


Getting Started with Tivoli Storage Manager: Implementation Guide

How to plan and set up your Tivoli Storage Manager server from scratch

Includes worksheets, scripts, and macros to make your job easier

Tips for both new and experienced administrators



Mathis Landzettel
Arnold Balingit
Ross Battaglia
Charlotte Brooks
J.P. Houle
Armando Lemos da Silva Filho
Rod MacLeod
Andy Pattinson
Patrick Randall
Phil Thomas



International Technical Support Organization

SG24-5416-01

**Getting Started with Tivoli Storage Manager:
Implementation Guide**

August 2000

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix C, "Special notices" on page 301.

Second Edition (August 2000)

This edition applies to Version 3 Release 7 of Tivoli Storage Manager for AIX, Solaris, HP-UX and NT, program 5697-TSM; Version 3 Release 1 of Tivoli Storage Manager for AS/400, program 5697-TSM; Version 3 Release 7 of Tivoli Storage Manager for MVS-OS/390, program 5697-TSO; and Version 3 Release 1 of Tivoli Storage Manager for VM, program 5697-TSO as well as complementary Tivoli Storage Management Solutions.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. 471F Building 80-E2
650 Harry Road
San Jose, California 95120-6099

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Figures	xi
Tables	xiii
Preface	xv
The team that wrote this redbook	xv
Comments welcome	xvii
Chapter 1. Tivoli Storage Manager implementation checklists	1
1.1 The big picture	1
1.2 Redbook support material	2
1.3 Planning checklist	5
1.4 Server implementation checklist	6
1.5 Client implementation checklist	8
1.6 Operations checklist	9
Chapter 2. Tivoli Storage Manager implementation planning	11
2.1 Client environment data	11
2.1.1 Client name	13
2.1.2 Contact information	14
2.1.3 Operating system	14
2.1.4 Total storage available	14
2.1.5 Total storage used	14
2.1.6 GB changed per backup	14
2.1.7 Number of files backed up	15
2.1.8 Number of backup versions kept	15
2.1.9 Data compression	15
2.1.10 Backup window times	16
2.1.11 Backup number of hours	16
2.1.12 Required recovery time frame	16
2.1.13 Tivoli Storage Manager restore time frame	17
2.1.14 GB copied per archive	17
2.1.15 Number of files archived	17
2.1.16 Number of archives kept	17
2.1.17 Archive frequency	17
2.1.18 Archive window times	17
2.1.19 Archive number of hours	18
2.1.20 Number of image backups	18
2.1.21 Image backup frequency	18
2.1.22 Number of backup sets	18
2.1.23 Backup set frequency	18
2.1.24 Policy domain	18
2.1.25 Client option set	18
2.2 Data retention requirements	19
2.2.1 Group name	19
2.2.2 Number of backup versions	19
2.2.3 Backup file retention period	19
2.2.4 Number of deleted versions	20
2.2.5 Last deleted file version retention period	20
2.2.6 Archive retention period	20
2.2.7 Offsite copies	20

2.2.8	Onsite collocation	20
2.2.9	Offsite collocation	21
2.2.10	Image backup retention	21
2.2.11	Backup set retention	21
2.3	Server architecture considerations	22
2.3.1	Server platform	22
2.3.2	System size	24
2.3.3	Multiple Tivoli Storage Manager servers	25
2.3.4	Server licensed features	26
2.3.5	Network	27
2.4	Disk considerations and sizing	28
2.4.1	The disk subsystem	28
2.4.2	Tivoli Storage Manager database	29
2.4.3	Tivoli Storage Manager recovery log	32
2.4.4	Primary disk storage pool	33
2.4.5	Device configuration table and volume history file	35
2.4.6	Total disk	36
2.5	Tape considerations and sizing	36
2.5.1	Tape devices	37
2.5.2	To library or not to library?	37
2.5.3	Number of tape drives	38
2.5.4	Number of tape volumes	38
2.5.5	OS/390 tape management	42
2.6	Administrator IDs	42
2.7	License considerations	43
2.7.1	Server type	44
2.7.2	Client connections	44
2.7.3	Network connections	44
2.7.4	OpenSystems environment	44
2.7.5	Tivoli Space Manager optional product	45
2.7.6	Tivoli Disaster Recovery Manager optional product	45
2.7.7	Advanced device support	45
2.7.8	Tivoli Data Protection features	45
2.8	Other considerations	45
Chapter 3. Tivoli Storage Manager server		47
3.1	Latest code updates	47
3.2	Code installation	48
3.3	Customization	48
3.3.1	Options file	48
3.3.2	Settings	51
3.4	Server implementation	52
3.4.1	Load sample scripts	52
3.4.2	Format server volumes	53
3.4.3	Label tape volumes	55
3.4.4	Windows NT configuration wizards	55
Chapter 4. Database and recovery log		57
4.1	Defining a new database volume	57
4.2	Defining a new recovery log volume	59
4.3	Setting the log mode to ROLLFORWARD	61
4.4	Defining database backup trigger	63
4.5	Setting the expansion trigger	63

4.5.1 Database space trigger	64
4.5.2 Recovery log space trigger	65
4.5.3 When SPACETRIGGER fails.	65
4.6 Mirroring	66
4.6.1 Database mirroring	67
4.6.2 Recovery log mirroring	68
4.7 Removing default volumes	70
4.7.1 Removing the default database volume.	70
4.7.2 Removing the default recovery log volume	71
4.8 Database backup	72
4.9 Additional commands	74
4.9.1 EXTEND DB	74
4.9.2 EXTEND LOG	74
4.9.3 REDUCE DB.	75
4.9.4 REDUCE LOG.	76
4.9.5 DELETE VOLHIST TYPE=DBBACKUP.	76
Chapter 5. Data storage	79
5.1 Storage pools	79
5.1.1 Primary storage pools	80
5.1.2 Copy storage pools	80
5.2 Storage pool hierarchy	81
5.3 Movement of data between storage pools.	82
5.3.1 Migration	82
5.3.2 MAXSIZE	84
5.4 Reclamation - how to use fewer tapes.	84
5.4.1 Single drive reclamation	86
5.4.2 Reclamation of offsite volumes	86
5.5 Reduce restore times	87
5.5.1 Collocation	87
5.5.2 Disk caching	88
5.6 RAID	89
5.7 Setting up storage pools	89
5.7.1 Defining a library — physical.	92
5.7.2 Defining a drive — physical.	92
5.7.3 Defining a device class — logical	93
5.7.4 Defining a storage pool — logical	95
5.7.5 Defining storage pool volumes — logical.	99
5.8 Recommended setup	100
5.8.1 Defining an automated library	101
5.8.2 Defining a drive in an automated library	101
5.8.3 Defining a device class to an automated library.	103
5.8.4 Defining storage pools.	104
5.8.5 Defining storage pool volumes	107
5.8.6 Deleting default storage pools.	110
5.9 Additional commands	110
5.9.1 Audit library.	111
5.9.2 Audit volume	111
5.9.3 Back up storage pool	112
5.9.4 Checkin libvolume	114
5.9.5 Checkout library volumes	114
5.9.6 Delete a library	115
5.9.7 Dismount a volume	117

5.9.8 Migration	118
5.9.9 Move data	119
5.9.10 Query content	119
5.9.11 Query occupancy	120
5.9.12 Rename a storage pool	120
5.9.13 SQL commands	120
Chapter 6. Data storage policy	123
6.1 Setting up data storage policies	124
6.1.1 Defining a policy domain	125
6.1.2 Defining a policy set	126
6.1.3 Defining management classes	126
6.1.4 Defining backup copy groups	128
6.1.5 Defining an archive copy group	132
6.2 Recommended setup	133
6.2.1 Defining policy domains	134
6.2.2 Defining policy sets	134
6.2.3 Defining management classes	135
6.2.4 Defining backup copy groups	136
6.2.5 Defining the archive copy group	137
6.2.6 Deleting the STANDARD policy domain	138
6.3 Verifying policy definitions	139
6.3.1 Backup copy groups	140
6.3.2 Archive copy groups	140
6.4 Validating and activating a policy set	140
6.4.1 Validating the recommended policy sets	141
6.4.2 Activating the recommended policy sets	141
6.5 Enforcing your policy	141
Chapter 7. User management	143
7.1 Administrators	143
7.1.1 Considerations	143
7.1.2 Default environment	143
7.1.3 Recommended administrators	144
7.1.4 Working with administrators	146
7.2 Client nodes	148
7.2.1 Considerations	148
7.2.2 Default environment	148
7.2.3 Working with client nodes	149
7.3 Client option sets	153
7.3.1 Considerations	153
7.3.2 Default environment	154
7.3.3 Recommended client option sets	154
7.3.4 Associating a client node with a client option set	154
7.3.5 Working with client option sets	155
Chapter 8. Licensing your Tivoli Storage Manager system	157
8.1 Licensed features	157
8.2 Registering licensed features	158
8.3 Saving your licenses	159
8.4 License compliance	159
8.5 Monitoring licenses	159
8.5.1 Displaying license information	159
8.5.2 Auditing licenses	160

8.5.3 Scheduling automatic license audits	160
Chapter 9. Administrative client	161
9.1 Code installation	162
9.1.1 Administrative client code	162
9.1.2 Web administrative interface	162
9.2 Customization	162
9.2.1 Administrative client	162
9.2.2 Web administrative interface	163
9.3 Interfaces	163
9.3.1 Command line interface	163
9.3.2 Web administrative interface	164
Chapter 10. Backup/archive client	169
10.1 Code installation	169
10.1.1 Backup/archive client code	170
10.1.2 Web backup/archive client	170
10.2 Customization	171
10.2.1 Backup/archive client	171
10.2.2 Web backup/archive client	175
10.3 Interfaces	175
10.3.1 Command line	176
10.3.2 GUI	193
10.4 Client scheduler	204
10.4.1 Starting the client scheduler	205
10.4.2 Stopping the scheduler	209
10.5 Web client usage	210
10.5.1 Remote client GUI	210
10.5.2 Client functions	210
10.5.3 Access authorization	210
10.5.4 Starting the web client	211
10.5.5 Stopping the remote agent services	213
Chapter 11. Scheduling	215
11.1 The wheel of life	215
11.2 Administrative schedules	216
11.2.1 Defining an offsite backup schedule	218
11.2.2 Defining the volume history schedule	220
11.2.3 Defining a migration schedule	221
11.2.4 Defining a reclamation schedule	222
11.2.5 Defining an expiration schedule	223
11.2.6 Defining a licensing audit schedule	224
11.3 Client schedules	224
11.3.1 Defining a client schedule	225
11.3.2 Associating a client with a schedule	226
11.3.3 Verifying the client schedules	227
Chapter 12. Operations	229
12.1 Server procedures	229
12.1.1 Starting the Tivoli Storage Manager server	230
12.1.2 Stopping the Tivoli Storage Manager server	233
12.2 Event monitoring	236
12.3 Daily sanity checks	239
12.3.1 Database and recovery log	239

12.3.2	Data storage	242
12.3.3	Client-server activity	245
12.4	Tape management	248
12.4.1	Labeling tapes	252
12.4.2	Onsite tape management to offsite	254
12.4.3	Offsite tape management to onsite	258
12.4.4	Reclaiming offsite tapes	260
12.4.5	Database backup management	260
12.5	Error conditions	263
12.5.1	Tivoli Storage Manager errors	263
12.5.2	Machine errors	264
Chapter 13. Performance considerations		269
13.1	Tivoli Storage Manager server	269
13.1.1	Self tuning parameters	269
13.1.2	General performance parameters	271
13.2	Client node	274
13.2.1	COMPRESSION	274
13.2.2	COMPRESSALWAYS	275
13.2.3	DIRMC	275
13.2.4	QUIET	275
13.2.5	LARGECOMMBUFFERS	275
13.2.6	TCPBUFFSIZE	275
13.2.7	TCPWINDOWSIZE	276
13.2.8	TXNBYTELIMIT	276
13.3	How to measure performance	276
Appendix A. Planning and sizing worksheets		279
Appendix B. Redbook support material — macros and scripts		285
B.1	Macros	285
B.1.1	Define administrators	285
B.1.2	Define client option sets	286
B.1.3	Define policy structure	287
B.1.4	Define schedules	289
B.1.5	Define server scripts	290
B.1.6	Create storage pools	291
B.1.7	Delete default storage pools	292
B.2	Server options files	292
B.2.1	AIX	293
B.2.2	MVS	294
B.2.3	Windows NT	296
B.3	Client options files	297
B.3.1	AIX	297
B.3.2	NetWare	298
B.3.3	Windows 95/98/NT	299
Appendix C. Special notices		301
Appendix D. Related publications		303
D.1	International Technical Support Organization publications	303
D.2	Redbooks on CD-ROMs	304
D.3	Tivoli Storage Management Publications	304

How to get IBM Redbooks	307
IBM Redbooks fax order form	308
Index	309
IBM Redbooks review	319

Figures

1. Redbook Tivoli Storage Manager environment	2
2. Client requirements	12
3. Tivoli Storage Manager database and its volumes	58
4. Tivoli Storage Manager recovery log and its volumes	60
5. Tivoli Storage Manager mirror view	66
6. How volumes and storage pools are arranged for disk and tape	79
7. Possible hierarchical arrangement of different storage devices	82
8. Migration between storage pools	84
9. Space reclamation of tape volumes	85
10. Storage pool collocation on a client level	88
11. Device support worksheet from Tivoli Storage Manager Web site	102
12. Storage pool hierarchy for our recommended setup	105
13. Policy relationships and resources	123
14. Data storage policy components	124
15. Include option example	127
16. ARCHMC example	127
17. DIRMC example	127
18. Data flow through copy groups	130
19. Recommended policy definitions	134
20. Web administrative interface logon screen	166
21. Tivoli Storage Manager Enterprise Console browser interface	168
22. Difference between a local and web-based restore	169
23. Web backup/archive client main window	170
24. Backup/archive client GUI: Windows	194
25. Backup/archive client GUI: UNIX	195
26. GUI backup window	196
27. GUI restore window	197
28. GUI backupset restore window from server data	200
29. GUI for a local backupset restore	201
30. GUI archive window	202
31. GUI retrieve window	203
32. Batch file for scheduler start on windows	206
33. Web backup/archive client main window	212
34. Web client login dialog box	213
35. Scheduling of operations	216
36. Administrative commands to back up storage pools and database	218
37. Script file for offsite backups	219
38. Launch server console in foreground	232
39. Onsite and offsite distinction	249
40. Tivoli Storage Manager tape processing	250
41. How Tivoli Storage Manager tapes are processed	251
42. Tivoli Storage Manager and tape lifecycle	252
43. Tivoli Storage Manager error entries in Windows Event Viewer	266
44. Detailed event information of a Tivoli Storage Manager error	267

Tables

1. Redbook support material	3
2. Planning checklist	5
3. Server implementation checklist	7
4. Client implementation checklist: server tasks	8
5. Client implementation checklist: client tasks	8
6. Daily operations checklist.	9
7. Client requirements worksheet.	13
8. Rules of thumb for selecting percentage of data changed.	14
9. Rules of thumb for selecting data compression	15
10. Typical data compression ratios.	16
11. Storage policy requirement worksheet.	19
12. Factors affecting collocation.	21
13. Tivoli Storage Manager server platform considerations.	22
14. Rough estimates of maximum network throughputs	28
15. Database worksheet - backup requirements only	31
16. Database worksheet - backup and archive requirements	32
17. Recovery log worksheet.	33
18. Primary storage pool worksheet.	35
19. Device configuration and volume history worksheets	36
20. Total Tivoli Storage Manager disk required worksheet	36
21. Tape drive configuration worksheet	37
22. Administrator IDs worksheet	43
23. License requirements worksheet	43
24. Server options file location	49
25. Server options: site dependent	50
26. Server options: performance	50
27. Security related settings.	51
28. Database and recovery log volumes allocation	67
29. Guide to setting up storage pools	91
30. Client requirements worksheet: example.	149
31. Tivoli Storage Manager licensed features	157
32. Administrative client options.	163
33. Web client options	175
34. Backup-Archive client interfaces by platform	175
35. Backup command examples: UNIX	178
36. Backup command examples: Windows	178
37. Restore command examples: UNIX.	184
38. Restore command examples: Windows.	185
39. Archive command examples: UNIX	189
40. Archive command examples: Windows	189
41. Retrieve command examples: UNIX	191
42. Retrieve command examples: Windows	191
43. Recommended schedules	217
44. Client requirements worksheet.	279
45. Storage policy requirements worksheet.	280
46. Database worksheet	280
47. Recovery log worksheet.	280
48. Primary storage pool worksheet.	281
49. Device configuration and volume history worksheet	281
50. Total Tivoli Storage Manager disk required worksheet	281

51. Tape drive configuration worksheet	282
52. Administrator IDs worksheet	282
53. License requirements worksheet	282

Preface

This redbook provides material describing how to implement and operate Tivoli Storage Manager for people who already have a conceptual understanding of Tivoli Storage Manager but have little or no practical experience of Tivoli Storage Manager implementations.

Support materials in the form of example option files and macros are included where appropriate, and checklists lead you step-by-step through an implementation of the basic functions of Tivoli Storage Manager:

- Planning and sizing
- Product implementation
- Customization
- Operation
- Performance and tuning

This practical guide is intended for the following audience:

- New Business Partners and others that are aiming for Tivoli Storage Manager Certification
- System administrators, new to Tivoli Storage Manager, who are asked to commence a basic Tivoli Storage Manager implementation for the very first time
- Tivoli Storage Manager administrators who want to learn more about the basic Tivoli Storage Manager components and their implementation

A companion redbook to this is *Tivoli Storage Management/ADSM Concepts*, SG24-4877, which covers the features and functions of Tivoli Storage Management products. That redbook would be a useful general introduction for people who have had no previous exposure to Tivoli Storage Manager.

Tivoli Storage Manager is the follow-on product to IBM's ADSTAR Distributed Storage Manager (ADSM). It is a part of Tivoli Storage Management Solutions. For information about Tivoli Storage Management Solutions, please visit the Tivoli Storage Management Web site at:

<http://www.tivoli.com/products/solutions/storage>

The team that wrote this redbook

This redbook was produced by two teams of specialists from around the world working at the International Technical Support Organization San Jose Center.

Arnold Balingit is an IT Specialist in IBM Philippines. He has seven years experience in the IT field. He holds a degree in Computer Engineering from Adamson University. His areas of expertise include Tivoli Storage Manager implementation on a variety of supported platforms, project management, enterprise wide implementation of Systems Management Solutions (Tivoli), and development of Disaster Recovery Plans (DRPs).

Ross Battaglia is the Director of Technical Services at ETI, in Boca Raton, Florida, an IBM Business Partner, and Developer for the Tandem Tivoli Storage

Manager client. He has 20 years of experience in the data processing field. He holds a degree in Computer Science from Northeastern University. His areas of expertise include MVS systems programming, Storage Solutions, Tivoli Storage Manager, Disaster Recovery, and Data Communications. He has written extensively on product evaluations and third party software.

Charlotte Brooks is a Tivoli Storage Management Project Manager at the International Technical Support Organization, San Jose Center. She has ten years of experience with IBM as a technical specialist in the fields of RS/6000 and Storage Management. Before joining the ITSO this year, Charlotte was Technical Support Manager for Tivoli Storage Manager in the Asia Pacific region.

J.P. Houle is a Senior IT Specialist with IBM Canada in Ottawa, Ontario. He is IBM Certified in AIX, Tivoli Storage Manager and Tivoli Framework products and has been working in the UNIX field for eight years. He holds a bachelor's degree in Mechanical Engineering with a concentration in Computer Integrated Manufacturing from Carleton University in Ottawa, Ontario, Canada.

Mathis Landzettel was a Project Leader at the International Technical Support Organization, San Jose Center. He joined IBM in 1994 after completing his degree in mathematics at the Technical University of Darmstadt. He writes extensively and teaches IBM classes worldwide on all areas of Tivoli Storage Manager. Before joining the ITSO in 1998, Mathis worked in the Tivoli Storage Manager development department in Mainz, Germany, as a software test team leader. Mathis is now with the SAN Interoperability Lab, also in Mainz.

Armando Lemos da Silva Filho is the Information Security Officer at Deutsche Bank Brazil, UNIX Level 2 support and also an IBM certified specialist in Tivoli Storage Manager. He has 12 years of experience in the data processing field for Mainframe and Midrange systems. He has been working with Tivoli Storage Manager, AIX, and NT for the past four years. He holds a degree in Computer Science from Pontificia Universidade Catolica in Sao Paulo. His areas of expertise include Hypertext Technology, Computer Security, Operating Systems and Disaster Recovery.

Rod MacLeod is an MVS Team Leader for TELUS Communications in Alberta, Canada. He has 19 years of experience in the data processing field. He holds a degree in Computation Science from the University of Saskatchewan. His areas of expertise include MVS systems programming, performance, capacity planning, and data storage. He has been working with Tivoli Storage Manager for over three years. He has written articles and presented at conferences regarding his S/390 Parallel Sysplex implementation experiences.

Andy Pattinson is a Storage Consultant for Storm UK, an IBM business partner and European storage distributor in the north of England. He has 17 years experience in computer systems and storage. He is a Tivoli Storage Manager and NT certified. His areas of expertise include storage planning and hardware configuration, in addition to Tivoli Storage Manager implementation.

Patrick Randall is a Distributed Storage Software Specialist at the International Technical Support Organization, San Jose Center. He has written eight redbooks on Tivoli Storage Manager, teaches IBM classes worldwide on all areas of distributed storage, and is a consultant in Disaster and Business Recovery.

Before joining the ITSO in July 1996, Patrick worked in IBM UK Business Recovery Services as a Solutions Architect.

Phil Thomas is a senior Technical consultant for Aspect Computing, an IBM Business Partner in Australia. He has 22 years of experience in the Information Technology field. He holds a degree in Science from Sydney University and is a Tivoli Storage Manager certified specialist. Apart from Tivoli Storage Manager solution design and implementation, his areas of expertise include AIX, MVS, systems design, and Assembler programming. He has taught extensively in Asia and Australia.

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

Elsa Martinez
International Technical Support Organization, San Jose Center

Yvonne Lyon
International Technical Support Organization, San Jose Center

Emma Jacobs
International Technical Support Organization, San Jose Center

Gabriele Velez
International Technical Support Organization, Rochester Center

Mike Kaczmarek
IBM Storage Systems Division, Tucson

Tricia Jiang
Tivoli Systems, San Jose

Comments welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in “IBM Redbooks review” on page 319 to the fax number shown on the form.
- Use the online evaluation form found at ibm.com/redbooks
- Send your comments in an Internet note to redbook@us.ibm.com

Chapter 1. Tivoli Storage Manager implementation checklists

The objective of this redbook is to provide material describing how to implement and operate Tivoli Storage Manager Version 3.7. We assume a basic knowledge of Tivoli Storage Manager, which you can gain by reading the companion redbook, *Tivoli Storage Management/ADSM Concepts*, SG24-4877.

This chapter provides an overview of our redbook Tivoli Storage Manager environment and implementation checklists for planning, installing, and operating that environment.

Our environment provides an integrated solution which incorporates complementary client and server options, basic performance recommendations, and operational processes. In our experience, this environment has been shown to satisfy the most common customer requirements while also forming a sound basis for extension.

The checklists provide step-by-step processes to plan and implement a Tivoli Storage Manager environment. Although geared towards our redbook environment, the checklists can be used for any implementation. There are separate checklists for planning the environment, server implementation, client implementation, and daily operations.

We have provided planning sheets, option files, and administrative macros to help plan and implement your Tivoli Storage Manager environment. Appendix A, "Planning and sizing worksheets" on page 279 and Appendix B, "Redbook support material — macros and scripts" on page 285 provide copies of those materials.

1.1 The big picture

Any Tivoli Storage Manager solution consists of a number of pieces that are crafted to satisfy a particular set of requirements. These solution pieces include definitions for data storage management, policy management, user management, and operational management.

The difficulty is in determining how to craft each of these pieces to complete the solution jigsaw puzzle. This is complicated by the vast number of options and variations that are possible with Tivoli Storage Manager. We can help you.

We have developed a functional Tivoli Storage Manager environment which has been shown to satisfy a number of key customer requirements. Those key requirements are:

- Multiple backup copies of files to be kept
- Second copy of backup data to be kept offsite
- Restore time to be minimized
- High level of automation

Our environment also incorporates basic performance tuning recommendations and operational procedures for onsite-offsite tape movement. It forms a sound platform for future development.

Figure 1 shows the data storage management perspective of our Tivoli Storage Manager environment. The figure shows the flow of data to and from the onsite storage pools and offsite copy pools. Some key features of this environment are:

- Separate storage pools for client directory information and client data
- Client data written to a disk storage pool, then migrated to tape storage pool
- Duplicate copies of onsite data created for offsite storage
- Mirrored Tivoli Storage Manager database and recovery log

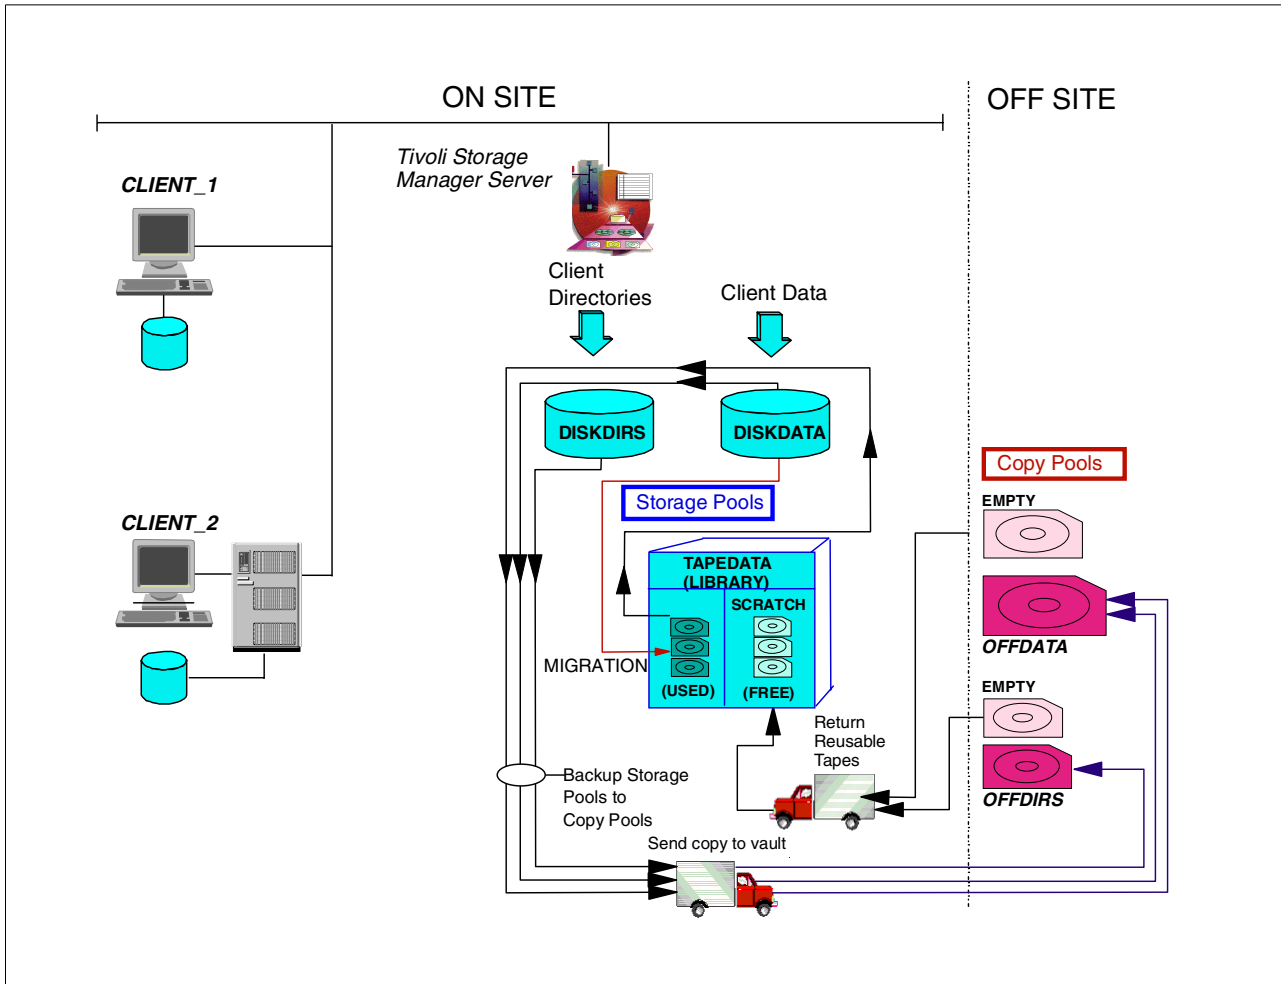


Figure 1. Redbook Tivoli Storage Manager environment

The two primary disk storage pools hold client directory information (DISKDIRS) and client data (DISKDATA). The remaining storage pool (TAPEDATA) is on tape and holds only client data. A copy of client data is stored in one offsite storage pool (OFFDATA), and a copy of the client directory data is stored in another storage pool (OFFDIRS).

1.2 Redbook support material

We provide worksheets, option files, and administrative macros to help plan and implement your Tivoli Storage Manager environment.

The worksheets are used during the planning phase of a Tivoli Storage Manager implementation which we use and discuss in Chapter 2, “Tivoli Storage Manager implementation planning” on page 11.

The option files provided are examples for customizing your Tivoli Storage Manager server and clients. If you want to use them in your implementation, you need to modify the example files to fit into your environment, and replace the existing options files with these.

The administrative macros are used to reduce some steps in your Tivoli Storage Manager implementation. Again, you need to modify some of the macros to meet specific requirements in your Tivoli Storage Manager environment.

The redbook support material is available in softcopy on the internet from the redbooks Web server. Point your Web browser to:

<ftp://www.redbooks.ibm.com/redbooks/SG245416>

Alternatively, you can get to the same Web page at:

<http://www.redbooks.ibm.com>

Select Additional Materials and click the suggested link (or follow the instructions given, since the Web pages change frequently!), and then select SG245416.

We provide two files. Each file contains all the support material in compressed format. The UNIX platform file is named `sg245416.tar`, and the Windows platform file is named `sg245416.zip`.

Table 1 lists the contents of the support material from those files.

Table 1. Redbook support material

File name	Contents
<code>dsmopt.aix</code>	Client user options file for AIX
<code>dsmopt.nw</code>	Client options file for NetWare
<code>dsmopt.win</code>	Client options file for Windows 95/NT
<code>dsmserv.aix</code>	Server options file for AIX
<code>dsmserv.mvs</code>	Server options file for MVS
<code>dsmserv.win</code>	Server options file for Windows NT
<code>dsmsys.aix</code>	Client system options file for AIX
<code>mac.admins</code>	Administrative macro to define administrators
<code>mac.optionsets</code>	Administrative macro to define client option sets
<code>mac.schedules</code>	Administrative macro to define administrative and client schedules
<code>mac.scripts</code>	Administrative macro to define server scripts
<code>mac.stgcreate</code>	Administrative macro to create storage pools
<code>mac.stgdelete</code>	Administrative macro to delete default storage pools
<code>mac.policy</code>	Administrative macro to define policy domains, policy sets, management classes, and copy groups.

File name	Contents
planning.123	Planning spreadsheets (Lotus 123 format)
plan_sampledata.123	Planning spreadsheets (Lotus 123 format) with some sample data
planning.xls	Planning spreadsheets (Microsoft Excel format)
plan_sampledata.xls	Planning spreadsheets (Microsoft Excel format) with some sample data
readme.1st	Contents of Support Materials

We recommend that you download the support material files into a separate directory on a system from which you can run an administrative command client. In our experience, the implementation works best when you choose a UNIX or Windows platform for that system.

You will probably need to customize the macros to meet your requirements. If you want to use our redbook support material, you have to execute the macro files in Tivoli Storage Manager. We assume the administrative client is running in a Windows environment and the macro files are located in directory C:\redbook. You must invoke the administrative command line client with the ITEMCOMMIT parameter, otherwise the macro command will fail with a series of error messages.

The following command shows how to execute the macro mac.admins to create the recommended administrative IDs. To execute other macros, you need to specify the appropriate macro file.


```

C:\Program Files\Tivoli\TSM\baclient> dsmadm -itemcommit
Tivoli Storage Manager
Command Line Administrative Interface - Version 3, Release 7, Level 1.0
(C) Copyright IBM Corporation, 1990, 1999, All Rights Reserved.

Enter your user id: admin

Enter your password: *****

Session established with server A78AXFWY: Windows NT
Server Version 3, Release 7, Level 1.0
Server date/time: 02/07/2000 17:30:22 Last access: 02/07/2000 17:20:47

tsm: A78AXFWY> macro c:\redbook\mac.admins

ANS8000I Server command: 'register admin sysadmin sysadmin contact='System Administrator''
ANR2068I Administrator SYSADMIN registered.
ANS8000I Server command: 'grant authority sysadmin classes=system'
ANR2076I System privilege granted to administrator SYSADMIN.
ANS8000I Server command: 'register admin support support contact='System Support''
ANR2068I Administrator SUPPORT registered.
ANS8000I Server command: 'grant authority support classes=system'
ANR2076I System privilege granted to administrator SUPPORT.
ANS8000I Server command: 'register admin reporter reporter contact='System Reporting''
ANR2068I Administrator REPORTER registered.
ANS8000I Server command: 'register admin helpdesk helpdesk contact='Client Administrator''
ANR2068I Administrator HELPDESK registered.
ANS8000I Server command: 'query admin'

Administrator      Days Since      Days Since      Locked?      Privilege Classes
Name               Last Access    Password Set
-----
ADMIN              <1             <1             No           System
CLIENT            <1             <1             No           Client Owner
HELPDESK          <1             <1             No
REPORTER          <1             <1             No
SERVER_CONSOLE    <1             <1             No           System
SUPPORT           <1             <1             No           System
SYSADMIN          <1             <1             No           System

```

1.3 Planning checklist

Proper planning for your Tivoli Storage Manager environment is critical to the success of the implementation. You should review all of Chapter 2, “Tivoli Storage Manager implementation planning” on page 11.

The tasks contained in the Tivoli Storage Manager planning checklist are shown in Table 2. You should complete all these tasks before implementing the Tivoli Storage Manager environment.

Table 2. Planning checklist

✓	Tasks	Refer to
	Download redbook support materials	1.2 on page 2
	Complete client requirements worksheet	2.1 on page 11
	Complete data retention worksheet	2.2 on page 19
	Choose server platform	2.3.1 on page 22
	Size Tivoli Storage Manager server	2.3.2 on page 24
	Complete licensing worksheet	2.3.4 on page 26
	Determine network load	2.3.5 on page 27

✓	Tasks	Refer to
	Complete Tivoli Storage Manager database worksheet	2.4.2 on page 29
	Complete Tivoli Storage Manager recovery log worksheet	2.4.3 on page 32
	Complete Tivoli Storage Manager storage pool worksheet	2.4.4 on page 33
	Complete Tivoli Storage Manager disk worksheet	2.4.6 on page 36
	Determine tape library	2.5.2 on page 37
	Determine number of tape drives	2.5.3 on page 38
	Determine number of tape volumes	2.5.4 on page 38
	Complete administrator worksheet	2.6 on page 42

1.4 Server implementation checklist

The server checklist identifies those tasks you must complete to set up the redbook Tivoli Storage Manager server environment. The tasks contained in the checklist are shown in Table 3. These tasks are performed by either the system administrator or Tivoli Storage Manager administrator.

The checklist consists of a series of tasks which must be performed sequentially. Each task in the table has a reference to another section in this redbook. The referred section contains the specific details on how to complete that task. For some tasks we additionally refer to the macro file we provide as described in 1.2, “Redbook support material” on page 2.

Table 3. Server implementation checklist

✓	Tasks	Refer to	Macro
	Download latest server code fixes	3.1 on page 47	
	Install base server code and latest fixes	3.2 on page 48	
	Update server options file	3.3.1 on page 48	
	Install Web administrative client	9.1.2 on page 162	
	Create database volumes	4.1 on page 57	
	Create recovery log volumes	4.2 on page 59	
	Mirror database	4.6.1 on page 67	
	Mirror recovery log	4.6.2 on page 68	
	Remove default database volumes	4.7.1 on page 70	
	Remove default recovery log volumes	4.7.2 on page 71	
	Set up server licensing	8.2 on page 158	
	Define tape libraries	5.8.1 on page 101	
	Define tape drives	5.8.2 on page 101	
	Define device classes	5.8.3 on page 103	
	Change server run-time settings	3.3.2 on page 51	
	Set recovery log mode	4.3 on page 61	
	Set up database trigger	4.4 on page 63	
	Set up database expansion trigger	4.5.1 on page 64	
	Set up recovery log expansion trigger	4.5.2 on page 65	
	Define storage pools	on page 104	stgcreate
	Define storage pool volumes	5.8.5 on page 107	
	Remove default storage pools	5.8.6 on page 110	stgdelete
	Define policy domains	6.2.1 on page 134	policy
	Define policy sets	6.2.2 on page 134	policy
	Define management classes	6.2.3 on page 135	policy
	Define backup copy groups	6.2.4 on page 136	policy
	Define archive copy groups	6.2.5 on page 137	policy
	Activate new policy	6.4.2 on page 141	
	Remove default policy management definitions	6.2.6 on page 138	
	Define administrator IDs	7.1.3 on page 144	admins
	Define administrative schedules	11.2 on page 216	schedules
	Define client schedules	11.3.1 on page 225	schedules
	Create client option sets	7.3.3 on page 154	optionsets

1.5 Client implementation checklist

The client implementation checklist consists of two parts that identify those tasks you must complete to set up the redbook Tivoli Storage Manager client environment. The tasks contained in the checklists are shown in Table 4 and Table 5.

Each checklist consists of a series of tasks which must be performed sequentially. Each task in the table has a reference to another section in this redbook. The referred section contains the specific details on how to complete that task.

The first checklist consists of tasks performed at the Tivoli Storage Manager server. These tasks are performed by the Tivoli Storage Manager administrator.

Table 4. Client implementation checklist: server tasks

✓	Tasks	Refer to
	Register client node	7.2.3.1 on page 149
	Associate client nodes with schedules	11.3.2 on page 226
	Associate client nodes with client option set	7.3.4 on page 154
	Grant authority for Web client access	7.2.3.3 on page 150
	Define event logging	12.2 on page 236

The second checklist consists of tasks performed at the Tivoli Storage Manager client. These tasks are performed by the administrator of that client system.

Table 5. Client implementation checklist: client tasks

✓	Tasks	Refer to
	Download latest client code	3.1 on page 47
	Install client code	10.1 on page 169
	Update environment	10.2.1.1 on page 171
	Update client options files	10.2.1.2 on page 172
	Start backup-archive client	10.3.2.1 on page 193
	Implement scheduler	10.4.1.2 on page 207
	Implement Web client	10.5.4 on page 211

1.6 Operations checklist

The operations checklist consists of those tasks you should complete on a daily basis. The tasks contained in the checklist are shown in Table 6. Each task in the table has a reference to another section in this redbook. The referred section contains the specific details on how to complete that task. This checklist does not include the tasks we recommend scheduling on a daily basis. See Chapter 11, “Scheduling” on page 215 for more information on scheduled operations.

Table 6. Daily operations checklist

✓	Tasks	Refer to
	Check completion of client and administrative events	12.3.3.2 on page 246
	Bring free offsite volumes back to onsite	12.4.3 on page 258
	Send copy tapes to offsite location	12.4.2 on page 254
	Send Tivoli Storage Manager database copy to offsite	12.4.5 on page 260
	Check database and log utilization	12.3.1 on page 239
	Monitor number of scratch tapes available	12.3.2.4 on page 243
	Check client restartable restores	12.3.3.4 on page 247

Chapter 2. Tivoli Storage Manager implementation planning

A successful implementation of Tivoli Storage Manager benefits enormously from planning prior to attempting to set up the environment.

In this chapter, we present a number of planning worksheets that lead you through gathering the client requirements and the data retention requirements in an orderly way. We assume you are somewhat familiar with Tivoli Storage Manager concepts and terms.

We provide planning sheets, option files, server scripts, and administrative macros to help plan and implement your Tivoli Storage Manager environment. Blank worksheets are provided in Appendix A, “Planning and sizing worksheets” on page 279. Appendix B, “Redbook support material — macros and scripts” on page 285 contains information about how to download those support materials and what those materials provide.

The redbook support material is available in softcopy on the Internet from the Redbook Web server. Point your Web browser to:

<ftp://www.redbooks.ibm.com/redbooks/SG245416>

Alternatively, you can get to the same Web page at:

<http://www.redbooks.ibm.com>

Select Additional Materials and click the suggested link (or follow the instructions given, since the Web pages change frequently!), and then select SG245416/.

2.1 Client environment data

Tivoli Storage Manager exists to provide services to clients, so it makes sense to begin by gathering data about the client environment. A Tivoli Storage Manager client is the machine from which Tivoli Storage Manager backs up or archives data. It could be various flavors of a workstation, laptop computer, file server, or possibly a database or application server. Even though you may know the machine as a server, to Tivoli Storage Manager it is a client. The Tivoli Storage Manager server refers to the machine where the Tivoli Storage Manager server code runs. The Tivoli Storage Manager server stores and manages all the data backed up from clients.

UNIX Users:

For UNIX clients, Tivoli Storage Manager offers an extra level of protection called Logical Volume Backup. We will be using the term *Image backup* and *Logical Volume backup* interchangeably throughout the book to refer to this feature.

At this point, it is very important that you consider your backup and restore requirements, so that you can match the Tivoli Storage Manager features to your needs. Since statistics show that the majority of user restore requests are for single files or other small amounts of data, the Tivoli Storage manager file-level

backup is the foundation of any recover strategy. However, there are additional methods which can more completely cover your data restore requirements.

Tivoli Storage Manager provides the following backup methods:

- **File-level backup:** good for small and ad-hoc file or directory restores, but may not be the optimal solution for large-volume restores.
- **Logical volume backup:** good for large full restores, but only available on UNIX and not for single file restores.
- **Backup sets:** good for remote and/or portable restores.

Figure 2 shows some possible generic scenarios. Of course, you can combine those in different ways to fit your own requirements.

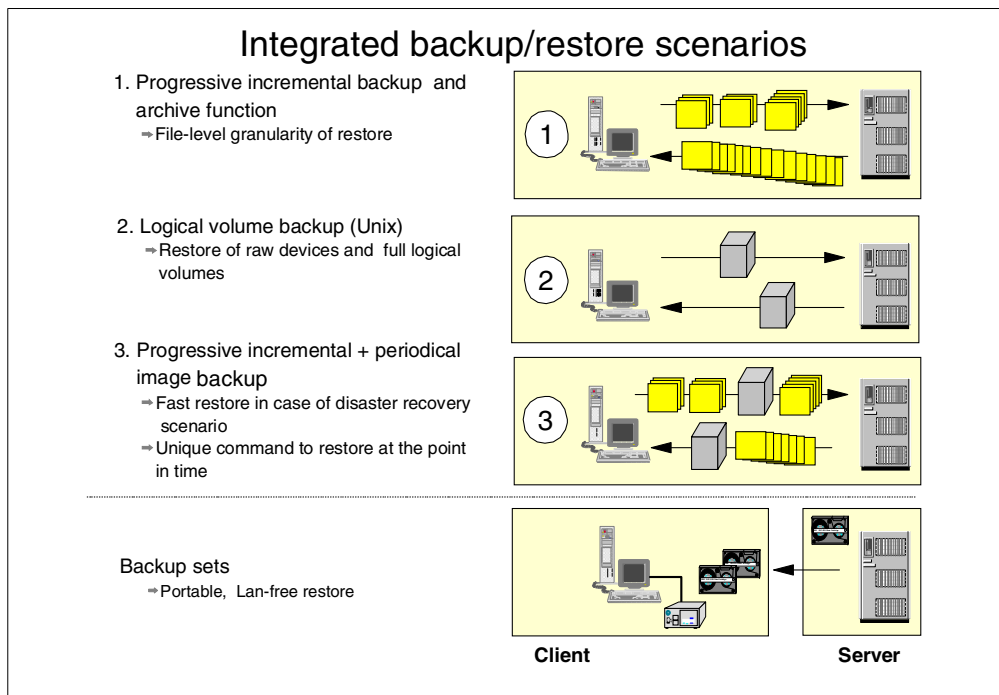


Figure 2. Client requirements

Due to the nature and complexity of customer environments and requirements, our planning description will give the building blocks for your own environment definition. The tables and calculations take into account the backup, archive, image and backup set requirements that should be suitable for a real-life situation. In any case, you must evaluate if all such requirements are really important in your case and decide whether or not to implement them.

Complete Table 7 using a column for data from each client being considered for this Tivoli Storage Manager implementation. Analyzing client data helps you make decisions about the Tivoli Storage Manager server environment. The information collected in the table is used for node definitions and for calculating database, recovery log, and storage pool sizes.

This table is presented in a portrait orientation in the book due to space considerations. If you have a large number of nodes or use a spreadsheet

version, you may find the table more workable by transposing it to a landscape orientation.

Table 7. Client requirements worksheet

	Client 1	Client 2	Client 3
Client name	RRomano	Rod Server	File Server 2
Contact information	Ralph at 5570	Server Group	UNIX Group
Operating system	Windows 95	Windows NT	AIX 4.3
Total storage available (GB)	3	10	33
Total storage used (GB)	0.75	8	25
GB changed per backup	0.01	0.8	1.5
Number of files backed up	5000	15000	2.5GB
Number of backup versions kept	2	2	7
Data compression	0.5	0.5	0.66
Backup window times	16:30 - 08:00	20:00 - 07:00	23:30 - 04:30
Backup number of hours	15.5	11	5
Required recovery time frame	48 hours	24 hours	10 hours
Tivoli Storage Manager restore time frame	24 hours	6 hours	8 hours
GB copied per archive	1	8	5
Number of files archived	5000	45000	17500
Number of archives kept	2	12	12
Archive frequency	6 months	Monthly	Monthly
Archive window times	Year end	Month end	Month end
Archive number of hours	24 hours	10 hours	24 hours
Number of image backups	N/A	N/A	12
Image backup frequency	N/A	N/A	Monthly
Number of backup sets	2	6	6
Backup set frequency	6 months	bi-monthly	bi-monthly
Policy domain	Workstn	Server	Server
Client option set	Windows	Windows	AIX

2.1.1 Client name

Enter the name Tivoli Storage Manager will use for each client. Each name must be unique. We recommend using the hostname for the Tivoli Storage Manager client name, so various groups such as Help Desk personnel or end users can easily correlate the client node to the Tivoli Storage Manager name without having to look in a translation table.

