the Net-SNMP server on RSM.

The Net-SNMP server receives all of the SNMP managed devices traps, filters them out (eliminating duplicate and less significant events), and then forwards only enterprise significant events (the orange arrows labelled as SNMP traps filtered) to the IBM Director Server, which in turn presents these significant events to the operator.



As depicted by the pink arrow (labelled Alerts e-mail filtered), the Net-SNMP server can also dispatch e-mail alerts to designated customer personnel.

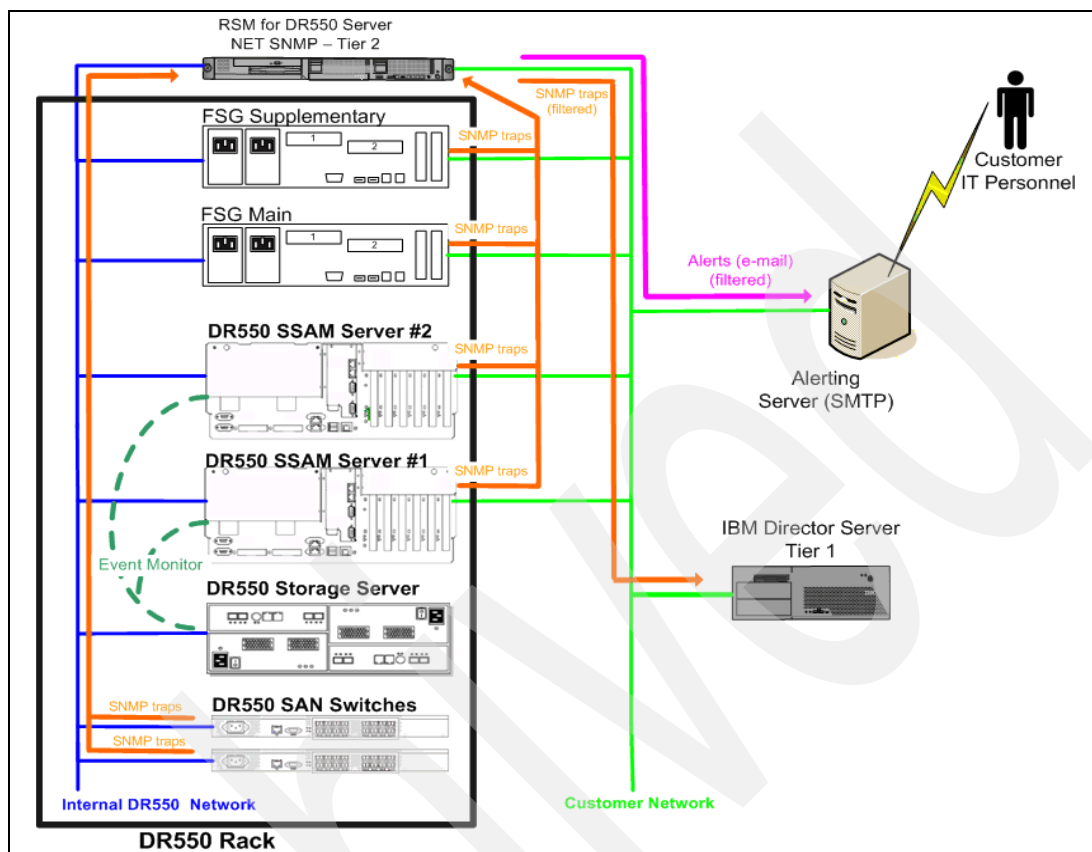


Figure 10-2 Two tier SNMP manager example environment on DR550

### 10.2.1 DR550 SSAM Server error notification and monitoring

The DR550 Storage Servers (p52A) can be monitored using different methods. The AIX operating system includes an SNMP agent that can be used for SNMP monitoring of the DR550 Storage Server. This standard SNMP agent is capable of returning requested information in response to SNMP get as well as sending an SNMP trap for some predefined conditions, such as agent start (Cold Start Trap).

### 10.2.2 SNMP setup for error notification on DR550 SSAM Server

The AIX operating system includes a number of SNMP agents that implement support for different versions of the SNMP protocol. Before the SNMP agent can be used, you must choose the version. By default, AIX comes with an SNMP agent configured for SNMP Version 3. Some other components, such as the SSAM application, currently relies on the SNMP DPI® (Distributed Protocol) interface supported only in SNMP Version 1.

**Note:** AIX SNMP agent Version 1 must be configured in order to have support for SSAM SNMP subagents.

## 10.3 Setting up trap handlers on RSM Server for DR550

Using a two-tier architecture for the SNMP management permits you to segregate functions of different management applications. The Tier 1 SNMP manager primary function is to facilitate the interaction with system administrators or operators by presenting enterprise significant events only through a graphical interface. Every event sent to the Tier 1 SNMP manager (IBM Director in our example) must be acted upon by an operator or administrator.

In contrast, the Tier 2 SNMP manager, installed on the RSM server for DR550, has no direct interaction with users. Instead, it is used as a primary event processing engine. SNMP traps from all SNMP-enabled devices and applications are forwarded to the Tier 2 SNMP manager, which processes these events. The Tier 2 SNMP engine should be equipped with all the necessary tools to perform the following functions:

- ▶ Event filtering and duplicate detection
- ▶ Enterprise significant event selection and forwarding to Tier 1 SNMP manager (IBM Director)
- ▶ Event attribute mapping
- ▶ Integration with non-SNMP-enabled alert notification applications, such as e-mail or paging

The event filtering and duplicate detection function automatically discards duplicate events within a specific time interval and filters out events that are less important. Discarded and filtered out events can still be logged by the Tier 2 SNMP manager (RSM server for DR550) but they are not forwarded to the Tier 1 SNMP manager (IBM Director).

Note that some SNMP enabled components are themselves capable of event filtering. An example of such a component is the SSAM server, which offers a mechanism for enabling and disabling specific event types. If event filtering is implemented at the SNMP subagent level, only a subset of all events is forwarded to the (Tier 2) SNMP manager, which results in reduced load on the Tier 2 SNMP manager.

**Tip:** Event filtering should be implemented as close to a source of the event as possible.

Enterprise significant event selection and forwarding to the Tier 1 SNMP manager function prioritizes incoming events and forwards only those events with the specified higher severity to the Tier 1 SNMP manager.

Event attribute mapping allows you to adequately format event attributes for presentation to the user through a GUI. Event attributes such as “Severity” from different subsystems should be mapped to the standard set of the Tier 1 SNMP manager (IBM Director), which includes the following event severities:

- ▶ FATAL
- ▶ CRITICAL
- ▶ MINOR
- ▶ WARNING
- ▶ HARMLESS
- ▶ UNKNOWN

Event severity is just an example of an attribute that must be mapped to correspond to the attribute set of the Tier 1 SNMP manager.

The Tier 2 SNMP manager should also be capable of forwarding event information to a non-SNMP application, such as e-mail or paging, in order to notify support personnel.

Software product selection for implementing Tier 2 SNMP manager could depend on a number of factors, such as previous experience with other products, cost, complexity, and so on. In our experiments, we used the NET-SNMP product, which comes pre-installed with the Linux operating system on the RSM server. It includes all the functional capabilities required for the implementation of a Tier 2 SNMP manager. Furthermore, if you implement an RSM server for DR550, this server could be the logical choice for where to implement the Tier 2 SNMP manager, because the operating system has all the required components pre-installed and the server has connections to both internal and external DR550 networks. In addition, having the Tier 2 SNMP manager connected to the internal DR550 network will enable monitoring of DR550 SAN switches, which only have network connections to the internal DR550 Ethernet switches.

### 10.3.1 Initial configuration

We used Net-SNMP as our Tier 2 SNMP manager. If you are using SLES 10, there is Net-SNMP already installed by default on the RSM server for DR550. If you use a different Linux, please verify Net-SNMP is installed before you start configuration.

The first step is to configure the Net-SNMP notification receiver “snmptrapd” daemon. By default, the notification receiver will not accept any incoming traps. Example 10-1 shows the configuration example that would allow the notification receiver daemon to receive traps from any device using the community names “public” and “private”

*Example 10-1 NET-SNMP server /etc/snmp/snmptrapd.conf*

---

```
authCommunity log public
authCommunity log private
~
"/etc/snmp/snmptrapd.conf"
```

---

The last step of configuring Net-SNMP is to configure the Net-SNMP SNMP agent. The configuration file for the Net-SNMP SNMP agent defines SNMP community attributes such as the range of IP addresses allowed to use a specific community name. In Example 10-2, only devices connected to 100.100.51.0/24 subnet will be allowed to have read-write access to MIB variables using the community “public”.

*Example 10-2 NET-SNMP server /etc/snmp/snmptrapd.conf*

---

```
syslocation Room 12321
syscontact snmpmaster (snmpmaster@company.com)
rwcommunity private 100.100.51.0/24
rocommunity public
"/etc/snmp/snmpd.conf"
```

---

Optionally, you can update the /etc/hosts file with the IBM Director IP address. This makes it easier to handle the IP environment. Example 10-3 shows you the /etc/hosts from the RSM Server for DR550.

*Example 10-3 NET-SNMP Server /etc/hosts*

---

```
10.100.51.10 RSMDR.de.ibm.com RSMDR
10.100.51.11 IBM Director
```

---

### 10.3.2 Setting up the trap handler for AIX

The first step is to configure the Net-SNMP server to accept AIX traps from the DR550 SSAM Server.

Example 10-4 and Example 10-5 illustrate an AIX Errorlog trap handler script that implements the following functions

- ▶ Event filtering and duplicate detection. Only the event filtering rules FATAL, MINOR, CRITICAL, and WARNING events can be forwarded to the IBM Director SNMP manager. An event is discarded as a duplicate if a previous event of the same type was logged less than 3600 seconds prior to the arrival of the current event. Two events are considered to be of the same type if the host name, severity, and resource name are identical.
- ▶ Enterprise significant event selection and forwarding to IBM Director SNMP manager. Only FATAL, MINOR, CRITICAL and WARNING events that are not duplicates are forwarded to IBM Director SNMP manager.
- ▶ Event attribute mapping. Severity attribute mapping rules:

AIX ERRORLOG CLASS "H" TYPE "PERM"	-> IBM Director "FATAL"
AIX ERRORLOG CLASS "S" TYPE "PERM"	-> IBM Director "CRITICAL"
AIX ERRORLOG CLASS "O" TYPE "INFO"	-> IBM Director "CRITICAL"
AIX ERRORLOG CLASS "H" TYPE "PEND"	-> IBM Director "CRITICAL"
AIX ERRORLOG CLASS "H" TYPE "PERF"	-> IBM Director "MINOR"
AIX ERRORLOG CLASS "H" TYPE "TEMP"	-> IBM Director "WARNING"
AIX ERRORLOG CLASS "H" TYPE "UNKN"	-> IBM Director "UNKNOWN"
AIX ERRORLOG CLASS "H" TYPE "INFO"	-> IBM Director "HARMLESS"

Integration with non-SNMP-enabled alert notification applications is implemented by sending information about enterprise significant events to system administrator e-mail (e-mail must be stored in the sysContact MIB variable on SSAM server).

*Example 10-4 Defining trap handler for AIX ERRORLOG SNMP traps in "/etc/snmp/snmptrapd.conf"*

---

```
traphandle RISC6000CLSMUXPD-MIB::ibm /root/snmp/AIXERRLOGtraphandler
```

---

*Example 10-5 Sample AIX ERRORLOG SNMP trap handler*

---

```
#!/bin/ksh

DUPLICATE_INTERVAL=3600
SNMP_TMP=/root/snmp
TIER1_SNMP_MANAGER=100.100.51.11

read HOST_NAME
read IP_ADDRESS

vars=

while read oid val
do

case "$oid" in
"IBMADSM-MIB::ibmProd.191.1.6.1")
 SEVERITY=`echo $val | awk '{ print $2}'`;
 MESSAGE_ID=`echo $val | awk '{ print $1}' | sed -e 's/\\/\\/';`
 RESOURCE_NAME=`echo $val | awk '{ print $6}'`;
 AIXERRORLOG_MESSAGE=$val;
 ;;
*)
```

---

```

;;
esac

if ["$vars" = ""]
then
 vars="$oid = $val"
else
 vars="$vars, $oid = $val"
fi
done

if [-f $SNMP_TMP/.deduplication/$HOST_NAME_SEVERITY_MESSAGE_ID_RESOURCE_NAME]
then
 _LAST=`cat
$SNMP_TMP/.deduplication/$HOST_NAME_SEVERITY_MESSAGE_ID_RESOURCE_NAME`
 _CURRENT=`date +%s`
 _ELAPSED=`expr $_CURRENT - $_LAST`
else
 _CURRENT=`date +%s`
 _ELAPSED=$_CURRENT
 date +%s >
$SNMP_TMP/.deduplication/$HOST_NAME_SEVERITY_MESSAGE_ID_RESOURCE_NAME
fi

if [$_ELAPSED -gt $DUPLICATE_INTERVAL]
then
 if [[$SEVERITY == 'WARNINIG' || $SEVERITY == 'MINOR' || $SEVERITY ==
'CRITICAL' || $SEVERITY == 'FATAL']]
 then
 echo _ELAPSED: $_ELAPSED >> $SNMP_TMP/traplog.log
 EMAIL_ADDRESS=`snmpget -v 1 -c public $HOST_NAME SNMPv2-MIB::sysContact.0 |
awk -F('{' '{ print $2 }' | sed -e 's/)//g' `
 NOTIFICATION="Severity: $SEVERITY; SUBSYSTEM_ID: AIX_ERRORLOG; MESSAGE:
$AIXERRORLOG_MESSAGE"
 echo $NOTIFICATION | mail -s "AIX ERRORLOG SNMP Trap information"
$EMAIL_ADDRESS
 # forwarding snmp trap to Tier 1 SNMP manager
 snmptrap -v 1 -c public $TIER1_SNMP_MANAGER
.iso.org.dod.internet.private.enterprises.ibm "" 6 1 ""
.iso.org.dod.internet.private.enterprises.ibm.191.1.6.1 s "$NOTIFICATION"
 date +%s >
$SNMP_TMP/.deduplication/$HOST_NAME_SEVERITY_MESSAGE_ID_RESOURCE_NAME
 else
 echo Non-critical event. Discarding.
 fi
else
 echo Duplicate detected
fi

```

---

The AIX ERRORLOG trap handler script is provided as an example only. The sample script demonstrates the approach of accessing trap information and processing this information. Once the information is processed by the Tier 2 SNMP manager, the forwarded enterprise significant events in Tier 1 SNMP manager GUI will look as shown in Figure 10-3.

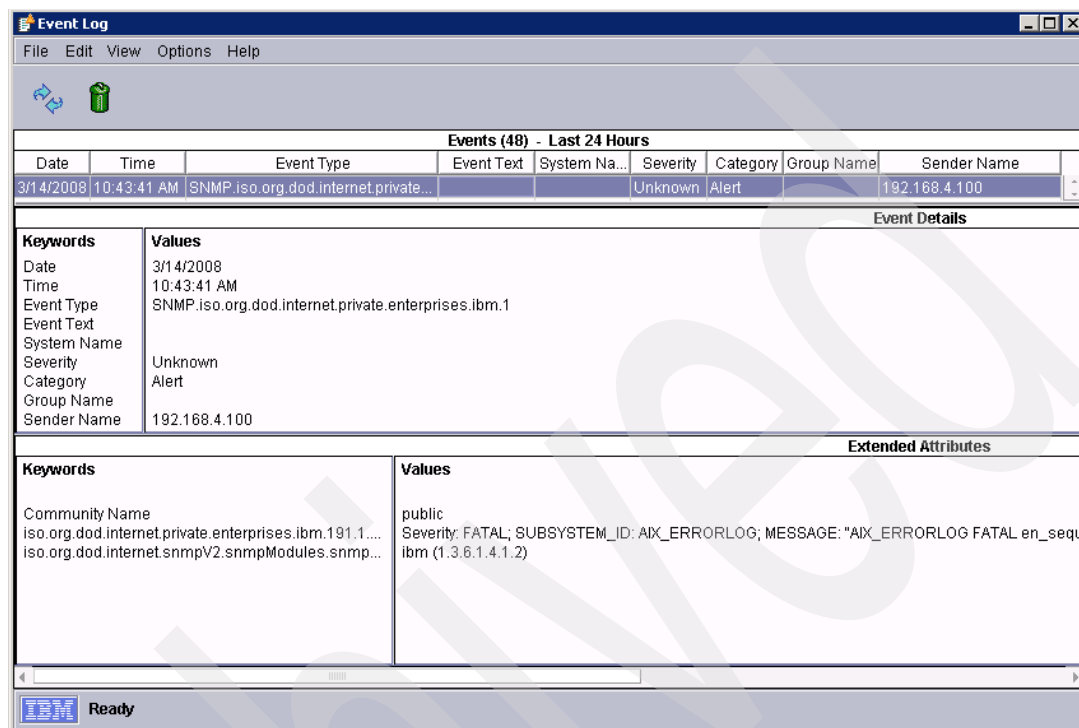


Figure 10-3 IBM Director SNMP Manager AIX ERRORLOG event

### 10.3.3 Setting up the trap handler for HACMP

The first step in configuring Net-SNMP server to accept SNMP traps from the HACMP Cluster Manager application is to copy the HACMP MIB definition file to the directory on the NET-SNMP server where all other MIB files are stored. The SSAM SNMP MIB file is called hacmp.my and is stored in the /usr/es/sbin/cluster/ directory. The destination directory file name on NET-SNMP server is /usr/share/snmp/mibs/RISC6000CLSMUXPD-MIB.txt. This is required to follow NET-SNMP naming standards for MIB files. By default NET-SNMP will not load any new MIB files unless it is specified explicitly what additional MIB needs to be loaded. The line shown in Example 10-6 must be added to /etc/snmp/snmp.conf file on the NET-SNMP server to load the SSAM SNMP MIB file.

*Example 10-6 NET-SNMP server /etc/snmp/snmp.conf*

```
mibs
+/usr/share/snmp/mibs/RISC6000CLSMUXPD-MIB.txt:/usr/share/snmp/mibs/ADMSERV-MIB.t
xt
~
"/etc/snmp/snmp.conf"
```

The “+” sign before the file name specifies that the new MIB file will be loaded in addition to all standard MIB files loaded by the server. If more than one additional MIB file has to be loaded, multiple file names can be specified in the “mibs” directive separated by a “:” sign.

To test the newly created configuration on the Net-SNMP server, run the SNMP GET query shown in Example 10-7 from the operating system prompt on the Net-SNMP server. In this example, 100.100.51.122 is the IP address of the AIX server where the HACMP Cluster Manager is running.

---

*Example 10-7 Verifying communication between Net-SNMP and HACMP SNMP subagent*

---

#on Net-SNMP server:

```
snmpwalk -v 1 -c public 100.100.51.122 RISC6000CLSMUXPD-MIB::cluster
RISC6000CLSMUXPD-MIB::clusterId.0 = INTEGER: 1077722580
RISC6000CLSMUXPD-MIB::clusterName.0 = STRING: "drs450"
RISC6000CLSMUXPD-MIB::clusterConfiguration.0 = ""
RISC6000CLSMUXPD-MIB::clusterState.0 = INTEGER: up(2)
RISC6000CLSMUXPD-MIB::clusterPrimary.0 = INTEGER: 3
RISC6000CLSMUXPD-MIB::clusterLastChange.0 = INTEGER: 1205773232
RISC6000CLSMUXPD-MIB::clusterGmtOffset.0 = INTEGER: 25200
RISC6000CLSMUXPD-MIB::clusterSubState.0 = INTEGER: stable(32)
RISC6000CLSMUXPD-MIB::clusterNodeName.0 = STRING: "drs_engine1"
RISC6000CLSMUXPD-MIB::clusterPrimaryNodeName.0 = STRING: "drs_engine1"
RISC6000CLSMUXPD-MIB::clusterNumNodes.0 = INTEGER: 2
RISC6000CLSMUXPD-MIB::clusterNodeId.0 = INTEGER: 3
RISC6000CLSMUXPD-MIB::clusterNumSites.0 = INTEGER: 0
```

---

If you did not receive a valid response from the HACMP SNMP subagent (cluster Id, cluster Name), you must resolve the issue before continuing with the Net-SNMP server configuration. See “Set up the SSAM application for monitoring through SNMP” on page 428 for more information about configuring the SNMP trap destination. HACMP and SSAM application SNMP monitoring processes are dependant upon the same SNMP agent configuration file: /etc/snmpd.conf on both HACMP cluster nodes.

The second step in the Net-SNMP server configuration process is to create trap handling scripts for all supported HACMP event types. In the current implementation of HACMP Cluster Manager, there are 64 trap types defined in the MIB file. In order for Net-SNMP to process HACMP traps, the trap handle scripts for all 64 types have to be configured. It might not be practical to write an individual trap handle script for every type of HACMP trap. In our example, we illustrate how it can be done with multiple trap handle scripts linked to a single script for simplicity, as shown in Example 10-8. The following command can be used on a Net-SNMP server to generate a list of HACMP SNMP traps defined in MIB file:

```
grep TRAP-TYPE /usr/share/snmp/mibs/RISC6000CLSMUXPD-MIB.txt|awk '{ print $1 }'
```

---

*Example 10-8 Net-SNMP server script generating trap handle definitions*

---

```
for i in `grep TRAP-TYPE /usr/share/snmp/mibs/RISC6000CLSMUXPD-MIB.txt | \
awk '{ print $1 }' ` ; do echo traphandle RISC6000CLSMUXPD-MIB::$i\
/root/snmp/$i_HACMPtraphandler >> /tmp/traphandle_snmp.tmp ; done
```

---

Newly generated trap handle script definitions are stored in /tmp/traphandle\_snmp.tmp. Each of the 64 lines in this file should have the following format:

```
traphandle RISC6000CLSMUXPD-MIB::trapClusterState \
/root/snmp/trapClusterState_HACMPtraphandler
```

If you choose to use a different location for your trap handle script, modify the script shown in Example 10-8 on page 401 accordingly. After reviewing the generated file, new definitions can be added to the Net-SNMP notification receiver daemon configuration file with the following command:

```
echo "" >> /etc/snmp/snmptrapd.conf ;\
cat /tmp/traphandle_snmp.tmp >> /etc/snmp/snmptrapd.conf
```

The next step is to create the main trap handle script to which all other scripts will be linked. The script listed in Example 10-9 for a HACMP trap handler implements the following functions:

- ▶ Event filtering and duplicate detection. Event filtering is not implemented in this example. An event is identified as a duplicate if a previous HACMP event of the same or different type from the same sender was logged less than 600 seconds prior to the arrival of the current event.
- ▶ Enterprise significant event selection and forewarning to IBM Director SNMP manager. All events are treated as enterprise-significant. All events that are not duplicates are forwarded to the IBM Director SNMP manager. (Effectively, the script only forwards the first HACMP trap information to IBM Director; all other traps within a 10 minute interval from the same server are logged and discarded as duplicates.)
- ▶ Event attribute mapping. Each HACMP trap has a unique attribute set. It might not be practical to implement the attribute mapping for the entire set of attributes defined in HACMP MIB. There is also no “Severity” attribute defined for SNMP traps. All traps are processed as though defined with Severity=“CRITICAL”.
- ▶ Integration with non SNMP-enabled alert notification applications is implemented by sending information about enterprise significant events to system administrator e-mail (e-mail must be stored in sysContact MIB variable on SSAM server).

*Example 10-9 Sample AIX ERRORLOG SNMP trap handler “/root/snmp/HACMPtraphandler”*

---

```
#!/bin/ksh

trapname=`basename $0 | awk -F '_' '{ print $1 }'`

DUPLICATE_INTERVAL=600
SNMP_TMP=/root/snmp
TIER1_SNMP_MANAGER=100.100.51.11

read HOST_NAME
read IP_ADDRESS

vars=

while read oid val
do
 if ["$vars" = ""]
 then
 vars="$oid = $val"
 else
 vars="$vars, $oid = $val"
 fi
done

if [-f $SNMP_TMP/.deduplication/$HOST_NAME_HACMP]
then
 _LAST=`cat $SNMP_TMP/.deduplication/$HOST_NAME_HACMP`
 _CURRENT=`date +%s`
```



```

 _ELAPSED=`expr $_CURRENT - $_LAST`
else
 _CURRENT=`date +%s`
 _ELAPSED=$_CURRENT
 date +%s > $SNMP_TMP/.deduplication/$HOST_NAME_HACMP
fi

if [$_ELAPSED -gt $DUPLICATE_INTERVAL]
then
 echo _ELAPSED: $_ELAPSED >> $SNMP_TMP/traplog.log
 EMAIL_ADDRESS=`snmpget -v 1 -c public $HOST_NAME\
SNMPv2-MIB::sysContact.0 | awk -F '(' '{ print $2 }' | sed -e 's/)//g'`
 # CLINFO=`snmpwalk -v 1 -c public $HOST_NAME\
RISC6000CLSMUXPD-MIB::cluster`

 NOTIFICATION="Severity: CRITICAL; $vars"

 echo $NOTIFICATION | mail -s "HACMP $trapname Trap information from
$HOST_NAME" $EMAIL_ADDRESS
 # forwarding snmp trap to Tier 1 SNMP manager
 snmptrap -v 1 -c public $TIER1_SNMP_MANAGER
.iso.org.dod.internet.private.enterprises.ibm "" 6 1 ""\
.iso.org.dod.internet.private.enterprises.ibm.191.1.6.1 s "$NOTIFICATION"
 date +%s > $SNMP_TMP/.deduplication/$HOST_NAME_HACMP
else
 echo Duplicate detected
 echo HACMPTRAPNAME: $trapname >> $SNMP_TMP/traplog.log
fi

```

---

After /root/snmp/HACMPtraphandler is created, you must verify that it has the appropriate execute permission set and must only be writable by root. Use the script in Example 10-10 to create links for individual trap handle scripts to this file.

*Example 10-10 Net-SNMP server script creating symbolic link definition file for HACMP trap handle scripts*

---

```

for i in `grep TRAP-TYPE /usr/share/snmp/mibs/RISC6000CLSMUXPD-MIB.txt | \
awk '{ print $1 }'` ; do echo ln -s /root/snmp/HACMPtraphandler \
/root/snmp/$i_HACMPtraphandler >> /tmp/traphandle_links.tmp ; \
done

```

---

Review and run the symbolic link definition file with the **ksh /tmp/traphandle\_links.tmp** command. If you choose to use a different location for your trap handle script, modify the script accordingly.

This example illustrates how the notification process works when several HACMP events take place. To initiate HACMP events, we use HACMP monitoring functionality to detect an outage and restart the application server (the SSAM application in the example shown in Example 10-11).

*Example 10-11 Verifying delivery of SNMP trap generated by HACMP to Net-SNMP*

---

```

#SSAM server:
dsmadm
IBM Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 5, Level 0.0
(c) Copyright by IBM Corporation and other(s) 1990, 2007. All Rights Reserved.

```

Enter your user id: admin

Enter your password:

Session established with server TSM: AIX-RS/6000

Server Version 5, Release 5, Level 0.0

Server date/time: 03/18/08 15:21:26 Last access: 03/18/08 15:18:50

tsm: TSM>halt

ANR2234W This command will halt the server; if the command is issued from a remote client, it may not

be possible to restart the server from the remote location.

Do you wish to proceed? (Yes (Y)/No (N))Y

#

#Net-SNMP server:

/usr/sbin/snmptrapd -f -Leo

NET-SNMP version 5.3.0.1

Mar 17 17:01:26 ibmrsm snmptrapd[28503]: 2008-03-17 17:01:26 localhost [127.0.0.1] (via

UDP: [100.100.51.122]:161) TRAP, SNM

unity public RISC6000CLSMUXPD-MIB::risc6000clsmuxpd Enterprise Specific Trap

(RISC6000CLSMUXPD-MIB::trapEventNewPrimary)

51:21.55 RISC6000CLSMUXPD-MIB::nodeName.0 = STRING: "drs\_engine1"

RISC6000CLSMUXPD-MIB::clusterName.0 = STRING

RISC6000CLSMUXPD-MIB::clusterPrimaryNodeName.0 = STRING: "drs\_engine1"

RISC6000CLSMUXPD-MIB::eventCount.0 = Counter32: 43

Mar 17 17:01:27 ibmrsm snmptrapd[28503]: 2008-03-17 17:01:27 localhost [127.0.0.1] (via

UDP: [100.100.51.122]:161) TRAP, SNM

unity public RISC6000CLSMUXPD-MIB::risc6000clsmuxpd Enterprise Specific Trap

(RISC6000CLSMUXPD-MIB::trapClusterStable) Up

:23.48 RISC6000CLSMUXPD-MIB::nodeName.0 = STRING: "drs\_engine1"

RISC6000CLSMUXPD-MIB::clusterName.0 = STRING: "drs45

00CLSMUXPD-MIB::netName.0 = STRING: "(unknown)" RISC6000CLSMUXPD-MIB::eventCount.0 =

Counter32: 44

the output is truncated ...

---

The output of the **snmptrapd** command is truncated due to the large number of lines. HACMP initiated 12 different events, each of which triggered an HACMP trap. Refer to the HACMP log files for more information about HACMP cluster events, in particular, see the /tmp/hacmp.out file on both cluster nodes.

The administrator whose e-mail address is stored in the sysContact MIB variable on the cluster member initiating the cluster events will receive an e-mail message similar to the one shown in Example 10-12.

---

*Example 10-12 Notification message sent by trap handle script in response to HACMP trap*

---

To: ssam\_admin@company.com

Subject: HACMP trapClusterSubState Trap information from engine

Message body:

Severity: CRITICAL; DISMAN-EVENT-MIB::sysUpTimeInstance = 0:21:43:11.81,

SNMPv2-MIB::snmpTrapOID.0 = RISC6000CLSMUXPD-MIB::trapClusterSubState,

RISC6000CLSMUXPD-MIB::clusterSubState.0 = unstable, RISC6000CLSMUXPD-MIB::clusterId.0 =

1077722580, RISC6000CLSMUXPD-MIB::clusterNodeId.0 = 3,

SNMP-COMMUNITY-MIB::snmpTrapAddress.0 = 127.0.0.1, SNMP-COMMUNITY-MIB::snmpTrapCommunity.0 = "public"

---

AIX HACMP trap handler scripts are provided as an example only. The sample scripts demonstrate the approach of accessing and processing trap information. Once the information is processed by the Tier 2 SNMP manager, the forwarded enterprise significant events in Tier 1 SNMP manager GUI will look as shown in Figure 10-4.

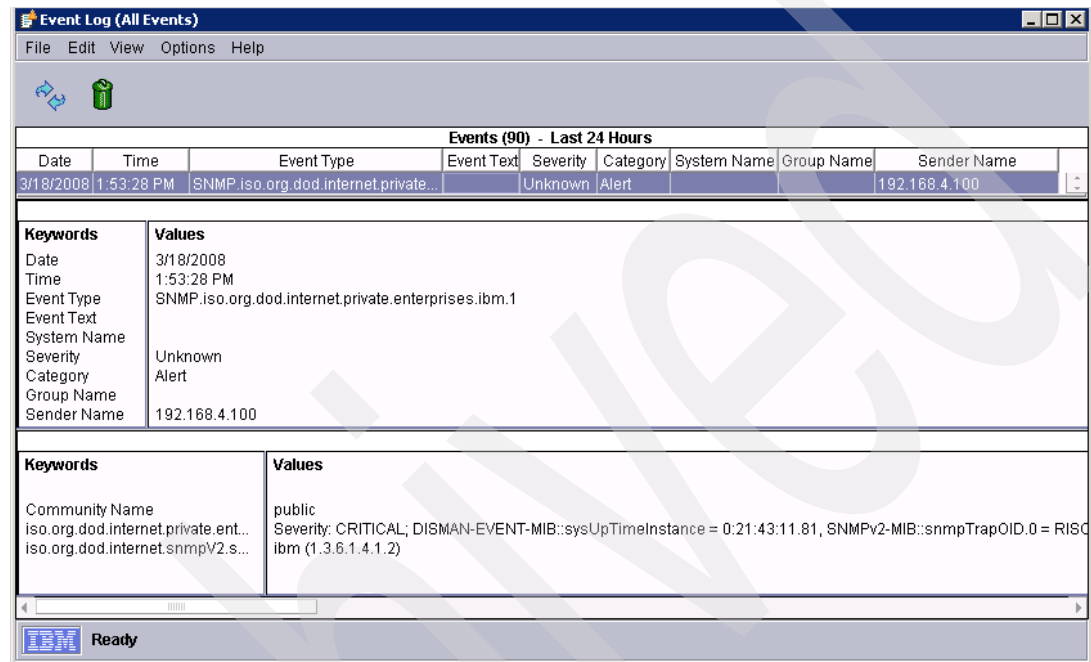


Figure 10-4 IBM Director SNMP Manager HACMP event

### 10.3.4 Setting up a trap handler for SSAM

The first step in configuring a Net-SNMP server to accept SNMP traps from the SSAM application is to copy the SSAM MIB definition file to the directory (on the NET-SNMP server) where all the other MIB files are stored.

The SSAM SNMP MIB file is called `adsmserve.mib` and is stored in the `/usr/tivoli/tsm/server/bin/` directory. The destination directory file name on NET-SNMP server is `/usr/share/snmp/mibs/ADSMSEV-MIB.txt`. This is required to follow the NET-SNMP MIB file name conventions. By default, NET-SNMP will not load any new MIB files unless it is specified explicitly what additional MIBs need to be loaded. The line shown in Example 10-13 must be added to `/etc/snmp/snmp.conf` file on the NET-SNMP server to load the SSAM SNMP MIB file. A "+" sign before the file name specifies that the new MIB file will be loaded in addition to all the standard MIB files loaded by the server.

*Example 10-13 NET-SNMP server `/etc/snmp/snmp.conf`*

```
mibs +/usr/share/snmp/mibs/ADSMSEV-MIB.txt
~
"/etc/snmp/snmp.conf"
```

---

To test the newly created configuration, run the SNMP GET query from the operating system prompt on the Net-SNMP server, as shown in Example 10-14. In this example 100.100.51.122 is the IP address of AIX server where SSAM application is running.

*Example 10-14 Verifying communication between Net-SNMP and SSAM SNMP subagent*

---

```
snmpget -v 1 -c public 100.100.51.122 IBMADSM-MIB::ibmAdsmServerHeartbeat.1
IBMADSM-MIB::ibmAdsmServerHeartbeat.1 = STRING: Mar 12 18:28:59 2008
```

---

If you did not receive a valid response from the SSAM SNMP subagent (GMT time stamp of the last SSAM application heartbeat), you must resolve the issue before continuing with the Net-SNMP server configuration.

The Net-SNMP notification receiver daemon can be started after completing the verification steps. To start the Net-SNMP notification receiver daemon in foreground mode, run the following command at the operating system prompt on Net-SNMP server:

```
/usr/sbin/snmptrapd -f -Leo
```

Specified startup options instruct the daemon to run in foreground mode to allow viewing incoming events in real time as well as logging trap information in the system log file. Once the process is started, it is ready to accept SNMP traps from the SSAM SNMP subagent. SSAM messages now will be forwarded to Tier 2 SNMP manager through SNMP traps according to the SSAM event rules discussed in “Configuring SSAM application for SNMP trap forwarding” on page 430.

*Example 10-15 Verifying delivery of a SNMP trap generated by SSAM to Net-SNMP*

---

SSAM server:

```
main:~ # dsmadm
IBM Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 3, Level 0.0
(c) Copyright by IBM Corporation and other(s) 1990, 2004. All Rights Reserved.

Enter your user id: admin

Enter your password:

Session established with server TSM: AIX-RS/6000
 Server Version 5, Release 5, Level 0.0
 Server date/time: 03/12/08 13:06:57 Last access: 03/12/08 12:53:34

tsm: TSM> migrate stgp ARCHIVEPOOL
ANR4921E MIGRATE STGPOOL: Primary storage pool ARCHIVEPOOL does not have a target pool.
ANS8001I Return code 11.
tsm: TSM>quit
```

ANS8002I Highest return code was 11.

Net-SNMP server:

```
/usr/sbin/snmptrapd -f -Leo
NET-SNMP version 5.3.0.1
2008-03-12 12:04:56 localhost [127.0.0.1] (via UDP: [100.100.51.122]:161) TRAP, SNMP v1,
community public
 IBMADSM-MIB::ibm.11.9 Enterprise Specific Trap (3) Uptime: 1 day, 0:36:08.67
 IBMADSM-MIB::ibmAdsmMIBMessages.3.0 = STRING: "ANR4921E MIGRATE STGPOOL: Primary
storage pool ARCHIVEPOOL does not have a target pool."
```

```
2008-03-12 12:04:56 localhost [127.0.0.1] (via UDP: [100.100.51.122]:161) TRAP, SNMP v1,
community public
 IBMADSM-MIB::ibm.11.9 Enterprise Specific Trap (4) Uptime: 1 day, 0:36:08.67
 IBMADSM-MIB::ibmAdsmMIBMessages.4.0 = STRING: "ANR2017I Administrator ADMIN issued
command: ROLLBACK ~"
```

---

The SSAM error ANR4921E is generated by the SSAM application in response to the **migrate stgp ARCHIVEPOOL** command, which is a normal behavior, and in this case, is only used for demonstration purposes.

The Net-SNMP notification receiver daemon successfully received two traps: IBMADSM-MIB::ibm.11.9 Enterprise Specific Trap (3) and IBMADSM-MIB::ibm.11.9 Enterprise Specific Trap (4). The trap type 3 and 4 in this case represent SSAM messages severity "ERROR" and "INFO" respectively. Each trap also has an IBMADSM-MIB::ibmAdsmMIBMessages string variable that is assigned the original SSAM message in both received traps.

Once the SNMP trap is received at the Tier 2 SNMP manager, it has to be processed. Processing of trap information is done by an external program that is invoked by the Net-SNMP notification receiver daemon. The program processing trap information in this book is referred to as a *trap handler*. To configure the trap handler for SSAM generated traps, the `/etc/snmp/snmptrapd.conf` file has to be modified to include a `traphandler` directive, as shown in Example 10-16.

*Example 10-16 Defining trap handler for SSAM generated traps in "/etc/snmp/snmptrapd.conf"*

---

```
traphandle IBMADSM-MIB::ibm.11.9 /root/snmp/SSAMtraphandler
```

---

The sample SSAM trap handler script shown in Example 10-17 implements the following functions:

- ▶ Event filtering and duplicate detection. Only the event filtering rules FATAL and CRITICAL events are forwarded to the Tier 1 SNMP manager. An event is discarded as a duplicate if a previous event of the same type was already logged less than 3600 seconds prior to the arrival of the current event.
- ▶ Enterprise significant event selection and forwarding to Tier 1 SNMP manager. Only FATAL and CRITICAL events that are not duplicates are forwarded to Tier 1 SNMP Manager (IBM Director).
- ▶ Event attribute mapping. The severity attribute mapping rules are:
  - SSAM "SEVERE" -> IBM Director "FATAL"
  - SSAM "ERROR" -> IBM Director "CRITICAL"
  - SSAM "WARNINIG" -> IBM Director "WARNINIG"
  - SSAM "INFO" -> IBM Director "HARMLESS"
- ▶ Integration with non SNMP-enabled alert notification applications is implemented by sending information about enterprise significant events to the system administrator e-mail (the e-mail address must be stored in the `sysContact` MIB variable on SSAM server).

*Example 10-17 Sample SSAM SNMP trap handler*

---

```
#!/bin/ksh

DUPLICATE_INTERVAL=3600
SNMP_TMP=/root/snmp
TIER1_SNMP_MANAGER=100.100.51.11

read HOST_NAME
read IP_ADDRESS
```

```

vars=

while read oid val
do

case "$oid" in
 "IBMADSM-MIB::ibmAdsmMIBMessages.1.0")
 SEVERITY=FATAL
 SSAM_MESSAGE_ID=`echo $val | awk '{ print $1}' | sed -e s/\"//g`;
 SSAM_MESSAGE=$val
 ;;
 "IBMADSM-MIB::ibmAdsmMIBMessages.2.0")
 SEVERITY=WARNINIG
 SSAM_MESSAGE_ID=`echo $val | awk '{ print $1}' | sed -e s/\"//g`;
 SSAM_MESSAGE=$val
 ;;
 "IBMADSM-MIB::ibmAdsmMIBMessages.3.0")
 SEVERITY=CRITICAL
 SSAM_MESSAGE_ID=`echo $val | awk '{ print $1}' | sed -e s/\"//g`;
 SSAM_MESSAGE=$val
 ;;
 "IBMADSM-MIB::ibmAdsmMIBMessages.4.0")
 SEVERITY=HARMLESS
 SSAM_MESSAGE_ID=`echo $val | awk '{ print $1}' | sed -e s/\"//g`;
 SSAM_MESSAGE=$val
 ;;
*)
 ;;
esac
if ["$vars" = ""]
then
 vars="$oid = $val"
else
 vars="$vars, $oid = $val"
fi
done
if [-f $SNMP_TMP/.deduplication/$SSAM_MESSAGE_ID]
then
 _LAST=`cat $SNMP_TMP/.deduplication/$SSAM_MESSAGE_ID`
 _CURRENT=`date +%s`
 _ELAPSED=`expr $_CURRENT - $_LAST`
else
 _CURRENT=`date +%s`
 _ELAPSED=$_CURRENT
 date +%s > $SNMP_TMP/.deduplication/$SSAM_MESSAGE_ID
fi

if [$_ELAPSED -gt $DUPLICATE_INTERVAL]
then
 if [[$SEVERITY == 'CRITICAL' || $SEVERITY == 'FATAL']]
 then
 echo _ELAPSED: $_ELAPSED >> $SNMP_TMP/traplog.log
 EMAIL_ADDRESS=`snmpget -v 1 -c public $HOST_NAME SNMPv2-MIB::sysContact.0 |
awk -F'(' '{ print $2 }' | sed -e 's/)//g' `
 NOTIFICATION="Severity: $SEVERITY; SUBSYSTEM_ID: SSAM; MESSAGE:
$SSAM_MESSAGE"
 echo $NOTIFICATION | mail -s "SSAM SNMP Trap information" $EMAIL_ADDRESS
 # forwarding snmp trap to Tier 1 SNMP manager

```

```

snmptrap -v 1 -c public $TIER1_SNMP_MANAGER
.iso.org.dod.internet.private.enterprises.ibm "" 6 1 ""
.iso.org.dod.internet.private.enterprises.ibm.191.1.6.1 s "$NOTIFICATION"
date +%s > $SNMP_TMP/.deduplication/$SSAM_MESSAGE_ID
else
 echo Non-critical event. Discarding.
fi
else
 echo Duplicate detected
fi

```

This SSAM trap handler script is provided as an example only. The sample script demonstrates the approach of accessing trap information and processing this information. Once the information is processed by the Tier 2 SNMP manager, the forwarded enterprise significant events in Tier 1 SNMP manager GUI will look as shown in Figure 10-5.

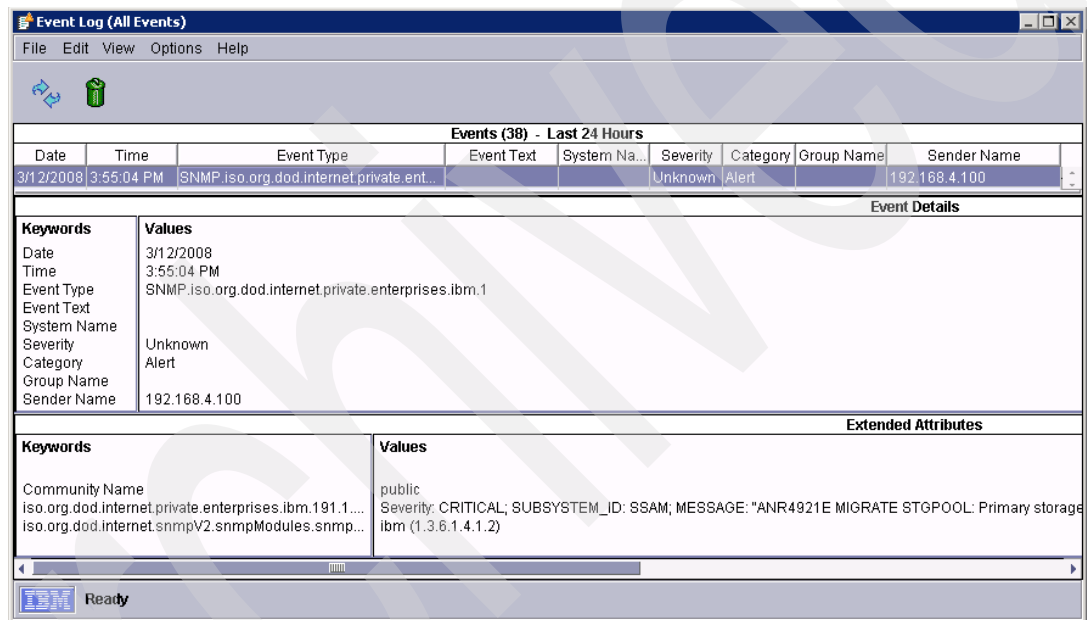


Figure 10-5 IBM Director SNMP Manager SSAM event

### 10.3.5 Setting up a trap handler for DR550 Storage Controller

The first step in configuring a Net-SNMP server to accept SNMP traps from Storage Servers is to copy the MIB definition files to the directory where all the other MIB files are stored on the Net-SNMP server. The destination file name on the Net-SNMP server is /usr/share/snmp/mibs/SM9R3A-MIB.txt. MIB definition file names must be changed to be consistent with the Net-SNMP naming standards for MIB files. The MIB file for DS4000 is included in the downloadable package for the SMclient. MIB can also be downloaded from the following URL:

<https://www-304.ibm.com/systems/support/storage/disk/ds4200/stormgr1.html>

**Important:** Do not install the SMclient for AIX from the URL listed above, as the version of the code is not compatible with DR550.

The line shown in Example 10-18 must be added to the /etc/snmp/snmp.conf file on the NET-SNMP server to load the SNMP MIB file.

*Example 10-18 NET-SNMP server /etc/snmp/snmp.conf*

---

```
mibs
+/usr/share/snmp/mibs/RISC6000CLSMUXPD-MIB.txt:/usr/share/snmp/mibs/ADSMSESV-MIB.t
xt:/usr/share/snmp/mibs/SM9R3A-MIB.txt
~
"/etc/snmp/snmp.conf"
```

---

The “+” sign before the file name specifies that the new MIB file will be loaded in addition to all standard MIB files loaded by the server. If more than one additional MIB file has to be loaded, multiple file names can be specified in the mibs directive separated by a “:” sign.

The second step in the Net-SNMP server configuration process is to create a trap handling script. In Example 10-19, we create a trap handling script for storageArrayCritical, which is defined in the SM9-R3MIB MIB.

*Example 10-19 Defining trap handler for Storage Server SNMP traps in “/etc/snmp/snmptrapd.conf”*

---

```
traphandle SM9-R3MIB::storageArrayCritical /root/snmp/DS4200traphandler
```

---

The script listed in Example 10-20 for a Storage Server trap handler implements the following functions:

- ▶ Event filtering and duplicate detection. No event filtering is implemented for SNMP traps generated by the Storage Server. All events are assigned “CRITICAL” severity
- ▶ An event is identified as a duplicate if a previous event with the same deviceUserLabel is received from the same IP address within less than 3600 seconds prior to the arrival of the current event.
- ▶ Enterprise significant event selection and forwarding to IBM Director SNMP manager. All events that are not duplicates are forwarded to IBM Director SNMP manager.
- ▶ Event attribute mapping. The severity attribute mapping rule is Storage Server SMclient “CRITICAL” -> IBM Director “CRITICAL”.
- ▶ Integration with non SNMP-enabled alert notification applications is implemented by sending information about enterprise significant events to the system administrator’s e-mail (the e-mail address must be stored in the sysContact MIB variable on the SSAM server running SMclient).

*Example 10-20 Sample Storage Server trap handler “/root/snmp/DS4200traphandler”*

---

```
#!/bin/ksh

DUPLICATE_INTERVAL=3600
SNMP_TMP=/root/snmp
TIER1_SNMP_MANAGER=100.100.51.11

read HOST_NAME
read IP_ADDRESS
vars=

while read oid val
do

case "$oid" in
"SM9-R3MIB::deviceUserLabel")
```



```

 deviceUserLabel=$(echo $val | sed -e 's/"//g' | awk '{ print $1 }')
 ;;
"SM9-R3MIB::deviceErrorCode")
 deviceErrorCode=$(echo $val | sed -e 's/"//g' | awk '{ print $1 }')
 ;;
"SM9-R3MIB::eventTime")
 eventTime=$(echo $val | sed -e 's/"//g' | awk '{ print $1 }')
 ;;
"SM9-R3MIB::trapDescription")
 trapDescription=$(echo $val | sed -e 's/"//g' | awk '{ print $1 }')
 ;;
"SM9-R3MIB::componentLocation")
 componentLocation=$(echo $val | sed -e 's/"//g' | awk '{ print $1 }')
 ;;
*)
 ;;
esac

if ["$vars" = ""]
then
 vars="$oid = $val"
else
 vars="$vars, $oid = $val"
fi
done

if [-f $SNMP_TMP/.deduplication/$HOST_NAME_deviceUserLabel]
then
 _LAST=`cat $SNMP_TMP/.deduplication/$HOST_NAME_deviceUserLabel`
 _CURRENT=`date +%s`
 _ELAPSED=`expr $_CURRENT - $_LAST`
else
 _CURRENT=`date +%s`
 _ELAPSED=$_CURRENT
 date +%s > $SNMP_TMP/.deduplication/$HOST_NAME_deviceUserLabel
fi

if [$_ELAPSED -gt $DUPLICATE_INTERVAL]
then
 echo _ELAPSED: $_ELAPSED >> $SNMP_TMP/traplog.log

 EMAIL_ADDRESS=`snmpget -v 1 -c public $HOST_NAME SNMPv2-MIB::sysContact.0 |
awk -F(' ' '{ print $2 }' | sed -e 's/)//g' `
 NOTIFICATION="Severity: CRITICAL;
deviceUserLabel: $deviceUserLabel;
deviceErrorCode: $deviceErrorCode;
eventTime: $eventTime;
trapDescription: $trapDescription;
componentLocation: $componentLocation;"

 echo $NOTIFICATION | mail -s "DS4200 SNMP Trap information" $EMAIL_ADDRESS
 echo $NOTIFICATION
 echo EMAIL_ADDRESS: $EMAIL_ADDRESS
 # forwarding snmp trap to Tier 1 SNMP manager
 snmptrap -v 1 -c public $TIER1_SNMP_MANAGER
.iso.org.dod.internet.private.enterprises.ibm "" 6 1 ""
.iso.org.dod.internet.private.enterprises.ibm.191.1.6.1 s "$NOTIFICATION"
 date +%s > $SNMP_TMP/.deduplication/$HOST_NAME_deviceUserLabel
else

```

echo Duplicate detected

fi

---

Example 10-21 shows the verification of the delivery of the SNMP trap.

*Example 10-21 Verifying delivery of SNMP trap generated by the SMclient to Net-SNMP server*

---

```
on drs_engine2:
SMclient
Select the SNMP tab from the Edit → Alert Destinations menu window at the DS4000 Storage
Manager client Enterprise Management window. Highlight trap destination IP address
Click on “Test” button
Click “OK”.

on RSMserver:
ibmrsm# /usr/sbin/snmptrapd -f -Leo
NET-SNMP version 5.3.0.1
2008-03-30 21:51:34 localhost [127.0.0.1] (via UDP: [100.100.51.123]:36575) TRAP, SNMP v1,
community public
 SM9-R3MIB::sm9-R3 Enterprise Specific Trap (SM9-R3MIB::storageArrayCritical)
Uptime: 0:00:00.00
 SM9-R3MIB::deviceHostIP = IPAddress: 0.0.0.0 SM9-R3MIB::deviceHostName = STRING:
" SM9-R3MIB::deviceUserLabel = STRING: "DRS_FastT1
" SM9-R3MIB::deviceErrorCode = STRING: " "
SM9-R3MIB::eventTime = STRING: "Sun, 30 Mar 2008 22:51:09 MDT "
SM9-R3MIB::trapDescription = STRING: "Alert Test Message
"SM9-R3MIB::componentType = STRING: "
" SM9-R3MIB::componentLocation = STRING: " "
Cannot find module (FC-MGMT-MIB): At line 14 in /usr/share/snmp/mibs/FCIP-MIB.txt
Did not find 'FcNameIdOrZero' in module #-1 (/usr/share/snmp/mibs/FCIP-MIB.txt)
Severity: CRITICAL; deviceUserLabel: DRS_FastT1; deviceErrorCode: ; eventTime: Sun;;
trapDescription: Alert; componentLocation: ; EMAIL_ADDRESS: storage_admin@company.com
```

---

The administrator whose e-mail address is stored in sysContact MIB variable on SSAM server running SMclient will receive an e-mail message similar to the one shown in Example 10-22.

*Example 10-22 Notification message in response to “storageArrayCritical” Storage Server trap*

---

```
To: storage_admin@company.com
Subject: DS4200 SNMP Trap information
Message body:
Severity: CRITICAL; deviceUserLabel: DRS_FastT1; deviceErrorCode: ; eventTime: Sun;;
trapDescription: Alert; componentLocation:
```

---

Storage Server trap handling scripts are provided as an example only. Once the information is processed by the Tier 2 SNMP manager, the forwarded enterprise significant events in Tier 1 SNMP manager GUI will look as shown in Figure 10-6 on page 413.

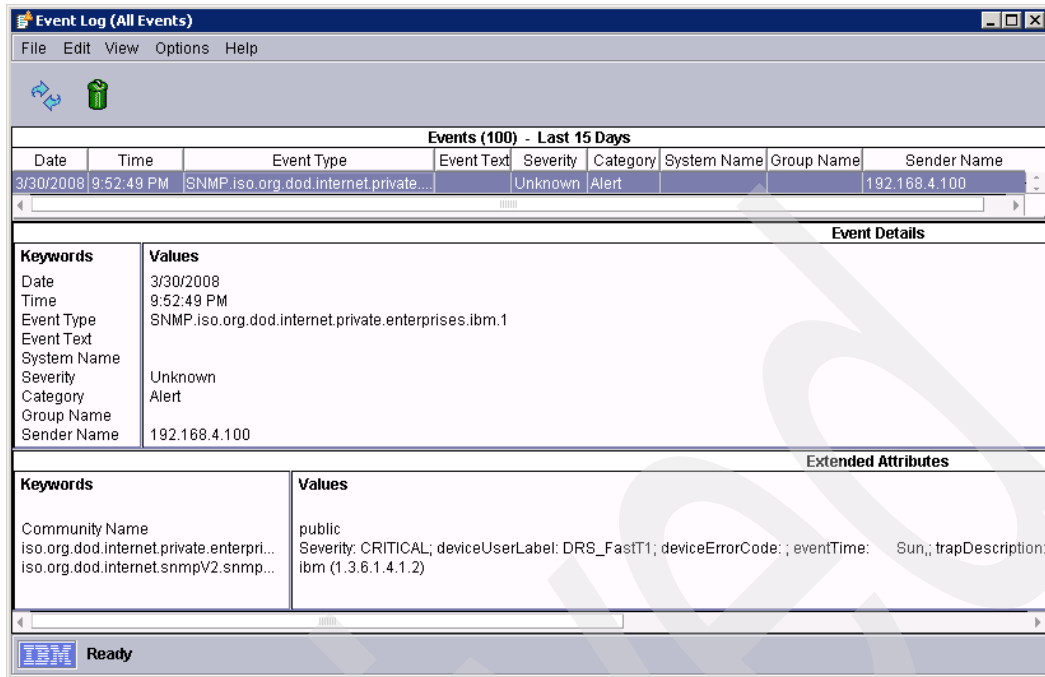


Figure 10-6 IBM Director SNMP Manager Storage Server event

### 10.3.6 Setting up a trap handler for SAN switches

The first step in configuring the Net-SNMP server for accepting SNMP traps from SAN switches is to copy MIB definition files to the directory where all other MIB files are stored on the Net-SNMP server. The destination directory file name on Net-SNMP server is /usr/share/snmp/mibs. MIB definition file names must be changed to be consistent with Net-SNMP naming standards for MIB files.

Table 10-2 shows the SAN switch MIB files.

Table 10-2 SAN switch MIB files

MIB	Brocade file name	Net-SNMP file name
FE-MIB	FE_RFC2837.mib	FE_RFC2837-MIB.txt
SW-MIB	SW_v5_6.mib	SW_v5_6-MIB.txt
FA-MIB	FA_v3_0.mib	FA_v3_0-MIB.txt
FICON®-MIB	FICON_v5_0.mib	FICON_v5_0-MIB.txt
HA-MIB	HA_v5_1.mib	HA_v5_1-MIB.txt
FCIP-MIB	fcip.mib	FCIP-MIB.txt

The line shown in Example 10-23 must be added to the `/etc/snmp/snmp.conf` file on the NET-SNMP server to load the Brocade SNMP MIB file. The “+” sign before the file name specifies that the new MIB file will be loaded in addition to all standard MIB files loaded by the server. If more than one additional MIB file has to be loaded, multiple file names can be specified in the `mibs` directive and separated by a “.” sign.

*Example 10-23 NET-SNMP server /etc/snmp/snmp.conf*

---

```
mibs
+/usr/share/snmp/mibs/RISC6000CLSMUXPD-MIB.txt:/usr/share/snmp/mibs/ADSMSESV-MIB.t
xt:/usr/share/snmp/mibs/FA_v3_0-MIB.txt:/usr/share/snmp/mibs/FE_RFC2837-MIB.txt:/u
sr/share/snmp/mibs/HA_v5_1-MIB.txt:/usr/share/snmp/mibs/SW_v5_6-MIB.txt:/usr/share
/snmplib/FCIP-MIB.txt
~
"/etc/snmp/snmp.conf"
```

---

To test the newly created configuration on the Net-SNMP server, run the SNMP GET query from the operating system prompt on the Net-SNMP server, as shown in Example 10-24. In this example, 192.168.4.31 is the IP address of the SAN switch.

*Example 10-24 Verifying communication between Net-SNMP and SAN Switch SNMP agent*

on Net-SNMP server:

---

```
ibmrsm# snmpget -v 1 -c public 192.168.4.31 sysContact.0
SNMPv2-MIB::sysContact.0 = STRING: SAN Administrator (san_admin@company.com)
```

---

If you did not receive a valid response from the SAN switch SNMP agent, you must resolve the issue before continuing with the Net-SNMP server configuration.

The second step of the Net-SNMP server configuration process is to create a trap handling script, as shown in Example 10-25. In this example, we only create trap handling script for `swEventTrap` defined in the SW-MIB MIB.

*Example 10-25 Defining trap handler for SAN Switch SNMP traps in "/etc/snmp/snmptrapd.conf"*

---

```
traphandle SW-MIB::swEventTrap /root/snmp/swEventTrap_SANSwtraphandler
```

---

The script listed in Example 10-26 on page 415 for a SAN switch trap handler implements the following functions:

- ▶ Event filtering and duplicate detection. Event filtering is implemented at the switch SNMP agent level as well as in the trap handling script. The SAN switch SNMP agent can be configured to only forward traps of a specified severity level. See Table 10-3 on page 438. Also, the SAN switch MIB capabilities can be customized to disable certain MIBs. See Table 10-4 on page 439. In addition, the trap handling script only accepts “critical”, “error”, and “informational” trap severities.
- ▶ An event is identified as a duplicate if a previous trap of the same severity is received from the same IP address within less than 3600 seconds prior to the arrival of the current event.
- ▶ Enterprise significant event selection and forwarding to IBM Director SNMP manager. All events that are not duplicates and not filtered out as non-critical events are forwarded to IBM Director SNMP manager.
- ▶ Event attribute mapping. The severity attribute mapping rules are:
 

SAN Switch “critical”	-> IBM Director “FATAL”
SAN Switch “error”	-> IBM Director “CRITICAL”
SAN Switch “informational”	-> IBM Director “HARMLESS”

- Integration with non SNMP-enabled alert notification applications is implemented by sending information about enterprise significant events to the system administrator e-mail (the e-mail address must be stored in sysContact MIB variable on SAN Switch).