By allowing the Tivoli Storage Manager client name to default to the machine name on Windows clients, you are able to roll out a standard Tivoli Storage Manager options file to numerous clients automatically without having to modify the options file for each client.

2.1.2 Contact information

Enter information identifying the contact person or group responsible for this client.

2.1.3 Operating system

Identify the operating system and level the client is using. Any clients using operating systems not supported by Tivoli Storage Manager will have to be handled separately.

2.1.4 Total storage available

Calculate the total amount of usable disk storage in GB available on the client. This number is the amount of disk storage seen by the client file system and does not contain the actual amount installed and used by a RAID or mirroring implementation.

2.1.5 Total storage used

Calculate the total amount of disk storage in GB currently in use or expected to be used on the client. If this number is unknown, use the *Total storage available* from above.

2.1.6 GB changed per backup

Calculate the amount of storage, data, or files changed between backup cycles. It indicates how long the client will be busy backing up, and how robust a network is required to complete all the backups each backup cycle. It is used to estimate disk and tape storage requirements on the Tivoli Storage Manager server. Table 8 presents some typical percentages of data changed for various sizes and types of data.

If your data change percentage is known, then enter that number in the worksheet. Otherwise, use Table 8 to estimate a rate, keeping in mind that your numbers may vary. A high estimate is better than one that is too low.

Table 8. Rules of thumb for selecting percentage of data changed

Configuration	Percentage of Data Changed (%)
Large, busy file server	10
Smaller, less busy file server	5
Workstation	1
Database using utilities or Tivoli Data Protection Agent	10-20
Database not using utilities or Tivoli Data Protection Agent	100

2.1.7 Number of files backed up

Calculate the total number of files to be backed up for each client. It is used to estimate disk and tape storage requirements on the Tivoli Storage Manager server.

If the number of files is unknown, two values are possible. Enter either an estimate of the number of files in this field, or enter 10% of the *Total Storage Used* field, in GB.

2.1.8 Number of backup versions kept

Determine the number of changed copies that you want to keep of a file that exists on the client when the backup task runs. How many different versions of that file do you want to be able to restore? For example, if backup runs every night and a file changes every day, and you want to be able to restore any version up to one week ago, then you would choose 7 as the number of backups to keep. Remember that the more versions you choose to keep, the more database and storage pool space you will need to configure.

2.1.9 Data compression

Estimate a data compression rate. To do that, you must first decide whether to use data compression on the client.

Tivoli Storage Manager allows a client to compress data before sending it to the Tivoli Storage Manager server. To decide whether to use Tivoli Storage Manager compression, take into account the speed of your network, the speed of the client CPU, and whether the Tivoli Storage Manager server storage devices are capable of compression. Table 9 provides some criteria for selecting data compression.

Table 9. Rules of thumb for selecting data compression

Configuration	Compression
Dial-up connection	Yes
Your network is approaching capacity	Yes
ATM network	No
CPU model of the client is less than Pentium II 100 Mhz (or equivalent)	No
CPU model of the client is greater than Pentium II 100 Mhz (or equivalent)	Yes
No hardware compression on tape devices	Yes
Small capacity or slow response storage devices	Yes

Compression takes time, so you need to decide whether the compression is helping or hurting the total elapsed time of your operation. In general, fast networks such as ATM networks do not benefit from compression because network bandwidth is usually not saturated. Conversely, older client CPU models may run slowly because the CPU cannot compress data fast enough to keep the network connection busy.

Many sites have various combinations of these configurations, such as slow networks, slow CPUs and small/slow devices. In this case, make a judgement call about using Tivoli Storage Manager compression. Performing tests may show you the most efficient mode for your environment.

Data compresses to varying degrees depending on the content of the files. Data composed of text, or many repeated characters like blanks, compresses well. Data that is already compressed, and data consisting of random characters like executables does not compress well, and may actually grow.

If you decide to use compression, enter your data reduction rate. Use Table 10 to estimate the data reduction if your actual ratio is unknown.

If you decide not to use compression, or are unsure whether to use compression, enter one (no compression) in the *Data Compression* field.

If you do not use client compression and your tape device has a hardware compression capability, then the data will still end up compressed when it reaches a tape volume. If you do use client compression, then you should disable any hardware compression capability on your tape devices, since data in general cannot be 'doubly compressed'.

Table 10. Typical data compression ratios

	Compression Ratio	Data Reduction
Database data	3:1 - 4:1	0.66 - 0.75
Print and file server data	2:1	0.5
Executables, compressed data, encrypted data	1:1	1

2.1.10 Backup window times

Enter the times of the day between which Tivoli Storage Manager must start and complete its backup cycle. This window is influenced when end-user client usage drops off, by availability requirements, and by network capacity usage time frames.

2.1.11 Backup number of hours

Calculate the number of hours in the backup window.

2.1.12 Required recovery time frame

Identify the time frame in hours, which you have agreed to with the customer, to recover a client completely. This is the time from when the data is lost to the time the data is usable again. It includes the time to fix or replace the machine and the time to restore the data.

Recovery time frames are vitally important, as this is the whole reason you are backing up.

This field documents your service level agreement with the customer.

2.1.13 Tivoli Storage Manager restore time frame

Calculate the time frame, in hours, allotted to Tivoli Storage Manager to restore (possibly all) the data to the client. This number is used to size factors affecting the restore process such as network throughput and the number of tape drives required.

To calculate the *Tivoli Storage Manager Restore Time Frame*, subtract from the *Required Restore Time Frame*, above, the maximum time required to prepare the client machine for a restore. In reality this number will vary depending upon the complexity of the restore. In the worst case (a disaster), the time to prepare the machine may include contacting support, fixing or replacing the machine, installing an operating system, installing the Tivoli Storage Manager client code, and connecting to the network. In the best case (just deleted data), the time to prepare the machine may be the same as the *Required Restore Time Frame*.

A full Tivoli Storage Manager restore time frame typically takes significantly more time than a full Tivoli Storage Manager backup time frame. We have observed restore times of 110% or more over a full file-level backup.

2.1.14 GB copied per archive

Calculate or estimate the amount of data to be archived during each archive session. Archives target specific data files. Typically, you do not archive whole systems.

2.1.15 Number of files archived

Calculate or estimate the number of files archived in each archive session. You can calculate this number by using an operating system utility to verify the size of a directory tree or an entire filesystem.

2.1.16 Number of archives kept

Identify how many archives will be kept. For example, if an archive is performed monthly, you may want to keep 12 copies of that archive. Note that this number is basically influenced by the retention period for that archive, since archives, unlike backups, do not have versions which expire.

2.1.17 Archive frequency

Determine how often you want to perform an archive function. Archives are typically run less frequently than incremental backups. Time frames such as monthly or yearly are common.

2.1.18 Archive window times

Enter the times between which Tivoli Storage Manager must start and complete its archive cycle for this client. This window is influenced by when end-user client usage drops off, by availability requirements, by network capacity usage time frames, or by when the data becomes available for archive.

Archive time frames may be less time sensitive than incremental backups if they are going against copies of data or against data already processed. Often it is not possible to automatically schedule this archive function using Tivoli Storage Manager schedules, due to the dependency on other events external to Tivoli Storage Manager.

This number may be a specific time frame, such as between 23:00 and 04:00, or a more general time frame, such as after month-end processing finishes.

2.1.19 Archive number of hours

Identify the number of hours available to complete the archive function.

2.1.20 Number of image backups

Identify the number of eligible UNIX images that you want to backup. For a workstation that doesn't change very frequently, a good number could be six to twelve per machine.

2.1.21 Image backup frequency

Identify how often you want to perform image backups. A good standard will be on a monthly basis for less critical and static machines, with more frequent backups for more critical and dynamic machines.

2.1.22 Number of backup sets

Determine the number of backup sets that you want to create. This number must be balanced against your full recovery time and the change rate for data. Consider a more frequent backup set for large filesystems.

2.1.23 Backup set frequency

Determine how often you will need to generate a backup set. Although a backup set is made of client data, it is a server initiated process which will require some scratch tapes in the automated library or some guidelines for the operations team to run.

2.1.24 Policy domain

Leave this field blank. During the planning phase, this is unknown. We use this field later in 1.4, "Server implementation checklist" on page 6 to contain the policy domain chosen for this client. For example, if you are implementing our redbook configuration, the policy domains would be `SERVER` or `WORKSTN`.

For more information on policy domains, see Chapter 6, "Data storage policy" on page 123.

2.1.25 Client option set

Leave this field blank. During the planning phase, this is unknown. We use this field later in 1.5, "Client implementation checklist" on page 8 to contain the client option set chosen for this client. For example, if you are implementing our redbook configuration, the client options set names would be `AIX` or `WINDOWS` or `NETWARE`.

For more information on client option sets, see 7.3, "Client option sets" on page 153.

2.2 Data retention requirements

In this section we identify the requirements for managing the data received from the clients. Categorize your data into a small number of groups with similar requirements. This table provides information to create copy groups under Management Classes in Tivoli Storage Manager, to calculate storage pool sizes, and to calculate the number of tapes required to hold the data.

Complete a column in Table 11 for each different group. We show two example groups.

Table 11. Storage policy requirement worksheet

	Example 1	Example 2
Group name	Workstn	Server
Number of backup versions	2	7
Backup file retention period	90 days	NOLIMIT
Number of deleted versions	1	2
Last deleted file version retention period	60 days	180 days
Archive retention period	180 days	360 days
Offsite copies	Yes	Yes
Onsite collocation	No	Yes
Offsite collocation	No	No
Image backup retention	30 days	30 days
Backup set retention	180 days	60 days

2.2.1 Group name

Choose a descriptive name for each different categorized group of data. In our example, the *Workstn* group is used for files from a workstation, while the *Server* group is used for all data from a large file server.

2.2.2 Number of backup versions

Determine the number of changed copies that you want to keep of a file that exists on the client when the backup task runs. How many different versions of that file do you want to be able to restore? For example, if backup runs every night and a file changes every day, and you want to be able to restore any version up to one week ago, then you would choose seven as the number of backups to keep.

You will use the field *Number of Backup Versions* as a basis to group your data storage requirements into management classes.

2.2.3 Backup file retention period

Determine the number of days you want to keep a backup version of a file (other than the current version). There are two options: keep the backup version for a set number of days; or specify NOLIMIT, which implies that you want Tivoli Storage Manager to retain all backup versions, other than the most recent

version, indefinitely (the most recent active version is already stored indefinitely by default).

2.2.4 Number of deleted versions

Determine how many versions of a file to keep after the file has been deleted from the original file system. This parameter comes into force during the first backup cycle after the file has been deleted. For example, assume you are keeping seven versions of a file as specified above, and you have set this parameter to one. When the next backup cycle runs after the file has been deleted off the client, Tivoli Storage Manager will flag the six oldest backup versions of the file for deletion and just keep the most current backup version.

2.2.5 Last deleted file version retention period

Determine the number of days you want to keep the last backup version of a file after it has been deleted from the client. There are two options: keep the last backup version for a set number of days; or NOLIMIT, which implies that you want to keep the backup version indefinitely.

For example, if you are keeping one version of a deleted file, and you set this parameter to 60, then 60 days after this file is noticed by Tivoli Storage Manager as being deleted from the client file system, the one remaining backup version will be deleted from Tivoli Storage Manager.

2.2.6 Archive retention period

Determine how long you want to keep a file that is archived. Many sites set up a limited number of data groups with standard archive retention periods, such as seven days, 31 days, 180 days, 365 days, or seven years. Nonstandard requests for archive retention periods are slotted into the next larger retention period group. This reduces management complexity at the expense of keeping some data longer than actually required. If every nonstandard request is honored, the number of groups quickly becomes unmanageable. On the other hand, you can use the backup set feature to retain all nonstandard backup requirements, or even just use backup sets instead of archive.

2.2.7 Offsite copies

Determine if you want to send a copy of the data offsite. Copying data to a removable device like tape allows the data to be taken offsite. An offsite copy along with other procedures provides recoverability in the event that the Tivoli Storage Manager server becomes unusable, or that data on the Tivoli Storage Manager server becomes corrupted.

Enter Yes to use offsite copies or No to not use offsite copies.

2.2.8 Onsite collocation

Determine if you want to use onsite collocation.

Tivoli Storage Manager uses collocation to dedicate as few tapes as required to hold all of one client's files. Collocation reduces elapsed time for multiple file restores and full client restores at the expense of using more tapes, potentially increasing backup times and increasing Tivoli Storage Manager management time for migration and for storage pool copies.

Collocating by client allows as many clients to be restored simultaneously as you have tape drives. If you have stringent restore requirements and sufficient tape drives, then collocation makes good sense.

Table 12 highlights some factors affecting whether to use collocation.

Table 12. Factors affecting collocation

	Onsite Collocation	Offsite Collocation
Short restore window	Yes	Yes
Less than 10 Clients, each 1 GB storage	No	No
More than 50 clients, each 1 GB storage	No	No
More than 50 clients, each 20 GB or more	Yes	Yes/No
Limited disk or tape resources	No	No
Workstations	No	No
Database, print, and file servers	Yes	Yes

Tivoli Storage Manager implementations that are small (small number of clients managing a small amount of data) do not see much benefit from collocation, due to the small number of tapes required for any restore. Large Tivoli Storage Manager implementations often collocate both onsite and offsite files due to the amount of data required for the restore of a client.

Tape drive capacity is a consideration for collocation. To use the minimum number of tapes, each tape must be used to its maximum capacity. Smaller capacity tapes will tend to fill completely even with small clients. In this case, collocation may be useful. With large capacity tape devices and collocation, a small client may not be able to fill a tape. With a large number of clients, significant numbers of tape volumes may be required.

See 5.5.1, “Collocation” on page 87 for more information on collocation.

2.2.9 Offsite collocation

Determine if you want to use offsite collocation. See 2.2.8, “Onsite collocation” on page 20 for considerations.

2.2.10 Image backup retention

Determine how long you want to keep an image backup. Consider your restore timeframe and balance the criticality of a full filespace restore, compared with single or a small number of files. Image backups can be very useful for quick restore of large filesystems. However, this process will take longer, the more changes there have been to the filesystem since the last image backup. A rule of thumb would be to keep at least one weekly image for small servers and a monthly image for bigger servers (or more frequently if change rates in the filesystem are high).

2.2.11 Backup set retention

A backup set execution creates a copy of the client node's previously backed up active files and stores them on sequential media. This has an impact on the number of tapes that you may need, specially if you want to retain those backup

sets for long periods of time or even if you want to have one copy onsite and one for offsite purposes.

Determine how long you want to keep the backup sets. Use a small retention if your data changes frequently and you do not need to keep it for long periods. You can use a longer retention for special cases or for legal requirements.

2.3 Server architecture considerations

Having gathered information on the total client environment, you can now make decisions about the architecture of the Tivoli Storage Manager server environment. This section deals with issues related to the Tivoli Storage Manager server.

2.3.1 Server platform

A Tivoli Storage Manager server runs on several platforms. How do you choose one platform over another? With only minor differences, a Tivoli Storage Manager server provides the same functionality on every platform. The major differences between Tivoli Storage Manager server platforms relate to capacity, cost, installation, operation, supported devices, and installed user base. Each of these factors is explained below. Table 13 summarizes these considerations in choosing a Tivoli Storage Manager server platform.

An additional platform possibility is the IBM 3466 Network Storage Manager. It is an integrated, pre-configured Storage Manager approach that minimizes much of the configuration and setup of the Tivoli Storage Manager server, as well as the maintenance needed to keep hardware and software up to date. The sections 2.1, "Client environment data" on page 11 and 2.2, "Data retention requirements" on page 19 may be of use in sizing this device. See *A Practical Guide to Network Storage Manager*, SG24-2242, for further information about Network Storage Manager.

Table 13. Tivoli Storage Manager server platform considerations

	Windows NT	UNIX	AS/400	MVS
Installed user base	Medium	AIX - Large; HP-UX, Sun Solaris - Small	Small	Small
Cost	Low	Medium	Medium - High	High
Capacity	Low - Medium	Medium - High	Medium - High	Medium - Very high
Platform installation	Simple	Simple	Medium	Complex
Operation	Simple	Complex	Complex	Very complex
Supported devices	Many	Many	Limited	Limited

2.3.1.1 Installed user base

The number of Tivoli Storage Manager servers installed for a particular platform is a consideration. IBM ships new functions to the most popular Tivoli Storage Manager server platforms (AIX, MVS, and Windows NT) first. Currently Tivoli

Storage Manager server platform updates on AS/400 are done once per year due to the AS/400 software release cycle and modification cycle.

The more popular a platform, the more customers are exercising the software, leading to a lower probability that you will find a unique problem. More popular platforms have a larger base of knowledgeable people to call upon for advice and to draw from for new hires.

2.3.1.2 Cost

Cost is a very dynamic area to discuss in a static manual, so we shall only generalize! Check for special promotions and other discounts before committing to acquiring a particular platform configuration.

Cost is further subdivided into platform costs and Tivoli Storage Manager software license costs. Platform costs include the cost to acquire all the hardware and software to run the platform exclusive of the Tivoli Storage Manager software license. It ranges from very low for Windows NT to extremely high for MVS. The high cost of installing a new MVS platform usually precludes it from being selected to run a Tivoli Storage Manager server exclusively.

Tivoli Storage Manager license costs vary considerably, with Windows NT costs being the lowest, followed by AIX, HP-UX, Sun Solaris, and AS/400 costs. All of these are one-time charges to purchase the product. MVS licenses are available for a one-time charge or as a monthly license fee. You may want to calculate the break-even point.

Tivoli Storage Manager server costs include a license for only one client by default. To manage more clients, more client licenses must be purchased and registered on the server. The actual client code is free.

2.3.1.3 Capacity

The Tivoli Storage Manager server has the capacity to manage a basically unlimited number of clients and an unlimited amount of data. The platform the Tivoli Storage Manager server software is running on does limit what Tivoli Storage Manager can administer. Various platforms have different capacities in regard to the CPU power it can deliver to Tivoli Storage Manager, the number of devices it can attach, and the throughput it can deliver.

Choose your platform with growth in mind. Moving from a small platform to a larger platform of the same server type, such as from a small AIX box to a larger AIX box, is relatively simple. Starting at the top end of a server type and moving to another server type, such as from Windows NT to AIX, involves exporting and importing each client separately. Although the procedure is straightforward, it can be time-consuming and labor-intensive.

2.3.1.4 Platform installation

Installation consists of the platform installation and the Tivoli Storage Manager server code installation. Platform installation consists of hardware installation and configuration, and of operating system installation and configuration. Installation of each platform requires specialized knowledge that will not be covered here.

The Tivoli Storage Manager server code installation varies by platform in the specifics, but generally follows a similar procedure. Installation on Windows NT can be easier due to the Windows NT Wizards that have been provided.

Installation of the Tivoli Storage Manager server on other platforms is not difficult for an administrator familiar with the platform.

2.3.1.5 Operation

Operation of a platform varies from almost completely automatic on Windows NT, to very complex on MVS, with the UNIX and AS/400 platforms somewhere in the middle. The availability of tools to assist in managing the operation of the various platforms is nearly opposite, in that the MVS environment has a rich, powerful assortment of tools, while Windows NT is lacking in this regard.

Operation of Tivoli Storage Manager itself varies only in the way some operating system-specific Tivoli Storage Manager commands are issued on each platform. It is important to look at the skills available among your staff for a particular operating system platform. If there are more people familiar with a particular platform, then it will be easier to maintain Tivoli Storage Manager in this environment.

2.3.1.6 Supported devices

There are a wide variety of supported devices on the Windows NT and UNIX platforms, including disk drives, tape drives, optical drives and automated tape libraries. MVS and AS/400 are limited to their standard choice of devices, but these devices generally have tremendous capacity. For a list of supported devices on each Tivoli Storage Manager server platform, see the URL:

http://www.tivoli.com/support/storage_mgr/adsercli.htm

A concern with the smaller platforms is the ability to attach the required amount of devices as the environment grows. On larger platforms this concern is usually reduced.

2.3.1.7 Recommendations

If you already have an MVS or AS/400 system with capacity available, these platforms are a good choice. Otherwise, they are probably out of consideration due to their complexity, cost, or specialized nature.

If your site already runs HP-UX or Sun Solaris, these are acceptable choices.

If you have experience with Windows NT, and the Tivoli Storage Manager implementation is small and will remain small, choose Windows NT.

For sites with little or no commitment to other platforms, choose AIX. It is robust enough to handle most implementations while remaining relatively cost effective. AIX-based Tivoli Storage Manager systems are very well supported by IBM in both the operating system and in Tivoli Storage Manager. There are large numbers of Tivoli Storage Manager installations to gather knowledge from. The AIX platform (and others) scales from small to very large Tivoli Storage Manager systems.

2.3.2 System size

Choosing the correct platform CPU size and memory requirements is an inexact art. There is very little information to guide you in selecting a Tivoli Storage Manager platform. As you would expect, the risk of choosing an inappropriate platform size increases as the size of the Tivoli Storage Manager implementation increases. Small Tivoli Storage Manager implementations are at less risk of

choosing an incorrect platform size, and the incremental cost to scale up or down is small. Many sites start small and grow into larger systems. However, this is of little help if you are starting large.

The Tivoli Storage Manager server is CPU-, I/O-, and memory-intensive. CPU is a function of the number of files to manage and how your platform processes I/O. A large number of small files is more CPU-intensive than a small number of large files. As the number of files and the amount of data to be moved increases, each backup, migration, storage pool copy, and expiration process will use more CPU to maintain the database entries. Tivoli Storage Manager takes advantage of multiple processors. In our experience, MVS platforms seem to be more CPU-intensive than UNIX platforms. MVS sites should be aware that Tivoli Storage Manager can use significant amounts of CPU. We have seen Tivoli Storage Manager among the top five users of CPU on MVS systems. AS/400 real case examples have shown that although Tivoli Storage Manager can take 10% to 40% of CPU resources (depending on machine model), it does not seem to be a bottleneck for such systems, especially when you consider that these are peak values (when the server is running a backup/restore operation, or when it is running a server process).

I/O is the major part of Tivoli Storage Manager processing. Tivoli Storage Manager does very little else. Backups and restores are I/O intensive. Database updates and retrievals are I/O intensive. Storage pool management is I/O intensive. The I/O subsystem needs to be robust enough to handle this load. As the number of files and the amount of data climb, the need for a larger, faster I/O subsystem increases. Separate controllers or adapters for disk and tape devices become essential as the load increases.

Memory is used to cache database entries, among other things. As the number of files being managed increases (and thus the database size increases), the amount of memory that Tivoli Storage Manager requires increases. MVS users note that Tivoli Storage Manager likes to keep a large part of its address space in real storage. Since memory is relatively cheap, and you will never regret having too much memory, we recommend starting with a minimum of 256 MB on most platforms.

2.3.3 Multiple Tivoli Storage Manager servers

When first setting up a Tivoli Storage Manager environment, we recommend implementing a single Tivoli Storage Manager server. Once experience has been gained, and the implementation has grown enough to be reaching the capacity of the current server model family, a second server may be considered. Note that we recommend upgrading your current server hardware to its next larger model before considering a second server, because this will keep the management overhead smaller. Tivoli Storage Manager can handle very large amounts of data or clients in one implementation. Currently, we have seen implementations with Tivoli Storage Manager internal server databases of 80 GB and larger (admittedly on very powerful, large platforms).

2.3.3.1 Reasons to consider multiple servers

Multiple Tivoli Storage Manager servers can be configured to provide some redundancy and disaster recoverability in the event of a Tivoli Storage Manager server outage. For example, a company with two well-connected sites A and B, may decide to install a Tivoli Storage Manager system in each site. The system in

site A would back up the client data from site B and vice versa. The loss of site A would mean that the Tivoli Storage Manager system in site B (which holds the backup data for site A) could immediately start restoring the client systems onto replacement equipment. The server at site A (which was lost) could be recovered to the server at site B.

The Virtual Volumes and Enterprise Administration capabilities of Tivoli Storage Manager make managing multiple servers easier by centralizing some administration functions and allowing changes to be replicated on some or all systems.

For very large or critical clients such as a large, enterprise-wide business intelligence complex, a dedicated Tivoli Storage Manager server (either on the same system or a different one) might be the best solution.

In installations where network connectivity is slow or expensive, placing a Tivoli Storage Manager server close to the client(s) usually makes sense. For example, for a business that has multiple file servers in each of a number of cities interconnected by a slow network, it may be appropriate to install a Tivoli Storage Manager server in each city.

2.3.3.2 Disadvantages of multiple servers

Multiple servers increase costs. Two small server CPUs may be more expensive than one larger CPU of the same power. Where one automated tape library may be enough, multiple servers may require multiple automated libraries. Every Tivoli Storage Manager server requires a Tivoli Storage Manager server license.

Management of a multiple server environment is more complex, costly and time consuming than a single environment. Installation and maintenance procedures have to be repeated on each server. Confusion about where data is stored, in the event of a restore, may result. Some of these disadvantages can be reduced by using the Enterprise Administration feature.

2.3.4 Server licensed features

The base Tivoli Storage Manager server license provides the backup and restore function, archive and retrieve function, logical volume backup and instant archive and rapid recovery functions, as well as enterprise management. It also supports Virtual Volume operations to use the TARGET server as device class for primary storage pool for the SOURCE server. One single client (which can only be installed on the server machine itself) is supported for basic backup and restore functions.

Additional licenses are available for extra clients, network connected clients, and extended device support, as well as Open systems (S/390 UNIX client and AIX AFS/DFS client).

Optional products for Tivoli Storage Manager include Tivoli Space Manager and Tivoli Disaster Recovery Manager. Neither of those products are covered in this book, and we recommend that you become familiar with the basic Tivoli Storage Manager operation prior to implementing these options.

2.3.5 Network

The network connection plays an integral part in providing service to the customer. If the other components of the solution have been correctly sized, it will often turn out that the network is the performance bottleneck. The network typically consists of a combination of network interface cards, hubs, routers, gateways, wire, and software. Tivoli Storage Manager server software, Tivoli Storage Manager client software, the server platform and the client platform all have at least minimal monitoring or management capabilities. Often small networks may have no more than this (that is, no additional network management hardware or software), with limited or no network administration expertise. This makes the network a weak link in the overall management of the implementation.

Network design, implementation, and operation is beyond the scope of this book. However, we cover some basic recommendations.

2.3.5.1 Network topology

Network topologies such as Ethernet, Token ring, FDDI, and ATM all work well with Tivoli Storage Manager. Each has its strengths and weaknesses. In general, choose the fastest network topology you can afford. Alternatively, the new SAN technologies can be implemented to take full advantage of tape library resources, freeing part of the ordinary network consumption to business applications.

To estimate the speed of the network required, use the following calculation.

Using Table 7 on page 13, calculate the total amount of data to be transferred during a backup window. For each client:

1. Multiply the *GB changed per backup* field by (1 - *Data compression* field).
2. Sum this number for all clients, giving the total data to be transferred.
3. Divide the total data transferred by the number of hours in your backup window to give the transfer rate in GB/hr.

For example, using the sample data in this table, calculate the network speed required as follows:

1. $(.01 * .5) + (.08 * .5) + (1.5 * .33)$
2. $.005 + .04 + .51 = .555$ GB
3. Assume smallest backup window of 5 hours, gives $.555 / 5 = .111$ GB/hr.

To calculate archive speed, use the same calculation as above, but substitute *GB copied per archive* for *GB changed per backup*. If backups and archives will occur during the same time period, these two figures must be added together. Keep in mind that Tivoli Storage Manager traffic probably will not be the only traffic on this network, unless you decide to create a dedicated network for backup and restore operations.

Use this number to decide what type of network speed is required to ensure that your backups finish in the backup window. Table 14 provides some very rough estimates of network throughput for various network topologies, but bear in mind that as a network approaches its potential capacity, throughput tends to drop substantially. The table provides only an indication of topologies that may or may not be capable of providing the required capacity. It is not meant to indicate actual

throughput rates. Careful testing in your own environment is highly recommended.

Table 14. Rough estimates of maximum network throughputs

Network Topology	MB/sec	GB/hr
10 Mb/sec Ethernet	1	3.6
100 Mb/sec Ethernet	10	36
Gigabit Ethernet	100	360
16 Mb/sec Token Ring	1.6	5.8
FDDI	10	36
ATM 155 Mb/sec	15.5	55.8

Using the same methodology, work out a rough estimate of how long a restore will take. Use the *Total Storage Used* field instead of the *GB Changed per Backup* field. This estimate is rough because Tivoli Storage Manager management time and Tivoli Storage Manager compression time (if used) are not considered.

2.3.5.2 Communication protocol

Most network protocols such as TCP/IP and NETBIOS are supported by Tivoli Storage Manager. TCP/IP is the most common communication method and possibly the easiest to set up from a Tivoli Storage Manager perspective. Certain functions, such as server prompted mode and the web clients, require TCP/IP.

2.3.5.3 Network names

The Tivoli Storage Manager server machine requires a name that clients can use to point to the Tivoli Storage Manager server. If TCP/IP is used, create a Domain Name Server (DNS) entry for Tivoli Storage Manager itself, as well as a DNS entry for the machine's hostname(s). For example, create DNS name TSMSRV1 and use it in the client option files. If Tivoli Storage Manager needs to be moved to another machine, only the DNS entry needs to be changed, instead of editing the Tivoli Storage Manager options file for each client. Additionally, this allows for the addition of other Tivoli Storage Manager servers. However, using DNS may impact backup availability if for some reason the DNS service is down (the DNS name will not get resolved unless availability has been built into the DNS setup).

2.4 Disk considerations and sizing

Tivoli Storage Manager requires disk to operate: it needs disk to hold the database, logs, and usually the primary storage pools. In this section we talk about how to choose a disk subsystem, and how to determine the correct amount of disk storage capacity to acquire.

2.4.1 The disk subsystem

In general, choose the fastest disk subsystem that you can afford. Slow disks may not be a hindrance for a small implementation, but as the site grows, disk access becomes a large percentage of the overall elapsed time to complete a task. Size the disk subsystem for growth because the vast majority of Tivoli Storage Manager implementations grow substantially. Choose a disk model that meets your present estimated needs and has room for expansion. Multiple I/O

paths and hotswappable components should also be considered, both for performance and availability.

See 5.6, “RAID” on page 89 for a discussion about choosing RAID devices.

2.4.2 Tivoli Storage Manager database

The Tivoli Storage Manager database size is based upon how many files are managed with Tivoli Storage Manager, and whether the files are in a primary storage pool or a copy storage pool. The database holds two types of data: entries for backups and entries for archives. The database also holds items such as entries for image backups, server scripts and the volume history, but typically, these are insignificant in sizing the database.

Backup sizing calculates the size of the Tivoli Storage Manager database holding backup entries. Archive sizing calculates the size of the Tivoli Storage Manager database holding archive entries. Use either or both depending on the data you will be storing in Tivoli Storage Manager. If both are used, then add the calculated database sizes together to arrive at the total database size. If you are planning to use backup sets, remember that these are not tracked in the Tivoli Storage Manager database, so do not impact our calculation. They do affect the tape library size if the media used is retained online. Image backups are entered in the database, although only one entry for each backup is required. The database space requirements for image backup entries will therefore be small compared to the backups and archives and we will not include this in our calculation.

Our calculations are based on the following assumptions: each backed up or archived file version uses 600 bytes of space in the server database. If offsite copies are used as well, this takes an extra 200 bytes of space. Where we did not know the number of files to be kept, but only the total storage, we estimate 5% of the total server storage requirement for onsite copies and 1% for offsite.

2.4.2.1 Backup sizing

To estimate the Tivoli Storage Manager database size for backup data, use the data collected in Table 7 on page 13 and Table 11 on page 19.

The calculation is based upon two types of numbers: the actual number of file versions backed up; and when this is not known, by a percentage of the data backed up. Steps 1, 2, 3, and 6 are based on the number of file versions. Steps 4 and 7 are based on the estimated percentage of data method.

1. Sum the field *Number of Files Backed Up* for all clients, leaving out any fields containing estimates in GB.
2. Multiply this number by the *Number of Versions Kept*, giving a total number of files backed up.
3. Multiply this number by 600 bytes, to give bytes used in the database for all known files backed up.
4. Sum of all the estimated files in GB multiplied by the *Number of Versions Kept*. Take 5% of this number to give the estimated bytes for backed up files.
5. Add bytes for known files backed up to estimated bytes backed up to give total bytes for backed up files.

6. If copy storage pools are used (and we strongly recommend this), multiply the total number of files backed up and calculated in Step 2 above by 200 bytes, giving the bytes for known copy storage pool files.
7. Sum of all estimated files in GB multiplied by the *Number of Versions Kept*. Take 1% of this number to give the estimated bytes for copy storage pool files.
8. Add bytes for known copy storage pool files to estimated bytes for copy storage pool files to give total bytes for copy storage pool files.
9. Add the total bytes for backed up files to the total bytes for copy storage pool files to give the total bytes calculated for the database.
10. Calculate 135% of total bytes calculated for the database to give the database size. This is for overhead and for growth.

For example, using the sample data in Table 7 on page 13 and Table 11 on page 19, calculate the sample database size as follows:

1. $5,000 + 15,000 = 20,000$ files
2. $20,000 * 2 = 40,000$ files
3. $40,000 * 600 = 24,000,000$ bytes (24MB)
4. $2,500 \text{ MB} * 7 * 0.05 = 875 \text{ MB}$
5. $24 \text{ MB} + 875 \text{ MB} = 899 \text{ MB}$ for backed up files
6. $40,000 * 200 = 8,000,000$ bytes (8 MB)
7. $2,500 \text{ MB} * 7 * 0.01 = 175 \text{ MB}$
8. $8 \text{ MB} + 175 \text{ MB} = 183 \text{ MB}$ for copy storage pool files
9. $875 \text{ MB} + 183 \text{ MB} = 1058 \text{ MB}$
10. $1058 \text{ MB} * 1.35 = 1428 \text{ MB}$ database size for backup files

2.4.2.2 Archive sizing

To estimate the Tivoli Storage Manager database size for archive data, use the data collected in Table 7 on page 13 and Table 11 on page 19.

The calculation is based upon the number of files archived:

1. Sum the field *Number of Files Archived* for all clients.
2. Multiply the *Number of files archived* by the *Number of archives kept* times the yearly retention ratio (i.e., desired monthly retention/12 months), giving a total number of files archived.
3. Multiply this number by 600 bytes to give total bytes for archived files.
4. If copy storage pools are used (and we strongly recommend this), then multiply the total number of files archived and calculated in Step 2 above by 200 bytes, giving the total bytes for copy storage pool files.
5. Add the total bytes for archived files to the total bytes for copy storage pool files to give the total bytes calculated for the database.
6. Calculate 135% of total bytes calculated for the database to give the database size for archives. This is for overhead and for growth.

For example, using the sample data in Table 7 on page 13 and Table 11 on page 19, (for simplicity we only present here the two Server class machines - the

process is similar for the Workstn system) calculate the sample database size for archive as follows:

1. $45,000 + 17,500 = 62,500$ files
2. $6,500 * 12 * (12/12) = 750,000$ files archived
3. $750,000 * 600 = 450,000,000$ bytes (450 MB)
4. $750,000 * 200 = 150,000,000$ bytes (150 MB)
5. $450 \text{ MB} + 150\text{MB} = 600 \text{ MB}$
6. $716 \text{ MB} * 1.35 = 810 \text{ MB}$ database size for archive data

2.4.2.3 Identify database volumes

The total required database size including both backup and archive requirements will be $1428 + 810 = 2238 \text{ MB}$

We recommend using the Tivoli Storage Manager mirroring function for the database instead of a hardware mirror or operating system mirror, because Tivoli Storage Manager has additional functions to handle error conditions that may affect the mirrored copy.

If you are using Tivoli Storage Manager mirroring, you need to plan for the mirror copy by doubling the amount of disk for the database.

Various file systems have different maximum capacities, so the database may have to be split across numerous volumes to make up your database size. To ensure no single point of failure when mirroring the database, each volume in a mirrored set should be placed on a separate drive and even controller.

Note:

Our recommended configuration will set up triggers to additionally extend the database whenever necessary. This is documented in Chapter 4, “Database and recovery log” on page 57.

By summing up backup and archive database sizes, you will have a full, consolidated database for the calculated timeframe. You can start with a smaller configuration initially but leave enough spare disk space for growth.

You should round up database volume sizes to multiples of 4 MB plus 1 MB for overhead.

Complete Table 15 with the database file names and the volume names for your primary database and copy database. This table includes the space for the backup requirements only.

Table 15. Database worksheet - backup requirements only

Database volume	Filename (Primary)	Size (MB)	Filename (Copy)	Size (MB)
Vol1	/tsm/database/primary/file01	717	/tsm/database/copy/file01	717
Vol2	/tsm/database/primary/file02	717	/tsm/database/copy/file02	717

Database volume	Filename (Primary)	Size (MB)	Filename (Copy)	Size (MB)
	Total	1434	Total	1434

You can also use a full database allocation, using both backup and archive requirements, as shown in Table 16:

Table 16. Database worksheet - backup and archive requirements

Database volume	Filename (Primary)	Size (MB)	Filename (Copy)	Size (MB)
Vol1	/tsm/database/primary/file01	561	/tsm/database/copy/file01	561
Vol2	/tsm/database/primary/file02	561	/tsm/database/copy/file02	561
Vol3	/tsm/database/primary/file03	561	/tsm/database/copy/file03	561
Vol4	/tsm/database/primary/file04	561	/tsm/database/copy/file04	561
	Total	2244	Total	2244

2.4.3 Tivoli Storage Manager recovery log

The size of the recovery log depends on the amount of data changed between Tivoli Storage Manager database backups. The larger the amount of data, the larger the recovery log needs to be. Either a full or an incremental Tivoli Storage Manager database backup (in roll-forward mode) resets the recovery log back to empty. If the recovery log fills up completely, Tivoli Storage Manager stops and you have to manually increase the size of the recovery log. This may take some time but can usually be avoided with adequate precautions (for example, by monitoring and planning for growth).

Note:

Our recommended configuration will setup triggers to additionally extend the recovery log whenever necessary. This is documented in Chapter 4, "Database and recovery log" on page 57.

You should round up recovery log volume sizes to multiples of 4 MB plus 1 MB for overhead.

To estimate the size of the recovery log, multiply the database size by the percentage of data that changes each backup cycle. Double this number to allow for two backup cycles to occur without a database backup. This gives a starting point for the recovery log.

For example, if the database size is 1434 MB, and 5% of the data changes every backup cycle, then the estimated size for the recovery log would be $1434 \text{ MB} \times 0.05 \times 2 = 143.4 \text{ MB}$ (145 MB for better allocation, if using a single volume).

As with the database, we recommend using the Tivoli Storage Manager mirroring function for the recovery log instead of a hardware mirror or operating system mirror.

If you are using Tivoli Storage Manager mirroring, you need to plan for the mirror copy by doubling the amount of disk for the recovery log.

Various file systems have different maximum capacities, so the recovery log may have to be split across numerous volumes to make up your total recovery log size. To ensure no single point of failure when mirroring the recovery log, each volume in a mirrored set should be placed on a separate drive and even controller.

Complete Table 17 with the recovery log file names for your primary and copy recovery log.

Table 17. Recovery log worksheet

Log Volume	Filename (Primary)	Size (MB)	Filename (Copy)	Size (MB)
Vol1	/tsm/log/primary/ file01	144	/tsm/log/copy/file01	144
	Total	288	Total	288

2.4.4 Primary disk storage pool

In our scenario, when a client node backs up, archives, or migrates data (during HSM function), the data is stored in the same primary storage pool. You could also use separate storage pools for a backup, archive and HSM data for improved controls and manageability of production data.

To size the disk storage pool, calculate the amount of data that will be backed up during one cycle and add in a proportion of the amount of archive data transferred to the server. This amount (plus a contingency for growth) is the recommended size the storage pool needs to be.

If you are also using space management (HSM) features, the rate of data migration to the server is hard to predict, so you need to get experience of your particular environment to make an accurate sizing.

Primary storage pools (usually disk devices) can be made larger or smaller than the recommended size, depending on the resources available. A larger pool size would allow for more than one backup cycle of data to remain on disk, thus improving restore times. It also allows for spare capacity for an unexpectedly large amount of backup data to prevent server migration from running during backup. A smaller primary storage pool uses less disk, but runs the risk that the pool will fill up during back up. This is not functionally a problem, since the backup and migration to the next storage pool can execute concurrently, however performance will degrade.

We recommend using a primary disk storage pool at the recommended size to reduce interference from migration while backup is running.

Backup disk sizing

To estimate a primary storage pool size if running backup cycles only, perform the following steps:

1. Using Table 7 on page 13, multiply the *GB changed per backup* by (1 - the *Data compression* rate) to give the total bytes transferred for each client.
2. Sum the total bytes transferred for all clients to give the total bytes transferred per backup cycle.
3. Add 15% to total bytes transferred per backup cycle to give the storage pool size. This allows for variability in the size and number of files per backup.

For example, using the sample figures in Table 7 on page 13, the *GB changed per backup* are 0.01, 0.8, and 1.5 while the *Data compression* figures are 0.5, 0.5, and 0.66 respectively.

1. Multiplying 0.01 by (1 - 0.5) gives 0.005, and multiplying 0.8 by (1 - 0.5) gives 0.4, and multiplying 1.5 by (1 - 0.66) gives 0.5.
2. Summing 0.005, 0.4, and 0.5 gives 0.905 GB.
3. Adding 15% gives a storage pool size of 1.04 GB.

Archive disk sizing

In most environments, archive disk sizing is less critical than for backup. This is because backup operations tend to run much more frequently (usually nightly) and to a stricter timeframe. Archives may run less frequently, and on weekends when the overall workload is lighter. Since in these circumstances, the impact of running concurrent archive and storage pool migration might not be so critical, it is not normally necessary to use the full archive size in calculating storage pool requirements. Another factor to be considered for sites which are doing frequent archiving of large amounts of data, where the storage pool and processing impact are great, is the possibility of substituting backup set operations for archive. The advantage of this is that generating a backup set requires the Tivoli Storage Manager server only - the client is not involved in any way since the backup set is created using data which has already been sent to the server in normal backup operations.

To additionally increase the storage pool to hold archive data as well, use the following steps:

1. Using Table 7 on page 13 and the field *GB copied per archive*, group all machines that require simultaneous archive operations during one common timeframe (for example, every month). Select the biggest group, giving *peak archive size*.
2. Take 10% of *peak archive size*, giving *archive storage size*.

For example, using the sample figures in Table 7 on page 13:

1. Assuming a monthly/month end timeframe as our baseline, we have *Client 2* (8 GB) and *Client 3* (5 GB), which equals 13 GB.
2. Taking 10% from all archive storage required during the timeframe, $13 \text{ GB} * 0.10 = 1.3 \text{ GB}$ for archive disk storage.

Image disk sizing

If you are planning to use image backups for your UNIX machines, you would consider sizing the disk storage pool to hold the filespace you want to backup.

This is because image backups are single objects and therefore, the server will require that size for storing data (or at least part of it, assuming compression is enabled)

For example, assuming that the client filesystems eligible for image backup are /oralog (200 MB), /finsys (1.5 GB), /oradata (1 GB) and /findata (1 GB), then the disk image requirements should be at least 3.7 GB to hold those filesystems in disk without having to use tape immediately. This is especially true if you are running parallel backup operations (that is, executing multiple concurrent `backup image` commands). Alternatively, since the disk storage requirements are so high for this operation, you could consider sending these backups straight to tape, provided that you have enough tape drives and/or backup window to co-exist with normal backup and archive operations as well.

Note:

You must sum up all disk storage requirements (backup, archive, image) to have your final disk storage size. You can of course, start with smaller numbers and evaluate future growth.

We recommend that the disk storage pools be allocated on fault tolerant hardware devices such as RAID 5 devices. If you are using hardware or operating system mirroring, you need to plan for the mirror copy by doubling the amount of disk for the primary storage pool.

Various file systems have different maximum capacities, so the primary storage pool may have to be split across numerous volumes to make up your total primary storage pool size. We recommend that the disk storage pools be placed on their own disk devices and controller separate from the database and the recovery log, if possible.

Complete Table 18 with the primary storage pool file names and volume names for your primary storage pool.

Table 18. Primary storage pool worksheet

Filename	Size (MB)
/tsm/stgpool/file01	400
/tsm/stgpool/file02	400
/tsm/stgpool/file03	400
/tsm/stgpool/file04	400
/tsm/stgpool/file05	400
/tsm/stgpool/file06	400
Total	2,400

2.4.5 Device configuration table and volume history file

The device configuration table and the volume history table also require disk space, but typically, they are very small. The device configuration table contains entries for defined device classes and definitions for drives and libraries. Every volume that is used by Tivoli Storage Manager is tracked in the volume history

database, including the volume identifier for the database backups. The volume history information is periodically copied out to a volume history file that you can specify with the `VOLUMEHISTORY` option in the `dsmserv.opt` file.

We recommend that you define at least two copies of both the device configuration table and the volume history file, in case one becomes unusable due to a hardware or software failure.

Complete Table 19 with the device configuration and volume history filenames and sizes.

Table 19. Device configuration and volume history worksheets

Name	Size (MB)
/tsm/files/primary/devconfig1.out	0.5
/tsm/files/copy/devconfig2.out	0.5
/tsm/files/primary/volhist1.out	0.5
/tsm/files/copy/volhist2.out	0.5
Total	2

2.4.6 Total disk

Total disk refers only to the numbers discussed here. If you are using mirroring or some other version of RAID, you need to take that into consideration separately. The disk required to run the server platform operating system efficiently also has not been considered.

Tivoli Storage Manager code requirements for disk vary depending on the server platform and release level. We use an estimate of 150 MB as the time of writing. Table 20 summarizes the disk requirements for the Tivoli Storage Manager server as we have planned it.

Table 20. Total Tivoli Storage Manager disk required worksheet

	Size (MB)
Tivoli Storage Manager code (dependent on platform)	150
Tivoli Storage Manager database	2,888
Tivoli Storage Manager recovery log	576
Primary storage pools	2,400
Device configuration table and volume history table	2
Other (RAID, Operating system)	0
Total	6,016

2.5 Tape considerations and sizing

Most Tivoli Storage Manager systems use tape as the ultimate destination for storing their data. There are a variety of tape subsystems available for most Tivoli Storage Manager server platforms. In this section, we discuss various tape configurations and how to size these configurations.

In general, choose the biggest, fastest, most automated tape drive solution with the greatest number of drives you can afford.

Optical disks are an alternative to tape, but we recommend using tape for most Tivoli Storage Manager implementations, because tape is generally faster and more convenient to work with.

Planning for and development of manual procedures in Tivoli Storage Manager may be required to identify and move volumes from copy storage pool to the offsite location and back. Alternatively, Tivoli Disaster Recovery Manager can be used to automate these procedures.

Complete the worksheet in Table 21 using the information from the sections below. The data collected will be used when sizing tape libraries and then when defining them to Tivoli Storage Manager.

Table 21. Tape drive configuration worksheet

	Option
Library model	IBM 3570 MP
Number of drives	2
Drive model	3570
Number of onsite tape volumes	32
Number of offsite tape volumes	48
Number of database volumes	6
Number of scratch tapes	13
Number of backup set tape volumes	4
Total tape volumes required	103

2.5.1 Tape devices

Tape drives come in 4mm, 8mm, DAT, 3570, 3590, and other device types. Each type of drive has a different data capacity, performance, cost, and reliability characteristics. Although data capacity and cost per megabyte stored are important, reliability is much more important. Having saved money buying tapes is small consolation when you are unable to restore your customer billing database due to a tape error.

In general, tape drives where the tape touches the read/write heads, such as 4mm and 8mm, tend to be less reliable (and slower) than tape drives where the tape does not touch the read/write heads, such as 3570 and 3590. If you do implement drives that touch the read/write heads, plan to replace the tapes at regular intervals.

2.5.2 To library or not to library?

Tivoli Storage Manager can become tape-intensive. The larger the number and size of clients you are backing up, the greater the load placed on the tape drives. As this load increases, manual tape handling can become time consuming and can create a bottleneck in system performance.

An automated tape library permits a Tivoli Storage Manager implementation to run a nearly lights-out operation. Manual intervention should only be required to insert and remove tapes for offsite storage. A tape library having a bar code reader for tape labels reduces errors.

When choosing a tape library, pay attention to the library cycle time, including mount, search, transfer, and dismount time. For busy libraries, a long cycle time can adversely affect performance of Tivoli Storage Manager.

We recommend an automated tape library with a bar code reader for all but the smallest of Tivoli Storage Manager implementations.

2.5.3 Number of tape drives

In attempting to reduce costs, many Tivoli Storage Manager customers acquire only one or at most two tape drives. One-drive systems are possible but are subject to outages if the sole tape drive ever fails. Such systems will require more intervention during server operations like reclamation and when creating offsite data copies. Two-drive systems allow for quicker reclamation and pool copies and reduced drive outages, but restores coming from tape may be delayed due to other competing tape activity. Systems with three or more tape drives can handle restores from tape occurring while tape reclamation or other tape processing is in progress.

A critical consideration for the number of tape drives is how quickly you need to restore one or more clients. More tape drives allow you to restore more clients in a given time. Collocation allows as many clients to be restored simultaneously as you have tape drives. If you have stringent restore requirements, collocation and multiple tape drives make sense. You could also consider features like backup sets and logical volume backups for highly time-dependant servers to improve their restore window.

To calculate the rate of a restore operation, divide the amount of storage to be restored by the sustained data rate of the tape drive (not the instantaneous, or burst, rate) quoted by the manufacturer. If this number does not allow you to meet your service levels, collocation, more tape drives, faster tape drives, a new service level, or another backup strategy may be required.

We recommend a minimum of two tape drives on any Tivoli Storage Manager system, with at least three tape drives being preferable.

2.5.4 Number of tape volumes

Tape volumes are used to store onsite copies of data, offsite copies of data and database backups. Additional tape volumes are required because all tape volumes are not filled to capacity. Some volumes are required to stock a scratch pool so that mounts for unused tape volumes can be satisfied. If you are considering backup sets and image backups, you may need to add some more tapes to your scratch pool.

2.5.4.1 Onsite volumes

You must consider how many tapes you will need to hold your backup, archive and image data in a specified timeframe. Each of those have different expiration requirements, and therefore, the calculation may not be linear. In any case, keep

in mind that it is best to assume a middle case / worst case scenario, so that you do not run out of tapes.

Backup tapes

To calculate the number of onsite tape volumes required for backup operations, carry out the following calculations:

1. If this is a sequential storage pool (tape device), multiply the primary storage pool size by the *Number of backup versions* from Table 11 on page 19 to give versions pool size.
2. Add the sum of all *Total storage used* fields for each client from Table 7 on page 13 to the versions pool size, giving tape pool size.
3. Divide the tape pool size by the device capacity to find the number of tape cartridges required.
4. Add 50% to cater to tapes that are in *Filling* status to give the total cartridges required for onsite tapes.
5. If using collocation, normally there should be at least as many tape cartridges as there are clients. Consider tape native capacity as a rounding factor.

We do not use the compressed capacity of the tape here because we factored the client compression rate into the calculation of the storage pool size. If data is compressed at the client, it will not receive any benefit from hardware compression done by the tape drive.

Example:

1. If the primary storage pool is 1.04 GB, and the number of versions kept is 7, then multiplying $1.04 * 7$ gives a 7.3 GB pool size.
2. If all the clients are using 0.75, 8, and 25 GB, then $0.75+8+25+7.3$ equals 41.05 GB.
3. If the tape device has a capacity of 5 GB, then dividing 41.05 by 5 gives 9 cartridges required to store all the data.
4. $9 * 1.5$ gives 14 total cartridges required for onsite tapes.
5. If we use collocation and in this example there are only three clients, we have enough tape volumes.

Archive tapes

1. Using Table 7 on page 13, multiply the *GB copied per archive* by $(1 - \text{the Data compression rate})$ giving transferred archive data.
2. For each client, multiply transferred archive data by *Number of archives kept* in a year times the yearly retention ratio (that is, desired monthly retention/12 months), and sum up giving total archive data.
3. Divide total archive data by native tape capacity, giving the total number of tapes required.

Example:

1. The clients in our table have 1, 8 and 5 GB of archive data. The compression ratios are 0.5, 0.5 and 0.66. This equals $1*(0.5)$, $8*(0.5)$, $5*(0.34)$ which equals 0.5, 4.0, 1.7 GB.

2. The number of archives kept are 2, 12 and 12. The yearly retention ratios are 6/12, 12/12 and 12/12. This equals $0.5 \times 2 \times 0.5 + 4.0 \times 12 \times 1 + 1.7 \times 12 \times 1 = 68.9$ GB.
3. If the tape device has a capacity of 5 GB, then dividing 68.9 GB by 5 gives 14 cartridges required to store all archive data.

Image backup tapes

Image requirements are calculated on total allocated filesystem space. For simplicity sake, we are not considering any compression rating (you can use the compression ratio for the client on the sum of filesystems), and therefore the number is already overestimated without any need for extra tapes:

1. Consider all client filesystems eligible for image backup. Sum them all, giving transferred image size per client.
2. Using Table 7 on page 13, Multiply the *Number of image backups* by transferred image size per client times yearly retention ratio (that is, desired monthly retention/12 months), and sum them all giving total image size.
3. Divide total image size by native tape capacity, giving the total number of tapes required.

Example:

1. In our example, we assume that the sum of all of Client 3 eligible filesystems for logical volume backup would be 10 GB.
2. The retention ratio is 1/12. Assuming previous calculation, this gives $12 \times 10 \times 0.16 = 20$ GB.
3. If the tape device has a capacity of 5 GB, then dividing 20 GB by 5 gives 4 cartridges required to store all image backup data (assuming no compression).

Therefore, the total onsite tape requirements for this example would be:

$$14 + 14 + 4 = 32 \text{ tapes.}$$

2.5.4.2 Offsite volumes

You will probably have more tapes in your offsite pool than onsite, because of less frequent reclamation and partially filled tapes.

To estimate the number of tape cartridges required for offsite copies, use the number of onsite tape volumes calculated in the previous section. As a rough guide, add 50% to estimate the number of tapes required. In our example, this comes to 48 offsite tapes. Keep in mind that collocation may be set on for either, both, or none of onsite and offsite tape pools, which will also affect tape volume requirements.

2.5.4.3 Database tape volumes

Each database backup requires at least one tape volume. We recommend backing up the database every day and keeping these backups for at least five days.

To calculate the database tape volumes required:

1. Divide the database size calculated in 2.4.2, "Tivoli Storage Manager database" on page 29 by the tape device capacity and round up to the nearest

whole number to give the number of tape volumes required for one database backup.

2. Multiply this number by six (five copies plus the copy just being made) to give the total number of tape volumes required for database backups.

For example, If the database size is 1428 MB, and the device capacity is 5 GB:

1. $1.428 \text{ GB} / 5 = 0.07 \text{ GB}$ rounds up to 1 tape volume.
2. $1 \text{ tape} * 6 \text{ versions} = 6 \text{ tape volumes}$.

2.5.4.4 Scratch volumes

A scratch (or empty) volume is required every time Tivoli Storage Manager wants to write to an unused tape.

To estimate the number of scratch volumes required:

1. Total the number of tape volumes required for onsite tapes, offsite tapes and database backup tapes.
2. Calculate 15% of this number to allocate for scratch tape volumes.

We have a total of 86 onsite, offsite and scratch tapes. 15% of this is 13 tapes.

2.5.4.5 Backup set volumes

It is worth calculating backup set tape space requirements separately from regular file-based processing due to the nature of backup sets. Note that although the backup set is a server initiated procedure, our calculations are based on previous file-backup operations. Note that this step may not be necessary, since you may create backup sets onto disk and copy them onto media such as CD or use another tape format which is common to both the server and the client. For this calculation we are assuming the backup set consists of the entire client system's file, however you may only use a subset of its filespaces.

To estimate the number of backup set volumes required:

1. Using Table 7 on page 13, multiply the *Total storage used* by $(1 - \text{the Data compression rate})$ to give the total gigabytes for each client. Round up to the next multiple of the native tape capacity, since backup sets, like server database backups, cannot stack onto tapes.
2. Multiply *Number of backup sets* by the total gigabytes for that client to give the client backup set size.
3. Multiply client backup set size by the yearly retention ratio (that is, desired monthly retention/12 months) for that client. Sum all client backup set sizes to give the total backup set requirement.
4. Divide the total backup set requirement by the device capacity to find the number of tape cartridges required.

Example:

1. Multiplying 0.75 by $(1 - 0.5)$ gives 0.375, and multiplying 8 by $(1 - 0.5)$ gives 4, and multiplying 25 by $(1 - 0.66)$ gives 8.5. Assuming native tape capacity of 5Gb, this rounds up to 5, 5, 10 GB respectively.
2. From previous calculations, client backup set sizes are 2x5, 6x5, 6x10.
3. Assuming retention periods of 6/12, 2/12 and 2/12, this gives $10 * 0.5 + 30 * 0.16 + 60 * 0.16 = 20 \text{ GB}$ as backup set space requirement.

4. If the tape device has a capacity of 5 GB, then dividing 20 by 5 gives 4 cartridges required to store all backup set data during the specified timeframe.

Adding up all these numbers as stored in Table 21 on page 37 gives a total of 103 tape volumes required. You then need to consider this number against the library configurations available for your chosen technology. For example, suppose you were considering a tape library which came in models with 96, 180 and 256 slots, you would select the 180-slot model. This would give some space for future growth. It should be noted that Tivoli Storage Manager setups rarely if ever shrink!

2.5.5 OS/390 tape management

On all platforms except for AS/400 and OS/390 (or MVS), Tivoli Storage Manager provides its own tape management functions. For AS/400 and OS/390, it uses the functions provided within the operating system or other tape library management systems. For more information on tape management in Tivoli Storage Manager on AS/400, see the redbook *A Practical Guide to Implementing Tivoli Storage Manager on AS/400*, SG24-5472.

Tivoli Storage Manager uses data set names to identify various types of Tivoli Storage Manager data sets. The data set name prefix is set by the device class parameter PREFIX. Each device class can have a different data set name prefix. The data set name suffix is fixed by Tivoli Storage Manager for various data types. The suffix .DBB indicates a database backup data set. The suffix .BFS indicates an onsite or offsite data copy.

Most tape library management systems on OS/390 use the data set name to identify tapes to be taken offsite. Since Tivoli Storage Manager uses the data set name PREFIX.BFS for both onsite and offsite copies, the tape management system has no way to identify tapes that must be moved offsite.

To choose another data set name for offsite copies, create another device class for the offsite copies and choose a different prefix. Set the tape library management system to trigger on this different PREFIX.BFS data set name and offsite copies will be identified automatically.

Tivoli Storage Manager allows an external data manager (EDM) to control tapes. To inform the EDM when a tape goes into a scratch status, you can use the DFSMSshm ARCTVEXT parameter. Include the DELETIONEXIT ARCTVEXT parameter in the Tivoli Storage Manager server options file. For more information, see *Tivoli Storage Manager for MVS and OS/390 Administrator's Guide*, GC35-0377.

If your tape management system uses program names to identify External Data Managers, the Tivoli Storage Manager program name is ANRSERV.

2.6 Administrator IDs

Identify who will be the Tivoli Storage Manager administrator(s) in your organization.

A Tivoli Storage Manager administrator controls Tivoli Storage Manager resources. There can be numerous administrators with varying levels of authority.

It is possible to use the web backup-archive client to perform backup, restore, archive, and retrieve operations on the behalf of other users using only a Web browser such as Netscape or Internet Explorer. Help desk personnel can use this to perform these client tasks for their end users without having to logon to the client machine.

If your Tivoli Storage Manager installation is large or widely dispersed, you can delegate some authority to administrators based on policy domains or storage pools. Therefore, a workstation administrator could look after setting data retention criteria for workstation data only (assuming the correct policy domains were set up). A UNIX administrator could be given Tivoli Storage Manager authority over UNIX data retention criteria only.

Since Tivoli Storage Manager logs all commands issued by administrators and it has no limit on the number of administrators, do not share administrator IDs. Sharing administrator IDs reduces the accountability of each ID, and therefore of all the people sharing the ID. Conversely, numerous administrator IDs may give too many people too much authority.

Table 22 suggests several administrator IDs you may want to implement.

Table 22. Administrator IDs worksheet

Functions	Administrator ID	Authority
Server console	SERVER_CONSOLE	System
System administrator	sysadmin	System
System support	support	System
System reporter	reporter	none
Client administrator	helpdesk	Node

See 7.1, “Administrators” on page 143 for more information about administrators. We recommend you delete the default installed administrator, `admin`, to prevent the possibility of this ID being misused. Many sites leave this ID with its default password, creating a big security hole for any malicious person with basic Tivoli Storage Manager knowledge!

2.7 License considerations

Licenses are used by Tivoli Storage Manager in three ways. First, the Tivoli Storage Manager server must be licensed to obtain the server code and to legally operate the server. Second, a license is required for each client that Tivoli Storage Manager will manage. Third, optional features must be licensed to enable the feature on the Tivoli Storage Manager server.

Complete Table 23 with the information below for each Tivoli Storage Manager server. See 8.1, “Licensed features” on page 157 for more information.

Table 23. License requirements worksheet

	Required
Server type	Windows NT
Client connections	1, 250

	Required
Network connections	Yes
Open Systems Environment clients	No
Space management	No
Tivoli Disaster Recovery Manager	No
Advanced device support	Yes
Tivoli Data Protection Features	
TDP for Oracle	No
TDP for MS SQL	No
TDP for MS Exchange	No
TDP for Lotus Notes	No
TDP for Lotus Domino	No
TDP for Informix	No
TDP for R/3	No

2.7.1 Server type

Enter the Tivoli Storage Manager server platform type you are licensing. For example, use Windows NT, AIX, MVS or OS/390, VM, HP-UX, Sun Solaris, or AS/400.

2.7.2 Client connections

Enter one or more of the following to add up to the number of Tivoli Storage Manager backup-archive clients you want to administer.

- A x 1, where A is the number of single licenses required.
- B x 5, where B is the number of 5 client licenses required.
- C x 10, where C is the number of 10 client licenses required.
- D x 50, where D is the number of 50 client licenses required.

2.7.3 Network connections

Enter Yes to license a network communication method like TCP/IP. A network license is not required for the following:

- Administrative client web interface (HTTP communication method)
- Administrative client access
- Use of a Tivoli or an SNMP receiver
- Backing up the server system itself using the named pipe or shared memory communication method

2.7.4 OpenSystems environment

Enter Yes if you want to license the OS/390 UNIX Systems Services client.

2.7.5 Tivoli Space Manager optional product

Enter Yes to license the Tivoli Space Manager.

2.7.6 Tivoli Disaster Recovery Manager optional product

Enter Yes if you want to use the Tivoli Disaster Recovery Manager.

2.7.7 Advanced device support

Enter the number and type of advanced device support licenses required. Look at the URL http://www.tivoli.com/support/storage_mgr/adsercli.htm for more information about what devices are supported.

2.7.8 Tivoli Data Protection features

Enter Yes if you need to license any of the Tivoli Data Protection features for your installation.

2.8 Other considerations

There are numerous other topics to be considered when planning a Tivoli Storage Manager installation. Many of these topics are outside the scope of this book, but we are mentioning them for completeness.

- **Staffing:** Staffing requirements need to be addressed. The various functions such as operations, technical support, administration and help desk may all be performed by one person in a small site. Larger sites may find a more specialized approach useful. To provide backup coverage, two people per function is always a good idea.
- **Lead time:** Some tasks such as installing a tape exit in MVS may have considerable lead times before the change can be made. We have highlighted some of these, but check with your technical support group and your change management group for their guidelines.
- **Monitoring:** You may want to consider monitoring your Tivoli Storage Manager server system using a product such as Tivoli Enterprise or Tivoli Decision Support. We have highlighted some suggestions for this, but there are many more items that you may want to monitor. Monitoring also includes monitoring the health of the Tivoli Storage Manager software. Numerous queries are useful for displaying information about the Tivoli Storage Manager system and its workings. We have included some basic possibilities.
- **Chargeback:** Some Tivoli Storage Manager installations charge for their services. This is possible using the accounting records and site specific programs. Some items you may want to consider charging for include bytes stored, CPU time per client, or tapes used.
- **Code refreshes:** New client code typically has been released every three months. With installations greater than about 50 clients, keeping up with these refreshes of client code becomes a problem. Set up a procedure for tracking which clients are running which release of Tivoli Storage Manager code. Design your client installs to be as generic and as similar as possible. If an automated software install process is available, consult with the process owners regarding the best practices to use in setting up Tivoli Storage Manager clients.

- **Export/Import:** It is possible to export a client definition and all of its related data from one Tivoli Storage Manager server and import it into another Tivoli Storage Manager server. This facility is useful for moving clients from one server platform to another server on the same or different platform. With a large number of clients, or clients with a large amount of data, the export and import can take a significant amount of time, in the order of 24 to 48 hours. In these cases, planning and coordination needs to be done as to when the exporting server ceases backups and the importing server starts backups.
- **Server scripts:** These are very useful for issuing Tivoli Storage Manager commands repeatedly, or with some rudimentary logic around it.
- **SQL queries:** These are powerful queries you can run against the Tivoli Storage Manager database to extract information.
- **Problem determination:** Diagnosing and resolving problems are tasks you will have to do on a regular basis. In general, once you have determined you have a problem, install the latest level of Tivoli Storage Manager client and Tivoli Storage Manager server code and try to recreate the problem. If it still exists, you should engage formal software support procedures. Alternatively, you may search or post your problem to the Tivoli Storage Manager listserv list to see if someone else has experienced the problem or who may be able to offer some suggestions.
- **Disaster recovery:** We recommend planning and testing for disaster recovery be done on a regular basis. The optional product Tivoli Disaster Recovery Manager assists in gathering, maintaining and recommending information and planning pertinent to disaster recovery.

Chapter 3. Tivoli Storage Manager server

In this chapter, we explain steps relating to the implementation of a Tivoli Storage Manager server. We cover the topics of code installation, options file customization, and other server implementation tasks that are required for creating this redbook's Tivoli Storage Manager environment. See 2.3, "Server architecture considerations" on page 22 for planning considerations.

3.1 Latest code updates

Tivoli Storage Manager server and client code fixes and enhancements are released on a regular basis. The fixes are available from IBM via the Internet or CD-ROM. If you have a fast reliable Internet connection, then we recommend that you download the code. You should be aware that the fixes can be quite large as they are full or near-full code replacements. If you are unable to download from the Internet, you can order the fixes through your usual IBM service channel.

The Tivoli Storage Manager has links to the latest Tivoli Storage Manager server fixes, Tivoli Data Protection for application fixes, Tivoli Storage Manager client code, and important download information at:

http://www.tivoli.com/products/index/storage_mgr/index.html

Each link leads you to sets of files on the ftp server describing and containing the fixes at:

<ftp://service.boulder.ibm.com/storage/tivoli-storage-management/maintenance/>

The names and content of these files may vary slightly across platforms:

- README.ftp

This file contains the download and install instructions.

- README.1st

This file contains important information that is not yet available in manuals. This includes information such as description of code fixes, enhancements, limitations, and additions or corrections to the hardcopy publications.

- Code image(s)

These files contain the actual code fixes. You should download and unpack these files exactly as instructed in the `README.ftp` file.

If the fix is unavailable for download or your environment is not suitable for downloading large files, you can use the fix number to order the code on CD-ROM from IBM support.

We recommend that you keep a copy of the latest server, agent, and client fix files on a suitable file server at your site. This allows easier distribution and code installation, especially for clients. We recommend that for non-MVS systems, you use the Tivoli Storage Manager server to store the code.

3.2 Code installation

You should use the instructions in the associated *Quick Start* manual for your chosen Tivoli Storage Manager server platform to install the Tivoli Storage Manager code. The *Quick Start* manual is shipped with the install media. The latest version of this manual is also available from the Tivoli Storage Manager Web site in either HTML or PDF format. We recommend that you download the manual to ensure that you are working with the latest information. The URL for Tivoli Storage Manager publications is:

http://www.tivoli.com/support/storage_mgr/pubs/admanual.htm

AIX:

Chapter one of *Tivoli Storage Manager for AIX Quick Start*, GC35-0367, contains the steps to install the Tivoli Storage Manager code in a section titled, Installing and Setting Up Tivoli Storage Manager. You should install all the components. You should not install the sample command scripts. The sample scripts and our recommended scripts will be installed later.

MVS:

The Program Directory contains the steps to install all the server and administrator client code. *Tivoli Storage Manager for MVS and OS/390 Quick Start*, GC35-0376, does not contain the steps to install the code but does contain steps for setting up and customizing a Tivoli Storage Manager environment.

Windows NT:

Chapter one of *Tivoli Storage Manager for Windows NT Quick Start*, GC35-0379, contains the steps to install all the server and client code. You should accept all the default options.

3.3 Customization

The Tivoli Storage Manager server has a number of options and settings which control its operation. You specify the options in the server options file either by manually editing it or by using the GUI interface provided on some platforms. Tivoli Storage Manager uses the options specified in the options file at server start-up. You can specify many server settings via an administrative interface.

3.3.1 Options file

Tivoli Storage Manager provides a server options file with a set of default options to start the server. You can modify server options by using a text editor. On some platforms, Tivoli Storage Manager provides a server options file editor to perform this function. The supplied file contains information on what options and option values can be specified. You can display the current server options via the `query options` command.

We provide recommended options files for the various server platforms, because the layout of the supplied options file is not easy to understand. See B.2, “Server options files” on page 292, which contains our server option files. The file options are presented in a more logical sequence than those provided during the install. We assume that TCP/IP is the network protocol and that administrators can access the server via a browser such as Netscape or Internet Explorer.

Table 24 shows the default location of the Tivoli Storage Manager server options file by server platform. We recommend that you take a backup copy of the existing options file before you replace it with your updated version.

Table 24. Server options file location

Platform	File Location
AIX	/usr/tivoli/tsm/server/bin/dsmserv.opt
HP/UX	/opt/tivoli/tsm/server/bin/dsmserv.opt
MVS	As specified by DDname OPTIONS
Solaris	/opt/tivoli/tsm/server/bin/dsmserv.opt
Windows NT	C:\Program Files\Tivoli\tsm\server\dsmserv.opt

Although the number of server options is very large, there are only a small number that need to be changed for each server. These options fall into the following categories:

- Communication options
- Site-dependent options
- Performance options

Note:

At server initialization, the server reads the server options file. If you update a server option by editing the file, you must remember to stop and restart the server to activate the updated server options file.

3.3.1.1 Communication options

If your network protocol is not TCP/IP, then we recommend that you update the Communications Options section of our options file to support your protocols. A server can support multiple communication methods. The *Quick Start* manual provides the details on what options and values to specify.

3.3.1.2 Site-dependent options

If your locale is not the USA, then we recommend that you update the Date, Number, Time, and Language Options section of the options file to support your requirements. The appropriate Administrator’s Reference manual for your operating system platform (listed in D.3, “Tivoli Storage Management Publications” on page 304) provides the details on which options and values to specify. You must specify the value for each of these options that was determined during the planning phase. Some options do not apply to some server platforms.

Refer to the Administrator's Reference manual for specific details. Table 25 shows those server options.

Table 25. Server options: site dependent

Option	Description
DATEFORMAT	Specifies the format by which dates are displayed by the server
LANGUAGE	Specifies the language used for messages
NUMBERFORMAT	Specifies the format by which numbers are displayed by the server
TIMEFORMAT	Specifies the format by which times are displayed by the server
DEVCONFIG	The file in which Tivoli Storage Manager stores a backup copy of the device configuration information
VOLUMEHISTORY	The file in which Tivoli Storage Manager stores a backup copy of the server sequential volume history information
MIRRORWRITE log sequential	Avoids potential log data corruption when operating system fails during log update
ENABLE3590LIBRARY	Support for IBM 3590 tape drives in an IBM 349x automated library

3.3.1.3 Performance options

Table 26 identifies performance options for which we recommend different values than the default. Refer to the *Administration Reference* manual for specific details.

Table 26. Server options: performance

Server Option and Value	Description
EXPINTERVAL 0	Specifies the interval, in hours, between automatic inventory expiration runs by Tivoli Storage Manager. When set to 0, inventory expiration is disabled and can only be run manually.
TXNGROUPMAX 256	Specifies the number of files that are transferred as a group between a client and the server between transaction commit points
MOVEBATCHSIZE 256	Specifies the number of files that are to be moved and grouped together in a batch, within the same server transaction
MOVESIZETHRESH 500	Specifies a threshold for the amount of data moved as a batch, within the same server transaction
LOGPOOLSIZE 2048	Specifies the size of the recovery log buffer pool size
BUFPOOLSIZE 32768	Specifies the minimum size of the database buffer pool
SELFTUNEBUFPOOLSIZE YES	Dynamically adjust the database buffer pool
TCPWINDOWSIZE 32	Specifies the size of the buffer used when sending or receiving data.

3.3.2 Settings

Tivoli Storage Manager provides default server run-time settings. These settings are stored in the Tivoli Storage Manager database and persist across server restarts. You specify the server settings via the administrative interface. You can display the current server settings via the `query status` command.

The default values for the server settings are generally acceptable. There are a few settings for which we recommend values other than the defaults. The settings are changed via the `set` command and fall into the following categories:

- Security related settings
- Our recommended settings

3.3.2.1 Security related settings

Table 27 shows the various default security related settings for the server. We recommend that you select values for these settings which reflect the security policy at your site.

Table 27. Security related settings

Setting and Value	Description
AUTHENTICATION ON	Whether administrators and client nodes must use a password to access the server
INVALIDPWLIMIT 3	Maximum number of failed logon attempts before a node or administrator is locked
MINPWLENGTH 6	Minimum length of a password
PASSEXP 90	Password expiration period
WEBAUTHTIMEOUT 0	Time-out interval for the Web administrative interface

For example, consider the security policy at a company which specifies that all passwords must be changed every 90 days and must be a minimum of six characters in length. If three invalid passwords are entered in response to a password prompt, that user must have the system access revoked. Administrative users accessing the server via the Web interface are never to be timed out.

The policy is implemented by the following administrative commands:

```
tsm: TSM010> set authentication on
ANR2095I Authentication parameter set to ON.

tsm: TSM010> set passexp 90
ANR2092I Password expiration period set to 90 days.

tsm: TSM010> set minpwlength 6
ANR2138I Minimum password length set to 6.

tsm: TSM010> set invalidpwlmit 3
ANR2175I Invalid password limit set to 3 attempts.

tsm: TSM010> set webauthtimeout 0
ANR2198I SET WEBAUTHTIMEOUT: Web authentication time-out set to 0 minutes.
```

3.3.2.2 Recommended server settings

We recommend changing some of the default server settings.

The accounting server setting determines whether an accounting record is created every time a client node session ends. The default is not to create these records. We recommend that accounting be switched on to collect these records. For MVS systems, this information is recorded in the Systems Management Facility (SMF). For other systems, this information is written to a file which contains text records that can be viewed directly or can be imported into a spreadsheet program.

All server activity is recorded in an activity log. This log is located in the Tivoli Storage Manager database and contains text messages. The log is pruned automatically every night at midnight. The default retention period for the log is only one day. This does not provide an adequate amount of information for production systems. We recommend that you specify a value of between seven and 14 days.

The administrative command to collect accounting information, keep server activity log records for 10 days, and the results of those commands should look like this:

```
tsm: TSM010> set accounting on
ANR2091I Accounting parameter set to ON.

tsm: TSM010> set actlogretention 10
ANR2090I Activity log retention period set to 10 days.
```

3.4 Server implementation

In this section, we explain how to load the Tivoli Storage Manager supplied sample scripts, load the redbook environment scripts, format server volumes, label tape volumes, and use the Windows NT install wizards.

3.4.1 Load sample scripts

Tivoli Storage Manager is shipped with sample command scripts that can be loaded into the database and run from a Tivoli Storage Manager administrative client, administrative Web interface, or server console. They can also be included in Tivoli Storage Manager administrative command schedules. The sample scripts, in `scripts.smp`, are primarily SELECT queries, but also include scripts that define volumes, extend the database and recovery log, and backup storage pools.

You can load the sample scripts into the database by using either the DSMSEV LOAD command or the `macro` administrative command. When running the administrative client, it is required that the sample scripts file be accessible from the local machine. The DSMSEV LOAD command can only be used when the server is not running, while the `macro` command can only be used when the server is running. We recommend that you use the `macro` command method.

On non-MVS platforms, the sample scripts file is named `scripts.smp` and is in the install directory for the server. On the MVS platform, the sample scripts file is a member of the Tivoli Storage Manager SAMPLIB data set named ANRSCRIPT.

In the following example, we assume that the administrative client is running in an AIX environment and the sample scripts file has been put into the directory `/usr/tivoli/tsm/server/bin`. You must invoke the administrative command line client with the `ITEMCOMMIT` parameter, otherwise the macro command will fail with a series of error messages. The commands to create the sample scripts and resulting output looks similar to this:

```
root@ /usr/tivoli/tsm/server/webimages > dsmadmc -itemcommit
Tivoli Storage Manager
Command Line Administrative Interface - Version 3, Release 7, Level 1.0
(C) Copyright IBM Corporation, 1990, 1999, All Rights Reserved.

Enter your user id: admin

Enter your password:*****

Session established with server TSM010: AIX-RS/6000
Server Version 3, Release 7, Level 1.0
Server date/time: 02/08/00 16:33:18 Last access: 02/08/00 16:26:12

tsm: TSM010> macro /usr/tivoli/tsm/server/webimages/scripts.smp
NS8000I Server command: 'del script q_locked'
ANR1455E DELETE SCRIPT: Command script Q_LOCKED does not exist.
ANS8001I Return code 11.
ANS8000I Server command: 'del script q_inv_attempt'
ANR1455E DELETE SCRIPT: Command script Q_INV_ATTEMPT does not exist.
ANS8001I Return code 11.
...
...
ANS8000I Server command: 'upd script fsnobackup ' backup_end is null''
ANR1456I UPDATE SCRIPT: Command script FSNOBACKUP updated.
```

3.4.2 Format server volumes

The Tivoli Storage Manager database, recovery log, and storage pools consist of storage volumes. On UNIX server platforms, you have the choice of using raw logical volumes or operating system files for these volumes. On all other platforms, you can only use operating system files. In all cases where operating system files are used for server volumes, a formatting process is required to create and prepare the file for a particular type of usage. We do not discuss using raw logical volumes for server volumes further in this book.

The process of adding a new server volume consists of two steps:

1. Format the volume for the volume type using the `DSMFMT` utility.
2. Define the volume to Tivoli Storage Manager using the appropriate `define` command.

In a non-MVS environment, those two steps can be combined into a single process: you format the volume and immediately define it in Tivoli Storage Manager by using one single storage volume command using the `FORMATSIZE` option. You format and define new database volumes by using the `define dbvolume` command, new recovery log volumes by using the `define logvolume` command, and storage pool volumes by using the `define volume` command. We recommend that you use this method. See sections 4.1, “Defining a new database volume” on page 57, 4.2, “Defining a new recovery log volume” on page 59, and 5.8.5, “Defining storage pool volumes” on page 107 for examples.

In an MVS environment, the only way to format server volumes is to use the DSMFMT utility. This utility runs outside of the Tivoli Storage Manager server environment. Sample JCL to create and format server volumes is in the Tivoli Storage Manager SAMPLIB data set in members named ANRALLO1 and ANRALLO2. The server volume is allocated as a VSAM linear data set. Once the volume has been formatted, it must be defined to Tivoli Storage Manager for use.

3.4.2.1 Define database volume

For database volumes, the allocated file size is always a multiple of 4MB (plus an additional 1MB for overhead). The command to format a 100MB database volume on an AIX system, and the resulting output from that command, should look like this:

```
root@ / > dsmfmt -db /tsm/database/file07.db 101
AIX Server DSMFMT Extent/Volume Formatting Program

Licensed Materials - Property of IBM

5697-TSM (C) Copyright IBM Corporation 1990, 1997. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

Allocated space for /tsm/database/file07.db: 105906176 bytes
```

The command to define the formatted database volume to Tivoli Storage Manager, and the resulting output from that command, should look like this:

```
tsm: TSM010> define dbvol /tsm/database/file07.db
ANR2240I Database volume /tsm/database/file07.db defined.
```

3.4.2.2 Define recovery log volume

For recovery log volumes, the allocated file size is always a multiple of 4MB (plus an additional 1MB for overhead). The command to format a 24MB recovery log volume on an AIX system, and the resulting output from those commands, should look like this:

```
root@ / > dsmfmt -m -log /tsm/log/file02 25
AIX Server DSMFMT Extent/Volume Formatting Program

Licensed Materials - Property of IBM

5697-TSM (C) Copyright IBM Corporation 1990, 1997. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

Allocated space for /tsm/log/file02: 26214400 bytes
```

The command to define the formatted recovery log volume to Tivoli Storage Manager, and the resulting output from that command, should look like this:

```
tsm: TSM010> define logvol /tsm/log/file02
ANR2260I Recovery log volume /tsm/log/file02 defined.
```

3.4.2.3 Define storage pool volume

The command to format a 200MB storage pool volume on an AIX system, and the resulting output from that command, should look like this:

```
root@ / > dsmfmt -m -data /tsm/stgpool/backup1 200
AIX Server DSMFMT Extent/Volume Formatting Program

Licensed Materials - Property of IBM

5697-TSM (C) Copyright IBM Corporation 1990, 1997. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

Allocated space for /tsm/stgpool/backup1: 209715200 bytes
```

The command to define the formatted storage pool volume in the storage pool named *DISKDATA*, and the resulting output from that command, should look like this:

```
tsm: TSM010> define vol diskdata /tsm/stgpool/backup1
ANR2206I Volume /tsm/stgpool/backup1 defined in storage pool DISKDATA (device class DISK).
```

3.4.3 Label tape volumes

Tapes used by the Tivoli Storage Manager server must have labels written at the start of each volume. Labels are internal records that uniquely identify the tape volume. In an MVS environment, the use of standard labelled tape volumes provides the label information. In non-MVS environments, Tivoli Storage Manager must write the labels.

In an MVS environment, any tapes used by Tivoli Storage Manager must have MVS standard labels. Tivoli Storage Manager does not write those labels; they must be generated by some other method. The remainder of this section does not apply to MVS environments.

In non-MVS environments, you provide the unique labels on tape volumes by using either the DSMLABEL utility or the `label libvolume` administrative command. The DSMLABEL utility is run outside the Tivoli Storage Manager environment, while the `label libvolume` command is run inside the Tivoli Storage Manager environment. Both methods work with all tape drives, whether they are stand-alone units or are in libraries. We recommend that you use the administrative command `label libvolume` to label volumes because it writes the labels, does not require removing a drive from Tivoli Storage Manager control, and, in the case of a tape library, registers the volume with the server.

See 12.4.1, “Labeling tapes” on page 252 for examples of using the `label libvolume` command for both manual and automatic libraries.

3.4.4 Windows NT configuration wizards

This section applies to the Windows NT Tivoli Storage Manager server only.

The Tivoli Storage Manager server utilities are available through the *Tivoli Storage Manager* program group. It provides access to various configuration wizards, an administrative interface, a server console, Web sites, and other useful utilities. In particular, the initial configuration wizards provide a structured way to implement a Tivoli Storage Manager environment. Each of the configuration wizards is also available independently within the Server Utilities.

Documentation on the use of the initial configuration wizards is in the chapter titled, *Performing the Initial Configuration*, of the book *Tivoli Storage Manager for Windows NT Quick Start*, GC35-0379.

The wizards provide a good front end to some steps of the Tivoli Storage Manager configuration. They are especially useful in the areas of licensing, services configuration, device configuration, volume formatting and media preparation. The configuration wizards are easier to use than the equivalent administrative interfaces. The wizards hide the details of the command interface and provide good help. However, you still need some Tivoli Storage Manager knowledge to create a good working environment.

The wizards and server utilities do not cover all aspects of an implementation and give potentially unexpected results in other areas. The wizards do not cover definition of items such as administrative schedules, client option sets, or administrators. The node configuration wizard can easily generate multiple domains, the names of which you have no control in specifying.

The wizards are not suited to bulk entry of definitions. This task is better handled through the use of administrative macros or scripts. A good example of this is the definition of our redbook administrative schedules. The definitions require one administrative command or dozens of window interactions. Our recommendation is to use the wizards with care.

Chapter 4. Database and recovery log

This chapter covers the steps you must run to set up and change your database and recovery log settings to a recommended configuration.

The database and the recovery log are closely related. The Tivoli Storage Manager database contains information needed for server operations and information about client data that has been backed up, archived, and space-managed. The database contains pointers to the locations of all backup versions and client files in the Tivoli Storage Manager storage pools. The client data itself is stored in storage pools, not in the database.

We assume that you have planned your database and recovery log sizes. See 2.4, “Disk considerations and sizing” on page 28 for planning considerations. If you have not planned the values that your installation requires to run Tivoli Storage Manager, we strongly recommend that you do so before continuing.

Use the information you gathered from the planning chapter to issue your own commands. Because the database is highly dependent on site-specific values, we cannot predict all possible combinations. Therefore, we came up with a sequence of commands, so that all you have to do is change the values for those that best fit your environment.

Platform specific commands:

Some minor changes are required in UNIX and Windows because of the different filesystem naming conventions of those platforms (a filesystem in UNIX is a drive in Windows). Because the examples in this section are based on an AIX system, if you have a Windows Tivoli Storage Manager version, use the Windows directory names instead. For other platforms, see the *Tivoli Storage Manager Administrator's Reference* for that server.

4.1 Defining a new database volume

Tivoli Storage Manager tracks all volumes defined to the database as one logical object. For performance reasons, you can split the database volumes into different disks, so that the actual read and write operations are balanced across the disks available. For example, Figure 3 on page 58 shows a database that consists of four volumes (`file01` through `file04`). The volumes could exist on the same disk, or be spread across any number of actual disk drives. Although the volumes are separate files from an operating system point of view, Tivoli Storage Manager tracks the database as a single logical image.

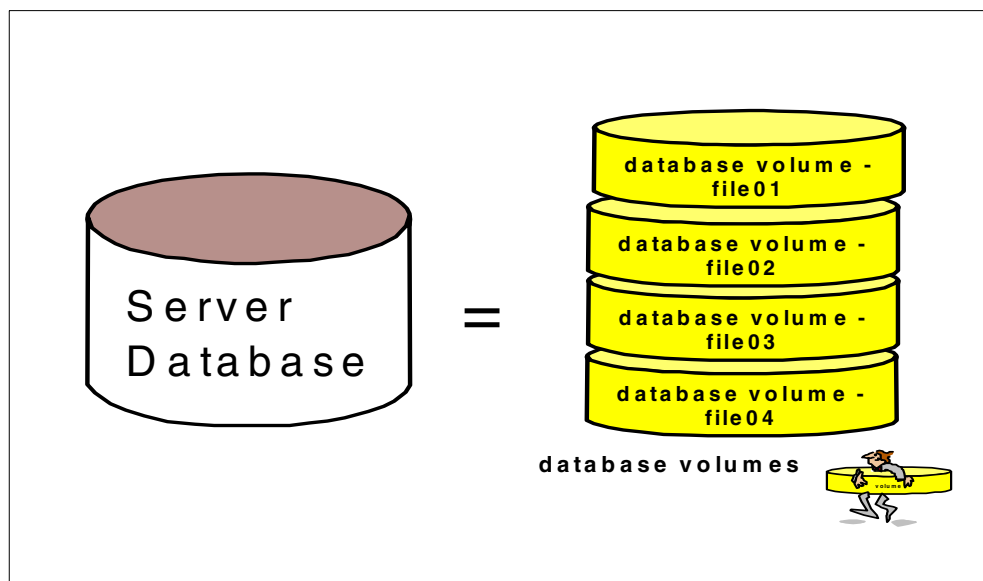


Figure 3. Tivoli Storage Manager database and its volumes

All database volumes must reside on random access media. You can have Tivoli Storage Manager create the volume before it is assigned.

There are two methods of successfully allocating a new database volume. The most common and easiest to use is to allocate a new Tivoli Storage Manager database volume by using the `define dbvolume` command using the `FORMATSIZE` option from either the Tivoli Storage Manager console or an administrative client. We recommend that you use this one-step method. The second method is to use the `DSMFMT` utility command to allocate the volume file and then the `define dbvolume` command to define the database volume to Tivoli Storage Manager. For further information about the `DSMFMT` utility, see 3.4.2, “Format server volumes” on page 53.

The following command allocates space and defines a new database volume in a single operation. It creates a volume `/tsm/database/primary/file01` of size 176 MB. However, we have to specify 177 MB, since the database volume requires 1 MB extra space for overhead.

```

tsm: TSM010> define dbvolume /tsm/database/primary/file01 formssize=177 wait=yes

ANR0984I Process 9 for DEFINE DBVOLUME started in the FOREGROUND at 17:29:14.
ANR2240I Database volume /tsm/database/primary/file01 defined.
ANR0986I Process 9 for DEFINE DBVOLUME running in the FOREGROUND processed 1
items for a total of 185,597,952 bytes with a completion state of SUCCESS at
17:29:27.

tsm: TSM010> query db

Available Assigned Maximum Maximum Page Total Used Pct Max.
Space Capacity Extension Reduction Size Usable Pct Util Pct
(MB) (MB) (MB) (MB) (bytes) Pages Pages Util Pct
-----
192 4 188 0 4,096 1,024 305 29.8 29.8

tsm: TSM010> query dbvolume

Volume Name Copy Volume Name Copy Volume Name Copy
(Copy 1) Status (Copy 2) Status (Copy 3) Status
-----
/tsm/database/p- Sync'd Undefined Undefined
rimary/file01
/usr/tivoli/tsm- Sync'd Undefined Undefined
/server/bin/db-
.dsm

```

UNIX Users:

The database name is case-sensitive. If you want to create a database volume named `file01`, you must type the name in lower case.

4.2 Defining a new recovery log volume

The recovery log contains information about updates that have not yet been committed to the database. For example, when a client is running backups, all transactional data that controls this operation is first written to the recovery log.

Tivoli Storage Manager tracks all volumes defined to the recovery log as one logical object. For performance reasons, you can split the recovery log volumes into different disks, so that the actual read and write operations are balanced across the disks available. For example, Figure 4 on page 60 shows a recovery log that consists of two volumes (`file01` and `file02`). The volumes could exist on the same disk, or be spread across any number of actual disk drives. Although the volumes are separate files from an operating system point of view, Tivoli Storage Manager tracks the recovery log as a single logical image.

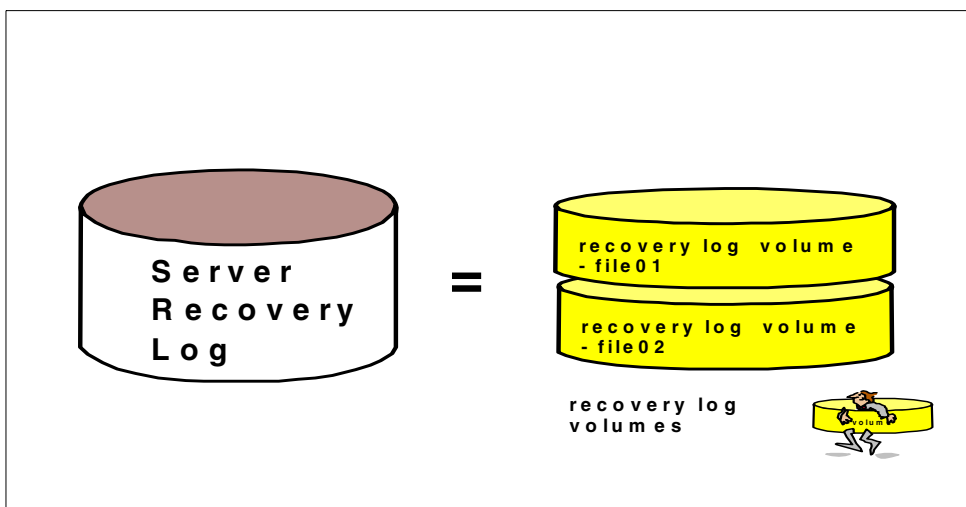


Figure 4. Tivoli Storage Manager recovery log and its volumes

All recovery log volumes must reside on random access media. You can have Tivoli Storage Manager create the volume before it is assigned.