*Example 10-26 Sample SAN switch trap handler "/root/snmp/swEventTrap\_SANSWtraphandler"*

---

```
#!/bin/ksh

trapname=`basename $0 | awk -F '_' '{ print $1 }'`

DUPLICATE_INTERVAL=3600
SNMP_TMP=/root/snmp
TIER1_SNMP_MANAGER=100.100.51.11

read HOST_NAME
read IP_ADDRESS
vars=

echo SANSW trap name: $trapname
while read oid val
do

 if ["$vars" = ""]
 then
 vars="$oid = $val"
 else
 vars="$vars, $oid = $val"
 fi
 OID_TRUNCATED=$(echo $oid | awk -F '.' '{ print $1}');
 echo OID_TRUNCATED=$OID_TRUNCATED VAL=$val

case "$OID_TRUNCATED" in
 "SW-MIB::swEventLevel")
 SEVERITY=$val
 ;;
 *)
 ;;
esac

done
echo SEVERITY: $SEVERITY

if [-f $SNMP_TMP/.deduplication/$HOST_NAME_trapname_SEVERITY]
then
 _LAST=`cat $SNMP_TMP/.deduplication/$HOST_NAME_trapname_SEVERITY`
 _CURRENT=`date +%s`
 _ELAPSED=`expr $_CURRENT - $_LAST`
else
 _CURRENT=`date +%s`
 _ELAPSED=$_CURRENT
 date +%s > $SNMP_TMP/.deduplication/$HOST_NAME_trapname_SEVERITY
fi

if [$_ELAPSED -gt $DUPLICATE_INTERVAL]
then
if [[$SEVERITY == 'critical' || $SEVERITY == 'error' || $SEVERITY == 'informational']]
then
 case "$SEVERITY" in
 "critical")
 ID_SEVERITY=FATAL
```

```

 ;;
 "error")
 ID_SEVERITY=CRITICAL
 ;;
 "informational")
 ID_SEVERITY=HARMLESS
 ;;
 *)
 ;;
 esac

 echo _ELAPSED: $_ELAPSED >> $SNMP_TMP/traplog.log
 EMAIL_ADDRESS=`snmpget -v 1 -c public $HOST_NAME SNMPv2-MIB::sysContact.0 |
 awk -F(' ' '{ print $2 }' | sed -e 's/)//g' `
 NOTIFICATION="Severity: $ID_SEVERITY; HOST_NAME: $HOST_NAME; $vars";

 echo $NOTIFICATION | mail -s "SAN Switch SNMP Trap information"
$EMAIL_ADDRESS
 # forwarding snmp trap to Tier 1 SNMP manager
 snmptrap -v 1 -c public $TIER1_SNMP_MANAGER
.iso.org.dod.internet.private.enterprises.ibm "" 6 1 ""
.iso.org.dod.internet.private.enterprises.ibm.191.1.6.1 s "$NOTIFICATION"
 date +%s > $SNMP_TMP/.deduplication/$HOST_NAME_$_trapname_$_SEVERITY
 else
 echo Non-critical event. Discarding.
 fi
else
 echo Duplicate detected
fi

```

---

Example 10-27 shows the verification of the delivery of the SNMP trap.

*Example 10-27 Verifying delivery of SNMP trap generated by SAN Switch agent to Net-SNMP server*

---

```

#on RSMserver (session #1):
ssh admin@192.168.4.32
admin@192.168.4.32's password:
Permission denied, please try again. ^C
#on RSMserver (session #2):
ibmrsm# /usr/sbin/snmptrapd -f -Leo
NET-SNMP version 5.3.0.1
2008-03-30 19:55:07 192.168.4.32(via UDP: [192.168.4.32]:1024) TRAP, SNMP v1, community
private
 SW-MIB::sw Enterprise Specific Trap (SW-MIB::swEventTrap) Uptime: 3:15:32.29
 SW-MIB::swEventIndex.138 = INTEGER: 138 SW-MIB::swEventTimeInfo.138 = STRING:
2008/03/31-02:44:48 SW-MIB::swEventLevel.138 = INTEGER: i
nformational(4) SW-MIB::swEventRepeatCount.138 = INTEGER: 1 SW-MIB::swEventDescr.138 =
STRING: SEC-1193 Security violation: Login failure attempt via TELNET/SSH/RSH. IP Addr:
192.168.4.100 SANSW trap name: swEventTrap

```

---

The administrator whose e-mail address is stored in the sysContact MIB variable on the SAN switch will receive an e-mail message similar to the one shown in Example 10-28 on page 417.

*Example 10-28 Notification message sent by trap handle script in response to "swEventTrap" SAN Switch trap*

---

To: san\_admin@company.com  
 Subject: SAN Switch SNMP Trap information  
 Message body:  
 Severity: HARMLESS; HOST\_NAME: 192.168.4.32; DISMAN-EVENT-MIB::sysUpTimeInstance = 9:23:55:33.44, SNMPv2-MIB::snmpTrapOID.0 = SW-MIB::swEventTrap, SW-MIB::swEventIndex.200 = 200, SW-MIB::swEventTimeInfo.200 = 2008/03/30-20:06:20, SW-MIB::swEventLevel.200 = informational, SW-MIB::swEventRepeatCount.200 = 1, SW-MIB::swEventDescr.200 = SEC-1193 Security violation: Login failure attempt via TELNET/SSH/RSH. IP Addr: 192.168.4.100, SW-MIB::swSsn.0 = 101498Z, SNMP-COMMUNITY-MIB::snmpTrapAddress.0 = 192.168.4.31, SNMP-COMMUNITY-MIB::snmpTrapCommunity.0 = "private"

---

The SAN Switch trap handling scripts are provided as an example only. The sample scripts demonstrate the approach of accessing and processing trap information. Once the information is processed by the Tier 2 SNMP manager, the forwarded enterprise significant events in Tier 1 SNMP manager GUI will look as shown in Figure 10-7.

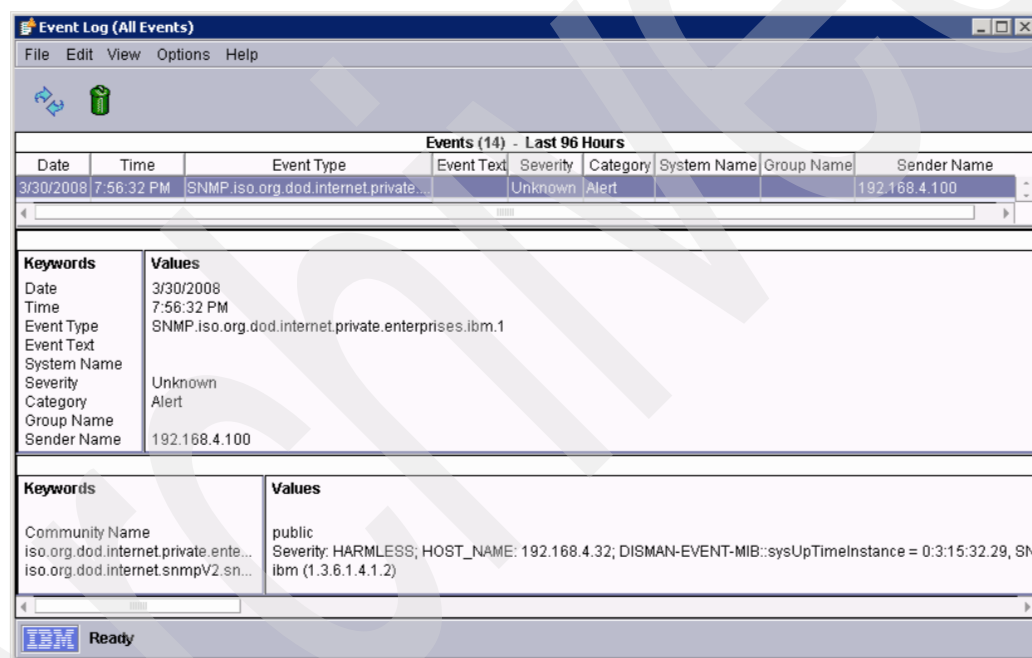


Figure 10-7 IBM Director SNMP Manager SAN Switch event

## 10.4 Configure SNMP monitoring for DR550 SSAM Server

There are three main applications installed on the DR550 SSAM Server. These are AIX, HACMP, and SSAM. Setting them up for monitoring is described in the following sections.

### 10.4.1 Configure SNMP monitoring for AIX

On AIX, the SNMP agent is installed as `/usr/sbin/snmpd`, which is a symbolic link to an actual executable command. By default, this link points to `snmpdv3ne`. If SSAM SNMP support is required, the AIX agent has to be reconfigured for SNMP Version 1 support, as shown in Example 10-29.

*Example 10-29 Switching AIX SNMP agent symbolic link to SNMP Version 1*

---

```
snmpv3_ssw -l
In /etc/rc.tcpip file, comment out the line that contains: snmpmibd
In /etc/rc.tcpip file, remove the comment from the line that contains: dpid2
Stop daemon: snmpd
Make the symbolic link from /usr/sbin/snmpd
to /usr/sbin/snmpdv1
Make the symbolic link from /usr/sbin/clsnmp
to /usr/sbin/clsnmpne
Start daemon: dpid2
Start daemon: snmpd
```

---

The symbolic link `/usr/sbin/snmpd` is now pointing to the SNMP Version 1 executable `/usr/sbin/snmpdv1`.

**Note:** You need to remove the comment from the line that contains `dpid2` in the `/etc/rc.tcpip` file manually in order to enable the automated startup of the `dpid2` process.

#### snmpd operation

The `snmpd` daemon is controlled by the system resource controller (SRC) on AIX. The command is invoked from `/etc/rc.tcpip` upon system initialization, as shown in Example 10-30.

*Example 10-30 AIX commands to analyze SNMP agent process status*

---

```
ps -ef | grep snmpd | grep -v grep
root 638984 503826 0 16:50:03 - 0:00 /usr/sbin/snmpd
lssrc -s snmpd
Subsystem Group PID Status
snmpd tcpip 638984 active
```

---

You can stop and start `snmpd` by issuing the following AIX commands:

```
stopsrc -s snmpd
startsrc -s snmpd
```

To control the operations of the AIX SNMP agent, you can customize the `/etc/snmpd.conf` configuration file. The `snmpd` daemon must be restarted in order to have the changes of the configuration file take effect. The `snmpd` daemon sends `warmStart` trap requests to the SNMP manager (if configured) upon refresh. The `snmpd` daemon sends `coldStart` trap requests to the SNMP manager upon restart.



### **Customization example for /etc/snmpd.conf**

The snmpd daemon reads the configuration file named /etc/snmpd.conf upon startup. The default /etc/snmpd.conf file provides you with many useful comments. By default, the snmpd daemon accepts requests from any IP-accessible hosts.

As explained in “Generic SNMP security” on page 392, we do not recommend this configuration. Therefore, you should restrict access in a way that only the specified servers are allowed access by setting up an IP address range on the community definition line.

For example, if you want to restrict the access to community “local” for the local host only (the IP address of the local loopback interface is always 127.0.0.1), you can define the community name “local” as follows:

```
community local 127.0.0.1 255.255.255.255
```

If you do not define any other communities, then you can send SNMP get requests from the local host only, and the snmpd agent will refuse any SNMP set requests.

If you want to restrict the access of a part of the MIB sub-tree, you can include the following two lines in /etc/snmpd.conf:

```
community public 0.0.0.0 0.0.0.0 readOnly 1.17.3
view 1.17.3 1.3.6.1.2.1.1.1
```

These lines instruct the agent to allow access to the object, specified by the view directive, from any IP hosts with the community named public.

To allow clients from a limited network address range to send SNMP get requests, you can include the following two lines in /etc/snmpd.conf:

```
community neighbor 192.168.1.0 255.255.255.0 readOnly 1.17.4
view 1.17.4 1.3.6.1.2.1
```

This configuration allows clients on the network address 192.168.1.0 to send SNMP get requests to the MIB-2 sub-tree.

### **Relationship among SNMP-related processes**

In AIX 5L Version 5.1 (and later), snmpd has the following two roles:

- ▶ SNMP agent
- ▶ SNMP Multiplex interface (SMUX) server

The SNMP agent is listening on UDP port 161. The SMUX server is listening on TCP port 199. A SMUX peer, named dpid2, is implemented as a SMUX-DPI2 converter. The Distributed Protocol Interface Version 2.0 (DPI2) interface is used to connect DPI2 subagents to the DPI2 server dpid2. The dpid2 daemon is an example of a DPI2 subagent. dsmsnmp is the DPI2 subagent that allows you to integrate SSAM SNMP functionality into AIX MIB. The dpid2 daemon supports integration of multiple subagents with AIX SNMP agent. Another example of an SNMP subagent is clsmuxpd, which is part of AIX HACMP subsystem (applies only to the DR550 DR2 model).

AIX includes several SNMP-related daemon processes, including dpid2. The dpid2 daemon processes are launched by the /etc/rc.tcpip startup script after the snmpd daemon process is started. The dpid2 daemon process are controlled by the System Resource Controller (SRC).

If you stop the AIX snmpd, then you should also stop the dpid2 process, because these processes rely on the presence of snmpd. If you do not stop these daemons, they might fill up the /var file system by expanding the size of log files. These are located in the /var/tmp directory.

For further information about SMUX, refer to *AIX 5L Version 5.3 Communications Programming Concepts*, SC23-4894. The sample source file is also available in the `/usr/samples/snmpd/smux` directory.

For further information about DPI2, refer to the files in the `/usr/samples/snmpd/dpi2` directory.

**Note:** The SNMP environment for AIX is subject to change in accordance with future AIX development.

## How to manage the MIB object definition file

The `snmpinfo` command reads the object definition file (`/etc/mib.defs`). The `mosy` command reads the Abstract Syntax Notation One (ASN.1) definitions of System Management Interface (SMI) and MIB modules and produces objects definition files in specific formats to be used by the `snmpinfo` command.

The following example creates an objects definition file to be used by the `snmpinfo` command:

```
mosy -o /etc/mib.defs /usr/samples/snmpd/smi.my /usr/samples/snmpd/mibII.my
```

Where:

- ▶ `/usr/samples/snmpd/smi.my`  
Defines the ASN.1 definitions by which the SMI is defined, as in RFC 1155.
- ▶ `/usr/samples/snmpd/mibII.my`  
Defines the ASN.1 definitions for the MIB II variables, as defined in RFC 1213.

Note this command invocation will replace the original `/etc/mib.defs` file.

## SNMP subagents

The AIX `snmpd` can send SNMP trap requests, but does not initiate the requests. In order to send traps or add your own MIBs, you need to implement your own SNMP subagents.

The subagent can be any software that is able to run on AIX 5L Version 5.2 (or later) and that can initiate SNMP traps. Follow the specific installation guide that is shipped with the subagent to install and configure the AIX SNMP environment.

## Setting up the SSAM server for monitoring the AIX error log

The AIX error logging facility records hardware and software failures in the error log for information purposes or for fault detection, corrective action, and notification.

Every time an error condition is logged, it gets assigned a number of attributes according to the structure of a corresponding object in the AIX Object Database Manager (ODM). The ODM object class name used by the AIX error logging facility is `errnotify`. To see the objects currently configured in the `errnotify` object class, use the AIX command `odmget errnotify`.

The Error Notification object class specifies the conditions and actions to be taken when errors are recorded in the system error log. Each time an error is logged, the error notification daemon determines if the error log entry matches the selection criteria of any of the Error Notification objects. If a match exists, the daemon runs the programmed action, also called a notify method, for each matched object. The Error Notification object class is located in the `/etc/objrepos/errnotify` file. Error Notification objects are added to the object class by using Object Data Manager (ODM) commands. Only processes running with the root user authority can add objects to the Error Notification object class. In this example, we describe how to set up automatic notification for “Hardware” event classes and some “Software” event classes. Objects in the Error Notification object class also include the “event type” attribute. This

attribute will be used for attribute mapping between the AIX Error Log facility and the IBM Director SNMP manager.

The procedure in Example 10-31 describes how to set up automatic notification methods to forward information logged in the AIX Error Log facility to the Net-SNMP manager. By modifying the following procedure, you can track other errors that are significant to you.

With root authority, make a backup copy of the /etc/objrepos/errnotify ODM file. You can name the backup copy anything you choose. In the following example, the backup copy appends the errnotify file name with the current date:

```
cd /etc/objrepos; cp errnotify errnotifycurrent_date
```

Use your favorite editor to create a file named /tmp/snmpnotify.add that contains the stanzas shown in Example 10-31.

---

*Example 10-31 ODM object definitions for enabling SNMP error notification*

---

```
errnotify:
 en_pid = 0
 en_name = ""
 en_persistenceflg = 1
 en_label = ""
 en_crcid = 0
 en_class = "H"
 en_type = "PEND"
 en_alertflg = ""
 en_resource = ""
 en_rtype = ""
 en_rclass = ""
 en_symptom = ""
 en_err64 = ""
 en_dup = ""
 en_method = "/usr/sbin/snmpttrap -c public -h snmp_manager -m AIX_ERRORLOG CRITICAL
en_sequence_number `echo $1 | /usr/bin/sed -e \"s/ /\n\"` en_resource_name `echo $6 |
/usr/bin/sed -e \"s/ /\n\"`

errnotify:
 en_pid = 0
 en_name = ""
 en_persistenceflg = 1
 en_label = ""
 en_crcid = 0
 en_class = "H"
 en_type = "PERF"
 en_alertflg = ""
 en_resource = ""
 en_rtype = ""
 en_rclass = ""
 en_symptom = ""
 en_err64 = ""
 en_dup = ""
 en_method = "/usr/sbin/snmpttrap -c public -h snmp_manager -m AIX_ERRORLOG MINOR
en_sequence_number `echo $1 | /usr/bin/sed -e \"s/ /\n\"` en_resource_name `echo $6 |
/usr/bin/sed -e \"s/ /\n\"`

errnotify:
 en_pid = 0
 en_name = ""
 en_persistenceflg = 1
 en_label = ""
```

```

en_crcid = 0
en_class = "H"
en_type = "PERM"
en_alertflg = ""
en_resource = ""
en_rtype = ""
en_rclass = ""
en_symptom = ""
en_err64 = ""
en_dup = ""
en_method = "/usr/sbin/snmptrap -c public -h snmp_manager -m AIX_ERRORLOG FATAL
en_sequence_number `echo $1 | /usr/bin/sed -e \"s/ /\\"` en_resource_name `echo $6 |
/usr/bin/sed -e \"s/ /\\"`

errnotify:
en_pid = 0
en_name = ""
en_persistenceflg = 1
en_label = ""
en_crcid = 0
en_class = "H"
en_type = "TEMP"
en_alertflg = ""
en_resource = ""
en_rtype = ""
en_rclass = ""
en_symptom = ""
en_err64 = ""
en_dup = ""
en_method = "/usr/sbin/snmptrap -c public -h snmp_manager -m AIX_ERRORLOG WARNING
en_sequence_number `echo $1 | /usr/bin/sed -e \"s/ /\\"` en_resource_name `echo $6 |
/usr/bin/sed -e \"s/ /\\"`

errnotify:
en_pid = 0
en_name = ""
en_persistenceflg = 1
en_label = ""
en_crcid = 0
en_class = "H"
en_type = "UNKN"
en_alertflg = ""
en_resource = ""
en_rtype = ""
en_rclass = ""
en_symptom = ""
en_err64 = ""
en_dup = ""
en_method = "/usr/sbin/snmptrap -c public -h snmp_manager -m AIX_ERRORLOG UNKNOWN
en_sequence_number `echo $1 | /usr/bin/sed -e \"s/ /\\"`
s/ /\\"` en_resource_name `echo $6 | /usr/bin/sed -e \"s/ /\\"`

errnotify:
en_pid = 0
en_name = ""
en_persistenceflg = 1
en_label = ""
en_crcid = 0
en_class = "H"
en_type = "INFO"

```

```

en_alertflg = ""
en_resource = ""
en_rtype = ""
en_rclass = ""
en_symptom = ""
en_err64 = ""
en_dup = ""
en_method = "/usr/sbin/snmptrap -c public -h snmp_manager -m AIX_ERRORLOG HARMLESS
en_sequence_number `echo $1 | /usr/bin/sed -e \"s/ /\` en_resource_name `echo $6 |
/usr/bin/sed -e \"s/ /\`"

```

errnotify:

```

en_pid = 0
en_name = ""
en_persistenceflg = 1
en_label = "CORE_DUMP"
en_crcid = 0
en_class = "S"
en_type = "PERM"
en_alertflg = ""
en_resource = ""
en_rtype = ""
en_rclass = ""
en_symptom = ""
en_err64 = ""
en_dup = ""
en_method = "/usr/sbin/snmptrap -c public -h snmp_manager -m AIX_ERRORLOG CRITICAL
en_sequence_number `echo $1 | /usr/bin/sed -e \"s/ /\` en_resource_name `echo $6 |
/usr/bin/sed -e \"s/ /\`"

```

errnotify:

```

en_pid = 0
en_name = ""
en_persistenceflg = 1
en_label = "JFS_FS_FULL"
en_crcid = 0
en_class = "O"
en_type = "INFO"
en_alertflg = ""
en_resource = ""
en_rtype = ""
en_rclass = ""
en_symptom = ""
en_err64 = ""
en_dup = ""
en_method = "/usr/sbin/snmptrap -c public -h snmp_manager -m AIX_ERRORLOG CRITICAL
en_sequence_number `echo $1 | /usr/bin/sed -e \"s/ /\` en_resource_name `echo $6 |
/usr/bin/sed -e \"s/ /\`"

```

errnotify:

```

en_pid = 0
en_name = ""
en_persistenceflg = 1
en_label = "J2_FS_FULL"
en_crcid = 0
en_class = "O"
en_type = "INFO"
en_alertflg = ""
en_resource = ""
en_rtype = ""

```

```

en_rclass = ""
en_symptom = ""
en_err64 = ""
en_dup = ""
en_method = "/usr/sbin/snmptrap -c public -h snmp_manager -m AIX_ERRORLOG CRITICAL
en_sequence_number `echo $1 | /usr/bin/sed -e \"s/ /\\"` en_resource_name `echo $6 |
/usr/bin/sed -e \"s/ /\\"`"

```

---

Run the command **odmadd /tmp/snmpnotify.add** after the file is saved. Note that the last line of every stanza starting with “en\_method” must be a continuous line until the end of the stanza with no <CR> before the end of the line.

**Note:** Do not run **odmadd /tmp/snmpnotify.add** more than once to avoid entering the same object definitions multiple times into the AIX ODM.

To verify that the new object are added successfully, you can use the command **odmget errnotify**.

In the definition of the event notification method in all the stanzas, we used the community name “public” and SNMP trap destination server “snmp\_manager”. We highly recommend to not use the default community name. Also, snmp\_manager is the host name that must be resolvable into the Net-SNMP server IP address. You can use any other host name in accordance with your host name naming standards. If you choose to use another name, you have to update stanza definitions for all the Error Notification objects prior to running the **odmadd /tmp/snmpnotify.add** command. If you decide to change the stanza definitions after new objects were added, the backup copy of the /etc/objrepos/errnotify ODM file can be copied back to restore the Error Notification definition to its original state.

**Note:** If you put the snmp\_manager host name into the /etc/hosts file on one HACMP cluster member, you must ensure that the same update is done on the second cluster node and that the /etc/hosts files are identical on both nodes.

The AIX command **/usr/sbin/snmptrap** is used for generating a notification SNMP trap to report an event to the Net-SNMP manager. The generated trap is using the .iso.org.dod.internet.private.enterprises.ibm enterprise specific trap, which is defined in the standard Net-SNMP MIB. Therefore, we do not need to add any new MIB files to be able to process the trap information.

Example 10-32 illustrates how the new notification process works when one of the AIX processes is abnormally terminated (core dump condition).

*Example 10-32 Verifying delivery of SNMP trap generated by AIX Error Log to Net-SNMP*

---

```

SSAM_server:
ksh &
kill -BUS 327770
[1] + Stopped (SIGTTIN) ksh &
ps
 PID TTY TIME CMD
 327770 pts/0 0:00 ksh
 327748 pts/0 0:00 -ksh
 667724 pts/0 0:00 -ksh
 676080 pts/0 0:00 ps
kill 327770
#
[1] + Bus error(core dump) ksh &

```

Net-SNMP server:

```
/usr/sbin/snmptrapd -f -Leo
NET-SNMP version 5.3.0.1
2008-03-14 09:57:54 ibmrsm.site [100.100.51.10] (via UDP: [100.100.51.122]:62432) TRAP,
SNMP v1, community public
RISC6000CLSMUXPD-MIB::ibm Enterprise Specific Trap (1) Uptime: 7 days, 0:38:19.00
IBMADSM-MIB::ibmProd.191.1.6.1 = STRING: "AIX_ERRORLOG CRITICAL en_sequence_number
39324 en_resource_name SYSPROC "
```

SSAM server:

```
errpt -l 39324 -a
```

```

LABEL: CORE_DUMP
IDENTIFIER: 40E9A4E1
```

```
Date/Time: Fri Mar 14 10:56:57 MDT 2008
Sequence Number: 39324
Machine Id: 00063634D700
Node Id: drs_engine1
Class: S
Type: PERM
Resource Name: SYSPROC
```

```
Description
SOFTWARE PROGRAM ABNORMALLY TERMINATED
```

```
Probable Causes
SOFTWARE PROGRAM
```

```
User Causes
USER GENERATED SIGNAL
```

```
Recommended Actions
CORRECT THEN RETRY
```

```
Failure Causes
SOFTWARE PROGRAM
```

```
Recommended Actions
RERUN THE APPLICATION PROGRAM
IF PROBLEM PERSISTS THEN DO THE FOLLOWING
CONTACT APPROPRIATE SERVICE REPRESENTATIVE
```

```
Detail Data
SIGNAL NUMBER
10
USER'S PROCESS ID:
327770
FILE SYSTEM SERIAL NUMBER
1
INODE NUMBER
2
PROCESSOR ID
-1
CORE FILE NAME
/core
PROGRAM NAME
ksh
```

```
STACK EXECUTION DISABLED
0
COME FROM ADDRESS REGISTER
??
ADDITIONAL INFORMATION
sh_done_1 138
??
??
??
Unable to generate symptom string.
```

---

The last command in this example, **errpt -1 39324 -a**, provides more details about the AIX event. The sequence number “39324” is reported in the SNMP trap as “en\_sequence\_number 39324”.

## 10.4.2 Configure SNMP monitoring for HACMP

By design, HACMP provides recovery for various failures that occur within the HACMP cluster. For example, HACMP can compensate for a network interface failure by swapping in a standby interface. As a result, it is possible that a component in the cluster has failed and that you are unaware of the fact. The danger here is that, while HACMP can survive one or possibly several component failures, each failure that escapes your notice threatens the cluster’s ability to keep providing a highly available environment, as the redundancy of cluster components is diminished. To avoid this situation, you should customize your system by adding event notification to ensure that no cluster events are left unnoticed. The Net-SNMP server (RSM for DR550) can be configured to receive and process HACMP generated SNMP traps.

### HACMP SNMP subagent startup

Starting with HACMP Version 5.3, the SNMP subagent is no longer using a separate clsmuxpd daemon to implement the SNMP subagent. Instead, this functionality is included in the Cluster Manager clstrmgrES daemon. This daemon is started automatically during system startup. The following SNMP-related configuration files are updated on the system during the HACMP software installation process. The following entry is added to the end of the /etc/snmpd.conf file to include the HACMP MIB controlled by the Cluster Manager:

```
smux 1.3.6.1.4.1.2.3.1.2.1.5 clsmuxpd_password # HACMP/ES for AIX clsmuxpd
```

The /etc/snmpd.peers file configures snmpd SMUX peers. During installation, HACMP adds the following entry to include the clsmuxpd password in this file:

```
clsmuxpd 1.3.6.1.4.1.2.3.1.2.1.5 "clsmuxpd_password" # HACMP/ES for AIX
```

**Note:** After a typical HACMP installation, all SNMP agent components are configured. There should be no need to customize these settings.

## 10.4.3 Configure SNMP monitoring for SSAM

For IBM System Storage Archive Manager, you can use SNMP traps and IBM Tivoli Storage Manager Operational Reporting to monitor the SSAM application and produce reports.



## IBM System Storage Archive Manager and SNMP

Figure 10-8 on page 428 shows the different elements and types of communication among the SNMP components.

The SSAM application provides processes and an SNMP subagent to send event information to the AIX SNMP agent (the AIX SNMP agent is a prerequisite).

- ▶ The SNMP manager and AIX SNMP agent communicate with each other through the SNMP protocol.
- ▶ The AIX SNMP agent communicates with the SSAM subagent. The AIX SNMP agent sends queries to the subagent and receives traps that inform the SNMP manager about events taking place on the Tivoli application monitored through the subagent. The SNMP agent and subagent communicate through the Distributed Protocol Interface (DPI). Communication takes place over a stream connection, which typically is a TCP connection, but could be another type of stream-connected transport mechanism.

The SSAM application supports the SNMP protocol to implement the following:

- ▶ Set up an SNMP heartbeat monitor to regularly check that the SSAM application is running.
- ▶ Send messages known as traps to an SNMP manager, such as NET-SNMP, IBM Director, NetView®, Tivoli Enterprise Console, or others.
- ▶ Run SSAM scripts and retrieve output and return codes.

The SSAM application will send trap messages for configurable events. All Tivoli Storage Manager messages are explained in *IBM Tivoli Storage Manager Messages*, SC32-0140. The classes of messages that can be sent to SNMP traps or disabled are the following:

- ▶ INFO Information messages (type of I)
- ▶ WARNING Warning messages (type of W)
- ▶ ERROR Error messages (type of E)
- ▶ SEVERE Severe error messages (type of S) enabled

In addition, particular messages can be enabled or disabled in the SSAM application. Disabling a specific message type prevents SSAM application from forwarding this type of event to the SSAM SNMP subagent and consequently sending those messages to the Net-SNMP manager.

The management information base (MIB) file, which is shipped with SSAM code, defines the variables that run server scripts and return the server script results. You must register SNMPADMIN, which is the SSAM administrator, using the administrative client. The SSAM SNMP subagent uses this account to run SSAM scripts in response to an authorized SNMP get request.

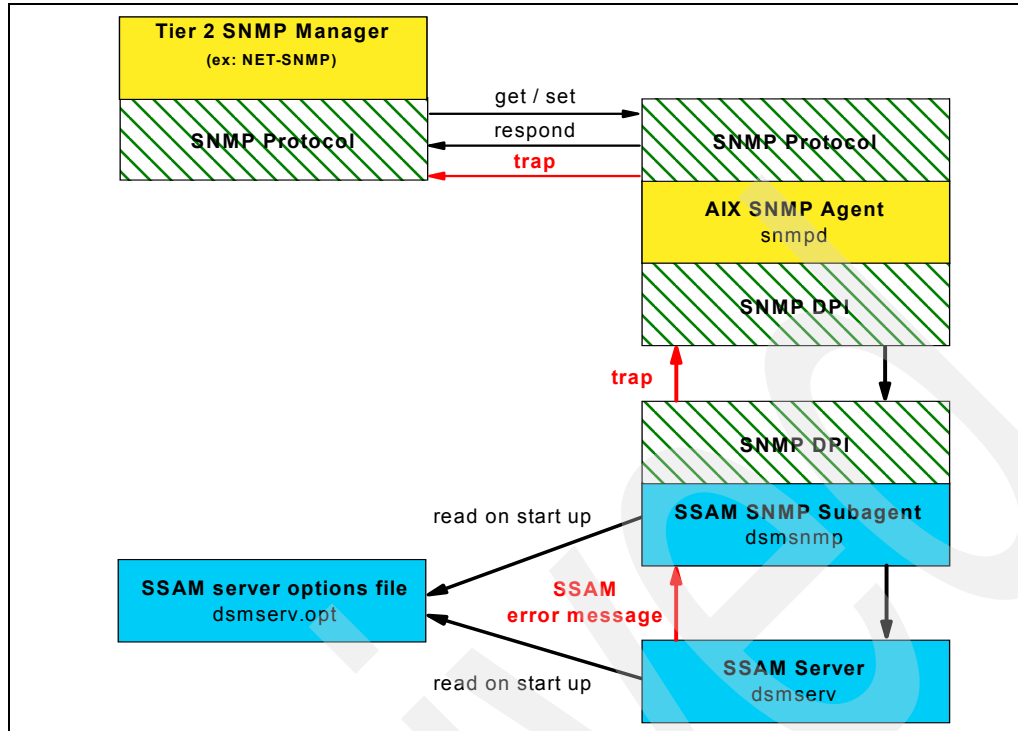


Figure 10-8 IBM Tivoli Storage Manager SNMP communication

## Set up the SSAM application for monitoring through SNMP

To set up the SSAM application monitoring through SNMP, do the following:

1. Configure the AIX SNMP agent (snmpd), as described in the documentation for that agent.
  - a. Stop the AIX SNMP agent (snmpd) with the AIX command `stopsrc -s snmpd`. Check that the status is inoperative with the AIX command `lssrc -s snmpd`.
  - b. The AIX SNMP agent is configured by customizing the file `/etc/snmpd.conf`. A configuration might look like Example 10-33.

### Example 10-33 `/etc/snmpd.conf`

```
logging file=/usr/tmp/snmpd.log enabled
logging size=0 level=0
community public
community private 127.0.0.1 255.255.255.255 readWrite
community private 100.100.51.10 255.255.255.255 readWrite # Tier 2 SNMP manager
community private 100.100.51.11 255.255.255.255 readWrite # Tier 1 SNMP manager
view 1.17.2 system enterprises view
trap public 100.100.51.10 1.2.3 fe # Tier 2 SNMP manager
smux 1.3.6.1.4.1.2.3.1.2.1.2 gated_password # gated
smux 1.3.6.1.4.1.2.3.1.2.2.1.1.2 dpid_password #dpid
snmpd smuxtimeout=200 #muxatmd
smux 1.3.6.1.4.1.2.3.1.2.3.1.1 muxatmd_password #muxatmd
smux 1.3.6.1.4.1.2.3.1.2.1.5 clsmuxpd_password # HACMP/ES for AIX clsmuxpd
syslocation "room 2222"
syscontact ssam_admin(ssam_admin@company.com)
```

The trap statement in `/etc/snmpd.conf` defines the system to which the AIX SNMP agent forwards traps that it receives. If you are only configuring SNMP trap forwarding, this is the only entry you need to change.

- c. Start the AIX SNMP agent again to read the new configuration by using the AIX command `startsrc -s snmpd`.
  - d. Ensure that the DPI agent (`dpid2`) has been started. You can check the status with the AIX command `lssrc -s dpid2`. The status must be active (rather than inoperative for a DPI agent that is down). If the DPI agent is not running, start it with the AIX command `startsrc -s dpid2`.
  - e. Check the proper operation of the snmp daemon by running the `tail -f /usr/tmp/snmpd.log` command. The `snmpd.log` is the log file specified in `snmpd.conf`. Verify that no exceptions are logged and that the daemon started with no errors.
2. Create a copy of the IBM Tivoli Storage Manager server options file (`dsmserv.opt`). Depending on the DR550 configuration, you can find the server options file in the following directories:
    - Single-node configuration: `/usr/tivoli/tsm/server/bin`
    - Dual-node configuration: `/tsm/files`Change to the appropriate directory with the AIX command `cd` and use the AIX command `cp dsmserv.opt dsmserv.opt.factory` to create the copy.
  3. Ensure the `dsmsnmp` process is not started by running `ps -ef | grep dsmsnmp`. If one or more `dsmsnmp` processes are started, kill these processes (`kill process-id`)
  4. Modify the server options file `dsmserv.opt` to specify the SNMP communication method.

**Tip:** The `dsmserv.opt` file includes special explanatory notes for all of its parameters. Read them carefully before changing any values.

You must specify the `COMMMETHOD` option for SNMP to enable the protocol for the Tivoli Storage Manager server and SNMP subagent. Because the SNMP Agent and the SNMP Subagent on the DR550 are both running on the local host, the `SNMPSUBAGENT` option and the `SNMPSUBAGENTHOST` option can specify the loopback address (127.0.0.1) for communication.

Use the `SNMPHEARTBEATINTERVAL` option and the `SNMPMESSAGECATEGORY` option to configure sending traps. Decide how often Tivoli Storage Manager should send a heartbeat to report its operating status. Also, decide how the different severities of the SSAM messages will be presented through SNMP traps.

Example 10-34 shows an excerpt of the `dsmserv.opt` file with the SNMP settings.

*Example 10-34 dsmserv.opt with SNMP settings (excerpt)*

<code>commmethod</code>	<code>snmp</code>
<code>snmpsubagent</code>	<code>hostname 127.0.0.1 communityname public timeout 600</code>
<code>snmpsubagentport</code>	<code>1521</code>
<code>snmpsubagenthost</code>	<code>127.0.0.1</code>
<code>snmpheartbeatinterval</code>	<code>5</code>
<code>snmpmessagecategory</code>	<code>SEVERITY</code>

5. Start the IBM Tivoli Storage Manager SNMP subagent (dsmsnmp) by running the dsmsnmp executable. Here are the AIX commands, depending on your DR550 configuration:

- Single-node configuration:

```
cd /usr/tivoli/tsm/server/bin
./dsmsnmp &
```

- Dual-node configuration:

```
export DSMSEV_CONFIG=/tsm/files/dsmserv.opt
cd /usr/tivoli/tsm/server/bin
./dsmsnmp &
```

The IBM Tivoli Storage Manager SNMP subagent starts and shows a message similar to the one in Example 10-35.

*Example 10-35 Start message of dsmsnmp on Data Retention 550 dual-node configuration*

---

```
SNMP Subagent Program for AIX-RS/6000
Tivoli Storage Manager (C)
Copyright IBM Corporation 1990, 2007.
All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication
or disclosure restricted by GSA ADP Schedule Contract with IBM Corporation.
Licensed Materials - Property of IBM.
Version 5, Release 5, Level 0.0
IBM Product ID 5608-ISM, 5608-ISX

ANR0900I Processing options file /tsm/files/dsmserv.opt.
ANR4691I DPI subagent (IBM ADSM subagent): connected, ready to receive requests.
ANR8200I TCP/IP Version 4 driver ready for connection with clients on port 1521.
```

---

Ensure the IBM Tivoli Storage Manager SNMP subagent is running. Use the AIX command `ps | grep dsmsnmp`. It should return one or more processes.

## Configuring SSAM application for SNMP trap forwarding

Do the following steps:

1. Ensure that the DPI agent (dpid2), the AIX SNMP agent (snmpd), and the IBM Tivoli Storage Manager SNMP subagent (dsmsnmp) will be automatically started after system shutdown. These agents should be present prior to the Tivoli Storage Manager server start.

**Tip:** For the DR550 dual-node configuration, you need to start and stop the IBM Tivoli Storage Manager SNMP subagent (dsmsnmp) through the HACMP application server startup script.

2. Stop the IBM Tivoli Storage Manager server (stopserver), if running, and start it again (startserver) to read the new configuration and begin communication through the configured TCP/IP port with the subagent.
3. Enter the Administrative console (dsmadm) and check which events are already enabled or disabled for SNMP by entering the command:

```
q eventrules
```

4. Begin event logging for the SNMP receiver and enable events to be reported to SNMP. For example, issue the following commands on the Tivoli Storage Manager server:

```
begin eventlogging snmp
```

```
enable event snmp [info, warning,error,server,all,event_name]
```

You will have to decide on what categories of the SSAM message should be forwarded to the SNMP manager. If all messages are forwarded, the number of SNMP traps delivered to the SNMP manager might quickly become unmanageable.

**Tip:** You might consider using the `ibmtsm.mac` file, which is installed in the SSAM server bin directory as a starting point.

This macro file defines rules for enabling all messages with the SEVERE and ERROR severity levels as well as some more important events with INFO and WARNING severities. The macro also disables less significant SEVERE and ERROR messages, such as syntax errors. Before executing this macro in your SSAM environment, it has to be customized for your specific needs.

5. Log off.

### ***Automating the SSAM SNMP subagent startup***

The SSAM SNMP subagent (`dsmsnmp`) must be started before or after the SSAM application server is started. It is valuable to automate the start of the SSAM SNMP subagent. Therefore, several UNIX script files have been created to start, stop, and display the status of the SSAM SNMP subagent on a system.

For a dual-node HACMP system, it is desirable to start the Tivoli Storage Manager SNMP subagent automatically on the active node and stop the Tivoli Storage Manager SNMP subagent when a node failover happens. The preferred method is to implement the start-script as an application server within the HACMP resource group. The following steps outline the process to set up an application server to start the Tivoli Storage Manager SNMP subagent:

1. Change the application server start script on all nodes in `/usr/bin` to include the SSAM SNMP subagent start and stop procedures.
2. Ensure that the startup script is identical on both cluster nodes.
3. Synchronize and verify the cluster configuration.
4. Test the SSAM SNMP subagent startup process by moving a resource group from one cluster node to another

The script to stop the Tivoli Storage Manager SNMP subagent shall be implemented within the stopserver script, so every time the server stops the SSAM SNMP subagent is terminated. Note that it is crucial to terminate the Tivoli Storage Manager SNMP subagent because it keeps some files open on the shared file space `/tsm` resulting in errors when the HACMP software unmounts the file system. SSAM SNMP subagent termination shall be done immediately after the SSAM server process is stopped.

The scripts outlined below will be installed at each node in the directory /usr/bin. Ensure that the scripts have execute-permission.

► **Starting the System Storage Archive Manager SNMP subagent**

The script shown in Example 10-36 starts the SSAM SNMP subagent if it is not started already. So, it first checks whether the Tivoli Storage Manager SNMP subagent is running; if not, it starts it. Otherwise, it will present a message indicating that the subagent is already running.

*Example 10-36 Start SNMP agent*

---

```
#start TSM SNMP subagent

dsmsnmp=`ps -ef | grep dsmsnmp | grep -v grep | awk '{print $2}'`
if [["$dsmsnmp" = ""]] then
 echo "TSM SNMP Sub-Agent not started - starting Agent"
 if [[-d /tsm]] then
 echo "Exporting DSMSEV_CONFIG=/tsm/files/dsmserv.opt"
 export DSMSEV_CONFIG="/tsm/files/dsmserv.opt"
 fi
 /usr/tivoli/tsm/server/bin/dsmsnmp &
else
 echo TSM SNMP Sub-Agent already STARTED pid: $dsmsnmp
fi
```

---

► **Display status of System Storage Archive Manager SNMP subagent**

Example 10-37 checks whether or not the Tivoli Storage Manager SNMP subagent is started by looking for dsmsnmp processes. It prints related messages to the screen.

*Example 10-37 Display status of SNMP agent*

---

```
#show whether dsmsnmp is started
pids=`ps -ef | grep dsmsnmp | grep -v grep | awk '{print $2}'`

if [["$pids" == ""]] then
 echo "Tivoli Storage Manager SNMP Sub-Agent not started - starting Agent"
else
 echo TSM SNMP Sub-Agent already STARTED pid=$pids
fi
```

---

**Stopping the System Storage Archive Manager SNMP subagent**

The script shown in Example 10-38 terminates the SSAM SNMP subagent where appropriate. It first checks whether the SSAM SNMP subagent is running. If this is true, it obtains all process IDs and sends a SIGTERM signal to these processes. Otherwise, it will present a message indicating that the agent is not running.

Note that there might be more than one process started for the SSAM SNMP subagent. This seems to be normal.

*Example 10-38 Stop SNMP agent*

---

```
stop the TSM SNMP Subagents dsmsnmp

pids=`ps -ef | grep dsmsnmp | grep -v grep | awk '{print $2}'`
```

```

if [["$pids" = ""]] then
 echo "TSM SNMP Subagent not started."
else
 echo TSM SNMP Subagent Processes: $pids
 for p in $pids
 do
 echo "Terminating process ID $p."
 kill $p
 done
fi

```

---

## Verifying the SSAM SNMP subagent functionality

After all SNMP agents and subagents are configured and started as described earlier, a test SNMP get query can be used to verify the connectivity and functionality of all the components (see Example 10-39).

### *Example 10-39 Retrieving the SSAM application server name*

```

snmpinfo -m get -v -h 127.0.0.1 ibmAdsmMServerName.0
ibmAdsmMServerName.0 = " TSM"

```

---

The command **snmpinfo** is a standard AIX utility built for querying and changing SNMP MIB variables. In Example 10-39, the command is run on an HACMP cluster node where the SSAM application is running (communication through the loopback IP address 127.0.0.1). In this example (Example 10-40), **snmpinfo** is used to query the `ibmAdsmMServerName` SSAM MIB variable. If any error message is returned instead of the TSM string, the root cause of the error must be identified and fixed before continuing with configuration of other SNMP components.

### *Example 10-40 Retrieving SSAM application server heartbeat*

```

snmpinfo -m get -v -h 127.0.0.1 ibmAdsmServerHeartbeat.1
ibmAdsmServerHeartbeat.1 = " Mar 11 17:18:59 2008

```

---

**Note:** The heartbeat value is returned for the GMT time zone.

Because the SNMP environment has weak security, you should consider not granting SNMPADMIN any administrative authority. This restricts SNMPADMIN to issuing only SSAM queries. Although a password is not required for the subagent to communicate with the server and run scripts, a password should be defined for SNMPADMIN to prevent access to the server from unauthorized users. See Example 10-41.

### *Example 10-41 Running SSAM application server script*

```

dsmadm
IBM Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 5, Level 0.0
(c) Copyright by IBM Corporation and other(s) 1990, 2007. All Rights Reserved.

```

```

Enter your user id: admin
Enter your password:

```

```

Session established with server TSM: AIX-RS/6000
 Server Version 5, Release 5, Level 0.0
 Server date/time: 03/11/08 11:27:29 Last access: 03/11/08 11:26:22
tsm: TSM>register admin SNMPADMIN s8cr8t

```

```

ANR2068I Administrator SNMPADMIN registered.

tsm: TSM>grant authority SNMPADMIN class=operator
ANR2082I Operator privilege granted to administrator SNMPADMIN.

tsm: TSM>define script QUERY_DB "QUERY DB" descr="Demo script for testing SNMP subagent
functionality"
ANR1454I DEFINE SCRIPT: Command script QUERY_DB defined
tsm: TSM>quit

ANS8002I Highest return code was 0.

snmpinfo -m set -c private -h 127.0.0.1 ibmAdsmServerScript1.1=QUERY_DB
1.3.6.1.4.1.2.6.135.1.2.1.3.1.2.1 = "QUERY_DB"

snmpinfo -m get -v -h 127.0.0.1 -c private ibmAdsmM1ReturnValue.1
ibmAdsmM1ReturnValue.1 = "
Available Assigned Maximum - Maximum - Page S- Total Us- Used Pag- Pct - Max.-
Space (M- Capacity Extension Reduction ize (b- able Pag- es Util Pct-
 B) (MB) (MB) (MB) ytes) es es Util Util

 300,000 300,000 0 299,932 4,096 76,800,00 82,330 0.1 0.1
 0
ANR1462I RUN: Command script QUERY_DB completed successfully.
"

```

A SSAM script can also be defined to accept input parameters, in which case these parameters will have to be set in SSAM MIB prior to the script invocation (see Example 10-42).

---

*Example 10-42 Running SSAM application server script with input parameters*

---

```

dsmadm
IBM Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 5, Level 0.0
(c) Copyright by IBM Corporation and other(s) 1990, 2007. All Rights Reserved.

Enter your user id: admin

Enter your password:

tsm: TSM>define script QUERY_IO "select sum(BYTES) as RETRIEVED_IN_$1_H from summary where
activity='RETRIEVE' and START_TIME>(CURRENT_TIMESTAMP - ($1 hours))" descr="Demo script
with parameters for testing SNMP subagent functionality"
ANR1454I DEFINE SCRIPT: Command script QUERY_IO defined.
tsm: TSM>quit

ANS8002I Highest return code was 0.

snmpinfo -m set -c private -h 127.0.0.1 ibmAdsmServerScript1.1=QUERY_IO
1.3.6.1.4.1.2.6.135.1.2.1.3.1.2.1 = "QUERY_IO"

snmpinfo -m set -c private -h 127.0.0.1 ibmAdsmM1Parm1.1="48"
1.3.6.1.4.1.2.6.135.1.2.1.3.1.3.1 = "48"

snmpinfo -m get -v -h 127.0.0.1 -c private ibmAdsmM1ReturnValue.1
ibmAdsmM1ReturnValue.1 = "
RETRIEVED_IN_48_H

26736789

```



ANR1462I RUN: Command script QUERY\_IO completed successfully.

---

The values assigned to `ibmAdsmM1Parm1.1` and `ibmAdsmServerScript1.1` will not be preserved after the SSAM SNMP subagent is restarted or if the SSAM application is moved from one HACMP node to another. Another SNMP GET query that can be used to obtain all SSAM related SNMP MIB variables is `snmpinfo -m dump -v -h 127.0.0.1 ibmAdsm`. The sample output of this command is not provided due to its size.

## 10.5 Configure SNMP monitoring for DR550 Storage Controller

The DR550 Storage Controller is a DS4200 Storage subsystem.

The Event Monitor is a feature of the DS4200 Storage Manager that monitors system health, and uses the Storage Manager client alerting feature to automatically notify a configured recipient when problems occur.

SNMP trap is one of the available alert notification options. Messages are only sent as a result of critical events in the storage system. This means it is not possible to configure the type of events; it is only for critical events.

### Configure SNMP alerts on DR550 Storage Server

To set up the alert environment and specify destinations, proceed as follows:

1. From each DR550 SSAM Server node, start the SMclient application. Use the AIX user `dr550` to log in. Then you have to run `su -` and switch to root. After you run `startx`, you get the graphical interface of AIX; run `SMclient`, which displays the DS4000 Storage Manager Enterprise Management window, and you will see a window similar to the one shown in Figure 10-9.

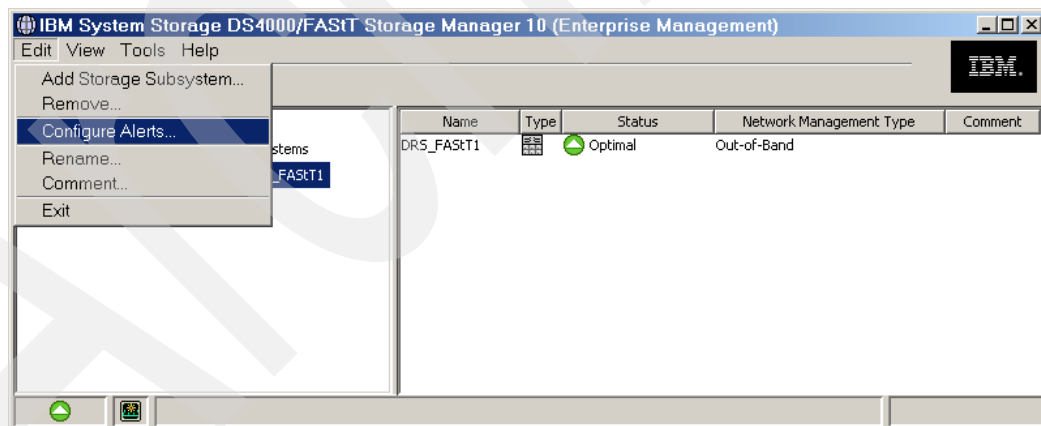


Figure 10-9 IBM DS4000 Storage Manager client Enterprise Management window

From the DS4000 Storage Manager for DR550 Enterprise Management window, select **Edit** → **Configure Alerts** in the upper left side of the window.

## SNMP alert destination

Do the following steps:

1. In the DS4000 Storage Manager client Enterprise Management window (Figure 10-10), select **Edit** → **Alert Destinations**, and then set the **SNMP**.
2. Enter the SNMP community name in the appropriate field. The SNMP community name is set in the NMS configuration file by a network administrator. The default is public.
3. Enter the trap destination IP or host name in the appropriate field. The SNMP trap destination is the IP address from the RSM Server or DR550. In our example, we will be using the Net-SNMP server (RSM Server for DR550) IP address at the SNMP trap recipient (100.100.51.10).

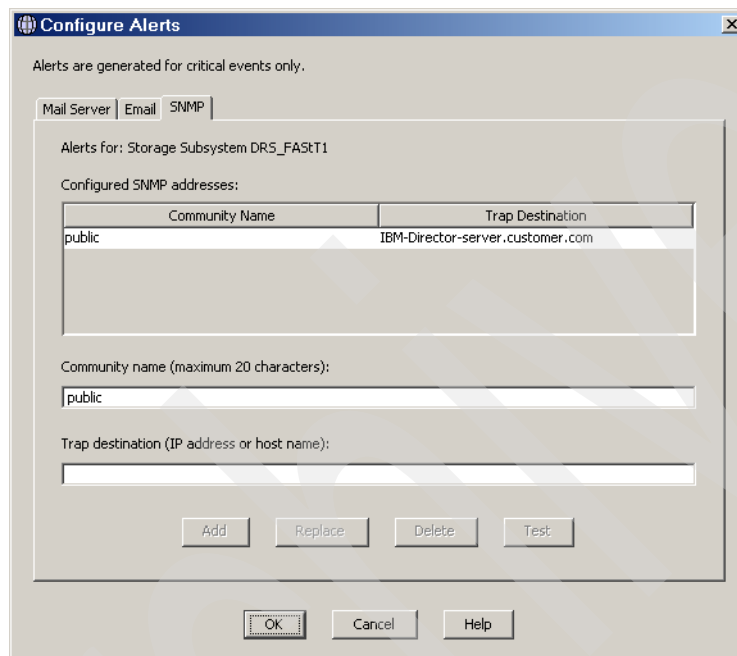


Figure 10-10 SNMP trap destination

4. Click **Add**.
5. To validate an SNMP address, type the address in the text box, and then click **Test**. Check the SNMP destination for SNMP alert validation.
6. Click **OK**.

## 10.6 Configure SNMP monitoring for SAN Switch

The DR550 SAN Switch (2005-B16) is capable of sending notifications through the Simple Network Management Protocol (SNMP). In order for the Net-SNMP server to receive a trap generated by the DR550 SAN Switch, it must use an existing network connection to the internal DR550 Ethernet switches. In our example, the Net-SNMP application is configured on the RSM server utilizing its connections to both internal DR550 Ethernet switches.

## Configure SNMP on DR550 SAN Switches

If your DR550 is equipped with a SAN switch or switches, the connection to the DR550 Ethernet switches will be established and a default IP address will be assigned to the SAN switch or switches. To access the administrative interface on a SAN switch, a command-line tool such as ssh, telnet, or rsh can be used. In our example, we will be using the ssh client on the RSM server for that purpose. The SNMP agent configuration of the SAN switch is managed with the **snmpConfig** administrative command. This command can be used to customize the community string, trap recipient's IP address, access control parameters, system location, system description, system contact, and enable or disable authentication failure traps. An illustration is shown in Example 10-43.

*Example 10-43 Configuring SAN switch systemGroup SNMP parameters*

---

```
ibmrsm # ssh admin@192.168.4.31
admin@192.168.4.32's password:
IBM_2005_B16:admin> snmpConfig --set systemGroup
Customizing MIB-II system variables ...
At each prompt, do one of the following:
 o <Return> to accept current value,
 o enter the appropriate new value,
 o <Control-D> to skip the rest of configuration, or
 o <Control-C> to cancel any change.
To correct any input mistake
<Backspace> erases the previous character,
<Control-U> erases the whole line,
sysDescr: [Fibre Channel Switch.] DR550 FC switch1
sysLocation: [End User Premise.] room 2222
sysContact: [Field Support.] SAN Administrator (san_admin@company.com)
authTrapsEnabled (true, t, false, f): [false] t
Committing configuration...done.
IBM_2005_B16:admin> snmpConfig --show systemGroup
 sysDescr = DR550 FC switch1
 sysLocation = room 2222
 sysContact = SAN Administrator (san_admin@company.com)
 authTraps = 1 (ON)
IBM_2005_B16:admin>
```

---

Authentication traps are disabled by default. In this example, the default behavior which suppresses authentication trap generation is changed to enabled (authTraps = 1 (ON)). You will have to decide whether to enable or disable this type of traps in your environment. The examples of SNMP trap handlers provided in this section assume that authentication traps are enabled.

In Example 10-44, 192.168.4.100 is the IP address of the RSM Ethernet interface configured in the same subnet as the Ethernet interface to the SAN switch. Each SAN switch can be configured to send snmpv1 traps to up to six trap recipients. For each recipient (identified by its IP address and an individual community string), an event severity level can be configured. In our example, we configure a single trap recipient IP address with the “public” community string and severity level 5. All other recipients are left unconfigured (IP address 0.0.0.0). By specifying the event severity, you can limit the number of traps that will be forwarded to a specific trap recipient. The specified number will allow forwarding traps with a severity less than or equal to the number specified. SW\_v5 MIB defines the event trap and severity levels shown in Table 10-3.

*Example 10-44 Configuring SAN switch snmpv1 SNMP parameters*

```
ibmrsm # ssh admin@192.168.4.31
admin@192.168.4.32's password:
IBM_2005_B16:admin> snmpConfig --set snmpv1
SNMP community and trap recipient configuration:
Community (rw): [Secret C0de] private
Trap Recipient's IP address : [0.0.0.0] 192.168.4.100
Trap recipient Severity level : (0..5) [0] 5
Community (rw): [OrigEquipMfr]
Trap Recipient's IP address : [0.0.0.0]
Community (rw): [private]
...
Trap Recipient's IP address : [0.0.0.0]
Committing configuration...done.
```

*Table 10-3 DR550 SAN Switch event trap and severity levels*

Event trap level	Event severity level
None	0
Critical	1
Error	2
Warning	3
Informational	4
Debug	5

By default, a trap recipient severity level is set to “0” (none), which effectively disables trap forwarding to a recipient. In our example, we set the severity level to “5”, which enables all traps to be forwarded to a recipient. The setting that enables all trap forwarding is only used for demonstration purposes and might not be practical for an implementation in a real production environment (see Example 10-45).

*Example 10-45 Configuring SAN switch access control SNMP parameters*

```
ibmrsm # ssh admin@192.168.4.31
admin@192.168.4.32's password:
IBM_2005_B16:admin> snmpConfig --set accessControl
SNMP access list configuration:
Access host subnet area : [192.168.4.100]
Read/Write? (true, t, false, f): [true] t
Access host subnet area : [0.0.0.0]
Read/Write? (true, t, false, f): [true]
```

```

...
Access host subnet area : [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Committing configuration...done.
IBM_2005_B16:admin> snmpConfig --show accessControl

SNMP access list configuration:
Entry 0: Access host subnet area 192.168.4.100 (rw)
Entry 1: No access host configured yet
...
Entry 5: No access host configured yet
IBM_2005_B16:admin>

```

It is not necessary to update SNMP access control parameters unless you plan to use SNMP for other functions in addition to monitoring. In this example, Net-SNMP server IP address 192.168.4.100 is granted read-write access to SNMP variables of the DR550 SAN switch. Not all MIB variables are defined with read-write access. Examples of read-write variables are swDomainID, swBeaconOperStatus, and swAdmStatus. By setting the values of the MIB variables, it is possible, for example, to change the switch operational status or operational status of the switch beacon (see Example 10-46).

*Example 10-46 Modifying operational status of the switch beacon*

```

ibmrsm # snmpget -v 3 -u snmpadmin1 -c public 192.168.4.31 SW-MIB::swBeaconAdmStatus.0
SW-MIB::swBeaconAdmStatus.0 = INTEGER: off(2)
ibmrsm # snmpset -v 3 -u snmpadmin1 -c public 192.168.4.31 swBeaconAdmStatus.0 = 1
SW-MIB::swBeaconAdmStatus.0 = INTEGER: on(1)

```

**Note:** MIB variables can only be modified using SNMP Version 3 protocol (-v 3 parameter). SNMP Version 1 cannot be used for updating SNMP variables. You must assign an authentication pass phrase and a privacy pass phrase used for encrypted SNMPv3 messages.

The SNMP agent on the SAN switch can be configured to enable or disable a specific MIB. Table 10-4 includes the list of supported MIBs and the default status of the MIB.

*Table 10-4 DR550 SAN Switch MIB capability*

MIB	Enabled	Brief description
FE-MIB	YES	Fibre Channel Fabric Element MIB
SW-MIB	YES	MIB for Brocade's Fibre Channel Switch.
FA-MIB	YES	Fibre Alliance FC Management Framework Integration MIB
FICON-MIB	NO	MIB module support for FICON
HA-MIB	YES	High Availability MIB
FCIP-MIB	NO	MIB specific to FCIP devices

FICON and FCIP functionality are not used in the DR550 product. FICON-MIB and FCIP-MIB MIBs should remain disabled. Each enabled MIB implements a number of SNMP traps. These traps can also be disabled individually. In our example, we only use traps defined in SW-MIB, which by default has all traps enabled. To customize the list of enabled MIBs as well as individual traps, use the following SAN switch administrative command:

```
snmpConfig --set mibCapability
```

## 10.7 Configure SNMP monitoring for FSG

The following sections detail the procedure to enable SNMP on the DR550 FSG.

### Prerequisites

- ▶ You need an IBM Director server in the network.
- ▶ You must know the IP address of your SNMP server (IBM Director Server) and the configured SNMP protocol version (V1 or V2c.)