There are two methods of successfully allocating a new recovery log volume. The most common and easiest to use is to allocate a new Tivoli Storage Manager recovery log volume by using the `define logvolume` command with the `FORMATSIZE` option from either the Tivoli Storage Manager console or an administrative client. We recommend that you do this one-step method. The second method is to use the `DSMFMT` utility command to allocate the volume file and then the `define logvolume` command to define the recovery log volume to Tivoli Storage Manager. For further information about the `DSMFMT` utility, see Chapter 3.4.2, “Format server volumes” on page 53.

The following command allocates space and defines a new recovery log volume in a single operation. It creates a volume `/tsm/log/primary/file01` of size 36 MB. However, we have to specify 37 MB, since the recovery log volume requires 1 MB extra space for overhead.


```

tsm: TSM010> define logvol /tsm/log/primary/file01 formssize=37 wait=yes
ANR0984I Process 10 for DEFINE LOGVOLUME started in the FOREGROUND at 17:34:02.
ANR2260I Recovery log volume /tsm/log/primary/file01 defined.
ANR0986I Process 10 for DEFINE LOGVOLUME running in the FOREGROUND processed 1
items for a total of 38,797,312 bytes with a completion state of SUCCESS at
17:34:05.

```

```

tsm: TSM010> query log

```

Available Space (MB)	Assigned Capacity (MB)	Maximum Extension (MB)	Maximum Reduction (MB)	Page Size (bytes)	Total Usable Pages	Used Pages	Pct Util	Max. Pct Util
44	8	36	0	4,096	1,536	322	21.0	22.0

```

tsm: TSM010> query logvol

```

Volume Name (Copy 1)	Copy Status	Volume Name (Copy 2)	Copy Status	Volume Name (Copy 3)	Copy Status
/tsm/log/primary/file01	Sync'd		Undefined		Undefined
/usr/tivoli/tsm-server/bin/log.dsm	Sync'd		Undefined		Undefined

UNIX Users:

The recovery log name is case-sensitive. If you want to create a recovery log volume named `file01`, you must type the name in lower case.

4.3 Setting the log mode to ROLLFORWARD

Use the `set logmode` command to set the mode for saving recovery log records. The log mode determines how long Tivoli Storage Manager saves records in the recovery log and the kind of database recovery you can use. The two log modes are NORMAL and ROLLFORWARD.

- **NORMAL:** Tivoli Storage Manager saves only those records needed to restore the database to the point of the last backup (point-in-time recovery). Tivoli Storage Manager deletes any unnecessary records from the recovery log. Changes made to the database since the last backup cannot be recovered. Any database backups created by issuing the `backup db` command can only be used to perform point-in-time recovery. In NORMAL log mode, you may need less space for the recovery log, because Tivoli Storage Manager does not keep records already committed to the database.
- **ROLLFORWARD:** Tivoli Storage Manager saves all recovery log records that contain changes made to the database since the last time it was backed up. Tivoli Storage Manager deletes recovery log records only after a successful database backup. The recovery log records can be used to restore a database to its most current state (roll-forward recovery) after loading the most current database backup series. A database backup series created in ROLLFORWARD mode can be used for either point-in-time recovery or roll-forward recovery. We recommend that you enable ROLLFORWARD log mode if your site requires a high level of availability to the Tivoli Storage Manager server. ROLLFORWARD log mode may require a significant amount of space to record all activity.

To configure the log mode for ROLLFORWARD, issue the `set logmode` command:

```
tsm: TSM010> set logmode rollforward
ANR2294I Log mode set to ROLLFORWARD.
```

You can check that the command was successful by using the `query status` command and checking the Log Mode field:

```
tsm: TSM010> query status
Storage Management Server for AIX-RS/6000 - Version 3, Release 7, Level 1.0

Server Name: TSM010
Server host name or IP address:
Server TCP/IP port number:
Server URL:
Crossdefine: Off
Server Password Set: No
Server Installation Date/Time: 01/24/00 16:44:02
Server Restart Date/Time: 02/08/00 14:23:42
Authentication: On
Password Expiration Period: 90 Day(s)
Invalid Sign-on Attempt Limit: 3
Minimum Password Length: 6
WEB Admin Authentication Time-out (minutes): 0
Registration: Closed
Availability: Enabled
Accounting: On
Activity Log Retention Period: 10 Day(s)
Activity Summary Retention Period: 30 Day(s)
License Audit Period: 30 Day(s)
Last License Audit: 01/24/00 16:44:17
Server License Compliance: Valid
Central Scheduler: Active
Maximum Sessions: 25
Maximum Scheduled Sessions: 12
Event Record Retention Period: 10 Day(s)
Client Action Duration: 5 Day(s)
Schedule Randomization Percentage: 25
Query Schedule Period: Client's Choice
Maximum Command Retries: Client's Choice
Retry Period: Client's Choice
Scheduling Modes: Any
Log Mode: RollForward
Database Backup Trigger: Not Defined
BufPoolSize: 2,048 K
Active Receivers: CONSOLE ACTLOG
Configuration manager?: Off
Refresh interval: 60
Last refresh date/time:
```

Note that if you are currently not using the Tivoli Storage Manager database in roll-forward mode, but wish to change this, the recovery log pool size will need to be increased. To estimate the new value, reset the cumulative consumption value using the administrative command `reset logconsumption` and then monitor the cumulative consumption over a number of days. Divide the cumulative consumption by the number of days since you reset the value, to get a representative value. A safe size for the log pool should be around 30-40% larger than this figure.

4.4 Defining database backup trigger

Use the `define dbbackuptrigger` command to define settings for the database backup trigger. The database backup trigger determines when Tivoli Storage Manager automatically runs a full or incremental backup of the Tivoli Storage Manager database and deletes any unnecessary recovery log records.

Tivoli Storage Manager uses the settings you specify with this command only when the log mode is set to `ROLLFORWARD`, which you previously configured with the `set logmode` command. With the `define dbbackuptrigger` command, you specify the percentage of the assigned capacity of the recovery log that can be used before Tivoli Storage Manager begins a backup of the database. The actual percentage that you should choose is highly dependent on the planning considerations in Chapter 2, “Tivoli Storage Manager implementation planning” on page 11. We recommend that you use the 75% percentage as a starting point and use the default number of incrementals (six). You must monitor your environment’s activity to make sure that you do not trigger database backups too frequently. If this happens, you should consider increasing the recovery log size.

To set the limit of 75% for the recovery log to start a `backup db` and run up to six incremental database backups before a full database backup, issue the `define dbbackuptrigger` command:

```
tsm: TSM010> define dbbackuptrigger devclass=3570 logfullpct=75 numincr=6
ANR2282I Database backup trigger defined and enabled.
```

You can check that the command was successful by issuing the `query dbbackuptrigger` command:

```
tsm: TSM010> query dbbackuptrigger format=detail
      Full Device Class: 3570
      Incremental Device Class: 3570
      Log Full Percentage: 75
      Incrementals Between Fulls: 6
      Last Update by (administrator): ADMIN
      Last Update Date/Time: 02/08/00 17:40:53
```

4.5 Setting the expansion trigger

Tivoli Storage Manager lets you fully automate the process of increasing the database and recovery log. For example, assume that you have a 200 MB database and a 100 MB recovery log. You want to increase the size of the database by 25% when 85% is in use, but not to more than 1 GB. You also want to increase the recovery log by 30% when 75% is in use, but not to more than 500 MB.

Note:

Setting a space trigger does not mean that the percentage used in the database and recovery log will always be less than the value specified with the FULLPCT parameter. Tivoli Storage Manager checks utilization when database and recovery log activity results in a “commit”. Deleting database volumes and reducing the database does not cause the trigger to activate. Therefore, the utilization percentage can exceed the set value before new

MVS Users:

Database expansion triggering is not fully supported in Tivoli Storage Manager MVS. In the MVS environment, the `define spacetrigger` command does not cause new volumes to be created. It causes the database and recovery log to be extended if space is available. Event logging allows messages to display the amount of new space that is needed to satisfy the space trigger utilization parameter. You can use these messages to initiate automatic expansion or to complete an allocation job (refer to ANRFMT1 and ANRFMT, sample jobs that are located in the ASAMPLIB library).

4.5.1 Database space trigger

To define a new space trigger for the database in the `/tsm/database/primary` directory (Tivoli Storage Manager generates the volume names), issue the `define spacetrigger db` command. The `expansionprefix` parameter allows you to specify where you want the created volumes to be stored in the filesystem.

```
tsm: TSM010> define spacetrigger db fullpct=85 spaceexpansion=25 \  
cont> expansionprefix=/tsm/database/primary/ maximumsize=1000  
  
ANR2274I Data Base Space trigger defined and enabled.
```

The MAXIMUMSIZE limit for the database is 9,999,999 MB. A value of zero, (0) or omitting this parameter disables checking for maximum size. If later on, you need to change the trigger definition, you can use the `update spacetrigger` command.

Whenever the system detects that the database needs extra space, it triggers the expansion as shown in the `query actlog` command:

```
tsm: TSM010> query actlog  
...  
  
02/08/00 17:58:30 ANR4414I Data Base Space expansion Is Needed, 4 Megabytes needed.  
02/08/00 17:58:30 ANR4412I Data Base and Recovery Log Space expansion triggered.  
02/08/00 17:58:31 ANR2240I Database volume /tsm/database/primary/D0151600.DBV defined.  
02/08/00 17:58:31 ANR2240I Database volume /tsm/database/primary/D0151600.DBV defined.  
02/08/00 17:58:31 ANR2248I Database assigned capacity has been extended.  
02/08/00 17:58:31 ANR4415I Data Base and Recovery Log Space Expansion Completed.  
  
..
```

4.5.2 Recovery log space trigger

To define a new space trigger for the recovery log in the `/tsm/log/primary` directory (Tivoli Storage Manager generates the volume names), issue the following `define spacetrigger log` command. The `expansionprefix` parameter allows you to specify where you want the created volumes to be stored in the filesystem.:

```
tsm: TSM010> define spacetrigger log fullpct=75 spaceexpansion=30 \  
cont> expansionprefix=/tsm/log/primary/ maximumsize=500  
  
ANR2279I Recovery Log Space trigger defined and enabled.
```

You can use a `MAXIMUMSIZE` from 9 MB through 5000 MB (5GB) for the recovery log. A value of zero, (0) or omitting this parameter disables checking for maximum size. If later on, you need to change the trigger definition, you can use the `update spacetrigger` command.

Whenever the system detects that the recovery log needs extra space, it triggers the expansion as shown in the `query actlog` command:

```
tsm: TSM010> query actlog  
...  
02/08/00 18:06:48 ANR4413I Recovery Log Space expansion Is Needed, 4 Megabytes needed.  
02/08/00 18:06:48 ANR4412I Data Base and Recovery Log Space expansiontriggered.  
02/08/00 18:06:50 ANR2260I Recovery log volume /tsm/log/primary/L8002600.LOG defined.  
02/08/00 18:06:50 ANR2260I Recovery log volume /tsm/log/primary/L8002600.LOG defined.  
02/08/00 18:06:50 ANR0984I Process 22 for EXTEND LOG started in the BACKGROUND  
02/08/00 18:06:50 ANR2268I Recovery log assigned capacity has been extended.  
02/08/00 18:06:50 ANR4415I Data Base and Recovery Log Space Expansion Completed.  
..
```

4.5.3 When SPACETRIGGER fails

Although Tivoli Storage Manager can expand either the database or the recovery log space or both if needed, you must pay special attention when using mirroring, because you can only give one directory location for the `expansionprefix` parameter in the `define spacetrigger` command. Therefore, if you are running Tivoli Storage Manager with software mirrored volumes (DB or LOG), you may get undesirable volumes in a single disk (making mirroring ineffective). To correct this, you must reallocate the mirror copies in the correct place.

Here is an example of database triggering, which leads to an allocation of both image and copy to the same location (`/tsm/database/primary`). In this case, the database volume, `/tsm/database/primary/C02372600.DBV`, must be reallocated to another filesystem, so that mirroring is still adequate:

```
02/08/00 18:18:51 ANR4414I Data Base Space expansion Is Needed, 4 Megabytes needed.  
02/08/00 18:18:51 ANR4412I Data Base and Recovery Log Space expansion triggered.  
02/08/00 18:18:52 ANR2240I Database volume /tsm/database/primary/D1372600.DBV defined.  
02/08/00 18:18:52 ANR2240I Database volume /tsm/database/primary/D1372600.DBV defined.  
02/08/00 18:18:53 ANR2241I Database volume copy /tsm/database/primary/C2372600.DBV defined.  
02/08/00 18:18:54 ANR2248I Database assigned capacity has been extended.  
02/08/00 18:18:54 ANR4415I Data Base and Recovery Log Space Expansion Completed.
```

4.6 Mirroring

The following scenario shows the importance of mirroring in the recovery process. As the result of a sudden power outage, a partial page write occurs. The recovery log is now corrupted and not completely readable. Without mirroring, transaction recovery operations cannot complete when the server is restarted and the database would need to be restored from its most recent backup. However, if the recovery log is mirrored and a partial write is detected, a mirror volume can be used to construct valid images of the missing pages.

Although you can mirror either the database or the recovery log, we recommend that you mirror both. This gives better availability should you need to recover from a failing disk, as you can see in Figure 5.

In this example, the Tivoli Storage Manager database volumes are split into two different disks, so if a system failure occurs on one, Tivoli Storage Manager can still function. The same thing happens for the recovery log, with each volume duplicated. Note that Tivoli Storage Manager still has one database object and one recovery log object (with a mirrored copy).

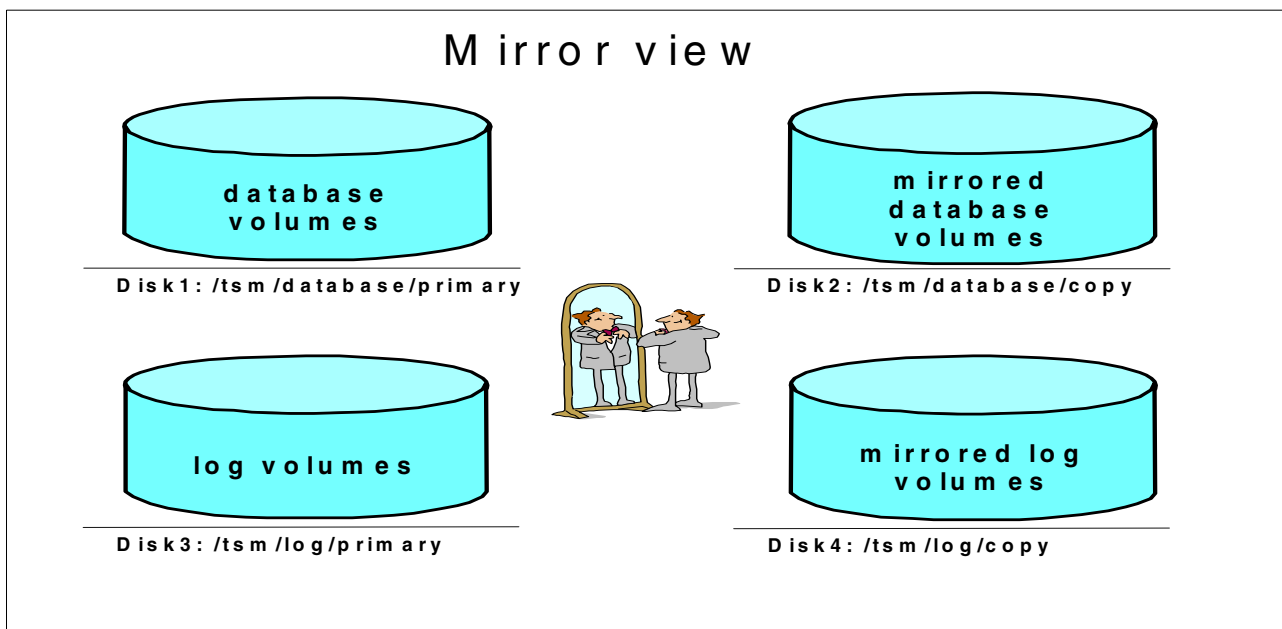


Figure 5. Tivoli Storage Manager mirror view

Although we have used the directories primary and copy for our mirrored volumes, this nomenclature is irrelevant to Tivoli Storage Manager. That is, all copies of a volume are equal and the server can continue operating in the event of failure of any member of the mirrored set.

Tivoli Storage Manager mirrored volumes must have at least the same capacity as the original volumes. This means that if your Tivoli Storage Manager database is made of two volumes of 100 MB each, then you will need at least two extra 100 MB allocated volumes for the mirrored volumes. If you create volumes larger than necessary, Tivoli Storage Manager gives you a warning message (ANR2253W for

database and ANR2273W for the log), but it still allows you to use the allocated volume (although the extra capacity will be wasted).

You must separate each of the mirrored volumes onto separate disks so that you do not lose data if the disk fails. We recommend that you use four disks for the database and recovery log volumes to keep each set of volumes on a separate disk. For example, on UNIX, you have all primary database volumes on one disk, /tsm/database/primary, and their mirrored volumes on another disk, /tsm/database/copy, and the similar distribution for the recovery log volumes and their mirrored copies. Table 28 shows a sample database and recovery log allocation for UNIX and Windows, respectively.

Table 28. Database and recovery log volumes allocation

	UNIX location	Windows location
Database volumes	/tsm/database/primary	D:\tsm\database
Mirrored database volumes	/tsm/database/copy	E:\tsm\database
Log volumes	/tsm/log/primary	F:\tsm\log
Mirrored log volumes	/tsm/log/copy	G:\tsm\log

AIX Users:

We recommend that you use Tivoli Storage Manager mirroring rather than AIX mirroring. If you use AIX mirroring, you may have a problem with raw volumes, but not with Journaled File System (JFS) files. AIX tracks mirroring activity by writing control information to the first 512 bytes of the USER area in a raw volume. This is not a problem for database and recovery log volumes, but Tivoli Storage Manager control information is also written in this area. If AIX overwrites Tivoli Storage Manager control information when raw volumes are mirrored, Tivoli Storage Manager may not be able to vary the volume online.

4.6.1 Database mirroring

When you first install Tivoli Storage Manager, it does not mirror the database. In our example, the query dbvolume command shows the default database (/usr/tivoli/tsm/server/bin/db.dsm) and one additional volume (/tsm/database/primary/file01) created in the previous sections without any mirror images in place:

```

tsm: TSM010> query dbvol

Volume Name      Copy      Volume Name      Copy      Volume Name      Copy
(Copy 1)         Status   (Copy 2)         Status   (Copy 3)         Status
-----
/tsm/database/p-  Sync'd   /usr/tivoli/tsm-  Undef-   /usr/tivoli/tsm-  Undef-
rimary/file01    ined    /server/bin/db-   ined     /server/bin/db-   ined
               .dsm

```

To mirror the database, you must first create all of the new database volumes by using the DSMFMT utility. We create the new copy in /tsm/database/copy.

```
root@ /tsm/database > dsmfmt -db /tsm/database/copy/file01 201
AIX Server DSMFMT Extent/Volume Formatting Program
```

Licensed Materials - Property of IBM

5697-TSM (C) Copyright IBM Corporation 1990, 1997. All rights reserved.
 U.S. Government Users Restricted Rights - Use, duplication or disclosure
 restricted by GSA ADP Schedule Contract with IBM Corporation.

Allocated space for /tsm/database/copy/file01: 210763776 bytes

The next step is to define all database volumes in Tivoli Storage Manager, using the `define dbcopy` command:

```
tsm: TSM010> define dbcopy /tsm/database/primary/file01 /tsm/database/copy/file01
ANR2241I Database volume copy /tsm/database/copy/file01 defined.
```

After a database volume copy is defined, Tivoli Storage Manager synchronizes the volume copy with the original volume. This process can range from minutes to hours, depending on the size of the volumes and performance of your system. After synchronization is complete, the volume copies are mirror images of each other.

You can request information about mirrored database volumes by using the `query dbvolume` command. The next screen shows you the database volume synchronized (Sync'd) with its mirrored volume copy:

```
tsm: TSM010> query dbvolume
```

Volume Name (Copy 1)	Copy Status	Volume Name (Copy 2)	Copy Status	Volume Name (Copy 3)	Copy Status
/tsm/database/p- rimary/file01	Sync'd	/tsm/database/c- opy/file01	Sync'd		Undef- ined
/usr/tivoli/tsm- /server/bin/db- .dsm	Sync'd		Undef- ined		Undef- ined

4.6.2 Recovery log mirroring

When you first install Tivoli Storage Manager, it does not mirror the recovery log. In our example, the command `query logvolume` shows the default recovery log (`/usr/tivoli/tsm/server/bin/log.dsm`) and one additional volume (`/tsm/log/primary/file01`), created in the previous sections, with no mirror in place:


```
tsm: TSM010> query logvol
```

Volume Name (Copy 1)	Copy Status	Volume Name (Copy 2)	Copy Status	Volume Name (Copy 3)	Copy Status
/tsm/log/primary/ file01	Sync'd		Undef- ined		Undef- ined
/usr/tivoli/tsm- /server/bin/lo- g.dsm	Sync'd		Undef- ined		Undef- ined

To mirror the recovery log, you must first create all of the new recovery log volumes by using the DSMFMT utility. Because we want to mirror the /tsm/log/primary/file01 volume, we allocate a new copy for it under /tsm/log/copy/file01:

```
root@ /tsm/log/primary > dsmfmt -log /tsm/log/copy/file01 101  
AIX Server DSMFMT Extent/Volume Formatting Program
```

Licensed Materials - Property of IBM

5697-TSM (C) Copyright IBM Corporation 1990, 1997. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

Allocated space for /tsm/log/copy/file01: 105906176 bytes

The next step is to define all recovery log volumes in Tivoli Storage Manager, using the `define logcopy` command:

```
tsm: TSM010> define logcopy /tsm/log/primary/file01 /tsm/log/copy/file01
```

```
ANR2261I Recovery log volume copy /tsm/log/copy/file01 defined.
```

After a volume copy is defined, Tivoli Storage Manager synchronizes the volume copy with the original volume. This process can range from minutes to hours, depending on the size of the volumes and performance of your system. After synchronization is complete, the volume copies are mirror images of each other.

You can request information about mirrored recovery log volumes by using the `query logvolume` command. For example, the next screen shows a recovery log with its mirrored images. The recovery log volume is synchronized (Sync'd) with its mirrored volume copy:

```
tsm: TSM010> query logvol
```

Volume Name (Copy 1)	Copy Status	Volume Name (Copy 2)	Copy Status	Volume Name (Copy 3)	Copy Status
/tsm/log/primary/ file01	Sync'd	/tsm/log/copy/f- ile01	Sync'd		Undef- ined
/usr/tivoli/tsm- /server/bin/lo- g.dsm	Sync'd		Undef- ined		Undef- ined

Because our recommended configuration removes the default recovery log volume (`/usr/tivoli/tsm/server/bin/log.dsm`), there is no need to mirror it. For further details, see 4.7.2, “Removing the default recovery log volume” on page 71.

4.7 Removing default volumes

During installation, default Tivoli Storage Manager database and recovery log volumes are created, usually in the server installation directory (`/usr/tivoli/tsm/server/bin` on AIX; or `C:\Program Files\Tivoli\tsm\server` on NT). We recommend that you remove these initial database and recovery log volumes so that you have complete control over the new database and recovery log distribution.

Be sure that you have defined a new database and recovery log volumes as explained in the previous sections before continuing. Otherwise, you can damage your Tivoli Storage Manager installation and it will not start.

4.7.1 Removing the default database volume

To remove the default database volume, you must first reduce its size by using the `reduce db` command, so that there is enough spare space in the database logical volume. Remember that the Tivoli Storage Manager database is a logical object consisting of one or more volumes. Use the `reduce db` command to reduce the amount of space that can be used by the database previously allocated to Tivoli Storage Manager:

```
tsm: TSM010> reduce db 160
ANR2248I Database assigned capacity has been reduced.
```

Note:

You can only delete a database volume if there is enough space for the data movement. Therefore, the value in the maximum extension (displayed in the `query database` command) must be at least the size of the volume you want to remove. Otherwise, you will get an error message (ANR2434E).

You can use the `query dbvolume` command to see the full name of the file. You can now delete the default database volume by using the `delete dbvolume` command:

```

tsm: TSM010> query dbvolume

Volume Name      Copy      Volume Name      Copy      Volume Name      Copy
(Copy 1)         Status    (Copy 2)         Status    (Copy 3)         Status
-----
/tsm/database/p- /usr/tivoli/tsm- /usr/tivoli/tsm- /usr/tivoli/tsm- /usr/tivoli/tsm- /usr/tivoli/tsm-
rimary/file01   Sync'd    rimary/file01   Sync'd    rimary/file01   Sync'd
/usr/tivoli/tsm- /usr/tivoli/tsm- /usr/tivoli/tsm- /usr/tivoli/tsm- /usr/tivoli/tsm- /usr/tivoli/tsm-
/server/bin/db- /server/bin/db- /server/bin/db- /server/bin/db- /server/bin/db- /server/bin/db-
.dsm            .dsm            .dsm            .dsm            .dsm            .dsm

tsm: TSM010> delete dbvol /usr/tivoli/tsm/server/bin/db.dsm
ANR2244I Delete process initiated for database volume
/usr/tivoli/tsm/server/bin/db.dsm (process id 32).
ANS8003I Process number 32 started.

```

When you ask Tivoli Storage Manager to delete a database volume, it moves all valid database data that is still on the volume to one of the empty database volumes (in our example, `/tsm/database/primary/file01`).

The next step is to physically erase the file from the operating system. If you do not do this, you will be keeping unused space in your system. This example assumes the use of the default AIX installation directory (`/usr/tivoli/tsm/server/bin`). In this case, use the UNIX `remove` command to delete the unused file:

```

root@ /usr/tivoli/tsm/server/bin > rm db.dsm

```

4.7.2 Removing the default recovery log volume

To remove the default recovery log volume, you must first reduce its size by using the `reduce log` command, so that there is enough room in the recovery log logical volume. Remember that the Tivoli Storage Manager recovery log is a logical object made of one or more volumes. Use the `reduce log` command to reduce the amount of space that can be used by the recovery log previously allocated to Tivoli Storage Manager:

```

tsm: TSM010> reduce log 36

ANR2270I Recovery log assigned capacity has been reduced.

```

Note: You can only delete a recovery log volume if there is enough space for the data movement. Therefore, the value in the maximum extension (displayed in the `query log` command) must be at least the size of the volume you want to remove. Otherwise, you will get an error message (ANR2445E).

You can use the `query logvolume` command to see the full name for the file. You may now delete the default recovery log volume by using the `delete logvolume` command:

```
tsm: TSM010> query logvolume
```

Volume Name (Copy 1)	Copy Status	Volume Name (Copy 2)	Copy Status	Volume Name (Copy 3)	Copy Status
/tsm/log/primary/file01	Sync'd	/tsm/log/copy/file01	Sync'd		Undefined
/usr/tivoli/tsm/server/bin/log.dsm	Sync'd		Undefined		Undefined

```
tsm: TSM010> delete logvolume /usr/tivoli/tsm/server/bin/log.dsm
ANR2264I Delete process initiated for recovery log volume
/usr/tivoli/tsm/server/bin/log.dsm (process id 34).
ANS8003I Process number 34 started.
```

When you ask Tivoli Storage Manager to delete a recovery log volume, it moves all valid recovery log data that is still on the volume to one of the empty log volumes (in our example, `/tsm/log/primary/file01`).

The next step is to physically erase the file from the operating system. If you do not do this, you will be keeping unused space in your system. This example assumes the use of the default AIX installation directory (`/usr/tivoli/tsm/server/bin`). In this case, use the UNIX remove command to delete the unused file:

```
root@ /usr/tivoli/tsm/server/bin > rm log.dsm
```

4.8 Database backup

It is important to run regular Tivoli Storage Manager database backups. If the database becomes damaged or lost, you can only restore it by using the `DSMSERV RESTORE DB` command to perform Tivoli Storage Manager recovery.

If you are following the steps in this chapter, the `DBBACKUPTRIGGER` will already start database backups as needed. It is a good idea to run a `FULL` database backup after those database configuration changes we applied in this chapter though, because the trigger may not start in the next few hours. You will also backup the database at least once every night as a scheduled operation (see Chapter 11, “Scheduling” on page 215 for more information). The triggered database backups are mainly regarded as a ‘safety valve’ used in case of unusual activity on the server.

Use the `backup db` command to back up a Tivoli Storage Manager database to sequential access storage volumes. You can use this command to run one of the following types of backup:

- Full backup (`TYPE=FULL`): Copies the entire Tivoli Storage Manager database.
- Incremental backup (`TYPE=INCREMENTAL`): Copies only those database pages that have been added or changed since the last time the database was backed up.
- DBSnapshot (`TYPE=DBSNAPSHOT`): Specifies that you want to run a full snapshot database backup. The entire contents of the database are copied

and a new snapshot database backup is created without interrupting the existing full and incremental backup series for the database.

To start a FULL backup:

```
tsm: TSM010> backup db type=full devclass=3570 wait=yes  
  
ANR0984I Process 35 for DATABASE BACKUP started in the FOREGROUND at 19:21:53.  
ANR2280I Full database backup started as process 35.  
ANR4554I Backed up 512 of 979 database pages.  
ANR4550I Full database backup (process 35) complete, 979 pages copied.  
ANR0985I Process 35 for DATABASE BACKUP running in the FOREGROUND completed  
with completion state SUCCESS at 19:22:14.
```

To start an out-of-band, database snapshot backup:

```
tsm: TSM010> backup db type=dbsnapshot devclass=c3570 wait=yes  
  
ANR0984I Process 39 for DATABASE BACKUP started in the FOREGROUND at 13:30:55.  
ANR2287I Snapshot database backup started as process 39.  
ANR4554I Backed up 512 of 409 database pages.  
ANR4554I Backed up 1024 of 409 database pages.  
ANR4550I Full database backup (process 39) complete, 409 pages copied.  
ANR0985I Process 39 for DATABASE BACKUP running in the FOREGROUND completed  
with completion state SUCCESS at 13:31:17.
```

Note:

The `devclass` parameter may have a different value for your installation. You can run the `query devclass format=detail` command to check the available device classes and `query library` to check the corresponding libraries available for use. You can run the `backup db` command at any time, without having to stop the server.

Note:

We recommend that you send the most up-to-date Tivoli Storage Manager database backup offsite, so that you have a recovery position for your Tivoli Storage Manager environment.

Although our Redbook material uses the FULL + Incremental features to handle Tivoli Storage Manager database backups, you should consider also running a DBSNAPSHOT and keep the FULL + Incremental database backups onsite for faster recovery. The DBSNAPSHOT backup, being an out of band backup and not dependent on any incremental backups, can be taken offsite. Keep in mind that if you want to change your running environment from FULL + Incremental to couple with the new DBSnapshot features, we recommend that you revise all relevant automated scripts, so that they reflect the new command options, especially if you use Tivoli Disaster Recovery functions.

4.9 Additional commands

Although the recommended settings enable database triggering, you may need to occasionally perform some database maintenance operations in Tivoli Storage Manager. This section shows the possible commands to increase or decrease both database and recovery log space. You should only need to run these commands if you want to redesign the volume layouts or if you add new volumes to Tivoli Storage Manager.

4.9.1 EXTEND DB

Use the `extend db` command to increase the amount of space that can be used by the database within all the database volumes previously allocated to Tivoli Storage Manager.

```
tsm: TSM010> query db
```

Available Space (MB)	Assigned Capacity (MB)	Maximum Extension (MB)	Maximum Reduction (MB)	Page Size (bytes)	Total Usable Pages	Used Pages	Pct Util	Max. Pct Util
20	8	12	0	4,096	2,048	763	37.3	50.3

```
tsm: TSM010> extend db 8
```

ANR2248I Database assigned capacity has been extended.

```
tsm: TSM010> query db
```

Available Space (MB)	Assigned Capacity (MB)	Maximum Extension (MB)	Maximum Reduction (MB)	Page Size (bytes)	Total Usable Pages	Used Pages	Pct Util	Max. Pct Util
20	16	4	8	4,096	4,096	765	18.7	25.2

4.9.2 EXTEND LOG

Use the `extend log` command to increase the amount of space that can be used by the recovery log within all the recovery log volumes previously allocated to Tivoli Storage Manager.

```

tsm: TSM010> query log

Available Space Capacity Maximum Extension Maximum Reduction Page Size Total Usable Pages Used Pages Pct Util Max. Pct Util
(MB) (MB) (MB) (MB) (bytes) Pages
-----
36 8 28 4 4,096 1,536 75 4.9 21.3

tsm: TSM010>extend log 8

ANR2269I Recovery log extension process initiated (process ID 36).

ANS8003I Process number 36 started.
tsm: TSM010> query log

Available Space Capacity Maximum Extension Maximum Reduction Page Size Total Usable Pages Used Pages Pct Util Max. Pct Util
(MB) (MB) (MB) (MB) (bytes) Pages
-----
36 16 20 12 4,096 3,584 76 2.1 2.1

tsm: TSM010> query actlog
...
02/08/00 19:27:57 ANR0984I Process 36 for EXTEND LOG started in the
BACKGROUND at 19:27:57.
02/08/00 19:27:57 ANR2269I Recovery log extension process initiated (process
ID 36).
02/08/00 19:27:58 ANR0307I Recovery log extend in progress; 4 megabytes of 8
formatted.
02/08/00 19:28:00 ANR0307I Recovery log extend in progress; 8 megabytes of 8
formatted.
02/08/00 19:28:00 ANR2268I Recovery log assigned capacity has been extended.
02/08/00 19:28:00 ANR0988I Process 36 for EXTEND LOG running in the
BACKGROUND processed 8,388,608 bytes with a completion
state of SUCCESS at 19:28:00.

```

4.9.3 REDUCE DB

Use the `reduce db` command to decrease the amount of space that can be used by the database. To reduce the capacity of the database, you must reduce the database in 4 MB increments. If you do not specify the reduction in 4 MB increments, Tivoli Storage Manager rounds the number to the next 4 MB partition. For example, if you reduce the database by 1 MB, Tivoli Storage Manager reduces the capacity of the database by 4 MB.

```

tsm: TSM010> query db

Available Space Capacity Maximum Extension Maximum Reduction Page Size Total Usable Pages Used Pages Pct Util Max. Pct Util
(MB) (MB) (MB) (MB) (bytes) Pages
-----
20 16 4 8 4,096 4,096 765 18.7 25.2

tsm: TSM010> reduce db 4

ANR2250I Database assigned capacity has been reduced.

tsm: TSM010> query db

Available Space Capacity Maximum Extension Maximum Reduction Page Size Total Usable Pages Used Pages Pct Util Max. Pct Util
(MB) (MB) (MB) (MB) (bytes) Pages
-----
20 12 8 4 4,096 3,072 764 24.9 33.6

```

4.9.4 REDUCE LOG

Use the `reduce log` command to decrease the amount of space that can be used by the recovery log. To decrease the capacity of the recovery log, you must reduce the recovery log in 4 MB increments. If you do not specify the reduction in 4 MB increments, Tivoli Storage Manager rounds the number to the next 4 MB partition. For example, if you reduce the recovery log by 1 MB, Tivoli Storage Manager reduces the capacity of the recovery log by 4 MB.

```
tsm: TSM010> query log
```

Available Space (MB)	Assigned Capacity (MB)	Maximum Extension (MB)	Maximum Reduction (MB)	Page Size (bytes)	Total Usable Pages	Used Pages	Pct Util	Max. Pct Util
36	16	20	12	4,096	3,584	76	2.1	2.1

```
tsm: TSM010> reduce log 8
```

ANR2270I Recovery log assigned capacity has been reduced.

```
tsm: TSM010> query log
```

Available Space (MB)	Assigned Capacity (MB)	Maximum Extension (MB)	Maximum Reduction (MB)	Page Size (bytes)	Total Usable Pages	Used Pages	Pct Util	Max. Pct Util
36	8	28	4	4,096	1,536	77	5.0	5.0

4.9.5 DELETE VOLHIST TYPE=DBBACKUP

A common problem reported by new Tivoli Storage Administrators is that they run out of tapes. Upon analysis, they discover many tapes have been used by old server database backups. The problem seems to be compounded by the fact that each database backup requires its own tape - it is not possible to 'stack' multiple backups onto a single tape. You should think of this as a security feature. Database backups are extremely important and you would not want to be in the situation where you had your last seven days of database backups contained on one tape which, at the crucial moment when a restore was required, turned out to be unreadable. That is why Tivoli Storage Manager requires you to use a new tape each time you backup the database, regardless of whether it is a full or incremental backup. But the problem remains, how to free up the obsolete database backup tapes for re-use? The answer is the `delete volhistory` command. Use this command to delete sequential volume history information collected by the server when the information is no longer needed. Depending on the options chosen, this can be for sequential volumes used for database backups, storage pools or other purposes.

When the records that contain volume history information about volumes that are not in storage pools (volume types DBSNAPSHOT, BACKUPFULL, BACKUPINCR, and EXPORT) are deleted, the volumes return to scratch status if they were acquired by Tivoli Storage Manager as scratch volumes. For scratch volumes with device type FILE, the files are deleted.

When the records that contain volume history information about volumes in storage pools are deleted, the volumes themselves are not affected and remain in the Tivoli Storage Manager database.

In 11.2.2, “Defining the volume history schedule” on page 220 we describe setting up a schedule to take care of this automatically.

```
tsm: TSM010> query volhist type=dbbackup

      Date/Time: 02/07/00 16:48:17
      Volume Type: BACKUPFULL
      Backup Series: 1
      Backup Operation: 0
      Volume Seq: 1
      Device Class: 3570
      Volume Name: 080CED
      Volume Location: NOTMOUNTABLE
      Command:

      Date/Time: 02/08/00 17:40:54
      Volume Type: BACKUPFULL
      Backup Series: 2
      Backup Operation: 0
      Volume Seq: 1
      Device Class: 3570
      Volume Name: 081DF3
      Volume Location:
      Command:

tsm: TSM010> del volhist type=dbbackup todate=today

Do you wish to proceed? (Yes/No) y
ANR2467I DELETE VOLHISTORY: 1 sequential volume history entries were
successfully deleted.

tsm: TSM010> query volhist type=dbbackup

      Date/Time: 02/08/00 17:40:54
      Volume Type: BACKUPFULL
      Backup Series: 2
      Backup Operation: 0
      Volume Seq: 1
      Device Class: 3570
      Volume Name: 081DF3
      Volume Location:
      Command:
```

Note:

You should not run the `delete volhist` command to delete obsolete database backups if you are using Tivoli Disaster Recovery Manager, since that product handles the expiry of this data itself.

Chapter 5. Data storage

Tivoli Storage Manager represents data storage with administrator-defined Tivoli Storage Manager objects: storage pools and storage pool volumes.

In the following sections, we discuss how Tivoli Storage Manager can automatically manage these pools and how you can define and check that the command used has performed the changes you require.

We also discuss some Tivoli Storage Manager concepts related to storage pools: how to use fewer tapes (reclamation), and how to reduce a client's restore time (collocation). We provide a brief outline of protecting disk storage by using some form of redundant array of independent disks (RAID).

5.1 Storage pools

A storage pool is a collection of storage pool volumes; each storage pool represents one type of media - see Figure 6. For example, a storage pool for a 4mm, digital audio tape device (DAT), represents collections of only 4mm tapes, and a storage pool for an IBM 3590 represents collections of only 3590 tapes. A storage pool created on a disk has files formatted as volumes under Tivoli Storage Manager which are collectively arranged as a storage pool.

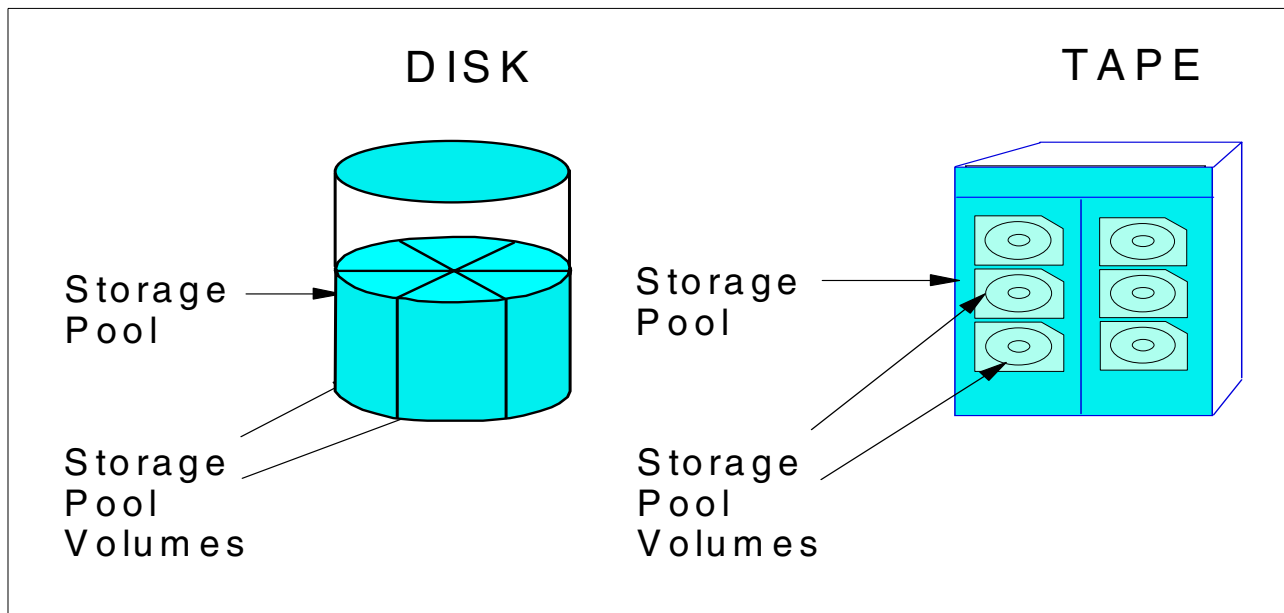


Figure 6. How volumes and storage pools are arranged for disk and tape

You can add or remove volumes without interrupting server operations. In this way you can increase or decrease the size of a storage pool dynamically without any interruption to the Tivoli Storage Manager service.

The two main categories of devices supported for storage pools are random access and sequential devices.

1. Random access devices refer to magnetic disks.
2. Sequential devices refer to tape devices and/or optical devices that require manual configuration.

Storage pools for Tivoli Storage Manager are defined on a wide range of different devices that are attached locally to the server or accessed via a Storage Area Network (SAN).

For a complete up-to-date list of different devices that can be attached to the server platform you are going to use, the best information source is the Tivoli Storage Manager Web site. This can be found at the URL:

http://www.tivoli.com/support/storage_mgr/adsercli.htm

Find your Tivoli Storage Manager server operating system type and click on the link to get to the most current device support list. Current SAN supported devices and configuration can be viewed at:

http://www.tivoli.com/support/storage_mgr/san/overview.html

Tivoli Storage Manager has two types of storage pools:

- Primary storage pools
- Copy storage pools

5.1.1 Primary storage pools

When a client node backs up, archives or migrates data (via hierarchical space management, or HSM), the data is stored in a primary storage pool.

When a user tries to restore, retrieve or export file data, the requested file is obtained from a primary storage pool if possible. Primary storage pool volumes are always located onsite.

A primary storage pool can use random access storage (DISK device class) or sequential access storage (for example, tape, optical or FILE device classes).

5.1.2 Copy storage pools

A copy storage pool provides an additional level of protection for client data. It is created by the administrator backing up a primary storage pool. The copy storage pool contains all versions of all files, exactly as they appear in the primary storage pool.

A copy storage pool provides recovery from partial and complete failures in a primary storage pool. An example of a partial failure is a single tape in a primary storage pool being lost or defective. When a client attempts to restore a file which was on this volume, the server will automatically switch to the right copy storage volume if it is available to seamlessly restore the client's data. If a complete primary storage pool is destroyed, for example in a major disaster, the copy storage pool is used to recreate the primary storage pool.

A copy storage pool can use only sequential access storage (for example, tape optical or FILE device classes).

Copy storage pool volumes can be moved offsite and still be tracked by Tivoli Storage Manager. This is done by using the access mode of OFFSITE to ensure that Tivoli Storage Manager does not request a volume mount. Moving copy storage volumes offsite provides a means of recovering from disaster at your primary site. Tivoli Disaster Recovery Manager helps automate offsite media tracking and storage pool recovery.

Copy pools are not part of the storage migration hierarchy. Files are not migrated to or from copy storage pools. The only way to store files in copy storage pools is by using the `BACKUP STGPOOL` command.

5.2 Storage pool hierarchy

Tivoli Storage Manager allows you to configure storage pools to provide the best combination of performance throughput and data permanence. In most cases, keeping client data on tape or optical media is a requirement. However, taking the backups direct to tape may not give the best performance, especially where there are many clients to backup concurrently, and/or many small files are being backed up. Therefore, Tivoli Storage Manager provides a storage pool hierarchy, whereby a client backs up to storage pool usually on disk. When this storage pool fills up, Tivoli Storage Manager will automatically move or 'migrate' files to the next storage pool in the hierarchy (usually on tape or optical) while the client continues its backup operation. This migration process is controlled by high and low thresholds set on the storage pool.

The management class determines where the client data enters the storage hierarchy. Figure 7 shows a configuration where 5 storage pools are defined. The default management class sends backups first to storage pool A which are then migrated by the server to pool B and finally to pool C. The other management class uses storage pool D for space managed files, whereas backups are sent directly to Tape_Pool. This would be appropriate for large client files (for example, application database files) which can exploit the streaming performance capacity of the tape device.