### Procedure

Do the following steps:

1. Log into the main DR550 FSG using the fsgadm user ID and password. Then issue `su -` to switch to root.
2. Change to the `/etc/snmp` directory by entering the following command:  

```
cd /etc/snmp
```
3. Open the `snmpd.conf` file using an editor like `vi`, as shown in Figure 10-11:  

```
vi snmpd.conf
```

```
Please see /usr/share/doc/packages/net-snmp/EXAMPLE.conf for a
more complete example and snmpd.conf(5).
#
Writing is disabled by default for security reasons. If you'd like
to enable it uncomment the rwcommunity line and change the community
name to something nominally secure (keeping in mind that this is
transmitted in clear text).
#
don't use ' < >' in strings for syslocation or syscontact
Note that if you define the following here you won't be able to change
them with snmpset
syslocation Server Room
syscontact Sysadmin (root@localhost)
#
These really aren't meant for production use. They include all MIBS
and can use considerable resources. See snmpd.conf(5) for information
on setting up groups and limiting MIBS.
rocommunity public 127.0.0.1
rwcommunity mysecret 127.0.0.1
~
~
~
~
"snmpd.conf" 19L, 811C 19,1 All
```

Figure 10-11 Edit `snmpd.conf` file

4. Change the following values:

syslocation = location of the machine (Server Room)

syscontact = a responsible contact person for this machine

rocommunity = community name & IP address of your SNMP server

Example:

syslocation Room 1110

syscontact Mr. Ron (Ron@company.com)

rocommunity public 10.10.10.24

5. Save the file.

6. Copy the snmpd.conf file to the /usr/share/snmp directory by typing the following command:

```
cp snmpd.conf /usr/share/snmp/
```

7. Type the following command to import your configuration file:

```
snmpconf
```

You should see the output shown in Figure 10-12.

```
main:/ # snmpconf

The following installed configuration files were found:

1: ./snmpd.conf
2: /etc/snmp/snmpd.conf

Would you like me to read them in? Their content will be merged with the
output files created by this session.

Valid answer examples: "all", "none", "3", "1,2,5"

Read in which (default = all):
```

Figure 10-12 The snmpconf - select file

Select the /etc/snmp/snmpd.conf file to update by typing 2 and pressing Enter.

8. The next menu is displayed, as shown in Figure 10-13. Select the snmpd.conf file by typing 1 and pressing Enter.

```
Read in which (default = all): 2

I can create the following types of configuration files for you.
Select the file type you wish to create:
(you can create more than one as you run this program)

1: snmpd.conf
2: snmptrapd.conf
3: snmp.conf

Other options: quit

Select File:
```

Figure 10-13 Select snmp file type to create

9. The next menu is displayed, as shown in Figure 10-14. Select the option to update Trap Destinations by typing 5 and pressing Enter.

```
Select File: 1

The configuration information which can be put into snmpd.conf is divided
into sections. Select a configuration section for snmpd.conf
that you wish to create:

 1: Access Control Setup
 2: Extending the Agent
 3: Monitor Various Aspects of the Running Host
 4: Agent Operating Mode
 5: Trap Destinations
 6: System Information Setup

Other options: finished
Select section:
```

Figure 10-14 Select trap destination

- 10.. The next menu(Figure 10-15) lets you select the trap destination.

```
Select section: 5

Section: Trap Destinations
Description:
 Here we define who the agent will send traps to.

Select from:

 1: A SNMPv1 trap receiver
 2: A SNMPv2c trap receiver
 3: A SNMPv2c inform (acknowledged trap) receiver
 4: A generic trap receiver defined using snmpcmd style arguments.
 5: Default trap sink community to use
 6: Should we send traps when authentication failures occur

Other options: finished, list
Select section:
```

Figure 10-15 Trap catcher screen

11. Select SNMPv1 or SNMPv2c by typing either 1 or 2 and pressing Enter.

The choice is based on the protocol used by the SNMP catcher installed. Refer to your documentation for the SNMP catcher software.

We selected 2 in our example (Figure 10-16 on page 443). Your version of SNMP catcher might be different.



```

Select section: 2

Configuring: trap2sink
Description:
 A SNMPv2c trap receiver
 arguments: host [community] [portnum]

A host name that should receive the trap:

```

Figure 10-16 Trap receiver host name

12. Set up the IP address, community and port by completing the following steps (Figure 10-17):

- a. Type the IP address of the server running the SNMP catcher software and press Enter.
- b. Type the community name and press Enter.
- c. Type the port number the SNMP catcher software is listening on and press Enter.

```

A host name that should receive the trap: 127.0.0.1
The community to be used in the trap sent [optional]: DR550 FSG
The port number the trap should be sent to [optional]: 5989

Finished Output: trap2sink 127.0.0.1 "DR550 FSG" 5989

Section: Trap Destinations
Description:
 Here we define who the agent will send traps to.

Select from:

 1: A SNMPv1 trap receiver
 2: A SNMPv2c trap receiver
 3: A SNMPv2c inform (acknowledged trap) receiver
 4: A generic trap receiver defined using snmpcmd style arguments.
 5: Default trap sink community to use
 6: Should we send traps when authentication failures occur

Other options: finished, list

Select section:

```

Figure 10-17 Community IP address and port

While our example shows localhost IP address, the actual IP address will not be 127.0.0.1. It should be the IP address of the IBM Director in the network (Figure 10-18).

```
A host name that should receive the trap: 127.0.0.1
The community to be used in the trap sent [optional]: DR550 FSG
The port number the trap should be sent to [optional]: 5989

Finished Output: trap2sink 127.0.0.1 "DR550 FSG" 5989

Section: Trap Destinations
Description:
 Here we define who the agent will send traps to.

Select from:

 1: A SNMPv1 trap receiver
 2: A SNMPv2c trap receiver
 3: A SNMPv2c inform (acknowledged trap) receiver
 4: A generic trap receiver defined using snmpcmd style arguments.
 5: Default trap sink community to use
 6: Should we send traps when authentication failures occur

Other options: finished, list

Select section:
```

Figure 10-18 Finish SNMP configuration

13.Exit by typing finished and pressing Enter.

14.Exit by typing finished and pressing Enter again (Figure 10-19)

```
 1: Access Control Setup
 2: Extending the Agent
 3: Monitor Various Aspects of the Running Host
 4: Agent Operating Mode
 5: Trap Destinations
 6: System Information Setup

Other options: finished

Select section: finished

I can create the following types of configuration files for you.
Select the file type you wish to create:
(you can create more than one as you run this program)

 1: snmpd.conf
 2: snmptrapd.conf
 3: snmp.conf

Other options: quit

Select File:
```

Figure 10-19 Quit SNMP configuration task

15.Finally, exit by typing quit and pressing Enter (Figure 10-20 on page 445).

```

Select File: quit

Error: An snmpd.conf file already exists in this directory.

'overwrite', 'skip', 'rename' or 'append'? : overwrite

The following files were created:

 snmpd.conf

These files should be moved to /usr/share/snmp if you
want them used by everyone on the system. In the future, if you add
the -i option to the command line I'll copy them there automatically for you.

Or, if you want them for your personal use only, copy them to
/root/.snmp . In the future, if you add the -p option to the
command line I'll copy them there automatically for you.

main:/ #

```

Figure 10-20 Confirm SNMP configuration changes

If the program states that there is an existing file, type `overwrite` and press Enter.

Repeat steps 2-15 for the supplementary FSG, if you have one.

The configuration is finished and the FSG nodes can now send their SNMP traps to the specified SNMP server.

### ***Additional IBM Director server tasks***

You need to do the following additional tasks:

1. Go to the IBM Director Server and, from the console, select the function to discover the IBM DR550 FSGs.
2. Select the IBM DR550 FSG and then select the request access function.
3. Enter the user ID and password of the FSG when prompted.

### ***Testing your environment***

If you have redundant power supplies, pull the power cable of the power supply<sup>1</sup> at the IBM DR550 FSG. An alert should display at the IBM Director Server within minutes

Archived

## Tape attachment

In this chapter, we explain how to use tape devices to enhance the standard capabilities of DR550 Models DR1 and DR2. The versatile possibilities of IBM System Storage Archive Manager (SSAM) allow for a complete hierarchical storage management solution that can benefit from tape media for long-term archiving. In the short term, removable media such as tape is also the answer when you need to back up data archived on DR550 Models DR1 and DR2 disks.

## 11.1 Tape device attachment

IBM System Storage Archive Manager (SSAM) uses tape media for the following purposes:

- ▶ Migrating archived data

Migrating data off of the primary disk storage pool over time or after a certain percentage of the pool capacity has been reached onto a tape storage pool can tremendously extend, at a reasonable cost, the storage capacity of the DR550. The nature of archived data is that it is accessed more frequently shortly after its creation, but less and less frequently as time passes, which increases the need to move data over time from disks to less expensive media such as tape. The inevitable longer access time to retrieve a document in the future is tolerable in most cases.

- ▶ Backing up archived data

Keeping a backup of the archived data in a remote tape storage pool protects it from disasters that could happen to the disk storage pool at the site where the DR550 is located. With data being an essential asset for many enterprises, having the capability to recover from a disaster by restoring from the tape backup pool is invaluable.

In DR550 environments with Enhanced Remote Mirroring (ERM), the tape attachment is mainly used to enhance storage capacity for long term storage.

- ▶ Backing up the SSAM database

Like the actual data, the SSAM database needs protection from the very same scenarios just described. Without the SSAM database, access to the archived data is impossible. Therefore, backing up the database to tape (preferably to a remote site) is as vital as the backup of the original data.

In DR550 environments with Enhanced Remote Mirroring (ERM) enabled, it is also critical to back up the SSAM database to tape in order to protect against possible database corruption.

We document two examples showing you how to attach and use tape in the context of the DR550 archiving solution. The first example explains the basics of attaching a 3494 Enterprise Tape Library with 3592-J1A or 3592-E05 (TS1120) Enterprise Tape Drive technology and shows the creation of a simple backup pool to hold a copy of the archived data. The second example, based on a IBM System Storage TS3500 Tape Library with 3588 Ultrium 3 Tape Drive technology and using 3589 Write Once Read Many (WORM) media, illustrates the creation of a migration storage pool and a backup copy pool and shows you how to create additional SSAM database backups. For the latter, regular 3589 rewritable media is used, because WORM media is not economically feasible for database backups that roll over frequently.

The placement of a tape library depends on its intended use. Migration does not protect against failures, but rather expands the capacity. In this case, the library does not necessarily need to be placed in a different, remote site. In fact, if the migrated data resides in a remote library, both sites would have to be protected against disasters. On the other hand, keeping backups is indeed intended to prevent the loss of data in case of an outage or disaster at the primary site where the DR550 resides. In this case, the tape library should be physically separated from the DR550 and placed in a protected section or location. Ideally, if both methods are used (migration and backup), we recommend the installation of two libraries, one locally for migration and the second remote for backups. The example in this chapter, however, does not focus on these considerations.

With regards to the DR550, it should be pointed out that the term *Write Once Read Many* (WORM) is appropriate when the media is non-rewritable and non-erasable, such as WORM optical and WORM tape. (Therefore, using the term “WORM disk” can be misleading in that in today’s technology, the disk does not possess these attributes.)

**Note:** We recommend WORM-capable tape drive technology as the solution of choice to complement the DR550 Models DR1 and DR2. The TS1120 Tape Drive is one example of WORM capability, along with a variety of 3592 media, among different storage capacities. This technology is available for two IBM automated library products: the 3494 Enterprise Tape Library and the TS3500 UltraScalable Library. Existing libraries can be upgraded in the field to incorporate 3592 technology, thus providing a high degree of investment protection.

The Ultrium 3, 400 GB WORM Data Cartridge adds WORM support for the IBM System Storage LTO Ultrium 3 family of products. Note that existing LTO tape drives require a firmware upgrade for full product support.

## IBM tape drive and optical library overview

The following IBM tape drives, tape libraries, and optical libraries can be attached to the DR550.

### IBM tape drives

- ▶ TS1120 (supports Drive Encryption and dual drive path)
- ▶ LTO Generation 3 tape drive
- ▶ LTO Generation 4 tape drive (supports Drive Encryption)

### IBM tape libraries

- ▶ TS3100 (for LTO 3 and LTO 4 tape drives)
- ▶ TS3400 (only for TS1120 tape drives)
- ▶ TS3500 (for TS1120, LTO 3 and LTO 4 tape drives)
- ▶ TS3310 (for LTO 3 and LTO 4 tape drives)
- ▶ TS3200 (for LTO 3 and LTO 4 tape drives)
- ▶ 3494 Enterprise Library (only for TS1120 tape drives)

Go to <http://www.ibm.com/systems/storage/tape/index.html> if you want to know more about the models and features of IBM tape drives and libraries.

### Optical library

- ▶ IBM 3996

Refer to <http://www.ibm.com/systems/storage/optical/3996/index.html> to get more information about the models and features of the 3996.

**Note:** It might be necessary to update the device driver, SSAM version, or tape drive firmware to get full product support.

## 11.2 Planning tape attachment

To efficiently plan for tape attachment to the DR550, you should understand:

- ▶ Tape device and media technology, and product names
- ▶ Available functionality
- ▶ Number of tape libraries and tape drives required
- ▶ Available storage capacity

The tape devices can be used to strengthen data integrity and to prepare for disaster recovery. Tape is an ideal medium for these tasks because it can easily be moved to an off-site location. Another reason is the cost/MB ratio of tape media, which is still less expensive than disk media even with the SATA disk devices.

### IBM System Storage Archive Manager tape pools

When using SSAM, the technical reasons to establish a storage hierarchy, which includes disk and tape, are based on the different functions the product offers:

- ▶ Backup of storage pools (copy pools)
- ▶ Data migration
- ▶ SSAM Database backup (DBB)

### Supported tape devices

SSAM supports manual and automated tape devices:

- ▶ *Manual tape devices* are devices operated by the administrator because they do not have any automated functionality or the hardware necessary for automation. For example, any stand-alone tape drive is considered to be a manual tape device. The tapes are mounted and dismounted by the administrator, and the storage of tape volumes is under the control of the administrator.
- ▶ *Automated tape devices* have the hardware (such as cartridge accessor, storage slots, and input/output slots) and functionality to operate without administrator intervention. Mounting and dismounting tape volumes or storage of volumes within the library is fully automated. Whenever possible, we recommend that you choose automated tape devices over manual tape devices.

Tape devices are defined to SSAM through library and drive definitions. Each physical library (of whatever tape technology) is associated with or mapped to a *tape device class definition*. The device class definition informs SSAM about the type of drive being used, for example, the format and capacity. Tape drives within a large tape library can be logically grouped to meet performance requirements for different groups of data. This is illustrated in 11.3, “IBM tape libraries and drives: examples” on page 454.

**Tip:** SSAM supports an extensive list of tape drives, autoloaders, libraries, and optical devices. For a full list of supported devices, refer to the following Web site:

[http://www.ibm.com/software/sysmgmt/products/support/IBM\\_TSM\\_Supported\\_Devices\\_for\\_AIXHPSUNWIN.html](http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_AIXHPSUNWIN.html)

### Storage pools

Tape storage pools are typically used within SSAM for both primary and copy storage pools. To create copies of a primary object, Tivoli Storage Manager needs to back up the primary object. This process can be automated to create copies on a daily basis.



## Migration

The physical location of an object within the storage pool hierarchy has no effect on its retention policies. Migrating objects to another storage media such as tape can free up storage space on higher-performance devices such as disks.

## SSAM database backup (DBB)

The backups of the SSAM database do not belong to a storage pool, and they cannot be copied. The SSAM Server includes database backups to disk as a preconfigured feature. With the attachment of tape devices, you can also back up the database to tape. This provides additional security, and can be scheduled to run automatically every day.

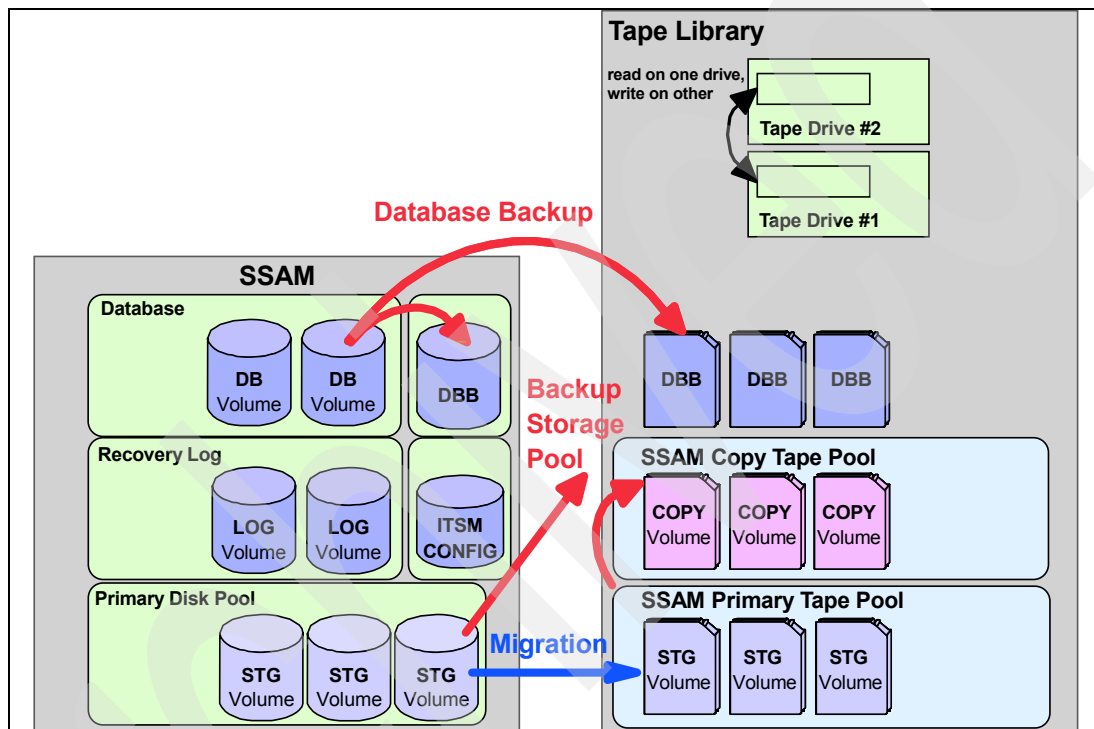


Figure 11-1 Use of tape attachment for DR550

## Beyond the DR550

Obviously, there might be other reasons to attach tape devices to the DR550. For example, if you want to create system images of the operating system with the SMIT mksysb or savevg functions, or if you want to use the Tivoli Storage Manager BA client that is part of the SSAM Server to back up the operating system and files on the system to an additional, externally located Tivoli Storage Manager server. You also might need a tape drive to perform administrative tasks such as loading software images or updates on the SSAM Server.

We assume that you do such backups through the network connection of the SSAM Server. With this connection, you can use AIX products, such as Network Installation Management (NIM) products, Tivoli Storage Manager for System Backup and Recovery, or the Tivoli Storage Manager backup-archive client.

The Tivoli Storage Manager backup-archive client can be connected to any other available server within your environment outside the DR550. This other server can also share (use) the library and tape devices of the DR550.

The IBM Tivoli Storage Manager for System Backup and Recovery product leverages the functionality of mksysb and savevg and can help you prepare for recovery on the AIX operating system in case of a major failure or disaster. Moreover, the product can be combined with NIM to create boot images, system backups, or even data volume backup through the network. When you need to recover any of these backups, the network is used again (such as network bootable system images). Using the Tivoli Storage Manager backup-archive client together with Tivoli Storage Manager for System Backup and Recovery defines a disaster recovery solution for the DR550. The target devices for the storage of those backups belong to any other Tivoli Storage Manager server within your environment on the same network as the SSAM Server. To learn more about IBM Tivoli Storage Manager for System Backup and Recovery, visit:

<http://www.ibm.com/software/tivoli/products/storage-mgr-sysback/>

As part of IBM Tivoli Storage Manager Extended Edition that ships with the DR550, you get the Disaster Recovery Manager (DRM). This feature is not configured, since it leverages the use of tape storage devices. Disaster Recovery Manager helps with the preparation tasks for disaster recovery. One of these tasks is to keep track of disaster recovery media, in particular, IBM Tivoli Storage Manager database backups (DBB) and copy pool volumes, both residing on tape media. DRM helps to manage the off-site storage of these tape media. It also creates a disaster recovery plan file that includes essential details regarding the most current IBM Tivoli Storage Manager installation. Altogether, maintaining the recovery plan file (which contains contact names and telephone numbers for critical people and their backups) plus DBB and copy pool tape media in an off-site location will ease the recovery in time for problems, errors, or even disasters.

**Tip:** The IBM System Storage Archive Manager Extended Edition within the DR550 includes the Disaster Recovery Manager (DRM). We suggest that you combine tape attachment with the use of DRM as a disaster recovery solution for the DR550.

It is beyond the scope of this book to describe those products, features, or processes in detail.

### 11.2.1 Hardware

The DR550 Model DR1 provides two internal Fibre Channel Ports in the DR550 SSAM Server for tape support. For tape devices or library, use Port 2 in Slot 1 and Port 2 in Slot 4.

The required hardware to attach three Fibre Channel-based tape devices per switch is provided with the DR550 Model DR2. These tape-ready configurations come with small form-factor pluggables (SFPs) for the DR550 SAN Switch (2005-B16). Zoning definitions within the switch have already been preloaded in the factory. Basically, only the required Fibre Channel cable connections between the tape devices/library and the SAN switch have to be established in order for the p5 servers to be able to communicate with the devices. Ports 5, 6, and 7 on the DR550 SAN Switch are usable for tape devices/library.

If libraries other than the IBM 3494 or TS3500 will be used, additional tasks might be required to set up a control path between the engine or engines and the library. In this case, consult the documentation or vendor for these libraries.

If you want to attach a optical library, such as an IBM 3996 or tape device with a LVD SCSI interface, you must order the optional LVD SCSI Adapter (feature code 3560) for the DR550 Models DR1 and DR2, which will be installed in slot 3 of the DR550 SSAM Server.

## 11.2.2 Software

When the hardware is in place, additional software (device drivers) is typically required for the operating system to communicate with the devices. In addition, the application (SSAM) has to be configured to use the tape devices. The DR550 SSAM Server has preinstalled device drivers for the IBM tape products: 3494 Enterprise Tape Library, TS3500 Tape Library, and 3592 Enterprise Tape Drive. Therefore, installation of this software is normally not required, only verification and then SSAM integration of the devices, which will be described later in this chapter. However, it might be necessary to install a newer version of the software components to benefit from further enhancements.

### Device drivers

Device drivers have to be installed on each of the DR550 SSAM Servers for the tape library and the tape drives. Depending on the devices, this can be only one driver, separate drivers for library and drives, or sometimes, the SSAM device driver. Installing the device drivers means making the devices available within the operating system. It does not mean making the devices available for use by any application.

The general procedure for the installation of tape device drivers is as follows:

1. Attach the devices to your system, which includes cabling.
2. Check the SSAM (Tivoli Storage Manager V5.5.0) guides for the recommended device drivers.
3. Obtain the most current version of the device drivers.
4. Obtain the latest installation instructions and user guides.
5. Install all the necessary device drivers.
6. Configure the devices within AIX.
7. Determine the device names within AIX.

Because the devices are used exclusively by SSAM, check the IBM System Storage Archive Manager (Tivoli Storage Manager V5.5.0) publications first to see what the recommended device drivers are for your devices. Even if your devices come with their own device drivers and are shipped by the manufacturer, you have to follow the recommendations in the SSAM documentation.

For all IBM tape drives and IBM tape libraries, we recommend using the IBM tape device driver, which is already preinstalled in the DR550 Models DR1 and DR2. It might be necessary to update the IBM tape device driver to get support for newer hardware. You can download the latest version of the IBM tape device driver at:

<ftp://ftp.software.ibm.com/storage/devdrv/AIX>

Refer to the *Installation and User's Guide* for more information about the update and the usage of the IBM tape device driver, which can be found at:

[ftp://ftp.software.ibm.com/storage/devdrv/Doc/IBM\\_Tape\\_Driver\\_IUG.pdf](ftp://ftp.software.ibm.com/storage/devdrv/Doc/IBM_Tape_Driver_IUG.pdf)

To install the recommended device drivers, follow the appropriate installation instructions of the device drivers and the latest readme files. Read carefully about prerequisites, restrictions, or other important hints, such as how to configure the devices after installation. Check if your operating system has all the prerequisites. If not, update the operating system.

After device driver installation, you can configure the devices within AIX. This task is described in the installation guides for the particular device drivers and devices. The device driver software package might include tools to test the devices. Use these tools carefully. We suggest that you do this only in certain situations, such as problem analysis or troubleshooting.

Determine the device names of the tape devices. This is necessary to configure the devices within applications, such as the test tools or SSAM. Make sure that the operating system on both DR550 SSAM Servers (in dual-node configurations) identifies the devices with the same device name for each of the devices. This is important in order to find the correct device in an HACMP configuration.

## SSAM tape devices

Before you can use tape devices (libraries and drives) with IBM System Storage Archive Manager, you must do the following:

1. Make sure that the devices are available within AIX.
2. Define the following within SSAM:
  - Library for the drives
  - Path from server to library
  - Drives
  - Path from server to drives
  - Device class
  - Storage pool associated with the device class
3. Within SSAM, include the storage pool in your storage hierarchy.

Step 1 is done within AIX; this ensures that the devices are available. If you have no available devices within AIX, you cannot configure SSAM to use these devices, because the devices must be available during configuration. To perform the other two steps, you can use the administrative client command line or the Administration Center.

**Tip:** Refer to the following Web site to get detailed information about how to configure your supported IBM or third-party tape drive or library with SSAM (Tivoli Storage Manager V5.5.0):

[http://www.ibm.com/software/sysmgmt/products/support/IBM\\_TSM\\_Supported\\_Devices\\_for\\_AIXHPSUNWIN.html](http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_AIXHPSUNWIN.html)

## 11.3 IBM tape libraries and drives: examples

This section shows how to attach a tape library with two tape drives to the DR550.

This example is based on attaching two different IBM Automated Tape Libraries, one with the IBM System Storage Enterprise Tape Drives TS1120 (IBM type 3592-E05), and the other with the IBM LTO Ultrium 3 WORM-Capable Tape Drives 3588 Model F3A. We selected these devices because they support WORM functionality and, in the case of the TS1120, redundant Fibre Channel attachment and hardware encryption. However, the technical aspects of this illustration remain the same for many other devices. Instead of the IBM Automated Tape Libraries, you can use simpler options, such as LTO libraries or stand-alone tape drives.

**Note:** The IBM Ultrium 3 LTO drives supports WORM functionality. However, it might be necessary to update the drive, library microcode, and the driver as well as the SSAM server software to level 5.3.1.2 or higher.

### 11.3.1 IBM System Storage TS1120 Tape Drive

The IBM System Storage TS1120 Tape Drive offers a solution to address applications that need high capacity, fast access to data, or long-term data retention. It is supported in IBM tape libraries or frames that support stand-alone installations. It is designed to help reduce the complexity and cost of the tape infrastructure.

The TS1120 Tape Drive uses the existing 3592 media, which is available in rewritable or Write Once Read Many (WORM) media to store 100 GBs, 500 GBs, or 700 GBs, depending on the cartridge type. The 3592 JA/JW media helps reduce resources to lower total costs, while the 3592 JJ/JR media is designed to support applications that require rapid access to data.

The TS1120 tape drive supports a native data transfer rate of up to 100 MBps. In open system environments, where data typically compresses at 2:1, the TS1120 tape drive can transfer data up to 200 MBps. This can help reduce backup and recovery times or require fewer resources to support the environment.



Figure 11-2 IBM System Storage TS1120 Tape Drive

Some features of the IBM System Storage TS1120 Tape Drive are:

- ▶ Supports IBM Servers and selected open system platforms
- ▶ Supported on existing IBM automation like TS3500 or 3494
- ▶ Offers native data transfer rate of up to 100 MBps
- ▶ Supports up to 1.5 TBs on a standard length cartridge
- ▶ Supports fast access with short length cartridge
- ▶ Supports drive encryption (only drives with the ENC sticker at the rear of the drive)

The TS1120 comes with dual-ported switched fabric 4 Gbps Fibre Channel interfaces for attachment to multiple servers or a single server with redundancy.

### 11.3.2 IBM System Storage Enterprise 3592 WORM Tape Cartridges

The 3592 WORM Tape Cartridge is designed to prevent the alteration or deletion of stored data, while allowing data to be appended to existing cartridges. This is achieved by advanced security features in the cartridge and on the media. Figure 11-3 shows the IBM System Storage Enterprise 3592 WORM Tape Cartridges.



Figure 11-3 IBM System Storage Enterprise 3592 WORM Tape Cartridges

Some features of the 3592 WORM Tape Cartridges are:

- ▶ Advanced fourth generation metal particle formulation in a dual-layer coating on a half-inch wide tape.
- ▶ A precision timing-based server helps enable high-track densities, high data rates, data access performance, high reliability, and stop-start performance.
- ▶ Modified cartridge design and construction help improve pin retention, hub and clutch engagement, spool alignment, and tape stacking within the cartridge for improved reliability and durability.
- ▶ A cartridge memory chip that stores access history and media performance information. These records are used by the IBM Statistical Analysis and Reporting System (SARS) to help diagnose and isolate tape errors.

### 11.3.3 IBM System Storage Enterprise Automated Tape Library (3494)

We refer to the IBM System Storage Enterprise Automated Tape Library as the *3494 Tape Library*, *3494*, or *IBM 3494 Tape Library*.

The IBM 3494 Tape Library provides a modular tape automation solution for multiple computing environments. Some of the features of the 3494 include:

- ▶ The 3494 has a data storage capacity of up to 124 terabytes (TBs) of noncompacted data and up to 374 TBs of compacted data.
- ▶ The 3494 supports the IBM System Storage Enterprise Tape Drive 3592 Model J1A or the TS1120.
- ▶ The 3494 provides a high availability model, which contains dual library managers and dual accessors for reduced service interventions and greater availability. Also, an optional feature enables two active accessors for increased performance.
- ▶ The 3494 can mount up to 610 cartridges per hour with dual active accessors.
- ▶ The 3494 provides an easy-to-use graphical user interface (GUI) for operational control and setup of the library. It also provides a Web-based Specialist for viewing and administering the library from a remote location.



The L22 Frame is the minimum configuration and the base of a 3494, which should include 3592 Tape Drives. It contains a tape subsystem, library manager, cartridge accessor, convenience input/output (I/O) station, accessor rail, and cartridge storage cells.

To augment the capacity of the 3494, you can add one or more IBM System Storage Enterprise Tape Drive Expansion Frames D14, D22, or D24.

If you already use a 3494 Tape Library with different drive types, you can share this 3494 Tape Library with the DR550. Figure 11-4 shows the Enterprise Tape Library with one expansion frame.



Figure 11-4 IBM System Storage Enterprise Tape Library with one expansion frame

### 3494 host attachment

The 3494 provides multiple host connectivity options. Here, we describe how the 3494 attaches to the DR550. (See Figure 11-5 on page 458.)

The 3494 Tape Library depicted contains two 3592 (or TS1120) Tape Drives. The SAN zoning in this case is based on a port zoning on the SAN switch. To separate disk and tape traffic, be sure to select the ports reserved for tape on the SAN Switch (ports 4, 5, 6, and 7).

The control path to the 3494 Tape Library is established through a LAN connection. Therefore, the tape library is connected to the same network as the both DR550-DR2. In this case, both DR550 SSAM Servers will have access to the tape library after this network connection is configured. This is important to guarantee access to the library in case of an HACMP failover.

To allow both DR550 SSAM Servers (hosts) access to the library, you must configure hosts at the library manager graphical user interface (Add LAN Host to Library window). This process is described in *IBM System Storage Enterprise Automated Tape Library 3494 Operator Guide*, GA32-0449.

Within the library manager, go to the Add LAN Host to Library window. This enables you to configure a LAN-attached host for communication with the 3494. You can configure up to 32 LAN host ports, for example, when you want to share the tape library with other applications.

To configure the tape library for use with the DR550 SSAM Server, you must, at a minimum, add a host (whose I/O address is the IP address of your HACMP cluster service). Use the host name of the HACMP cluster service as it is specified in the /etc/hosts file in the DR550 SSAM Server. We also recommend that you add the IP boot addresses of the p5 servers. This will enable you to access the tape library without HACMP running, that is, without the cluster started. This is useful for problem management and troubleshooting.

Figure 11-5 shows the physical connection between the 3494 Tape Library, its drives, and the DR550 SSAM Server. The HBA of the DR550 SSAM Server is solely for the tape device attachment. The HBAs connected to the disk storage subsystem are not shown. The tape drives are connected directly to the SAN switch (using ports 5 and 6 on each switch). The tape library has no Fibre Channel connection; the only connection is through the LAN-attached library manager.

With this configuration, you have the ability to use all Tivoli Storage Manager features (database backup, backup storage pool, and migration), and you can use the tape devices after any HACMP takeover without manual intervention.

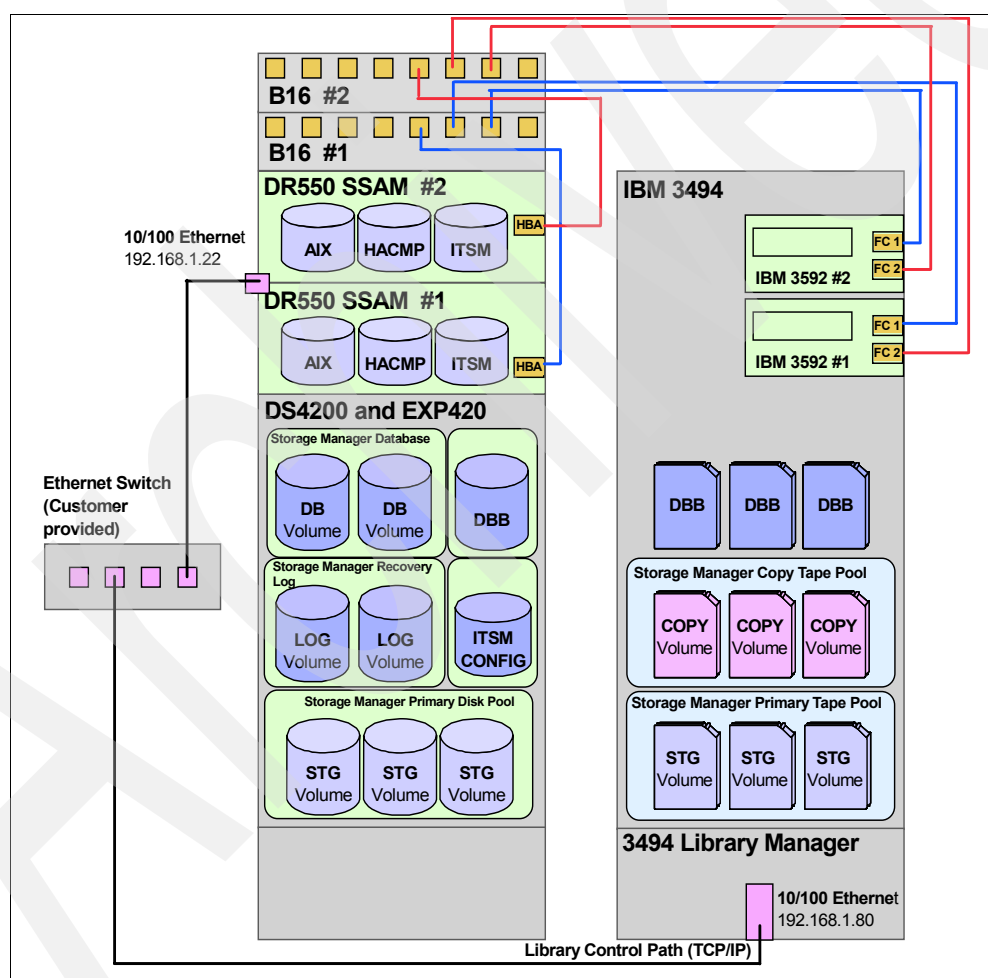


Figure 11-5 IBM 3494 attachment to the DR550-DR2 dual node



### 11.3.4 IBM System Storage 3588 Model F3A Ultrium 3 WORM Tape Drive

For the remainder of this chapter, the IBM System Storage Enterprise Tape Drive 3588 will be referred to as the *3588 Tape Drive* or *IBM 3588*. Figure 11-6 shows a picture of the IBM 3588 Tape Drive.




Figure 11-6 IBM System Storage Enterprise 3588 Tape Drive

The IBM 3588 Tape Drive Model F3A is an IBM LTO Ultrium 3 Tape Drive that combines tape reliability and performance at open systems prices.

Some features of the 3588 tape drive are:

- ▶ Up to 80 MBps native transfer rate.
- ▶ 400 GB native capacity, up to 800 GB with onboard compression (2:1).
- ▶ Includes a 2 Gbps Fibre Channel interface attachment.
- ▶ Mounts in a 3584 Tape Library Model L52, L32, D52, or D32.
- ▶ Sixteen-channel read-while write speeds up the operation of the drive and ensures data integrity.
- ▶ Support for most major operating systems.
- ▶ Supports Ultrium 3, 2, and 1 media (Ultrium1 = read-only).
- ▶ Support for Ultrium 3 400 GB WORM media and functionality.

**Note:** The use of 400 GB WORM media and functionality is now supported and available. For currently installed 358x products, with IBM System Storage LTO Ultrium 3 Tape Drives, a drive microcode update and library firmware update might be required. However, no hardware additions or features are necessary to enable LTO 3 WORM capability.

See the following Web site for information and the required microcode:

<http://www.ibm.com/servers/storage/support/allproducts/downloading.html>

### 11.3.5 IBM System Storage 3589 Ultrium 3 WORM Tape Cartridge

The IBM System Storage 3589 Tape Cartridge Models include a Ultrium 3 400 GB WORM Data Cartridge with the following features:

- ▶ Cartridge memory, built into every data cartridge, which helps enhance functionality and media reliability by storing access history and media performance information for use by the tape drive every time the cartridge is accessed.
- ▶ Half-inch particle tape with 400 GB native capacity in a single compact cartridge.
- ▶ Pre-labeling, with the ability to specify a starting volume serial and color-coding, which helps save time and labor when adding new cartridges to your existing inventory.

These new Ultrium 3 Cartridges were designed for applications, such as archiving and retention, as well as those applications requiring an audit trail. Figure 11-7 shows the IBM System Storage 3589 WORM Tape Cartridges.



Figure 11-7 IBM System Storage 3589 WORM Tape Cartridges

IBM has taken several steps to reduce the possibility of tampering with the information:

- ▶ The bottom of the cartridge is molded in a different color than rewritable cartridges.
- ▶ The special cartridge memory helps protect the WORM nature of the media.
- ▶ A unique format is factory-written on each WORM cartridge.

### 11.3.6 IBM System Storage TS3500 Tape Library

We refer to the IBM System Storage TS3500 Tape Library as the *TS3500 Tape Library*, *3584 Tape Library*, *3584*, *TS3500*, or *IBM 3584 Tape Library*.

The IBM System Storage TS3500 Tape Library is a highly scalable, automated tape library for midrange to enterprise open systems environments, combining IBM automation and drive technology. The IBM TS3500 is designed to handle backup, archive, and disaster recovery data storage functions with ease:

- ▶ The TS3500 has a data storage capacity of up to 2752 terabytes (TBs) of noncompacted data and up to 5504 TBs of compacted data (2:1 compression).
- ▶ The TS3500 supports the IBM System Storage Enterprise Tape Drive TS1120 and IBM LTO Ultrium 4, 3, or 2, even in a mixed configuration.
- ▶ Supports TS1120 and TS1040 tape drive encryption for data protection.
- ▶ The TS3500 can attach up to 15 expansion frames with up to 12 drives per frame (192 total per library) and up to a total maximum number of storage slots of 6887.
- ▶ The patented multipath architecture is designed to help increase configuration flexibility with logical library partitioning, while enabling system redundancy to help increase availability.

- ▶ The new IBM Advanced Library Management System (ALMS) is available on the TS3500 Tape Library, incorporating dynamic storage management, which is designed to enable the user to dynamically create and change logical libraries and configure any drive into any logical library.
- ▶ Different frame types for TS1120 or 3592 and LTO tape drive technology:
  - L23: Base frame for TS1120 or 3592. D23: expansion frame for TS1120 or 3592
  - L53: Base frame for LTO. D53: Expansion frame.
- ▶ HA1: Service bay frame for High Availability.



*Figure 11-8 IBM System Storage TS3500 Tape Library base frame*

### **TS3500 host attachment**

The TS3500 has a different library concept than the 3494. Although the 3494 Library requires a dedicated connection (LAN or RS232; also referred to as outboard library management) for library control commands, such as moving a cartridge from a storage slot to a tape drive, the TS3500 does not. The TS3500 is a SCSI medium changer library; the SCSI protocol specification defines medium changer commands that can be issued through the same physical connection as the tape commands. To control the TS3500, a single Fibre Channel connection from the host to a tape drive is required. When transporting the SCSI protocol commands, in our case, over the Fibre Channel connection, the drive is addressed from a host through LUN 0, for example, for the first drive within a library SCSI ID 0, LUN 0. To issue a library command, the host uses address SCSI ID 0, LUN 1. The command is transported through the same physical connection from the HBA to the tape drive, which in turn passes the request through an internal RS232 connection to the library controller component. Responses from the controller follow the same path through the drive back to the HBA and the issuing application.

The patented IBM multipath architecture exploits this approach. Because any of the drives within a library can become a control path drive with LUN 1 enabled, this provides several features:

► Logical partitioning

Having more than one control path enables the physical library components to be split up in logical partitions. Each logical partition, containing one or more drives, several storage slots, and cartridges, can be assigned to a separate host system and application, sharing the library robotics. Figure 11-9 illustrates the multipath architecture and the logical partitioning.

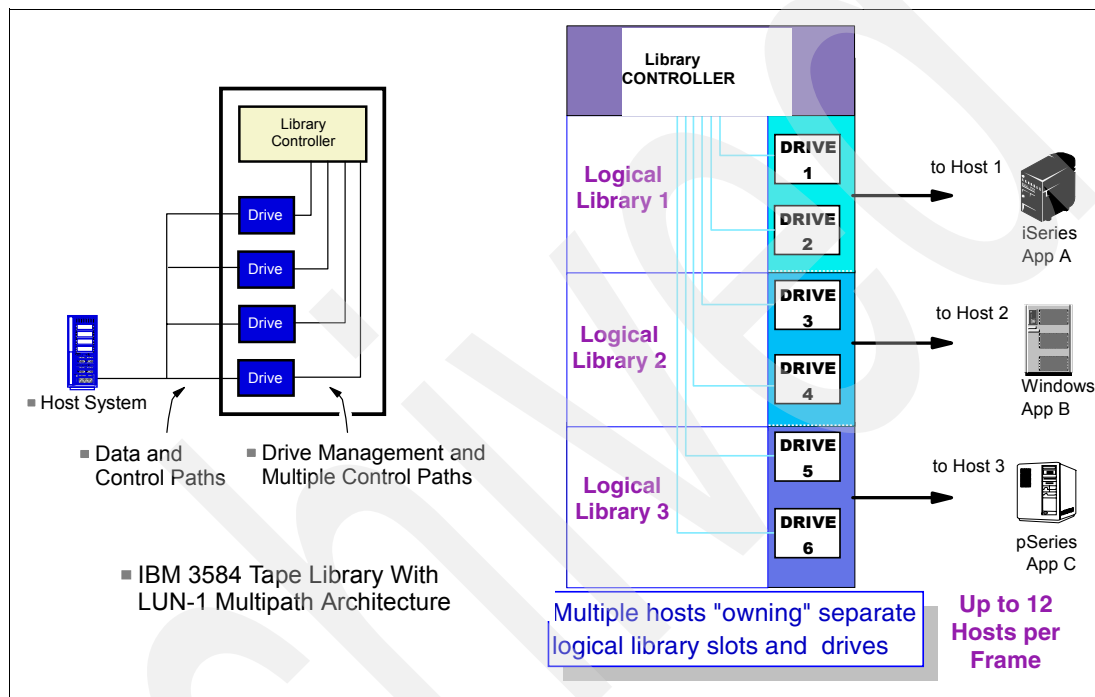


Figure 11-9 IBM TS3500 multipath architecture and logical partitioning

► Redundancy

A second control path defined from the same host and probably through a separate fabric eliminates the risk of an outage. The Atape driver for AIX provides automatic and transparent control path failover functionality, so SSAM could continue to operate the library even if a failure of one component in the primary control path (for example, the HBA, switch, cable, SFP, or tape drive) occurs. The TS3500 requires a chargeable license code to support this feature (see Figure 11-10 on page 463).

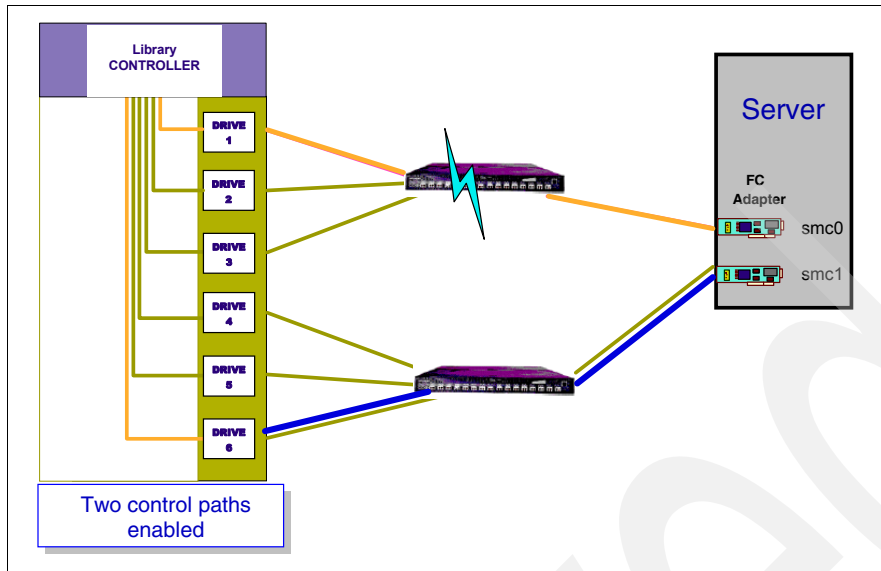


Figure 11-10 IBM TS3500 control path failover

For further information about the TS3500 multipath architecture, see *IBM System Storage Tape Library Guide for Open Systems*, SG24-5946, available at:

<http://www.redbooks.ibm.com/abstracts/sg245946.html>

Figure 11-11 on page 464 illustrates the required Fibre Channel for tape attachment from the DR550-DR2 dual node to the TS3500 Library equipped with two 3592-J1A or TS1120 Tape Drives. No dedicated control path cabling is necessary, because library control commands are sent through the Fibre Channel connections through the drives to the controller.

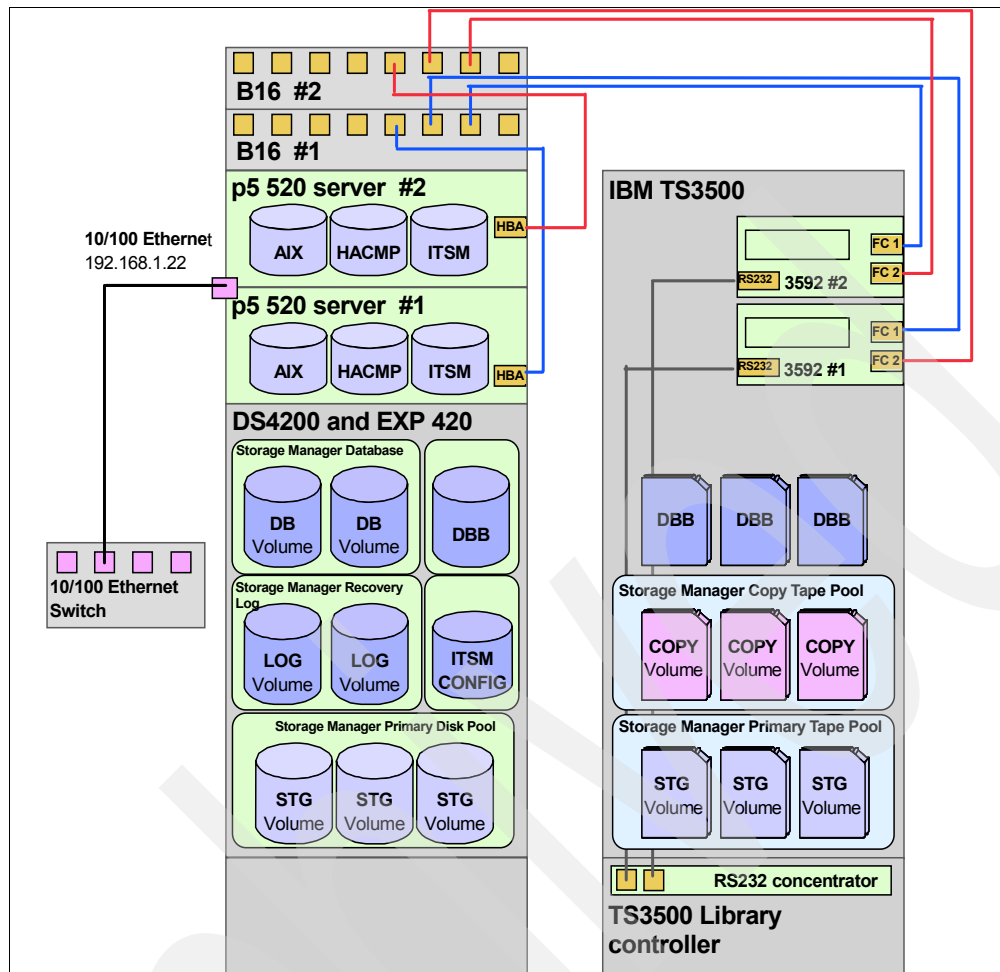


Figure 11-11 IBM TS3500 attachment to the DR550-DR2 dual node

### 11.3.7 Device driver verification for 3592 and 3494

Installation of device drivers for the tape library and the tape drives on the host that will use these devices is required. *Host*, in this context, is any of the DR550 SSAM Servers within the DR550. Although SSAM provides its own drivers for some devices, for the 3494 Library and the 3592 Tape Drives, the IBM device drivers shipped with the devices have to be used:

- ▶ The device driver package required for AIX for the 3494 Tape Library is called *atlld*. The package includes the necessary software for the operating system to communicate with the library manager and a utility program (*mtlib*), which provides a command-line interface to the library. This utility cannot only be used to test library communications and verify other functionality, but also can be scripted, for example, to automate monitoring tasks.
- ▶ The device driver *Atape* is necessary under AIX to operate the 3592-J1A Tape Drives or the TS1120. In addition to the driver files, the package contains a utility program called *tapeutil*. Again, this tool can be used to verify proper communication, to download firmware to drives, or to perform basic tape commands, such as load and unload, locate, or rewind. The *tapeutil* utility can be started from the AIX command prompt in an interactive, menu-driven mode, or it can be called with parameters in CLI mode for scripting purposes.

The device driver software is preinstalled on the DR550 SSAM Server in manufacturing. However, if newer versions of the software become available with fixes or support for new functionality, you might want to perform an update. We recommend that you first uninstall the current version and then install the new release. For the procedures to install and uninstall the software packages, follow the device driver installation guide. The current version of the installation guide, *IBM Tape Device Drivers Installation and User's Guide*, GC27-2130, is available at:

[ftp://ftp.software.ibm.com/storage/devdrv/Doc/IBM\\_Tape\\_Driver\\_IUG.pdf](ftp://ftp.software.ibm.com/storage/devdrv/Doc/IBM_Tape_Driver_IUG.pdf)

This FTP site also provides the latest releases of the driver packages available for AIX:

<ftp://ftp.software.ibm.com/storage/devdrv/AIX/>

To verify the 3494 Tape Library driver installation, log on to the DR550 SSAM Server and issue the command `lslpp -l atldd.driver`. The output should be similar to that shown in Example 11-1.

*Example 11-1 AIX command lslpp -l atldd.driver*

Fileset	Level	State	Description
-----			
Path: /usr/lib/objrepos atldd.driver	6.5.4.0	COMMITTED	IBM Automated Tape Library Device Driver

**Important:** The device driver package `atldd` has support for 3592 Tape Drives starting with Version 5.4.8.0 or later. Make sure that you use one of these versions.

For the Atape driver, type `lslpp -l Atape.driver` and verify that the result is similar to that shown in Example 11-2.

*Example 11-2 AIX command lslpp -l Atape.driver*

Fileset	Level	State	Description
-----			
Path: /usr/lib/objrepos Atape.driver	10.3.9.0	COMMITTED	IBM AIX Enhanced Tape and Medium Changer Device Driver

**Important:** To support 3592 Tape Drives, the Atape driver must be at Version 8.3.1.0 or later.

Although the device driver packages are preinstalled in the DR550 SSAM Server, a few customization tasks have to be performed in order to operate the 3494 Tape Library in a customer environment.

The following steps are required on each engine of the DR550 to establish communication between engine or engines and the 3494 Tape Library. It is assumed that the Fibre Channel cables from the HBA or HBAs to the DR550 SAN Switch and the cables from the tape drives to the switch are connected.

Complete the following steps:

1. Prepare the library daemon. The `atldd` driver, after installation, provides a daemon, known to AIX specifically as the IBM library manager control point daemon (*lmcpd*). For this daemon to operate appropriately, tape libraries have to be defined in a configuration file. This file, `ibmatl.conf`, is placed in the `/etc` directory when the software package is installed.

The *IBM Tape Device Drivers Installation and User's Guide*, GC27-2130, also describes how to customize this file. The file has to be edited to contain a statement of the form *libraryname IP-address* to reflect the IP address of the 3494 library manager in your environment.

Example 11-3 shows a typical example of the configuration file.

Example 11-3 Custom /etc/ibmatl.conf file

---

```
This is the file which defines the 3494 libraries and how they are attached.
The format of this file is:
#
Library name address symbolic name
3494lib 100.100.1.80 atl3494
```

---

2. Create the 3494 library manager control point (*lmcp*). This procedure is also described in *IBM Tape Device Drivers Installation and User's Guide*, GC27-2130.

To create the library manager control point, start SMIT and select Devices → Tape Drive → Add a Tape Drive and press Enter. Scroll down to almost the bottom of the list to highlight the line *atl library LAN/TTY Library Management Control Point* and press Enter. On the *Add an LMCP Logical Device* window, press F4 for the Library Name; the list of library names is derived from your configuration file */etc/ibmatl.conf*. (See Figure 11-12.) The remaining fields can be left blank; the Logical Name of LMCP will be automatically assigned (typically *lmcp0* for the first control point in the system) after pressing Enter.

Add an LMCP Logical Device

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

	[Entry Fields]
Logical Name of LMCP (optional)	[]
* Library Name (F4 to list library names)	[3494lib]
Command Timeout in Minutes	[]

Figure 11-12 SMIT Add an LMCP Logical Device

Check that the result is OK and exit SMIT by pressing F10.

3. Verify that the FC HBA is present and has the status available in AIX with the **lsdev -Cc adapter | grep fcs2** command. (Note that *fcs0* and *fcs1* are the resources for disk attachment.)
4. Configure the tape devices within AIX by entering the **cfgmgr** command. The command should complete without notification within a few seconds.

**Tip:** If **cfgmgr** does not find the connected Fibre Channel tape devices, issue the following command to re-establish the link between the DR550 SSAM Server Host Bus Adapter and the DR550 SAN switch: run **rmdev -d1 fcs2 -R**, then run **cfgmgr** again. Make sure that no other devices are connected to *fcs2* when you run this command.

5. Verify the hardware installation with the AIX command **lsdev -Cc tape**.

The result of this AIX command looks similar to the output shown in Example 11-4 on page 467.



*Example 11-4 AIX command lsdev -Cc tape (result)*

---

```
lsdev -Cc tape
lmcp0 Available LAN/TTY Library Management Control Point
rmt0 Available OC-08-01 IBM 3592 Tape Drive (FCP)
rmt1 Available OC-08-01 IBM 3592 Tape Drive (FCP)
```

---

6. Verify that the library manager control point daemon is running by using the AIX command **ps -ef | grep lmcpd**. The output should look similar to that shown in Example 11-5.

*Example 11-5 AIX command ps -ef | grep lmcpd (result)*

---

```
ps -ef | grep lmcpd
root 225424 1 0 09:13:48 - 0:00 /etc/lmcpd
root 245970 225424 0 09:13:48 - 0:00 /etc/lmcpd
```

---

7. Test the communication to the 3494 library manager using the **mtlib** utility. Type **mtlib -l /dev/lmcp0 -qL**, press Enter, and an output similar to that shown in Example 11-6 should be displayed (note that the output has been truncated).

*Example 11-6 AIX command mtlib -l /dev/lmcp0 -qL (result)*

---

```
Library Data:
operational state.....Automated Operational State
functional state.....00
input stations.....1
output stations.....1
input/output status.....ALL input stations empty
 ALL output stations empty
 Bulk input/output allowed

machine type.....3494
sequence number.....61378
number of cells.....885
available cells.....762
subsystems.....1
convenience capacity.....10
accessor config.....01
accessor status.....Accessor available
 Gripper 1 available
 Gripper 2 not installed
 Vision system operational
```

---

8. Verify the communication to one or all of the installed tape drives with the **tapeutil** utility. Type, for example, **tapeutil -f /dev/rmt0 inquiry**, and press Enter. The response of the drive should look similar to the result in Example 11-7, which shows that it is an IBM 3592 Model J1A device and its serial number, here 78-01363.

*Example 11-7 AIX command tapeutil -f /dev/rmt0 inquiry (result)*

---

```
tapeutil -f /dev/rmt0 inquiry
Issuing inquiry...
```

Inquiry Data, Length 56

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F	
0000	-	0180	0302	3300	1000	4942	4D20	2020	2020								0123456789ABCDEF
0010	-	3033	3539	324A	3141	2020	2020	2020	2020								[...3...IBM ]
0020	-	3035	3039	3638	3030	3030	3037	3830	3133								[0509680000078013]
0030	-	3633	2030	0900	4081												[63 0..@ ]

---

For further details about the capabilities of the mtlb and tapeutil utilities, refer to the *IBM Tape Device Drivers Installation and User's Guide*, GC27-2130.

Installation and verification of the device driver software for proper communication within AIX with the 3494 library manager and the 3592 Tape Drives is now complete.

### 11.3.8 Device driver verification for 3588 and TS3500

Using the same physical connection for the library medium changer and tape drive commands eliminates the need to have a separate device driver for the tape library installed. Instead, only Atape is required to operate a TS3500 Library with 3588 Tape Drives. Atape is preinstalled in the DR550 SSAM Server, so the following steps demonstrate how to verify the correct communication with the TS3500 library controller and the 3588 Tape Drives, after the physical Fibre Channel connectivity of the HBA and the tape devices has been established.

Complete the following steps:

1. Verify the proper installation of Atape using the `lslpp -l Atape.driver` command at the AIX prompt. Example 11-2 on page 465 shows what the output should look like.
2. Verify that the FC HBA is present and has the status available in AIX with the `lsdev -Cc adapter | grep fcs2` command. (Note that fcs0 and fcs1 are the resources for disk attachment.)
3. Configure the tape devices within AIX by entering the `cfgmgr` command. The command should complete without notification within a few seconds.
4. Verify the configuration by typing the `lsdev -Cc tape` command. The result should be the same as the result shown in Example 11-8.

*Example 11-8 AIX command lsdev -Cc tape (result)*

---

```
lsdev -Cc tape
rmt0 Available 0C-08-01 LTO Ultrium Tape Drive (FCP)
rmt1 Available 0C-08-01 LTO Ultrium Tape Drive (FCP)
smc0 Available 0C-08-01 IBM 3584 Library Medium Changer (FCP)
```

---

5. The test for proper communication with the library controller is performed with the tapeutil utility, as for the 3592 Tape Drives; Atape is the common driver for both. Typing the `tapeutil -f /dev/smc0 inquiry` command returns the TS3500 Library (3584-L32) vital product data, as displayed in Example 11-9.

*Example 11-9 AIX command tapeutil -f /dev/smc0 (result)*

---

```
tapeutil -f /dev/smc0 inquiry
Issuing inquiry...

Inquiry Data, Length 58

 0 1 2 3 4 5 6 7 8 9 A B C D E F 0123456789ABCDEF
0000 - 0880 0302 3500 2002 4942 4D20 2020 2020 [...5. .IBM]
0010 - 3033 3538 344C 3232 2020 2020 2020 2020 [03584L32]
0020 - 3430 3930 3133 3030 3030 3031 3334 3030 [5040130000013400]
0030 - 3637 2030 0000 0000 0000 [52 0.....]
```

---

The command in the previous step has been passed through drive rmt0 to the library controller and back. Therefore, the verification of communication to the drive is basically needless, but can, of course, be performed for completeness with the resource name

/dev/rmt0. We recommend that you check the remaining rmt resources with the **inquiry** command. Refer to Example 11-7 on page 467 for the proper results.