Storage Pool Assignment

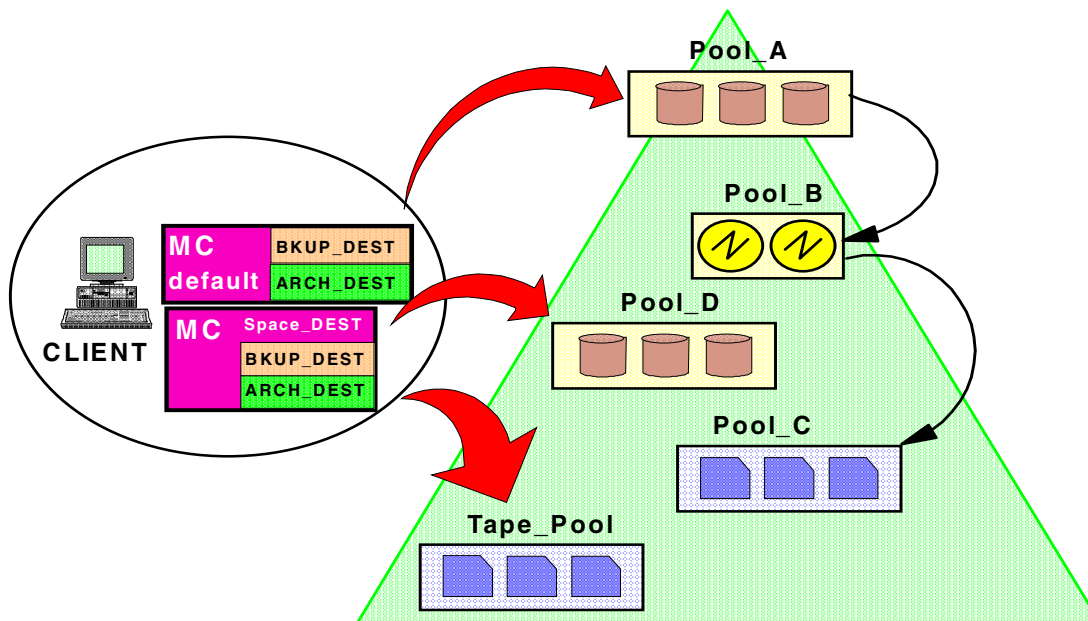


Figure 7. Possible hierarchical arrangement of different storage devices

5.3 Movement of data between storage pools

There are two controls available to help you automatically control the space in the storage pools:

- Migration
- MAXSIZE

5.3.1 Migration

Migration helps you control the amount of free space within a storage pool. You can define a high and a low migration threshold for each storage pool in the storage hierarchy. These thresholds tell Tivoli Storage Manager when to move data from one storage pool to another as illustrated in Figure 8.

When the amount of data in a storage pool reaches the high threshold, Tivoli Storage Manager moves client data to the next storage pool, until the storage pool reaches the low threshold. Tivoli Storage Manager's objective is to clear as much space as quickly as possible. To do this, Tivoli Storage Manager picks the client node that has backed up or migrated the largest single filespace or has archived files that occupy the most space and migrates *all* files from *every* filespace belonging to the client to the next storage pool. If the utilization of the storage pool which was at the high threshold has then dropped below the low threshold, then the migration process terminates, otherwise Tivoli Storage

Manager then picks the next client node that has the largest single filespace or largest amount of archived files and migrates it. Since all the filespace occupied by a client are migrated together, this operation not only clears space quickly but also keeps a client's data close in the hierarchy. This mechanism in turn reduces the number of media mounts which would be required during client data restoration.

For random access storage pools only, you can specify the number of processes that are used for migrating files from this storage pool. This parameter is optional and the default value is 1. You can use a migration process value equal to or less than the number of drives in the next storage pool you are migrating to; these processes are performed in parallel to improve data transfer rates. Be aware that if you have two tape drives in the storage pool you are migrating to, and if you have set this parameter to 2, this will interfere with any other requests for this pool until the migration process is completed. If your disk pool is sufficiently sized and you have disk caching on, or a migration delay value set, then this should present no problem.

You can specify the number of days to delay migration for files in a storage pool, this ensures files stay in a pool for a minimum number of days. Setting this parameter may necessitate a larger space for your disk storage pool, since the setting will probably not be able to be complied with if the storage pool fills up. The benefit of setting this parameter is that it increases the chance of finding files for restore in a disk storage pool rather than on tape, thereby speeding up the restore operation.

The migration process is normally controlled by administrative schedules which set the migration threshold values, as described in 11.2.3, "Defining a migration schedule" on page 221.




Figure 8. Migration between storage pools

5.3.2 MAXSIZE

If you assign a MAXSIZE value to a storage pool, files can be placed in that storage pool provided their size is equal to or less than the assigned MAXSIZE. If compression is used, the uncompressed size of the file is used in the comparison. If a file is too large for the first storage pool it will go straight to the next storage pool defined. This is another way to improve performance by not wasting time writing large files to the first pool, only to trigger an immediate migration when the storage pool is filled.

This parameter is not required. If you do not assign a MAXSIZE, and a file entering the storage pool causes it to exceed its high threshold, migration will occur to move data to the next storage pool.

5.4 Reclamation - how to use fewer tapes

Reclamation is the process which allows partially or completely empty tapes to be recycled back into the scratch pool for reuse.

How do tape volumes become empty? Tivoli Storage Manager keeps only up the maximum defined number of versions of each file. Each time it does an incremental backup for a client node, it checks to see if the number of versions stored of a particular file exceeds the management class setting. For example, suppose the management class specifies to keep 3 versions of a file and that the file changed and was backed up on Monday, Wednesday and Thursday. On

Friday, if the file has changed again, it is backed up, making 4 versions in all. Therefore, the oldest version (from Monday) gets marked for expiration.

Files marked for expiration will be physically deleted when the next server inventory expiration process runs. It is common that a tape volume will have files that will expire on different dates. Therefore as these files reach their expiry date and the expiration process occurs, either manually or as a scheduled operation, there will be empty spaces on the tape volume. Fragmentation may also appear on the volume through file deletion on the client (since VERDELETED in the management class is usually less than VEREXISTS) or filespace deletion. These empty spaces cannot be re-used because it is only possible to write on the end of a sequential volume. Fragmentation is undesirable on tape because it makes restores slower through the requirement to skip empty spaces, and it also increases the total number of tapes required for data storage.

The amount of space that can be reclaimed on a volume increases as files on the volume are deleted. When the percentage of space that can be reclaimed on a volume rises above the value for the reclaim parameter, the volume is reclaimed. Active files on the volume are rewritten to other volumes in the storage pool, making the original volume available for new files. See Figure 9, which shows reclamation on two fragmented volumes to produce a 95% full volume and a free volume.

It is recommended that you specify a value of 50 percent or greater for this parameter so that files stored on two volumes can be combined onto a single output volume.

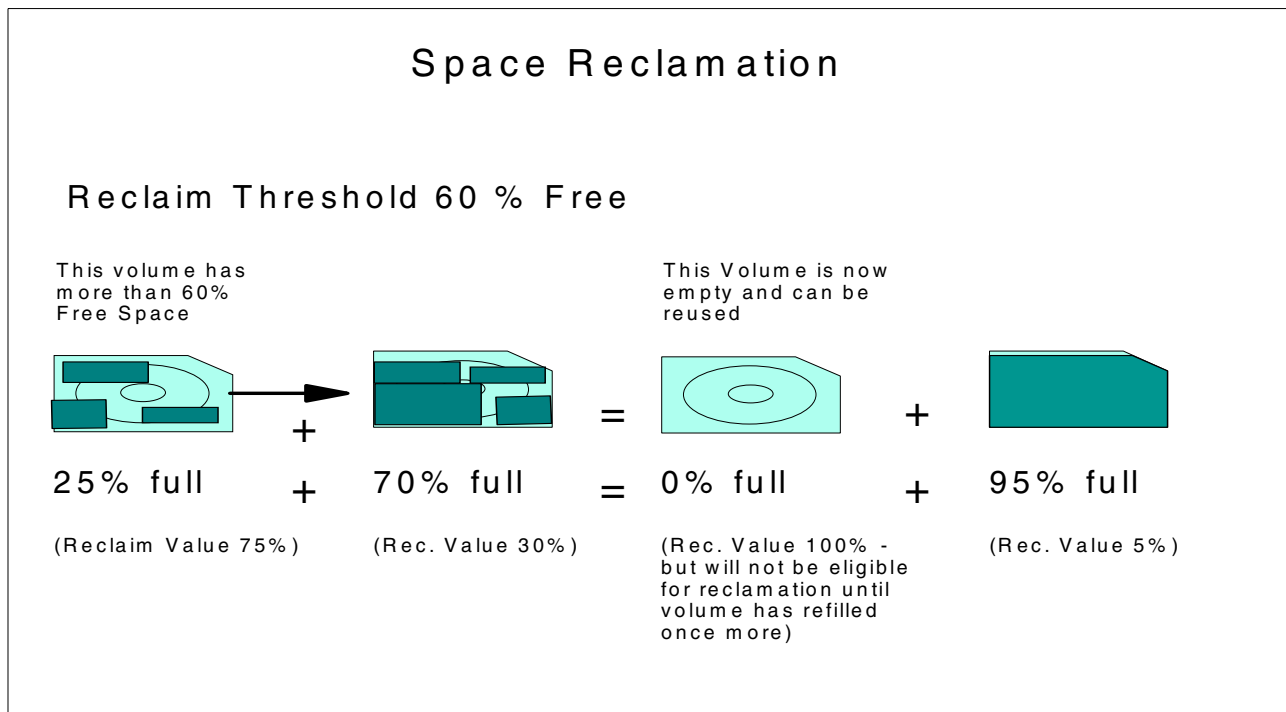


Figure 9. Space reclamation of tape volumes

5.4.1 Single drive reclamation

Reclamation requires two or more drives to work most efficiently. Nevertheless, reclamation can be performed by a single drive in Tivoli Storage Manager by specifying the RECLAIMSTGPOOL parameter. This parameter allows another storage pool to be used as the holding area for the sequential volume being consolidated.

The storage pool specified as the reclaim storage pool must be a primary sequential storage pool on the system. Typically, it is a new primary storage pool created for this purpose. The only disk pool allowed is one whose devtype=file.

When the amount of reclaimable space on a media volume exceeds the reclamation threshold, Tivoli Storage Manager automatically begins the reclamation process. The volume to be reclaimed is mounted in the drive, and the active data is moved to the reclaim storage pool. If the reclaim storage pool is filled, the volume being reclaimed is dismounted, and a new volume in the same tape pool is mounted. The reclaimed data in the reclaim pool is then migrated to that tape volume. Once this process is complete, it repeats until all valid data has been reclaimed from the source volume being reclaimed.

If the reclaim storage pool is not filled before the source volume is emptied, another source volume is mounted, and reclamation continues. This process continues until either the reclaim storage pool is filled or all expired data within the storage pool has been removed.

When defining the reclaim storage pool you must also define the NEXTSTGPOOL parameter pointing back to the pool being reclaimed. Therefore, the reclaimed data can be migrated back to the original storage pool. Refer to Figure 10.

5.4.2 Reclamation of offsite volumes

Care must be taken in reclaiming offsite copy storage volumes.

Tivoli Storage Manager cannot physically move the data from one of these volumes to another, because they are in an offsite vault and therefore not available in the library. Tivoli Storage Manager manages reclamation for an offsite copy pool by obtaining the active files from a primary storage pool or from an onsite volume of a copy pool. These files are then written to a new volume in the copy pool, and the database is updated. A message is then issued that an offsite volume was reclaimed. The new volume will get moved to the offsite location, and the offsite volumes, whose active data it has combined, will be moved back to the scratch pool onsite, when the reuse delay parameter has been satisfied.

However, if a disaster occurs, we could have a problem. If the reclaimed volumes have already been returned to the scratch pool and the new volume has not yet been taken offsite, we have lost any offsite backup. This is the reason for the *delay reuse* period for tapes returning to the scratch pool.

This is achievable by controlling when reclamation will occur by scheduling it properly. Basically, you must ensure that you have performed a backup on all the onsite storage pools to their offsite copy pools, then perform a database backup so the offsite database copy points to all the files in their new locations and then you can perform reclamation. The reuse delay period must be long enough to

ensure the reclaimed volumes do not return to the scratch pool before another database backup is taken.

5.5 Reduce restore times

Storage pool configuration can reduce restore times with two techniques:

- Collocation by minimizing the number of tape volumes used to store a client's data.
- Disk caching by restoring data from disk even if it has already been migrated.

5.5.1 Collocation

Collocation is a process in which the server attempts to keep all files belonging to a client node on the smallest number of sequential access storage volumes. It applies to sequential access storage pools only.

You can choose to have the Tivoli Storage Manager server collocate at the client or filesystem level. To have collocation set at the client level, set the `collocate` option for the storage pool to YES. If this is set, the server attempts to put data for one node on as few volumes as possible. If sufficient volumes are available, then each volume will contain data from only one client. To have collocation set at the filesystem level, set `collocate` to `filesystem`. If this is set, the server attempts to put data for one node and filesystem on one volume. If a node has multiple filesystems, the server attempts to place data for each filesystem on different sequential volumes in the storage pool.

When collocation is set to NO, the server attempts to use all available space on each volume before selecting a new volume. While this process provides better utilization of individual volumes, user files can become scattered across many volumes, so, to restore a client completely may require many volume mounts.

Collocation can only guarantee to keep every client's data on a separate tape if there at least as many tapes defined in the storage pool (or available as scratch tapes) as there are clients who will write to that pool. For example, if there are 5 clients who write data to a collocated storage pool, at least 5 tapes must be available to guarantee strict collocation. In a situation where there were 10 clients writing to a newly defined, collocated storage pool, and the `MAXSCRATCH` value was 5, then 2 clients would use each tape volume.

By using collocation, you reduce the number of volume mount operations required when users restore or retrieve many files from the storage pool. Collocation therefore improves access time for these operations—see Figure 10. The downside is that more tape mounts and volumes will be required for the backup operation.

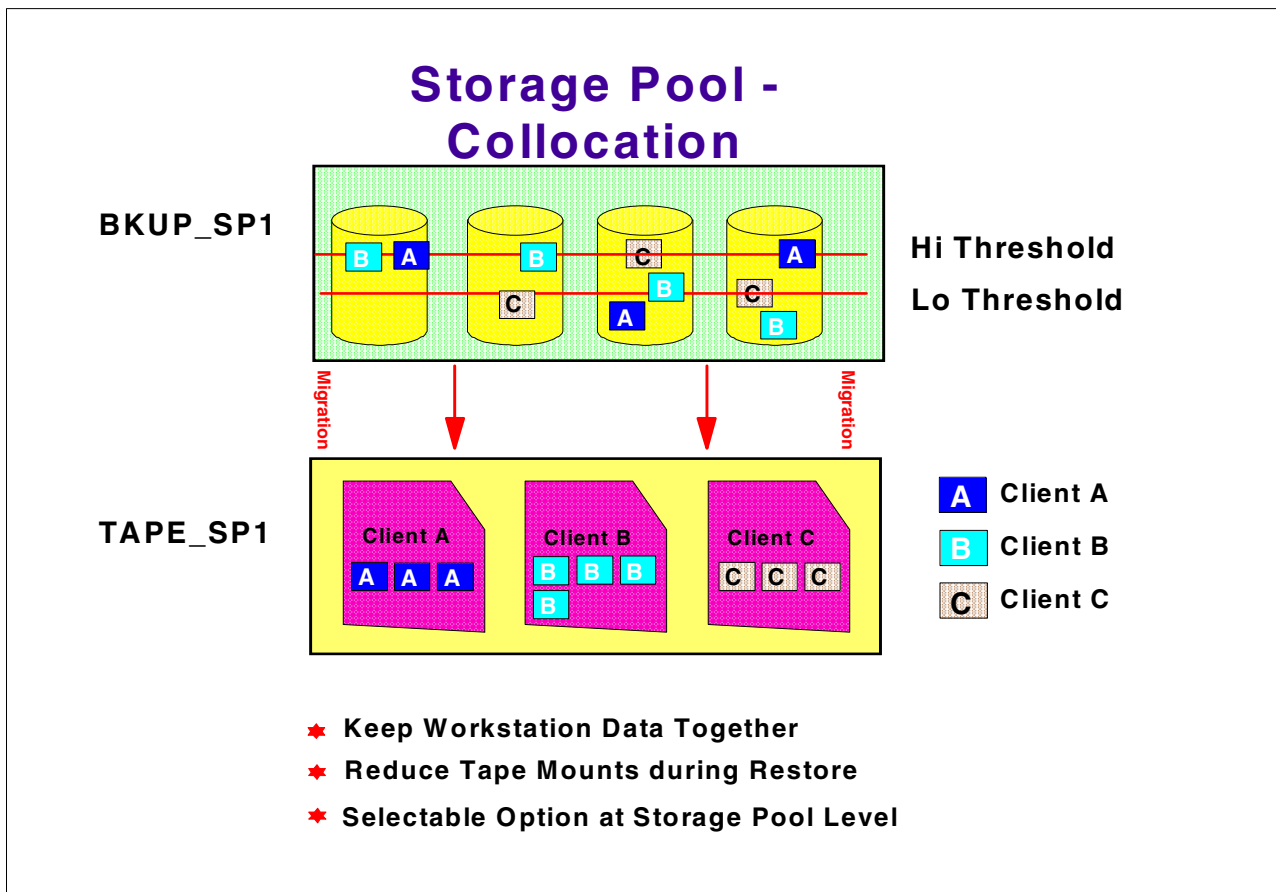


Figure 10. Storage pool collocation on a client level

5.5.2 Disk caching

Disk caching is used for disk storage pools only. When cache is enabled, the migration process leaves behind duplicate copies of the files on disk after the server migrates these files to subordinate storage pools in the storage pool hierarchy. The copies remain in the disk storage pool, but in a cached state, so that subsequent retrieval requests can be satisfied quickly. However, if space is needed to store new data in the disk storage pool, the space occupied by cached files can be immediately reused for the new data. When space is needed, the server reclaims space by writing over the cached files. Files that have the oldest retrieval date and occupy the largest amount of disk space are overwritten first.

When caching is used, the occupancy value that the server compares against the migration thresholds does not include space occupied by cached copies of files. The space utilization of the pool, however, does include the space used by any cached files in the storage pool. If you update a storage pool from `cache=yes` to `cache=no`, the cached files will not disappear immediately. The occupancy value `%util` remains the same. The cache space will be reclaimed over time as the server needs the space, and no additional cached files will be created.

Do not use cache if you have limited database space. When you use cache, more database space is needed because the server has to keep track of both the cached copy of the file and the new copy in the subordinate storage pool.

5.6 RAID

To protect data in a disk type storage device, we recommend that you use some form of redundant array of independent disks (RAID).

RAID can be implemented at either a hardware or software level. Protecting data stored on a disk is a subject in itself, and below we have just touched on some of the possibilities you may want to look into further. The best solution for you will depend on many factors.

A hardware RAID solution can be beneficial to reduce the performance penalty on the server, when implementing RAID. It can also provide disks and controllers that can be exchanged on the fly, should they fail. This solution can also protect against failed controllers, power supply units, and the input power supply itself, and it can provide a spare disk to automatically take over, should a disk fail. You may also consider implementing disk storage that is shared among a number of servers with the ability to allocate and increase storage for each server, as the need arises, without any server downtime.

A software type RAID solution can also provide levels of redundancy with a much lower initial outlay if implemented carefully, using a number of physical disks.

There are also different ways of transmitting data to the physical disks, such as SCSI, SSA and fiber, which all have advantages and disadvantages.

Whichever path you choose, we recommend that you implement one of the three following RAID levels:

- RAID 1 Mirroring
- RAID 1 + 0 Mirror and Stripe
- RAID 5 Distributed Parity

RAID can only provide protection from hardware failure. Application software failure is still possible and is not prevented by a RAID solution. For performance reasons, you should not position storage pools, Tivoli Storage Manager database and recovery log volumes on the same RAID 5 array.

5.7 Setting up storage pools

To define a storage pool initially, you need to judge how large it needs to be. See 2.4, “Disk considerations and sizing” on page 28 and 2.5, “Tape considerations and sizing” on page 36 for planning considerations. If your initial calculations are inaccurate or data to be backed up grows over time, the storage pool can be increased in size easily and dynamically.

When defining a new device to the Tivoli Storage Manager server, the device must first be configured under the operating system. Once we have a device name the operating system understands, we can refer to this, when configuring the device under Tivoli Storage Manager. Physical devices are treated under Tivoli Storage Manager as two types: Random Access Devices and Sequential devices, and these are connected locally to the server.

- Random access devices

Random access devices refer to magnetic disk devices that are used for two main purposes:

- To store the database and recovery log.
- To store client data that has been backed up or archived from client nodes. The client data is stored in disk storage pools.

Tivoli Storage Manager stores data on magnetic disk in two ways:

- Random access volumes.
- Storing data sequentially using a FILE device class.

The physical part of random access devices do not need defining to Tivoli Storage Manager; they can be referred to directly using a drive/path reference. The device type of DISK used is also predefined in Tivoli Storage Manager.

- Sequential storage devices

Sequential storage devices under Tivoli Storage Manager refer to tape libraries, optical devices and write once read many (WORM) devices.

The drive and library must first be defined under the operating system. In the following example on AIX, we define an ADIC VLSDLT library, the drive is set at SCSI ID 6, the library is set at SCSI ID 3. First you have to find out the address of the SCSI card and the attached device by issuing the following command:

```
root@ / > lsdev -Cc tape
mt1 Available 00-00-0s-6,0 Other SCSI Tape Drive
```

Now that you know the address of the SCSI card 00-00-0s and the ID and logical unit number (LUN) of the drive 6,0 you can define the drive and library to AIX using the Tivoli Storage Manager driver.

To define the drive, start the system management interface tool (SMIT) and go to **Devices** → **Tivoli Storage Manager Devices** → **Tape Drive** → **Add a Tape Drive**. Pick the `ADSM-SCSI-MT` device listed. Pick the card with address 00-00-0s. Supply the connection address in the brackets [6,0]

To define the library, in SMIT go to **Devices** → **Tivoli Storage Manager Devices** → **Library** → **Add a Library**. Pick the `ADSM-SCSI-LB` device listed. Pick the card with the address 00-00-0s. Supply the connection address in the brackets [3,0].

Windows NT Users:

On Windows NT, you define drives and libraries under **Settings** → **Control Panel** → **Tape Devices** → **Drivers**.

To configure the device under Tivoli Storage Manager, five components must be configured. These can be thought of in two component classes, physical and logical. These components must be configured in the order of the following five steps:

- Physical
 1. The library: AUTOMATED or MANUAL.
 2. The drive (or drives, if there is more than one in your library).
- Logical
 3. The device class: Here you specify the device type and associate the device class name with the device type and the library name. Here the device type is from a Tivoli Storage Manager defined list, and the command is common to all Tivoli Storage Manager server platforms, for example 3570, 4MM, or DLT.
 4. The storage pool: This specifies the device class to use for the storage pool you are creating and includes specific storage pool rules for handling data.
 5. The volumes: This is where you label the volumes and either define them as scratch volumes or as private volumes belonging to a particular storage pool.

To define storage pools to Tivoli Storage Manager successfully, there is an order that must be followed. If you try to create a Tivoli Storage Manager object that has a parameter referring to another item, that item must already exist, and therefore it should be created first.

Table 29 details which steps you have to complete in order to configure each component for the device and storage pool you wish to set up. It also shows the reference sections in this book for completing these tasks.

Table 29. Guide to setting up storage pools

Order of Definition	Component Class	Component	To Define the Type of Device Below, Follow Each Component Section				
			Manual Library	Automated Library	Disk Storage Pool	Sequential Disk Storage Pool	Copy Storage Pool
1	Physical	Library	5.7.1	5.8.1	N/A	N/A	N/A
2		Drive	5.7.2	5.8.2	N/A	N/A	N/A
3	Logical	Device Class	5.7.3.1	5.8.3	5.7.3.2	5.7.3.3	5.7.3.4
4		Storage Pool	5.7.4.1		5.7.4.2	5.7.4.3	5.7.4.4
5		Volume	5.7.5.1	5.8.5	5.7.5.2	N/A	5.7.5.3

5.7.1 Defining a library — physical

In this section we refer to a manual library. See 5.8.1, “Defining an automated library” on page 101 for an automated library.

First of all, you need to define a library, even if you are using a single tape drive. The following command defines a manual library named M8MM. There is no device name associated with a manual library because the library part can be thought of as virtual.

```
tsm: TSM010> define library m8mm libtype=manual
ANR8400I Library M8MM defined.

tsm: TSM010> query library m8mm

Library Name: M8MM
Library Type: MANUAL
Device:
Private Category:
Scratch Category:
External Manager:
Shared: No
```

The private and scratch categories apply to 3494/3495 type libraries only. The external manager is the path to the location of the external library manager. If you have one, this is where some external media management software identifies the appropriate drive for media access operations. You would need to specify LIBTYPE=EXTERNAL for this parameter to be required.

A library is defined as shared if it is to be shared with another Tivoli Storage Manager server over a storage area network (SAN) or a dual SCSI connection to library drives. When you select this parameter, you must also specify the PRIMARYLIBMANAGER (primary library manager). Setting up SAN shared libraries is discussed in detail in the redpaper *Tivoli Storage Manager SAN Tape Library Sharing*, REDP0024.

5.7.2 Defining a drive — physical

In this section we refer to a manual library. See 5.8.1, “Defining an automated library” on page 101 for an automated library.

After defining the library, you should now define a drive for the library to use.

The drive to use for the manual library must be specified, using its device number that the operating system recognizes it by. The drive therefore must be configured under the operating system.

Windows NT Users:

Run the Tivoli Storage Manager Server Utilities, and under Device Information, check the Tivoli Storage Manager Device Name, for example mt5.0.0.2. Depending on the device type, you need to install additional device drivers. The ADSMSCSI service must already be running and you may need to disable the Windows native device driver first. Check the QuickStart manual for more details.

AIX Users:

Run the following AIX command to find the device name, for example /dev/mt5:

```
root@ / > lsdev -C -c tape
```

Depending on the device type, you may need to install Atape device driver, or you may use the Tivoli Storage Manager inbuilt tape device drivers.

The following command defines a drive named DRV0 in the manual library on AIX named M8MM. The physical tape drive identifier is /dev/mt0 because it is an 8mm tape unit which uses the Tivoli Storage Manager inbuilt tape device drivers. If it was a device which used the IBM device driver known as Atape, for example IBM 3590 or 3570 drives, the identifier would be of the form /dev/rmtx:

```
tsm: TSM010> define drive m8mm drv0 device=/dev/mt0
ANR8404I Drive DRV0 defined in library M8MM.
```

```
tsm: TSM010> query drive
```

Library Name	Drive Name	Device Type	Device	ON LINE
M8MM	DRV0	8MM	/dev/mt0	Yes

Having defined the physical parts of the library to Tivoli Storage Manager, we can now define the logical part to complete the process. There are three steps to this:

- Device Class
- Storage Pool
- Volumes

5.7.3 Defining a device class — logical

A device class can be thought of as a software driver. It is defined so Tivoli Storage Manager can communicate with the physical drive correctly.

5.7.3.1 Manual library

To find the different device types available, look in the Administrators Reference manual or use the online help facility by issuing the `help define devclass` command. This is the online help facility available in Tivoli Storage Manager to give detailed information on every parameter of a command.

The following command defines a device class named C8MM that uses a predefined Tivoli Storage Manager device type of 8MM in the library named M8MM. When a tape is sitting idle in a drive in the library, it remains there for 5 minutes (`MOUNTRETENTION=5`). It uses a tape format of 8500 that gives a capacity of 5GB on a 112m tape.

```
tsm: TSM010> define devclass c8mm devtype=8mm library=m8mm format=8500 estcap=5g \
cont> mountretention=5
ANR2203I Device class C8MM defined.
```

```
tsm: TSM010> query devclass
```

Device Class Name	Device Access Strategy	Storage Pool Count	Device Type	Format	Est/Max Capacity (MB)	Mount Limit
C8MM	Sequential	0	8MM	8500	5,120.0	DRIVES
DISK	Random	4				

The best place to look for information on tape format parameters is in the appropriate Administrators Reference manual or the server README files for devices newly available with Tivoli Storage Manager. A more common setting is (including taking hardware compression into account) that it can be supported by the drive on which a volume is mounted.

Once you have set up the device classes to match the physical devices available to your Tivoli Storage Manager server, you use these classes when defining the storage pools.

5.7.3.2 Disk storage pool

We do not need to define the device class for a disk storage pool, because these always use the Tivoli Storage Manager predefined device class of DISK.

5.7.3.3 Sequential disk storage pool

This is where a disk can simulate a sequential device, such as tape.

If you are using a single drive and want to perform reclamation, then the pool type you specify as the reclaim storage pool must be a primary sequential storage pool. This is one instance where you may want to use a sequential disk storage pool with a device type of FILE.

A copy pool must also use sequential access storage and can be defined on a disk using this device type.

To configure a sequential storage pool, you first need to define a device class of FILE, and tell Tivoli Storage Manager which directory to use and what the maximum size of this file can be.

The following command defines a device class named CFILE using a Tivoli Storage Manager predefined device type of FILE with a maximum capacity of 12 MB.

```
tsm: TSM010> define devclass cfile devtype=file directory=/tsm/stgpool/ \
cont> maxcapacity=12m
ANR2203I Device class CFILE defined.
```

```
tsm: TSM010> query devclass cfile
```

Device Class Name	Device Access Strategy	Storage Pool Count	Device Type	Format	Est/Max Capacity (MB)	Mount Limit
CFILE	Sequential	0	FILE		12.0	1

5.7.3.4 Copy storage pool

A copy storage pool can use only sequential access storage (for example, tape, optical, or file device classes). The device class you should define has to represent the type of device you are using for your copy pool. If you are using the same library as one already defined, then you can use the same device class that has been defined for it or define a separate one. If you are using a separate library or an MVS server platform, you need to define a separate device class.

Each separate library needs at least one unique device class to identify it.

5.7.4 Defining a storage pool — logical

Tivoli Storage Manager has two types of storage pools:

- Primary storage pools
- Copy storage pools

A primary storage pool can use random access storage (DISK device class) or sequential access storage (for example, tape, optical or file device classes). A copy storage pool can use only sequential access storage.

5.7.4.1 Primary storage pool for a manual library

Following the device class, you can now define the storage pool you are going to use with the library.

The following command defines a storage pool named TAPEDATA using the device class named C8MM. The high migration parameter is set to 100 to stop migration because there is no NEXTSTGPOOL for this storage pool to migrate to. The MAXSCRATCH value of 10 means that when 10 scratch tapes have been used, the storage pool indicates that it is full — this may not be desirable. However, this parameter is used to estimate the total number of volumes in the pool and the corresponding estimated capacity of the pool. MAXSCRATCH is a required value for sequential type storage pools. It can be increased at any time if more volumes are needed in the storage pool, however it is not a good practice to set MAXSCRATCH higher than the actual number of volumes available.

```
tsm: TSM010> define stgpool tapedata c8mm highmig=100 maxscratch=10
ANR2200I Storage pool TAPEDATA defined (device class C8MM).
```

```
tsm: TSM010> query stgpool tapedata
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
TAPEDATA	C8MM	0.0	0.0	0.0	100	70	

5.7.4.2 Primary storage pool for disk

Using the Tivoli Storage Manager predefined device class of DISK, you can define a primary disk storage pool.

The following command defines a storage pool named DISKDATA with a high migration threshold of 70% and a low migration threshold of 30%. It uses the predefined Tivoli Storage Manager device class of DISK.

```
tsm: TSM010> define stgpool diskdata disk description="Data Storage" \
cont> high=70 low=30
```

```
ANR2200I Storage pool DISKDATA defined (device class DISK).
```

```
tsm: TSM010> query stgpool diskdata
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKDATA	DISK	0.0	0.0	0.0	70	30	

5.7.4.3 Primary storage pool for sequential disk

This is where a disk can simulate a sequential device, such as tape.

The following command defines a storage pool named DISKFILE using a device class named CFILE. Because it simulates a sequential storage pool, the MAXSCRATCH parameter must be included, and because there is no next storage pool in this case, migration is disabled.

```
tsm: TSM010> define stgpool diskfile cfile maxscratch=100 highmig=100
ANR2200I Storage pool DISKFILE defined (device class CFILE).
```

```
tsm: TSM010> query stgpool diskfile
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKFILE	CFILE	0.0	0.0	0.0	100	70	

Operation compared to a 'real' sequential access storage pool

When using a sequential access storage pool with a FILE device class, it looks and operates slightly differently compared to a storage pool mapped to a physical tape or optical device. When disk is used, the server opens a file when one is required, and so the storage pool shows 0.0 Estimated Capacity until data has been sent to it. Following a client backup operation to this storage pool, we can see the capacity increase. Checking the volume used for this operation, the server has created a file named 00000005 with an extension of .BFS to store a client's data. Because of this difference in operation, volumes do not have to be defined for this type of diskpool — the server creates files itself, within the overall maximum capacity of the device class. You would only manually define volumes, rather than allowing them to be automatically assigned, if for some reason you wanted to have control over the volume names. Note: There are suffixes other than .BFS that may be created when using a sequential pool for other types of operation, such as backing up the server database or a primary storage pool.

```
tsm: TSM010> query stgpool diskfile
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKFILE	CFILE	1,200.0	0.3	1.0	100	70	

```
tsm: TSM010> query volume stgpool=diskfile
```

Volume Name	Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Volume Status
/tsm/stgpool/00000005.B-FS	DISKFILE	CFILE	12.0	29.8	Filling

Single drive reclamation using disk for reclaim storage pool

Reclamation can be performed by a single drive in Tivoli Storage Manager by specifying the RECLAIMSTGPOOL parameter. This parameter points to another storage pool that can be used as the holding area for the data being consolidated. The holding area then refers back to the original pool being reclaimed as its next pool to migrate to. Reclamation was discussed already in Section 5.4, "Reclamation - how to use fewer tapes" on page 84.

The following command updates the existing storage pool named TAPEDATA to use the disk storage pool named DISKFILE to hold data temporarily when performing reclamation.

```
tsm: TSM010> update stgpool tapedata reclaimstgpool=diskfile
ANR2202I Storage pool TAPEDATA updated.
```

```
tsm: TSM010> query stgpool tapedata format=detail
```

```
Storage Pool Name: TAPEDATA
Storage Pool Type: Primary
Device Class Name: C8MM
Estimated Capacity (MB): 0.0
Pct Util: 0.0
Pct Migr: 0.0
Pct Logical: 100.0
High Mig Pct: 100
Low Mig Pct: 70
Migration Delay: 0
Migration Continue: Yes
Migration Processes:
Next Storage Pool:
Reclaim Storage Pool: DISKFILE
Maximum Size Threshold: No Limit
Access: Read/Write
Description:
Overflow Location:
Cache Migrated Files?:
Collocate?: No
Reclamation Threshold: 60
Maximum Scratch Volumes Allowed: 10
Delay Period for Volume Reuse: 0 Day(s)
Migration in Progress?: No
Amount Migrated (MB): 0.00
Elapsed Migration Time (seconds): 0
Reclamation in Progress?: No
Volume Being Migrated/Reclaimed:
Last Update by (administrator): ADMIN
Last Update Date/Time: 02/09/00 15:06:48
```

The disk storage pool with a device class of FILE, named DISKFILE, is then updated to point back to the pool being reclaimed.

```
tsm: TSM010> update stgpool diskfile nextstgpool=tapedata  
ANR2202I Storage pool DISKFILE updated.
```

```
tsm: TSM010> query stgpool diskfile
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKFILE	CFILE	1,200.0	0.3	1.0	100	70	TAPEDATA

5.7.4.4 Copy storage pool

The copy storage pool is defined so that it can be used to backup the primary storage pools. This pool has a reuse delay of five days before reclaimed offsite tape volumes are moved back to the scratch pool. This is to allow for a long company shutdown, such as Easter holidays, where the offsite data must remain intact in case a disaster occurs onsite and it is required. You must set this value equal to the longest period of time before normal operations resume for you. Normal operations (database and primary pools) are backed up and volumes are taken offsite.

5.7.5 Defining storage pool volumes — logical

Tivoli Storage Manager is dynamic; you can add or remove volumes without interrupting server operations. For example, if you install a new type of device as a disk pool, the Tivoli Storage Manager administrator can move the storage pool data from the old pool to the new pool without shutting down the server. Or, if you have to add space to a storage pool, you can easily define new volumes and thereby expand the size of the storage pool without disrupting service.

Library volumes are volumes checked in to an automated library, including scratch volumes and private volumes.

5.7.5.1 Manual library

There are two methods of labeling a volume in a manual library. Described here is the more efficient one-step process, using the `label libvolume` command. The second method uses the DSMLABEL utility to create a volume and the `define volume` command from within Tivoli Storage Manager to define the volume to a storage pool.

You must use a unique label for every tape or optical volume defined in Tivoli Storage Manager. If the tapes you purchase come with physical external labels, we highly recommend you set the Tivoli Storage Manager internal labels to match.

Mount the volume you want to label in the library. The following command labels a tape in the library named M8MM with a label V26116 and overwrites any existing label and data on the tape.

```
tsm: TSM010> label libvolume m8mm v26116 overwrite=yes
ANS8003I Process number 3 started.

tsm: TSM010> query process

Process Process Description Status
Number
-----
      3 LABEL LIBVOLUME      ANR8804I Labelling volume V26116 in library
                               M8MM.
```

Note:

Although Tivoli Storage Manager allows you to use a volume identification larger than six characters, we strongly recommend that you use a maximum of six alphanumeric characters for the label (also known as VOLSER or Volume ID). This is to be compatible with other ANSI styled label systems.

5.7.5.2 Disk storage pool

Disk volumes can be defined to a Tivoli Storage Manager storage pool using either a one-step or two-step process. Described here is the more efficient one-step process. The two-step process requires preparing the disk volume using the DSMFMT utility and then defining the volume to the storage pool using the `define volume` command.

The following command defines a volume that is located in `/tsm/stgpool/` and that is named `diskdata1.dsm` to the storage pool DISKDATA. It also prepares the volume and formats it as a background process of size 50 MB.

```
tsm: TSM010> define volume diskdata /tsm/stgpool/diskdata1.dsm formatsize=50
ANR2491I Volume Creation Process starting for /tsm/stgpool/diskdata1.dsm,
Process Id 6.
```

This results in our storage pool, DISKDATA, looking like this:

```
tsm: TSM010> query stgpool diskdata

Storage   Device   Estimated   Pct   Pct   High   Low   Next
Pool Name Class Name Capacity   Util Migr Mig  Mig  Storage
-----
DISKDATA  DISK     50.0       0.0  0.0   70    30
-----

tsm: TSM010> query volume stgpool=diskdata

Volume Name           Storage   Device   Estimated   Pct   Volume
Pool Name             Class Name Capacity   Util   Status
-----
/tsm/stgpool/diskdata1.- DISKDATA  DISK     50.0       0.0  On-Line
dsm
```

Defining additional volumes to the storage pool is how you increase the size of the storage pool dynamically with no interruption to the Tivoli Storage Manager service.

Note:

If we had defined a "Next Storage Pool" to migrate to, this storage pool would have to be defined first.

5.7.5.3 Copy storage pool

If you are not using a disk for your copy storage pool, then you need to define volumes for your copy storage pool. This is done by using the methods described for defining volumes in a manual library or an automatic library (5.7, "Setting up storage pools" on page 89).

5.8 Recommended setup

Chapter 2, "Tivoli Storage Manager implementation planning" on page 11, may have referred you to this section of the data storage chapter. To complete this section successfully, the component parts of your Tivoli Storage Manager setup must be followed, in order, from 5.8.1, "Defining an automated library" on page 101 through 5.8.5, "Defining storage pool volumes" on page 107.

5.8.1 Defining an automated library

Automated libraries require a device name to be associated with them. This is the ID of the robot that moves the tapes within the library.

The library must be specified, using its device name that the operating system recognizes it by. The library therefore must already be configured under the operating system.

Windows NT Users:

Run the Tivoli Storage Manager Server Utilities. Under Device Information, check the Tivoli Storage Manager Device Name, for example lb0.0.0.0.

AIX Users:

Run the following AIX command to find the device name, for example, /dev/lb0:

```
root@ / > lsdev -C -c library -H
```

The following command defines our recommended solution, an automated library, of library type SCSI, named L3570, that was seen by our AIX system at /dev/rmt1.smc.

```
tsm: TSM010> define libr l3570 libtype=scsi device=/dev/rmt1.smc
ANR8400I Library L3570 defined.

tsm: TSM010> query library l3570

Library Name: L3570
Library Type: SCSI
Device: /dev/rmt1.smc
Private Category:
Scratch Category:
External Manager:
Shared: No
```

5.8.2 Defining a drive in an automated library

Libraries with multiple physical drives must have each drive defined to Tivoli Storage Manager and also the element address specified. The element address gives a unique name to each part of the library that concerns Tivoli Storage Manager. There are element addresses for each physical drive, the robot, and each cartridge slot within the library.

5.8.2.1 Element information

To find element information look at the Tivoli Web site:

http://www.tivoli.com/support/storage_mgr/adsercli.htm

Click on the supported device list for your operating system and click on the library you are attaching. The device support worksheet, shown in Figure 11, is a useful document to complete and file as part of your installation.

MAGSTAR MP 3570 Devices

Device	SCSI ID	Device Name
Tape drive 1 (element 16)	_____	_____
Tape drive 2 (element 17)	_____	_____
Robot	_____	_____

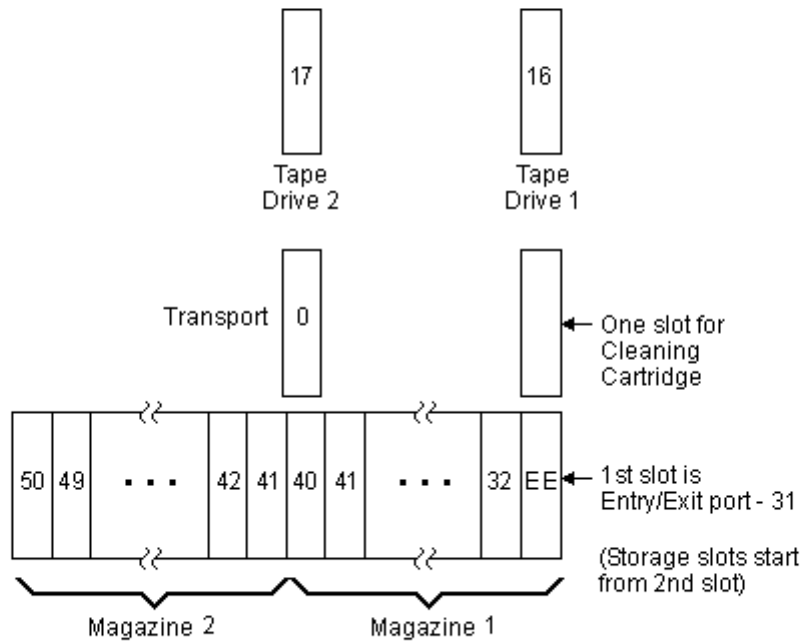


Figure 11. Device support worksheet from Tivoli Storage Manager Web site

5.8.2.2 Drive

The following command defines a drive named rmt1 and rmt2 in library named L3570. The drives are at /dev/rmt1 and /dev/rmt2 and have a element addresses of 16 and 17, respectively.

For libraries with a single drive, the element address is an optional parameter.

```

tsm: TSM010> define drive l3570 rmt1 device=/dev/rmt1 element=16
ANR8404I Drive RMT1 defined in library L3570.

tsm: TSM010> define drive l3570 rmt2 device=/dev/rmt2 element=17
ANR8404I Drive RMT2 defined in library L3570.

tsm: TSM010> query drive l3570 format=detail

                Library Name: L3570
                Drive Name: RMT1
                Device Type: 3570
                Device: /dev/rmt1
                ON LINE: Yes
                Element: 16
                Allocated to:
                Last Update by (administrator): ADMIN
                Last Update Date/Time: 02/09/00 16:09:05
Cleaning Frequency (Gigabytes/ASNEEDED/NONE): NONE
                Library Name: L3570
                Drive Name: RMT2
                Device Type: 3570
                Device: /dev/rmt2
                ON LINE: Yes
                Element: 17
                Allocated to:
                Last Update by (administrator): ADMIN
                Last Update Date/Time: 02/09/00 16:09:36
Cleaning Frequency (Gigabytes/ASNEEDED/NONE): NONE

```

5.8.3 Defining a device class to an automated library

For our recommended solution we have set up two device classes here. One is called 3570 for the library we are using in our example, and an optional one is called COFFSITE for the offsite storage pools.

Note:

The library L3570 must have two drives for these two device classes to work.

The following command defines a device class named C3570 for library L3570 using a Tivoli Storage Manager predefined device type of 3570. A tape sitting idly in one of the drives in the library remains there for 15 minutes before being automatically dismounted.

```

tsm: TSM010> define devclass c3570 devtype=3570 format=drive library=l3570 \
cont> mountret=15
ANR2203I Device class C3570 defined.

tsm: TSM010> query devclass

```

Device Class Name	Device Access Strategy	Storage Pool Count	Device Type	Format	Est/Max Capacity (MB)	Mount Limit
C3570	Sequential	0	3570	DRIVE	0.0	DRIVES
DISK	Random	1				

Carefully consider the value you use for mount retention. If your setup requires frequent mounts of different volumes, using an automatic library, then set a low value. A value of 0 may even be beneficial for drives in constant use. Collocation

increases the number of tape mounts required during backup. At the other end of the scale, with a manual drive and only one volume in daily use, console messages requesting a volume to be mounted are unwanted and may lead to processing delays until someone can respond to the message; therefore a high value should be used. The value range is 0-9999 minutes. Other factors that may influence your choice are: the time it takes to mount a volume on the library you are using, and regular access to the same volume.

We now set up the optional separate device class for offsite storage pools. The device class you should define has to represent the type of device you are using for your copy pool. If you are using the same library as one already defined, then you can use the same device class that has been defined for it — we could have used C3570 in our example. If you are using a separate library or an MVS server platform, you need to define a separate device class.

Note for MVS Users:

Most tape library management systems on MVS or OS/390 use the dataset name to identify tapes to be taken offsite. Tivoli Storage Manager uses the format `<prefix>.BFS` for both onsite and offsite dataset names in one device class. To choose another dataset name for offsite copies, create another device class for the offsite copies and choose a different prefix. Set the tape library management system to trigger on this offsite name, and offsite copies will be identified automatically.

The following command defines a device class named COFFSITE for library L3570 using a Tivoli Storage Manager predefined device type of 3570.