## 11.4 Integrating devices into System Storage Archive Manager

The library and tape devices are now configured, verified, and can be defined to the SSAM Server. First, we go through the steps required to configure SSAM to use the 3592 Tape Drives and the 3494 library manager. Second, we illustrate the configuration of the 3588 Tape Drives and the companion TS3500 library controller. For the first procedure, we use the administrative command line; the second procedure is shown using the Administration Center Web interface.

Furthermore, in the 3494 example, we show the creation of a backup storage pool, which can be used in the SSAM Server to copy objects of the primary disk storage pool ARCHIVEPOOL. The TS3500 example goes further and describes three scenarios that benefit from the tape attachment:

- ▶ One scenario with a sequential storage pool that is used to migrate data off the primary ARCHIVEPOOL based on the occupancy level of this pool, literally extending the primary's pool capacity
- ▶ A second scenario with a tape copy pool to incorporate periodic incremental backups to provide disaster protection
- ▶ A third scenario to generate additional SSAM database backups on tape, also with the intent to have a resource from which to restore in case of a disaster

### 11.4.1 Using the IBM SSAM Administration Center

Beginning with IBM System Storage Archive Manager Version 5.3, the administrative Web interface has been replaced with the Administration Center. The Administration Center is also a Web-based interface that can be used to centrally configure and manage SSAM Version 5.3 or later servers. The former (Version 5.2 and prior) administrative Web interface is no longer supported.

The Administration Center and the Integrated Solutions Console (ISC) are preinstalled on the DR500 Models DR1 and DR2. The software for the ISC is also downloadable at <http://www.ibm.com/developerworks/autonomic/csa1.html> for various operating systems. For detailed instructions about installation and prerequisites for your operating system, refer to the *IBM Tivoli Storage Manager for AIX Installation Guide*, SC32-0134.

Figure 11-13 on page 470 shows the Storage Devices main window of the SSAM Administration Center Web interface. For most screen captures, we left out the menus at the top and on the left side of the window because they do not change.

To connect to the Administration Center Web interface, start a Web browser and start an http session to the IP address of the workstation where the Administration Center and Integrated Solutions Console are installed, using the port number specified when installing the Integrated Solutions Console:

`http://<ip_of_management_station>:8421/ibm/console`

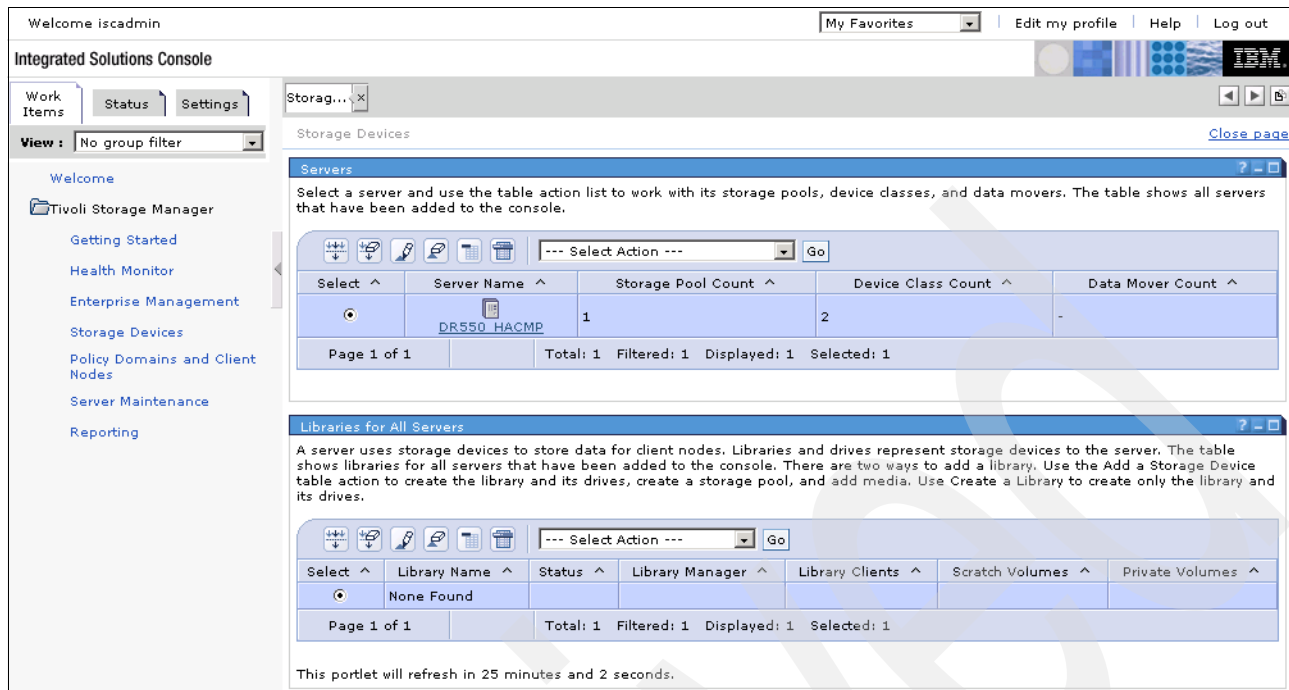


Figure 11-13 Administration Center running in the Integrated Solutions Console (ISC)

The Administration Center Web interface provides both context-sensitive and general help topics. For context-sensitive help, click the question mark “?” in the upper right-hand corner of the frame in which you are working. For more general help, select **Help** in the upper right-hand corner of the page.

## 11.4.2 Using the SSAM command line

In the examples we provide, we illustrate the use of the Administration Center Web interface as well as the command line. To use the command line, we recommend that you install the SSAM Console on your management station. The command line is also available through the Administration Center.

You may also install and use the SSAM administrative client on your management station as the command-line interface. Refer to the appropriate installation guide of the SSAM (Tivoli Storage Manager) Client for your operating system.

If you do not want to install any of the administrative clients on a management workstation, you can log on to the DR550 SSAM Server directly as user `dr550` through a local keyboard or as user `dr550adm` through a ssh connection. Invoke the SSAM command line by typing `dsmdmc` at the command line.

When using the command line, you can obtain help for any command by typing the word `help` in front of the command, for example, `help define library`. For a general overview of the command line functions and syntax, just type `help`.

## 11.4.3 Defining 3592 and 3494 devices to SSAM

In this first example, we use the command line to define 3592 and 3494 devices to SSAM:

1. Make sure that the devices are available to AIX; use the `lsdev -Cc tape` command.
2. Start a SSAM administrative command line.

3. Define a 3494 library named 3494LIB.

Use the command **define library 3494lib libtype=349x** and check for the successful result of the command, which is ANR8400I Library 3494LIB defined.

4. Define a path from the server to the library.

Use the following command:

```
define path dr550_hacmp 3494lib srctype=server desttype=library
device=/dev/lmcp0
```

where dr550\_hacmp is the name of your DR550 Tivoli Storage Manager server instance.

Check for the successful completion of the command:

ANR1720I A path from DR550\_HACMP to 3494LIB has been defined.

If you get an error message such as ANR8418E DEFINE PATH: An I/O error occurred while accessing library 3494LIB, you must check the general availability of the library or check that you have configured the correct device (/dev/lmcp0). The DEVICE parameter specifies the device special file for the LMCP.

5. Define the drives in the library.

Use the following two commands to define two drives named drive01 and drive02 that belong to 3494LIB:

```
define drive 3494lib drive01
define drive 3494lib drive02
```

Check for the successful completion of the command:

ANR8404I Drive DRIVE01 defined in library 3494LIB

If you get an error message such as ANR8444E DEFINE DRIVE: Library 3494LIB is currently unavailable, you must check the general availability of the library (not drive) or check that you have configured the correct library manager control point (/dev/lmcp0) within the appropriate configuration file.

6. Define a path from the server to each drive.

Use the following two commands:

```
define path server1 drive01 srctype=server desttype=drive library=3494lib
device=/dev/rmt0
```

```
define path server1 drive02 srctype=server desttype=drive library=3494lib
device=/dev/rmt1
```

where server1 is the name of your DR550 SSAM Server instance.

The DEVICE parameter gives the device special file name for the drive, such as /dev/rmt0. You can obtain the device special file names from AIX.

You should get a message like:

ANR1720I A path from DR550\_HACMP to 3494LIB DRIVE01 has been defined.

7. Define a device class for the drives.

Use the following command:

```
define devclass 3494_3592class library=3494lib devtype=3592 format=drive
```

This example uses format=drive (Recording Format) as the recording format, because both drives associated with the device class use the same recording format.

To define a device class for the 3592 WORM media instead of the rewritable media, define a device class including the parameter worm=yes (described later in this chapter).

8. Verify the definitions.

To verify your definitions, issue the following commands:

```
query library
query drive
query path
query devclass
```

To get more detailed output from these commands, you may append `f=d`, which is the short form of `format=detailed`.

#### 11.4.4 Integrating 3592 and 3494 into SSAM storage hierarchy

Now that the library and tape resources are defined to SSAM, we need to create storage pools.

To integrate the 3592-J1A (or TS1120) and 3494 into the SSAM storage hierarchy, complete the following steps:

1. Define a storage pool associated with the device class.

Define a primary storage pool named `3494_PRIM` and a copy storage pool named `3494_COPY`, both associated with the device class named `3494_3592CLASS`, by typing the following two commands:

```
define stgpool 3494_prim 3494_3592class pooltype=primary maxscratch=999
define stgpool 3494_copy 3494_3592class pooltype=copy maxscratch=999
```

For the `maxscratch` parameter, you can specify an integer from 0 to 100000000. By allowing the server to request scratch volumes, you avoid having to define each volume to use.

The value specified for this parameter is used to estimate the total number of volumes available in the storage pool and the corresponding estimated capacity for the storage pool. Scratch volumes are automatically deleted from the storage pool when they become empty.

2. Include the storage pools in your storage hierarchy.

To make use of the previously defined copy storage pool, for example, schedules are required to periodically initiate backups of the primary pool. Type:

```
backup stgpool 3494_prim 3494_copy
```

We have now completed the example for 3494. Next, we explain the same steps on the basis of the Administration Center Web interface as well as the creation of scripts and schedules.

#### 11.4.5 Defining 3588 and TS3500 devices to SSAM

To define the 3588 and TS3500 devices to SSAM, complete the following steps:

1. Make sure that the devices are available to AIX by issuing the `lsdev -Cc tape` command.
2. Start a SSAM Administration Center Web interface.

Start by choosing **Storage Devices** from the main menu on the left side of the Administration Center Web interface, as shown in Figure 11-14 on page 473.

3. Start the storage device wizard by selecting **Servers** → **Add a storage device**. Click **Next** to start the wizard.

Figure 11-14 shows a drop-down list with the available device types for new storage devices. In our example, we use the LTO device type because our 3588 drives use LTO compatible media, which include LTO4, LTO3, LTO2, and LTO1.

In this window, you could also define devices that are connected to other SSAM/Tivoli Storage Manager servers. These devices can be shared between SSAM/Tivoli Storage Manager servers defined in an Enterprise configuration.

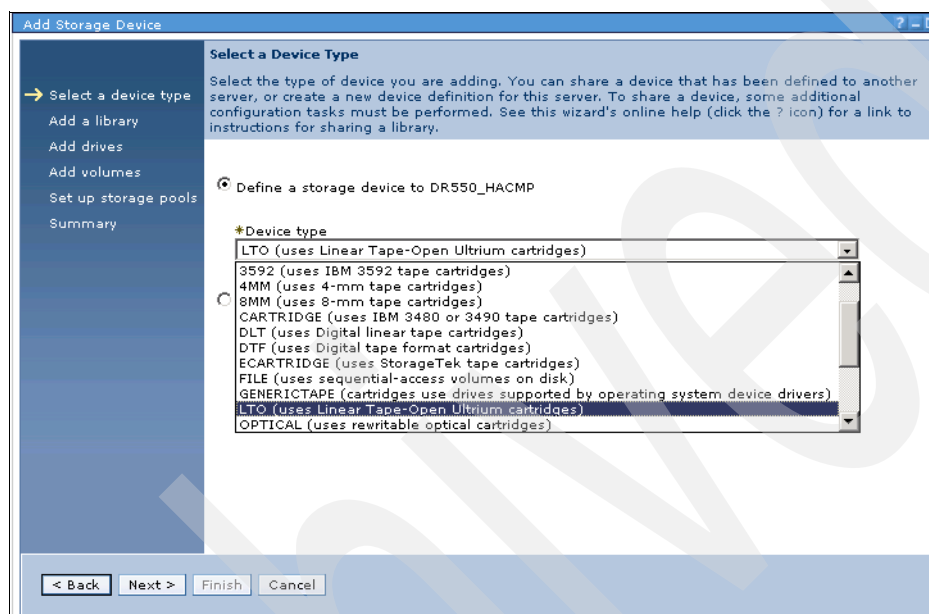


Figure 11-14 Storage device wizard: Select your device type

4. Define a TS3500 library named 3584LIB.

To define a TS3500 library named 3584LIB, type the name 3584lib into the Library name field and choose **SCSI** as the library type, as shown in Figure 11-15 on page 474.

Depending on the device type you chose from the previous menu, there are different library types available.

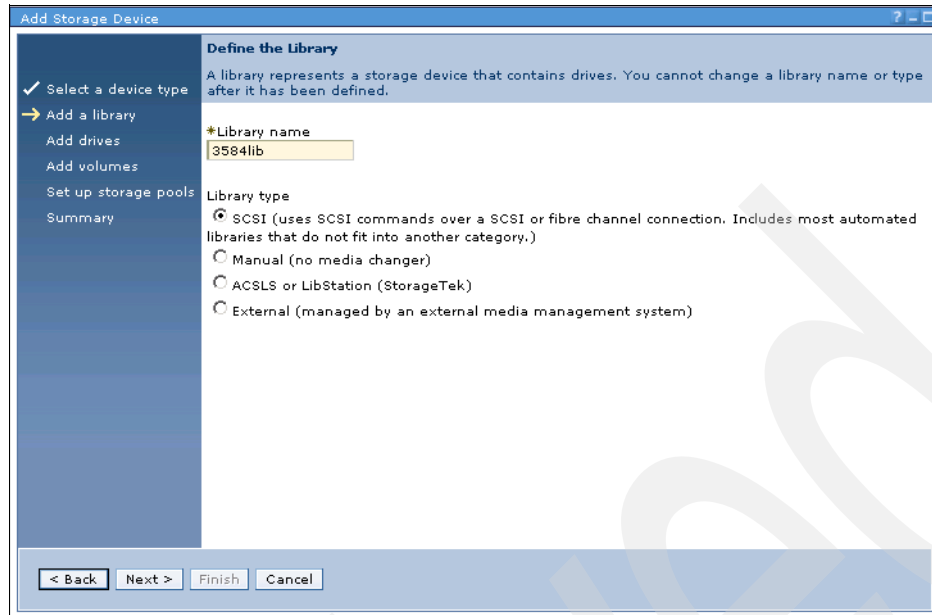


Figure 11-15 Storage device wizard: Select your library type

##### 5. Define library settings.

The *device special file name* is the device name that is used by the operating system to communicate with the library. In the case of our TS3500 library, this is /dev/smc0. Use the AIX command `lsdev -Cc tape` to find the correct device name.

Figure 11-16 also gives you the option to share the library with other SSAM/Tivoli Storage Manager servers, as mentioned in step 3.

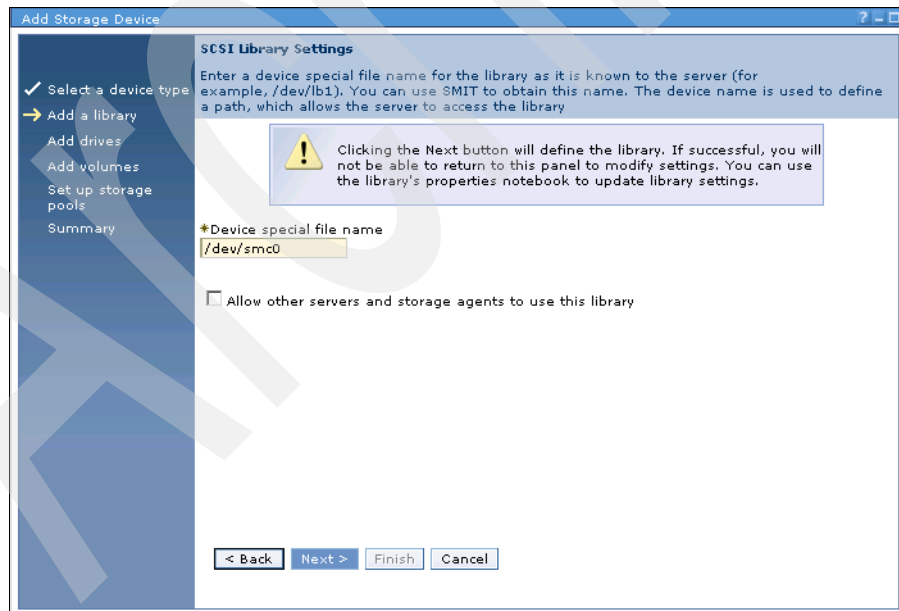


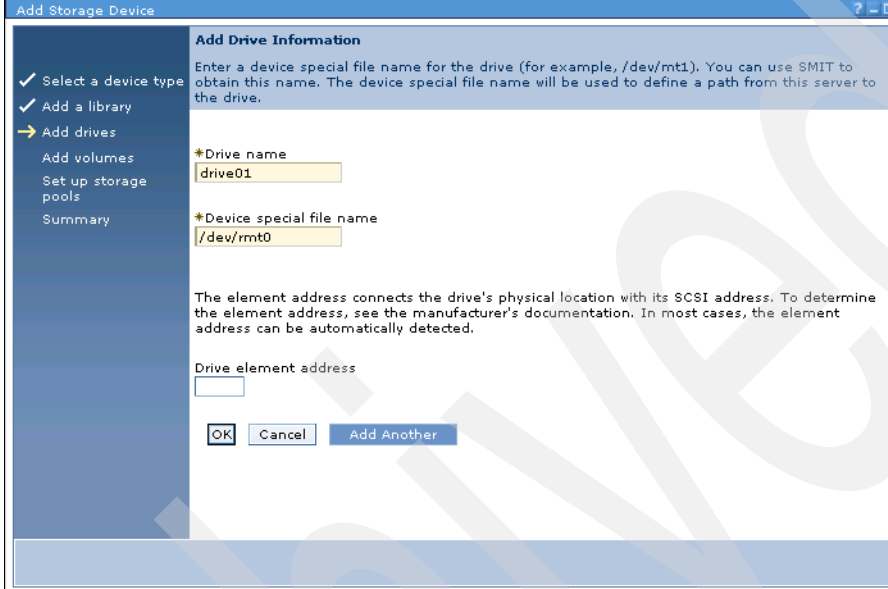
Figure 11-16 Storage device wizard: Select library settings

The next displayed window gives you an overview of the defined library and its parameters (Figure 11-17 on page 475).



6. Define the tape drives.

Select **Define Drives** → **Add Drive**, and click **GO**. Specify the drive name `drive01` and the device special file name `/dev/rmt0` for the first drive. In the case of the TS3500 library, the *drive element address* will be determined automatically by the SSAM server. Click **Add Another** to repeat this step for the second drive using `drive02` and `/dev/rmt1` as parameters. Click **OK** to define the drives. By clicking **Next**, you get an overview of the defined drives.



The screenshot shows a window titled "Add Storage Device" with a sidebar on the left and a main content area. The sidebar contains a list of steps: "Select a device type" (checked), "Add a library" (checked), "Add drives" (highlighted with a yellow arrow), "Add volumes", "Set up storage pools", and "Summary". The main content area is titled "Add Drive Information" and contains the following text: "Enter a device special file name for the drive (for example, /dev/mt1). You can use SMIT to obtain this name. The device special file name will be used to define a path from this server to the drive." Below this text are two input fields: "\*Drive name" with the value "drive01" and "\*Device special file name" with the value "/dev/rmt0". Below these fields is a paragraph of text: "The element address connects the drive's physical location with its SCSI address. To determine the element address, see the manufacturer's documentation. In most cases, the element address can be automatically detected." Below this text is an input field labeled "Drive element address". At the bottom of the main content area are three buttons: "OK", "Cancel", and "Add Another".

Figure 11-17 Storage device wizard: Define tape drives

7. Add Volumes.

This part of the wizard helps you discover and check-in the storage media, in our case, the tape volumes. We assume the use of an empty tape volume with barcode labels for this first test. You may add more volumes at a later time by invoking the Add volumes wizard again in the library properties drop-down menu (see Figure 11-18 on page 476).

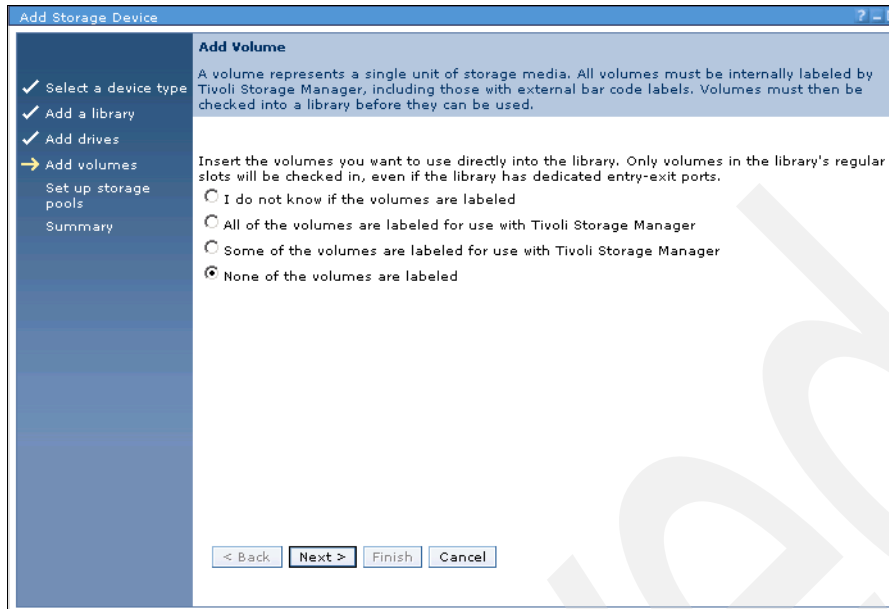


Figure 11-18 Storage device wizard: Add volumes

At this time, you should not have any volume known to the SSAM database, so you may proceed with the standard options. The wizard will start a background process that can be monitored by entering the command **query process** from a SSAM command-line interface (Figure 11-19).

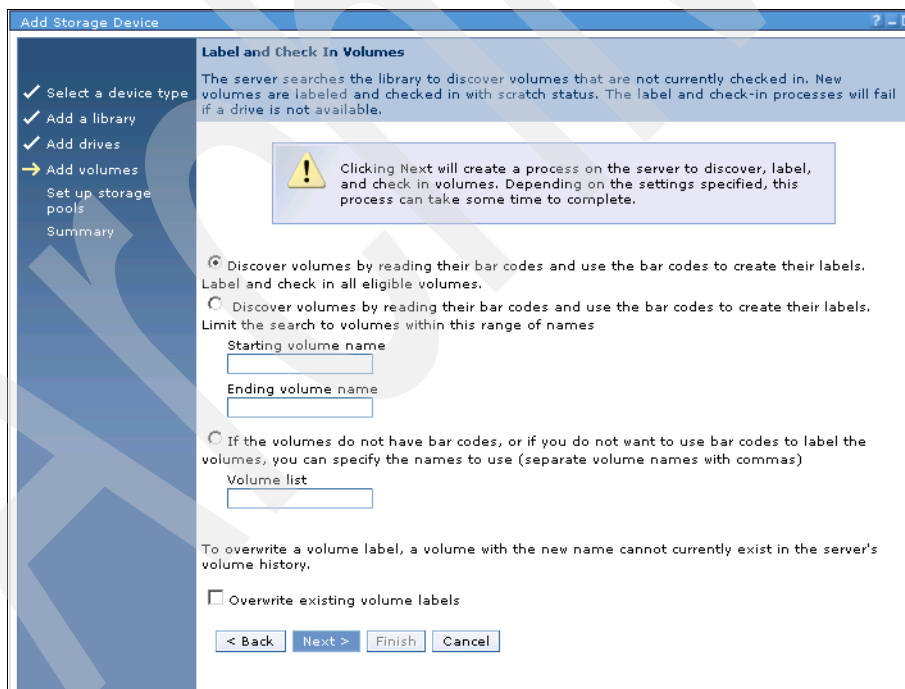


Figure 11-19 Storage device wizard: Volume options

## 8. Create storage pools.

In this step, we define a primary tape storage pool named `TAPE_POOL` with a maximum of 999 scratch volumes (Figure 11-20). The number of scratch volumes depends on your configuration. You can also define a copy pool, but we will skip this step for now because it is described later in this example.

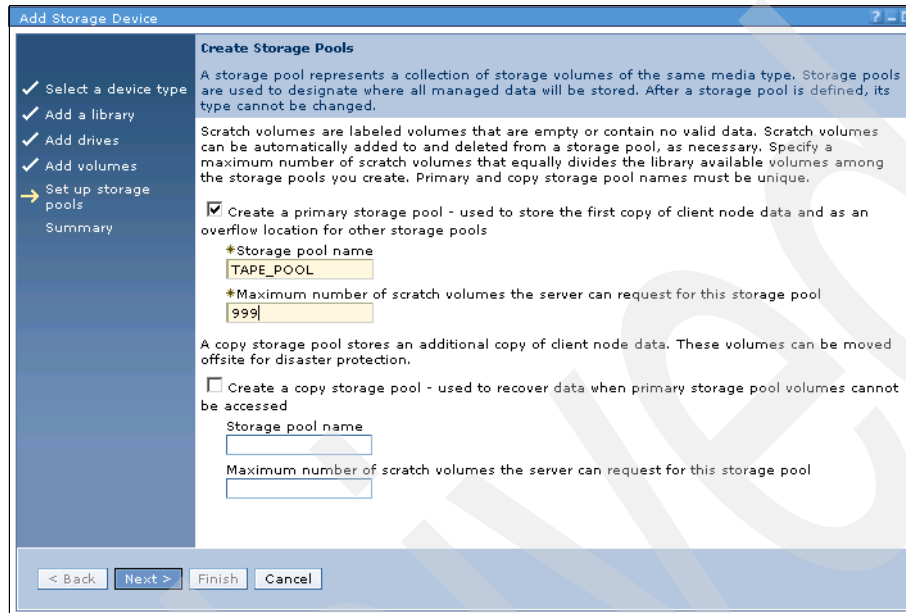


Figure 11-20 Storage device wizard: Creating storage pools

The Storage device wizard has finished the steps and shows a list of the defined devices (Figure 11-21).

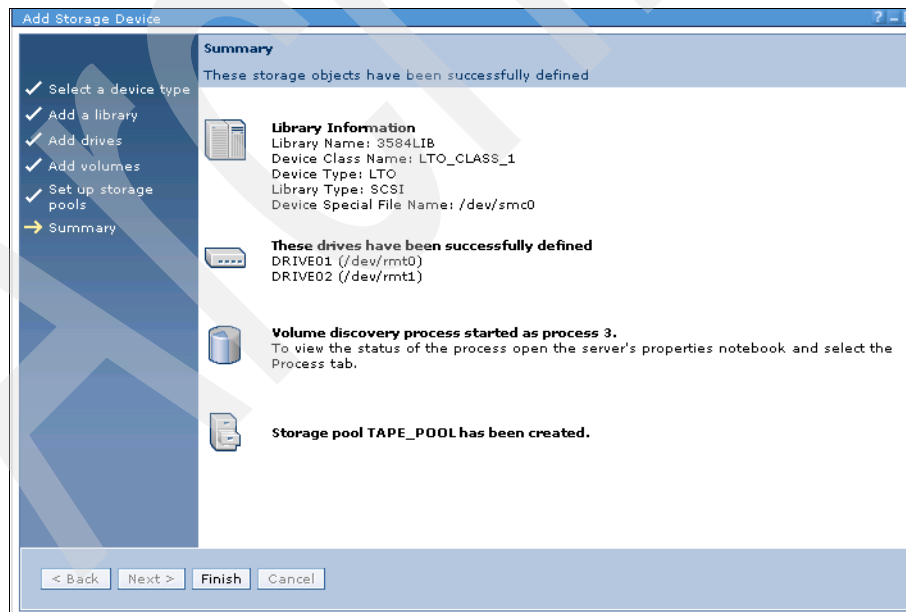


Figure 11-21 Storage device wizard: Final overview

In the storage devices main window, you should now see the previously defined library. By clicking the name of the library, you get the library properties overview, where you can adjust the library parameters and add or remove volumes. Figure 11-22 shows the general library properties box. Note that the serial number and the worldwide name of the library have been automatically detected.

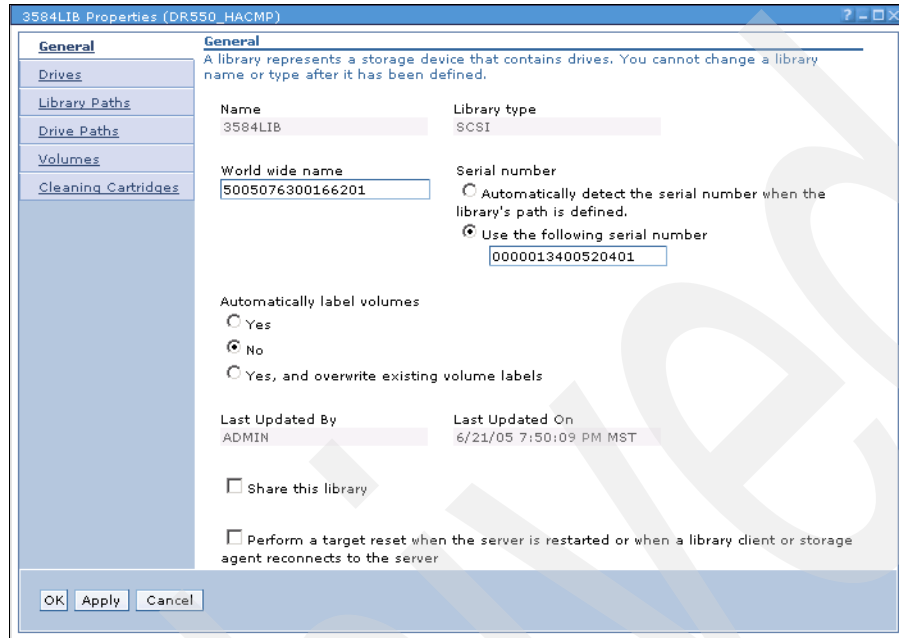


Figure 11-22 Library properties box

A cleaning frequency of NONE is preferred for the drives, because the 3592-J1A (or TS1120) and 3588 Tape Drives internally maintain a cleaning algorithm and request cleaning cartridges from the library when required.

Note that a standard LTO device class named LTO\_CLASS\_1 has already been defined and activated by the wizard in read/write mode.

9. Define an additional device class for the 3588 WORM media.

From the storage devices main window, select **Servers** → **View Device Classes**. You get a list showing the Device Classes defined for the SSAM Server. Select **Create a Device Class** from the drop-down menu of this list, as shown in Figure 11-23.

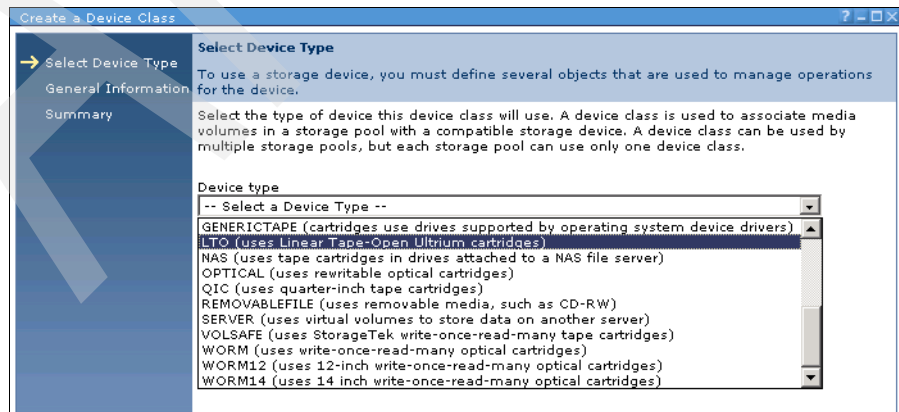


Figure 11-23 Device classes wizard: Select device type

Complete this step by selecting **LTO** as the device type. By clicking **Next**, you will see the device class properties box. Provide a useful name such as **LTO\_CLASS\_WORM** and select the previously defined library. Enable the **WORM** capability by checking the check box, as shown in Figure 11-24, and finish the wizard.

Figure 11-24 Device classes wizard: Properties

#### 10. Verify the definitions.

Verify your definitions by issuing the following commands at the command line:

```
query library
query drive
query path
query devclass
query stgpool
```

For a more detailed list, specify **f=d**, which is the short form of **format=detailed**.

Example 11-10 shows detailed information about the previously defined device class using the command **query devclass LTO\_CLASS\_WORM f=d**.

Example 11-10 Detailed informational output for a WORM-enabled device class

---

```
Device Class Name: LTO_CLASS_WORM
Device Access Strategy: Sequential
Storage Pool Count: 0
Device Type: LTO
Format: DRIVE
Est/Max Capacity (MB):
Mount Limit: DRIVES
Mount Wait (min): 60
Mount Retention (min): 60
Label Prefix: AD5M
Library: 3584LIB
Directory:
Server Name:
Retry Period:
Retry Interval:
Shared:
High-level Address:
Minimum Capacity:
WORM: Yes
Scaled Capacity:
Last Update by (administrator): ADMIN
```

---

## 11.4.6 Integrating 3588 and TS3500 into SSAM storage hierarchy

In this section, we describe three scenarios for integrating the 3588 and TS3500 into the SSAM storage hierarchy.

### Scenario One: Using (WORM) tape as the migration destination

For this scenario, complete the following steps in Figure 11-25:

1. Define a sequential access storage pool for migration-based thresholds, off the primary disk pool to the tape pool associated with the WORM tape device class.

Define a sequential access storage pool named `ARCHIVE_TAPES` associated with the device class named `LTO_CLASS_WORM` by typing the following command:

```
define stgpool ARCHIVE_TAPES LTO_CLASS_WORM pooltype=primary maxxcrash=999
```

If you are using the Web interface, go to the Storage devices main window and select **Servers** → **View Storage Pools** → **Create a Storage Pool** and follow the wizard. Fill in the storage pool name `ARCHIVE_TAPES` and an optional description. Choose **Primary, sequential access** as the storage pool type.

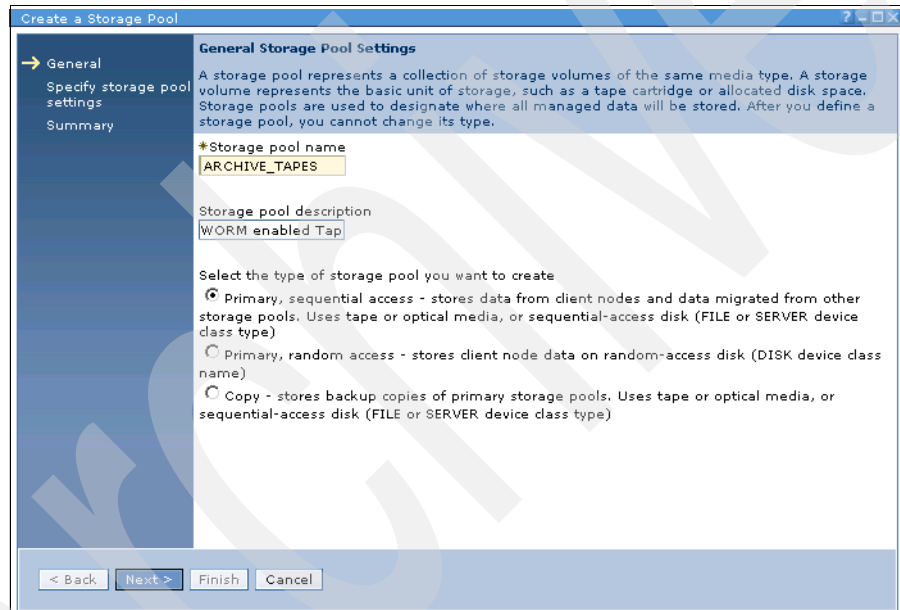


Figure 11-25 Storage pool wizard: Select name, description, and type

Finish the wizard by selecting the device class `LTO_CLASS_WORM` for the storage pool, as shown in Figure 11-26 on page 481, and confirm the overview of the created storage pool settings.

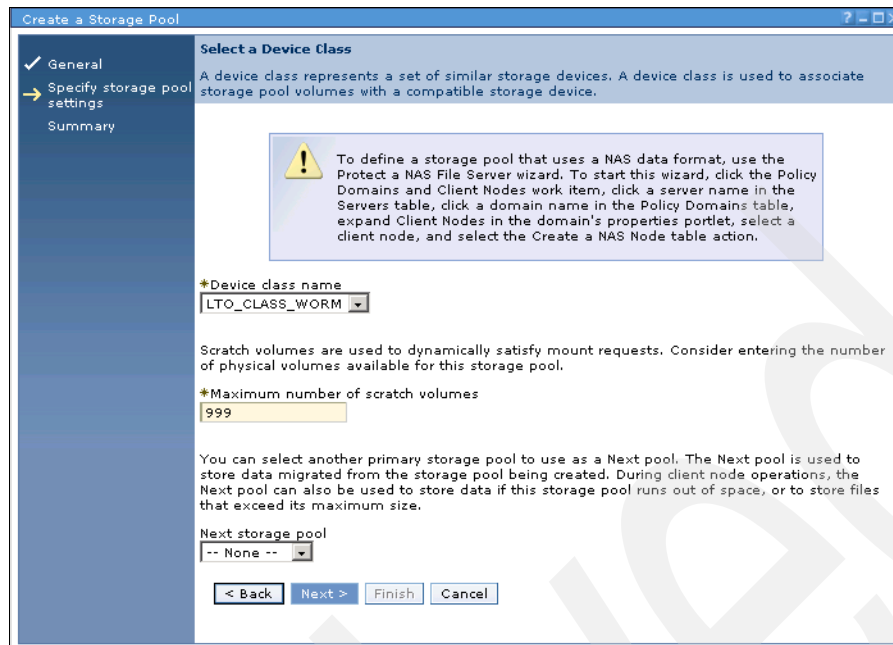


Figure 11-26 Storage pool wizard: Select the WORM-enabled device class

You should now see an overview of the defined storage pools, as shown in Figure 11-27. Note that the ARCHIVEPOOL is preconfigured for the SSAM Server with the appropriate capacity. In our environment, we recreated the ARCHIVEPOOL with a smaller size.

Select	Name	Device Class	Estimated Capacity (MB)	Percent Utilized	Next
<input type="radio"/>	ARCHIVEPOOL	DISK	3 G	18.9	
<input type="radio"/>	ARCHIVE_TAPES	LTO_CLASS_WORM	0.0 M	0.0	
<input type="radio"/>	TAPE_POOL	LTO_CLASS_1	0.0 M	0.0	

Page 1 of 1      Total: 3    Filtered: 3    Displayed: 3    Selected: 0

Figure 11-27 Defined storage pools for the SSAM Server

By clicking the storage pool name, you get the storage pool properties box. Set the values according to your needs. On the command line, issue the command **help update stgpool** for information about syntax and the use of parameters.

In the context of a WORM medium, the Delay Period for Volumes Reuse parameter does not make sense at first. By nature, a tape WORM media cannot be reused at all. However, the effect of this parameter is that an “empty” volume (all data expired from the SSAM point of view) is kept in the SSAM database in status PENDING until the delay period passes. After that, SSAM wipes out all references to this volume; it is unknown from now on. This status enables you to define a mechanism within SSAM to identify expired volumes, and, for example, convey them to a scrapping process. For example, a daily **query volume status=pending** command would display the “empty” tapes. Furthermore, these volumes should have been kept in pending status for as long as database backups are held. This allows a rollback to a previous version of the database in case of a disaster while still having access to the data on the WORM media, which would not be the case if the volumes had already been scrapped. Therefore, it is a good practice to set the Delay

Period for Volumes Reuse to a non-zero value; in the context of the DR550 SSAM Server, set this to three days, because database backups are kept for this period of time.

CRC Data is set to **YES**; this improves the data integrity for the copy objects. CRC Data specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. By setting CRC Data Validation to **YES**, data is stored that contains CRC information. When you schedule audit volume processing, you can continually ensure the integrity of data stored in your storage hierarchy. We assume that you always configure the CRC data validation on the DR550 SSAM Server, even if you never use the audit volume processing to validate the data.

**Tip:** Use the SSAM features such as Delay Period for Volumes Reuse and cyclic redundancy check (CRC) data validation to improve data integrity.

2. Include the sequential access storage pool in your storage hierarchy.

To use the previously defined storage pool `ARCHIVE_TAPES`, you need to include it in the existing storage hierarchy by updating the `ARCHIVEPOOL` to point to the new pool:

```
update stgpool archivepool nextstgpool=archive_tapes
```

If you are using the Web interface, open the `ARCHIVEPOOL` properties box from the Storage Pools overview and choose **ARCHIVE\_TAPES** from the **Next Storage Pool** drop-down menu, as shown in Figure 11-28.

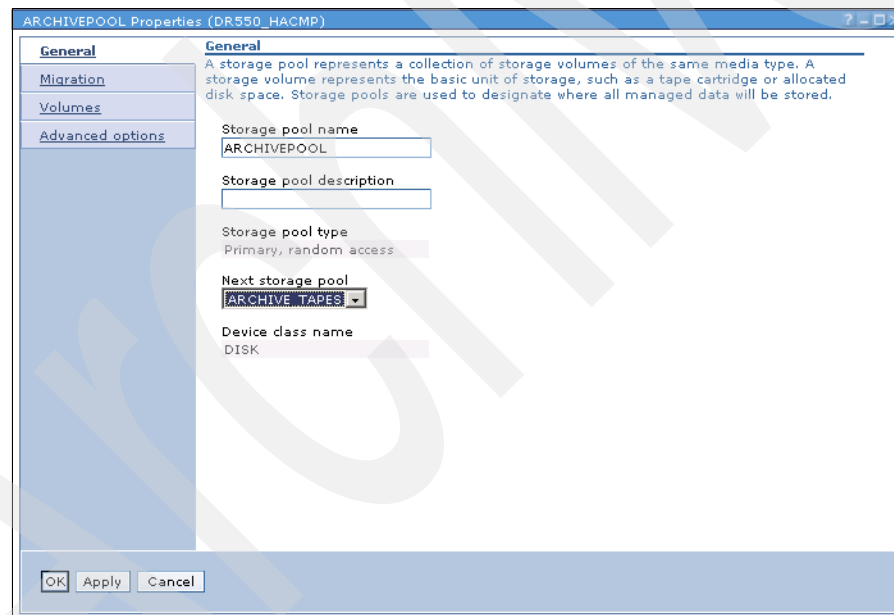


Figure 11-28 ARCHIVEPOOL properties box: Choose the next storage pool

The High Migration Threshold and Low Migration Threshold values are kept at the default values. Crossing the high threshold causes SSAM to start migrating data off this pool to the next specified pool until the percentage of occupancy (versus the total capacity of the storage pool) indicated by the low threshold has been reached. A good practice to control the occupancy level is to work with scripts and schedules to adjust these values dynamically according to, for example, storage capacity requirements or backup schedules. For example, set up a schedule to run every day to initiate the migration of all data from the disk pool to the tape pool on a daily basis. This schedule would call a script that sets both the low migration threshold and then the high migration threshold to zero. This will instantly result in the migration of all data off the disk pool into the tape pool. After



completion, the values will be set to the original values, again using a combination of scripts and schedules. The advantage of this method, rather than maintaining the migration thresholds constantly at the same level, is a guarantee that all objects will eventually migrate onto WORM tape. Otherwise, small objects run the risk of never migrating, because they could always be within a capacity level that is lower than the low migration threshold.

Cache Migrated Files controls whether or not objects are deleted from the ARCHIVEPOOL after successful migration to the ARCHIVE\_TAPES. Deleting objects releases the space in the primary pool, while caching them increases the hit ratio on disk and reduces the number of tape mounts required in case a object is accessed by users.

Migration Delay set to 0 means that objects can be migrated by SSAM according to the occupancy level. If, for example, a requirement exists to keep objects in the disk pool for at least one year to guarantee fast access times, this parameter would have to be set to a value of 365.

Migration Continue specifies if migration should continue disregarding and overriding the (non-zero) Migration Delay value. If you do not want this behavior, consider an appropriate sizing of the primary disk pool.

## Scenario Two: Using (WORM) tape as data backup destination

For this scenario, complete the following steps in Figure 11-29:

1. Define a copy storage pool for incremental backups of the archived data to WORM tape associated with the device class.

Define a primary copy storage pool named COPY\_TAPES associated with the device class named LTO\_CLASS\_WORM by typing the following command:

```
define stgpool copy_tapes lto_class_worm pooltype=copy maxscratch=999
reusedelay=3 crcdata=yes
```

If using the Web interface, go to the storage devices main window and select **Servers** → **View Storage Pools** → **Create a Storage Pool** and follow the wizard. Fill in the storage pool name COPY\_TAPES and add an optional description. Choose **Copy** as the storage pool type.

The screenshot shows a web-based wizard titled "Create a Storage Pool". On the left, a sidebar lists "General", "Specify storage pool settings", and "Summary", with "General" being the active tab. The main content area is titled "General Storage Pool Settings" and includes a descriptive paragraph about storage pools. Below this, there are two text input fields: "Storage pool name" with the value "COPY\_TAPES" and "Storage pool description" with the value "WORM enabled cop". A section titled "Select the type of storage pool you want to create" contains three radio button options: "Primary, sequential access", "Primary, random access", and "Copy". The "Copy" option is selected. At the bottom of the wizard, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

Figure 11-29 Storage pool wizard: Select name and description for copypool

Define LTO\_CLASS\_WORM as the device class and choose the maximum number of scratch volumes, as shown in Figure 11-30.

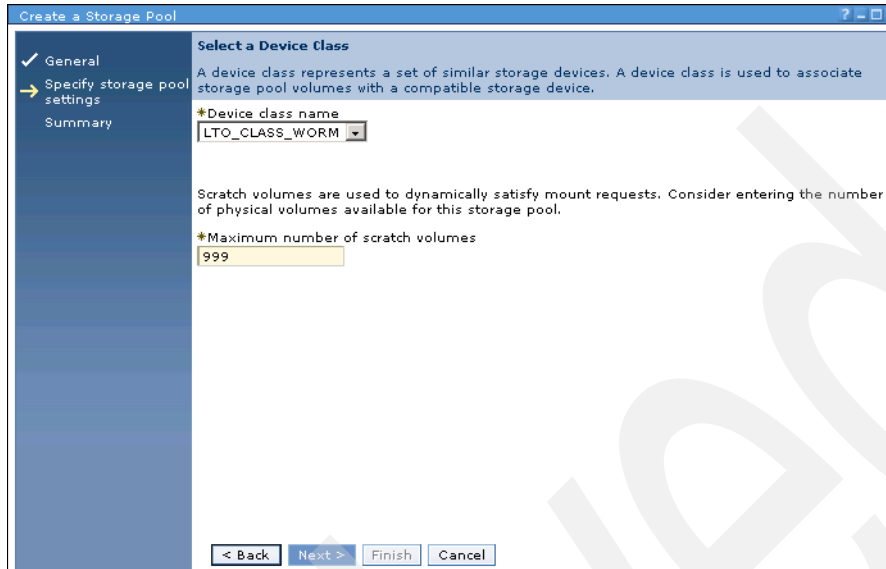


Figure 11-30 Storage pool wizard: Choose device class

Finish the wizard by confirming the summary.

This new copy storage pool can now be used as a target to incrementally back up data to the 3589 WORM media. The backup will be established with two scheduled processes: One copying data off the disk storage pool ARCHIVEPOOL and another copying data off the tape storage pool ARCHIVE\_TAPES, because any document or data object can be stored in either of the pools, depending on whether it has already been migrated or not.

2. Create a server command script to back up the primary pool to the copy pool.

To create a server command script named BASTGPOOL to initiate automatic backups of first the primary pool ARCHIVEPOOL and, after its completion, the sequential access pool ARCHIVEPOOLTAPE, type:

```
define script bastgpool
update script bastgpool "backup stgpool archivepool copy_tapes wait=yes"
update script bastgpool "backup stgpool archive_tapes copy_tapes wait=yes"
```

If using the Web interface, go to the storage devices main window and select **Servers** → **Server Properties** → **Scripts** → **Create Script**, as shown in Figure 11-31.

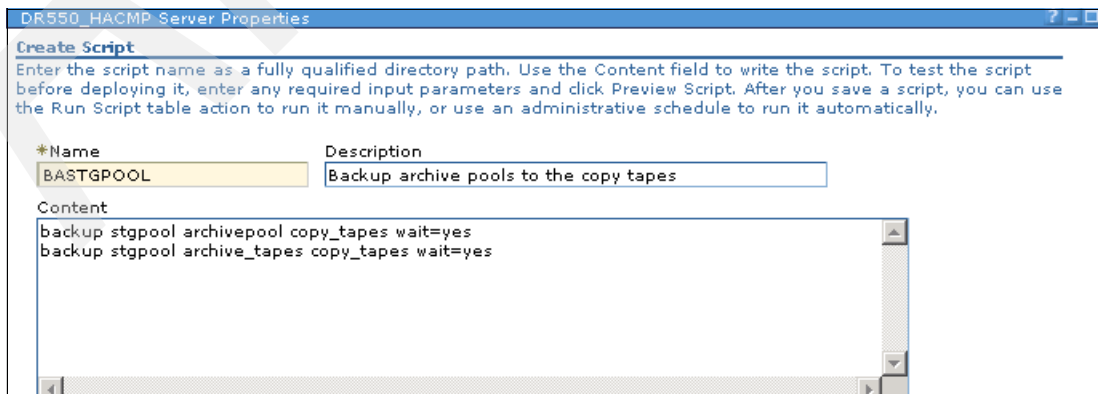


Figure 11-31 Define a new command script BASTGPOOL

Note that by using the Administration Center Web interface, you have the capability to test scripts (even with user-defined variables) before production use.

3. Create a schedule to execute a server command script.

Create a schedule named BASTGPOOL or execute the previously created server command script BASTGPOOL at the command line:

```
define schedule bastgpool cmd="run bastgpool" active=yes starttime="12:00:00"
```

If using the Web interface, select **Server Properties** → **Administrative Schedules** → **Create a Schedule**. Follow the wizard and provide BASTGPOOL as the schedule name, add a description, and enter **run BASTGPOOL** as the command to run in the schedule, as shown in Figure 11-32.

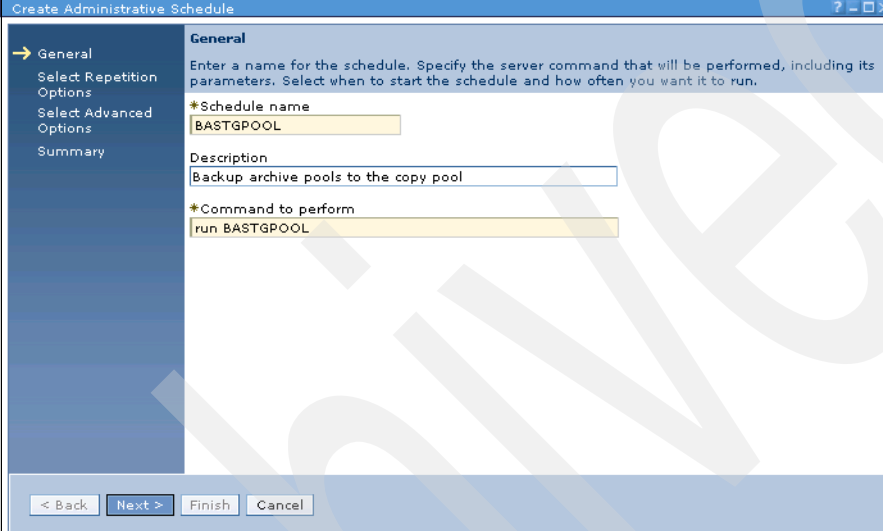
The screenshot shows the 'Create Administrative Schedule' wizard with the 'General' tab selected. The left sidebar contains links for 'General', 'Select Repetition Options', 'Select Advanced Options', and 'Summary'. The main area has a title bar 'Create Administrative Schedule' and a subtitle 'General'. Below the subtitle is a description: 'Enter a name for the schedule. Specify the server command that will be performed, including its parameters. Select when to start the schedule and how often you want it to run.' There are three input fields: '\*Schedule name' with the value 'BASTGPOOL', 'Description' with the value 'Backup archive pools to the copy pool', and '\*Command to perform' with the value 'run BASTGPOOL'. At the bottom are buttons for '< Back', 'Next >', 'Finish', and 'Cancel'.

Figure 11-32 Administrative schedule wizard: Choose name and command to run

Select the time, date, and repeat frequency of this schedule. See Figure 11-33.

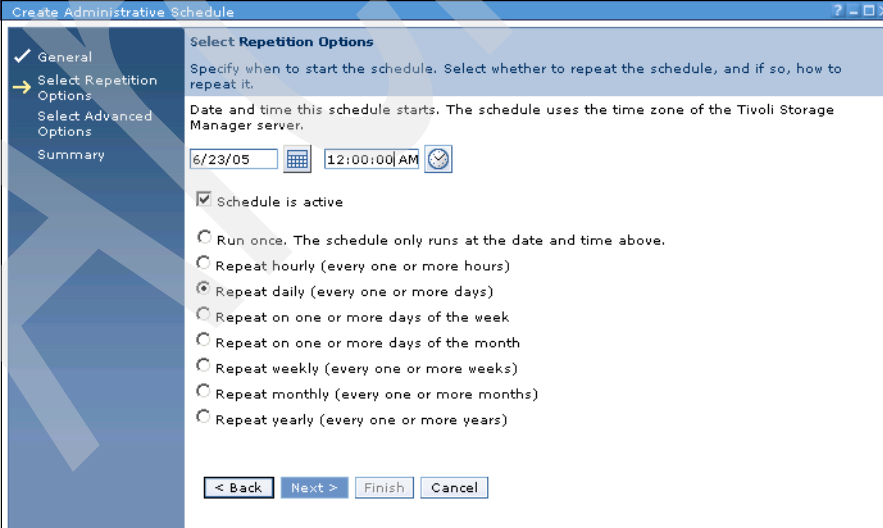
The screenshot shows the 'Create Administrative Schedule' wizard with the 'Select Repetition Options' tab selected. The left sidebar contains links for 'General', 'Select Repetition Options', 'Select Advanced Options', and 'Summary'. The main area has a title bar 'Create Administrative Schedule' and a subtitle 'Select Repetition Options'. Below the subtitle is a description: 'Specify when to start the schedule. Select whether to repeat the schedule, and if so, how to repeat it.' There are two input fields: 'Date and time this schedule starts. The schedule uses the time zone of the Tivoli Storage Manager server.' with the value '6/23/05' and '12:00:00 AM'. Below these are radio buttons for 'Schedule is active' (checked), 'Run once. The schedule only runs at the date and time above.', 'Repeat hourly (every one or more hours)', 'Repeat daily (every one or more days)', 'Repeat on one or more days of the week', 'Repeat on one or more days of the month', 'Repeat weekly (every one or more weeks)', 'Repeat monthly (every one or more months)', and 'Repeat yearly (every one or more years)'. At the bottom are buttons for '< Back', 'Next >', 'Finish', and 'Cancel'.

Figure 11-33 Define schedule, execution time, and repetition

Specify the options shown in Figure 11-34. End the wizard by committing the summary.

The screenshot shows a Windows-style dialog box titled "Create Administrative Schedule". On the left is a vertical navigation pane with four items: "General" (checked), "Select Repetition Options" (checked), "Select Advanced Options" (highlighted with a yellow arrow), and "Summary". The main area is titled "Advanced Schedule Options" and contains the following text: "If two or more schedules are set to start at the same time, the schedule with the highest priority starts first." Below this is a "Schedule priority (1 has the highest priority)" section with a dropdown menu showing the number "5". The "Schedule Expiration" section has two radio buttons: "Schedule never expires" (which is selected) and "Schedule expires on the following date" (with an empty text box and a calendar icon). At the bottom of the main area is a text box for a time limit, containing the number "1" and a dropdown menu set to "Hours". At the bottom of the dialog are four buttons: "< Back", "Next >", "Finish", and "Cancel".

Figure 11-34 Define further options

In our example, the SSAM Server will, from 6/23/05 (June 23, 2005) and forward, attempt to start the script BASTGPOOL on a daily basis at 12:00:00. If it cannot be started within the specified duration of one hour, the script will be skipped and thus not be executed until the next day. Once started, the primary pools will be backed up to the tape copy pool as specified in the server command script. The Start time has to be chosen individually to reflect the desired schedules and workload within the customer environment.

Note that the tape library has to have a sufficient number of cartridges and that these are checked into SSAM. The number of cartridges depends on the amount of data stored in the primary storage pool or pools being backed up. If not enough media are available, the schedule can be suspended by making it inactive with the following command:

```
update schedule bastgpool t=a active=no
```

If you are using the Web interface, select **Server Properties** → **Administrative Schedules**, select the **BASTGPOOL** script, and choose **Modify Schedule**. In the schedule properties notebook, uncheck the **Schedule is active** check box as shown in Figure 11-35 on page 487.

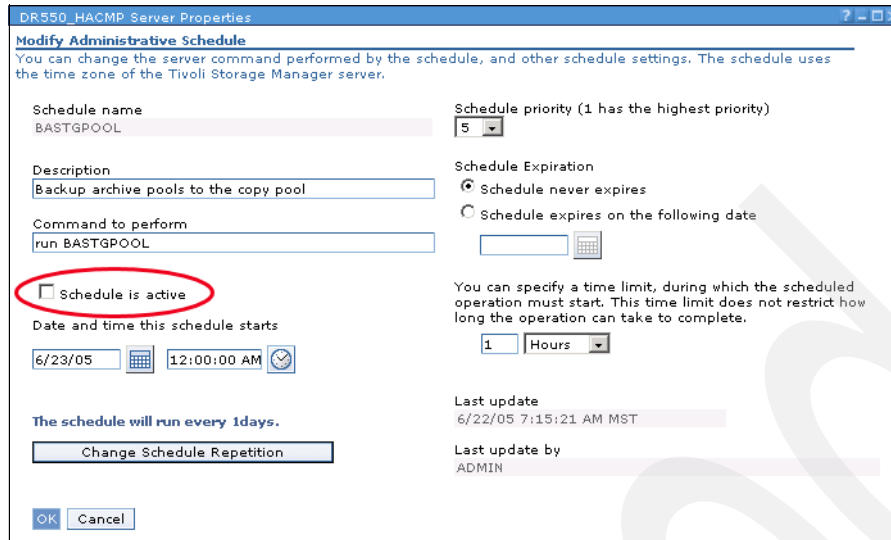


Figure 11-35 Temporary deactivate administrative schedule

**Tip:** Working with scripts instead of single schedules has a significant advantage: Single schedules will be executed based on their starting time, without depending on other schedules. Within a script, the parameter `wait=yes` enables you to initiate a process dependent on the previous one, which is often desired. In our example, we want to have the backup of the primary disk storage pool happen and complete first, before the backup of the sequential access storage pool is carried out. A script can include many and any kind of SSAM server commands, such as **disable session**, **expire inventory**, and **update stgpool**.

4. *Optional:* Run the script once to verify that it is working correctly.

The script can be started manually to verify that the desired backups are actually carried out. Note that this can be a time-consuming process, depending on the amount of data already stored in the SSAM primary storage pools. It also assumes that enough tape media are inserted in the library and available for SSAM use.

Start the script BASTGPOOL to initiate the backup of the primary storage pools:

```
run bastgpool
```

If using the Web interface, select **Server Properties** → **Scripts**, select the **BASTGPOOL** script, and choose **Run Script**. In the Run Script window, you should leave the check box **Show processing information in addition to script commands** checked and watch the script results. Click **Run Script** to start the process. See Figure 11-36.

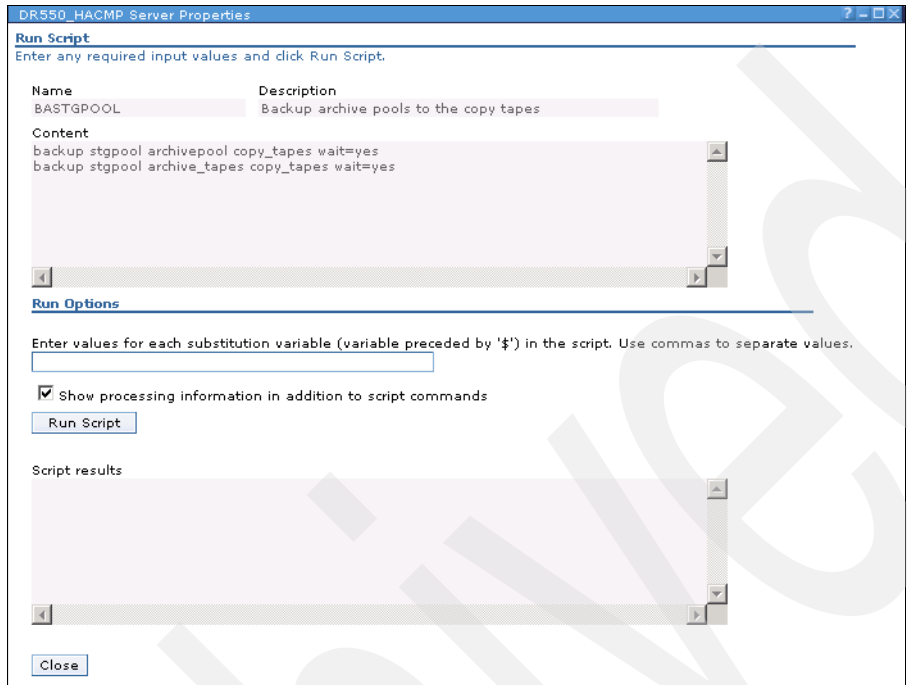


Figure 11-36 Run script: Watch the output in the script results box

The output of a successful backup process should look similar to the illustration shown in Figure 11-37.

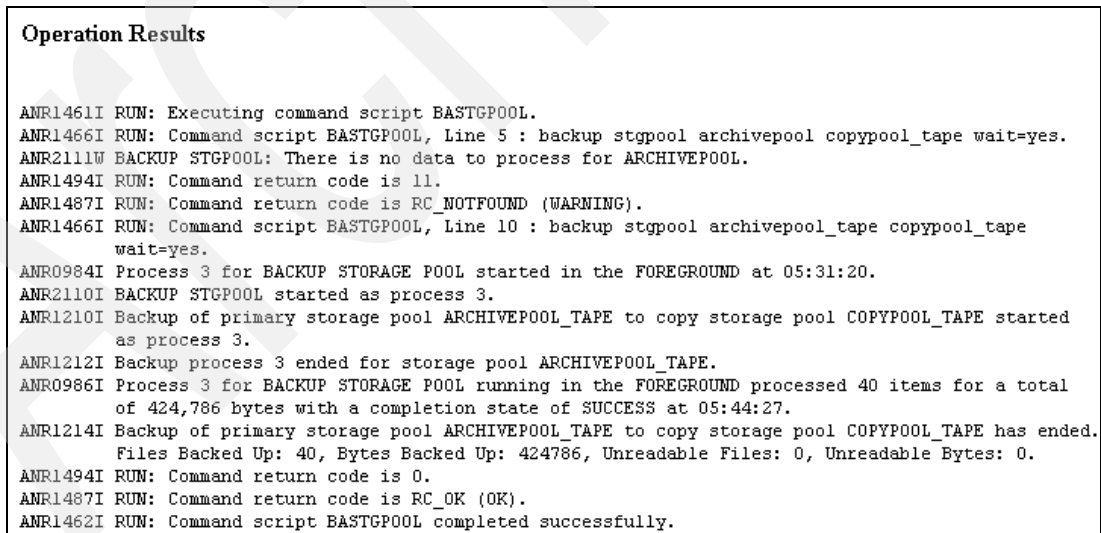


Figure 11-37 Output of successful backup process

### Scenario Three: Using (rewritable) tape for database backups

For this scenario, create a schedule to execute a SSAM database backup to rewritable tape media. For this purpose, we use the automatically generated read/write tape device class LTO\_CLASS\_1 generated by the initial device wizard we used in 11.4.5, “Defining 3588 and TS3500 devices to SSAM” on page 472.

Create a schedule named BADBTAPE to periodically generate backups of the Tivoli Storage Manager database onto rewritable tape media:

```
define schedule badbtape cmd="backup db devc=LTO_CLASS_1 type=full" active=yes
starttime="07:00:00"
```

If you are using the Web interface, select **Server Properties** → **Administrative Schedules** and run **Create a Schedule**. Follow the wizard and provide BADBTAPE as the schedule name, a description, and type **backup db devc=LTO\_CLASS\_1 type=full** as the command to run in this schedule, as shown in Figure 11-38.

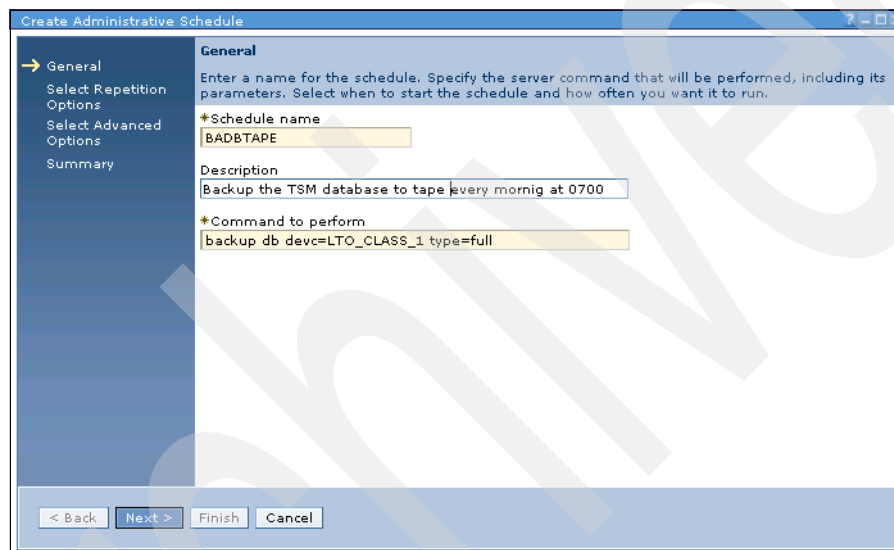


Figure 11-38 Create administrative schedule for daily database backups to tape

Accept the default settings for the remaining steps until you get to the summary. Then click **Finish** to complete the wizard.

Start the SSAM database backup once with the following command:

```
backup db devc=LTO_CLASS_1 type=full
```

Your database should now be successfully backed up to rewritable tape media. Check the activity log with the command **query actlog** for entries such as the following:

```
ANR4550I Full database backup (process 10) complete, 643 pages copied.
ANR0985I Process 10 for DATABASE BACKUP running in the BACKGROUND completed with
completion state SUCCESS at 22:47:18.
```

The command **query libvol** should show at least one rewritable tape volume with a status of DbBackup, as shown in Example 11-11.

*Example 11-11 query libvol command output*

Library Name	Volume Name	Status	Owner	Last Use	Home Element	Device Type
3584LIB	3DW585	Private(Worm)			1,026	LTO
3584LIB	3DW586	Scratch(Worm)			1,027	LTO
3584LIB	3DW587	Scratch(Worm)			1,028	LTO
3584LIB	3TX006L3	Private		DbBackup	1,025	LTO
3584LIB	3TX053L3	Scratch			1,029	LTO
3584LIB	3TX059L3	Scratch			1,030	LTO

This schedule initiates a full database backup onto rewritable tape media every day at 07:00:00. The preconfigured database backup onto specific disk space in the DS4200 starts at 06:00:00, and it will be completed by the time the backup to tape starts. If the script cannot be started within the specified duration of one hour, the script will be skipped and not be executed until the next day.

This step demonstrates how to integrate a database backup to tape into the preconfigured SSAM Server. A better practice is again to create a server command script, combining both the database backup to disk and then to tape, consecutively, using the wait=yes parameter, as described in the previous steps. It is up to the customer to customize the SSAM Server concepts according to the business needs and requirements of the company.



## DR550 backup and restore

In this chapter, we explain the basics of disaster recovery and the concept of the Tivoli Storage Manager Disaster Recovery Manager. We show how to use Disaster Recovery Manager (DRM) in combination with facilities to create an AIX image and to back up and restore an operating system and Tivoli Storage Manager server database.

See Chapter 13, “DR550 and Enhanced Remote Mirroring” on page 521 for Disaster Recovery and Business Continuity capabilities based on the DS4000 Enhanced Remote Mirroring feature.

## 12.1 What disaster recovery is

Disaster recovery is the process of restoring the operations of a business or organization in the event of a catastrophe. There can be many aspects related to the restoration, including facilities, equipment, personnel, supplies, customer services, and data. One of the most valuable business assets is the critical data that resides on the computer systems throughout the company. The recovery of this data needs to be a primary focus of the disaster recovery plan. IBM Tivoli Storage Manager, along with the Disaster Recovery Manager function included in IBM Tivoli Storage Manager Extended Edition, will assist you in the technical steps that you need to make your data available to users after a widespread failure.

Distributed data recovery restores data to workstations, application servers, and file servers in the event of data loss due to accidental erasure, media failures, sabotage, and natural disasters. It involves creating, managing, and recovering copies of distributed data. These copies should be taken off-site to minimize the chance that a disaster will destroy backup copies along with primary copies. Many data administrators also choose to keep backup copies on-site to expedite recovery from smaller media failures.

Disaster recovery requires, at a minimum, creating copies of primary data. Many businesses and backup products stop here. To achieve a complete recovery solution for distributed data, you must consider several additional features.

## 12.2 Disaster Recovery Manager

Disaster Recovery Manager (DRM) is a feature of IBM Tivoli Storage Manager Extended Edition and coordinates and automates the process of recovering from a disaster. It provides for off-site media management, automated restoration of the Tivoli Storage Manager server, and managed client recovery. It complements the already implemented robust protection features of Tivoli Storage Manager and automates many already facilitated protection functions. DRM automatically captures information required to recover the Tivoli Storage Manager server after a disaster. It assists in preparing a plan that allows recovery in the most expedient manner. This disaster recovery plan contains information, scripts, and procedures needed to automate server restoration and helps ensure quick recovery of your data after a disaster. DRM also manages and tracks the movement of off-site media to reduce the time required to recover in the event of a disaster. It is able to track media that are stored on-site, in-transit, or off-site in a vault, no matter whether it is a manual or electronic vault, so you can locate your data easily if disaster strikes. Figure 12-1 on page 493 shows the DRM media cycle.

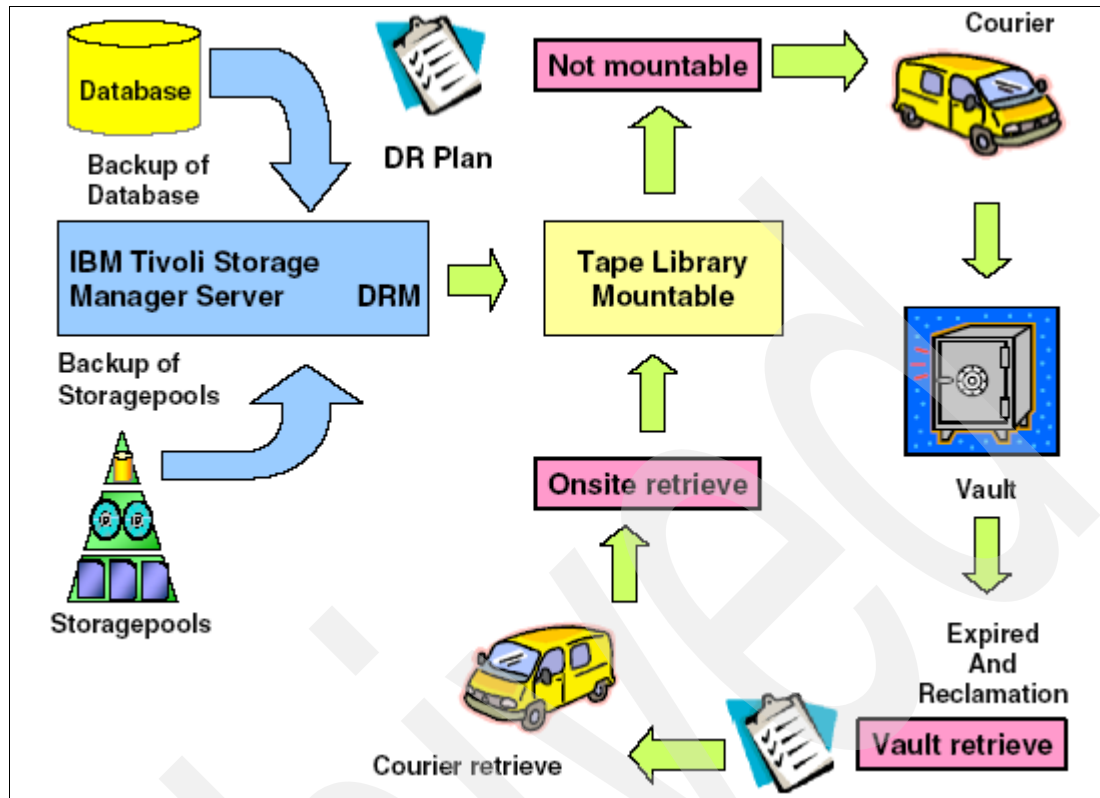


Figure 12-1 Disaster Recovery Manager off-site media tracking flow

DRM can also capture client recovery information. You can use this information to assist with identifying what clients need to be recovered, in what order, and what is required to perform the recovery, including data and media that are not managed by Tivoli Storage Manager. Client recovery is not considered in the context of the DR550, because regular Tivoli Storage Manager backup and archive clients cannot store data in the offering.

In a typical protected Tivoli Storage Manager environment, after each night's normal client backup, the copy storage pools are also updated with the newly arrived data. Then, a server database backup is done. The latest generated volumes are sent to a safe location, and a recovery plan file is regenerated by Disaster Recovery Manager to make sure it includes the latest information. As data expires from the on-site pools, it also expires from the off-site pools and unneeded database backups. Disaster Recovery Manager also tracks such media as they become empty so that you can report on free tapes that can be brought back on-site for reuse.

## Volume tracking

Disaster Recovery Manager provides several levels of volume tracking. Disaster Recovery Manager volume management includes:

- Identifying which off-site volumes are needed for a given recovery: Disaster Recovery Manager knows the volumes that are associated with each primary Tivoli Storage Manager server so that you can initiate a complete recovery of all storage pools, or only a partial recovery, depending on the extent of the disaster. You can also configure Disaster Recovery Manager to track volumes only from certain storage pools (this is useful, for example, to provide critical client nodes with full off-site protection, and other, less-critical nodes, with no off-site protection).

- ▶ Integrating with tape management systems: Because Disaster Recovery Manager is fully integrated with tape management, every time a new tape is created in the corresponding copy storage pools, it is automatically eligible for off-site movement.
- ▶ Recycling partially filled volumes: Off-site volumes are reclaimed just as on-site volumes are. Disaster Recovery Manager enables you to see which volumes have reached an empty state because of reclamation so that you can request that they are returned on-site. This feature is not applicable for WORM media pools, where space reclamation is not enabled.
- ▶ Tracking off-site volumes: This is one of Disaster Recovery Manager's strongest features. Disaster Recovery Manager manages tapes by assigning a special, predefined set of states to each off-site tape. Depending on where the tape should be, there are two possible directions for a tape: from on-site to off-site and from off-site to on-site. The first starts during normal backup processing to save up-to-date data to the copy storage pool. The tapes pass through a number of states in their journey from the production tape library to the safe vault. Then, time elapses while the tape remains off-site, ready to be used for a restore in the event of a disaster. During this time, data is gradually expiring from the tape. When the tape finally reaches its reclamation threshold, it is reclaimed by normal processes. Once empty, it moves in the reverse direction, that is, it is returned on-site for reuse. Again, with the use of WORM media and space reclamation turned off, the journey back on-site will only occur if a disaster recovery has to be performed.

## On-site to off-site

Tapes are transitioned in the following order: Mountable, NotMountable, Courier, and Vault.

1. **Mountable:** Newly created copy storage pool and database backup tapes are in a mountable state as long as they are still online in the tape library. An on-site disaster would destroy these tapes, because they have not yet been taken safely off-site. Therefore, Disaster Recovery Manager knows that these tapes are not part of the guaranteed recovery set when creating its recovery plan. The next step for these tapes is to remove them from the tape library and, using Disaster Recovery Manager commands, change them to the NotMountable state.
2. **NotMountable:** The volumes are still in the on-site location, but are not online to Tivoli Storage Manager. They could be in the data center, in a locked room near the tape library, or on a different floor waiting to be taken off-site. The actual physical location is not important for Tivoli Storage Manager. What is important is that volumes are still on-site and it is assumed that they would also be destroyed in a disaster. This state only changes when they are taken off-site, usually by courier pickup. A Disaster Recovery Manager command will change the state of the eligible volumes.
3. **Courier:** An intermediate situation in which you consider that the tapes are in transit. This means that you are still considering the physical transition from your on-site location to the vault as a potential risk. For example, suppose that you have a service agreement with an external company to deliver those tapes for you. Disaster Recovery Manager takes this into account if you create a recovery plan file just after changing the tape state to Courier. After you are sure that all tapes have reached their destination, it is safe to change to the next and final state.
4. **Vault:** After you have received acknowledgement that the volumes have been safely received, you can change them to the Vault state, which is the final state in sending a tape from an on-site location to an off-site location. After the volumes are off-site, they stay there until they are eventually reclaimed (or are required for an actual recovery). When they are empty, they are then ready to begin the reverse trip. Disaster Recovery Manager **query** commands enable you to see which volumes can be retrieved so that you can issue a list to your vault administrators.

## Off-site to on-site

Tapes are returned on-site in the following order: VaultRetrieve, CourierRetrieve, and OnsiteRetrieve.

1. **VaultRetrieve:** When all data in an off-site tape is no longer valid, the tape state is automatically changed to VaultRetrieve, meaning that it is available to be brought back on-site for usage. This happens for both Tivoli Storage Manager expired database backups and empty reclaimed copy storage pool tapes. Typically, you would send this list to your storage vault administrators, maybe once a week, so they know which tapes to find and send back to you. Tapes that are in VaultRetrieve status are still part of the safe recovery set until they are physically removed from the vault.
2. **CourierRetrieve:** Change a tape to this state when you know that it has been taken from the vault and is in transit back on-site. Similar to the Courier state, these tapes may or may not be preserved safely in the event of a disaster, so it is important to know which tapes they are. After the volumes are correctly acknowledged upon delivery on-site, they are ready for the next state.
3. **OnsiteRetrieve:** After the volumes are back on-site, you can use them as scratch tapes. Usually, they will be loaded back into the tape library.

**Tip:** Disaster Recovery Manager enables you to skip some intermediate transition states. For example, you could choose to change your tapes directly from the Mountable to the Vault state. However, we do not recommend that you do this, because it means less control over exactly which tapes are and are not safely available for recovery in the event of a disaster. Disaster Recovery Manager uses the **prepare** command to generate a plan file that will contain critical information needed for recovery.

## 12.3 Recovery strategy for the server

To make the creation and maintenance of the server disaster recovery plan easier, the **prepare** command automatically queries the required information from the Tivoli Storage Manager server and creates the recovery plan file. The **prepare** command can be scheduled using the Tivoli Storage Manager central scheduling capabilities.

### Auditable plan for Tivoli Storage Manager server

The recovery plan file contains the information and procedures necessary to assist with the recovery of the Tivoli Storage Manager server. The information in the plan file includes site-specific server recovery instructions and information as defined by the administrator (for example, contact names and telephone numbers for critical people and their backups).

The necessary items to recover a Tivoli Storage Manager server are:

1. List of Tivoli Storage Manager database backup and copy storage pool volumes required to perform the recovery (including the off-site location where the volumes reside)
2. Devices required to read the database backup and copy storage pool volumes
3. Space requirements for the Tivoli Storage Manager database and recovery log
4. Copy of Tivoli Storage Manager server options file, device configuration file, and volume history information file
5. Shell scripts (on UNIX) and Tivoli Storage Manager macros for performing server database recovery and primary storage pool recovery

## Off-site recovery media management

Knowing the location of off-site recovery media is critical to the successful implementation of a disaster recovery management plan. The off-site recovery media management function provides:

- ▶ Determination of which database and copy storage pool volumes need to be moved off-site and back on-site
- ▶ Automatic ejection of volumes from an automated library
- ▶ Tracking the media location and state in the Tivoli Storage Manager database

This function allows database backup volumes and copy storage pool volumes to be treated as logical collections that are selected to move off-site for safekeeping and on-site for use. The reclamation of off-site volumes includes the capability to specify the number of days to retain a Tivoli Storage Manager database backup series. After the expiration interval is reached, the data on the media is no longer considered to be valid. The media can then be reused (or discarded).

Figure 12-2 illustrates how your off-site data could be used to recover your environment. Note that V1 is the point in time requested; therefore, you cannot only rebuild the latest one, but also data from any specific point in time that you still have saved. The execution of the recovery scripts (which perform the Automatic Recovery Steps in the figure) starts after you have reinstalled the operating system and Tivoli Storage Manager server code on your replacement server hardware.

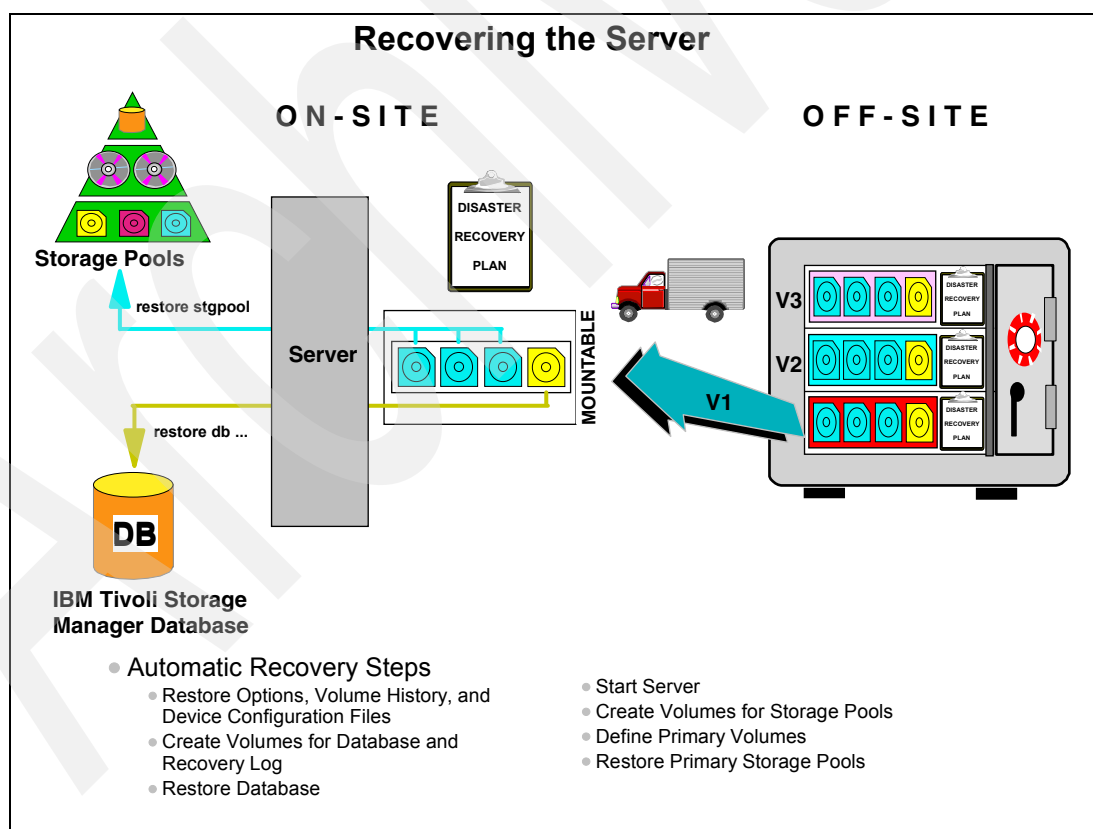


Figure 12-2 Restoring a Tivoli Storage Manager server

## Additional disaster recovery issues

Disaster recovery goes far beyond simple technical measures. To have a fully operational and prepared environment, you must also pay attention to additional issues, such as those described in the following sections.

### Hardware system requirements

Disaster Recovery Manager creates a recovery plan file based on the information and space allocation on the Tivoli Storage Manager production server machine. This means that you must evaluate whether to have a similar machine for off-site recovery and make the changes to fit the new environment.

### Additional operating system recovery steps

Depending on which operating system Tivoli Storage Manager is installed, you might need to send special DVD or tape images (for the specific OS recovery steps) to the off-site location. For example, this would be fully supported on an AIX machine by using the `mksysb` operating system command to produce a valid, bootable tape or DVD image of your present configuration. We describe this situation in 12.5, “Back up and restore the DR550 AIX environment on DVD” on page 501 to complete the disaster recovery concept for your DR550 installation.

### Recovery testing

You must test a recovery solution before you actually need it. A good approach is to create all documents, operating system tapes, special hardware requirements, and installation scripts, and send them to the off-site location labeled as a “Disaster Recovery starter kit.” Then, perform a complete recovery test once a year to ensure that the documents are accurate for recovery and incorporate any changes that were uncovered during your test.

You can find further information about disaster recovery concepts, and especially the DRM, in *IBM Tivoli Storage Management Concepts*, SG24-4877, available at:

<http://www.redbooks.ibm.com/abstracts/sg244877.html>

## 12.4 Using Disaster Recovery Manager on the DR550

In this section, we demonstrate how to use DRM on the DR550 to create a disaster recovery process. Detailed information about DRM in general is available in Chapter 24, “Using Disaster Recovery Manager”, of *IBM Tivoli Storage Manager for AIX Administrator's Guide*, GC32-0768, available at:

<http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/topic/com.ibm.itsmaixn.doc/anragd53.pdf>

These steps are required to prepare DRM for use within the DR550:

1. Create a directory in AIX for DRM instructions.
2. Within that directory, create example files with customer-specific and installation-specific information. The example file names have to follow a certain convention so that they are recognized by the Tivoli Storage Manager server to prepare the plan files:
  - `prefix.RECOVERY.INSTRUCTIONS.GENERAL`
  - `prefix.RECOVERY.INSTRUCTIONS.OFFSITE`
  - `prefix.RECOVERY.INSTRUCTIONS.INSTALL`
  - `prefix.RECOVERY.INSTRUCTIONS.DATABASE`
  - `prefix.RECOVERY.INSTRUCTIONS.STGPOOL`

*prefix* should be replaced by a meaningful word.

3. Create a directory for DRM plan files. In this directory, the Tivoli Storage Manager server will put a plan file based on the example files whenever the **prepare** command is executed on the server.
4. Configure DRM to use the new directories and example files to create plan files.

The following steps provide the detailed procedure to perform the previously described steps:

1. While logged on to AIX with root privileges, create a directory for DRM instructions:

```
mkdir /tsm/drm
mkdir /tsm/drm/instructions
```

2. Change to the new directory and create, with the editor of your choice (for example, vi), the first example file *prefix*.RECOVERY.INSTRUCTIONS.GENERAL. We use DRM as the *prefix* throughout the procedure.

```
cd /tsm/drm/instructions
vi DRM.RECOVERY.INSTRUCTIONS.GENERAL
```

Fill in the general instructions, as shown in Figure 12-3.

Recovery Instructions for Tivoli Storage Manager Server on system DR550

Joe Smith (wk 002-000-1111 hm 002-003-0000): primary system programmer  
Sally Doe (wk 002-000-1112 hm 002-005-0000): primary recovery administrator  
Jane Smith (wk 002-000-1113 hm 002-004-0000): responsible manager

Security Considerations: Joe Smith has the password for the Admin ID admin. If Joe is unavailable, you need to either issue SET AUTHENTICATION OFF or define a new administrative user ID at the replacement Tivoli Storage Manager server console.

Figure 12-3 Example of *DRM.RECOVERY.INSTRUCTIONS.GENERAL*

3. Create the file *DRM.RECOVERY.INSTRUCTIONS.OFFSITE* and fill in instructions about the off-site location, as shown in Figure 12-4.

Our offsite vault location is Ironvault, Safetown, Az. The phone number is 1-800-000-0008.

You need to contact them directly to authorize release of the tapes to the courier. Our courier's name is Fred Harvey. You can contact him at 1-800-444-0000. Because our vault is so far away, be sure to give the courier a list of both the database backup and copy storage pool volumes required. Fred is committed to returning these volumes to us in less than 12 hours.

Figure 12-4 Example of *DRM.RECOVERY.INSTRUCTIONS.OFFSITE*

4. Create the file *DRM.RECOVERY.INSTRUCTIONS.INSTALL* and fill in instructions about the server installation, for example, operating system restore, as shown in Figure 12-5 on page 499.



Most likely you will not need to reinstall the Tivoli Storage Manager server because we use mksysb to back up the rootvg volume group, and the Tivoli Storage Manager server code and configuration files exist in this group. However, if you cannot do a mksysb restore of the base server system, and instead have to start with a fresh AIX build, you may need to add Tivoli Storage Manager server code to that AIX system. The install volume for the Tivoli Storage Manager server is INS001. If that is lost, you will need to contact Copy4You Software, at 1-800-000-0000, and obtain a new copy. Another possibility is the local IBM Branch office at 555-7777.

Figure 12-5 Example of `DRM.RECOVERY.INSTRUCTIONS.INSTALL`

5. Create the file `DRM.RECOVERY.INSTRUCTIONS.DATABASE` and fill in instructions about where to restore the database, as shown in Figure 12-6.

You will need to find replacement disk space for the server database. We have an agreement with Joe Replace that in the event of a disaster, he will provide us with disk space.

Figure 12-6 Example of `DRM.RECOVERY.INSTRUCTIONS.DATABASE`

6. Create the file `DRM.RECOVERY.INSTRUCTIONS.STGPOOL` and fill in instructions about the priority for the storage pools, as shown in Figure 12-7.

Focus on ARCHIVEPOOL. This is the most important storage pool.

Figure 12-7 Example of `DRM.RECOVERY.INSTRUCTIONS.STGPOOL`

7. Create a directory for the plan files:

```
mkdir /tsm/drm/planfiles
```

8. Configure DRM to use the newly created directories and file structures by issuing the following commands:

- SET `DRMPLANPREFIX` to specify the RPF prefix
- SET `DRMINSTRPREFIX` to specify the user instruction file prefix
- SET `DRMDBBACKUPEXPIREDAYS` to define the database backup expiration
- SET `DRMFILEPROCESS` to specify file processing

The parameter `DRMDBBACKUPEXPIREDAYS` must be set to 3 Days to meet the preconfigured database backup strategy of the DR550. If desired, you can set the period of how long to keep database backups to a different value.

The `DRMFILEPROCESS` parameter must be set to Yes because the preconfiguration schedule writes database backups onto disk space, and this method will be maintained.

**Important:** If DRM is used, we strongly recommend that you control the database backup expiration with the SET `DRMDBBACKUPEXPIREDAYS` command instead of the preconfigured DELETE VOLHISTORY command.

9. With the use of DRM, two entities are now actively controlling the expiration of database backup volumes. The DELETE VOLHISTORY command removes the Tivoli Storage Manager record of the volume. This can cause volumes to be lost that were managed by the DRM. The recommended way to manage the automatic expiration of DRM database backup volumes is by specifying the DB Backup Series Expiration Days parameter, which has been done in the previous step. To turn off the automated DELETE VOLHISTORY, you must modify the associated schedule DELVOLH. You can delete it, but to set it to inactive is sufficient. To do this, issue the following command at the command-line interface:

```
update schedule delvolh active=no
```

DRM is now configured and ready for use. In the next sections, we describe how to query and modify the state of volumes with DRM commands.

### Querying DRM for volume states

To query the state of a volume from a DRM perspective, simply type **query drmedia** in the command-line interface. Figure 12-8 shows an example output of a DRM query.

Operation Results			
Volume Name	State	Last Update Date/T ime	Automated LibNa me
-----	-----	-----	-----
/tsmdbbkup/0072 2405.DBB	Mountable	11/17/04 13:13:24	
/tsmdbbkup/9755 3394.DBB	Mountable	10/12/04 10:56:34	
JW1042	Mountable	11/15/04 17:34:27	3584LIB
JW1043	Mountable	11/15/04 18:12:45	3584LIB
JW1044	Mountable	11/15/04 17:14:23	3584LIB
JW1045	Mountable	11/15/04 17:53:17	3584LIB
JW1046	Mountable	11/15/04 19:34:21	3584LIB

Figure 12-8 Example output of a DRM query

The output shows two database backups on disk and five 3592 cartridges, all with a state of Mountable. This means that all volumes can be accessed by Tivoli Storage Manager. Other possible states of a volume can be NotMountable, Courier, Vault, VaultRetrieve, CourierRetrieve, and Remote. The following list describes these states:

- ▶ **Mountable:** Volumes in this state contain valid data and are accessible for on-site processing.
- ▶ **NotMountable:** Volumes in this state are on-site, contain valid data, and inaccessible for on-site processing.
- ▶ **Courier:** Volumes in this state are being moved to an off-site location.
- ▶ **Vault:** Volumes in this state are off-site, contain valid data, and are inaccessible for on-site processing.
- ▶ **VaultRetrieve:** Volumes in this state are located at the off-site vault and do not contain valid data.
- ▶ **CourierRetrieve:** Volumes in this state are being moved back to the on-site location.
- ▶ **Remote:** Volumes in this state contain valid data and are located at the off-site remote server.

### Modifying the state of a volume

Use the **move drmedia** command to move a volume from a given state to a new state, for example, to move a tape cartridge out of the library into a vault. Use it to track database backup and copy storage pool volumes that are to be moved off-site and to identify the expired or empty volumes that are to be moved on-site. You can change volumes through each state, or you can use the **TOSTATE** parameter and skip states to simplify the movements. DRM takes care not only to track the state, or in other words, the current location of a volume, but also to initiate the required checkout command to physically move the cartridge to the I/O station of a library if the appropriate parameter **remove=yes** is specified.

In the following example, we remove a tape cartridge from the library and update the state directly to reflect the Vault state.

Initiate the movement of the media with, for example, **move drm JW1042 tostate=vault remove=yes**. A background process is started and carried out by the server. Once completed, you can verify the results in the activity log with **query actlog**. The window shown in Figure 12-9 shows the output of a subsequent **query drmedia** command after the move process completed successfully.

Operation Results			
Volume Name	State	Last Update Date/Time	Automated LibName
-----	-----	-----	-----
/tsmdbbkup/0072 2405.DBB	Mountable	11/19/04 16:48:54	
/tsmdbbkup/9755 3394.DBB	Mountable	10/12/04 10:56:34	
JW1042	Vault	11/17/04 17:01:22	3584LIB
JW1043	Mountable	11/15/04 18:12:45	3584LIB
JW1044	Mountable	11/15/04 17:14:23	3584LIB
JW1045	Mountable	11/15/04 17:53:17	3584LIB
JW1046	Mountable	11/15/04 19:34:21	3584LIB

Figure 12-9 Example output of a DRM query after a DRM move command

These examples are only a brief description of a few of the DRM capabilities, but they provide an idea of how you can use DRM to efficiently manage and track off-site media for a professional disaster recovery process. For further details about DRM, see *IBM Tivoli Storage Manager for AIX Administrator's Guide*, SC32-0117, available at:

<http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/topic/com.ibm.itsmaixn.doc/anragd54.pdf>

## 12.5 Back up and restore the DR550 AIX environment on DVD

We have described and discussed the options to protect archived data that, in the case of the DR550, reside on the RAID-protected DS4000 Storage Server and EXP Expansion Units. These options basically consist of standard Tivoli Storage Manager methods and tape techniques to protect against disasters. You can also use the same options to protect the Tivoli Storage Manager database, a critical and essential component of the DR550 solution. We describe how to back up that database not only to disk, but also to removable tape media. For the capability to fully restore the DR550, you also need a backup or image of the AIX installation in the AIX volume group *rootvg* and the Tivoli Storage Manager script files

required for HACMP (these are kept on a shared drive in the volume group *TSMApps* (see Figure 12-10)).

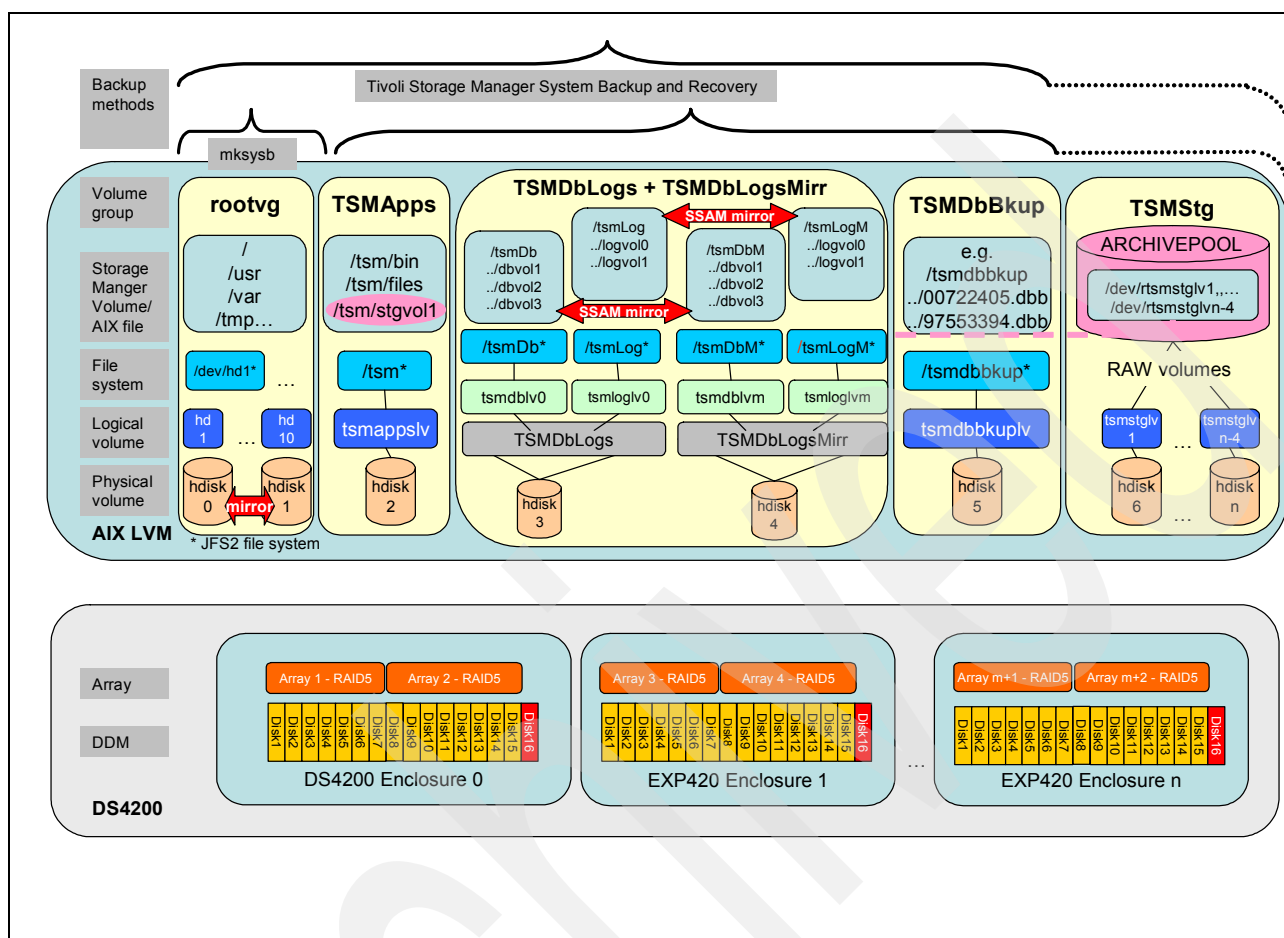


Figure 12-10 AIX volume group and DS4700 logical drive layout in the DR550 V4.0

The rootvg is mirrored between the two internal SCSI disk drives in the p5 520 server (hdisk0 and hdisk1); however, this does not provide sufficient protection for a system failure. You can accomplish backing up the rootvg, or in other words, creating a bootable image on different media, in various ways.

One way is to use the **mksysb** command to write onto an attached tape drive and then store the cartridge in a safe place. Or you can use enterprise strength products, such as IBM Tivoli Storage Manager for System Backup and Recovery, to create remote bootable images through the network.

The shared volume group *TSMApps* (needed for HACMP) resides on a DS4000 and thus is RAID-protected. Nevertheless, a disaster can destroy parts of the rack or the entire rack, making this volume group unusable. Using the AIX program **savevg**, you can create backups of volume groups on favored removable media, such as tape or DVD, providing the capability to store the backup off-site.

Tivoli Storage Manager for System Backup and Recovery (also called Tivoli Storage Manager SysBack™) is incorporating, among many additional functions, the advantages of the **mksysb** and **savevg** commands. With Tivoli Storage Manager SysBack, you can accomplish a backup image of rootvg, plus one or more additional volume groups, in one operation. Additionally, the restore features of Tivoli Storage Manager SysBack offer flexible options for how and

where to restore the data. Tivoli Storage Manager SysBack is part of the Tivoli Storage Manager product portfolio and requires a license. For further details about Tivoli Storage Manager System Backup and Recovery, refer to the following Web site:

<http://publib.boulder.ibm.com/infocenter/tivihelp/index.jsp?topic=/com.ibm.itsmsbr.doc/smsbr.html>

We explain an easy and convenient method based on **mksysb** and **savevg** to create a restorable system image of the rootvg and an image of the TSMApps volume group onto DVD-RAM media with Universal Disk Format (UDF).

**Hint:** Universal Disk Format is a feature introduced in AIX 5L Version 5.2.0. UDF uses less disk space to create the backup, because it writes the backup directly to the DVD. You will still have to create the backup file on the system, unless it is user-specified, and then it will be copied directly to the media. This saves space because there does not have to be a CD/DVD image created prior to writing to the media.

It is not necessarily required to generate images of rootvg and TSMApps on a frequent basis, as it is to generate images of the Tivoli Storage Manager database, for example. But, we advise you to create a backup occasionally and store the offline media in a safe place, especially after modifications to the AIX operating system, the Tivoli Storage Manager server application, device drivers, or just minor, but probably important changes, to configuration files (Tivoli Storage Manager, SNMP, SSH, and so on).

The following procedures describe how to create an image with **mksysb** on the AIX volume group rootvg that contains the operating system and store it on DVD-RAM media using the internal DVD drive that ships with each p5 520 server included in the DR550. You can apply this procedure to a single-node machine or to both engines of a dual-node configuration. A second procedure goes through similar steps to generate an image of the volume group TSMApps also onto DVD-RAM media using **savevg**.

## Creating an AIX system backup on a DVD-RAM

In this section, we describe how to create a root volume group backup on DVD-RAM with Universal Disk Format.

Prerequisites:

- ▶ Ensure that all applications have been stopped from accessing the DR550.
- ▶ On a dual-node system, ensure that HACMP cluster services have been stopped on both nodes. Refer to 4.2, “Starting and stopping IBM System Storage Archive Manager” on page 148.
- ▶ Ensure that the System Storage Archive Manager server is not running. Refer to 4.1, “Starting and stopping HACMP cluster services” on page 138.

To create a backup of the root volume group on DVD-RAM with UDF, use SMIT as follows:

1. Insert a blank DVD-RAM media into the DVD drive.
2. Enter the **smitty mkdvd** fast path. The system asks whether you are using an existing mksysb image. Scroll down until option 2 NO is highlighted and press Enter.
3. On the next window, scroll down to option 2 UDF (Universal Disk Format) and press Enter.
4. In the DVD-RAM Device field, press F4 to call up the selection list. A line such as /dev/cd0 IDE DVD-RAM Drive should be highlighted. Press Enter to apply this selection.

5. Fill in the fields, as shown in Figure 12-11, and press Enter.

Back Up This System to UDF DVD		
Type or select values in entry fields. Press Enter AFTER making all desired changes.		
[TOP]	[Entry Fields]	
DVD-RAM Device	[/dev/cd0]	+
mksysb creation options:		
Create map files?	yes	+
Exclude files?	no	+
Disable software packing of backup?	no	+
File system to store mksysb image (If blank, the file system will be created for you.)	[ ]	/
If file systems are being created:		
Volume Group for created file systems	[rootvg]	+
Advanced Customization Options:		
Do you want the DVD to be bootable?	yes	+
Install bundle file	[ ]	/
File with list of packages to copy to DVD	[ ]	/
Location of packages to copy to DVD	[ ]	+/
Customization script	[ ]	/
User supplied bosinst.data file	[ ]	/
Debug output?	no	+
User supplied image.data file	[ ]	/
[BOTTOM]		
F1=Help	F2=Refresh	F3=Cancel
F5=Reset	F6=Command	F7=Edit
F9=Shell	F10=Exit	Enter=Do
		F4=List
		F8=Image

Figure 12-11 Creating an AIX system backup to DVD-RAM

Use the following values:

- Select yes in the Create map files field. The File System to store mksysb image field is left blank. The **mkcd** command will then create the file system and remove it when the command completes.
  - The DVD has to be created as bootable. To do so, leave the value of yes unchanged (If you select no, you must boot from a product CD at the same version.release.maintenance level and then select to install the system backup from the system backup DVD).
  - The default for debugging for the **mkcd** command is set to no. It can be set to yes if problems occur with the backup process. The debug output goes to the smit.log file.
6. Wait for the backup process to complete and check for an OK status. The process takes about 1.5 hours to complete, depending on the size of the rootvg volume group. Figure 12-12 on page 505 shows the output of a successful system backup on a DVD-RAM.

```
COMMAND STATUS

Command: OK stdout: yes stderr: no

Before command completion, additional instructions may appear below.

[TOP]
Initializing mkcd log: /var/adm/ras/mkcd.log...
Verifying command parameters...
Creating image.data file...
Creating temporary file system: /mkcd/mksysb_image...
Creating mksysb image...

Creating list of files to back up.
Backing up 39415 files.....
39415 of 39415 files (100%)
0512-038 mksysb: Backup Completed Successfully.
Populating the CD or DVD file system...
Copying backup to the CD or DVD file system...
Building chrp boot image...
Removing temporary file system: /mkcd/mksysb_image...

[BOTTOM]

F1=Help F2=Refresh F3=Cancel F6=Command
F8=Image F9=Shell F10=Exit /=Find
n=Find Next
```

Figure 12-12 Successful AIX system backup output

7. Remove the DVD-RAM media from the device and label it accordingly, including the type of backup, the time stamp, and the engine from which it was taken. Store the media in a safe place.
8. Resume cluster operations if desired (refer to 4.2, “Starting and stopping IBM System Storage Archive Manager” on page 148).

### Restoring AIX from a bootable DVD-RAM system backup

In this section, we describe how to restore a root volume group backup from bootable DVD to a DR550 engine.

Prerequisites:

- Ensure that the engine is powered down.
- Have a restorable rootvg image backup DVD available.

To restore AIX from a bootable DVD-RAM system backup, complete the following steps:

1. Insert the DVD media into the DVD drive of the engine.
2. Log on to the HMC application and activate the engine (partition) by right-clicking the partition name. Select **Activate** from the menu and select the option for opening a console session. Click **Advanced** and select **SMS** from the Boot Mode list. Click **OK** twice to initiate the activation of the partition. A VTERM terminal screen opens, displaying the startup process.

3. Alternatively, you can use a terminal or notebook instead of the HMC. If you use the HMC, skip this step. If not, connect the terminal or a workstation with a terminal emulation program (for example, Hyperterm on Windows, settings 19200, 8N1, Xon/Xoff) to the front serial port of the engine that is going to be restored (using the provided cable RJ45-to-9pin and a null modem extension cable). Complete the following steps:
  - a. With the p5 520 connected to power, but with AIX shutdown, you see the Service Processor Menu, indicating that you have successfully established the serial connection. You are not required to log on to the Service Processor at this point. See Figure 12-13.

```

Welcome
Machine type-model: 9131-52A
Serial number: 10E89ED
Date: 2006-11-3
Time: 7:02:13
Service Processor: Primary
User ID:

```

Figure 12-13 Service Processor Login prompt

- b. Power on the p5 520 server by pressing the white button on the front of the server. The system starts its power-on self tests, displaying various 8-digit codes on the front display.
4. Wait a few minutes until IBM logos start scrolling over the window and a menu displays, as shown in Figure 12-14. At the menu prompt, type 1 to enter the System Management Services (SMS) menu. If you missed that point, and AIX is starting up, shut down AIX again after it has opened with **shutdown -F** and restart the procedure.

```

IBM IBM
IBM IBM

 1 = SMS Menu 5 = Default Boot List
 8 = Open Firmware Prompt 6 = Stored Boot List

memory keyboard network scsi

```

Figure 12-14 IBM welcome screen

When the system prompts you to type the password for the Service Processor's admin account, type admin.

5. On the SMS main menu, type 5 and press Enter to choose Select Boot Options.
6. On the next menu, type 1 and press Enter to choose Select Install/Boot Device.
7. On the next menu, type 3 and press Enter to select device type CD/DVD.
8. On the next menu, type 4 and press Enter to select media type IDE.
9. On the next menu, type 1 and press Enter to select device IDE CD-ROM.
10. On the next menu, type 2 and press Enter to select task Normal Mode Boot.
11. On the next window, type 1 for Yes to exit the SMS menus. A Starting Software Please Wait... statement displays, surrounded by IBM logos, followed by a Welcome to AIX message.
12. When prompted, type 2 and press Enter to use this terminal as the system console.
13. When prompted, type 1 and press Enter to use English during the installation.



14. At the next prompt, disregard the warning message and continue by typing 1 and then pressing Enter to select Continue with Install.
15. On the Welcome to Base Operating System Installation and Maintenance window, type 2 and press Enter to select Change/Show Installation Settings and Install.
16. On the System Backup Installation and Settings screen, leave the default settings (hdisk0, yes, no, yes, yes) unchanged and type 0 to select Install with the settings listed above.
17. A window similar to that shown in Figure 12-15 indicates the progress of the restore process.




Figure 12-15 Progress of restore process screen

18. After the restore process has completed successfully (approximately one hour from the beginning of this procedure), the system automatically reboots and is ready for operation.
19. Now you have to reimport the logical volumes from the DS4000:
 

```
/usr/sbin/importvg -V'81' -y TSMApps hdisk2
/usr/sbin/importvg -V'82' -y TSMDBLogs hdisk3
/usr/sbin/importvg -V'84' -y TSMDBLogsMirr hdisk4
/usr/sbin/importvg -V'83' -y TSMDBBkup hdisk5
/usr/sbin/importvg -V'85' -y TSMStg hdisk6
```
20. Remove the DVD media from the drive and resume cluster operations if desired. Refer to 4.1, “Starting and stopping HACMP cluster services” on page 138.

## Creating a backup of the TSMApps volume group on DVD-RAM

In this section, we describe how to create a TSMApps volume group backup on DVD-RAM with Universal Disk Format.

### Prerequisites:

- Ensure that all applications have been stopped from accessing the DR550.
- On a dual-node system, ensure that the HACMP cluster services have been stopped on both nodes. Refer to 4.1, “Starting and stopping HACMP cluster services” on page 138.
- Ensure that the Tivoli Storage Manager server is not running. Refer to 4.2, “Starting and stopping IBM System Storage Archive Manager” on page 148.

To create a TSMApps volume group backup on DVD-RAM with UDF, use SMIT as follows:

1. Insert a blank DVD-RAM media into the DVD drive.
2. Enter the **smitty savevgdvd** fast path. The system asks whether you are using an existing mksysb image. Scroll down until option 2 NO is highlighted and press Enter.
3. On the next screen, scroll down to option 2 UDF (Universal Disk Format) and press Enter.

4. In the DVD-RAM Device field, press F4 to call up the selection list. A line such as /dev/cd0 IDE DVD-RAM Drive should be highlighted. Press Enter to apply this selection.
5. In the Volume Group to back up field, press F4 to call up the selection list. Scroll down to highlight TSMApps and press Enter.
6. Fill in the fields, as shown in Figure 12-16, and press Enter.

Back Up a Volume Group to UDF DVD			
Type or select values in entry fields. Press Enter AFTER making all desired changes.			
DVD-RAM Device	[Entry Fields]		
	[/dev/cd0]	+	
* Volume Group to back up	[TSMApps]	+	
savevg creation options:			
Create map files?	yes	+	
Exclude files?	no	+	
Disable software packing of backup?	no	+	
File system to store savevg image (If blank, the file system will be created for you.)	[ ]	/	
If file systems are being created:			
Volume Group for created file systems	[rootvg]	+	
Advanced Customization Options:			
Debug output?	no	+	
F1=Help	F2=Refresh	F3=Cancel	F4=List
F5=Reset	F6=Command	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

Figure 12-16 Creating a Tivoli Storage Manager volume group backup to DVD

Use the following values:

- Select yes in the Create map files field. The File System to store mksysb image field is left blank. The **mkcd** command will then create the file system and remove it when the command completes.
  - The default for debugging for the **mkcd** command is set to no. It can be set to yes if problems occur with the backup process. The debug output goes to the smit.log file.
7. Wait for the backup process to complete and check for an OK status. The process takes a few minutes to complete, depending on the size of the TSMApps volume group. Figure 12-17 on page 509 shows the output of a successful backup on DVD.

```
COMMAND STATUS

Command: OK stdout: yes stderr: no

Before command completion, additional instructions may appear below.

Initializing mkcd log: /var/adm/ras/mkcd.log...
Verifying command parameters...
Creating information file for volume group TSMApps.
Creating temporary file system: /mkcd/mksysb_image...
Creating savevg image...

Creating list of files to back up.
Backing up 29 files
29 of 29 files (100%)
0512-038 savevg: Backup Completed Successfully.
Copying backup to the CD or DVD file system...
Removing temporary file system: /mkcd/mksysb_image...

F1=Help F2=Refresh F3=Cancel F6=Command
F8=Image F9=Shell F10=Exit /=Find
n=Find Next
```

Figure 12-17 Successful TSMApps volume group backup output

8. Remove the DVD-RAM media from the device, label it accordingly, including the type of backup, the time stamp, and the engine from which it was taken. Store the media in a safe place.
9. Resume cluster operations if desired. Refer to 4.1, “Starting and stopping HACMP cluster services” on page 138.

## Restoring the TSMApps volume group from DVD-RAM

In this section, we describe how to restore a backup of the volume group TSMApps from DVD to a DR550 engine.

### Prerequisites:

- ▶ Ensure that all applications have been stopped from accessing the DR550.
- ▶ On a dual-node system, ensure that the HACMP cluster services have been stopped on both nodes. Refer to 4.1, “Starting and stopping HACMP cluster services” on page 138.
- ▶ Ensure that the SSAM server is not running. Refer to 4.2, “Starting and stopping IBM System Storage Archive Manager” on page 148.
- ▶ Ensure that you have a restorable TSMAppss image backup DVD available.
- ▶ The volume group TSMApps must not exist on the system (lsvg).
- ▶ The hdisk2 has to be available and must not be assigned to a volume group (use **lsdev -Cc disk, lsvg** to confirm).

To restore the TSMApps volume group from DVD-RAM, complete the following steps:

1. Insert the DVD media into the DVD drive of the engine.
2. Enter the **smitty restvg** fast path. In the Restore DEVICE or FILE field, press F4 to call up the selection list. A line such as /dev/cd0 IDE DVD-RAM Drive should be highlighted. Press Enter to apply it.

3. Fill in the fields, as shown in Figure 12-18, and press Enter. Select yes when prompted to continue.

Remake a Volume Group		
Type or select values in entry fields. Press Enter AFTER making all desired changes.		
	[Entry Fields]	
* Restore DEVICE or FILE	[/dev/cd0]	+/
SHRINK the filesystems?	no	+
Recreate logical volumes and filesystems only?	no	
VOLUME names	[hdisk2]	+
(Leave blank to use the PHYSICAL VOLUMES listed in the vgname.data file in the backup image)		
Use existing MAP files?	yes	+
Physical partition SIZE in megabytes	[ ]	+#
(Leave blank to have the SIZE determined based on disk size)		
Number of BLOCKS to read in a single input	[ ]	#
(Leave blank to use a system default)		
Alternate vg.data file	[ ]	/
(Leave blank to use vg.data stored in backup image)		
F1=Help	F2=Refresh	F3=Cancel
F5=Reset	F6=Command	F7=Edit
F9=Shell	F10=Exit	Enter=Do
	F4=List	F8=Image

Figure 12-18 Restoring volume group TSMApps from DVD

4. A window opens showing the status of the restore process. Wait for the process to complete and check that the status is OK. A successful restore will display an output similar to that shown in Figure 12-19 on page 511.

```

Command: OK stdout: yes stderr: no

Before command completion, additional instructions may appear below.

[TOP]

Will create the Volume Group: TSMApps
Target Disks: hdisk2
Allocation Policy:
 Shrink Filesystems: no
 Preserve Physical Partitions for each Logical Volume: yes

/usr/sbin/mkvg: This concurrent capable volume group must be varied on manually.
TSMApps
loglv01
tsmappslv
New volume on /tmp/vgdata.295020/cdmount/usr/sys/inst.images/savevg_image:
Cluster size is 51200 bytes (100 blocks).
The volume number is 1.
The backup date is: Wed Nov 17 19:11:44 MST 2004
Files are backed up by name.
The user is root.
The number of restored files is 29.
x 170 ./tmp/vgdata/vgdata.files368890
x 170 ./tmp/vgdata/vgdata.files
x 2333 ./tmp/vgdata/TSMApps/filesystems
x 0 .
x 2281 ./tmp/vgdata/TSMApps/TSMApps.data
[MORE...26]

F1=Help F2=Refresh F3=Cancel F6=Command
F8=Image F9=Shell F10=Exit /=Find
n=Find Next

```

Figure 12-19 Output of successful restore process

5. Leave the smit menus by pressing F10 and verify that the file system of the restored volume group is accessible, for example, with the **ls -l /tsm** command.
6. Unmount and vary the volume group offline so that HACMP can start properly later. To accomplish these tasks, type **umount /tsm** and then **varyoffvg TSMApps**.
7. Remove the DVD media from the drive and resume cluster operations if desired. Refer to 4.1, “Starting and stopping HACMP cluster services” on page 138.

## 12.6 File System Gateway (FSG) backup and restore

In this section, we first provide an overview of the mechanisms included with the FSG software that enable the FSG to back up and restore information it maintains about the archived data (metadata). However, this procedure does not create a complete backup of the information needed to restore the FSG. For that reason, we discuss a procedure to back up and restore the FSG configuration information and other software components that you could have installed and configured on your FSG system. That would be the case with the IBM Director Agent, local user management information (user ID, passwords, and permissions), or the LDAP Client for Linux if you were using LDAP services as a centralized user management method.

## 12.6.1 Metadata backup and restore

The FSG metadata includes information describing protection attributes and retention periods that characterize the data managed by the FSG. The metadata also includes reference information about the archived data. The metadata is maintained locally at the FSG (on the local disk). Automatic backups of the metadata are taken at regular intervals and saved into the DR550 storage pools. In case the metadata was lost or corrupted at the FSG, the FSG software has the ability to automatically restore the last metadata backup from the DR550 storage pool. A restore of the metadata would also happen automatically if you had to reinstall the FSG from scratch (both FSG nodes in an HA configuration). If you reinstall both nodes of a high availability cluster, bring only one node (main, for example) online until the metadata has been fully restored.

**Important:** To make sure that the automatic backup/restore works properly on your FSG, the procedure should be tested right after completing your initial FSG installation. To verify the functionality and proper communication with the DR550 storage, issue the following command at the FSG:

```
/usr/local/drg/forcebackup
```

Note that the restore operation might require you to mount tapes if some data was already migrated to tape. It is therefore necessary to watch the SSAM activity log to ensure that all necessary tapes are accessible. This is because the FSG restore procedure reads some attributes from the archived data.

**Important:** During an automatic restore of the metadata, the FSG services are not available. The restore process can take a few hours. Wait until the restore is finished and the FSG processes have come up. Check by issuing the command `crm_mon -i 15`.

## 12.6.2 Back up and restore the FSG configuration

The metadata backup and restore does not include additional configuration data specific to your FSG installation. This additional data includes information such as the NFS or CIFS shares configuration files, network configuration information, or any additional software that you might have installed on your FSG.

Because there is no supported procedure at this time to back up and restore the whole FSG system, we recommend that you save a separate copy of the various configuration files that you might have changed. In 12.6.3, “Full FSG backup and restore” on page 513, we describe an as-is (that is, not formally supported) procedure to do a full backup of the FSG system.

**Important:** We recommend that you document every configuration changes you make at the File System Gateway not only at installation time, but also any changes made during maintenance or general administration of the system.

Table 12-1 on page 513 gives examples of configuration files that will be changed during the FSG configuration at your site. This is not an exhaustive list and you might have more or less configuration files impacted depending upon which additional functions or components you are using in your environment.