```
tsm: TSM010> define devclass coffsite devtype=3570 format=drive library=l3570
ANR2203I Device class COFFSITE defined.
```

```
tsm: TSM010> query devclass
```

Device Class Name	Device Access Strategy	Storage Pool Count	Device Type	Format	Est/Max Capacity (MB)	Mount Limit
C3570	Sequential	0	3570	DRIVE	0.0	DRIVES
COFFSITE	Sequential	0	3570	DRIVE	0.0	DRIVES
DISK	Random	1				

5.8.4 Defining storage pools

There are four primary storage pools and two copy pools to set up in our recommended solution. Client data backups go to storage pool DISKDATA and are then migrated to storage pool TAPEDATA before going offsite to copy pool OFFDATA. Client directory structures, which are much smaller in size, go to storage pool DISKDIRS and are then backed up to the offsite copy pool OFFDIRS. This is achieved by setting the DIRMC option in either the client options file or client option set. See Figure 12 for our recommended storage pool setup.

In 1.2, “Redbook support material” on page 2, we show how to load a predefined macro into Tivoli Storage Manager. The macro `mac.stgcreate`, that we provide to

create the storage pools in our redbook environment, is shown in B.1.6, “Create storage pools” on page 291.

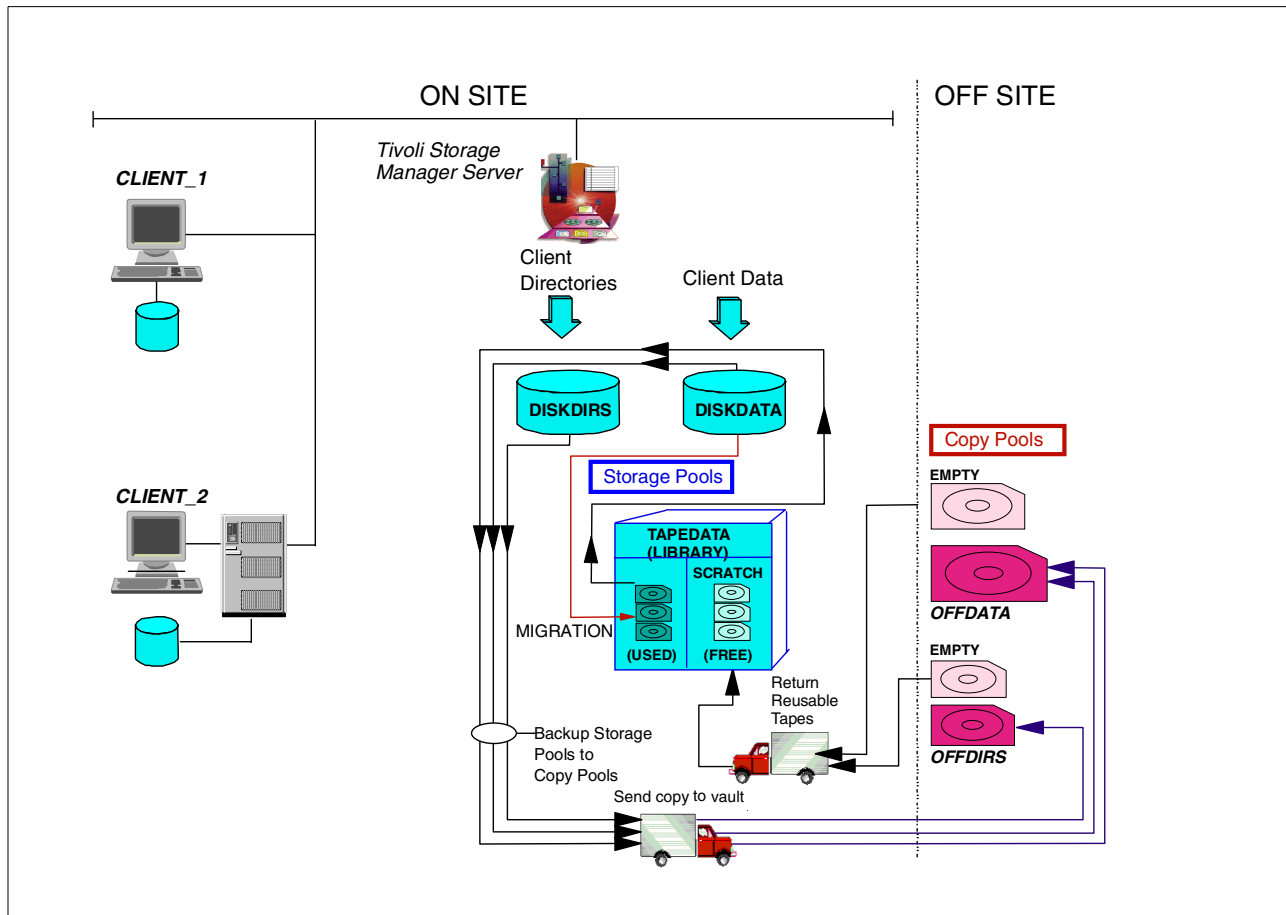


Figure 12. Storage pool hierarchy for our recommended setup

5.8.4.1 Primary storage pools

For our recommended solution we set up four primary storage pools: TAPEDATA, DISKDATA, DISKDIRS, and NONE. Storage pool TAPEDATA must be set up before storage pool DISKDATA. This is because DISKDATA refers to TAPEDATA as its next storage pool.

The following commands set up these storage pools:

1. A primary storage pool named TAPEDATA, with device class C3570. It has collocation and reclamation turned off, and has a delay of one day before empty tapes are reused. To define a storage pool for an automated library, use the same procedure as defining a pool for a manual library.
2. A primary storage pool named DISKDATA, with a device class of DISK. It has an upper threshold of 70% that must be reached for migration to start, and a lower threshold of 30% that is reached or surpassed before migration stops. The pool to which data migrated goes is called TAPEDATA. The disk pool is not using disk caching.
3. A primary storage pool named DISKDIRS, with a device class of DISK. It has migration turned off, since we want to keep the client directory structure

permanently on disk. This will allow fast reconstruction of the directory structure during a restore. The space needed to store directory information is small, so only a small-size volume will be needed in this storage pool. Note that only client operating systems where the directories have extended attributes need to store them in a storage pool, in particular NT. On UNIX clients, the directory entries are stored directly in the Tivoli Storage Manager server database. We recommend to use a special directory storage pool for purposes of control. If you do not use the DIRMC client option to specify where directories are stored, they will be sent to whichever management class specifies the longest retention period.

4. A dummy storage pool named NONE, which is just an empty definition. No data will be stored in this storage pool. There are no storage pool volumes assigned to it. See 6.2.3, "Defining management classes" on page 135 for an explanation for why we define this here.

```

tsm: TSM010> define stgpool tapedata c3570 highmig=100 maxscratch=1000 \
cont> collocate=no reclaim=100 reusedelay=1
ANR2200I Storage pool TAPEADATA defined (device class C3570).

tsm: TSM010> define stg diskdata disk nextstg=tapedata highmig=70 \
cont> lowmig=30 cache=no
ANR2200I Storage pool DISKDATA defined (device class DISK).

tsm: TSM010> define stgpool diskdirs disk highmig=100
ANR2200I Storage pool DISKDIRS defined (device class DISK).

tsm: TSM010> define stgpool none disk
ANR2200I Storage pool NONE defined (device class DISK).

tsm: TSM010> query stgpool

```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKDATA	DISK	0.0	0.0	0.0	70	30	TAPEADATA
DISKDIRS	DISK	0.0	0.0	0.0	100	70	
DISKFILE	CFILE	1,200.0	0.3	1.0	100	70	
NONE	DISK	0.0	0.0	0.0	90	70	
TAPEADATA	C3570	0.0	0.0	0.0	100	70	

5.8.4.2 Copy storage pools

For our recommended solution we have set up two copy storage pools: OFFDIRS and OFFDATA. These pools have a reuse delay of five days before reclaimed offsite tape volumes are moved back to the scratch pool. This is to allow for a long company shutdown, such as Easter holidays, where the offsite data must remain intact in case a disaster occurs onsite and it is required for the recovery. You must set this value equal to the longest period of time before normal operations resume for you. Normal operations are when the database and primary pools are backed up and volumes are taken offsite.

The following commands set up these storage pools:

1. A copy storage pool named OFFDIRS, with device class COFFSITE. Reclamation is switched off and the reuse delay is set for five days.
2. A copy storage pool named OFFDATA, with device class COFFSITE. Reclamation is switched off and the reuse delay is set for five days. The MAXSCRATCH value is set high, purposefully, to avoid a "storage pool full" message occurring.

```

tsm: TSM010> define stgpool offdirs coffsite pooltyp=copy reclaim=100 \
cont> maxscratch=10000 reusedelay=5
ANR2200I Storage pool OFFDIRS defined (device class COFFSITE).

tsm: TSM010> define stgpool offdata coffsite pooltyp=copy reclaim=100 \
cont> maxscratch=10000 reusedelay=5
ANR2200I Storage pool OFFDATA defined (device class COFFSITE).

tsm: TSM010> query stgpool pooltype=copy

```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
OFFDATA	COFFSITE	0.0	0.0				
OFFDIRS	COFFSITE	0.0	0.0				

5.8.5 Defining storage pool volumes

The following shows how to set up both tape storage pool volumes for an automated library and disk storage pool volumes.

5.8.5.1 Tape

Setting up storage pool volumes for an automated library is done in three stages:

1. Label volumes and check them into scratch pool
2. Check for outstanding mount requests
3. Reply to mount requests and issue labels

Similarly to creating volumes in a disk storage pool, the first step can be done in a one or two step process. The one-step process involves using the `label libvolume` administrative command which is available on all non VM, MVS or OS/390 Version 3 or higher platforms. To do it in two-steps, you use the operating system `dsmlabel` command, followed by the `checkin libvolume` administrative command. Both methods can be set up to automatically handle labeling a series of tapes, using the library's dedicated input ports or slots if provided, or searching the regular library slots for appropriate tapes. If your library has a barcode reader the process is even simpler as you can specify options on the `dsmlabel (-barcode)` or `label libvolume (LABELSOURCE=BARCODE)` commands that mean Tivoli Storage Manager will automatically use each tape's barcode label as its internal label instead of manually prompting for it.

If your library does not support a barcode reader or an automatic search, you may need to manually check for outstanding tape requests and issue a reply to the request number with the label. In this case, you or an operator would also need to be on hand to physically add and remove volumes from the library. You will also need to type in the required label for each volume, or redirect a list of volume names as input to the command.

Label volumes and check them in

The following command searches for tapes in the library named L3570 that are not already defined to Tivoli Storage Manager, and prompts you for a label. It does not overwrite any tapes that have been labeled already, and it does not re-label any tapes already checked in.

```
tsm: TSM010> label libvolume l3570 search=yes labelsource=prompt checkin=scratch \
cont> overwrite=no
ANS8003I Process number 8 started.
```

Check for outstanding mount requests

To see if there are any pending mount requests, use the `query request` command.

```
tsm: TSM010> query request
ANR8352I Requests outstanding:
ANR8809I 004: Please provide the label name for the volume in slot element 32
of library L3570 by issuing REPLY n LABEL=xxx within 59 minutes, where n is
the re quest ID and xxx is the desired label name.
```

Or, you can check the activity log to find out the number to reply to this request.

```
tsm: TSM010> query actlog begintime=-00:01

Date/Time          Message
-----
02/09/00  16:56:14 ANR8352I Requests outstanding:
02/09/00  16:56:14 ANR8809I 004: Please provide the label name for the volume
in slot element 32 of library L3570 by issuing REPLY n
LABEL=xxx within 59 minutes, where n is the re quest ID
and xxx is the desired label name.
```

Here we see the ANR8809I message, which is followed by the request number (004).

Reply to mount requests and issue labels

To label a tape we need to reply to this message. The following command replies to request number 4 and issues a label name of V00001 to the tape volume.

```
tsm: TSM010> reply 4 label=V00001
ANR8499I Command accepted.
```

Tivoli Storage Manager continues searching through the library for non-defined tapes and continues issuing requests for labels. For this example, we had three volumes, and cycling through the procedure above, we labeled and checked them into the scratch pool.

```
tsm: TSM010> query libvolume

Library Name  Volume Name  Status      Owner      Last Use    Home Element
-----
L3570         ALF007      Scratch
L3570         CLA001      Scratch
L3570         DAN001      Scratch
                                     32
                                     33
                                     35
```

Note:

If you label volumes incorrectly and wish to re-label them, you will need to check the volume out of the library.

5.8.5.2 Disk

The size of the volumes you create depends on the values you decided are necessary in your initial planning. The values used below are small and are intended to be used only as an example.

When deciding how many volumes to use, consider this. It may be better to use a large volume to cover most of the expected storage pool size and then some smaller volumes to allow yourself flexibility in reassigning volumes to different storage pools as the need arises.

An example of this strategy could be to increase the size of the archive pool for monthly archives. In this case, you could migrate the backup random access storage pool to its next sequential storage pool. Redefine some of the backup storage pool volumes as archive pool volumes, complete the monthly archive, and back up the pool to offsite. You can then reassign volumes to their original state. This is just one example of how you can more efficiently use the resources at your disposal, since having a number of volumes in your storage pool allows you to redefine the volume or volumes to cover the extra capacity you need.

An alternative strategy is to define a number of volumes, all of the same size, since when writing backup or archive data to a disk pool, Tivoli Storage Manager will automatically use all the storage pool volumes concurrently, rather than filling up each volume in turn.

The process of adding a new storage pool volume consists of two steps:

1. Format the volume using the DSMFMT utility
2. Define the volume to Tivoli Storage Manager using the `define volume` command

In a non VM, MVS or OS/390 Version 3 or higher platform, those two steps can be combined into a single process: you format the volume and immediately define it in Tivoli Storage Manager by using the `define volume` command using the `FORMATSIZE` option. We recommend that you use the one-step method if it is available.

The following commands show how to format and allocate a volume for each of two storage pools:

1. Format a data volume named `file01` under the `/tsm/stgpool` directory. The size is specified in MB and in this case is 200 MB. It then assigns the volume to the disk storage pool `DISKDATA`.
2. Format a data volume named `dirs01` under the `/tsm/stgpool` directory of size 64 MB. It then assigns the volume to the disk storage pool `DISKDIRS`.

```

tsm: TSM010> define volume diskdata /tsm/stgpool/file01 formatsize=200
ANR2491I Volume Creation Process starting for /tsm/stgpool/file01, Process Id 11.

tsm: TSM010> define volume diskdirs /tsm/stgpool/dirs01 formatsize=64
ANR2491I Volume Creation Process starting for /tsm/stgpool/dirs01, Process Id 12.

tsm: TSM010> query vol

```

Volume Name	Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Volume Status
/tsm/stgpool/dirs01	DISKDIRS	DISK	64.0	0.0	On-Line
/tsm/stgpool/file01	DISKDATA	DISK	200.0	0.0	On-Line

There is a guide for sizing the DISKDATA storage pool in the planning chapter, but the DISKDIRS storage pool must be treated differently. The size of this pool depends largely on the types of directories you backup. Network directories, because of the many trustee directory assignments, take more space than Windows NT directories, which have fewer options for access permissions on their directory structure. At the other end of the scale, UNIX directories without access control lists (ACLs) may all be saved directly in the Tivoli Storage Manager database, and will not increase the space utilization percentage of the DISKDIRS storage pool.

We recommend that you initially choose a capacity of 64 MB or 128 MB. The OFFDIRS storage pool is also small and is likely to use part of one volume. However, the advantages of fast restores—in case of disaster—outweigh this small inefficiency.

5.8.6 Deleting default storage pools

You can now delete the default storage pools to tidy up your installation, by using the `delete stgpool` command:

```

tsm: TSM010> delete stgpool backuppool
ANR2201I Storage pool BACKUPPOOL deleted.

tsm: TSM010> delete stgpool archivepool
ANR2201I Storage pool ARCHIVEPOOL deleted.

tsm: TSM010> delete stgpool spacemgpool
ANR2201I Storage pool SPACEMGPOOL deleted.

```

In 1.2, “Redbook support material” on page 2, we show how to load a predefined macro into Tivoli Storage Manager. The macro `mac.stgdelete`, which we provide to delete the default storage pools, is shown in B.1.7, “Delete default storage pools” on page 292. This macro deletes the four default storage pools automatically set up by the Windows NT installation; including an additional storage pool named DISKPOOL. (Note: A Tivoli Storage Manager install on a UNIX platform does not create this pool, so this line needs to be removed before running the macro.)

5.9 Additional commands

In the following we provide some examples of commands that have a direct relationship to data storage.

5.9.1 Audit library

The `audit library` command verifies that an automated library is in a consistent state with respect to the server's internal library volume inventory. You should use this command if someone accidentally or deliberately moves tapes around in their slots, or inserts or removes volumes without using the correct Tivoli Storage Manager commands. It is also useful as a general check, if you suspect that Tivoli Storage Manager's own inventory of what tapes are where does not match the real situation. If any inconsistencies are detected, the server updates its inventory to reflect the current state of the library. Missing volumes are deleted and moved volumes are updated with their new location. New volumes are not added, they must be checked in. This command does not start until all volumes in any of the tape drives are dismounted. You can force this using the `dismount volume` command. It then reads the labels from all the volumes in the library. This can be the label written on the media itself (default), or you can specify to use the barcode-read label if available by specifying `CHECKLABEL=BARCODE`. Be aware that the `audit library` command can be a time consuming task, especially for large capacity libraries, as it must mount and read the label on every volume.

The following command audits the library named L3570:

```
tsm: TSM010> audit library l3570
ANS8003I Process number 13 started.

tsm: TSM010> query proc

Process Process Description Status
Number
-----
13 AUDIT LIBRARY ANR8459I Auditing volume inventory for library
L3570.

02/09/00 17:18:50 ANR2017I Administrator ADMIN issued command: AUDIT LIBRARY
L3570
02/09/00 17:18:50 ANR0984I Process 13 for AUDIT LIBRARY started in the
BACKGROUND at 17:18:50.
02/09/00 17:18:50 ANR8457I AUDIT LIBRARY: Operation for library L3570
started as process 13.
02/09/00 17:25:42 ANR8461I AUDIT LIBRARY process for library L3570 completed
successfully.
02/09/00 17:25:42 ANR0985I Process 13 for AUDIT LIBRARY running in the
BACKGROUND completed with completion state SUCCESS at
17:25:42.
```

5.9.2 Audit volume

The `audit volume` command checks for inconsistencies between the database references for what files are stored on a volume, and what is actually found to be stored on the volume. It requires that the volume is mounted in a tape drive so its contents can be read. Compare this with the `query content` command (see 5.9.10, "Query content" on page 119) which displays only the database's view of the what is stored and does not require the volume to be mounted. The `audit volume` command is especially useful when integrity errors are detected on a volume, as you can determine what, if any, files on the volume are still readable and take appropriate action to restore it from a copy storage pool.

The following command audits a volume named ALF007 and logs discrepancies to the activity log.

```

tsm: TSM010> audit vol alf007 fix=no
ANR2310W This command will compare all inventory references to volume ALF007
with the actual data stored on the volume and will report any discrepancies;
the data will be inaccessible to users until the operation completes.

Do you wish to proceed? (Yes/No) y
ANR2313I Audit Volume (Inspect Only) process started for volume ALF007 (process
ID 15).
ANS8003I Process number 15 started.

```

Using parameter FIX=NO, which is the default, reports any discrepancies to the activity log only. No changes are made. The course of action you take, if problems are found, may differ, so this monitoring and reporting process should be used first before re-running with FIX=YES.

```

tsm: TSM010> audit vol alf007
ANR2310W This command will compare all inventory references to volume ALF007
with the actual data stored on the volume and will report any discrepancies;
the data will be inaccessible to users until the operation completes.

Do you wish to proceed? (Yes/No) y
ANR2313I Audit Volume (Inspect Only) process started for volume ALF007 (process
ID 17).
ANS8003I Process number 17 started.

tsm: TSM010> query proc

Process Process Description Status
Number -----
-----
17 Audit Volume Volume ALF007 (storage pool TAPEDATA), Files
(Inspect Only) Processed: 0, Damaged Files Found: 0, Partial
Files Skipped: 0. Current Physical File (bytes):
1,074,198
Current input volume: ALF007.
02/09/00 17:41:30 ANR2017I Administrator ADMIN issued command: AUDIT VOLUME alf007
02/09/00 17:41:31 ANR1199I Removable volume ALF007 is required for audit process.
02/09/00 17:41:31 ANR8324I 3570 volume ALF007 is expected to be mounted (R/W).
02/09/00 17:41:31 ANR0984I Process 17 for AUDIT VOLUME (INSPECT ONLY)
started in the BACKGROUND at 17:41:31.
02/09/00 17:41:31 ANR2313I Audit Volume (Inspect Only) process started for
volume ALF007 (process ID 17).
02/09/00 17:41:35 ANR2315I Audit volume process ended for volume ALF007;
1646 files inspected, 0 damaged files found and marked as
damaged.
02/09/00 17:41:35 ANR0987I Process 17 for AUDIT VOLUME (INSPECT ONLY)
running in the BACKGROUND processed 1646 items with a
completion state of SUCCESS at 17:41:35.

```

5.9.3 Back up storage pool

The backup `stgpool` command backs up a primary storage pool to a copy pool. If the primary storage pool is a sequential storage pool, then this command requires two physical drives since the input volume from the primary storage pool will be mounted in one, and the output volume from the copy storage pool will be mounted in the other. The two drives can be in the same or different tape libraries. To back up a random access storage pool to a sequential storage pool requires only one drive. Typically this process is automated: we describe how in 11.2.1, “Defining an offsite backup schedule” on page 218. Backing up from a disk primary pool is quicker than backing up from a tape storage pool. This is why we recommend you size your disk pools sufficiently to hold a normal night’s backup operation, since in this way, there will have been little or no migration to tape during the backup. Typically you would define one copy storage pool for a primary storage pool hierarchy. In our setup, both DISKDATA and TAPEDATA get

backed up to the copy storage pool OFFDATA. Since the backup stgpool operation is incremental, only new files are put into the copy storage pool. This means that if we first backup the storage pool DISKDATA, then TAPEDATA, any files which might have migrated from disk to tape after we backed up DISKDATA do not get copied a second time from the TAPEDATA storage pool.

The following command backs up a storage pool named TAPEDATA to a copy pool named OFFDATA.

```

tsm: TSM010> backup stgpool tapedata offdata
ANS8003I Process number 18 started.

tsm: TSM010> query proc

Process Process Description Status
Number
-----
18 Backup Storage Pool Primary Pool TAPEDATA, Copy Pool OFFDATA, Files
Backed Up: 0, Bytes Backed Up: 0, Unreadable
Files: 0, Unreadable Bytes: 0. Current Physical
File (bytes): 1,074,198

Current input volume: ALF007.

Waiting for mount of scratch volume (11
seconds).

```

Check the activity log to see the successful completion, as shown in this example.

```

tsm: TSM010> query actlog

02/09/00 17:54:40 ANR2017I Administrator ADMIN issued command: BACKUP
STGPOOL tapedata offdata
02/09/00 17:54:40 ANR0984I Process 18 for BACKUP STORAGE POOL started in the
BACKGROUND at 17:54:40.
02/09/00 17:54:40 ANR2110I BACKUP STGPOOL started as process 18.
02/09/00 17:54:40 ANR1210I Backup of primary storage pool TAPEDATA to copy
storage pool OFFDATA started as process 18.
02/09/00 17:54:40 ANR1228I Removable volume ALF007 is required for storage
pool backup.
02/09/00 17:54:40 ANR8324I 3570 volume ALF007 is expected to be mounted
(R/W).
02/09/00 17:55:12 ANR8329I 3570 volume ALF007 is mounted R/W in drive RMT1
(/dev/rmt1), status: IDLE.
02/09/00 17:55:12 ANR8334I 1 volumes found.
02/09/00 17:55:15 ANR2017I Administrator ADMIN issued command: QUERY REQ
02/09/00 17:55:15 ANR8346I QUERY REQUEST: No requests are outstanding.
02/09/00 17:55:15 ANR2017I Administrator ADMIN issued command: ROLLBACK
02/09/00 17:55:19 ANR2017I Administrator ADMIN issued command: QUERY DRIVE
02/09/00 17:55:35 ANR2017I Administrator ADMIN issued command: DEFINE DRIVE
l3570 rmt2 device=/dev/rmt2 element=17
02/09/00 17:55:48 ANR8404I Drive RMT2 defined in library L3570.
02/09/00 17:56:02 ANR8337I 3570 volume CLA001 mounted in drive RMT2
(/dev/rmt2).
02/09/00 17:56:02 ANR1340I Scratch volume CLA001 is now defined in storage
pool OFFDATA.
02/09/00 17:56:12 ANR1212I Backup process 18 ended for storage pool
TAPEDATA.
02/09/00 17:56:12 ANR0986I Process 18 for BACKUP STORAGE POOL running in the
BACKGROUND processed 1646 items for a total of 6,459,588
bytes with a completion state of SUCCESS at 17:56:12.
02/09/00 17:56:12 ANR1214I Backup of primary storage pool TAPEDATA to copy
storage pool OFFDATA has ended. Files Backed Up: 1646,
Bytes Backed Up: 6459588, Unreadable Files: 0, Unreadable
Bytes: 0.

```

5.9.4 Checkin libvolume

The `checkin libvolume` command checks in a library volume physically placed in the library, so that it can be seen by Tivoli Storage Manager.

The following command checks-in volume DAN001 as a scratch volume.

```
tsm: TSM010> checkin libv l3570 dan001 status=scratch
ANS8003I Process number 21 started.

tsm: TSM010> query req
ANR8352I Requests outstanding:
ANR8323I 011: Insert 3570 volume DAN001 R/W into entry/exit port of library
L3570 within 60 minute(s); issue 'REPLY' along with t he request ID when
ready.

tsm: TSM010> reply 11
ANR8499I Command accepted.

tsm: TSM010> query libvolume
```

Library Name	Volume Name	Status	Owner	Last Use	Home Element
L3570	ALF007	Private		Data	32
L3570	CLA001	Private		Data	33
L3570	DAN001	Scratch			35

5.9.5 Checkout library volumes

There are three parts to this operation:

- Check out the library volume.
- Check for outstanding mount requests.
- Reply to mount request.

Checkout library volumes

This command checks out library volume named DAN001 from a library named L3570.

```
tsm: TSM010> checkout libv l3570 dan001
ANS8003I Process number 22 started.

tsm: TSM010> query proc
```

Process Number	Process Description	Status
22	CHECKOUT LIBVOLUME	ANR8436I Checking out volume DAN001 from library L3570.

```
tsm: TSM010> query libv
```

Library Name	Volume Name	Status	Owner	Last Use	Home Element
L3570	ALF007	Private		Data	32
L3570	CLA001	Private		Data	33
L3570	DAN001	Scratch			35

Check for mount requests

To check for outstanding mount requests, issue a `query request` command or check the activity log.

```
tsm: TSM010> query request
ANR8352I Requests outstanding:
ANR8322I 012: Remove 3570 volume DAN001 from entry/exit port of library L3570;
issue 'REPLY' along with the request ID when ready.
```

Reply to mount request

Issue the reply command.

The `reply` command replies to request number 012 to confirm that volume DAN001 has been removed from the library. To re-label the tape only, there is no need to physically remove the volume from the library.

```
tsm: TSM010> reply 12
ANR8499I Command accepted.

tsm: TSM010> query libv
```

Library Name	Volume Name	Status	Owner	Last Use	Home Element
L3570	ALF007	Private		Data	32
L3570	CLA001	Private		Data	33

Re-labeling an incorrectly labeled tape

If you wish to re-label an incorrectly labeled tape, ensure that the parameter `OVERWRITE=YES` is set when executing the `label libvolume` command. The following command searches for and re-labels any library volumes not yet checked in to Tivoli Storage Manager, and overwrites an existing label.

```
tsm: TSM010> label libv 13570 search=yes labelsource=prompt checkin=scratch \
cont> overwrite=yes
ANS8003I Process number 23 started.
```

If you want to re-label just one volume, you may specify it explicitly rather than use the `SEARCH=YES` and the `LABELSOURCE=PROMPT` parameter, but you still need to check it out first.

5.9.6 Delete a library

To explain the delete command, we have an example for deleting a library. This shows how the different levels are interlinked and must be considered for successful deletion.

To delete a library you have to remove any references to it. There are many levels above a library that you may want deleted or updated to point to something else. You need to delete the following objects in the correct order:

1. Storage pool
2. Device class
3. Drive
4. Library

The following example explains how to delete library M8MM that we set up in 5.7.1, “Defining a library — physical” on page 92. We show how to delete the whole structure, including the storage pool initially set up for this library.

5.9.6.1 Delete storage pool

Before you delete the storage pool TAPEDATA, you need to remove any references to it. Check the hierarchical structure using `query stgpool` commands.

```
tsm: TSM010> query stgpool diskdata
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKDATA	DISK	200.0	0.0	0.0	0	0	TAPEDATA

```
tsm: TSM010> query stgpool tapedata
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
TAPEDATA	C8MM	0.0	0.0	0.0	90	70	

Update storage pool DISKDATA to remove the reference to TAPEDATA as its next storage pool. You cannot delete a storage pool that is defined as a NEXTSTGPOOL for another storage pool. Do not delete a storage pool that is specified as a destination for a management class or copy group in the ACTIVE policy set (see Chapter 6, “Data storage policy” on page 123 for more information).

```
tsm: TSM010> update stgpool diskdata next=""
ANR2202I Storage pool DISKDATA updated.
```

```
tsm: TSM010> query stgpool diskdata
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKDATA	DISK	200.0	0.0	0.0	0	0	

The following command deletes the storage pool named TAPEDATA.

```
tsm: TSM010> delete stgpool tapedata
ANR2201I Storage pool TAPEDATA deleted.
```

Note:

You can only delete a storage pool if there are no storage pool volumes assigned to it. If there are storage pool volumes assigned then you need to use the `delete volume` command to remove them first. With this command you can specify whether or not to discard any data stored on these volumes. Be very careful with this command! You can also use the `move data` command to move data from one storage pool to another. This is used if you do not want to lose the data from the storage pool which is being deleted.

5.9.6.2 Delete device class

Next you have to delete the device class of the library. The following commands query all defined device classes and then delete the device class named C8MM. You cannot delete the device class unless there are no storage pools defined with it. You also cannot delete the predefined device class DISK.

```
tsm: TSM010> query devclass
```

Device Class Name	Device Access Strategy	Storage Pool Count	Device Type	Format	Est/Max Capacity (MB)	Mount Limit
C3570	Sequential	0	3570	DRIVE	0.0	DRIVES
C8MM	Sequential	0	8MM	8500	5,120.0	DRIVES
DISK	Random	3				

```
tsm: TSM010> delete devclass c8mm
ANR2204I Device class C8MM deleted.
```

5.9.6.3 Delete drive

Next you have to delete the drive of the library. The following commands query all defined drives and delete the drive named DRV0 from library named M8MM.

```
tsm: TSM010> query drive
```

Library Name	Drive Name	Device Type	Device	ON LINE
M8MM	DRV0	8MM	/dev/mt0	Yes

```
tsm: TSM010> delete drive m8mm drv0
ANR8412I Drive DRV0 deleted from library M8MM.
```

5.9.6.4 Delete library

Now library M8MM can be deleted, since it no longer has any drives associated with it.

```
tsm: TSM010> delete library m8mm
ANR8410I Library M8MM deleted.

tsm: TSM010> query library m8mm
ANR2034E QUERY LIBRARY: No match found using this criteria.
ANS8001I Return code 11.
```

5.9.7 Dismount a volume

The following command dismounts the volume named CLA001 that is physically inside one of the library's drives.

```
tsm: TSM010> dismount volume cla001
ANR8499I Command accepted.

tsm: TSM010> query mount
ANR8331I 3570 volume CLA001 is mounted R/W in drive RMT2 (/dev/rmt2), status:
DISMOUNTING.
ANR8334I          1 volumes found.

tsm: TSM010> query mount
ANR2034E QUERY MOUNT: No match found using this criteria.
```

5.9.8 Migration

Migration is forced by updating the threshold parameters for migration on the storage pool to a value higher than its current utilization level. This will cause migration to start automatically. To clear all data from one pool to the next storage pool, you change the high and low values to 0.

Check storage pool migration values before you change them so that you know what to reset them back to after the migration process is finished.

```
tsm: TSM010> query stgpool
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKDATA	DISK	200.0	26.4	26.4	70	30	TAPEDATA
TAPEDATA	C3570	0.0	0.0	0.0	90	70	

The following command forces migration for storage pool DISKDATA. This is done by setting both storage pool parameters, the high threshold and the low threshold, to 0. Use the query process command to check on the progress. You also can search the activity log over the last hour for any string "migra".

```
tsm: TSM010> update stgpool diskdata hi=0 lo=0
ANR2202I Storage pool DISKDATA updated.
```

```
tsm: TSM010> query proc
```

Process Number	Process Description	Status
28	Migration	Disk Storage Pool DISKDATA, Moved Files: 0, Moved Bytes: 0, Unreadable Files: 0, Unreadable Bytes: 0. Current Physical File (bytes): 1,077,248 Current output volume: ALF007.

```
tsm: TSM010> query actlog search=migra*
```

Date/Time	Message
02/09/00 18:40:43	ANR0984I Process 28 for MIGRATION started in the BACKGROUND at 18:40:43.
02/09/00 18:40:43	ANR1000I Migration process 28 started for storage pool DISKDATA.
02/09/00 18:41:15	ANR1001I Migration process 28 ended for storage pool DISKDATA.
02/09/00 18:41:15	ANR0986I Process 28 for MIGRATION running in the BACKGROUND processed 1686 items for a total of 55,345,152 bytes with a completion state of SUCCESS at 18:41:15.
02/09/00 18:42:22	ANR2017I Administrator ADMIN issued command: QUERY ACTLOG search=migra*

The data has moved from the storage pool DISKDATA, to storage pool TAPEDATA. After migration has completed, you have to change the threshold values back to their original levels.

```

tsm: TSM010> query stgpool

Storage      Device      Estimated   Pct   Pct  High  Low  Next
Pool Name    Class Name  Capacity   Util  Migr  Mig  Mig  Storage
-----
DISKDATA     DISK        200.0      0.0   0.0   0    0    TAPEDATA
TAPEDATA     C3570      50,000,000 0.0   0.0   90   70

tsm: TSM010> update stgpool diskdata high=70 low=30
ANR2202I Storage pool DISKDATA updated.

```

5.9.9 Move data

The `move data` command moves the data from one volume to another within the same storage pool (this requires two drives to be available if it is a sequential storage pool), or to another storage pool.

The following command moves the data from volume ALF007 in the TAPEDATA storage pool to a volume in the DISKDATA.

```

tsm: TSM010> move data alf007 stg=diskdata
ANR2233W This command will move all of the data stored on volume ALF007 to
other volumes in storage pool DISKDATA; the data will be inaccessible to users
until the operation completes.

Do you wish to proceed? (Yes/No) y
ANS8003I Process number 29 started.

```

Here is an example where you may use this command: A volume showed a write error, but after checking for any inconsistencies between the volume and the database (using the `audit volume` command), the data on the volume was determined to be good. However, you do not want to use this volume any more. You may now move all the data from this volume to another volume in the same storage pool. The volume with an error can now be checked out of the library and discarded.

5.9.10 Query content

The `query content` command queries the contents of a volume and lists all the files contained on it. The output for this command can be very long.

```

tsm: TSM010> query content alf007

Node Name          Type Filespace  Client's Name for File
                   Name
-----
PALANA             Bkup /         / .sh_history
PALANA             Bkup /         / dsmerror.log
PALANA             Bkup /         / .dt/ errorlog
PALANA             Bkup /         /bruno/home2clau/06/

```

It can be reduced by specifying the number of files you want displayed by using the `COUNT` parameter. The following command lists the contents of volume ALF007, but restricts the output to files that are backed up to a copy storage pool, and only shows four lines of output. If the parameter was `COPIED=NO`, the output would only show files that were not backed up to a copy pool.

```
tsm: TSM010> query content alf007 copied=yes count=4
```

Node Name	Type	Filespace Name	Client's Name for File
PALANA	Bkup	/	/.sh_history
PALANA	Bkup	/	/dsmerror.log
PALANA	Bkup	/	/.dt/errorlog
PALANA	Bkup	/	/.netscape/cache/06/cache380288A600B44FA

5.9.11 Query occupancy

The following command displays information on where a client's filespace are stored and how much space they occupy.

```
tsm: TSM010> query occupancy palana
```

Node Name	Type	Filespace Name	Storage Pool Name	Number of Files	Physical Space Occupied (MB)	Logical Space Occupied (MB)
PALANA	Bkup	/	DISKFILE	151	3.57	3.57
PALANA	Bkup	/	OFFDATA	1,658	35.33	35.33
PALANA	Bkup	/	TAPEDATA	1,658	35.33	35.33
PALANA	Bkup	/usr	OFFDATA	28	17.34	17.34
PALANA	Bkup	/usr	TAPEDATA	28	17.34	17.34

The query occupancy command can also be used to determine whose filespace are occupying a certain device class or storage pool.

```
tsm: TSM010> query occupancy devclass=disk
```

Node Name	Type	Filespace Name	Storage Pool Name	Number of Files	Physical Space Occupied (MB)	Logical Space Occupied (MB)
PALANA	Bkup	/	DISKDATA	1	0.00	0.00
PALANA	Bkup	/usr	DISKDATA	5	60.35	60.35

5.9.12 Rename a storage pool

The rename stgpool command renames a storage pool. The following command renames SPACEMGPOOL to SPACEMGMT.

```
tsm: TSM010> rename stgpool spacemgpool spacemgmt
ANR2213I RENAME STGPOOL: Storage pool SPACEMGPOOL renamed to SPACEMGMT.
```

5.9.13 SQL commands

Tivoli Storage Manager uses a database that accepts structured query language (SQL) commands. Tivoli Storage Manager provides three system catalog tables:

- SYSCAT.TABLES: Contains information about all tables that can be queried.
- SYSCAT.COLUMNS: Describes the columns in each table.
- SYSCAT.ENUMTYPES: Defines the valid values for each enumerated data type.

For select operations, you can simply use the TABLES, COLUMNS and ENUMTYPES to display the database structure:

```
SELECT * FROM TABLES
SELECT * FROM COLUMNS
SELECT * FROM ENUMTYPES
```

There are many commands that can be generated in this way to create a customized query to the Tivoli Storage Manager database. There are a small number of useful commands listed here to start with.

Example 1

The following command lists all volumes from the database table VOLUMES which are in an error condition. In our example, this command fails, since we do not have any volumes in an error condition.

```
tsm: TSM010> select * from volumes where error_state='YES'
ANR2034E SELECT: No match found using this criteria.
ANS8001I Return code 11.
```

Example 2

The following command counts the number of volumes within each storage pool. This can be useful where large libraries with many volumes are concerned.

```
tsm: TSM010> select count(*),stgpool_name from volumes group by stgpool_name

  Unnamed[1]      STGPOOL_NAME
-----
           1      DISKDATA
           1      DISKDIRS
           1      DISKFILE
           2      OFFDATA
           1      TAPEDATA
```

Example 3

The following command lists how many files have been backed up from each node. If you replace BACKUPS with ARCHIVES you see the number of files archived from each node.

```
tsm: TSM010> select node_name,count(*) from backups group by node_name
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/No) yes

  NODE_NAME      Unnamed[2]
-----
AIXCLIENT1      57
ANDY             10522
NOGALES_EXCH    2
PAGOPAGO_EXCH   13
SCANDIUM_EXCH   1
WKS.23FFHEV    51
```

Note the warning which is displayed when a potentially complex or time/space-consuming query is issued. This allows you to decide to run the query at another time if server operations could be adversely affected.

Example 4

The following command shows the volumes in which all filespace from client PALANA are residing for the storage pool TAPEDATA .

```
tsm: TSM010> select distinct node_name,volume_name from volumeusage where \  
cont> node_name='PALANA' and stgpool_name='TAPEDATA'
```

NODE_NAME	VOLUME_NAME
-----	-----
PALANA	ALF007

For further information on SQL commands, use the `help select` command.

There are also some useful sample scripts available. See 3.4.1, “Load sample scripts” on page 52 for an explanation of how to load the sample scripts into the Tivoli Storage Manager database.

Chapter 6. Data storage policy

A data storage environment consists of three types of resources: machines, rules, and data. The machines are computers containing data that must be backed up and the rules specify how the backup copies are to be treated. Basically, a data storage policy defines the relationships between these three resources. Figure 13 illustrates this policy relationship. See 2.2, “Data retention requirements” on page 19 for further planning considerations.

Depending on how you need to manage your backup data, your data storage policy can be very simple or very complex. The simplest policy would have one set of rules that apply to all of the data in your organization while the most complex policy would be to have a unique set of rules for each individual file. Most installations fall somewhere between these two extremes.

Tivoli Storage Manager has entities that group and organize the resources and define relationships between them. A machine, or node in Tivoli Storage Manager terminology, is grouped together with other nodes into a policy domain. The domain links the nodes to a policy set which consists of management classes. A management class contains rules called copy groups that it links to the data. When the data is linked to particular rules, it is said to be bound to the management class that contains the rules.

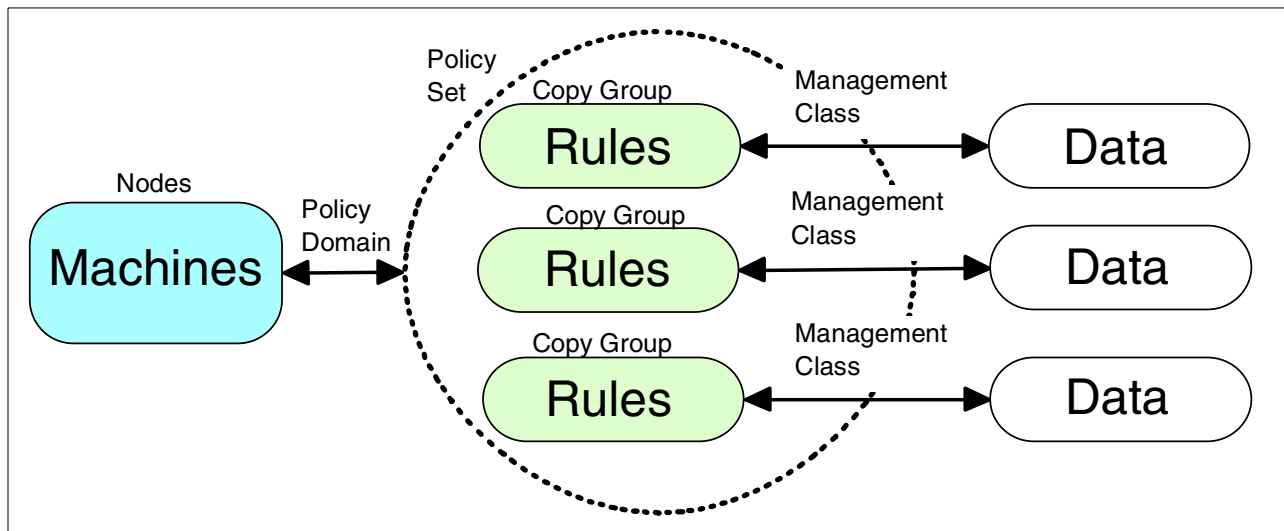


Figure 13. Policy relationships and resources

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. Our view of the policy components is illustrated in Figure 14.

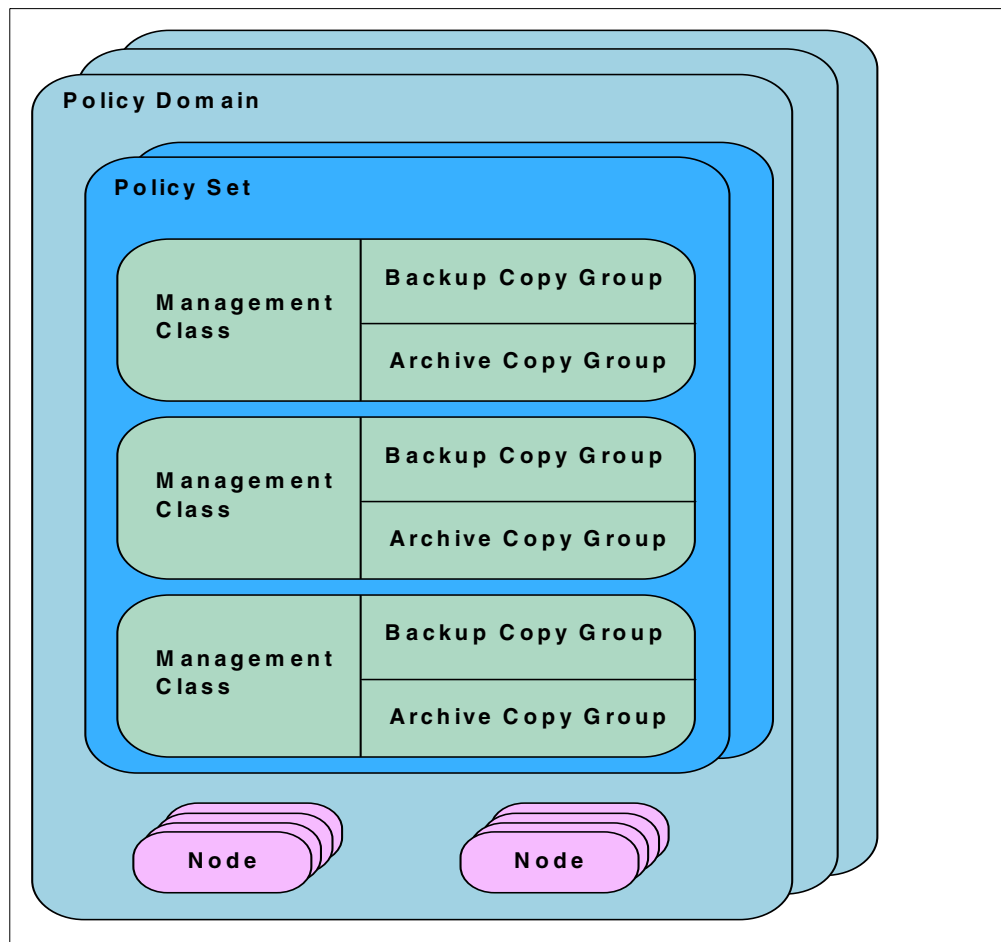


Figure 14. Data storage policy components

6.1 Setting up data storage policies

This section takes you through the steps necessary to define a data storage policy. The minimum steps required are:

- Define a policy domain
- Define a policy set
- Define a management class
- Define a backup copy group
- Define an archive copy group

Tivoli Storage Manager has a default policy, called STANDARD, that is configured when the product is installed. The STANDARD policy is good for product evaluation and testing, but the recommended configuration we have developed for this redbook is a better starting point. It is fully integrated with the rest of the Tivoli Storage Manager server components to get your system up and running as quickly as possible.

6.1.1 Defining a policy domain

A policy domain is a way to group Tivoli Storage Manager clients depending on how you want to treat their data. It allows you to logically group the machines in your organization according to:

- **Default policy:** The default set of rules to apply to the clients. The rules define the storage management policy, including how many copies of data to keep and how long to keep them.
- **Administrative control:** Access to the clients and their policy can be restricted to certain administrators.

Let us consider a typical organization consisting of UNIX and Windows NT machines. The UNIX machines are large file servers and need many copies of their data maintained for a long period of time. The UNIX support group is the only group authorized to access the UNIX machines. The UNIX policy domain would hold all of the UNIX machines, and would only be accessible to the UNIX Tivoli Storage Manager administrators. The default policy in this domain would apply only to the UNIX machines.

The Windows NT machines are workstations and need a few copies of their data maintained for a short period of time. The Windows NT support group is the only group qualified to access the Windows NT machines. The Windows NT policy domain would hold all of the Windows NT machines, and access to it would be restricted to the Windows NT support group. The default policy in this domain would only apply to the Windows NT machines.

This is a good example of how default policy and administrative control can be used to break up your organization into policy domains.

The `define domain` command creates new policy domains. The following command shows how to create a policy domain called `production` which is used to group the production servers for the ABC corporation.

```
tsm: TSM010> define domain production backret=100 archret=365 \  
cont> description="Production Servers for ABC Corporation"  
ANR1500I Policy domain PRODUCTION defined.
```

The Tivoli Storage Manager clients that belong to this domain share the same default policy, and access to them can be restricted to administrators authorized to manage the production servers. The nodes must be assigned to this policy domain when they are registered with the Tivoli Storage Manager server as described in 7.2.3.1, “Registering a client node” on page 149.

The `define domain` command also defines two parameters that are used in special cases to determine how data is handled:

- **BACKRET:** How long backup data is retained if it is not governed by the existing backup policies
- **ARCHRET:** How long archive data is retained if it is not governed by the existing archive policies

These two retention grace period parameters are used as a safety net for data that is ‘orphaned’ from a valid set of policies for some reason. This can happen if

a set of policies were accidentally deleted while still containing data. The BACKRET and ARCHRET parameters are used to manage this data.

Time-saving hint

When you use the `define domain` command, the policy domain is created empty — that is, you need to define the subcomponents (policy set, management classes, and copy groups) separately as described in the following section. However, if you use the `copy domain` command to copy an existing domain, then all its subcomponents are also created in the new domain automatically. This could be a big time saver if you have a number of domains with reasonably similar components. You can create one, then copy it as needed, and use the appropriate `define` commands to make any changes required.

6.1.2 Defining a policy set

The policy set is a grouping of management classes. There can be multiple policy sets within a policy domain, but only one policy set can be active at a time. This restriction has resulted in most installations using only one policy set per policy domain.

The ACTIVE policy set is a special entity in the policy domain. It exists in every domain and cannot be changed directly. To change the ACTIVE policy set, you must define your rules in a policy set which is subsequently validated and activated. The activation process takes a snapshot of the policy set and places it in the ACTIVE policy set. It is important to note that the ACTIVE policy set is a point-in-time snapshot of the originating policy set. Therefore, further changes to the originating policy set have no effect on the ACTIVE policy set unless and until the changed policy set is validated and activated.

The `define policyset` command creates a new policy set within a policy domain. Although the policy set can be named anything except the reserved name ACTIVE, it is common practice to give the same name to the policy domain and policy set for simplicity. The following command shows how to define a PRODUCTION policy set for the PRODUCTION policy domain.

```
tsm: TSM010> define policyset production production \  
cont> description="Production Servers for ABC Corporation"  
ANR1510I Policy set PRODUCTION defined in policy domain PRODUCTION.
```

6.1.3 Defining management classes

A management class contains specific rules and is the link between rules and data. The rules are contained in constructs called copy groups, one for rules governing backup data and one for rules governing archive data. A management class can have a backup copy group or an archive copy group or both.

A policy set may contain many management classes but only one is designated as the default. The default management class in the policy set is linked to any data in the domain that is not explicitly linked to another management class. Don't forget to designate one of your management classes in each policy set as the default, otherwise the `validate policyset` will fail.

Explicitly linking, or binding, files to management classes can be accomplished in three different ways depending on the type of data. Backup data is bound to a management class using the INCLUDE option of the Tivoli Storage Manager client's include-exclude list. Archive data is bound to a management class using the ARCHMC command line option. Directory data can be bound to a management class using the DIRMC client option.

Figure 15 shows an example of an INCLUDE statement that assigns the backup file `/home/admin/redbook.script` to a management class called REDBOOK while allowing the rest of the backup files in `/home/admin` to go to the default management class. These statements appear in the client options files. The binding is actually done during the backup operation. See "Include-exclude options" on page 174 and B.3, "Client options files" on page 297 for more examples of the INCLUDE statement.

```
include /home/admin/.../*
include /home/admin/redbook.script redbook
```

Figure 15. Include option example

Figure 16 shows an example of the ARCHMC command line option. The file `/home/admin/redbook.doc` will be bound to the management class REDBOOKARCHIVE. Without the ARCHMC option, this data would have been bound to the default management class. See 10.3.1.4, "Running archive operations" on page 189 for more examples of the ARCHMC option.

```
dsmc archive -archmc=redbookarchive /home/admin/redbook.doc
```

Figure 16. ARCHMC example

Figure 17 shows an example of the DIRMC option in a client options file. All of the directories from this client will be bound to the directory management class. Make sure that the retention period for directories is set longer than for the data itself. Without the DIRMC option, the directory information would go to the management class with the longest retention period. You can use this option to make your restores quicker by keeping your directory data on a primary disk storage pool instead of on tape. The directory information is the first thing that is rebuilt during a restore and by making it quickly accessible, your restore time will be improved. The recommended policy configuration includes a directory management class for the directory information, and the DIRMC option is specified in the recommended client option sets described in B.3, "Client options files" on page 297.

```
dirmc directory
```

Figure 17. DIRMC example

Windows NT Users:

You could consider not recovering the NTFS security information for large restores, especially if you are recovering to a different domain structure, and therefore, new security rules must be applied anyway.

The `define mgmtclass` command creates a new management class in the policy set and domain. The `assign defmgmtclass` command defines a management class as the default for the domain. The following commands show how to create the CUSTDATA management class in the PRODUCTION policy set within the PRODUCTION policy domain and how to assign it as the default management class.

```
tsm: TSM010> define mgmtclass production production custdata \  
cont> description="Custdata management class for ABC Corporation"  
ANR1520I Management class CUSTDATA defined in policy domain PRODUCTION, set  
PRODUCTION.  
  
tsm: TSM010> assign defmgmtclass production production custdata  
ANR1538I Default management class set to CUSTDATA for policy domain PRODUCTION,  
set PRODUCTION.
```

6.1.4 Defining backup copy groups

The backup copy group is concerned with two logical objects: the file, and the file copy. A file is the actual data on a node, while a file copy is a point-in-time copy of the file. Another way to think of it is that Tivoli Storage Manager contains file copies, and nodes contain files.

A file can be in one of two possible states: existing or deleted. When we talk about an existing file on a node, we mean a file that has been previously backed up and still exists on the node. A deleted file is a file that has been previously backed up and has been deleted from the node. This simple concept is important when discussing data storage rules.

A file copy can be in one of three states: active, inactive, or expired. An active file copy is the most current copy of the file, an inactive file copy is a previous copy of the file, and an expired file copy is a copy to be removed from the Tivoli Storage Manager server. A backup file copy is set to the expired state when it no longer conforms to the rules defined in the backup copy group.

Whether the file exists or is deleted, the file copy always passes through the same states in the same order. A file copy starts out as active, since it is the first copy of the file, and therefore the most current. Once the file changes and we take another file copy, the first file copy changes to inactive because we have a more recent one. Eventually, the first file copy expires based on one of two limits placed on it by our rules: number of copies or retention period.

The number of copies that we set in our rules specifies the total number of file copies to maintain in the Tivoli Storage Manager database. It is important to note that the specified number includes the active file copy. Thus, when we set the number of file copies to three, we are keeping one active copy and a maximum of two inactive copies. When the number of copies is exceeded, the oldest copy is removed from the database.

The retention period that we set in our rules specifies the length of time that we retain inactive file copies. Note that there is no retention period for active file copies; they exist as long as the file exists on the node.

Whether or not the file exists on the node affects which rules are used to expire the file copies. If the file exists, the following two backup copy group parameters are in effect:

- **VEREXISTS:** Specifies the number of file copies, or versions, to keep. This number includes active and inactive file copies.
- **RETEXTRA:** Specifies how long to keep inactive file copies. When a file changes from active to inactive, it is kept for these extra days and then removed. It is important to note that the retention period starts from when the file copy changes to inactive, and not from its original backup date. This parameter can also be set to **NOLIMIT** which means the inactive versions are only expired if superseded by other copies and never from a time limit.

If the file has been deleted, the active file copy is made inactive. At this point, there are only inactive file copies for this data in the Tivoli Storage Manager server, and the following parameters apply:

- **VERDELETED:** Specifies the number of file copies to keep after the file has been deleted. This must be equal to or less than **VEREXISTS**.
- **RETEXTRA:** Specifies how long to keep inactive file copies. This means the same as for existing file with the exception of the very last inactive remaining file copy.
- **REONLY:** Specifies how long to maintain the last file copy of the data. This is the number of days to keep the last remaining copy only, and does not apply to other inactive file copies which are still governed by the **RETEXTRA** parameter. This parameter can also be set to **NOLIMIT** which means the last remaining copy is never expired.

The backup copy group defines five other attributes that control the way that backup data is handled:

- **TYPE:** The **TYPE** parameter is used to differentiate between the two possible types of copy groups. In the case of a backup copy group, it is set to **BACKUP**.
- **DESTINATION:** The backup copy group specifies where to store the data sent to it from backup operations using the **DESTINATION** parameter. The copy group bridges the gap between data files and storage pools as illustrated in Figure 18. The figure shows different types of data flowing through the copy groups and into the storage pools. Note that there is not necessarily a one-to-one relationship between copy groups and storage pools: many copy groups could share the same storage pool. A copy group can only ever point to one storage pool and each backed up file can only be in one management class.

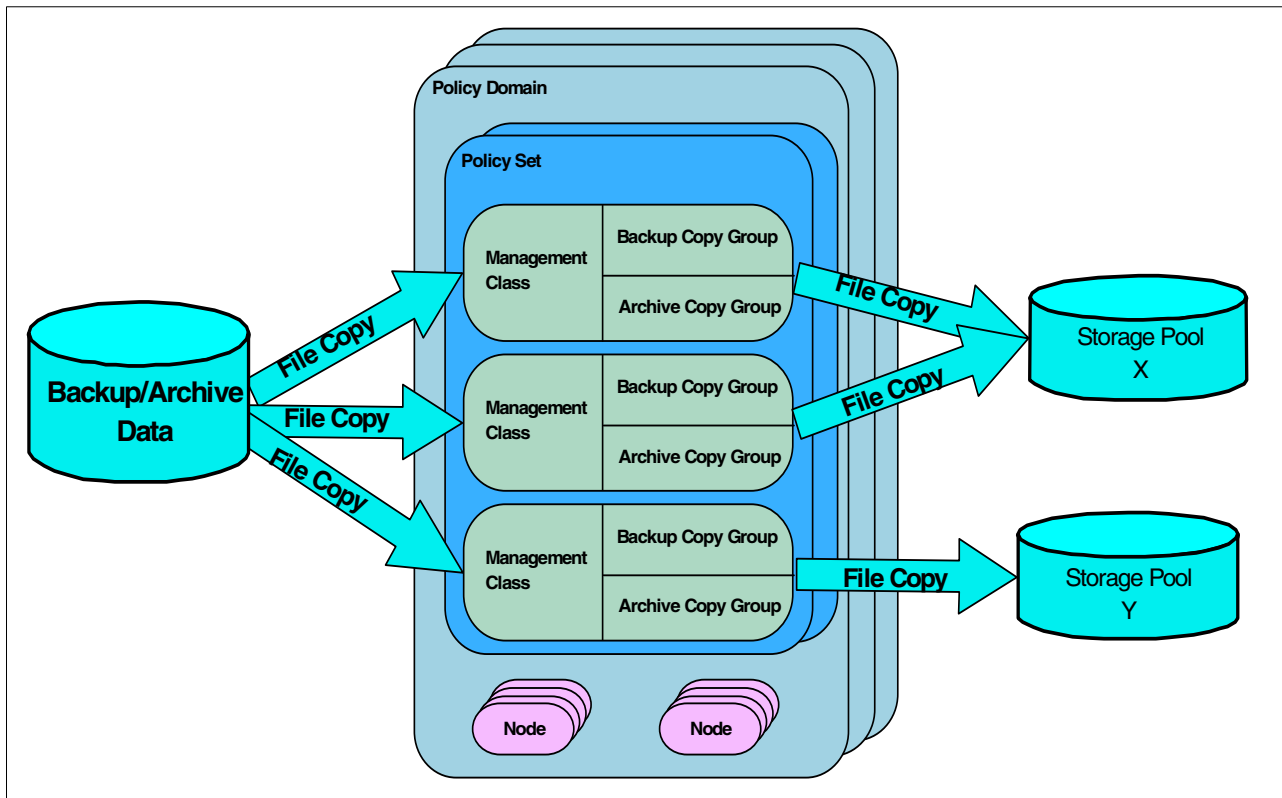


Figure 18. Data flow through copy groups

- **MODE:** The MODE parameter specifies how files are to be selected for incremental backup. Setting the mode to MODIFIED allows a file to be backed up only if it has changed since the last backup. The ABSOLUTE setting allows files to be backed up, regardless of whether they have changed or not. The latter value would only be used for special cases; the default value is MODIFIED in accordance with the progressive incremental backup paradigm of Tivoli Storage Manager.
- **FREQUENCY:** The FREQUENCY parameter specifies how often to allow a file to be incrementally backed up. A client selective backup, which backs up data regardless of whether it has changed or not, is not affected by this parameter. To incrementally back up a file from a node, three conditions must be satisfied:
 1. Include-exclude statements allow the file to be considered for backup.
 2. The file satisfies the MODE setting. That is, if the MODE is set to MODIFIED, the file must have changed to qualify for backup. If the MODE is set to ABSOLUTE, then the file is automatically allowed to be backed up.
 3. The difference between the server time and the active file copy timestamp must be greater than FREQUENCY setting. The frequency is converted to hours to compare to the timestamp difference.

The default value for FREQUENCY is 0 which means there is no minimum elapsed time specified between backups.

For example, take a file called `/home/admin/redbook.doc` that is eligible for backup in the include-exclude list, and that has changed since the last backup at 8 AM

this morning. The server time is 11 AM when an incremental backup is started, so the difference between the server time and the file copy time is three hours. If the frequency is set to one day, then 24 hours must pass between incremental backups before a file is backed up again. Therefore, the `/home/admin/redbook.doc` file is not backed up, since three hours is less than 24 hours.

- **SERIALIZATION:** The **SERIALIZATION** parameter specifies what to do with files that are modified during a backup operation. When we say that a file is modified during backup, we mean that it is modified after Tivoli Storage Manager examined it for its details, but before it was completely backed up to the server. This sort of backup is referred to as a "dirty backup" because the file is in an inconsistent state and may not restore properly. The **SERIALIZATION** parameter provides four options to deal with this issue:
 1. The **SHRSTATIC** setting specifies that a file is not backed up if it is modified during backup, but multiple attempts are made to back up the file. If the file is being modified through all of these attempts, the file is not backed up. The number of attempts can be controlled using the **CHANGINGRETRIES** option in the client options file. So, the file is backed up only if a clean backup can be obtained within the defined number of backup attempts.
 2. The **STATIC** setting specifies that a file is not backed up if it is modified during the first attempt at backup and no additional attempts are made. So, the file is backed up only if a clean backup can be obtained on the first attempt.
 3. The **SHRDYNAMIC** setting specifies that a file is backed up if it is modified during backup, but multiple attempts are made to back it up without modification first. If that cannot be done, then the file is backed up anyway.
 4. The **DYNAMIC** setting specifies that a file is backed up even if it is modified during backup. So, the file is backed up on the first attempt, regardless of whether it changes or not.

You should set these options carefully in your copygroups, driven by whether it will be possible to correctly restore a file which is possibly dirty during the backup. There are two additional points to note here. The first is that if an application locks a file exclusively then Tivoli Storage Manager will not be able to back it up, regardless of the **SERIALIZATION** setting, since it is not available to it or to any other application. The second is that many files which change or are locked by the application are part of a database (for example, Oracle, Informix, SAP R/3 or Microsoft SQL Server) or other application (such as Lotus Notes or Microsoft Exchange). If this is the case for your environment, we strongly recommend that you use an add-on product which will interface the backup of the database or application to Tivoli Storage Manager. This will ensure that open, updated or locked files are intelligently handled for both backup and restore. These products are available both from Tivoli and other vendors. The redbook *Tivoli Storage Management Concepts*, SG24-4877, has more information.

The `define copygroup` command creates a new copy group within the management class. Note that the command does not require a name for the backup copy group; it is always called **STANDARD** and cannot be changed. In any case, there is no reason to change this parameter, since it only has relevance within the Tivoli Storage Manager server itself and there can only ever be one copy group per type (backup or archive) in any management class.

The following command shows how to create a backup copy group for the management class CUSTDATA in the policy set PRODUCTION of the policy domain PRODUCTION. When using this copy group, the client checks for files that have been modified since the last backup process and tries to back them up to the DISKDATA primary storage pool. After a few retries, the process gives up on files that are being modified during backup. The client data is governed by the following rules:

- If the file exists on the client, then three versions of it are kept. The extra inactive versions will expire after 100 days.
- If the file no longer exists on the client, then the last copy is kept for 365 days.

```
tsm: TSM010> define copygroup production production custdata type=backup \  
cont> destination=DISKDATA frequency=1 verexists=3 verdeleted=1 \  
cont> retextra=100 reonly=365 mode=modified serialization=shrstatic  
  
ANR1530I Backup copy group STANDARD defined in policy domain PRODUCTION, set  
PRODUCTION, management class CUSTDATA.
```

6.1.5 Defining an archive copy group

An archive copy group consists of fewer parameters than the backup copy group. The most important of these are:

- **TYPE:** Parameter used to differentiate between two possible types of copy groups. In the case of an archive copy group, it is set to ARCHIVE.
- **DESTINATION:** Destination storage pool to use for data storage. Usually, this is the primary storage pool in a storage pool hierarchy including disk pools and tape pools.
- **RETVR:** How long to keep this archive copy. The default is 365 days, but it can be set to any period or to NOLIMIT.
- **SERIALIZATION:** How to deal with files that are being modified during processing. The options and implications are the same as for a backup copy group. SHRSTATIC means that the process tries multiple times to archive a file that is in use and then gives up. SHRDYNAMIC means that the file is archived up even though it is still in use. This method is not the recommended way, but it should be used for files that are always in use and need to be archived in any case.

The same considerations for backup copy group SERIALIZATION (locked files, applications like Oracle and R/3) apply here.

Archive data is assigned to the archive copy group of the default management class unless you explicitly specify otherwise on the client command line or GUI interface.

The `define copygroup` command is used to define the archive copy group. Note that the command does not require a name for the archive copy group; it is always called STANDARD and cannot be changed. In any case, there is no reason to change this parameter, since it only has relevance within the Tivoli Storage Manager server itself and there can only ever be one copy group per type (backup or archive) in any management class.

The following command shows how to define an archive copy group for the CUSTDATA management class in the PRODUCTION policy set within the PRODUCTION policy domain. The archive files are sent to the DISKDATA primary storage pool, and are kept for 365 days. If the operation cannot archive a file because it is in use, it tries a few times and then gives up.

```
tsm: TSM010> define copygroup production production custdata type=archive \  
cont> destination=DISKDATA retver=365 serialization=shrstatic  
  
ANR1535I Archive copy group STANDARD defined in policy domain PRODUCTION, set  
PRODUCTION, management class CUSTDATA.
```

6.2 Recommended setup

Figure 19 is a graphical representation of the key components of our recommended policy configuration. We define two policy domains, SERVER and WORKSTN. Both domains have *similar* policy sets and management classes, but their copy group details show that the SERVER domain has more copies, and longer retention periods than the WORKSTN domain. It is not necessary to define multiple domains, but this demonstrates that the policy is really defined in the copy groups, and that the rest of the constructs are used primarily for flexibility.

Warning:

Care must be taken if you ever need to migrate a client node to another policy domain. If a client had data backed up under the old domain, the data is rebound to the same management class in the new domain, if it exists. If the management class does not exist in the new domain, the data is rebound to the default management class of the new domain. If this management class does not have a backup copy group, the data is rebound to the retention grace period defined in the new domain. This can lead to undesired expiration of backup versions. This is yet another illustration of the importance of carefully pre-planning your Tivoli Storage Manager environment.

In 1.2, “Redbook support material” on page 2, we show how to load a predefined macro into Tivoli Storage Manager. The macro `mac.policy`, which we provide to define the recommended policy settings in our redbook environment, is shown in B.1.3, “Define policy structure” on page 287.

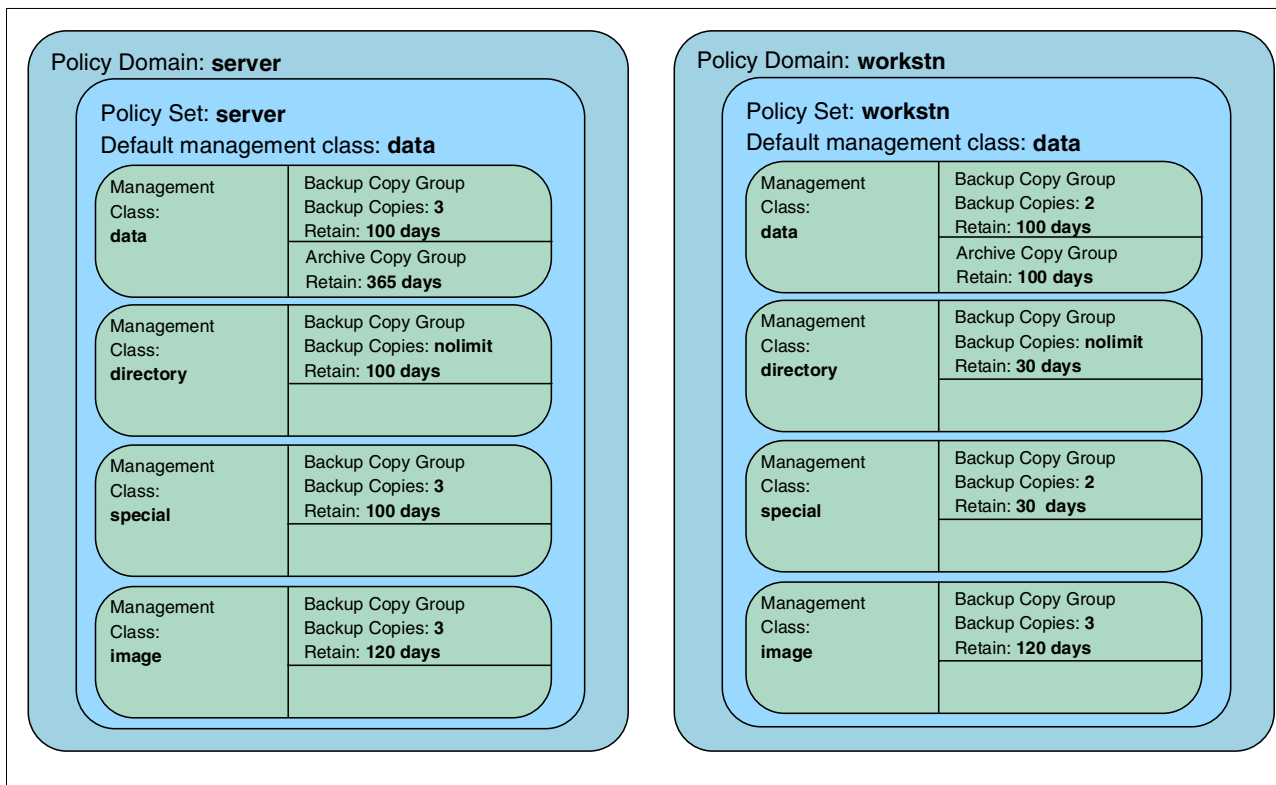


Figure 19. Recommended policy definitions

6.2.1 Defining policy domains

The recommended configuration consists of two policy domains that are used to separate server machines from workstation machines. The nature of the data found on each type of machine makes this a logical division of resources and follows the administrative boundaries used by most customers. This example defines the recommended policy domains:

```
tsm: TSM010> define domain server description="Server nodes" \
cont> backretention=100 archretention=365
ANR1500I Policy domain SERVER defined.

tsm: TSM010> define domain workstn description="Workstation nodes" \
cont> backretention=100 archretention=365
ANR1500I Policy domain WORKSTIN defined.
```

6.2.2 Defining policy sets

We recommend defining one policy set for each of the policy domains; for example:

```
tsm: TSM010> define policyset server server description="Server nodes"
ANR1510I Policy set SERVER defined in policy domain SERVER.

tsm: TSM010> define policyset workstn workstn description="Workstation nodes"
ANR1510I Policy set WORKSTIN defined in policy domain WORKSTIN.
```

6.2.3 Defining management classes

We have defined four management classes for each domain in the recommended configuration. These management classes are used to categorize the domain data:

- **DATA:** This is the default management class for both domains. It is used for most of the data in the environment.
- **DIRECTORY:** Directory structure information is assigned to this class to make it possible to improve restore time.
- **SPECIAL:** This management class is used to store files that are modified during backup. An application log file could be bound to this management class to make sure that a copy is taken. A typical example of this type of file is an application log file that receives messages or errors. While the application is running, it is likely that the file will be modified to additional messages or errors while the scheduled Tivoli Storage Manager backup or archive operations are run. Let's assume it is determined that it is OK to backup or archive this file, even if it is being updated, since the file will still be valid for restore. Since we do not want to stop the application so that the file is released cleanly for backup, we assign it to a management class whose specifications (in the backup/archive copy group definitions) include parameters for treatment of data that is being modified while a backup/archive operation is being performed. One of these options allows the file to be backed up even if it is being modified.
- **IMAGE:** This management class is used to store UNIX image backups, should you have such machines. This is useful if you want to handle the retention period differently for such operations. It is also useful if you further need to create a separate storage pool to store these backup, a tape storage pool for example.

Note:

The Logical Volume backup/restore function operates only with Tivoli Storage Manager client at level 3.7 or higher for AIX 4.2/4.3; SUN Solaris 2.6/2.7; and HP-UX 11.00.

If your mounted file system has nested mount points, unmount them before attempting a back up or restore. Image backups cannot be performed on nested mount points.

Space management clients are not covered in this redbook, but there is a space management parameter in the management class definition that we must change to avoid problems. The MIGDESTINATION parameter specifies the name of the storage pool for the migration of space managed files, and it defaults to a storage pool that we delete during the recommended configuration steps. We point it to a dummy storage pool called NONE, which allows us to validate and activate the policy set without receiving any error messages.

The management classes for the recommended configuration can be defined as follows:

```
tsm: TSM010> define mgmtclass server server data migdestination=NONE \  
cont> description="Default management class for server domain"  
ANR1520I Management class DATA defined in policy domain SERVER, set SERVER.  
  
tsm: TSM010> assign defmgmtclass server server data  
ANR1538I Default management class set to DATA for policy domain SERVER, set  
SERVER.  
  
tsm: TSM010> define mgmtclass server server directory migdestination=NONE \  
cont> description="Directory management class for server domain"  
ANR1520I Management class DIRECTORY defined in policy domain SERVER, set  
SERVER.  
  
tsm: TSM010> define mgmtclass server server special migdestination=NONE \  
cont> description="Special management class for server domain"  
ANR1520I Management class SPECIAL defined in policy domain SERVER, set SERVER.  
  
tsm: TSM010> define mgmtclass server server image migdestination=NONE \  
cont> description="Image management class for server domain"  
ANR1520I Management class IMAGE defined in policy domain SERVER, set SERVER.  
  
tsm: TSM010> define mgmtclass workstn workstn data migdestination=NONE \  
cont> description="Default management class for workstn domain"  
ANR1520I Management class DATA defined in policy domain WORKSTN, set WORKSTN.  
  
tsm: TSM010> assign defmgmtclass workstn workstn data  
ANR1538I Default management class set to DATA for policy domain WORKSTN, set  
WORKSTN.  
  
tsm: TSM010> define mgmtclass workstn workstn directory migdestination=NONE \  
cont> description="Directory management class for workstn domain"  
ANR1520I Management class DIRECTORY defined in policy domain WORKSTN, set  
WORKSTN.  
  
tsm: TSM010> define mgmtclass workstn workstn special migdestination=NONE \  
cont> description="Special management class for workstn domain"  
ANR1520I Management class SPECIAL defined in policy domain WORKSTN, set  
WORKSTN.  
  
tsm: TSM010> define mgmtclass workstn workstn image migdestination=NONE \  
cont> description="Image management class for workstn domain"  
ANR1520I Management class IMAGE defined in policy domain WORKSTN, set WORKSTN.
```

6.2.4 Defining backup copy groups

The reasons for creating two domains (WORKSTN and SERVER) are clearly illustrated in the backup copy group parameters. The two domains differ in the following ways:

- VEREXISTS: The WORKSTN domain only maintains two copies of existing data, as compared to three copies in the SERVER domain.
- RETEXTRA: The WORKSTN domain only keeps extra copies for 30 days, as compared to 100 days in the SERVER domain.
- RETONLY: The WORKSTN domain keeps the last copy of deleted data for 100 days, as compared to 365 days in the SERVER domain

An interesting parameter also illustrates the difference between the management classes:

- SERIALIZATION: The backup copy group within the special management class is set up to use SHR DYNAMIC instead of SHR STATIC. This management class should be assigned to log files and other cases where the files are always open.

- The IMAGE management class requires a SERIALIZATION setting as STATIC to work properly. This means that whenever an image backup operation is performed on a logical volume containing a filesystem, the client will attempt to unmount the filesystem first. This is the recommended method.

Note:

If some of your data require policy settings different than our recommended defaults—for example, if you need to keep 7 copy versions of certain files—then you need to tailor an additional management class for your purpose, and assign the data to it using the INCLUDE option. See 10.2.1.2, “Options file” on page 172 for further information.

We define a backup copy group for each of the management classes using the following commands:

```
tsm: TSM010> define copygroup server server data type=Backup destination=DISKDATA \
cont> frequency=1 verexists=3 verdeleted=1 retextra=100 retonly=365 \
cont> mode=modified serialization=shrstatic
ANR1530I Backup copy group STANDARD defined in policy domain SERVER, set
SERVER, management class DATA.

tsm: TSM010> define copygroup server server directory type=Backup destination=DISKDIRS \
cont> frequency=1 verexists=nolimit verdeleted=1 retextra=100 retonly=365
cont> mode=modified serialization=shrstatic
ANR1530I Backup copy group STANDARD defined in policy domain SERVER, set
SERVER, management class DIRECTORY.

tsm: TSM010> define copygroup server server special type=Backup destination=DISKDATA \
cont> frequency=1 verexists=3 verdeleted=1 retextra=100 retonly=365 \
cont> mode=modified serialization=shrdynamic
ANR1530I Backup copy group STANDARD defined in policy domain SERVER, set
SERVER, management class SPECIAL.

tsm: TSM010> define copygroup server server image type=backup destination=diskdata \
cont> frequency=1 verexists=3 verdeleted=1 retextra=120 retonly=365 \
cont> mode=modified serialization=static
ANR1530I Backup copy group STANDARD defined in policy domain SERVER, set
SERVER, management class IMAGE.

tsm: TSM010> define copygroup workstn workstn data type=Backup \
cont> destination=DISKDATA frequency=1 verexists=2 verdeleted=1 retextra=30 \
cont> retonly=100 mode=modified serialization=shrstatic
ANR1530I Backup copy group STANDARD defined in policy domain WORKSTN, set
WORKSTN, management class DATA.

tsm: TSM010> define copygroup workstn workstn directory type=Backup \
cont> destination=DISKDIRS frequency=1 verexists=nolimit verdeleted=1 retextra=30 \
cont> retonly=100 mode=modified serialization=shrstatic
ANR1530I Backup copy group STANDARD defined in policy domain WORKSTN, set
WORKSTN, management class DIRECTORY.

tsm: TSM010> define copygroup workstn workstn special type=Backup \
cont> destination=DISKDATA frequency=1 verexists=2 verdeleted=1 retextra=30 \
cont> retonly=100 mode=modified serialization=shrdynamic
ANR1530I Backup copy group STANDARD defined in policy domain WORKSTN, set
WORKSTN, management class SPECIAL.

tsm: TSM010> define copygroup workstn workstn image type=backup \
cont> destination=diskdata frequency=1 verexists=3 verdeleted=1 retextra=120 \
cont> retonly=365 mode=modified serialization=static
ANR1530I Backup copy group STANDARD defined in policy domain WORKSTN, set
WORKSTN, management class IMAGE.
```

6.2.5 Defining the archive copy group

For the recommended configuration, we define archive copy groups only for the default management classes. The data management class in the SERVER and

WORKSTN policy domains keeps archive copies for 365 days, and does not archive files that are unavailable after a few retries. If it later turned out that we had additional archive requirements, for example to keep certain files for seven years, we could easily define a suitable archive copy group in an existing or new management class.

```
tsm: TSM010> define copygroup server server data type=Archive destination=DISKDATA \  
cont> retver=365 serialization=shrstatic  
ANR1535I Archive copy group STANDARD defined in policy domain SERVER, set  
SERVER, management class DATA.  
  
tsm: TSM010> define copygroup workstn workstn data type=Archive \  
cont> destination=DISKDATA retver=100 serialization=shrstatic  
ANR1535I Archive copy group STANDARD defined in policy domain WORKSTN, set  
WORKSTN, management class DATA.
```

6.2.6 Deleting the STANDARD policy domain

You should remove the STANDARD domain information to clean up your environment. The `delete domain` command removes an existing policy domain. This command is very powerful in that it also removes the policy sets, management classes, and copy groups that belong to the domain. For this reason, the `delete domain` command should be used with caution. The following command shows how to remove a policy domain named STANDARD.

```
Tivoli Storage Manager> delete domain standard  
  
Do you wish to proceed? (Yes/No) yes  
ANR1501I Policy domain STANDARD deleted.
```

If policy domain STANDARD contains nodes and filespaces, this command fails. You must remove all registered nodes from a policy domain before deleting it. There are two ways to accomplish this task:

- Delete the nodes and their filespaces
- Move the nodes including their filespaces to another domain

6.2.6.1 Deleting a node

Deleting a node involves removing the node data with the `delete filespace` command, deleting any backupsets registered to the server with the `delete backupset` command and removing the node definition with the `remove node` command. The following commands show how to remove the filespaces and node definition for the node PALANA.

```

tsm: TSM010> delete filesystem palana *
ANR2238W This command will result in the deletion of all inventory references
to the data on filesystems that match the pattern * for node PALANA, whereby
rendering the data unrecoverable.

Do you wish to proceed? (Yes/No) y
ANS8003I Process number 35 started.

tsm: TSM010> delete backupset palana *
ANR2365I Backupset ARMANDO_BS.20548 for node PALANA has been deleted.

tsm: TSM010> remove node palana

Do you wish to proceed? (Yes/No) y
ANR2061I Node PALANA removed from policy domain STANDARD.
ANR2129I Administrative userid PALANA defined ONLY for authority over node
PALANA has been removed.

```

6.2.6.2 Moving a node

The `update node` command can be used to move a node to a new domain. The data is moved to the new domain during the next backup operation. The process of moving the data to the new domain's management classes is called rebinding. If the new domain has the same management class names as the old one, very little change takes place, as the data is affected more by the management class name than the policy domain name. However, if the new domain has completely different management class names, then the data is rebound to the default management class for the domain. The first backup after the domain change may take longer than normal, but it is a one time operation and returns to normal speed thereafter.

Warning:

Care must be taken when migrating to another policy domain. If a client had data backed up under the old domain, the data is rebound to the same management class of the new domain, if it exists. If the management class does not exist in the new domain, the data is rebound to the default management class of the new domain. If this management class does not have a backup copy group, the data is rebound to the retention grace period defined in the new domain. This can lead to undesired expiration of backup versions.

The following command shows how to move the node `devnode1` from its current domain to the `WORKSTN` domain.

```

tsm: TSM010> update node devnode1 domain=workstn
ANR2063I Node DEVNODE1 updated.

tsm: TSM010> query node devnode1

```

Node Name	Platform	Policy Domain Name	Days Since Last Access	Days Since Password Set	Locked?
DEVNODE1	(?)	WORKSTN	<1	<1	No

6.3 Verifying policy definitions

The best way to check the policy definitions is to examine the details of the copy groups using the `query copygroup` command. The output from this command

displays the policy domain, policy set, management class, and copy group names, as well as the copy group parameters. Note that it is really the copy group definitions that define the policy for the domain; the rest of the constructs between domain and copy group just provide flexibility in your configuration.

6.3.1 Backup copy groups

To check the settings for the recommended backup copy groups, use the query copygroup command to get the attributes displayed:

```
tsm: TSM010> query copy server server type=backup
```

Policy Domain Name	Policy Set Name	Mgmt Class Name	Copy Group Name	Versions Data Exists	Versions Data Deleted	Retain Extra Versions	Retain Only Version
SERVER	SERVER	DATA	STANDARD	3	1	100	365
SERVER	SERVER	DIRECTORY	STANDARD	No Limit	1	100	365
SERVER	SERVER	IMAGE	STANDARD	3	1	120	365
SERVER	SERVER	SPECIAL	STANDARD	3	1	100	365

```
tsm: TSM010> query copy workstn workstn type=backup
```

Policy Domain Name	Policy Set Name	Mgmt Class Name	Copy Group Name	Versions Data Exists	Versions Data Deleted	Retain Extra Versions	Retain Only Version
WORKSTN	WORKSTN	DATA	STANDARD	2	1	30	100
WORKSTN	WORKSTN	DIRECTORY	STANDARD	No Limit	1	30	100
WORKSTN	WORKSTN	IMAGE	STANDARD	3	1	120	365
WORKSTN	WORKSTN	SPECIAL	STANDARD	2	1	30	100

6.3.2 Archive copy groups

To check the settings for the recommended archive copy groups, use the query copygroup command to get a quick look at the attributes:

```
tsm: TSM010> query copygroup server server type=archive
```

Policy Domain Name	Policy Set Name	Mgmt Class Name	Copy Group Name	Retain Version
SERVER	SERVER	DATA	STANDARD	365

```
tsm: TSM010> quer copygroup workstn workstn type=archive
```

Policy Domain Name	Policy Set Name	Mgmt Class Name	Copy Group Name	Retain Version
WORKSTN	WORKSTN	DATA	STANDARD	100

6.4 Validating and activating a policy set

The last step in setting up your policy is to validate and activate your policy set. The commands are very straightforward and have few parameters.

6.4.1 Validating the recommended policy sets

The `validate policyset` command checks for completeness in the management class and copy group definitions. It validates these policies and makes them ready for activation.

To validate the recommended policy sets:

```
tsm: TSM010> validate policyset server server
ANR1515I Policy set SERVER validated in domain SERVER (ready for activation).

tsm: TSM010> validate policyset workstn workstn
ANR1515I Policy set WORKSTN validated in domain WORKSTN (ready for
activation).
```

6.4.2 Activating the recommended policy sets

The `activate policyset` command allows the specified policy set to be applied to the data within its policy domain.

The following example shows how to activate the recommended policy sets:

```
tsm: TSM010> activate policyset server server

Do you wish to proceed? (Yes/No) y
ANR1514I Policy set SERVER activated in policy domain SERVER.

tsm: TSM010> activate policyset workstn workstn

Do you wish to proceed? (Yes/No) y
ANR1514I Policy set WORKSTN activated in policy domain WORKSTN.

tsm: TSM010> query domain
```

Policy Domain Name	Activated Policy Set	Activated Default Mgmt Class	Number of Registered Nodes	Description
SERVER	SERVER	DATA	0	Server nodes
WORKSTN	WORKSTN	DATA	1	Workstation nodes

6.5 Enforcing your policy

Once the policy is defined, you want to enforce it using the `expire inventory` command. This command makes sure that any extra copies of data in your copy groups are removed from the database. It also takes care of data that is older than your specified retention period. The backup copy group parameters `VEREXISTS`, `VERDELETED`, `RETEXTRA`, and `RETONLY` are applied to backup data through the expiration process as well as the archive copy group parameter `RETVER`. The following command shows how to expire the database references for the Tivoli Storage Manager server.

```
tsm: TSM010> expire inventory

ANS8003I Process number 36 started.
```

This command can be computationally intensive; therefore, it should only be executed during quieter times on the Tivoli Storage Manager server.

The EXPINTERVAL parameter in the server options file specifies the number of hours between automatic expiration processing, and is initially configured for 24 hours. With automatic expiration processing enabled, the server runs inventory expiration at start-up and every 24 hours thereafter. We recommend setting the EXPINTERVAL to zero (disabling automatic expiration) and defining a daily administrative schedule to run this command at a convenient time.

Chapter 7. User management

In this chapter, we explain the creation and maintenance of users of a Tivoli Storage Manager system. There are two categories of users:

- Administrators
- Client nodes

Administrators manage Tivoli Storage Manager resources. They also administer Tivoli Storage Manager client nodes. One way of maintaining client nodes is to centrally define options that the clients will use during backup, restore, archive, and retrieve operations. This is done using client option sets.

7.1 Administrators

A Tivoli Storage Manager administrator manages Tivoli Storage Manager resources. The number of administrators and their level of privileges will vary according to your environment. Administrators with system privilege can perform any Tivoli Storage Manager function. Administrators with policy, storage, operator, or analyst privileges can perform subsets of Tivoli Storage Manager functions. All commands issued by administrators are logged to the server activity log. See 2.6, “Administrator IDs” on page 42 for planning considerations.

7.1.1 Considerations

The creation of an administrator is a two-step process:

1. Define the administrator ID.
2. Grant the necessary privileges.

The `register admin` command defines administrators. The only required parameters are the admin name, or user ID, and initial password. You should consider using the optional `CONTACT` parameter to distinguish administrators.

The `grant authority` command grants the necessary privileges to an administrator. The only required parameters are the admin name and the privileges being granted.

The `query admin` command displays information about one or more administrators.

7.1.2 Default environment

During the initial server startup, two administrators are automatically defined: `SERVER_CONSOLE` and `ADMIN`.

`SERVER_CONSOLE` is a special administrator associated with the server console that has system privileges. When the server is started in interactive mode, typically with the `dsmserv` command, a server console is started with a command prompt running with the assigned authority of the `SERVER_CONSOLE` ID. This is in contrast to running the server in background mode, either as an NT service or under UNIX using the `nohup` command. In that case, no command prompt is presented. You cannot update, lock, rename, or remove the `SERVER_CONSOLE` user ID from Tivoli Storage Manager. The `SERVER_CONSOLE` user ID does not have a password and cannot be given

one. Although you can reduce the authority of this administrator, we recommend that you do not. In an emergency situation, administrative commands can be issued from the server console to correct situations, such as forgotten administrator passwords and locked IDs. This assumes your server console runs in a secure environment. If you cannot secure your server console from unauthorized access, you should reduce the privilege of the SERVER_CONSOLE ID, perhaps to analyst. Then, you must carefully protect your administrator IDs which have system privilege. We recommend having at least two systems administrators. Tivoli Storage Manager will not allow you to remove either the SERVER_CONSOLE ID or the last administrator ID with system privilege.

ADMIN is an administrator that has system privileges and an initial password of ADMIN. This administrator is used to set up your Tivoli Storage Manager environment. For enhanced security after your environment has been created, we recommend that you delete this administrator.

7.1.3 Recommended administrators

We recommend that you define the following administrators to support your Tivoli Storage Manager environment:

- System
- Support
- Reporting
- Client

We recommend that you use the optional CONTACT parameter to distinguish administrators. Although we have recommended certain administrator names, you do not have to use them. You can use any name suitable for your purposes.

In 1.2, “Redbook support material” on page 2, we show how to load a predefined macro into Tivoli Storage Manager. The macro `mac.admins`, which we provide to define administrative IDs for our redbook environment, is shown in B.1.1, “Define administrators” on page 285.

7.1.3.1 System

Define an ID for the Tivoli Storage Manager administrator and possibly another for your backup Tivoli Storage Manager administrator. These administrators should have system privileges.

In our example, we define an administrator named SYSADMIN, and a backup named SYSADMIN2. The initial password for both administrators is AMANDA. The administrative commands to create our recommended administrators, and the results of those commands, should look like this:

```

tsm: TSM010> register admin sysadmin amanda contact='Tivoli Storage Manager Administrator'
ANR2068I Administrator SYSADMIN registered.

tsm: TSM010> grant authority sysadmin classes=system
ANR2076I System privilege granted to administrator SYSADMIN.

tsm: TSM010> register admin sysadmin2 amanda \
cont> contact='Tivoli Storage Manager Administrator (Alternate) '
ANR2068I Administrator SYSADMIN2 registered.

tsm: TSM010> grant authority sysadmin2 classes=system
ANR2076I System privilege granted to administrator SYSADMIN2.

```

7.1.3.2 Support

Define an ID for your Tivoli Storage Manager technical support person. This ID should have system privilege.

In our example, we define a support administrator named SUPPORT. The initial password for this administrator is KIRSTEN. The administrative commands to create our recommended administrator and the results of those commands should look like this:

```

tsm: TSM010> register admin support kirsten contact='Tivoli Storage Manager Support'
ANR2068I Administrator SUPPORT registered.

tsm: TSM010> grant authority support classes=system
ANR2076I System privilege granted to administrator SUPPORT.

```

7.1.3.3 Reporting

Define an ID for Tivoli Storage Manager reporting purposes. Although this administrator has no special privileges, it does have authority to issue QUERY commands. This is especially useful for producing regular reports through a method such as scripts. Disclosure of this administrator's password does not create a security exposure, because the administrator cannot change any Tivoli Storage Manager resources.

In our example, we define a support administrator named REPORTER. The initial password for this administrator is CAMERON. The administrative command to create our recommended administrator and the results of that command should look like this:

```

tsm: TSM010> register admin reporter cameron contact='Tivoli Storage Manager Reporting'
ANR2068I Administrator REPORTER registered.

```

7.1.3.4 Client

Define an ID for the Web administrative client access. This administrator would have client access authority to all nodes in the Tivoli Storage Manager environment. In particular, this administrator can perform restore operations on behalf of users. This is usually sufficient for smaller organizations. In larger organizations, you might require multiple administrators with this authority.

In our example, we define a client administrator named HELPDESK. The initial password for this administrator is ERIC2000. The administrative commands to

create our recommended administrator and the results of those commands should look like this:

```
tsm: TSM010> register admin helpdesk eric2000 \  
cont> contact='Tivoli Storage Manager Client Administrator'  
ANR2068I Administrator HELPDESK registered.  
  
tsm: TSM010> grant authority helpdesk classes=node node=*  
ANR2126I GRANT AUTHORITY: Administrator HELPDESK was granted ACCESS authority  
for client DEVNODE1.  
ANR2126I GRANT AUTHORITY: Administrator HELPDESK was granted ACCESS authority  
for client PAGOPAGO.  
ANR2126I GRANT AUTHORITY: Administrator HELPDESK was granted ACCESS authority  
for client FIJI.
```

The administrator is only granted authority to those nodes that exist at the time the command is issued. If you create additional nodes and you want administrators to have authority to those nodes, the `grant authority` command must be reissued.

7.1.4 Working with administrators

In this section we explain how to perform various processes related to administrators that you will find useful in your Tivoli Storage Manager environment:

7.1.4.1 Displaying administrators

You use the `query admin` command to display administrators. If no administrator name is specified, all administrators are displayed. If an administrator name is specified, only that administrator's information is displayed.

The administrative command to display the list of all administrators and the results of that command should look similar to this:

```
tsm: TSM010> query admin
```

Administrator Name	Days Since Last Access	Days Since Password Set	Locked?	Privilege Classes
ADMIN	<1	17	No	System
DEVNODE1	<1	<1	No	Client Owner
HELPDESK	<1	<1	No	Client Access
REPORTER	<1	<1	No	
SERVER_CONSOLE			No	System
SUPPORT	<1	<1	No	System
SYSADMIN	<1	<1	No	System
SYSADMIN2	<1	<1	No	System

The administrative command to display full details of our administrator named HELPDESK and the results of that command should look similar to this:

```

tsm: TSM010> query admin helpdesk format=detail

Administrator Name: HELPDESK
Last Access Date/Time: 02/10/00 10:41:50
Days Since Last Access: <1
Password Set Date/Time: 02/10/00 10:41:50
Days Since Password Set: <1
Invalid Sign-on Count: 0
Locked?: No
Contact: Tivoli Storage Manager Client Administrator
System Privilege:
Policy Privilege:
Storage Privilege:
Analyst Privilege:
Operator Privilege:
Client Access Privilege: DEVNODE1 PAGOPAGO FLJI
Client Owner Privilege:
Registration Date/Time: 02/10/00 10:41:50
Registering Administrator: ADMIN
Managing profile:
Password Expiration Period:

```

7.1.4.2 Changing an administrator password

An administrative user can change his or her own password using the `update admin` command. A user with system privilege can update any administrator's password. There is no way of determining the current password for an administrator — if it is forgotten or lost then it can be reset by another administrator.

The administrative command to change the password for the administrator named `SYSADMIN2` to `PEDRO2000` should look like this:

```

tsm: TSM010> update admin sysadmin2 pedro2000
ANR2071I Administrator SYSADMIN2 updated.

```

When this command is logged to the server activity log, the password in the log is replaced by asterisks.

7.1.4.3 Deleting the ADMIN administrator

After you have defined the set of system administrators, delete the installed system administration ID, `ADMIN`. Because you cannot delete an administrator if that administrator is currently accessing the server, you must logon to Tivoli Storage Manager as a different administrator to perform this task. The administrative session, the administrative command, and the results of that command should look similar to this:

```

root@ / > dsmdmc
Tivoli Storage Manager
Command Line Administrative Interface - Version 3, Release 7, Level 1.0
(C) Copyright IBM Corporation, 1990, 1999, All Rights Reserved.