Table 12-1 Some of the important config files

Type of configuration	File
NFS shares	/etc/exports
CIFS shares	/etc/samba/smb.conf
SNMP config	/etc/snmp/snmpd.conf /usr/shares/snmp/snmpd.conf
network configuration	/etc/resolv.conf /etc/hosts /etc/sysconfig/network/* ...
LDAP configuration	...
Active Directory configuration	...
drg config	/var/local/bundle-import/current/* (needs to restore in the .../import/* folder)
password file	/etc/passwd
samba (cifs) password file	/etc/samba/smbpasswd

**Important:** Do not forget to enable the autostart of these services in the different runlevels.

### 12.6.3 Full FSG backup and restore

In the process of writing this book, we tried many different products to create a full (easily restorable) backup of the FSG. We advise against using backup products (commercial or shareware) with the FSG. The main reason is that in most cases you would need to install additional SLES10 filesets, such as kernel sources and compilers, and recompile the OS kernel. Because the FSG is sold as a preconfigured component of the DR550 solution, you will no longer have a supported solution if you do this.

We explain in this section a procedure that will let you back up and restore a completely customized FSG by simply using the SLES 10 CD/DVD in recovery mode.

**Important:** Back up and restore the FSG is not officially supported in this release. The official procedure is to reinstall every component and redo all the configuration steps (see 6.7.8, “DR550 File System Gateway replacement” on page 272).

#### Provide a NFS share for the backup

We choose to create the backup on a NFS share to make the restore more efficient by having the FSG backup image available in the network.

To do so, you need to provide the necessary disk space on an external NSF file system to hold the full backup. The size of the backup file for one FSG node will be approximately 1.5 GB. You can create an NFS share for the backup or use an existing NFS server or NFS share.

### Prepare the NSF server

In our example, we used a SLES10 Linux server (different than the FSG server, of course) to provide the necessary backup space. We show the steps to execute to create and export an NFS share with SLES10 Linux below (the steps and commands would obviously be different if your NFS server is implemented on a different OS platform). The following commands depend on the operating system on your NFS server:

1. Run `mkdir /<nfs-sharename>`.
2. Run `vi /etc/exports`.
3. Add the line `/<nfs-sharename> *(rw, sync, anonuid=0)`.
4. Start or restart the NFS server. You can verify that the share is active with the command `exportfs`.

### Full FSG backup

The following steps must be executed on every FSG server that you want to back up:

1. Boot from the SUSE Linux SLES10 CD1/DVD and choose **Rescue System** at the boot menu. The system will continue the boot from CD/DVD. The boot is finished when the login prompt displays.
2. Log in as root; no password is required.
3. Enter `cat /proc/partitions` to check that all disks of the FSG are available. Figure 12-20 shows the list of disks that should be available.

major	minor	#blocks	name
8	0	71577600	sda
8	1	4200966	sda1
8	2	2104515	sda2
8	3	65264062	sda3
8	16	585728000	sdb
8	17	585713803	sdb1

Figure 12-20 FSG partitions

4. Configure one network adapter. Because the (discovery) sequence of the network interfaces in the rescue mode might be different than when booting in the normal mode, you have to discover which network interface is working. Execute the following commands using the IP address from the FSG bond0 interface:

```
ifconfig eth0 inet <FSG-ip-address> netmask 255.255.255.0
route add default gw <default-network-gateway ip-address>
ping <gw ip-address>
```

(It might take 30 seconds to get a response to the ping.)

**Note:** If the ping failed after 30 seconds, run the command `ifdown eth0`, try the procedure again starting with the `ifconfig` command, and use Ethernet interface `eth2` instead of `eth0`.



5. Because you booted from the CD/DVD, the (FSG) hard disks file systems were not mounted. They must be manually mounted. You need to mount the FSG file system to back up and the file system where you want to store the backup file. Enter the following commands:

```
mount <nfs-server ip address>:/<nfs-sharename> /media
```

(Here you must enter as nfs-sharename the one you had defined on the NFS server in “Prepare the NSF server” on page 514.)

```
mount /dev/sda1 /mnt
```

(This mounts the disk of the Linux / root file system to /mnt.)

```
mount /dev/sda3 /mnt/var/local
```

(This mounts the disk of the Linux /var/local file system to /mnt.)

Check that all three file systems are mounted correctly. Enter the **df -h** command.

6. Now you can back up all the data to the NFS shared directory with the **tar** command:

```
cd /mnt
```

```
tar czvf /media/backup_main.tgz *
```

(This backs up all files and will take approximately 10 minutes.)

7. Next you need to save the partition tables of the FSG disks to the NFS shared directory:

```
dd if=/dev/sda of=/media/sda_pt.bak bs=512 count=1
```

```
dd if=/dev/sdb of=/media/sdb_pt.bak bs=512 count=1
```

8. Once all the files have been successfully backed up, remove the SUSE Linux CD/DVD from the drive and reboot the FSG. If you have a high availability FSG cluster, you need to do the backup from both nodes!

## Full FSG restore

In case of a major failure, you may have to completely reinstall the FSG. The following steps show you how to install SLES10 and then restore from a full backup:

1. Reboot the FSG and press Ctrl-A during the initialization of the RAID controller to start the RAID adapter utility.
2. Select **Array Configuration Utility → Initialize Drives**, choose all drives, and initialize.
3. Recreate the arrays: The two 73 GB disks will be configured as RAID 1, and the other four 300 GB disks as a RAID 6 array. Assign the array names OS for the RAID 1 array and Cache for the RAID 6 array, because these are the original names used for the FSGs.
4. Press ESC to leave the utility and insert the SLES10 CD1 or DVD. The machine will boot from the SUSE SLES10 CD/DVD.
5. In the SLES10 boot menu, choose Rescue System.
6. Log in as root; no password is required.
7. Enter **cat /proc/partitions** to display the current server partitions.
8. Because the (discovery) sequence of the network interfaces in the rescue mode might be different than when booting in the normal mode, you have to discover which network interface is working. Execute the following commands using the IP address from the FSG bond0 interface:

```
ifconfig eth0 inet <FSG-ip-address> netmask 255.255.255.0
```

```
route add default gw <default-network-gateway ip-address>
```

```
ping <gw ip-address>
```

(It might take 30 seconds to get a response to the ping.)

**Note:** If the ping failed after 30 seconds, run the command **ifdown eth0**, try the procedure again starting with the **ifconfig** command, and use Ethernet interface eth2 instead of eth0.

9. If the ping was successful you can start to mount the NFS file system on which the backup file resides:

```
mount <nfs-server ip address>:/<nfs-sharename> /media
```

10. Next, restore the partition tables to the disks:

```
dd if=/media/sda_pt.bak of=/dev/sda bs=512 count=1
dd if=/media/sdb_pt.bak of=/dev/sdb bs=512 count=1
```

11. Reload the partition tables:

```
blockdev --rereadpt /dev/sda
blockdev --rereadpt /dev/sdb
```

12. Check that you can see the partitions by running the command **cat /proc/partitions**.

13. Next, create the file systems:

```
mke2fs -j /dev/sda1
mkswap /dev/sda2
mke2fs -j /dev/sda3
```

14. Mount these file systems:

```
mount /dev/sda1 /mnt
mkdir -p /mnt/var/local
```

15. To check that the mount points are correct, use the command **df -h**.

16. Now you can begin restoring the data:

```
cd /media
tar -x -z -C /mnt -v -f backup_main.tgz
```

(This will take approximately 15 minutes.)

17. Reinstall the bootloader:

```
chroot /mnt
mount /proc
mknod /dev/sda b 8 0
mknod /dev/sda1 b 8 1
mknod /dev/sda2 b 8 2
mknod /dev/sda3 b 8 3
grub-install --recheck /dev/sda
```

18. Finish the restore with the **exit** command, remove the SUSE Linux CD/DVD, and reboot the system.

19. At this point the Linux OS has been restored and now you need to prepare for restoring the second hard disk, which is the cache disk for the FSG. Issue the following command:

```
yast2 disk &
```

This opens the disk Expert Partitioner window, as shown in Figure 12-21 on page 517.

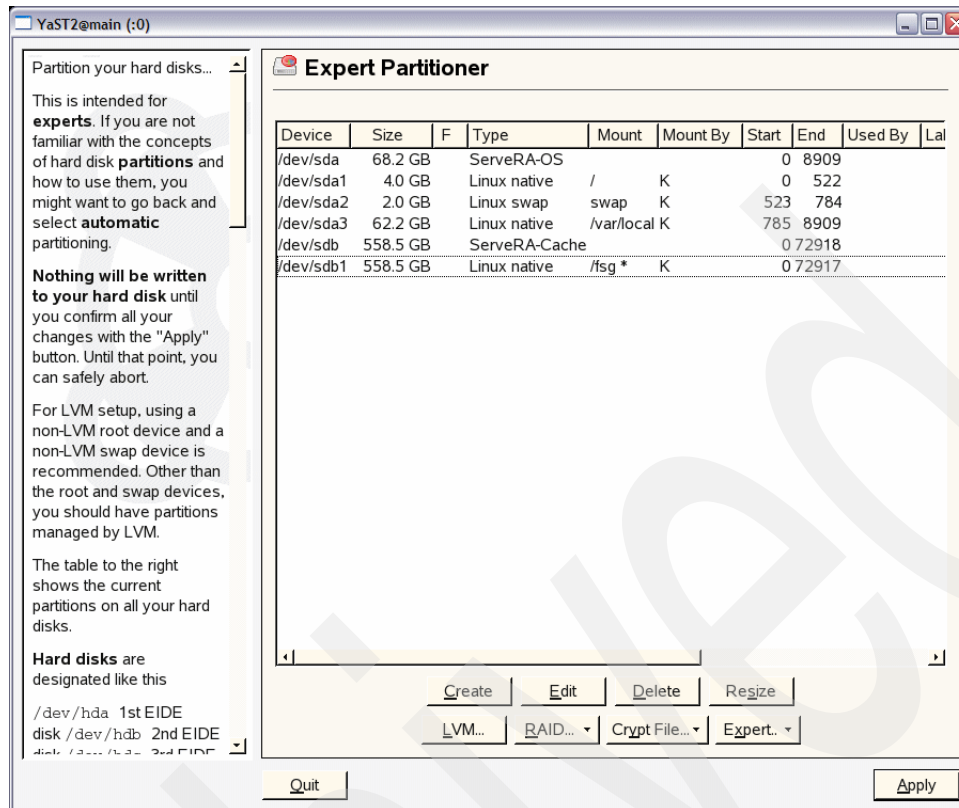


Figure 12-21 Expert Partitioner

20. Select the **/dev/sdb1 (/fsg)** partition and click **Edit**. This opens the Edit Existing Partition window, as shown in Figure 12-22.

21. Select the **Format** button and select **XFS** as the file system type.

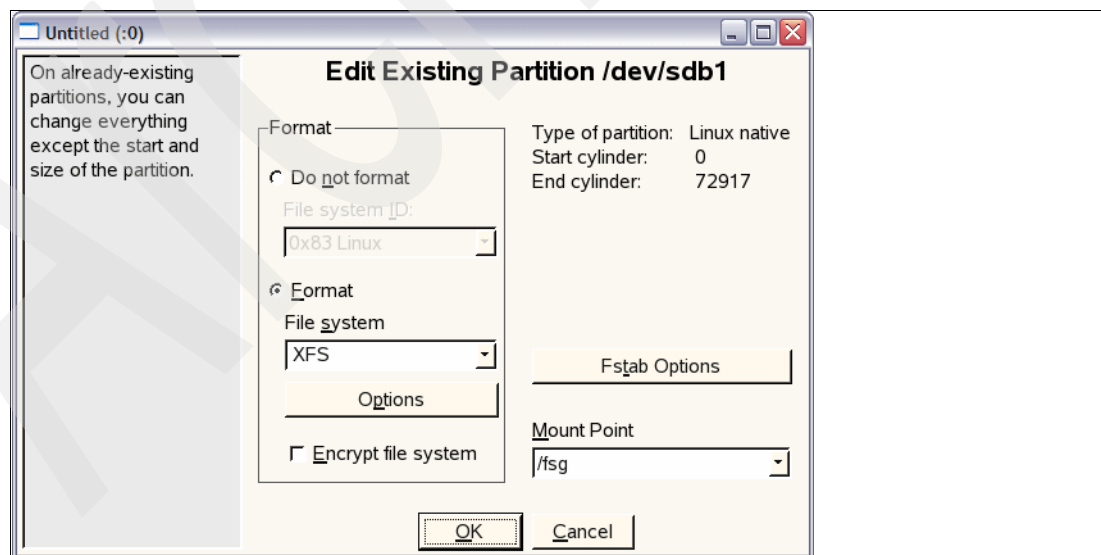


Figure 12-22 Format the /fsg Cache Partition

22. Click **Options** and change the values for the file system options, as shown in Figure 12-23. Click **OK** when finished.

Figure 12-23 File System Options for /fast

23. Click **Fstab Options** (see Figure 12-22 on page 517) and select the button **Mount by device name**. Enter `dmapi,mtp=/fsg` into the Arbitrary options field and click **OK**, as shown in Figure 12-24.

Figure 12-24 Fstab Options for /fsg

24. Click **Apply**.  
25. Click **Finish**.

## Clean up the FSG

As a last step, you need to delete all files in some directories. This cleans up the FSG metadata and the FSG will restore the latest (up-to-date) metadata from the other FSG, if it is a high availability cluster, or from the DR550 SSAM server for a stand-alone FSG configuration.

Delete all the files in these directories:

- ▶ /fsg
- ▶ /var/local/fsg
- ▶ /var/lib/heartbeat/crm
- ▶ /var/tmp

Proceed as follows:

1. Enter **init 1** to change the current runlevel and end all active processes.
2. Log on as root.
3. Delete the files:

```
cd /fsg
rm -r *
cd /var/local/fsg
rm -r *
cd /var/lib/heartbeat/crm
rm -r *
cd /var/tmp
rm -r *
```

Now you have to reboot the FSG again by entering the **reboot** command.

**Tip:** You can also use this clean up procedure if the FSG metadata or cache data were corrupted.

You have configured two disks, sda (disk1) and sdb (disk2). Table 12-2 shows you how to configure the file systems. Figure 12-23 on page 518 and Figure 12-24 on page 518 show you the settings for the /fsg.

Table 12-2 Configure file systems

Disk	Size	Type of file system	Mount point
sda	4 GB	ext3	/
sda	2 GB	swap	
sda	all the rest	ext3	/var/local
sdb	all	xfs	/fsg

Archived

## DR550 and Enhanced Remote Mirroring

The DR550 supports the remote site copy function, Enhanced Remote Mirroring (ERM). This chapter describes the concepts of ERM as they apply to the DR550 and provides details about how data is mirrored between sites.

The chapter also includes step-by-step procedures for using, monitoring, and recovering a DR550 with Enhanced Remote Mirroring installation.

Enhanced Remote Mirroring provides the technology that enables business continuity in the event of a disaster or unrecoverable error at one data retention subsystem. It achieves this by maintaining two copies of a data set in two different locations, enabling a second data retention subsystem to take over responsibility.

This chapter does not apply to users who plan to deploy a DR550 DR1 because the mirroring function is not available for the DR550 Model DR1. DR550 Mirroring is also not supported for DR550 systems, including more than one DR550 Storage Controller.

## 13.1 Enhanced Remote Mirroring overview

The DR550 Storage Controller provides an optional feature called Enhanced Remote Mirroring (ERM). The DR550 Storage Controller ERM feature in the context of DR550 is also referred to as DR550 Mirroring. The DR550 Storage Controller ERM feature can be ordered along with two DR550 systems. Therefore, the two DR550 systems, and in particular the DR550 Storage Controller disk systems, will be preconfigured for ERM in manufacturing. This means the mirroring feature is enabled and the mirror repository drives are pre-configured. When ordering two DR550 systems, including the mirroring option their storage capacity configuration must be exactly the same, whereas the node configuration (single and dual) may not be the same. It is also possible to upgrade an existing (non-mirrored) DR550 with the ERM feature.

The ERM option is used for online, real-time replication of data between storage subsystems over a remote distance (Figure 13-1). In the event of disaster or unrecoverable error at one storage subsystem, you can promote the second storage subsystem to take over responsibility for normal I/O operations.

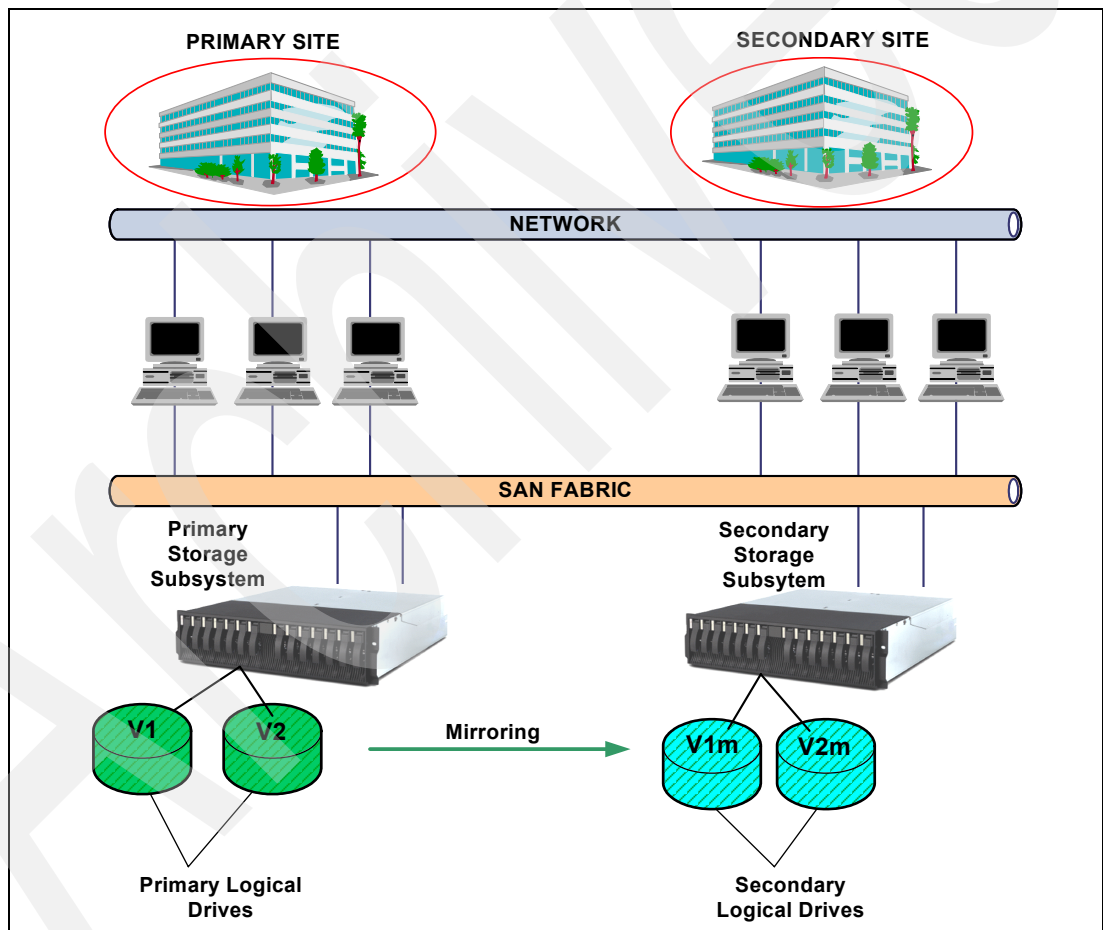


Figure 13-1 Enhanced Remote Mirroring

During operation, one DR550 is always active while the other DR550 is always standby. Both DR550 systems participating in a DR550 Mirroring configuration cannot be active at the same time. The host system(s) communicates with the primary DR550 system through the host network that connects both DR550 systems. The primary DR550 copies the data to the secondary DR550 system through the mirroring network based on Fibre Channel (SAN). The



secondary system cannot be addressed by the host system through the host network as long as it assumes the secondary role.

Enhanced Remote Mirroring offers three operating modes:

► Metro Mirror

Metro Mirror is a synchronous mirroring mode. Host write requests are written to the primary (local) storage subsystem and then transferred to the secondary (remote) storage subsystem. The remote storage controller reports the result of the write request operation to the local storage controller, which reports it to the host. This mode is called *synchronous*, because the host application does not get the write request result until the write request has been executed on both (local and remote) storage controllers.

► Global Copy

Global Copy is an asynchronous write mode. All write requests from a host are written to the primary (local) storage subsystem and immediately reported as completed to the host system. Regardless of when data was copied to the remote storage subsystem, the application does not wait for the I/O commit from the remote site. However, Global Copy does not ensure that write requests performed to multiple drives on the primary site are later processed in the same order on the remote site. As such, it is also referred to as *Asynchronous Mirroring without Consistency Group*.

**Important:** This mode is not supported with the DR550.

► Global Mirror

Global Mirror is an asynchronous write that ensures that the write requests are carried out in the same order at the remote site. This mode is also referred to as *Asynchronous Mirroring with Consistency Group*. This mode is supported on the DR550.

The Enhanced Remote Mirroring offers additional features that allow better design of business continuance solutions and easier maintenance.

***Secondary logical drive accessibility (not available on the DR550)***

With ERM, the mirrored drive (also called secondary logical drive) can be accessed in read-only mode.

**Note:** We do not recommend accessing the mirrored drive in read-only mode on the DR550. This is not supported because it may cause problems getting the drives back to read/write access in an emergency situation. There is also a risk of losing cluster stability by manually varying on volume groups (maintained by the cluster software) in read-only mode.

***Suspend Mirror and Resume Mirror capabilities***

This function allows you to suspend the mirroring process independently of the Mirroring Mode. While in Suspended State, the secondary subsystem no longer receives any write I/Os from the primary subsystem, and all data blocks that change are logged in the Mirroring Repository Volume at the primary site.

When you invoke the Resume Mirror function, it resynchronizes the changed data between the primary and the secondary logical drives without the need for a full resynchronization.

### **Change Write Mode option**

You can switch among the different mirroring modes at any time for an established mirror relationship. This is called *Dynamic Mode Switching*. You can switch among:

- ▶ Metro Mirroring (synchronous write mode)
- ▶ Global Copy (asynchronous write mode without Consistency Groups), which is not supported on DR550
- ▶ Global Mirroring (asynchronous write mode with Consistency Groups)

### **Test Mirror communication function**

After the mirroring relationship between the primary and the secondary logical drive is established, it is possible and easy to test the mirror communication status. Using the Storage Manager GUI client, you can perform the test with only one mouse click, and the result is graphically displayed (the system displays a picture of a green or red traffic light).

## **13.1.1 Requirements**

There are requirements to enable and use the Enhanced Remote Mirroring feature.

### **Enhanced Remote Mirroring Premium Feature**

An ERM Premium Feature License from IBM is required to enable the Premium Feature on *each* DR550 Storage Controller that is supposed to hold a primary or secondary mirrored logical drive. These licenses are included with the DR550 when purchased with the ERM option.

### **SAN environment**

The ERM Fibre Channel connection must be a dedicated connection between the subsystems. The ports used for the connection cannot receive I/Os from any host. These requirements are addressed through SAN zoning. The zone configuration on the Fibre Channel switches included with the DR550 separate the host ports from the subsystem mirroring ports and separate the mirroring ports between the redundant controllers. This means the mirroring ports for pair A (primary and secondary Controller A) are in a separate zone from the mirroring ports for pair B (primary and secondary Controller B). Zoning and site configuration for the DR550 are briefly discussed later in this chapter and covered in detail in 2.5.5, “DR550 SAN Switch configuration” on page 39.

### **Network environment**

To manage the Remote Mirroring connection, both DR550 Storage Controllers (local and remote) must be accessible from the same management workstation (typically one of the DR550 SSAM servers p52A nodes) over an Ethernet network. The mirroring network (data traffic) is based on Fibre Channel (SAN).

All networks must span both DR550 locations.

### **Failover requirements**

The failover process requires that the secondary system (DR550-B) is configured with the appropriate network addresses for the host network. It is possible to configure the secondary DR550 system with the same network configuration as the primary system or with a different set of TCP/IP addresses. When the same TCP/IP addresses are used for one system, you must ensure that only the primary DR550 is connected to the host network. The secondary DR550 must be disconnected from the host network or the DR550 SSAM server (p52A nodes) must be powered off. The advantage of this method is that no TCP/IP address change is required on the host system(s) during the failover procedure. When both DR550s use

different sets of TCP/IP addresses, a change in the host system(s) configuration files for the DR550 system is required. Usually this also requires a restart of the application on the host system(s).

### 13.1.2 ERM terminology

Before discussing the functionality and characteristics of ERM in more detail, we have included here, for reference, definitions of the terms used in the context of ERM.

The naming convention is generally based on the data flow direction in a mirror relationship.

**Primary Site** The Primary Site is the location of the Primary Storage Subsystem. This location, with its systems, provides the productive data to hosts and users. Data will be mirrored from here to the Secondary Site.

Primary Site is also referred to as the Local Site.

**Secondary Site** The Secondary Site is the location of Secondary Storage Subsystem. This location, with its systems, holds the mirrored data. Data is mirrored from the Primary Site to here.

Secondary Site is also referred to as the Remote Site.

**Primary Storage Subsystem**

This is the Storage Subsystem, which provides the productive data to hosts and users. Data will be mirrored from here to the Secondary Storage Subsystem.

Primary Storage Subsystem is also referred to as the Local Storage Subsystem.

**Secondary Storage Subsystem**

This is the Storage Subsystem, which holds the mirrored data. Data is mirrored from Primary Storage Subsystem to here.

Secondary Storage Subsystem is also referred to as the Remote Storage Subsystem.

**Mirroring Storage Controller Pair**

Because in a mirror relationship of a logical volume there are exactly two controllers involved, they are called a Mirroring Storage Controller Pair. Due to the mirroring software specification, the mirroring controller pair contains only one kind of controller. This means Controller A on the primary site is always connected to Controller A on the secondary site in a mirror relationship.

**Mirror Fibre Channel Connection**

A dedicated Fibre Channel connection for mirror I/Os from Primary Controller to Secondary Controller. The use of switches is mandatory for Mirror Fibre Channel Connection. Direct connections are not allowed. This connection must not be used for any other purpose.

**Primary Logical Drive and Secondary Logical Drive**

When you create a Enhanced Remote Mirroring relationship, a mirrored logical drive pair is defined and consists of a primary logical drive at the Primary Storage Subsystem and a secondary logical drive at a Secondary Storage Subsystem.

**Mirror Repository Logical Drive**

A Mirror Repository Logical Drive is a *special logical drive* in the Storage Subsystem created as a resource for the *controller owner* of the primary logical drives in a mirror relationship.

<b>Mirror Relationship</b>	Mirror Relationship is an established mirroring connection between two logical drives, defined on two different Storage Subsystems. It is defined through the corresponding Mirror Repository Drive, the mirrored logical drive pair, and controller ownership of the mirrored logical drives.
<b>Mirror Role</b>	There is always the Primary Role of a mirror relationship member that grants the write access to the Primary Logical Volume from the host systems and a Secondary Role of a mirror relationship member that prohibits the write access to the Secondary Logical Volume from the host systems.
<b>Role Reversal</b>	In the event of site failure or for maintenance tasks, you can change the role of a logical drive in a given mirror relationship from primary to secondary and vice versa and grant or deny (to the hosts) write access on this logical drive. This procedure is called Role Reversal.
<b>Full Synchronization</b>	When you first create the mirror relationship, all data is initially copied from the primary logical drive to the remote logical drive. This procedure is called Full Synchronization. Under normal circumstances, there is no longer a need for full synchronization, because data changes can be logged on the mirror repository logical drive. If all segments for the given logical drive are changed before the synchronization can take place (due to link error conditions, for example), then the full synchronization is obviously needed.

## 13.2 Primary and secondary logical drives

To create an Enhanced Remote Mirroring relationship, a mirrored logical drive pair is defined and consists of a primary logical drive at the primary storage subsystem and a secondary logical drive at a remote storage subsystem.

Figure 13-2 shows that both the primary and secondary logical drives are displayed at the primary site while only the secondary logical drive is displayed at the secondary site.



Figure 13-2 Primary and secondary volumes on primary (right) and secondary (left) sites

### 13.2.1 Logical drive roles

The primary or secondary role is determined at the logical drive level, not the DR550 Storage Controller subsystem level. Any given logical drive can be exclusively in either the primary or secondary role. In a DR550 ERM configuration, all of the logical drives on the subsystem at the active site are in a primary role and the logical drives on the subsystem at the backup site are in a secondary role.

### 13.2.2 Host accessibility of secondary logical drive

When you first create the mirror relationship, all data from the primary logical drive is copied to the remote logical drive (full synchronization). During full synchronization, the primary logical drive remains accessible for all normal host I/Os. The read/write behavior changes automatically if there is a role reversal of the logical drive (this applies in both directions: from secondary to primary and vice versa).

### 13.2.3 Mirrored logical drive controller ownership

The logical drives belonging to Controller A in the primary storage server must be mirrored to the logical drives owned by Controller A in the secondary subsystem. The same rule applies to the logical drives owned by Controller B.

A primary controller will only attempt to communicate with its matching controller in the secondary storage subsystem. The controller (A or B) that owns the primary logical drive determines the controller owner of the secondary logical drive. If the primary logical drive is owned by Controller A on the primary site, the secondary logical drive is therefore owned by Controller A on the secondary side. If primary Controller A cannot communicate with secondary Controller A, no controller ownership changes take place, and the remote mirror link is broken for that logical drive pair.

If an ownership change of the logical drive on the primary site occurs (caused by I/O path error or administrator interaction), an ownership change of the logical drive on the remote controller takes place with the first write request (from primary controller to the secondary controller) issued through the mirror connection. When the ownership transfer on the secondary side occurs, a “Needs Attention” status is *not* displayed. The logical drive ownership of the secondary controller cannot be changed by either the host or through the GUI: It is entirely controlled by the primary site.

## 13.3 Mirror repository logical drives

A mirror repository logical drive is a *special logical drive* in the storage subsystem created as a resource for the controller owner of the primary logical drive in a remote logical drive mirror. Two mirror repository logical drives (one for each controller in a subsystem) are automatically created when activating the ERM Premium Feature (Figure 13-3 on page 528). One mirror repository drive is created for each storage controller. The mirror repository logical drive stores (queues) the mirroring information, including information about remote write requests that are not yet written to the secondary logical drive. After a confirmation of a given write request has occurred, the corresponding entry stored in the mirror repository logical drive is removed.

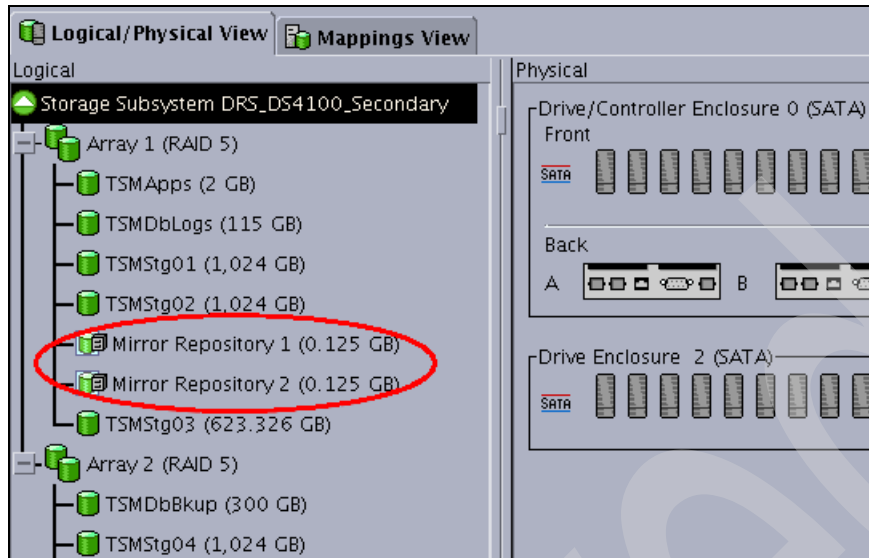


Figure 13-3 Mirror Repository logical drives after activating the ERM feature

The mirroring process for *all* primary drives defined on a storage controller is monitored by the corresponding controller's mirror repository drive.

#### Notes:

- ▶ Two mirror repository logical drives are created, one for each controller, in every storage subsystem that has Enhanced Remote Mirroring activated.
- ▶ No actual host data is written to the repository logical drive. It is only used for status and control data in relation to the Enhanced Remote Mirroring relationships.

The capacity is set at 128 MB for each logical drive. The segment size is set at 32 KB (or 64 blocks). The controller determines the modification priority. The drive size, the segment size, and modification priority for a mirror repository logical drive cannot be changed.

## 13.4 Mirror relationship

A mirror relationship is an established mirroring connection between two storage subsystems. It contains the definitions of mirror repository drives and mirroring pairs of the logical drives.

Before you can establish a mirror relationship between two storage subsystems, you must:

- ▶ Establish a proper network communication between the storage servers and the storage management server.
- ▶ Establish a proper Fibre Channel communication between the storage servers.
- ▶ Enable the Enhanced Remote Mirroring Premium Feature on both storage subsystems.
- ▶ Activate the ERM function on both storage subsystems. This step creates the mirror repository drives for all mirrored logical drives that will ever be created.

See 13.6, “SAN fabric connectivity” on page 536 and 13.7, “Enhanced Remote Mirroring on DR550: Step-by-step” on page 538 for information about ERM communication issues.

**Important:** You cannot create a mirror relationship if the primary logical drive contains unreadable sectors. Furthermore, if an unreadable sector is discovered during a mirroring operation, the mirror relationship will fail.

### 13.4.1 Remote Mirror status

The status of a Remote Mirror indicates whether or not the data on the primary logical drive is identical (fully synchronized) with data on the secondary logical drive. A mirror status is independent of the status of the actual logical drives in the mirrored pair. The primary and secondary logical drive icons in the Logical View indicate the state of each logical drive as well as the state of the Remote Mirror.

There are four states of a Remote Mirror: Synchronized, Synchronizing, Suspended, and Unsynchronized.

The mirror status is represented by different icons in the Logical View of the Storage Manager GUI. The icons depend on which storage subsystem you are monitoring and whether the storage subsystem contains the primary logical drive or the secondary logical drive. The table in Figure 13-4 on page 530 shows a summary of the mirrored drive status icons.

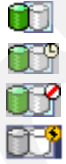










Remote Mirror status representation in logical view			
On Storage Subsystem containing the monitored Logical Drive		Primary LD Status	Secondary LD Status
Logical Drive Status	Mirrored Pair Status		
<b>Optimal</b>	Synchronized Synchronizing Unsynchronized Suspended		
<b>Degraded</b>	Synchronized Synchronizing Unsynchronized Suspended		
<b>Failed</b>	Synchronized Synchronizing Unsynchronized Suspended		
<b>Offline</b>	Synchronized Synchronizing Unsynchronized Suspended		
<b>Missing</b>	Synchronized Synchronizing Unsynchronized Suspended		
<b>Unresponsive</b>	Synchronized Synchronizing Unsynchronized Suspended		

Figure 13-4 Mirrored Drive status icons



## 13.5 Data replication process

This section describes how data is replicated between the DR550 Storage Controllers. Understanding how data flows between the subsystems is critical for setting the appropriate mirroring configuration and performing maintenance tasks.

We have previously mentioned that whenever a new mirror relationship is created, a *full synchronization* takes place between the primary and secondary logical drives. We assume here that a mirror relationship has already been established, and that the full synchronization is complete and successful. Remember also that data replication between the primary logical drive and the secondary logical drive is managed at the storage subsystem level. It is transparent to the attached host systems and applications.

### 13.5.1 Metro Mirroring (synchronous mirroring)

Metro Mirroring is a synchronous mirroring mode, meaning that the controller does not send the I/O completion to the host until the data has been copied to both the primary and secondary logical drives.

When a primary controller receives a write request from a host, the controller first logs information about the write request on the *mirror repository logical drive* (the information is actually placed in a queue). In parallel, it writes the data to the primary logical drive. The controller then initiates a remote write operation to copy the affected data blocks to the secondary logical drive at the remote site. When the remote write operation is complete, the primary controller removes the log record from the mirror repository logical drive, and the controller sends an I/O completion indication back to the host system.

**Note:** The owning primary controller only writes status and control information to the repository logical drive. The repository is not used to store actual host data.

On the DR550 Storage Controller, you can enable write caching. When write caching is enabled on either the primary or secondary logical drive, the I/O completion is sent when data is in the cache on the site (primary or secondary) where write caching is enabled. When write caching is disabled on either the primary or secondary logical drive, then the I/O completion is not sent until the data has been stored to physical media on that site.

Figure 13-5 depicts how a write request from the host flows to both controllers to provide an instant copy.

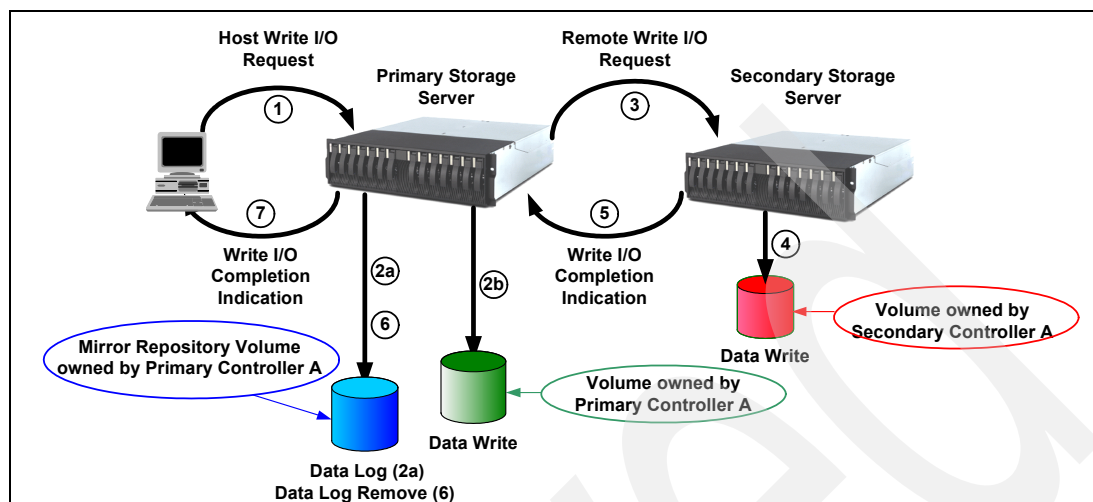


Figure 13-5 Metro Mirroring mode (synchronous mirroring) data flow

When a controller receives a read request from a host system, the read request is handled on the primary storage subsystem, and no communication takes place between the primary and secondary storage subsystems.

### 13.5.2 Global Copy (asynchronous mirroring without write consistency group)

Global Copy is an asynchronous write mode. All write requests from the host are written to the primary (local) logical drive and immediately reported as completed to the host system. Regardless of when data was copied to the remote Storage Subsystem, the application does not wait for the I/O write request result from the remote site. However, Global Copy does not ensure that write requests at the primary site are processed in the same order at the remote site. As such, it is also referred as asynchronous mirroring without write consistency group.

When a primary controller (the controller owner of the primary logical drive) receives a write request from a host, the controller first logs information about the write request on the mirror repository logical drive (the information is actually placed in a queue). In parallel, it writes the data to the primary logical drive (or cache). After the data has been written (or cached), the host receives an I/O completion from the primary controller. The controller then initiates a background remote write operation to copy the corresponding data blocks to the secondary logical drive at the remote site. After the data has been copied to the secondary logical drive at the remote site (or cached), the primary controller removes the log record on the mirror repository logical drive (delete from the queue).

When multiple mirror relationships are defined on the Storage Subsystem, the background synchronization of affected data blocks between the primary and secondary controller for the different relationships are conducted in parallel (a multi-threaded process). Thus, the write order for multiple volumes (for example, write requests to a Database Data Volume and a Database Log Volume on a Database Server) is not guaranteed with the Global Copy mode. For that reason, this mode is not supported on the DR550.

See Figure 13-6 on page 533 for a logical view of the data flow.

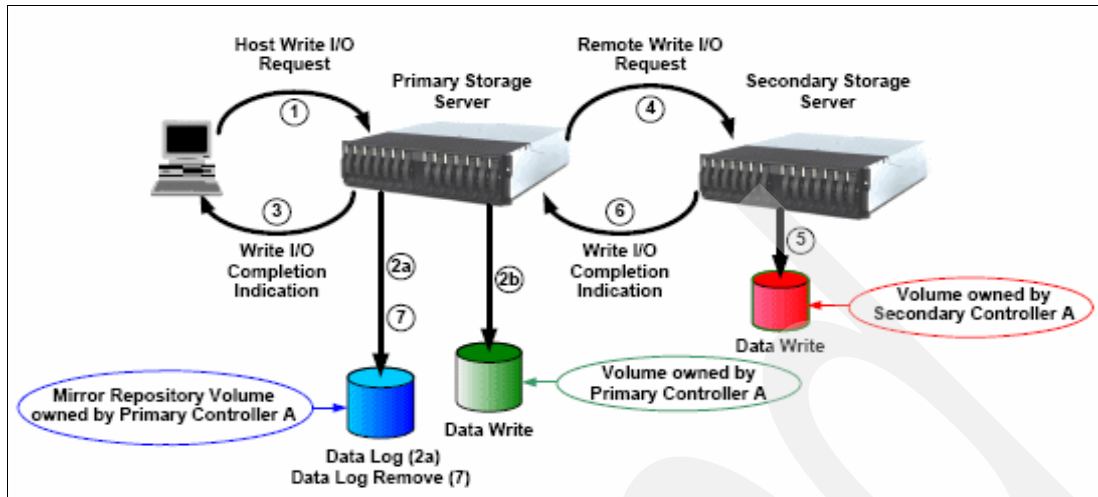


Figure 13-6 Global Copy Mode (Asynchronous Mirroring) data flow

When write cache is enabled on either the primary or secondary logical drive, the I/O completion is sent when data is in the cache on the site (primary or secondary) where write caching is enabled. When write caching is disabled on either the primary or secondary logical drive, then the I/O completion is not sent until the data has been stored to physical media on that site.

When a controller receives a read request from a host system, the read request is handled on the primary Storage Subsystem and no communication takes place between the primary and secondary Storage Subsystems.

### 13.5.3 Global Mirroring (asynchronous mirroring with consistency group)

Global Mirroring is an asynchronous write mode where the order of host write requests at the primary site is preserved at the secondary site. This mode is also referred as asynchronous mirroring with write consistency group.

To preserve the write order for multiple mirrored volumes, Global Mirroring uses the write consistency group functionality. It tracks the order of the host write requests, queues them, and sends them to the remote controller in the same order.

**Important:** Selecting write consistency for a single mirror relationship does not change the process in which data is replicated. More than one mirror relationship must reside on the primary Storage Subsystem for the replication process to change.

The volumes for which the write request order must be preserved have to be defined as members of a Write Consistency Group. The Write Consistency Group can be defined from the Storage Manager GUI.

When a primary controller (the controller owner of the primary logical drive) receives a write request from a host, the controller first logs information about the write on the *mirror repository logical drive*. It then writes the data to the primary logical drives. The controller then initiates a remote write operation to copy the affected data blocks to the secondary logical drives at the remote site. The remote write request order corresponds to the host write request order.

After the host write to the primary logical drive is completed and the data has been copied to the secondary logical drive at the remote site, the controller removes the log record from the mirror repository logical drive. Refer to Figure 13-7 for a logical view of the data flow.

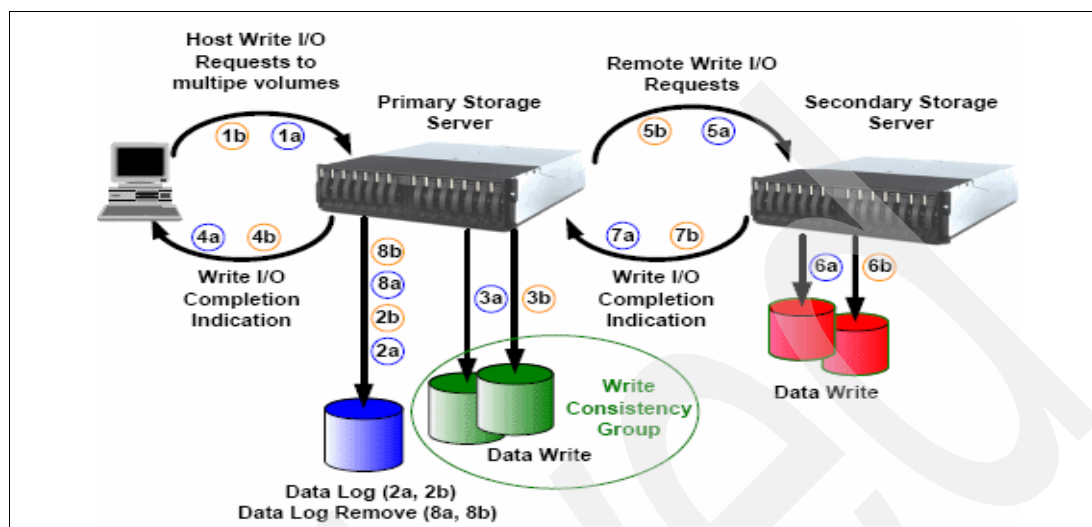


Figure 13-7 Global Mirroring logical data flow

When write caching is enabled on either the primary or secondary logical drive, the I/O completion is sent as soon as data is in the cache on the site (primary or secondary) where write caching is enabled. When write caching is disabled, then the I/O completion is not sent until the data has been stored to physical media.

**Note:** The Mirror Repository logical drive can queue a number of I/O requests (up to 128) with firmware 6.1x.xx.xx). Until the maximum number has been reached, the mirrored pair state is said to be in a *Synchronized* state. If the maximum number of unsynchronized I/O requests is exceeded, the state of the mirrored pair changes to *Unsynchronized*.

The host can continue to issue write requests to the primary logical drive, but remote writes to the secondary logical drive will not take place. The requests are stored in the Remote Mirror repository on the primary site (*delta logging*).

Whenever the data on the primary drive and the secondary drive becomes unsynchronized, the controller owner of the primary drive initiates a changed data synchronization.

### 13.5.4 Data resynchronization process

If a link interruption or logical drive error prevents communication with the secondary storage subsystem, the controller owner of the primary logical drive transitions the mirrored pair into an *Unsynchronized* status and sends an I/O completion to the host that sent the write request. The host can continue to issue write requests to the primary logical drive, but remote writes to the secondary logical drive will not take place. The requests are stored in the mirror repository on the primary site (*delta logging*).

When connectivity is restored between the controller owner of the primary logical drive and the controller owner of the secondary logical drive, a resynchronization takes place.

If the mirroring state is changed to a *Suspended* state, the host can also continue to issue write requests to the primary logical drive.

There are two essential differences between the *Unsynchronized* and the *Suspended* states. The first state causes an error condition indication (see 13.9, “Bringing up the secondary DR550: Step by step” on page 566) while the second state is an administrator-forced status change. Also, the behavior for data resynchronization is different for the two states.

When in a *Suspend* state, the administrator must manually resume the mirroring to return to a *Synchronized* state. The *Unsynchronized* state can either be manually or automatically transitioned into a *Synchronized* state.

**Notes:**

- ▶ Data Resynchronization is needed when a mirrored pair has become *Unsynchronized*.
- ▶ The *Suspended* state is a subset of the *Unsynchronized* state. The specification of the *Unsynchronized* and *Suspended* states let us differentiate between error (*Unsynchronized*) and maintenance (*Suspended*) conditions of a given mirrored logical drive pair.
- ▶ Only the blocks of data that have changed on the primary logical drive during *Unsynchronized* or *Suspended* state will be copied to the secondary logical drive.

Normally, when resuming a mirror relationship or reestablishing the communication between the subsystems, only changed data blocks are sent to the remote site. However, there are some cases in which full synchronization of the primary and secondary logical drives is necessary:

- ▶ Establishing a *new* mirror relationship between two given logical drives
- ▶ Any kind of total failure of the mirror relationship members
- ▶ Any kind of mirror repository logical drive failure
- ▶ Change of all data block track entries in the mirror repository logical drive while any kind of mirroring communication errors occurred
- ▶ Change of all data block track entries in the mirror repository logical drive in suspended state of the mirror relationship

**Note:** Information about changed data blocks (delta logging) is logged in the Mirror Repository Logical Drive. The resynchronization process uses this log to send only the changed data to the remote site. If during the interruption of the mirroring, all data blocks on the primary repository logical drive were changed, a full synchronization will take place.

### **Manual resynchronization**

Manual resynchronization is the recommended method for resynchronization of an *Unsynchronized* mirrored pair because it allows you to manage the resynchronization process in a way that provides the best opportunity for recovering data.

### **Automatic resynchronization**

Automatic resynchronization will start automatically after the controller detects that communication is restored for an unsynchronized mirrored pair. When the Automatic resynchronization option has been selected and a communication failure occurs between the primary and secondary storage subsystems, the controller owner of the primary logical drive will start resynchronizing the primary and secondary logical drives immediately after detecting that communication has been restored.

**Tip:** We recommend that you *do not* use automatic resynchronization.

**Important:** Any communication disruptions between the primary and secondary storage subsystem while resynchronization is underway can result in a mix of new and old data on the secondary logical drive. This would render the data unusable in a disaster situation.

### 13.5.5 Data synchronization priority

Data synchronization priority is a parameter in a mirror relationship that defines the proportion between system resources for host I/Os and system resources used to synchronize data in a given mirror relationship.

You can choose among five synchronization priorities for the primary controller owner, ranging from Lowest to Highest. The higher the priority, the greater performance impact there will be on host applications.

The synchronization priority is set when creating the Remote Mirror, but can be changed later using the Change Synchronization Settings option. Although the primary controller owner performs the full synchronization and uses the synchronization priority, it is set for both the primary and secondary logical drives.

**Important:** The Data Synchronization Priority parameter is only used during the data resynchronization process. Data resynchronization is needed every time a mirrored pair becomes *Unsynchronized*. Since the *Suspended* state is a sub-state of the *Unsynchronized* state, the Synchronization Priority is also used when resuming mirroring from a *Suspended* to a *Synchronized* state.

## 13.6 SAN fabric connectivity

The mirroring network between the mirrored DR550 systems requires an Inter Switch Link (ISL) between the SAN switches of the DR550 systems. This connection is based on Fibre Channel (SAN).

There are important requirements and rules to follow regarding Fibre Channel connections and SAN fabric attachment in order to correctly implement ERM. SAN planning is a critical task and must include SAN ports, SAN zoning, and cabling considerations.

This section reviews these SAN considerations and configuration.

Here we examine the SAN requirements from a general, conceptual standpoint. Detailed information and procedures for implementing the SAN fabric, configuring SAN switches, and SAN zones are beyond the scope of this book.

### SAN fabric configuration

Dedicated Remote Mirroring ports (A2 and B2 host side controller ports) must be attached to a SAN fabric with support for the Directory Service and Name Service interfaces. In other words, there must be at least one SAN Switch with SAN Zoning capability installed in the mirroring environment. Since ERM is typically used for providing a High Availability Solution, we strongly recommend that you use at least two switches in the SAN. See Figure 13-9 on page 539 and chapter 13.7, “Enhanced Remote Mirroring on DR550: Step-by-step” on page 538 for SAN fabric configuration examples.

**Tip:** IBM and other SAN switch vendors recommend that you configure two SAN fabrics with independent SAN zones for the highest level of redundancy when implementing ERM.

## SAN fabric zoning

It is also mandatory to keep ERM links and the host links in separate SAN zones. We recommend that you create four separate zones within the fabric:

- ▶ First for a host connection to Controller A
- ▶ Second for a host connection to Controller B
- ▶ Third for Controller A Remote Mirroring links between storage subsystems
- ▶ Fourth for Controller B Remote Mirroring links between storage subsystems

Figure 13-9 on page 539 shows an example of zoning a fabric for DR550 Storage Controller and Enhanced Remote Mirroring.

## DR550 Storage Controller Fibre Channel configuration for ERM

The Enhanced Remote Mirroring option requires two dedicated controller host port connections for each storage subsystem that will participate in Enhanced Remote Mirroring.

When the Enhanced Remote Mirroring option is *activated* on a storage subsystem, the A2 and B2 *host-side* controller ports become dedicated to Remote Mirror operations. See Figure 13-8 for the host side port locations on a DR550 Storage Controller mode 70. After ERM activation, the A2 and B2 host-side controller ports will no longer permit host system I/O requests. Persistent Reservations (if any) will also be removed from these ports. The A2 and B2 host-side controller ports will *only* be able to communicate to other storage subsystems that have the Enhanced Remote Mirroring option activated *and* are connected by the same fabric with proper SAN zoning configuration.

For more information about the SAN and zoning configuration of the DR550 with ERM, see 2.5.5, “DR550 SAN Switch configuration” on page 39.




Figure 13-8 Host-side ports for host and ERM connections on the DR550 Storage Controller

The distance between primary and remote storage subsystems is normally limited by the distance limits of Fibre Channel inter-switch links (ISL). See Table 13-1.

Table 13-1 Fibre Channel distance limits

Fiber cable type	Laser type	Distance limit in kilometers	Distance limit in miles
Single mode 9 micron	Long wave	10 km	6.25 miles
Multi mode 50 micron	Short wave	0.5 km	0.32 mile

**Important:** The maximum distance supported by SFPs impacts the maximum connection bandwidth. The SFP automatically changes the connection speed to 1 Gbps if the maximum distance for a 2 Gbps connection is exceeded.

For a short wave SFP, the maximum length of a 2 Gbps connection is 300 m; for the long wave SFP, it is 2 km. Refer to Table 13-2.

Table 13-2 Fiber types and connection bandwidth

Fiber type	Speed	Maximum distance
50 micron MMF (short wave)	1 Gbps	500 m
50 micron MMF (short wave)	2 Gbps	300 m
50 micron MMF (short wave)	4 Gbps	150 m
62.5 micron MMF (short wave)	1 Gbps	175 m/300 m
62.5 micron MMF (short wave)	2 Gbps	90 m/150 m

## 13.7 Enhanced Remote Mirroring on DR550: Step-by-step

This section presents a step-by-step procedural guide to common administration tasks for DR550 Storage Controller Enhanced Remote Mirroring on the DR550.

**Important:** To set up ERM and establish mirror relationships, application I/Os to the disks must be stopped. In other words, for the DR550, HACMP and SSAM (Tivoli Storage Manager) must be stopped and all AIX Volume groups must be varied off.

The tasks that are covered include:

- ▶ Establishing network access to the DR550 Storage Controllers
- ▶ Merging the primary and secondary fabrics
- ▶ Enabling and activating the Enhanced Remote Mirroring premium feature
- ▶ Creating Enhanced Remote Mirroring relationships
- ▶ Viewing Enhanced Remote Mirroring properties
- ▶ Changing Enhanced Remote Mirroring synchronization settings
- ▶ Removing Enhanced Remote Mirroring relationships

**Important:** Before setting up remote mirroring, make sure you do not have any data stored on the secondary DR550. Establishing a mirror relationship will reset the corresponding LUNs in the secondary DR550 and all data on the mirror LUNs will be lost.

Figure 13-9 on page 539 illustrates the FC connections with a single-node DR550.



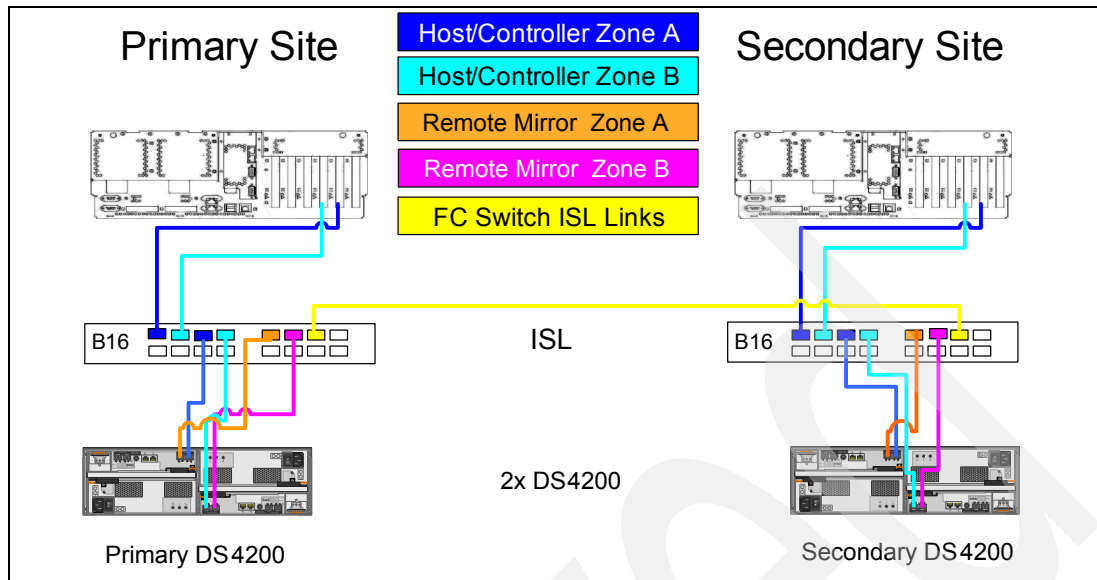


Figure 13-9 Enhanced Remote Mirroring

Figure 13-10 illustrates the FC connections with a dual-node DR550. In this case, we have two FC switches at each site and two Inter-Switch Links (ISL) across two separate fabrics, for a highly available configuration.

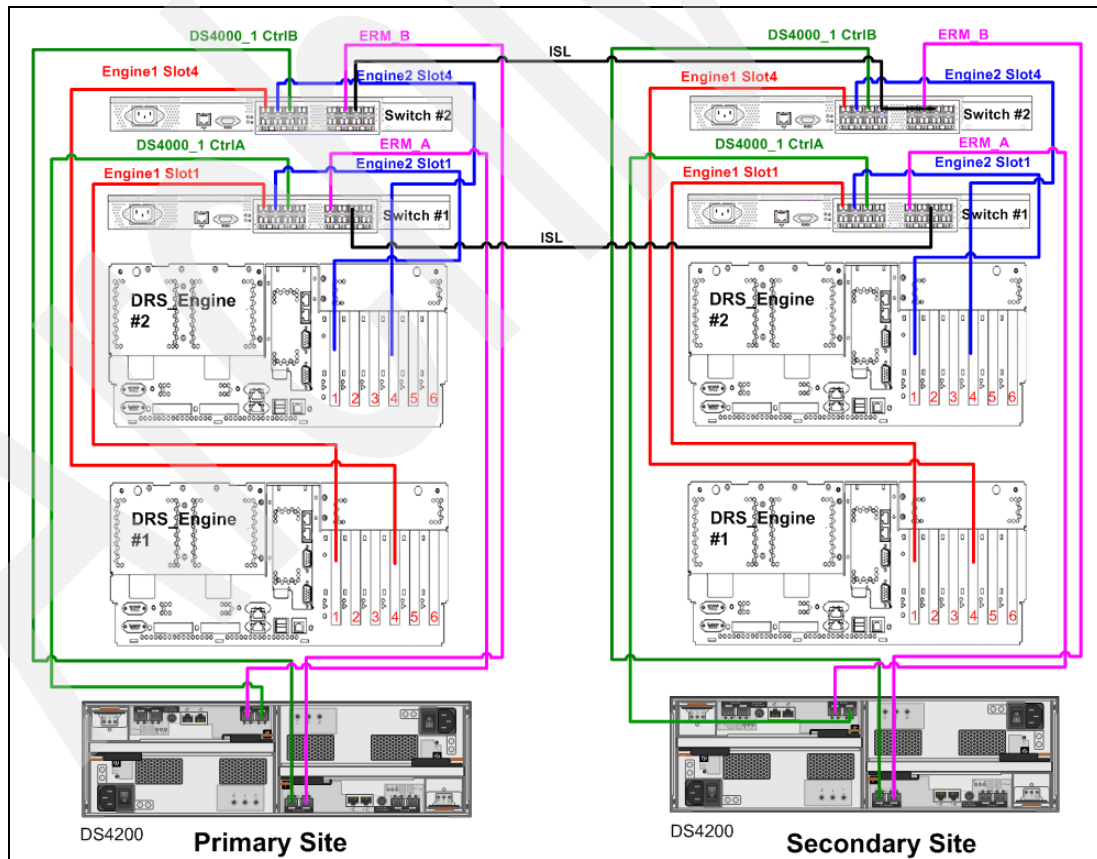


Figure 13-10 Remote Mirroring FC connections for dual-node DR550

### 13.7.1 Establishing network access to the DR550 Storage Controllers

In order to establish mirror pairs, the DR550 Storage Manager client must be able to access both of the DR550 Storage Controllers in the mirror relationship. Because the DR550 Storage Controllers are connected with all Ethernet Ports to the internal Ethernet Switches, some tasks are necessary to allow the DR550 Storage Manager to connect to both the primary and secondary DR550 Storage Controller at the same time.

Ethernet switches at the primary site must have a connection to the Ethernet ports of the secondary DR550 Storage Controller and vice versa. This is shown in Figure 13-11 on page 541.

By default, there are two Ethernet connections already established from both controllers to the internal Ethernet switches. When using ERM, you must remove the second Ethernet connection from every Controller (port E2), because this port is needed for ERM management and must be connected to the Remote Site.

The Ethernet connections shown in Figure 13-11 on page 541 offer redundancy because the SMCClient used to manage the mirror is available on each of the DR550 SSAM Servers (p52A). To remain compliant (with retention regulations), no outside connections should be made to the Ethernet switches. You must prevent any rogue management workstation to gain access to the mirroring Ethernet network and potentially be able to manipulate the DR550 Storage Controller system.

The default IP addresses on the DR550 Storage Controller are:

- ▶ Controller A Port 1: 192.168.4.101
- ▶ Controller A Port 2: 192.168.5.105
- ▶ Controller B Port1: 192.168.5.102
- ▶ Controller B port 2: 192.168.4.106

**Important:** Use a virtual private network (VPN), or private LAN to isolate the DR550 Storage Controller from the external LAN.

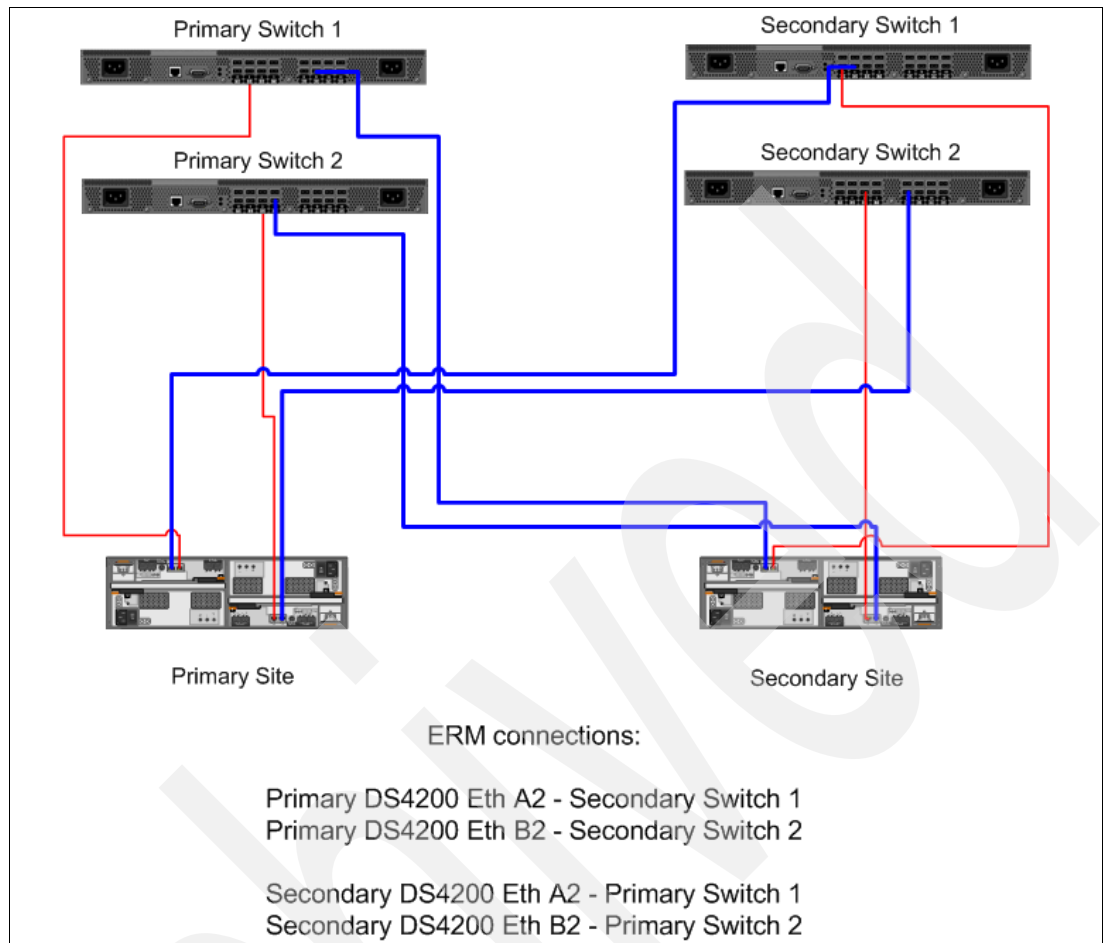


Figure 13-11 DR550 Storage Controller Ethernet connection ERM

### Rediscovering the DR550 Storage Controllers

Once connected to the Ethernet network, the DR550 Storage Controllers need to be rediscovered by the IBM DS4000 Storage Manager for DR550. From the main window of the DS4000 Storage Manager for DR550, select **Edit** → **Add Storage Subsystem** (Figure 13-12).

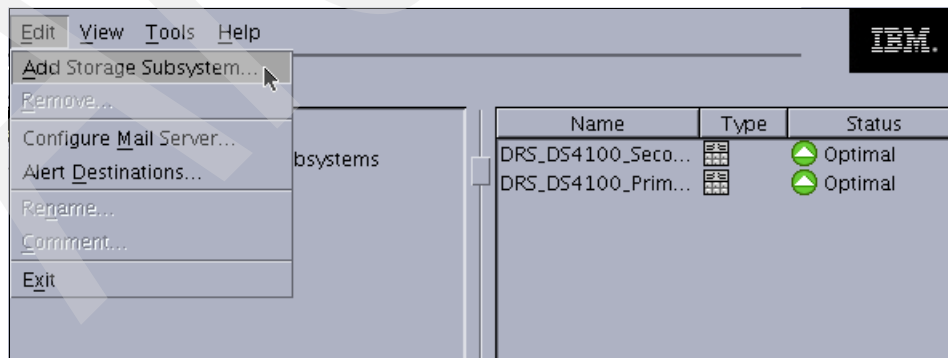


Figure 13-12 Rediscovering the DR550 Storage Controller

This will launch the Add Storage Subsystem dialog (Figure 13-13). Enter the IP addresses from the secondary Storage Subsystem and click **Add**. If network connectivity is working, the subsystems will be added to the list in the main window. This can be done on both sites.

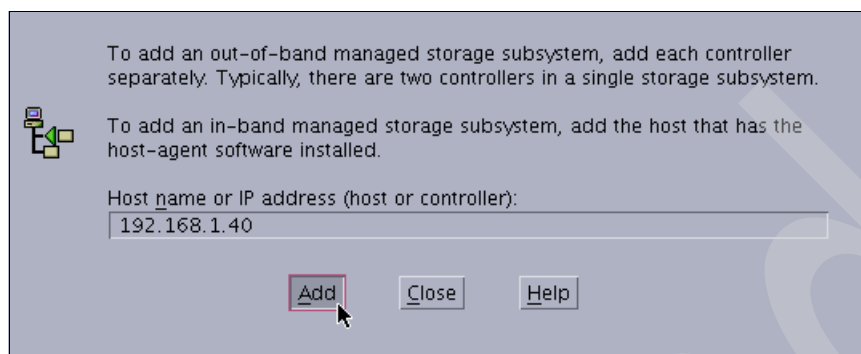


Figure 13-13 Add Storage Subsystem dialog

### 13.7.2 Merging the primary and secondary SAN fabrics

To allow the fabrics to merge, there are some important steps to take before connecting the Fibre Channel switches at the primary site to those at the secondary site. To perform these steps, it will be necessary to access the 2005 B16 FC switches. There are cables running from the San Switch's Ethernet ports to the internal Ethernet Switch. Please connect a mobile computer to any available port on the internal Ethernet Switch, change the IP address from your mobile computer according to the subnet, and point a browser at the appropriate Switch IP address. They are 192.168.4.31 and 192.168.4.32 for switches 1 and 2, respectively (counting from the bottom up in the rack). This will load the Switch Explorer Web GUI (Figure 13-14) from which you can launch the Admin Tools and Zone Admin as indicated.

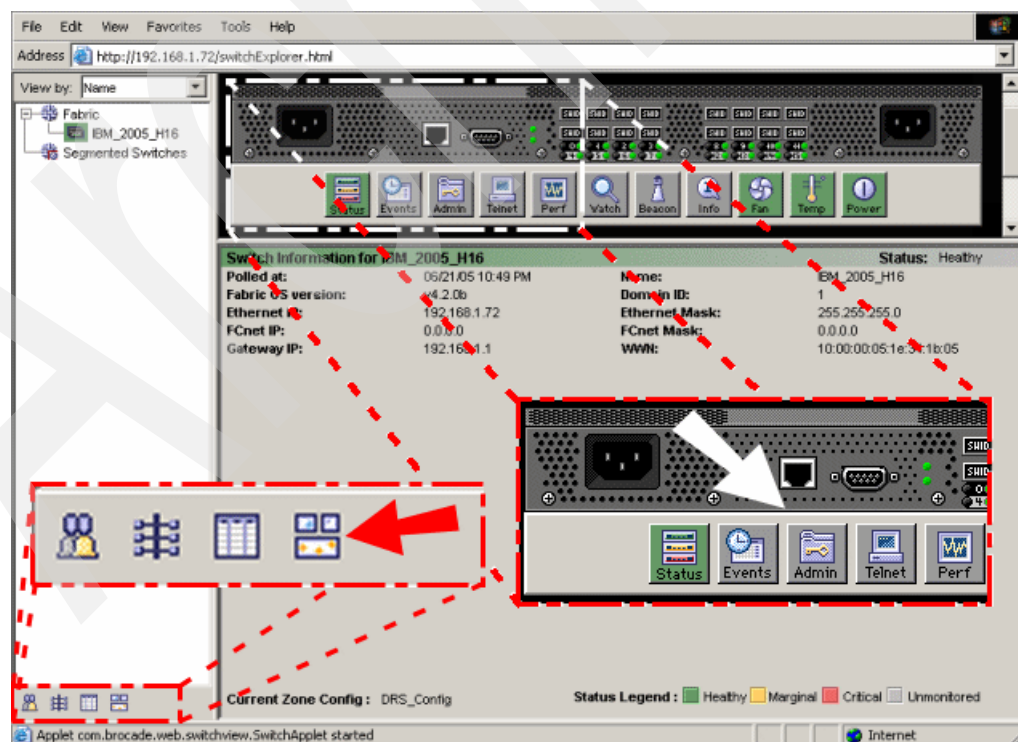


Figure 13-14 Switch Explorer: Admin Tools (white arrow) and Zone Admin (red arrow)

**Important:** The following procedures are *not* concurrent. The nodes will lose connection to the DR550 Storage Controller. In a single switch configuration, the node will have no paths to the external disk. It is mandatory that the node is shut down or that all of the SAN volumes are varied off. In a dual-switch configuration, this is not necessary, but we still recommend it.

## Changing the Domain ID

First, launch the Admin Tools. When prompted for login credentials, enter the default user ID admin and the default password password. Figure 13-15 shows the Admin Tools window, and the numbering for the following steps:

1. Under Switch Status, select **Disable**.
2. Click **Apply**.
3. Under Name and ID, change the **Domain ID** to 2.
4. Click **Apply**.
5. Under Switch Status, select **Enable**.
6. Click **Apply**.

After each Apply, a confirmation dialog will appear. Read the information and click **Yes**. This procedure changes the Domain ID for the switches at one site so there will be no Domain ID conflicts when the fabrics are connected.

The screenshot shows the 'Admin Tools' window for switch 'IBM\_2005\_H16'. The 'Name and ID' tab is active. The 'Domain ID' is set to 2. The 'Switch Status' is set to 'Enable'. The 'Apply' button is highlighted. Red circles and numbers indicate the steps: 1. 'Enable' radio button, 2. 'Apply' button, 3. 'Domain ID' field, 4. 'Apply' button, 5. 'Enable' radio button, 6. 'Apply' button. The 'Apply' button is circled in red and labeled '2 & 4 & 6'. The 'Domain ID' field is circled in red and labeled '3'. The 'Enable' radio button is circled in red and labeled '1 & 5'.

Figure 13-15 Admin Tools: Changing the Domain ID

## Updating the zoning

The SAN switches in the DR550 use port zoning, and as such, the definitions are dependent on the Domain ID of the switches. Now that the Domain ID has been changed, it is necessary to update the zone configuration. Also, all aliases, zones, and zone configurations need to be renamed to ensure there will be no conflicts with the zone configuration at the other site. Finally, zones for the ERM communication between the DR550 Storage Controller's need to be defined and added to the configuration.

Launch the Zone Admin (Figure 13-16) by clicking the icon indicated in Figure 13-14 on page 542.

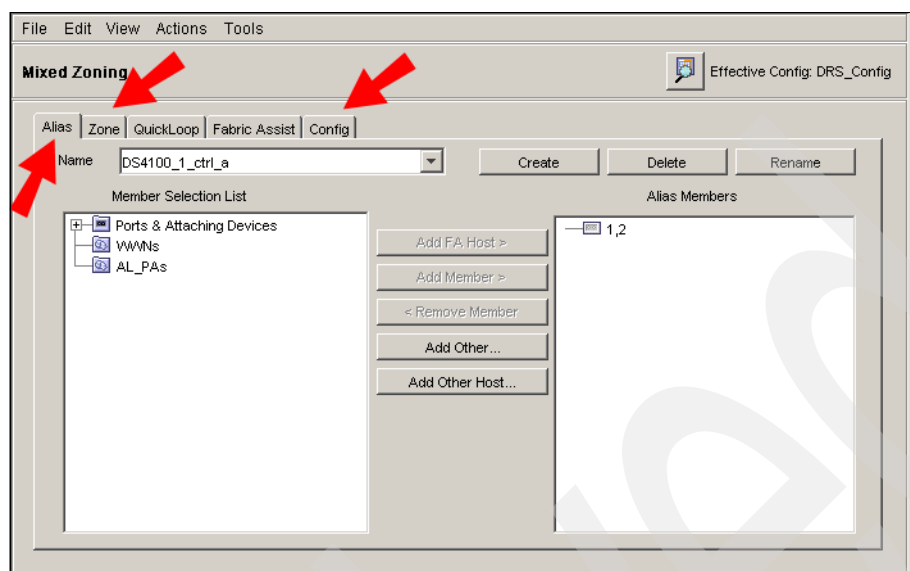


Figure 13-16 Zone Admin: Alias, Zone, and Config tabs

In port zoning, members of a zone are designated by two numbers separated by a comma (for example, 1,2). The number on the left indicates the Domain ID of the port's switch and the number on the right is the port number.

### Aliases

Under the Alias tab, there will be four aliases in the drop-down box. Select one at a time, and rename it by adding something meaningful to the end, such as \_site\_2 or \_remote. Also, remove the old alias member and replace it with the corresponding port from the Member Selection List, as shown in Figure 13-17 and Figure 13-18 on page 545.

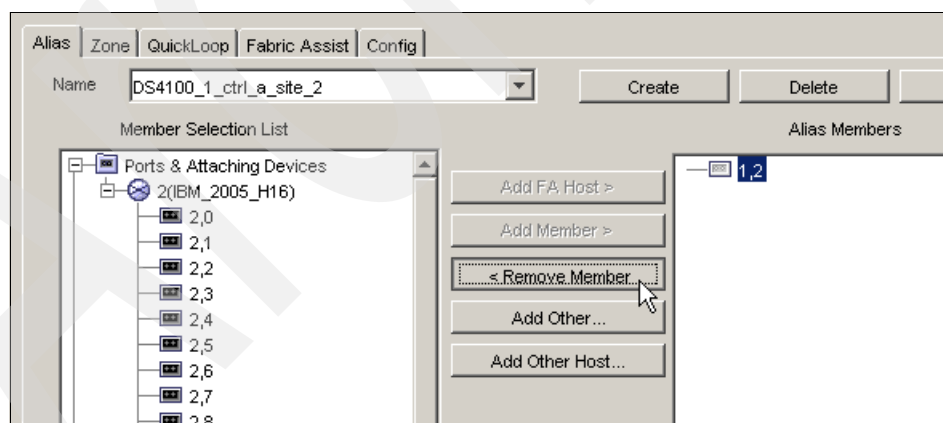


Figure 13-17 Removing the old alias member

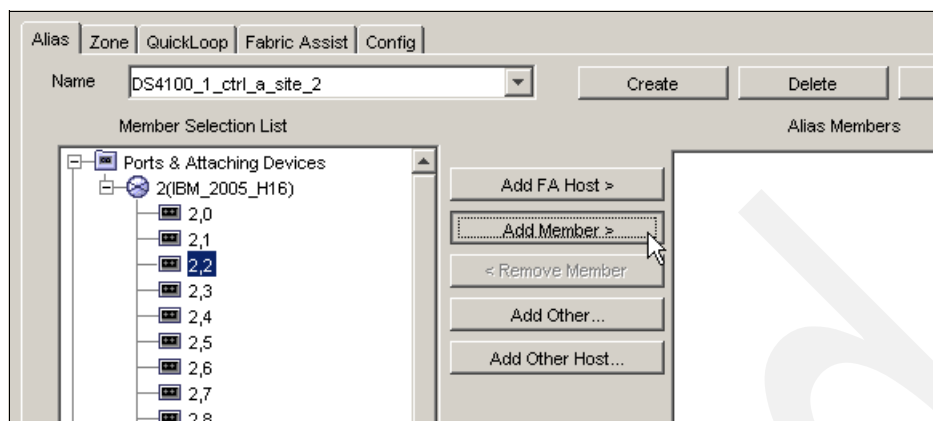


Figure 13-18 Adding the corresponding new member

## Zones

Under the Zone tab, there will be four zones in the drop-down box. Select one at a time, and rename it by adding something meaningful to the end, such as `_site_2` or `_remote`.

Then, add two new zones for ERM with meaningful names, such as `ERM_ctrl_a_zone` and `ERM_ctrl_b_zone`. Add ports 2,8 and 1,8 to the first, and 2,9 and 1,9 to the second. These zones will allow communication between the primary and secondary DR550 Storage Controllers.

## Zone configuration

Under the Config tab, add the two new ERM zones to the zone configuration. Then disable zoning by selecting **Actions** → **Disable Zoning** and rename the zone configuration by adding something meaningful to the end, such as `_site_2` or `_remote`.

Finally, save and reenable the configuration by selecting **Actions** → **Enable Config** and then **Actions** → **Save Config Only**.

In Table 13-3 and Table 13-4 on page 546, you can find Information about the port occupation and the zoning information for the DR550 DR2 single and dual node.

Table 13-3 FC switch port assignments

Port #	Single-Node FC Switch 1	Dual-Node FC Switch 1	Dual-Node FC Switch 2
0	Engine 1, slot 5	Engine 1, slot 5	Engine 1, slot 4
1	Engine 1, slot 4	Engine 2, slot 5	Engine 2, slot 4
2	DS4200 1, A1	DS4200 1, A1	DS4200 1, B1
3	DS4200 1, B1	DS4200 2, A1 (opt)	DS4200 2, B1 (opt)
4	Engine 1, slot 1 (opt)	Engine 2, slot 1 (opt)	Engine 1, slot 1 (opt)
5	Tape attachment	Tape attachment	Tape attachment
6	Tape attachment	Tape attachment	Tape attachment
7	Tape attachment	Tape attachment	Tape attachment
Additional B16 ports, included with Enhanced Remote Mirroring option only			
8	DS4200 1, A2	DS4200 1, A2	Unused
9	DS4200 1, B2	Unused	DS4200 1, B2

Port #	Single-Node FC Switch 1	Dual-Node FC Switch 1	Dual-Node FC Switch 2
10	ISL to Remote Site	ISL to Remote Site	ISL to Remote Site
11	Unused	Unused	Unused

Table 13-4 FC switch zoning

Zone #	Single-Node members	Dual-Node Sw1 members	Dual-Node Sw2 members
1	0 - Engine 1, slot 5 2 - DS4200 1, A1	0 - Engine 1, slot 5 2 - DS4200 1, A1 3 - DS4200 2, A1 (opt)	0 - Engine 1, slot 4 2 - DS4200 1, B1 3 - DS4200 2, B1 (opt)
2	1 - Engine 1, slot 4 3 - DS4200 1, B1	1 - Engine 2, slot 5 2 - DS4200 1, A1 3 - DS4200 2, A1 (opt)	0 - Engine 2, slot 4 2 - DS4200 1, B1 3 - DS4200 2, B1 (opt)
3	4 - Engine 1, slot 1 (opt) 5 - Tape attachment 6 - Tape attachment 7 - Tape attachment	4 - Engine 2, slot 1 (opt) 5 - Tape attachment 6 - Tape attachment 7 - Tape attachment	4 - Engine 1, slot 1 (opt) 5 - Tape attachment 6 - Tape attachment 7 - Tape attachment
Additional B16 zones, included with Enhanced Remote Mirroring option only			
4 (ERM)	8 - DS4200 1, A2 8 - DS4200 1, A2 Remote	8 - DS4200 1, A2 8 - DS4200 1, A2 Remote	Unused
5 (ERM)	9 - DS4200 1, B2 9 - DS4200 1, B2 Remote	Unused	9 - DS4200 1, B2 9 - DS4200 1, B2 Remote

For dual-node configurations, it is very important that switch 1 of the primary DR550 connects to switch 1 of the secondary and likewise for the second switches.

### 13.7.3 Enabling and activating Enhanced Remote Mirroring

Before you can use Enhanced Remote Mirroring, you must first purchase the option (feature key) from IBM for *all* storage subsystems participating in Remote Mirroring (it is included in the ERM option for the DR550). This section discusses the process for enabling, activating, and deactivating the Enhanced Remote Mirroring premium feature.

There are four possible states for the Enhanced Remote Mirroring option:

► Disabled/Deactivated

This is the default state if the feature key has not been installed.

► Enabled/Deactivated

This is the state when the feature key has been enabled, but Enhanced Remote Mirroring has not been activated. There are no changes made to the storage subsystem configuration and all host connectivity is normal.

► Enabled/Activated

This state makes controller ports A2 and B2 available for mirroring only. No host system I/Os are allowed. Up to 64 mirroring pairs on the DR550 Storage Controller can be created. This is the normal state for Enhanced Remote Mirroring operations.



► Disabled/Activated

This state is possible when the premium feature has been disabled before deactivation.

If mirrored pairs are present when this happens, an Out of Compliance error message will occur. Removing the mirrored pairs will remove this error. If the feature is Disabled/Activated, controller ports A2 and B2 will still be unavailable for normal host I/O activity. This is an undesirable state.

## Enabling the Enhanced Remote Mirroring premium feature

This section describes how to check the premium features list and enable Enhanced Remote Mirroring. Remember to enable the premium feature on all participating storage subsystems.

### Checking the premium features list

To check the premium option, follow these steps:

1. From the Subsystem Management window, select **Storage Subsystem** → **Premium Features** → **List**.

The List Premium Features dialog opens (Figure 13-19). It lists the premium features enabled and disabled on the storage subsystem.

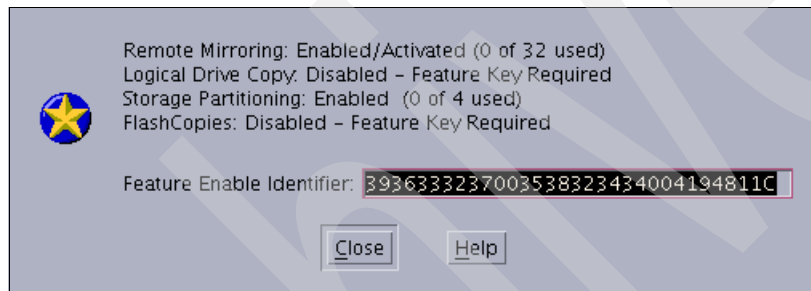


Figure 13-19 List premium features

### Enabling Enhanced Remote Mirroring feature

1. From the Subsystem Management window, select **Storage Subsystem** → **Premium Features** → **Enable**. The Select Feature Key File window opens. Choose the directory where you stored the \*.key file, then proceed. The Enable Premium Feature confirmation dialog displays. Click **Yes** to confirm.
2. To check that the Enhanced Remote Mirroring premium feature is enabled, you can view the icon in the lower left corner of the Subsystem Management window (Figure 13-25 on page 550). The red slash is removed.




Figure 13-20 Enhanced Remote Mirroring icon: Enabled/Deactivated

## Activating the Enhanced Remote Mirroring option

Use the Activate Remote Mirroring Wizard. Remember to activate the premium feature on all participating storage subsystems.

### Important:

- ▶ When the Remote Mirroring premium feature is activated, it creates two Mirror Repository logical drives per subsystem. Before activating the Remote Mirroring feature, verify that the number of configured logical drives on your storage subsystem is under the supported limit. This should normally be the case on any DR550.
- ▶ Ensure that the switch fabric is appropriately configured before beginning this procedure. Refer to 13.6, “SAN fabric connectivity” on page 536.
- ▶ Reservations on secondary logical drives are blocked. However, reservations on primary logical drives are allowed.

Activate the Enhanced Remote Mirroring feature as follows:

1. Select the **Storage Subsystem** → **Remote Mirroring** → **Activate** drop-down menu option, or select **Remote Mirroring** → **Activate** from the context menu. Refer to Figure 13-21.

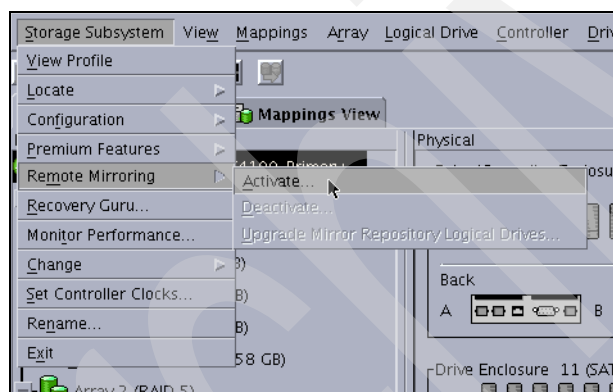


Figure 13-21 Activating Remote Mirroring

2. The Activate Remote Mirroring Wizard: Introduction dialog is opened (Figure 13-22 on page 549).

In this dialog, you choose the free capacity left for the mirror repository logical drives. Read the information in the window and click **Next** to proceed to the next window.

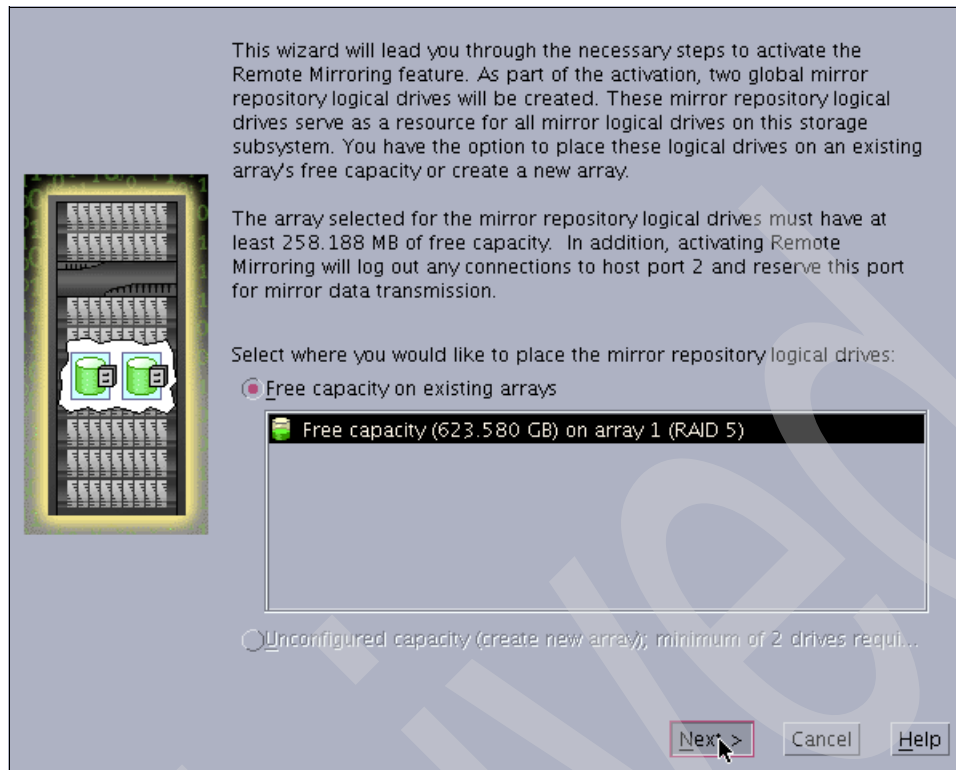


Figure 13-22 Activate Remote Mirroring Wizard: Introduction

3. After choosing where to create the mirror repository logical drives, the Activate Remote Mirroring Wizard Preview window appears (Figure 13-23).

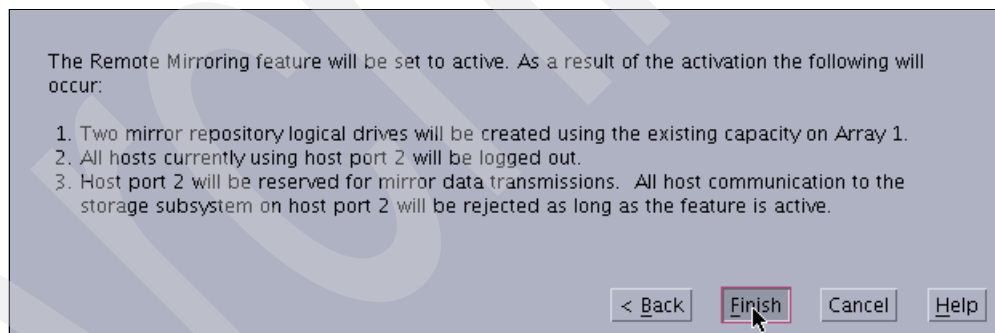


Figure 13-23 Remote Mirroring Wizard: Preview

4. Read the information and click **Finish** to continue.

5. The Activate Remote Mirroring Wizard Completed message is displayed (Figure 13-24). Click **OK** to close the dialog.

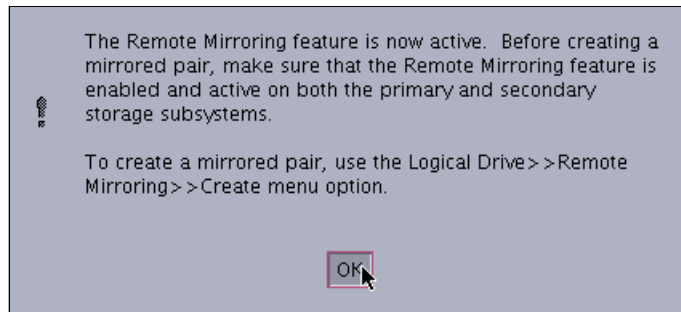


Figure 13-24 Activate Remote Mirroring Wizard: Completed

Two mirror repository logical drives, one for each controller, are created as a resource for the controller managing the Remote Mirror. Host port 2 on each controller is now dedicated to Remote Mirroring communication. The dedicated host ports will not accept read/write requests from a host application.

The Enhanced Remote Mirroring status icon at the lower left of the Subsystem Management window shows Enabled/Activated (Figure 13-25).




Figure 13-25 Enhanced Remote Mirroring icon: Enabled/Activated

### 13.7.4 Creating Enhanced Remote Mirroring relationships

Once you have enabled and activated the Enhanced Remote Mirroring option, you are ready to create Remote Mirror relationships. Before you create mirror relationships, the logical drive must exist at both the primary and secondary storage subsystems. The secondary logical drive on the remote storage subsystem must correspond exactly to the primary logical drive being mirrored. A mirrored pair needs to be established for each logical drive in the DR550.

The primary and secondary (remote) storage subsystems must be connected, as described in 13.6, “SAN fabric connectivity” on page 536.

The Create Remote Mirror wizard helps you quickly define a primary logical drive and a secondary logical drive and synchronize the data between the two logical drives.

Follow these steps:

1. To launch the Create Remote Mirror Wizard, select a logical drive in the Logical View of the designated primary storage subsystem and then select the **Logical Drive** → **Remote Mirroring** → **Create** pull-down menu option or select **Create Remote Mirror** from the context menu (Figure 13-26 on page 551).




Figure 13-26 Mirroring a logical drive

2. This launches the Create Remote Mirror: Introduction window (Figure 13-27).

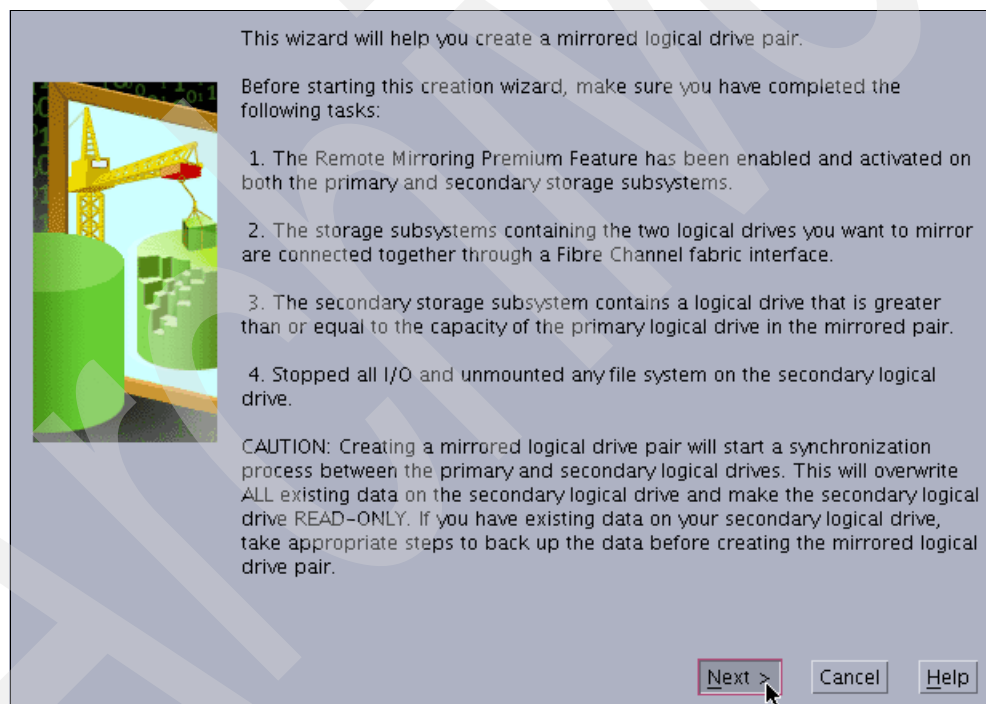


Figure 13-27 Create Remote Mirror: Introduction

3. Read the information displayed on the window and click **Next**.

- The next window displayed is the Create Remote Mirror Wizard: Selecting a storage subsystem for the Secondary Logical Drive (Figure 13-28).

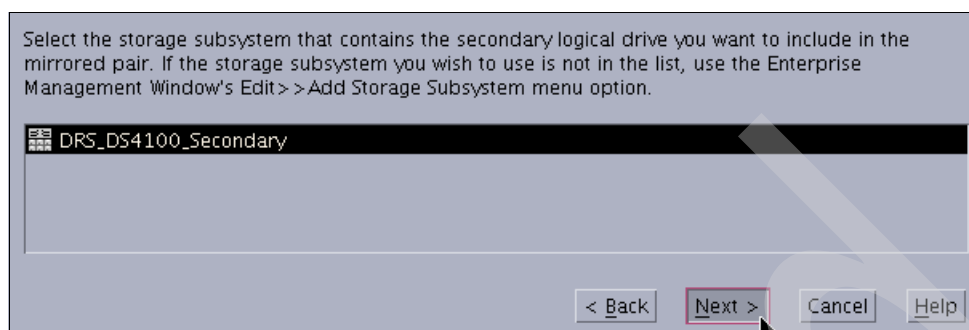


Figure 13-28 Create Remote Mirror: Select Storage Subsystem

- Select the storage subsystem that contains the secondary logical drive you want to include in the mirrored pair.
- Click **Next** to continue.
- The next window displayed lets you select a secondary logical drive for the Remote Mirrored pair (Figure 13-29).

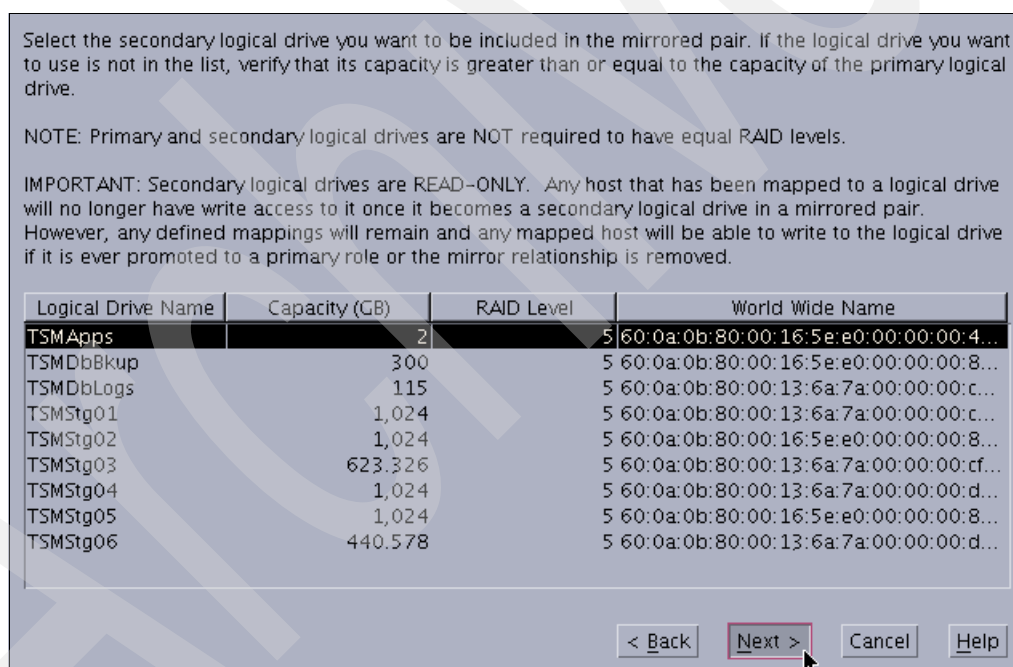


Figure 13-29 Create Remote Mirror: Select secondary logical drive

The logical drives shown in the table are sorted by capacity, starting with the logical drive nearest to the primary logical drive's capacity. Logical drives with identical capacity will be sorted alphabetically.

**Important:** The secondary logical drive at the remote site *must correspond exactly* to the primary logical drive on the primary DR550.

- Select the secondary logical drive to be included in the mirrored pair. Click **Next** to proceed.

- The next window prompts you to set the synchronization priority (Figure 13-30). Use this window to set the rate at which the controller owner of the primary logical drive will synchronize data with the secondary logical drive. The controller owner of the primary logical drive uses the synchronization priority, but a synchronization priority is also set on the secondary logical drive at this time (in case it is promoted to a primary logical drive, in a role reversal). Refer to 13.5.5, “Data synchronization priority” on page 536.

Figure 13-30 Create Remote Mirror: Set synchronization priority

- Select the synchronization priority and click **Next** to proceed (Figure 13-31).

**Note:** If you decide later that the synchronization priority is too low or too high, you can change it.

Figure 13-31 Create Remote Mirror - Write Mode Settings

11. Figure 13-31 on page 553 shows the Write Mode Settings window. Select the mode according to the mirroring requirements:
- Metro Mirror: Select **Synchronous write mode**.
  - Global Mirror: Select **Asynchronous write mode** and make sure the box labeled **Add to write consistency group** is checked.
  - Global Copy: This mode is *not supported* with the DR550.

The next window prompts you to confirm the creation of the remote mirror (Figure 13-32).

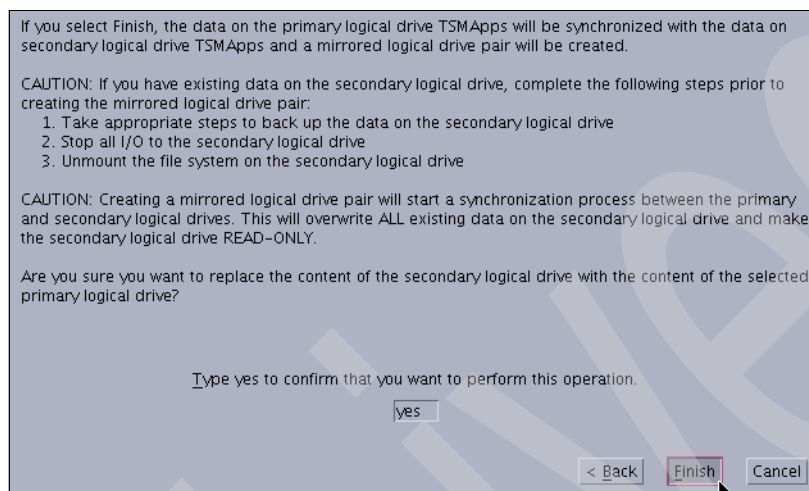


Figure 13-32 Create Remote Mirror: Preview

12. Select **Finish** to set the synchronization priority and finish the creation of the Remote Mirror.

The following message dialog is displayed (Figure 13-33):

The Remote Mirror was successfully created.




Figure 13-33 Create Remote Mirror: Creation successful

13. Click **Yes**.
14. Repeat the procedure for every logical drive on the primary storage subsystem. You will be taken back to the Create Remote Mirror: Select Primary Logical Drive window. If more capacity is added to the DR550 units, use this procedure to establish mirrored pairs for the new logical drives.



When you have no more mirrored pairs to create, the “Create Remote Mirror: Completed” message is displayed (Figure 13-34). Read the information and click **OK** to close the dialog.

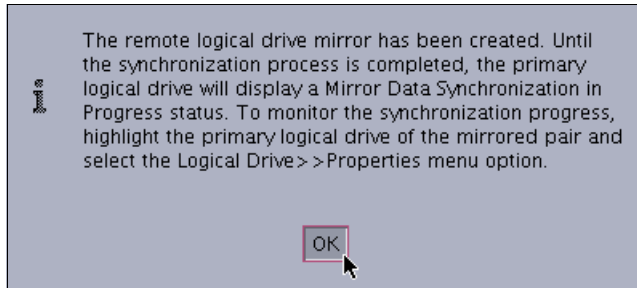


Figure 13-34 Create Remote Mirror: Completed

Looking at the logical view of the primary and secondary storage subsystems shows the synchronization in progress. On the primary subsystem, the secondary logical drive displays as a child of the primary logical drive, as indicated in Figure 13-35.

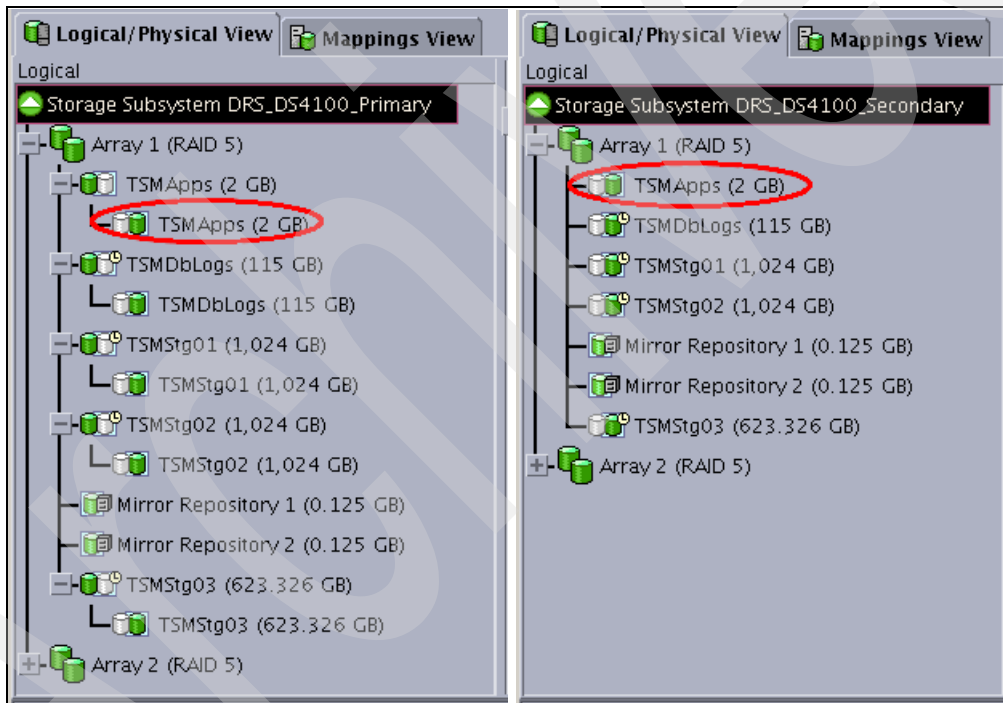


Figure 13-35 Primary and secondary storage subsystems logical view

**Important:** Do not restart the primary DR550 before all mirrored pairs are in the synchronized state. To check the status, refer to 13.7.5, “Viewing Enhanced Remote Mirroring properties and status” on page 556.

### 13.7.5 Viewing Enhanced Remote Mirroring properties and status

Select either the **Logical Drive** → **Properties** drop-down menu option, or **Properties** from the context menu to view the Logical Drive Properties dialog. If the Remote Mirroring premium feature is enabled and you have selected a primary logical drive or a secondary logical drive, then a Mirroring tab is available on this dialog that shows mirroring-related information (Figure 13-36).

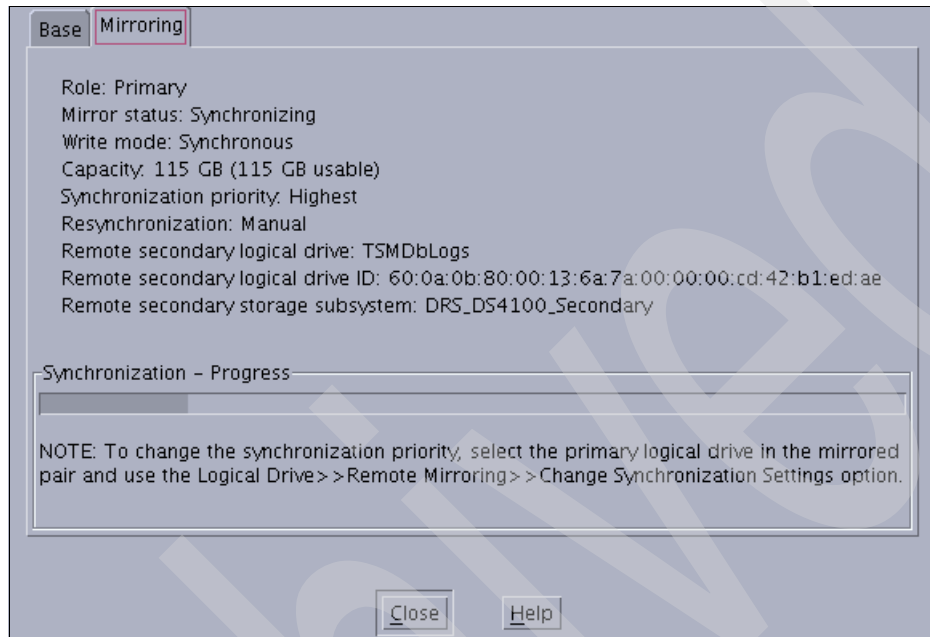


Figure 13-36 Viewing mirrored logical drive properties on the primary storage subsystem

The properties displayed include:

- ▶ Role: Primary or Secondary
- ▶ Mirror status (synchronized, synchronizing, suspended, or unsynchronized)
- ▶ Write Mode (always Synchronous for the DR550)
- ▶ Capacity, actual and usable
- ▶ Synchronization priority: Lowest, low, medium, high, or highest (not displayed for local or remote secondary logical drives)
- ▶ Resynchronization: Manual or Automatic
- ▶ Remote secondary (or primary) logical drive
- ▶ Remote secondary (or primary) logical drive ID
- ▶ Remote secondary (or primary) storage subsystem
- ▶ Synchronization progress

The progress bar at the bottom of the Mirroring tab of the Logical Drive Properties dialog displays the progress of a full synchronization. Select **Close** to exit the dialog.

Enhanced Remote Mirroring can also be monitored by viewing the Storage Subsystem Profile. To view the Storage Subsystem Profile, select either the **Storage Subsystem** → **View Profile** drop-down menu option, or **View Profile** from the context menu. The storage subsystem Profile dialog is displayed (Figure 13-37 on page 557).

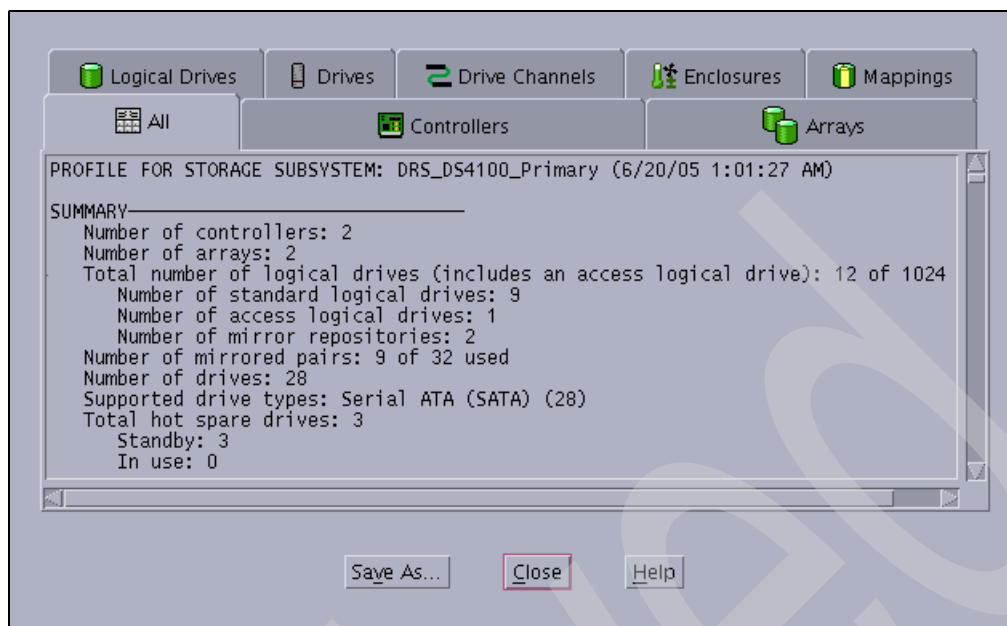


Figure 13-37 Storage Subsystem Profile: All tab

The All tab displays a summary of available Storage Subsystem Profile data and includes a summary of mirror repositories and the number of mirrored pairs in use.

The Logical Drives tab is divided into three sections with separate tabs providing access to information for each type of logical drive. The Repositories tab contains a section of logical drive information that displays mirror repository logical drive properties. See Figure 13-38.

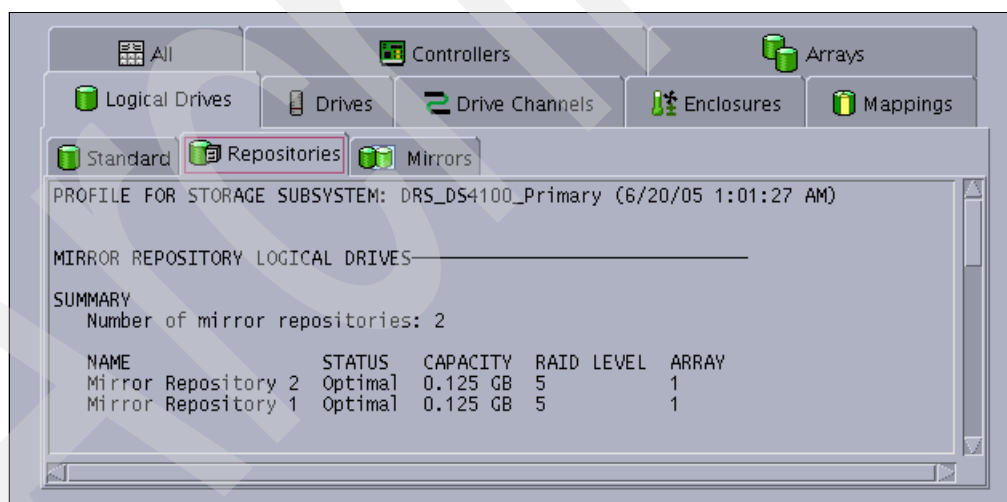


Figure 13-38 Storage Subsystem Profile: Repositories tab

The Mirrors tab is a section of logical drive information that displays Remote Mirror properties. The Mirrors tab is only available when the Remote Mirroring premium feature is enabled and activated. This tab includes a summary of all of the mirrored relationships followed by details on each mirrored pair. See Figure 13-39.

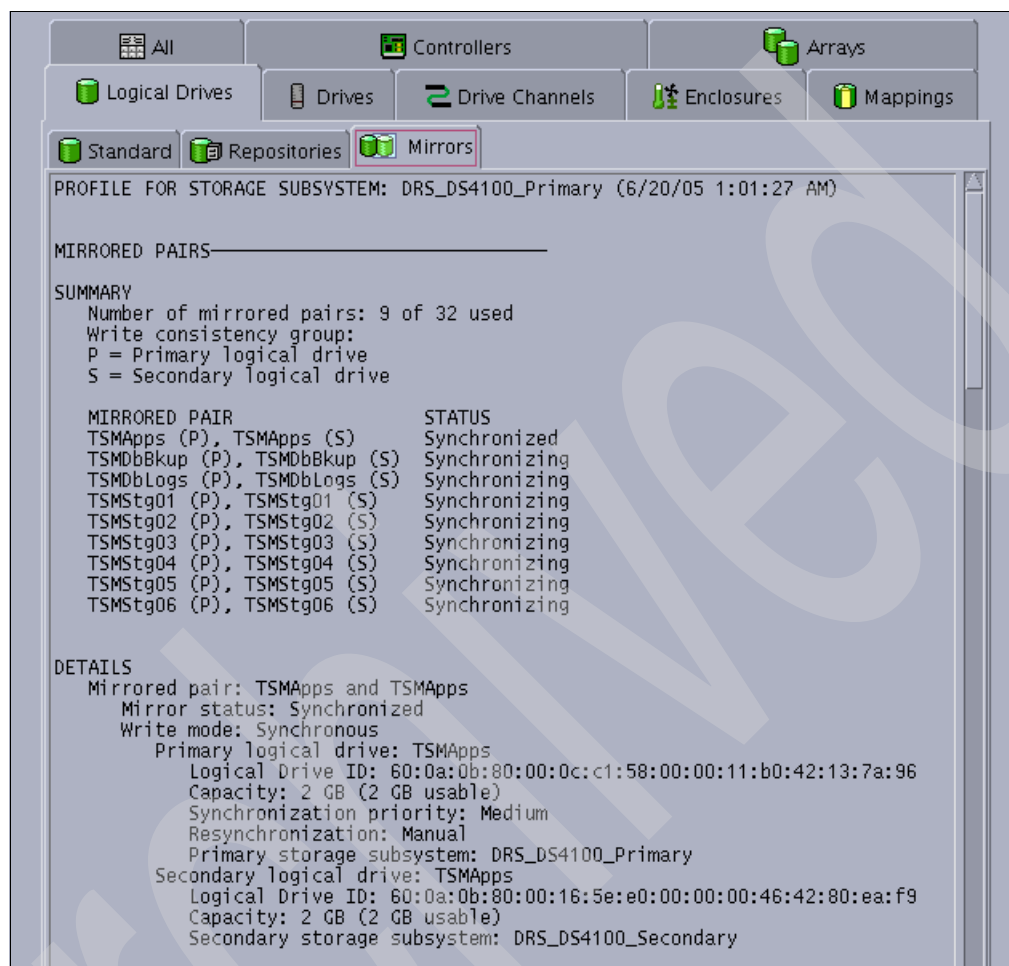


Figure 13-39 Storage Subsystem Profile: Mirrors tab

### 13.7.6 Mapping a secondary drive

A secondary mirror drive can be LUN-mapped to a host or host group like any standard virtual drive. The difference is that the mirrored drive is write-protected to any host system until it is promoted to a primary drive.

This option is useful because it enables the administrator to preconfigure the mapping of the secondary mirrored drives prior to the changing roles. This makes the transition in a site failover situation easier.

### 13.7.7 Suspend and resume a mirror relationship

The Suspend Mirror option allows you to stop the mirroring process independently of the current mirroring mode. In the Suspended state, all changed data blocks are logged (delta logging) in the mirror repository volume while the secondary logical drive does not receive any more write I/Os from the primary storage subsystem. Data transfer is suspended.

When the synchronous mirroring relationship on the DR550 (Metro Mirroring) is suspended, the primary storage controller reports the I/O completion to the hosts as soon as the write request has been logged to the mirror repository logical drive and written to the primary logical drive. No remote write request is initiated.

The Resume Mirror option is used to reenale the remote write I/O requests to the secondary logical drives. When a mirror relationship is suspended, the mirror repository logical drive collects the information about the changed data blocks. When the Resume Mirror command is issued, the primary controller starts transferring the changed data blocks to the remote controller.

### ***Suspend mirroring***

To suspend mirroring on the DR550:

1. Select a logical drive in the Logical View of the primary storage subsystem, and then select the **Logical Drive** → **Remote Mirroring** → **Suspend** drop-down menu option or select **Suspend Mirroring** from the context menu (Figure 13-40).

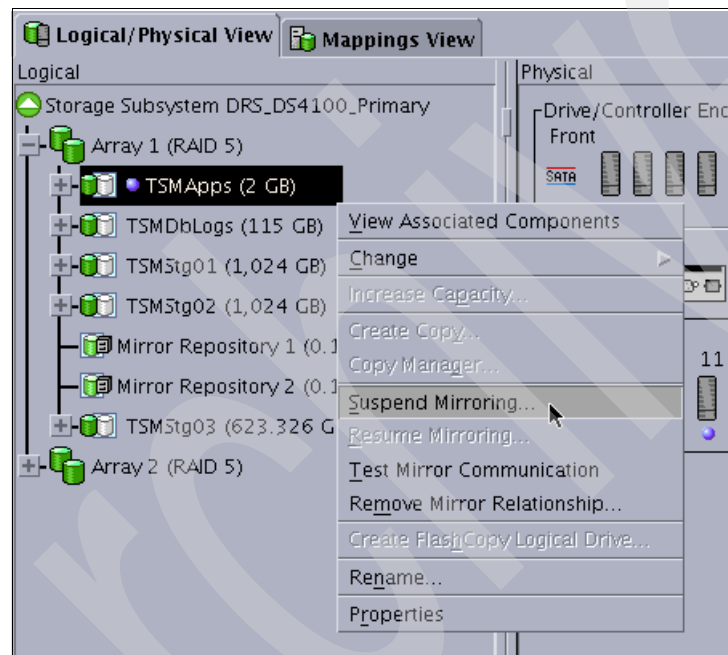


Figure 13-40 Suspend Mirroring

- The Suspend Mirrored Pair dialog is displayed. The dialog lists all primary logical drives in the storage subsystem. The primary logical drive that you initially selected is highlighted (see Figure 13-41). Select all of the primary logical drives by clicking **Select All** and click **Suspend** to proceed.

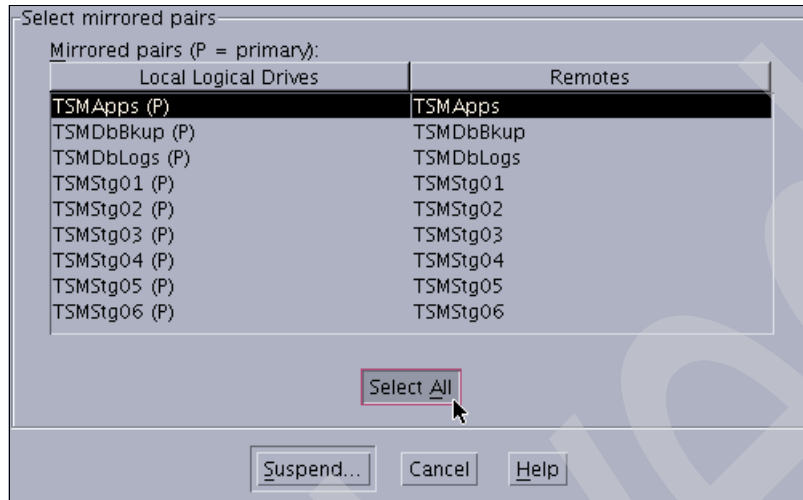


Figure 13-41 Suspend Mirroring: Select All mirrored pairs

- The Suspend Mirroring: Confirmation dialog is displayed (Figure 13-42).

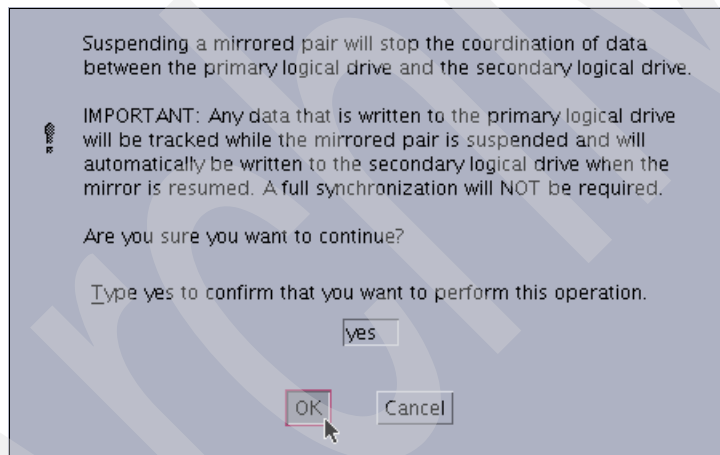


Figure 13-42 Suspend Mirroring: Confirmation dialog

- Type **Yes** into the text box and click **OK** to continue. The Suspend Mirrored Pair - Progress window displays. When all mirrored pairs are suspended, the progress window will indicate it has completed (Figure 13-43). Click **OK** to return to the Logical View.




Figure 13-43 Suspend Mirrored Pair: Completed

The suspended logical drives are indicated in the Logical View of the Storage Management window as shown in Figure 13-44 on page 561.

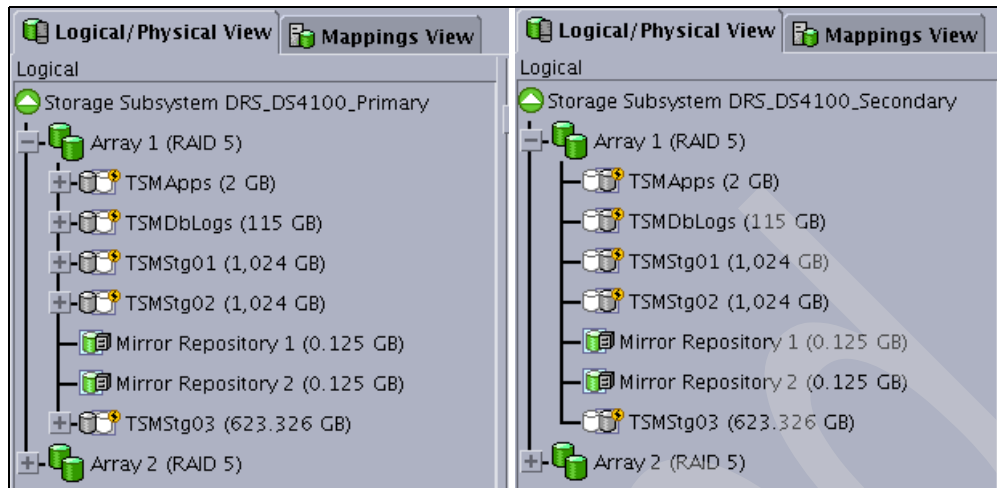


Figure 13-44 Suspended mirror pairs on primary (left) and secondary (right) storage servers

### Resume mirroring

To resume mirroring a logical drive:

1. Select the logical drive in the Logical View of the primary storage subsystem, and then select the **Logical Drive** → **Remote Mirroring** → **Resume** drop-down menu option, or select **Resume Mirroring** from the context menu (Figure 13-45).

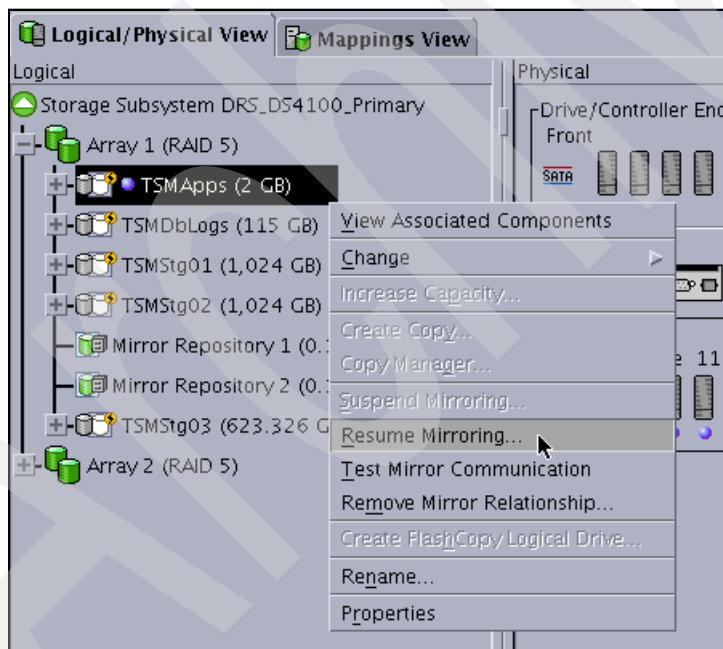


Figure 13-45 Resume mirroring on a logical drive

The Resume Mirror dialog is displayed. The window lists all primary logical drives in the storage subsystem. From here, the steps are identical to those for suspending the mirrored pairs.

## 13.8 ERM and disaster recovery

As modern business pressures increasingly require 24-hour data availability, system administrators are required to ensure that critical data is safeguarded against potential disasters. Additionally, storage administrators are searching for methods to migrate from one host to another or from one storage subsystem to another, with as little disruption as possible.

Remote Mirroring is one method that can be implemented to assist in business continuity and disaster recovery. Once critical data has been mirrored from a primary storage subsystem to a remote storage subsystem, primary and secondary logical drives can have their roles reversed so that the copied data can be accessed from the remote storage subsystem.

This section discusses how to reverse the roles of the primary and secondary logical drives. This section does *not* cover building and implementing a complete disaster recovery plan (other options available for the DR550 are also discussed in Chapter 12, “DR550 backup and restore” on page 491).

### 13.8.1 Role reversal concept

Role reversal promotes a selected secondary logical drive to become the primary logical drive of the mirrored pair. The roles of primary and secondary logical drives are naming conventions based on the direction of data flow. They differentiate as follows:

- ▶ Relevant administrative commands for ERM must be provided on the primary site.
- ▶ Mirror states are determined by the primary storage subsystem.
- ▶ Connection examination is provided by the primary storage subsystem.
- ▶ The secondary logical drive is not write-accessible.

A role reversal is performed using one of the following methods.

#### Changing a secondary logical drive to a primary logical drive

This option promotes selected secondary logical drives to become the primary logical drives of the mirrored pairs. If the associated primary logical drives can be contacted, the primary logical drives are automatically demoted to be the secondary logical drives. Use this option when a normally interruptible maintenance task on the primary site is needed or in the case of an unrecoverable failure to the storage subsystem that contains the primary logical drive, and you want to promote the secondary logical drive so that hosts can access data and business operations can continue.

##### Important:

- ▶ When the secondary logical drives become primary logical drives, the secondary DR550 will now be able to read or write to the logical drives.
- ▶ If a communication problem between the secondary and primary sites prevents the demotion of the primary logical drive, an error message is displayed. However, you are given the opportunity to proceed with the promotion of the secondary logical drive, even though this will lead to a Dual Primary Remote Mirror status condition. For more information, see “Forcing a secondary to change to a primary logical drive” on page 563.

Select a secondary logical drive in the Logical View, then select either the **Logical Drive → Remote Mirroring → Change → Role to Primary** drop-down menu option or **Change → Role to Primary** from the context menu (Figure 13-46 on page 563). This process needs to be repeated for each logical drive pair on the DR550.



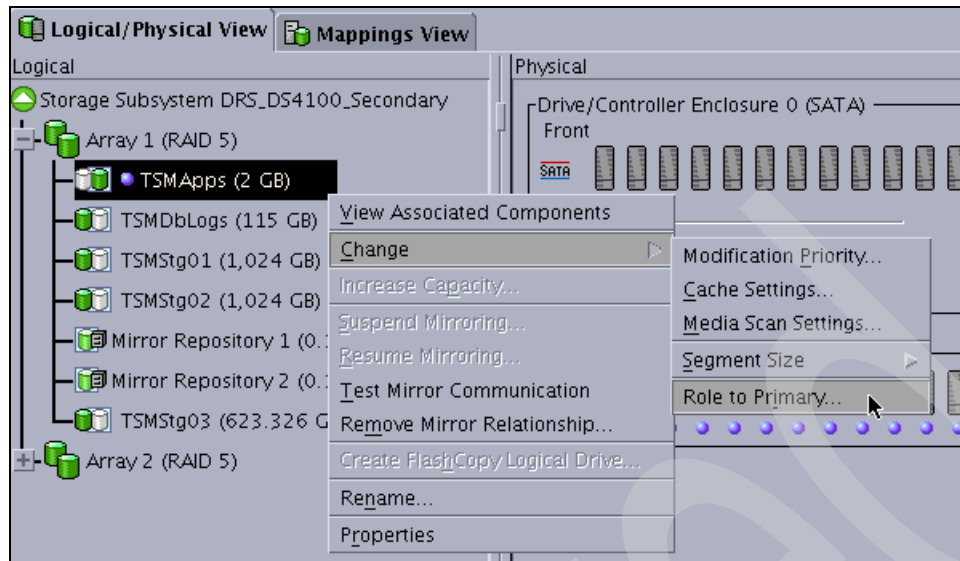


Figure 13-46 Role Reversal: Secondary to primary

The Change to Primary dialog is displayed. Select **Yes** to proceed. The secondary logical drive is promoted to be the primary logical drive in the Remote Mirror. If the controller owner of the primary logical drive can be contacted, the primary logical drive is automatically demoted to be the secondary logical drive in the Remote Mirror.

After this has been completed, bring up the secondary DR550 unit.

### Forcing a secondary to change to a primary logical drive

If, when attempting to promote a secondary logical drive to a primary logical drive, there is a communication failure between the primary and remote storage subsystems, then an error dialog will appear (Figure 13-47).

Select **Yes** on the Change to Primary: Error dialog to force the software to promote the selected secondary logical drive to a primary role.

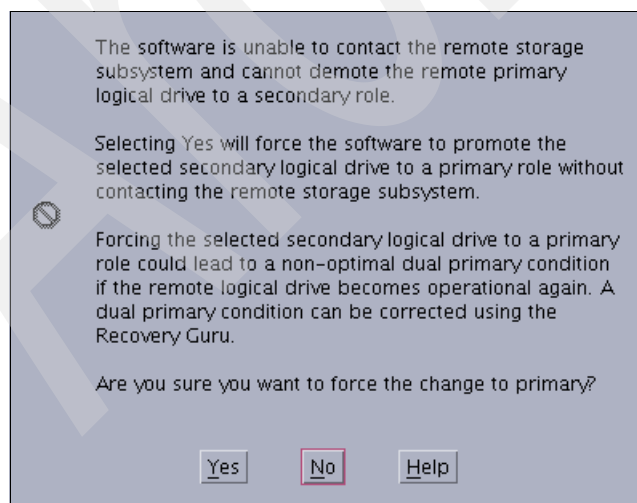


Figure 13-47 Force secondary to primary: Confirmation

Use this option when a catastrophic failure has occurred on the storage subsystem that contains the primary logical drive, the primary storage subsystem cannot be contacted, and you want to promote the secondary logical drive so that hosts can access data and business operations can continue. Or consider using this option when no catastrophic failure has occurred to the storage subsystem that contains the primary logical drive, but you want to perform a role reversal for some other reason and a communication problem between the local and remote sites is preventing the demotion of the primary logical drive at the remote site.

When the remote storage subsystem has recovered and communication problems have been resolved, a Dual Primary error condition will be raised. Click the **Recovery Guru** to resolve the condition (see Figure 13-48). Follow the procedure to remove and recreate the mirror relationship.

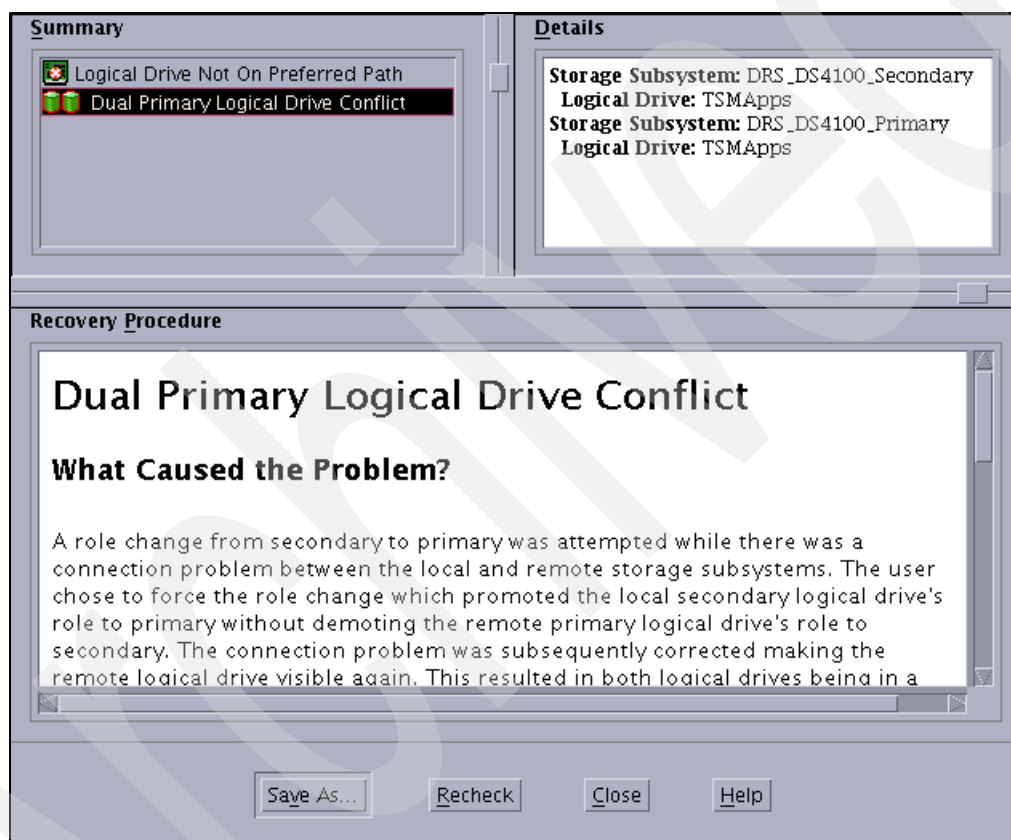


Figure 13-48 Recovery Guru: Dual primary logical drive conflict

To avoid the Dual Primary error condition and subsequent recovery steps when no catastrophic failure has occurred to the storage subsystem containing the primary logical drive, wait until the connection between the storage subsystems is operational to perform the role reversal.

### Changing a primary to a secondary logical drive

Use this option to perform a role reversal between the two paired logical drives in a Remote Mirror. This option demotes a selected primary logical drive to become the secondary logical drive of the mirrored pair.

**Important:** The primary DR550 will no longer be able to write to the logical drives.

To demote a primary logical drive to the secondary logical drive role, select the primary logical drive in the Logical View, then select either the **Logical Drive → Remote Mirroring → Change → Role to Secondary** drop-down menu option or select **Change → Role to Secondary** from the context menu. The Change to Secondary Confirmation dialog is displayed. Read the information and select **Yes** to proceed.

The primary logical drive is demoted to be the secondary logical drive in the Remote Mirror. If the controller owner of the secondary logical drive can be contacted, the secondary logical drive is automatically promoted to be the primary logical drive in the Remote Mirror.

If a communication problem between the primary and secondary sites prevents the promotion of the secondary logical drive, an error message is displayed. However, you are given the opportunity to proceed with the demotion of the primary logical drive, even though this will lead to a Dual Secondary Remote Mirror status condition.

### **Forcing a primary to become a secondary logical drive**

If, when attempting to demote a primary logical drive to a secondary logical drive, there is a communication failure between the primary and remote storage subsystems, then an error dialog will appear.

Select **Yes** on the Change to Secondary: Error dialog to force the software to demote the selected primary logical drive to a secondary role.

Consider using this option when you want to perform a role reversal and a communication problem between the local and remote sites is preventing the promotion of the secondary logical drive at the remote site.

## **13.8.2 Reestablishing Remote Mirroring after failure recovery**

When the damaged site is back online and properly configured, mirrored relationships can be resumed. Recreate a mirror relationship by completing the following steps:

1. Ensure that SAN connections and SAN zoning are properly configured.
2. From the active secondary site, define a mirror relationship using the logical drive on the recovered primary site as the secondary logical drive. For more information, see 13.7, “Enhanced Remote Mirroring on DR550: Step-by-step” on page 538.
3. Ensure that storage partitioning and host mapping are properly defined on the recovered primary site so that it can take over normal operations from the secondary site.
4. Ensure that the host software is properly configured so that the host systems at the recovered primary site can take over I/O from the secondary site host systems.
5. After the full synchronization has completed, perform a manual role reversal so that the recovered primary site now possesses the active primary logical drives, and the secondary logical drives now exist on the secondary site. For more information, see “Changing a primary to a secondary logical drive” on page 564.

## 13.9 Bringing up the secondary DR550: Step by step

Whether to recover from a catastrophic failure of the primary site or for maintenance purposes, the most important factor for a DR550 with an ERM configuration is the ability to resume business activities at the secondary site. This section walks you through the steps required to do this. As you will see, most of the individual steps have already been covered in this and earlier chapters. The two main categories of failure scenarios are:

- ▶ Controlled failures for maintenance purposes
- ▶ Unforeseen failures such as natural disasters, hardware failures, or human error

In either of these situations, almost all of the steps are the same with the exception of the role reversal on the secondary DR550 Storage Controller.

To return to the primary site once it is backed up, follow the same procedure. If the data has been lost at the primary site, you may have to reestablish the mirrored relationships and resynchronize the volumes according to 13.7.4, “Creating Enhanced Remote Mirroring relationships” on page 550 before proceeding with the steps below.

### 13.9.1 Single-node DR550

The steps to recover from a primary site failure to the secondary site may include:

- ▶ Preparing the secondary DR550 to recognize the mirrored drives
- ▶ Role reversal of secondary logical drives
- ▶ Powering on the DR550 node at the secondary site
- ▶ Changing secondary IP addresses
- ▶ Starting Tivoli Storage Manager server

#### Prepare the secondary DR550

This step should be performed *once* before failing over to the secondary DR550 (that is, before any controlled or unforeseen failures happen). We suggest you perform this step right after you initially configure and establish the mirror relationships.

Do the following:

1. Log on to the p5 server with root authority.
2. Export the original volume groups (TSMApps, TSMDbLogs, TSMDbBkup, and TSMStg) using the command **exportvg <volumegroup-name>**.
3. Delete all hdisks pertaining to the DR550 Storage Controller through the command **rmdev -d1 <hdisk>**.
4. Import the volume groups (TSMApps, TSMDbLogs, TSMDbBkup, and TSMStg) that are on the mirrored drives using the command:

```
importvg -V <major-number> -y <volume-group-name> <hdisk-name>.
```

The *major-number* of the volume group can be obtained from the p5 server at the primary site with the command **ls -l /dev/TSM**. This command lists the major number as shown in bold in the example below:

```
root@drs_engine1: /> ls -l /dev/TSM*
crw-r----- 1 root system 46, 0 Apr 30 2004 /dev/TSMApps
crw-r----- 1 root system 44, 0 Apr 30 2004 /dev/TSMDbBkup
crw-r----- 1 root system 43, 0 Apr 30 2004 /dev/TSMDbLogs
crw-r----- 1 root system 42, 0 Jul 26 18:54 /dev/TSMStg
```

The *volume-group-name* is the name of the volume group to be imported and the *hdisk-name* is the name of the hdisk on which this volume group resides. These parameters are identical to the primary site.

5. Vary on the volume groups (TSMApps, TSMDBLogs, TSMDBBkup, and TSMStg) that have been imported in the previous step by performing the command **varyonvg <vg-name>**.
6. Vary off the volume groups (TSMApps, TSMDBLogs, TSMDBBkup, and TSMStg) on all System p servers (use the command **varyoffvg <vg-name>**).
7. Start Tivoli Storage Manager (SSAM) on the secondary system. (Refer to 4.2, “Starting and stopping IBM System Storage Archive Manager” on page 148.)
8. Perform a checkout of the secondary DR550. (Refer to 4.4, “DR550 check procedures summary” on page 152.)
9. Once successfully checked out, stop SSAM on the secondary system. (Refer to 4.2, “Starting and stopping IBM System Storage Archive Manager” on page 148.)

### Role reversal of secondary logical drives

In a controlled failure, the logical drive roles can be reversed, as detailed in “Changing a secondary logical drive to a primary logical drive” on page 562. For maintenance on the primary DR550, you may also want to suspend the mirrored relationships between the sites while maintenance activities are performed.

In an unforeseen failure, the secondary logical drives will need to be forced to the role of primary. See “Forcing a secondary to change to a primary logical drive” on page 563 for details.

### Powering on the DR550 node at the secondary site

If the secondary DR550 node is kept powered off, you will need to power it on per the instructions in step 5 of the power on sequence in “DR550 Model DR2 single-node configuration” on page 79.

### Changing secondary IP addresses

If the secondary DR550 engine has been set up with a different IP address to avoid network conflict with the primary DR550, you will need to change the network configuration to match the primary node so your application servers (that is, DR550 client applications such as a Document Management System) will not need to be reconfigured. To do this, follow the steps in 3.6.4, “Configuring the IP network for a DR550 DR1 and DR2 single-node” on page 94.

### Starting Tivoli Storage Manager server

Finally, the Tivoli Storage Manager server will need to be started. Instructions about this can be found in 4.2, “Starting and stopping IBM System Storage Archive Manager” on page 148.

## 13.9.2 Dual-node (HACMP) DR550

The steps to recover from a primary site failure to the secondary site may include:

- ▶ Preparing the secondary DR550
- ▶ Role reversal of secondary logical drives
- ▶ Powering on the DR550 nodes
- ▶ Changing the secondary IP addresses
- ▶ Starting HACMP Cluster services

## Prepare the secondary DR550

This step should be performed *once* before failing over to the secondary DR550 (that is, before any controlled or unforeseen failures happen). We suggest you perform this step right after you initially configure and establish the mirror relationships.

Do the following:

1. Log on to the p5 server with root authority.
2. Export the original volume groups (TSMApps, TSMDBLogs, TSMDBBkup, and TSMStg) using the command **exportvg <volume-group-name>**.
3. Delete all hdisks pertaining to the DR550 Storage Controller through the command **rmdev -d1 <hdisk>**.

4. Import the volume groups (TSMApps, TSMDBLogs, TSMDBBkup, and TSMStg) that are on the mirrored drives using the command:

```
importvg -V <major-number> -y <volume-group-name> <hdisk-name>
```

The *major-number* of the volume group can be obtained from the p5 server at the primary site with the command **ls -l /dev/TSM**. This command lists the major number, as shown in bold in the example below:

```
root@drs_engine1: /> ls -l /dev/TSM*
crw-r----- 1 root system 46, 0 Apr 30 2004 /dev/TSMApps
crw-r----- 1 root system 44, 0 Apr 30 2004 /dev/TSMDBBkup
crw-r----- 1 root system 43, 0 Apr 30 2004 /dev/TSMDBLogs
crw-r----- 1 root system 42, 0 Jul 26 18:54 /dev/TSMStg
```

The *volume-group-name* is the name of the volume group to be imported and the *hdisk-name* is the name of the hdisk on which this volume group resides. These parameters are identical to the primary site.

5. Vary on the volume groups (TSMApps, TSMDBLogs, TSMDBBkup, and TSMStg) that have been imported in the previous step by performing the command **varyonvg <vg-name>**.
6. Perform steps 1 - 5 for the other p5 node.
7. Vary off the volume groups (TSMApps, TSMDBLogs, TSMDBBkup, and TSMStg) on both engines using the command **varyoffvg <vg-name>**.
8. Perform a HACMP synchronization and verification by running **smitty hacmp** and **selecting Extended Configuration → Extended Synchronization and Verification**.
9. Start Tivoli Storage Manager (SSAM) on the secondary system. (Refer to 4.2, “Starting and stopping IBM System Storage Archive Manager” on page 148.)
10. Perform a checkout of the secondary DR550. (Refer to 4.4, “DR550 check procedures summary” on page 152.)
11. Once successfully checked out, stop SSAM (Tivoli Storage Manager) on the secondary system. (Refer to 4.2, “Starting and stopping IBM System Storage Archive Manager” on page 148.)
12. Ensure that HACMP is not started. (Refer 4.4.4, “HACMP basic check” on page 154.)

## Role reversal of secondary logical drives

In a controlled failure, the logical drive roles can be reversed, as shown in “Changing a secondary logical drive to a primary logical drive” on page 562. For maintenance on the primary DR550, you may also want to suspend the mirror relationships between the site while you perform maintenance activities.

In an unforeseen failure, the secondary logical drives will need to be forced to the role of primary. See “Forcing a secondary to change to a primary logical drive” on page 563 for details.

### Powering on the DR550 nodes

If the secondary DR550 nodes are not already powered on, follow the instructions (sequence specific) in steps 6 and 7 of “DR550 Model DR2 dual-node configuration” on page 80.

### Changing the secondary IP address

We suggest that you set up the same IP addresses on the secondary DR550 engines, assuming the cluster is not connected to the Ethernet network at the same time as the primary DR550 cluster. You would just have to connect the secondary DR550 to the network when the primary has failed and has been disconnected.

### Starting HACMP cluster services

Finally, start up HACMP cluster services according to 4.1.2, “Starting cluster services” on page 140. This will also automatically start the Tivoli Storage Manager server.

## 13.10 ERM maintenance

This section explains controller behavior for some common failure situations. It should be used as a general rule for understanding how ERM handles the errors.

### Unsynchronized state indication

A failure can place the mirror relationship in an unsynchronized mirror state. This state is reported by the DS4000 Storage Manager for DR550 as a Needs Attention situation. You can get a further error description from the Recovery Guru diagnostic tool.

**Note:** Always use the Recovery Guru to check a Needs Attention situation. The Recovery Guru provides both an explanation of the error and a procedure and hints for recovery.

### Link interruptions

Loss of communication is most often caused by FC link failure.

#### ***Fibre Channel mirror link interruptions in synchronous write mode***

In synchronous mode, if the link is interrupted and the primary controller receives a write request from an attached host, the write request cannot be transferred to the secondary logical drive, and the primary and secondary logical drives are no longer appropriately mirrored. The primary controller transitions the mirrored pair into an *unsynchronized* state and sends an I/O completion to the primary host. The host can continue to write to the primary logical drive but remote writes will not take place.

When connectivity is restored between the controller owner of the primary logical drive and the controller owner of the secondary logical drive, depending on the configured resynchronization method, either an *automatic resynchronization* takes place or a *manual resynchronization* (recommended) must be performed. Only the data in changed blocks will be transferred during the resynchronization process. The status of the mirrored pair changes from an *Unsynchronized* state to a *Synchronization-in-Progress* state.

**Note:** A loss of communication between one of the controller pairs in the primary and secondary DR550 Storage Controller does not result in the automatic failover of the logical drives on those controllers to the alternate controller pair. The logical drive mirrors owned by that controller pair will break and become unsynchronized. The storage can be resynchronized once communication on the failed link is reestablished.

The only time logical drive ownership changes is on a host path failure. The secondary DR550 Storage Controller will also change ownership of the mirrored logical drives on the next I/O to match the primary DR550 Storage Controller.

### **Fibre Channel Mirror Link Test function**

A Fibre Channel Mirror Link Test function is included with Tivoli Storage Manager V9.12. To invoke the function, right-click the primary logical drive, and select **Test Mirror Communication**. The result displays, represented by the image of a traffic light.

### **Secondary logical drive error**

The primary controller also marks the mirrored pair as unsynchronized when a logical drive error on the secondary site prevents the remote write from completing. For example, an offline or a failed secondary logical drive can cause the Enhanced Remote Mirror to become *Unsynchronized*. The host can continue to write to the primary logical drive, but remote writes will not take place. Once the logical drive error is corrected (the secondary logical drive is placed online or recovered to an Optimal state), then the resynchronization process can begin. Depending on Synchronization Settings, the resynchronization can be started automatically or manually (recommended).

### **Primary controller failure**

If a remote write is interrupted by a primary controller failure before it can be written to the secondary logical drive, the primary storage subsystem provides controller ownership change from the preferred controller owner to the alternate controller in the storage subsystem, the first write request to the remote site changes the ownership of the secondary logical drive. Once the transition of the ownership is completed, the mirroring proceeds as usual.

### **Primary controller reset**

If a remote write is interrupted by a primary controller reset before it can be written to the secondary logical drive, there is normally no ownership change. After reboot, the controller reads information stored in a log file in the mirror repository logical drive and uses the information to copy the affected data blocks from the primary logical drive to the secondary logical drive. We highly recommend that you suspend mirror relationships before resetting the controller.

### **Secondary controller failure**

If the secondary controller fails, the primary site can no longer communicate with the secondary logical drive. The mirror state becomes *Unsynchronized*. The host can continue to write to the primary logical drive, but remote writes will not take place. After the secondary controller has been recovered following a failure, the primary controller changes the mirror state to Synchronizing; this occurs automatically or manually (recommended), depending on the synchronization settings.



## 13.11 Performance considerations

This section contains general performance considerations that apply to a DR550 remote mirror configuration.

### ***Synchronization priority***

The controller owner of a primary logical drive performs a synchronization in the background while processing local I/O write operations to the primary logical drive and associated remote write operations to the secondary logical drive.

Because the synchronization diverts controller processing resources from I/O activity, it can have a performance impact to the host application. The synchronization priority defines how much processing time is allocated for synchronization activities relative to system performance. The following priority rates are available:

- ▶ Lowest
- ▶ Low
- ▶ Medium
- ▶ High
- ▶ Highest

You can use the synchronization priority to establish how the controller owner prioritizes the resources needed for synchronization process relative to host I/O activity. The following rules are rough estimates for the relationships between the five synchronization priorities.

**Note:** Note that logical drive size can cause these estimates to vary widely.

A synchronization at the lowest, low, medium, and high priorities will take approximately eight, six, three and a half, and two times as long (respectively) as one at the highest priority. A synchronization at the highest priority will take approximately 6-10 hours for each terabyte of storage in the DR550 with a direct inter-switch link (ISL) averaging 40 MBps.

**Note:** The lowest priority rate favors system performance, but the synchronization will take longer. The highest priority rate favors the synchronization, but system performance may be compromised. Logical drive size and host I/O rate loads affect the synchronization time comparisons.

### ***Mirroring connection distance and performance***

The maximum distance that can be supported whether using short or long wave SFPs is greater in a 1 Gbps fabric than in a 2 Gbps (refer to Table 13-1 on page 537). However, using a 1 Gbps connection, rather than a 2 Gbps, negatively impacts the synchronization performance.

The major performance factor for long distance solutions is the IP connection bandwidth between the Fibre Channel Extenders. If we connect to a 100 Mbps IP network, we will be able to transfer about 10 MBs of data per second (assuming the connection is stable).

- ▶ Estimate the time needed for the first full synchronization.
- ▶ Check the amount of data that is changing on average on the primary logical drives over a certain period of time and compare it to the available connection bandwidth.
- ▶ Always keep some reserve for unexpected transfer rate peaks.

Archived

# Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

## IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 575. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *AIX 5L Version 5.3 Installation Guide and Reference*, SC23-4887-01
- ▶ *Content Manager Backup/Recovery and High Availability: Strategies, Options, and Procedures*, SG24-7063
- ▶ *Content Manager Implementation and Migration Cookbook*, SG24-7051
- ▶ *DS4000 Best Practices and Performance Tuning Guide*, SG24-6363
- ▶ *Implementing High Availability Cluster Multi-Processing (HACMP) Cookbook*, SG24-6769
- ▶ *IBM System p5 520 and 520Q Technical Overview and Introduction*, REDP-4137
- ▶ *IBM System Storage DS4000 and Storage Manager V10.10*, SG24-7010
- ▶ *Implementing an IBM/Brocade SAN*, SG24-6116
- ▶ *IBM System Storage Tape Library Guide for Open Systems*, SG24-5946
- ▶ *IBM Tivoli Storage Management Concepts*, SG24-4877
- ▶ *IBM Tivoli Storage Manager Implementation Guide*, SG24-5416
- ▶ *IBM Tivoli Storage Manager Versions 5.4 and 5.5 Technical Guide*, SG24-7447
- ▶ *Integrating IBM Tivoli Storage Manager Operational Reporting with Event Management*, REDP-3850

## Other publications

These publications are also relevant as further information sources:

- ▶ *AIX 5L Version 5.3 Communications Programming Concepts*, SC23-4894
- ▶ *Copy Services Users Guide - IBM System Storage DS4000 Storage Manager*, GC27-2172
- ▶ *Gigabit Ethernet-SX PCI-X Adapter and Dual Port Gigabit Ethernet-SX PCI-X Adapter Installation and Using Guide*, SA23-1302
- ▶ *Gigabit Fibre Channel PCI Adapter 2 Gigabit Fibre Channel Adapter Installation and Using Guide*, SA23-2550
- ▶ *HACMP for AIX 5L V5.4.1 Administration Guide*, SC23-4862
- ▶ *HACMP for AIX 5L V5.4.1 Planning Guide*, SC23-4861
- ▶ *IBM DB2 Content Manager OnDemand for Multiplatforms: Administration Guide*, SC18-9237
- ▶ *IBM eServer 7014 Series Model T00 and Model T42 System Rack Installation Guide*, SA38-0641

- ▶ *IBM System Storage DR550 File System Gateway Software Version 1.1.1 Installation Guide*, GC27-2123
- ▶ *IBM System Storage DR550 File System Gateway Software Version 1.1.1 Integration Guide*, GC27-2124
- ▶ *IBM System Storage DR550 File System Gateway Software Version 1.1.1 Maintenance Guide*, GC27-2126
- ▶ *IBM System Storage DR550 File System Gateway Software Version 1.1.1 User Guide*, GC27-2125
- ▶ *IBM System Storage DR550 Version 4.5 Installation Roadmap*, GC27-2175
- ▶ *IBM System Storage DR550 Version 4.5 Introduction and Planning Guide*, GA32-0577
- ▶ *IBM System Storage DR550 Version 4.5 Problem Determination and Service Guide*, GA32-0576
- ▶ *IBM System Storage DR550 Version 4.5 User's Guide*, GC27-2174
- ▶ *IBM System Storage DS4200 Storage Server Installation, User's and Maintenance Guide*, GC27-2048
- ▶ *IBM Tape Device Drivers Installation and User's Guide*, GC27-2130-00
- ▶ *IBM Tivoli Storage Manager for AIX Administrator's Guide*, SC32-0117
- ▶ *IBM Tivoli Storage Manager for AIX Administrator's Reference*, SC32-0123
- ▶ *IBM Tivoli Storage Manager for AIX Installation Guide*, SC32-0134
- ▶ *IBM Tivoli Storage Manager for AIX Quick Start Version 5.3*, GC32-0770
- ▶ *IBM Tivoli Storage Manager Basic Concepts Poster*, SC32-9464
- ▶ *IBM Tivoli Storage Manager Messages*, SC32-0140
- ▶ *IBM Tivoli Storage Manager Performance Tuning Guide*, SC32-9101
- ▶ *IBM Tivoli Storage Manager Using the Application Program Interface*, SC32-0147
- ▶ *IBM TotalStorage Enterprise Tape System 3592 Operator Guide*, GA32-0465
- ▶ *IBM TotalStorage SAN16B-2 Installation, Service, and User's Guide*, GC26-7753
- ▶ *Installation, User's and Maintenance Guide - DS4000 EXP420 Storage Expansion Enclosure*, GC27-2050
- ▶ *Tivoli Storage Manager Installing the Clients*, SH26-4119

## Online resources

These Web sites are also relevant as further information sources:

- ▶ IBM System Storage and TotalStorage products  
<http://www.ibm.com/systems/storage/index.html>
- ▶ IBM System Storage DR550  
<http://www.ibm.com/systems/storage/disk/dr/index.html>
- ▶ IBM TotalStorage DS4700 Storage Server  
<http://www.ibm.com/systems/storage/disk/ds4000/ds4700/>
- ▶ IBM TotalStorage FAStT EXP810 Storage Expansion Unit  
<http://www.ibm.com/systems/storage/disk/ds4000/exp810/>

- ▶ IBM Tivoli products  
<http://www.ibm.com/software/tivoli/products>
- ▶ Content Management  
<http://www.ibm.com/software/data/cm/>
- ▶ Regulatory compliance  
[http://www.ibm.com/software/data/cm/solutions\\_compliance.html](http://www.ibm.com/software/data/cm/solutions_compliance.html)

## How to get IBM Redbooks

You can search for, view, or download IBM Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

[ibm.com/redbooks](http://ibm.com/redbooks)

## Help from IBM

IBM Support and downloads

[ibm.com/support](http://ibm.com/support)

IBM Global Services

[ibm.com/services](http://ibm.com/services)

Archived

# Index

## Numerics

2005-B16 10, 16  
2229-FSG 11, 26  
2233-DR1 21  
2233-DR2 10  
30-net\_persistent\_names.rules 118  
3494 456–457, 460  
3588 tape drive 459  
3589 Tape Cartridge 460  
3592 165, 188, 456, 459–460  
7014 T00 10  
7014-S25 rack 22  
7014-T00 rack 12  
9131-52A 12

## A

AC power 69, 73  
accept date 90–91  
accept time 87  
Access Definition 342  
Activate 172  
activating the Enhanced Remote Mirroring option 548  
Active Directory 72, 283  
    accounts 286  
ADE Console 239  
ADE Console Interface 236  
Administration Center 68, 71, 161  
administrative interface 162  
Advanced Encryption Standard (AES) 21  
AES 21  
Agent  
    See IBM Director, Agent  
AIX Administrator  
    IBM Tivoli Storage Manager 497  
AIX error logging 420  
AIX Logical Volume Manager 49  
AIX user ID 31  
alias 543–544  
API 158  
API function 171, 173  
application encryption 181  
application managed encryption 180  
Application Monitoring 103  
Application program interface (API) 159  
Application Server Monitoring 103  
Archive 341  
archive copy group 167, 170–171, 177, 199, 201, 310, 331  
    reinit parameter 201  
Archive File 342  
archive retention 169  
    chronological 171  
    event based 171  
archiveretentionprotection 170

array 48, 190  
ARSMaint 337  
ASN 420  
Asynchronous Mirroring 523  
Atape 462, 464  
atldd 464–465  
atom 233  
attribute 232  
attribute XML files 233  
attributes 126, 306  
AUDIT VOLUME 177  
Auditable plan 495

## B

BA Client 204  
    install 205  
backup 448  
backup volume 189  
BM System Storage Archive Manager 20, 25  
bonding 117  
boot address 38, 97, 458  
bundle import 130  
bypass-err.log 267

## C

cable  
    type 537  
cabling  
    Ethernet 18  
    Fibre Channel 18  
cache 230  
cache data 336  
CAD 207  
call home 64  
CAT5 18, 67, 70  
CDE 47  
cfgmgr 466  
chronological archive retention  
    archive copy group 172  
chronological management class  
    first object 218  
chronological retention policy  
    simplified view 173  
CIFS 230  
CIFS share 254  
CIMOM 361  
cleaning frequency 478  
client 360  
Client Acceptor Daemon 207  
Client Acceptor Service 207  
client API 20  
client node 159  
climate control 67, 69  
clockd 138

- clsmuxpd 419
- clstart 139
- clstop 140
- clstrmgrES 139
- Cluster Manager 426
- cluster services 138, 140, 144
- command line 470
- command line client 170
- COMMETHOD 429
- Common Internet File System (CIFS) 27
- CommonStore for Exchange Server 305
- CommonStore for Lotus Domino 305
- community 419, 436
- community name 392
- compliance 3
- compliance data 1
- component
  - Backup 234
  - Client Services 234
  - Events 234
  - Replication 234
  - Resources 234
  - Storage 234
  - Timing 235
- components 232
- concurrent access 98
- configuration 10
- Configure CIFS share 254
- configure share 248
- Consistency Group 523, 533
- Contact us xiii
- container 233
- Content Engine 309
- Content Handle Release Inhibit (CHRI) 235–236
- content management
  - application 6, 188
- Content Manager 330
- Content Manager OnDemand 305
- Content Manager System Administration Client GUI 316
- control path failover 462
- controller 51
  - ownership 525
- controller ownership 527
- copy group 167, 172
  - archive 172, 199
- copy pool 3, 165, 450
- copy storage pool 450
  - tape 495
  - volume 495–496, 498, 501
- Core Services 361
- CRC 482
- crm\_mon 128, 131, 154
- crontab 187
- CRU 15
- C-SPOC utility 139–140
- current 127
- Cygwin/X 84

## D

- daemon 390, 418
- dapismp 154, 214, 216, 221
- data
  - retention-managed 3–4
- data archive 307
- data life cycle 4
- data mining 4
- data model 306
- data object 163
- data rendering 4
- Data Retention 1–2, 4–6, 19, 28, 160, 169, 173, 196, 214
  - Compliance 5
  - IBM Tivoli Storage Manager 158–159, 163, 165, 169, 175–176, 203
  - policy 88, 112, 176
  - protection 177
  - server 203
  - SSAM 170
  - System Storage Archive Manager 5, 170
- data shredding 1, 4, 21, 183
- data synchronization priority 536
- database backup (DB) 189, 194, 493, 495, 500
- datashred.sh 186
- Daylight Savings Time (DST) 88
- DB2 Content Manager 305
- DBBackuptrigger 189
- DBBKUP 189
- default management
  - class 168, 172, 198, 215–216, 218–219
- Delay Period for Volumes Reuse 481–482
- deletion
  - hold 21, 175, 200
  - release 175
- deletion protection 235, 239
- DELVOLH 189
- device class 164–165, 168, 194, 450, 454
  - DBBKUP 166, 189, 194
  - DISK 166, 168, 194
- device driver 453
- device type 165
  - SERVER 176
- directory 245, 276
- Directory Information Tree (DIT) 276
- directory tree 230
- Disaster Recovery
  - Manager command 494
  - Manager query command 494
  - Manager volume management 493
  - starter kit 497
- disaster recovery 492–497
- Disaster Recovery Manager (DRM) 491–492, 497
- disk drive 6, 12, 22, 24, 48, 54, 159, 192
- disk heartbeat 101
- disk space allocation 52
- distance limitations of Fibre Channel inter-switch links 537
- Distributed Protocol Interface (DPI) 427
- DNS 119
- Document Manager 305
- Domain ID 543–544
- dpid2 418



- DR550 31, 465
- DR550 Disk Controller 10
- DR550 DR2 Rack 12
- DR550 DR2 rack 10
- DR550 engine 505, 509
- DR550 Ethernet Switch 10
- DR550 Expansion Drawer 10, 13
- DR550 File System Gateway 11
- DR550 File System Gateway (FSG) 9, 25, 76
  - overview 230
- DR550 Model DR1 21
- DR550 Model DR2 10
- DR550 offering 9, 20–21
  - Overview 9
- DR550 SAN Switch 10, 16, 73
- DR550 SSAM Server 10, 12
- DR550 Storage Controller 10, 13, 63
- dr550adm 31
- drg 132, 135
- DRG profile 72, 237
- DRG profile editor 241
- drg restart 129
- DRG service 125, 232
- DRG service components 233
- drg stop 267
- DRG\_NODE 200
- DRGC 129
  - attributes 126
- DRGC bundle 126, 259
- DRG-DEFAULTMC 199
- DRG-DOMAIN 197
- DRG-NODE 114
- DRG-SYSMC 199
- DRIVEENCRYPTION 181
- DRM instruction 497
- drs\_cluster\_svc 38
- DS4000 Storage
  - Manager 48
    - Manager 8.47 Client application 45
    - Manager Client 20, 25, 50
    - Manager Version 8.47 20, 25
- DS4000 Storage Manager 20, 25
- DS4100 24
- DS4100 controller
  - failure 51
- DS4100 Logical Drive
  - layout 502
- DS4100 Storage
  - Server 159, 191, 501
- DS4200 13
- DS4200 Storage
  - Server 52
- DS4700 14, 24
- DS4700 Storage Server Model 70 13
- dsm.opt 204–205, 314
- dsmadm 170, 470
- dsmserv 89
- dsmserv.opt 429
- DST 88
- dual node

- configuration 35, 52, 188–189, 503
- configurations scale 10
- disk-and-tape-ready DR550 37
- setup 37
- system 503, 507, 509
- DVD drive 503, 505, 507, 509
  - blank DVD-RAM media 503, 507
  - DVD media 505, 509
- DVD-RAM 29, 509
- DVD-RAM Device 508

## E

- echo \$TZ 88
- ECM 303
- ECM portfolio 306
- Electronic Service Agent 64
- enable\_drg 125
- encryption 1, 179
  - application managed 180
  - tape drive 21
  - transparent 180
- encryption key 180
- ENCRYPTIONTYPE 179
- Enhanced Remote Mirroring 39
  - option 548
- Enhanced Remote Mirroring (ERM) 448, 521
- Enterprise Content Management (ECM) 303–304
- Environmental Service Module 14
- ERM
  - activating 546, 548
  - enabling 546
  - resume 561
  - suspend 560
- Errorlog 398
- ESM 14
- Ethernet switches 10
- event retver 201, 203
- Event-based 21
- eventlogging 430
- EXP100 unit 6, 12, 18
- EXP420 13–14
- expansion enclosure 15
- EXPINTERVAL 178
- expiration date 173–174, 177, 201, 220–223
- expiration processing 178
- expire inventory 179

## F

- fabric 17
- fabric configuration 536
- failback 232
- failover 144, 154
- fallback 129, 144
- FC HBA
  - failure 51
- feature key 546
- Federation Services 308
- fiber optic 33
- Fibre Channel 72–73

- connection 39
- Host Bus Adapter 51
- switch 15–16
- technology 28
- Fibre Channel (FC) 30
- Fibre Channel inter-switch link 537
- file system 164, 166, 189–190, 194, 504, 508
  - disk storage pool volumes 166
- File System Gateway (FSG) 56, 160
- File Transfer Protocol 30
- FileNet Business Process Manager 305
- FileNet Content Manager 305, 308
- FileNet Image Manager Active Edition 305
- FileNet P8 5, 305, 308
- FileNet Records Manager 305
- firewall 57, 257
- FOPT 236, 239
- forcebackup 132
- FPRF 236
- FSG 25
  - cache 230
  - maintenance 257
  - physical components 230
  - replacement 273
  - status 234
- FSG profile 114
- FSG start stop 151
- FSG status 154
- fsgadm 87, 135
- fsgtool 268
- fsid 251–252
- FTP 32
- full synchronization 523, 527

## G

- getent 282
- GID 286
- Gigabit Ethernet 33, 67
- Global Copy 523
- Global Mirror 523, 533

## H

- H08 Switch
  - licensed fabric management service 43
- HACMP 12, 19, 52, 88, 96–97, 138, 188, 458
  - shutdown 146
- HACMP cluster
  - service 503, 507, 509
- HACMP failover 457
- hacmpadm 31–32, 134
- halt 80–81, 150
- Hardware Management
  - Console 29
- heartbeat 19, 57, 76, 101, 231
- Hierarchical Storage Management (HSM) 165
- High Availability Cluster Multi-Processing (HACMP) 19
- HIPAA 2
- HMC application 505
- Hold 172

- Host Bus Adapters (HBA) 51
- host names 37–38
- host port 550
- Hot-spare 50
- Hyperterm 39

## I

- I/O ports 29
- IBM Content Manager 5, 305
- IBM Director 87
  - Agent 361
  - architecture 360
  - components 360
  - Core Services 361
  - Server 360
- IBM Optim 341
- IBM Statistical Analysis and Reporting System (SARS) 456
- IBM System Storage Archive Manager 20
  - main focus 20
- IBM System Storage Archive Manager (SSAM) 157
- IBM System x 3650 56
- IBM Tivoli Storage Manager 158, 164–165, 168–169, 174, 176, 201, 492–495, 497, 501
  - Administration Guide 195
  - Administrator 188, 574
  - API 28, 159, 204, 223
  - architecture overview 158
  - database 451, 496
  - Extended Edition 20, 176, 492
  - fundamental 158
  - HSM 165
  - macro 495
  - new features 214
  - overview 158
  - performance 195
  - policy 175
  - Policy concept 166
  - Recovery 190
  - Server 158–159, 163, 165, 194, 214, 217, 495
  - server code 496
  - server option 495
  - server status 170
  - server store 163
  - storage hierarchy 164
  - terminology 167
  - usage 191
  - Version 5.2.2 171, 175
- IBM Tivoli Storage Manager database
  - backup 52, 165, 191, 495–496
- IBM TotalStorage DR550 1
  - key software component 158
  - main focus 6
- IBM TotalStorage Enterprise Tape Drive 3592, See 3592
- ibmce 31
- ICMANNOTATION 323
- ICMBASE 323
- ICMBASESTREAM 324
- ICMBASETEXT 324
- ICMNOTELOG 324

- Identity Management 284
- ifcfg-bond0 118, 120
- import 127, 165
- imported 127
- indexed search 306
- Information Lifecycle Management (ILM) 1, 4
- Integrated Solution Console 161
- Integrated Solution Console (ISC) 68, 71
- inter-switch link 537
- invalid 127
- IP address 33, 39, 91, 93–94, 104, 214
- IP address takeover 100
- IP alias 101
- ISC 161–162
- item type 323

## J

- JFS2 189

## K

- kdestroy 299
- kill 89
- KVM 63, 82, 87

## L

- last match 243
- Last Update 215
- Last Update Date/Time 215
- LDAP 72, 245, 276
  - Bind Account 285
  - client 280
  - defining users and groups 278
- LDAP packages 277
- Level-1 agent 361
- Level-2 agent 361
- Library encryption 181
- library server 323
- library WEB interface 182
- Lightpath 26
- LMCP 471
- lmcpd 465
- Local Site 525
- logical drive 46
- logical network interface 99
- logical partition 462
- logical volume 54, 163–164, 166, 191–192, 510
  - Preserve Physical Partitions 511
- Logical Volume Manager (LVM) 53–54, 192
- lssrc -g cluster 147
- LVM 54, 192

## M

- main 113, 231
- managed system 360
  - See also* IBM Director, Agent
- management class 160, 167–168, 172, 177, 198, 201, 215, 222
  - retention policies 214

- Tivoli Storage Manager server 216
- management console 360
  - See also* IBM Director, Console
- Management Information Base (MIB) 390
- management server 360
  - See also* IBM Director, Server
- media management 496
- medium changer 468
- Metro Mirror 523
- Mgmt class 174, 215, 218, 220–223
  - override
    - event 218
- MIB 390, 392, 420, 427
- Microsoft Active Directory 245, 276
- Migrating data 448
- migration 451
- mining 4
- mirror relationship 528, 556
  - recreating 565
- mirror repository 525, 527–528, 549, 570
- mirror status 529
- mirroring 48
- mkcd command 504, 508
- mksysb 451–452
- mosy 420
- mount 257
- mtlib 464

## N

- NASD 5
- near-line storage 165
- netmask 99
- NET-SNMP 391
- Network File System (NFS) 27
- Network Information System (NIS) 276
- Network Installation Management (NIM) 451
- Network Time Protocol (NTP) 71
- New Folder 244
- NFS 230
- NFS export 248
- nfsserver 252
- NIM 452
- NMS 390
- No-Delete 235
- NOLIMIT 171, 174
- non-concurrent access 98
- NTP 71, 122–123
- ntpd 124

## O

- Object ID 220–222
- Object Server 330
- ODM (Object Database Manager) 420
- off-site (OS) 497
- OmniFind Enterprise Edition 305
- OnDemand 305, 307
  - Administrator 334, 336
  - Configurator 333
- on-site location 494

- Operational Reporting 223
- oplock 253
- oplocks 253
- Optim Archive 5
- orphaned files 270
- OU structure 283
- ownership 283, 527

## P

- PACS 25
- parallel loading 237
- partitioning 48, 59
- password 32, 68, 70, 133
- passwordaccess 207
- PBKI.xml 132
- PCA (Power Control Assembly) 18
- PCI-X adapter 29
- PCI-X slot 93
- permission 283
- Picture Archiving Communication Solutions (PACS) 1, 25
- PID 144
- policy 88, 112, 163, 176, 196
  - default settings 196
- policy domain 159, 167–168, 176, 196, 310, 331
  - APITEST1 215
  - STANDARD 196
- policy set 167–168, 170, 172, 176, 197, 201, 215
- post-installation 137
- POWER GXT135P Graphics Adapter 29
- power off 77, 80–81
- power on 77–78
- power rails 73
- power supply 67, 69
- Predictive Failure Analysis (PFA) 364
- preferred path 51
- primary logical drive 525–526
- Primary Site 525
- primary storage pool
  - archivepool 55, 193–194
  - recovery 495
- profile 556
- PuTTY 83, 87

## Q

- query process 151
- query sessions 151

## R

- rack 10
- Rack Security Kit 12
- RAID 64
- RAID 6 48, 59, 65
- RAID levels 48
- Records Manager 5, 305
- recovery 269
- recovery log 189–190
- recovery testing 497

- Redundant Array
  - of Independent Disks 48
- regulated information 2
- Release 172
- reliability, availability, and serviceability (RAS) 13
- Remote Client Agent Service 207
- Remote Mirror relationships 550
- Remote Mirror status 529
- Remote Support Manager (RSM) 64, 68, 70
- Remote Volume Mirror 48
  - status 556
- rendering 4
- replication errors 267
- replication group 232
- repository logical drive 528
- reservation 548
- Resource Manager 306
- resources 125
  - resources.main 125
  - resources.sngl 125
  - resources.suppl 125
- restored volume group
  - file system 511
- Resume Mirror 523, 559, 561
- resynchronization 535
- retention event 170, 173, 214
- retention initiation (RETINIT) 171–172, 174, 177, 199, 201, 203, 214–215, 220, 223
- retention life cycle 28
- retention period 2, 4, 21, 167, 230
  - chronological class 172
- retention policy 4, 164, 172–174, 188, 201, 203
- retention-managed data 1, 3–4
- RETMIN 170–173, 177
- RETVAR 170–171, 173, 177, 220
- revert 242
- RIO 29
- rlogin 32
- rmprofile 242
- role 526
- role reversal 562, 567
- RSCT 98, 101
- RSCT Resource Monitoring and Control 103
- rsh 32
- RSM 68
- RVM
  - distance limitations 537
  - mirror relationships 556
  - recreating a mirror relationship 565
  - status 556
  - storage subsystem profile 556
  - switch zoning 537
  - using 538

## S

- Samba 253
- SAN Switch 39
  - B16 31, 43
  - H08 40, 43
- Sarbanes Oxley Act 2

- SATA 3, 6, 14, 28, 165
- savevg 451–452
- schedule 162–163, 179, 189, 194
- schedule DELVOLH 500
- SCSI tape drive 24
- SDK 313
- SEC 2, 5, 28
- SEC/NASD 2
- secondary logical drive 526
- Secondary Site 525
- Secure Shell (SSH) 83
- security 68, 70, 133
- segment size 50
- sendmail 32
- Serial Advanced Technology Attachment (SATA) 14
- serial port 67, 70
- Server
  - See IBM Director, Server
  - server recovery strategy 495
  - service 232
  - service address 38, 97, 99
  - session 150
  - SFP 40
  - shredding 1, 4, 183
  - shutdown mode 146
  - Simple Network Management Protocol (SNMP) 389
  - single node 10
  - site preparation 63
  - SLES 27
    - firewall 57
  - SLP 361
  - SM Client 31, 45
  - small and medium business (SMB) 17
  - small form-factor pluggable (SFP) 51
  - smb.conf 253
  - smbd 253
  - smbpasswd 258
  - SMclient 46–47, 52, 191
  - SMIT 94, 103–104, 142
  - smitty 94, 104
  - smitty chtz date 88
  - SMUX 390, 419–420, 426
  - SNMP 389–390
  - SNMP agent 390–391
  - SNMP manage 390
  - SNMP subagent 426
  - SNMP trap 391, 393
  - SNMP trap request 390
  - SNMPADMIN 427
  - snmpd.conf 418
  - snmpinfo 420
  - SPCN 29
  - spooler 32
  - squash 245
  - SRC 139
  - SSAM 7, 157, 170, 177
    - archived data 177
    - data retention protection 176
    - database 52, 174, 189
    - database size 189
    - deletion protection 177
    - features 176
    - policy 170
    - server database 195
    - volume 54, 192
  - SSH 87
  - ssh 31, 84
  - standards
    - CIMOM 361
    - SLP 361
    - WMI 361
  - standby address 97
  - standby mode 80–81
  - startserver.foreground 89
  - startsrc 139
  - startx 82, 106–107
  - stat cache size 253
  - states 529
  - status 234
  - Storage Archive Manager Application Program Interface (SSAM API) 7
  - Storage Area Network (SAN) 43, 159
  - Storage Groups 321
  - storage hierarchy 164, 450, 454
  - Storage Manager 20, 25
  - storage pool 52, 159, 163–165, 168, 172, 189, 450, 454, 472, 483, 493–494, 499
    - complete recovery 493
  - Storage Profile 343
  - storage subsystem profile 556
  - striping 48
  - su 31, 95
  - supplementary 113, 231
  - SUSE Linux 26
  - Suspend Mirror 523, 558
  - switch zoning 537
  - synchronization 556, 570
    - priority 553
  - synchronization priority 553
  - synchronous mirroring 523
  - System encryption 181
  - System Resource Controller 138
  - System Storage Archive Manager
    - accounts 134
    - Administrator 188
    - API 20, 175
    - application 19
    - archive function 170
    - Client API 22
    - database 52, 189
- T**
  - tape 448
  - tape device 166, 454
  - tape drive 24
  - Tape drive encryption 180
  - tape library 159, 494–495
    - locked room 494
  - tape pool 450
  - tape volume 164, 166

- tape-based storage 5
- tapeutil 464
- TB capacity 49, 51–52
- TCP/IP address 37
- TCPServeraddress 111
- telnet 30, 32
- third-party certification 17, 42
- thresholds 480
- time server 123
- Time Zone 88
- time zone 87, 113, 122
- Tivoli Enterprise Console 427
- Tivoli Storage Manager
  - BA Client 204
- Tivoli Storage Manager for System Backup and Recovery 451
- Tivoli Storage Manager SysBack 502
  - further details 503
- Tivoli tab 343
- transparent encryption 180
- trap 390–391, 393
- TS1040 182
- TS1120 165, 181, 188, 448, 454
- TSMApps 507–508
- TSMApps volume group 503, 507, 509
- TSMDbBkup 52, 54, 191–192
- TSMDbLogs 54, 192

## U

- UDF 503
- Universal Disk Format (UDF) 503, 507
- UPS 64
- USB 29
- USB cable 18
- use SMIT 503, 507
- user accounts 30
- User ID
  - admin 31
- Users and Groups administration 247

## V

- vmstat 1 149
- volume group 54, 102, 191–192, 499, 501, 504, 508, 510–511
  - design 49
  - offline 511
  - TSMApps 55, 193, 502–503, 509–510
  - TSMStg 55, 193

## W

- Web Client 206
- winbind 127
- Windows Explorer 235
- WMI 361
- Workstation Collections 322
- WORM 21, 235, 449
- WORM Data Cartridge 449
- write cache 185

- write caching 531
- Write Once Read Many (WORM) 3, 6

## X

- X11Forwarding 84
- xedit 107
- xhost 84
- XML 27, 127
- X-server 83

## Y

- yast 116

## Z

- zone 537
- zoning 43, 524, 543



# IBM System Storage DR550 V4.5 Setup and Implementation

(1.0" spine)

0.875" <-> 1.498"

460 <-> 788 pages









# IBM System Storage DR550 V4.5 Setup and Implementation



**Redbooks**

## **New machine type model (2233)**

## **Enhanced RAS features**

## **Storage enhancements**

This IBM Redbooks publication explores the characteristics of the various DR550 models with details about how the different elements are initially configured. It explains and illustrates the additional configuration tasks required to deploy the 2233-DR1 and 2233-DR2 products in an existing network and storage environment.

This book also presents key features of IBM System Storage Archive Manager, which is the core of the DR550 and covers the DR550 File System Gateway option, which is designed to provide NFS and CIFS file system access to applications. It explains how the DR550 series integrates with IBM Enterprise Content Management products to deliver an efficient, fully managed data retention environment.

We also explain how to complement the offering with tape attachment and the use of WORM tape media for migration and backup purposes. The book also discusses disaster recovery considerations and has a chapter dedicated to the Enhanced Remote Mirroring option.

Finally, we cover call-home features and the use of SNMP and IBM Director for monitoring the DR550.

This publication is intended for those who want an understanding of the DR550 and also for readers who need detailed advice about how to configure and deploy the solution.

## **INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION**

## **BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:**  
[ibm.com/redbooks](http://ibm.com/redbooks)

SG24-7091-05

ISBN 0738431494