Enter your user id: sysadmin

Enter your password:*****

Session established with server TSM010: AIX-RS/6000
Server Version 3, Release 7, Level 1.0
Server date/time: 02/10/00 10:49:33 Last access: 02/10/00 10:34:26

tsm: TSM010> remove admin admin

Do you wish to proceed? (Yes/No) y
ANR2069I Administrator ADMIN removed.

```

As we said before, you cannot remove the last administrator with system privileges.

7.2 Client nodes

A Tivoli Storage Manager client must be registered with the Tivoli Storage Manager server before any backup and recovery operations can be performed for that client. Although Tivoli Storage Manager provides two modes, open and closed, for registering client nodes to the server, we recommend for security reasons that you do not change the closed mode default.

7.2.1 Considerations

The registration of a client node is a single-step process using the `register node` command. Registration requires a Tivoli Storage Manager node name and a client access password as a minimum. You should use the machine name for the Tivoli Storage Manager node name. If you define your own Tivoli Storage Manager policy domains, then you must also use the `DOMAIN` parameter. Further, we recommend that you also specify the `USERID` parameter.

The `DOMAIN` parameter specifies the name of the policy domain to which the node is assigned. If you do not specify a policy domain name, the node is assigned to the default policy domain (`STANDARD`). Within our recommended environment, there are two domains, `SERVER` and `WORKSTN`, and no `STANDARD` domain. Consequently, if you use our environment, you must specify the `DOMAIN` parameter when registering a client node.

The `USERID` parameter specifies the name of an ID who will be defined as an administrator with node owner authority to the client. Node authority allows an administrator to use the Web backup-archive client. When the `USERID` parameter is omitted, an administrator with the same name as the node is defined. This is the default. We recommend that you specify `USERID=NONE` for our environment and grant node authority to the `HELPDESK` administrator and other administrators explicitly through the `grant authority` command.

7.2.2 Default environment

Except for MVS, one client node is registered with the Tivoli Storage Manager server during the initial server startup. This client node is the Tivoli Storage

Manager server itself. The name of this client is CLIENT. The initial password is CLIENT. The node is assigned to the STANDARD domain. To comply with our recommendation of using the hostname for each client node, you should use the `rename node` command to change the name of the node from CLIENT to the hostname, and the `update node` command to assign it to a new policy domain.

7.2.3 Working with client nodes

In this section, we explain how to perform various processes related to client nodes that you will find useful in your Tivoli Storage Manager environment.

7.2.3.1 Registering a client node

You use the `register node` command to define a client to the server. In our environment, there are parameters (node name, node password, domain name) which you must supply, and an additional two parameters (contact, userid) which you should supply to register a client node.

Table 30 shows a layout you can use for documenting the information needed to register the clients, completed with a set of sample client information. We are using the domains defined in 6.2.1, “Defining policy domains” on page 134.

Table 30. Client requirements worksheet: example

Node name	Contact	Domain
PAGOPAGO	BRUNO M-6	SERVER
COCOS	phil@home	WORKSTN
FIJI	armando@iss	SERVER
PUTNEY	Ops x222	SERVER

Assuming a common password of LALACA, the commands to define this set of client nodes and resulting output should look like this:

```

tsm: TSM010> register node pagopago lalaca domain=server userid=none \
cont> contact="BRUNO M-6"
ANR2060I Node PAGOPAGO registered in policy domain SERVER.

tsm: TSM010> register node cocos lalaca domain=workstn userid=none \
cont> contact="phil@home"
ANR2060I Node COCOS registered in policy domain WORKSTN.

tsm: TSM010> register node fiji lalaca domain=server userid=none \
cont> contact="armando@iss"
ANR2060I Node FIJI registered in policy domain SERVER.

tsm: TSM010> register node putney lalaca domain=server userid=none \
cont> contact='Ops x222'
ANR2060I Node PUTNEY registered in policy domain SERVER.

```

7.2.3.2 Displaying client nodes

You use the `query node` command to display client nodes. If you do not specify a node name, all nodes are displayed. If you do specify a node name, only that node's information is displayed.

The administrative command to display the list of all nodes and the results of the command should look similar to this:

```
tsm: TSM010> query node
```

Node Name	Platform	Policy Domain Name	Days Since Last Access	Days Since Password Set	Locked?
COCOS	WinNT	WORKSTN	<1	<1	No
DEVNODE1	(?)	WORKSTN	<1	<1	No
FIJI	AIX	SERVER	<1	<1	No
PAGOPAGO	WinNT	SERVER	<1	<1	No
PUINEY	AIX	SERVER	<1	<1	No

The administrative command to display full details of the node named PAGOPAGO, and the results of that command, should look similar to this:

```
tsm: TSM010> query node pagopago format=detail
```

```
Node Name: PAGOPAGO
Platform: WinNT
Client OS Level: 4.00
Client Version: Version 3, Release 7, Level 1.0
Policy Domain Name: SERVER
Last Access Date/Time: 02/10/00 11:05:02
Days Since Last Access: <1
Password Set Date/Time: 02/10/00 10:58:16
Days Since Password Set: <1
Invalid Sign-on Count: 0
Locked?: No
Contact: BRUNO T-4
Compression: Client's Choice
Archive Delete Allowed?: Yes
Backup Delete Allowed?: No
Registration Date/Time: 02/10/00 10:58:16
Registering Administrator: SYSADMIN
Last Communication Method Used: Tcp/Ip
Bytes Received Last Session: 232
Bytes Sent Last Session: 339
Duration of Last Session: 30.78
Pct. Idle Wait Last Session: 99.73
Pct. Comm. Wait Last Session: 0.00
Pct. Media Wait Last Session: 0.00
Optionset:
URL:
Node Type: Client
Password Expiration Period:
Keep Mount Point?: No
Maximum Mount Points Allowed: 1
```

7.2.3.3 Granting access to client nodes

An administrator can perform all client operations on behalf of the user through the web backup/archive client. The administrator needs node authority to perform those operations. An administrative user with system privilege can grant that authority.

The administrative command to grant node authority to a client administrator named HELPDESK, and the results of that command, should look like this:

```

tsm: TSM010> grant authority helpdesk classes=node node=*
ANR2126I GRANT AUTHORITY: Administrator HELPDESK was granted ACCESS authority
for client COCOS.
ANR2126I GRANT AUTHORITY: Administrator HELPDESK was granted ACCESS authority
for client DEVNODE1.
ANR2126I GRANT AUTHORITY: Administrator HELPDESK was granted ACCESS authority
for client FIJI.
ANR2126I GRANT AUTHORITY: Administrator HELPDESK was granted ACCESS authority
for client PAGOPAGO.
ANR2126I GRANT AUTHORITY: Administrator HELPDESK was granted ACCESS authority
for client PUTNEY.

```

7.2.3.4 Changing passwords for client nodes

An administrative user with either system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned can change the password of a client node.

The administrative command to change the password for the client node named COCOS to ISLAND and the results of that command should look like this:

```

tsm: TSM010> update node cocos island
ANR2063I Node COCOS updated.

```

Note:

You can only change a node password if that node is not currently accessing the server.

7.2.3.5 Deleting filespaces for a client node

An administrative user with either system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned, can delete node filespaces. Filespaces contain a client's backup and archive data. There are one or more filespaces for each client.

The `delete filepace` command creates a server process that deletes one or more file spaces as a series of batch database transactions, therefore preventing a rollback or commit for an entire file space as a single action. If the `delete filepace` process is canceled or if a system failure occurs, a partial deletion can occur. In this case, a subsequent `delete filepace` command for the same node can delete the remaining data. You must delete all backup and archive file spaces that belong to a client node before you can delete that client node.

The administrative command to query the filespaces that belong to the client named COCOS, and the results of that command, should look similar to this:

```

tsm: TSM010> query filepace cocos

```

Node Name	Filespace Name	Platform	Filespace Type	Capacity (MB)	Pct Util
COCOS	\\a78axfwy- \c\$	WinNT	NTFS	4,008.4	27.0

The administrative command to delete the filespace \\a78axfwy\c\$ for the client node named COCOS, and the results of that command, should look similar to this:

```
tsm: TSM010> delete filespace cocos \\a78axfwy\c$
ANR2238W This command will result in the deletion of all inventory references
to the data on filespace that match the pattern \\a78axfwy\c$ for node COCOS,
whereby rendering the data unrecoverable.

Do you wish to proceed? (Yes/No) y
ANS8003I Process number 37 started.
```

Filespace deletions can run as either a background or foreground process. The previous example was run as a background process. You can monitor the completion of that process through the `query process` command. You can view messages from that process on the server activity log.

You can run filespace deletion in the foreground by adding the `WAIT=YES` parameter to the command. Foreground processes write messages directly to your session. If your filespace contains a large number of files, we recommend that you run the process in the background.

The administrative command to delete all filespace for the client node named COCOS in the foreground, and the results of that command, should look similar to this:

```
tsm: TSM010> delete filespace cocos * wait=yes
ANR2238W This command will result in the deletion of all inventory references
to the data on filespace that match the pattern * for node COCOS, whereby
rendering the data unrecoverable.

Do you wish to proceed? (Yes/No) y
ANR0984I Process 38 for DELETE FILESPACE started in the FOREGROUND at
11:21:27.
ANR0802I Delete Filespace * (backup/archive data) for node COCOS started.
ANR0800I DELETE FILESPACE: * for node COCOS started as process 38.
ANR0806I Delete Filespace * complete for node COCOS: 9 objects deleted.
ANR0987I Process 38 for DELETE FILESPACE running in the FOREGROUND processed 9
items with a completion state of SUCCESS at 11:21:27.
```

7.2.3.6 Deleting a client node

An administrative user with either system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned can delete a client node. Before you can remove a client node, you must delete all backup and archive file spaces that belong to that client node.

The administrative command to delete the client node named COCOS, and the results of that command, should look similar to this:

```
tsm: TSM010> remove node cocos

Do you wish to proceed? (Yes/No) y
ANR2061I Node COCOS removed from policy domain WORKSTN.
```

7.3 Client option sets

A Tivoli Storage Manager client session has a set of options which are used during the backup, archive, restore, or retrieve processes. Options can be specified in two ways:

1. In the client options file, this is mandatory during the set up of a client. A client options file is a set of Tivoli Storage Manager client options stored in one or two (UNIX clients only) files on the client.
2. Using the client option set, this is optional. A client option set is a set of Tivoli Storage Manager client options stored in the Tivoli Storage Manager database. An option set can be associated with one or more Tivoli Storage Manager clients, but a Tivoli Storage Manager client can be associated with only one option set.

7.3.1 Considerations

The options defined in a client option set are a subset of the available client options. Options such as communications are still stored on the client machine. When the same individual option is specified in both the local options file and the option set, the default is that the options file version is used. However, you can specify that individual options in an option set cannot be overridden in the client's local option file. Although include-exclude specifications cannot be overridden, you can specify the sequence in which the option set specifications are processed. Therefore, one set of default values can be defined for each type of client, and the client machines can still be customized, within acceptable limits.

We recommend that you use client option sets for ease of administration. Management of the environment is complex where the number of clients is growing and the number of options is increasing. The use of client option sets eases that administrative burden by centralizing the management of those options and clients. It is easier to update a client option set once, than to perform the same update to the local client options file on each node.

When you specify include-exclude options in both an options file and an option set, understanding how they relate to each other is critical to understanding what actually happens. The sequence is not obvious.

Include-exclude options in the client option set are additive options and will not override the include-exclude options in the client options file. The sequence number in the client option set determines the order in which the include-exclude statements are added to the existing include-exclude statements logically at the end of the include-exclude statements of the options file, in sequence order. Therefore, include-exclude options in a client option set will always be seen and processed during backup first, before any options set in the client options file.

For example, suppose you have the following specifications in a client option set:

```
include c:\test\* seq=1
exclude c:\working\* seq=2
```

And, suppose you have the following exclude statement in your options file:

```
exclude c:\test\*
include c:\working\*
```

The complete set of include-exclude specifications would be:

```
exclude c:\test\* (second last statement from the client options file)
include c:\working\* (last statement from the client options file)
include c:\test\* (sequence number 1 from the client option set)
exclude c:\working\* (sequence number 2 from the client option set)
```

Note:

Include-exclude statements are read and processed from the bottom to the top.

Processing using this set of include-exclude statements would result in the files under C:\working not being backed up, and those files under C:\test being backed up.

7.3.2 Default environment

There are no client option sets created by default in an installed Tivoli Storage Manager environment.

7.3.3 Recommended client option sets

In 1.2, “Redbook support material” on page 2, we show how to load a predefined macro into Tivoli Storage Manager. The macro `mac.optionsets`, which we provide to define client option sets for our redbook environment, is shown in B.1.2, “Define client option sets” on page 286.

In our environment, we use multiple client option sets. All option sets are built from the base option set which we called 'redbook'. This option set contains options other than the include-exclude specifications. All other option sets are based on client platform. Each is built from the base option set with the addition of specific include-exclude recommendations for that platform.

7.3.4 Associating a client node with a client option set

The client node definition should be updated to use a client option set. You need either system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node belongs to issue the required `update node` command.

You can use the `select` command to view the names of all the client option sets. Do not use the `query` command, as it generates too much output, and you cannot easily distinguish the client option set names.

The administrative command to show names of all client option sets and the results of that command should look similar to this:

```
tsm: TSM010> select * from cloptsets
```

OPTIONSET_NAME	DESCRIPTION	LAST_UPDATE_BY	PROFILE
AIX	AIX Clients	SYSADMIN	
NETWARE	Netware Clients	SYSADMIN	
REDBOOK	Redbook Base Set	SYSADMIN	
WINDOWS	Windows Clients	SYSADMIN	

The administrative command to associate a client option set named AIX with the client node named PUTNEY, and the results of that command, should look like this:

```
tsm: TSM010> update node putney cloptset=aix
ANR2063I Node PUTNEY updated.
```

7.3.5 Working with client option sets

In this section we explain how to perform various processes related to client option sets that you will find useful in your Tivoli Storage Manager environment.

7.3.5.1 Creating a client option set

Defining a client option set consists of two steps:

1. Creating the option set
2. Populating it with client options

You create a new client option set with the `define cloptset` command or you can clone an existing option set with the `copy cloptset` command. You can add an option to an option set with the `define clientopt` command. You can remove an option from an option set with the `delete clientopt` command. You need system privilege or unrestricted policy privilege to issue these commands.

The administrative commands to define a client option set named ALLNODES, populate it with options, and the results of those commands, should look like this:

```
tsm: TSM010> define cloptset allnodes
ANR2046I DEFINE CLOPTSET: Optionset ALLNODES defined.

tsm: TSM010> define clientopt allnodes compression yes
ANR2050I DEFINE CLIENTOPT: Option COMPRESSION defined in optionset ALLNODES.

tsm: TSM010> define clientopt allnodes maxcmdretries 4
ANR2050I DEFINE CLIENTOPT: Option MAXCMDRETRIES defined in optionset ALLNODES.

tsm: TSM010> define clientopt allnodes verbose ''
ANR2050I DEFINE CLIENTOPT: Option VERBOSE defined in optionset ALLNODES.
```

The administrative commands to clone the client option set named ALLNODES to another named CLONE, replace the verbose option with the quiet option, and the results of those commands, should look like this:

```
tsm: TSM010> copy cloptset allnodes clone
ANR2055I COPY CLOPTSET: Optionset ALLNODES copied to optionset CLONE.

tsm: TSM010> delete clientopt clone verbose
ANR2053I DELETE CLIENTOPT: Option VERBOSE, sequence number 0, has been deleted
from optionset CLONE.

tsm: TSM010> define clientopt clone quiet ''
ANR2050I DEFINE CLIENTOPT: Option QUIET defined in optionset CLONE.
```

7.3.5.2 Deleting a client option set

You need either system privilege or unrestricted policy privilege to delete a client option set with the `delete cloptset` command.

The administrative command to delete the option set named ALLNODES, and the results of that command, should look like this:

```
tsm: TSM010> delete cloptset allnodes  
Do you wish to proceed? (Yes/No) y  
ANR2048I DELETE CLOPTSET: Optionset ALLNODES deleted.
```

The association between a node and a client option set is removed automatically when that client option set is deleted. The node is not associated with another client option set until and unless you reissue the update node command.

Chapter 8. Licensing your Tivoli Storage Manager system

This chapter describes the tasks of licensing a Tivoli Storage Manager system and monitoring its compliance. See 2.7, “License considerations” on page 43 for planning considerations.

8.1 Licensed features

The base Tivoli Storage Manager server license supports an unlimited number of administrative clients, one backup/archive client using named pipes or shared memory protocol, and a selection of removable media devices. Table 31 lists all currently available licensed features you can add on to your base Tivoli Storage Manager server license.

Table 31. Tivoli Storage Manager licensed features

Licensed feature	License files
Additional backup-archive clients: The base Tivoli Storage Manager license supports one backup-archive client	1client.lic, 5client.lic 10client.lic 50client.lic
Network connections: To use a network communication method, for example, TCP/IP	network.lic
Open Systems Environment client support	opensys.lic
Hierarchical storage management (HSM) clients, also known as space-managed clients	spacemgr.lic
Tivoli Disaster Recovery Manager	drm.lic
Tivoli Data Protection for Oracle	oracle.lic
Tivoli Data Protection for MS SQL Server	mssql.lic
Tivoli Data Protection for MS Exchange	msexch.lic
Tivoli Data Protection for Lotus Notes	lnotes.lic
Tivoli Data Protection for Lotus Domino	domino.lic
Tivoli Data Protection for Informix	informix.lic
Tivoli Data Protection for R/3	sapr3.lic
Advanced device support	advdev.lic

The enrollment certificate files for all Tivoli Storage Manager licenses are on the Tivoli Storage Manager installation CD-ROM. You register those licenses that you need by issuing the `register license` command with the name of the enrollment certificate file. When registered, the licenses are stored in a file named `nodelock` in the current directory.

Working through a licensing example should help you to understand which licenses you will need for your configuration.

Consider the following scenario. You want to install your Tivoli Storage Manager server on a Windows NT system. There are different client platforms that want to

use the Tivoli Storage Manager service, namely, sixty AIX clients, five Sun Solaris clients, and ten Windows clients. An IBM 3575 library L06 with two drives is your preferred choice for library use. Which licenses are required?

The easiest way to work through this is by following each license type.

- Base license: 1 x Windows NT single server license, one server install for the Windows NT platform. The single server edition includes a license for one backup/archive client through the named pipes protocol, that is, on the same machine as the server only, and for administrators connected to the server by any method.
- Network connections for clients: 1 x Windows NT Network Enabler licenses clients to communicate with the Tivoli Storage Manager server using network methods (for example, TCP/IP, IPX and NetBIOS)
- Additional backup/archive clients: There are a total of 75 clients, 76 if you include the Tivoli Storage Manager server itself, which is covered by the single server license. Each client connection feature licenses one backup/archive client, regardless of the operating system of the client. When ordering Tivoli Storage Manager you would simply specify 75 client connections. To actually perform the licensing step, you would register 1 x 50 User registration, 2 x 10 user registration and 1x 5 user registration, making a total of 75 clients.
- Advanced device support: The library you want to use is an IBM 3575 library L06 with two drives. You need to find out whether this model is covered by the base supported devices list (no additional license required), or whether it will require the extended device support license. Check the URL:

http://www.tivoli.com/support/storage_mgr/adsercli.htm

Scroll down and click on the supported devices list for the Windows NT server platform. In our example, you need 1 x NT Extended Device Support for this model with two drives.

8.2 Registering licensed features

If you received a Tivoli Storage Manager evaluation (try-and-buy), the license registration is done during the installation of the server.

If you bought a base Tivoli Storage Manager server license, you can obtain licenses for licensed features, and register those licenses by specifying the license files using the `register license` command.

You can register any or all of these features. For example, if you want to add 10 additional backup-archive clients, issue the following command:

```
tsm: TSM010> register license file(10client.lic)
ANR2852I Current license information:
ANR2835I Server is licensed for 92 clients.
ANR2861I Server is licensed to support NETWORK connections.
ANR2869I Server is licensed for Advanced Device Support.
ANR2853I New license information:

ANR2835I Server is licensed for 102 clients.
ANR2861I Server is licensed to support NETWORK connections.
NR2869I Server is licensed for Advanced Device Support.
```

You can also register a license by specifying the product password that is included in the license certificate file.

8.3 Saving your licenses

When license registration is complete, the licenses are stored in a file named NODELOCK in the server start directory.

Save the CD-ROM containing your enrollment certificate files if you need to register your licenses again for any of the following reasons:

- The server is corrupted.
- The server has moved to a different machine.
- The NODELOCK file is destroyed or corrupted.

8.4 License compliance

If license terms change, for example, if a new license is defined for the server, the server conducts an audit to determine if the current server configuration conforms to the license terms.

The server also periodically audits compliance with the license terms. The results of this audit are used to check and enforce license terms. If 30 days have elapsed since the previous license audit, the administrator cannot cancel the audit.

If the server uses a licensed feature but the license is not registered, the function fails. When you issue a command associated with an unlicensed feature, Tivoli Storage Manager does not issue a warning message, and the command fails.

If a Tivoli Storage Manager system exceeds the terms of its license agreement, one of the following occurs:

- The server issues a warning message indicating that it is not in compliance with the licensing terms.
- Operations fail because the server is not licensed for specific features.

In either case, you must contact your Tivoli account representative or authorized reseller to modify your agreement.

8.5 Monitoring licenses

There are two commands to monitor the license registration on your Tivoli Storage Manager system: `query license` and `audit license`.

8.5.1 Displaying license information

Use the `query license` command to display details of your current licenses and determine licensing compliance.

```

tsm: TSM010> query license

                Last License Audit: 01/24/00   16:44:17
                Registered Client Nodes: 4
                Licensed Client Nodes: 102
                Are network connections in use ? : No
                Are network connections licensed ? : Yes
Are Open Systems Environment clients registered ? : No
Are Open Systems Environment clients licensed ? : Yes
                Is space management in use ? : No
                Is space management licensed ? : Yes
                Is Tivoli Disaster Recovery Manager in use ? : No
Is Tivoli Disaster Recovery Manager licensed ? : Yes
                Is Advanced Device Support required ? : No
                Is Advanced Device Support licensed ? : Yes
                Is TDP for Oracle in use ? : No
                Is TDP for Oracle licensed ? : Yes
                Is TDP for MS SQL Server in use ? : No
Is TDP for MS SQL Server licensed ? : Yes
                Is TDP for MS Exchange in use ? : No
                Is TDP for MS Exchange licensed ? : Yes
                Is TDP for Lotus Notes in use ? : No
                Is TDP for Lotus Notes licensed ? : Yes
                Is TDP for Lotus Domino in use ? : No
Is TDP for Lotus Domino licensed ? : Yes
                Is TDP for Informix in use ? : No
                Is TDP for Informix licensed ? : Yes
                Is TDP for SAP R/3 in use ? : No
                Is TDP for SAP R/3 licensed ? : Yes
                Server License Compliance: Valid

```

8.5.2 Auditing licenses

An administrator can monitor license compliance by issuing the administrative command `audit licenses`. This command is used to compare the current configuration with the current licenses.

```

tsm: TSM010> audit license

ANR2817I AUDIT LICENSES: License audit started as process 40.
ANS8003I Process number 40 started.

```

Note:

During a license audit, the server calculates, by node, the amount of backup, archive, and space management storage in use. This calculation can take a great deal of CPU time and can stall other server activity. Use the `NOAUDITSTORAGE` server option to specify that storage is not to be calculated as part of a license audit.

8.5.3 Scheduling automatic license audits

Use the `set licenseauditperiod` command to specify the number of days between automatic audits performed by the Tivoli Storage Manager server.

```

tsm: TSM010> set licenseauditperiod 30
ANR2814I SET LICENSEAUDITPERIOD: License audit period changed to 30 days.

```

Chapter 9. Administrative client

The Tivoli Storage Manager administrative client provides an interface to a server that allows an administrator to manage and control Tivoli Storage Manager resources. This chapter discusses how to install a Tivoli Storage Manager administrative client and what options to configure. Some examples are shown of the interfaces for various types of access.

There are two types of administrative clients in Tivoli Storage Manager:

- The administrative command line interface
- The web administrative interface

MVS Users:

A TSO administrative interface is available on MVS platforms, but is not discussed here. See the *Tivoli Storage Manager for MVS and OS/390 Quick Start*, GC35-0376 for more information.

There is also a client web interface which is used for performing operations for a particular client node, such as backup and restore. The web administrative interface, on the other hand allows an administrator to manage and control the actual server resources.

The two interfaces (web administrative interface and client web interface) are frequently confused because they have similar names, are accessed by similar means, and both use signon ids that are defined as Tivoli Storage Manager administrators.

The web administrative interface connects to a Tivoli Storage Manager server via the Web to present an interface to a true Tivoli Storage Manager administrator session. It is used by a Tivoli Storage Manager administrator to configure, monitor, and control a Tivoli Storage Manager server just as they would use the command line interface. The web interface has some features in common with the native Windows GUI interface which used to be shipped with older product releases; however this component is no longer available or supported.

The web backup/archive client connects to a Tivoli Storage Manager backup/archive client via the web to present an interface that remotely performs the functions of the backup/archive client on that client. It is used by Tivoli Storage Manager administrators, help desk personnel, or end users to manually back up, restore, archive, or retrieve files on a remote client without having to physically visit the client. To enforce security, the web backup/archive client connects to a backup/archive client using a Tivoli Storage Manager administrator ID with NODE authority. This eliminates unique security definitions on every client for the Tivoli Storage Manager administrators, help desk personnel, or end users. For more information about the web backup/archive client, see 10.5, "Web client usage" on page 210.

9.1 Code installation

Tivoli Storage Manager server and client code fixes and enhancements are released on a regular basis. The fixes are available from IBM via the Internet or CD-ROM. See 3.1, “Latest code updates” on page 47 for further information.

This section describes the installation procedures for the following:

- Administrative command line interface
- Web administrative interface

Use of any administrative client requires defining an administrator ID with the appropriate authority. See 7.1.3, “Recommended administrators” on page 144 for information on creating an administrator ID.

9.1.1 Administrative client code

Typically, the administrative client code is installed on a remote workstation. Installation of the Tivoli Storage Manager administrative client code is an option that can be selected during the installation of the backup/archive client code. For information on installing your version of the client software, see *Tivoli Storage Manager Installing the Clients*, SH26-4102.

9.1.2 Web administrative interface

Installation of the web administrative interface happens automatically during the installation of the Tivoli Storage Manager server code. To enable the web administrative interface, place the following options in the server option file `dsmserv.opt` on the Tivoli Storage Manager server:

- `COMMETHOD HTTP` to select the communication method
- `HTTPPORT <an unused portnumber>` : the default is 1580, to select the TCP/IP port that Tivoli Storage Manager will use to communicate with the web administrative interface

Note:

At server initialization, the server reads the server options file. If you update a server option by editing the file, you must remember to stop and restart the server to activate the updated server options file.

The web administrative interface requires a Java 1.1.6-capable web browser such as Microsoft Internet Explorer 4.01 or Netscape 4.06 on the administrator’s workstation. Using a web browser means that Tivoli Storage Manager code is not required on the web browser workstation.

9.2 Customization

This section describes the various settings for the administrative interfaces.

9.2.1 Administrative client

Use of the command line interface requires a client options file to be present on the client.

Client options for an administrative client (other than the web administrative interface) are set in the client option files `dsm.opt`. Assuming a TCP/IP communications method, the only required option is the `TCPSERVERNAME`. We strongly recommend using a DNS name instead of the dotted IP address (`xxx.xxx.xxx.xxx`) for the `TCPSERVERNAME`. A dotted IP address is inflexible, because it creates management issues if the Tivoli Storage Manager server is moved, or the addressing scheme is changed. Other useful parameters include setting the date, time, number and language format options. Table 32 summarizes some of the most common administrative client options.

Table 32. Administrative client options

	Option	Notes
<code>TCPSERVERNAME</code>	kindu	The DNS name or IP address
<code>TCPPORT</code>	1500	Default
<code>DATEFORMAT</code>	3	YYYY-MM-DD
<code>TIMEFORMAT</code>	1	hh:mm:ss
<code>LANGUAGE</code>	AMENG	American English
<code>NUMBERFORMAT</code>	1	1,000.00

9.2.2 Web administrative interface

No setup is required to use the web administrator client, assuming that you have a compatible Web browser and TCP/IP available on the administrative workstation.

9.3 Interfaces

This section discusses the command line and the web administrative interfaces available with Tivoli Storage Manager, and also the web backup/archive client as an administrative interface for the backup/archive client.

9.3.1 Command line interface

The command line interface has the benefit of being a character mode interface, and therefore is well suited for those users who want to type the commands. You may also consider using it when you cannot access the web administrative interface.

To start the command line interface, you must execute the administrative client program by typing `DSMADMC` at the operating system prompt. You are then prompted for an administrative user ID and password. Without authentication, you cannot logon to the Tivoli Storage Manager server. The following shows the start of an administrative client session on UNIX:

```

root@ / > dsmadm
Tivoli Storage Manager
Command Line Administrative Interface - Version 3, Release 7, Level 1.0
(C) Copyright IBM Corporation, 1990, 1999, All Rights Reserved.

Enter your user id: sysadmin

Enter your password: *****

Session established with server TSM010: AIX-RS/6000
Server Version 3, Release 7, Level 1.0
Server date/time: 02/10/00 14:31:42 Last access: 02/10/00 14:30:57

tsm: TSM010>

```

You then get the prompt `tsm: SERVERNAME>` to type the administrative Tivoli Storage Manager commands. Once an administrative command is executed, you always return back to the `tsm: SERVERNAME>` prompt. This interactive method is also called *loop* mode. If you use the interactive mode, you do not have to re-enter your password with each command.

Another method to start an administrative session is the so-called *batch* mode. You simply type DSMADM together with the administrative ID, the password, and the actual command that you want to execute. In this mode, the administrative client processes the command and returns to the operating system prompt.

Help is available from the `tsm: SERVERNAME>` command prompt by typing `help xxx` where `xxx` is blank, a command, or a Tivoli Storage Manager message number.

9.3.2 Web administrative interface

In Tivoli Storage Manager, the web administrative interface is the primary interface for administration. It incorporates Java applets and provides the following functions:

- Tivoli Storage Manager enterprise console

The enterprise console displays Tivoli Storage Manager as a single entity in the enterprise. It is an interface that integrates Tivoli Storage Manager server and client functions for the administrator as a single application to manage Tivoli Storage Manager in a distributed environment
- Server function

The web administrative interface is a Tivoli Storage Manager server function and consists of an integrated web server and a new HTTP server communications protocol.
- Browser requirements

The web administrative interface can be used from any workstation running an HTML 3.0 compliant Web browser with support for Java 1.1.6 or higher. Netscape 4.06 has the required Java support. Microsoft's Internet Explorer 4.01 has the required Java support.

- Security

Security for the web administrative interface is based on enhanced logging of Web-based sessions, password restrictions, and administrator lockout capabilities. In addition, Tivoli Storage Manager now has a Tivoli Storage Manager Secure Web Administration Proxy package, which enhances the web operations as well as potentially improving administrative client performance.

- Administrative functions

The web administrative interface can be used to perform all administrative functions graphically, and it provides a command line and server event viewer.

The web administrative interface connects to a Tivoli Storage Manager server through the Web to present an interface to a true Tivoli Storage Manager administrator session. Using a Web browser means that you do not have to have any Tivoli Storage Manager code on the Web browser workstation.

To access the web administrator interface, start your web browser. In the Location or Address field, enter the URL:

```
http://<servername>:1580
```

Where, <servername> is the TCP/IP DNS name or dotted IP address of your Tivoli Storage Manager server machine, and 1580 is the TCP/IP port number set by the HTTPPORT option.

Figure 20 shows the web administrative interface logon screen.

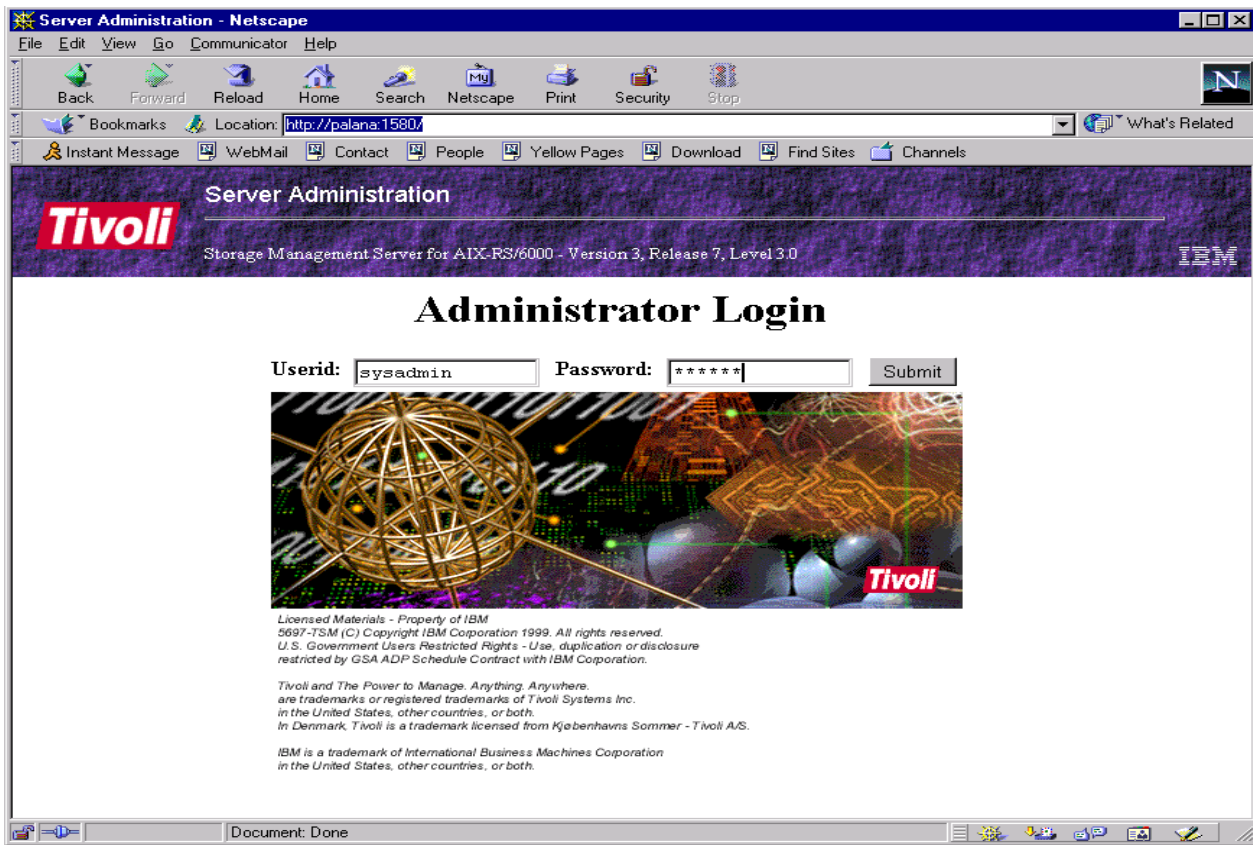


Figure 20. Web administrative interface logon screen

When an administrator logs in to the server through the web administrative interface, the Tivoli Storage Manager enterprise console is displayed. The enterprise console consists of three frames. Additional command line and event viewers can be displayed.

- **Banner Frame:** The banner frame indicates which server you are connected to (server name, platform, and version, release, level) and the fact that you are using the Tivoli Storage Manager enterprise console. It also indicates the name you are logged on to as administrator.

The only operable control on the banner frame is a selection button that you can use to choose whether you want to display the command line and the event viewer. You can turn those windows off or on, depending on what you want to see in your browser display. Another option offers you the ability to log off (forcing re-authentication).

- **Tree Frame:** A collapsible tree is used for navigation. The content and layout are determined by the “view” that you have chosen. The initial branches of the tree are the views from which you can choose. Views are simply paradigms that are used to navigate to the objects in the enterprise that you want to use. Available views are operation, network, configuration, and object.
- **Detail Frame:** The detail frame displays detailed information about the item that you have selected from the tree frame. All operations that are available to manipulate the selected item are provided in the form of a pull-down menu in the top-right corner of the detail frame. For example, when a server group is selected, its attributes are displayed in the detail frame. Operations are provided for adding servers to the group, removing servers from the group, and defining new server groups.
- **Command Line:** An administrative command line can be displayed by selecting it from the selections pull-down menu in the banner frame.
- **Event Viewer:** A server event viewer can be displayed by selecting it from the selections pull-down menu in the banner frame.

Figure 21 shows the Tivoli Storage Manager enterprise console with the Server Status screen and an open drop-down box with the Set password expiration option selected. Also shown are the administrative command line window and the server event viewer at the bottom of the screen.

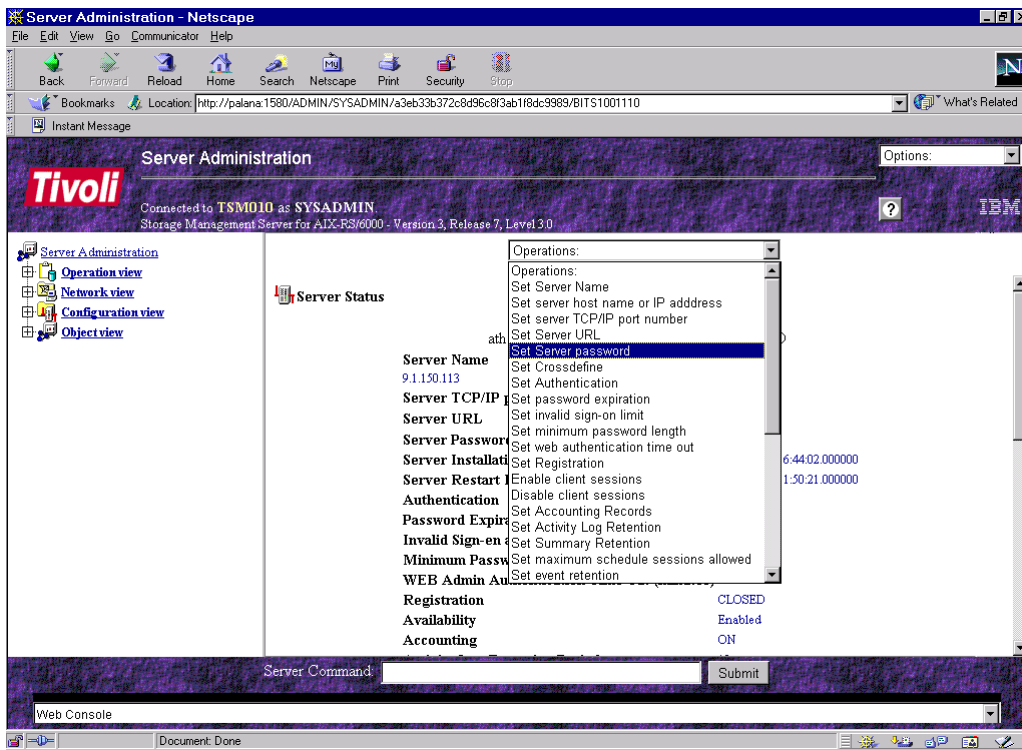


Figure 21. Tivoli Storage Manager Enterprise Console browser interface

The web administrative interface has a drop-down box in the upper right corner of the screen that allows you to perform three functions: show a command line, show an event console line, and log off the Tivoli Storage Manager administrator. In the command line area, you can issue any Tivoli Storage Manager command that could be issued from the command line interface. The event console line displays the command issued as a result of performing a function using the web administrative interface. It maintains a history of the commands that have been issued in this session. The command line and the event console line can be displayed simultaneously.

The web administrator session must re-authenticate after the time set in the Tivoli Storage Manager server setting WEBAUTHTIMEOUT expires. At the next interaction with the Tivoli Storage Manager server, you are prompted to enter your administrator ID and password. This time frame is based on wall clock time, not on the time since the last transaction. For example, if WEBAUTHTIMEOUT 10 is set, then every 10 minutes you will be prompted to re-authenticate, whether you have been active in those 10 minutes or not. Setting WEBAUTHTIMEOUT to zero indicates no timeout will occur.

We recommend that you use the web administrative interface, because it is platform-independent, it is location-independent (you only need a Web browser) and it allows you to perform all Tivoli Storage Manager administrator functions.

Chapter 10. Backup/archive client

This chapter covers the steps you need to perform after installing the code from the installation media. It shows you how to use the backup/archive client interactively by using the command line interface and the GUI. It also shows you how to configure the backup/archive client to work either as a foreground program or as a background process, so that you can automate the backup processes using the scheduler facility. See 2.1, “Client environment data” on page 11 for planning considerations.

10.1 Code installation

There are two types of Tivoli Storage Manager client. The first is the native backup/archive client. This code is installed on every Tivoli Storage Manager client node and provides the local interface to backup and restore operations on that node. The client code provides at least a command-line interface and on most operating system platforms, a native GUI, too (see 10.3, “Interfaces” on page 175 for details of platform support). This provides you with both a command-line and GUI interface to backup and restore files from that system only. The second type of client is the web client which provides a remote interface to backup and restore operations on a Tivoli Storage Manager client node. This client is particularly suitable for help desk operations. The help desk administrator can perform restores on behalf of the end users as required without having to be physically present at that client and only needs the Internet browser installed on his system. The client who is receiving the remote operations needs to install the backup/archive client code, including the web client component. Figure 22 shows the difference between a traditional restore using the native backup/archive client and using the web client.

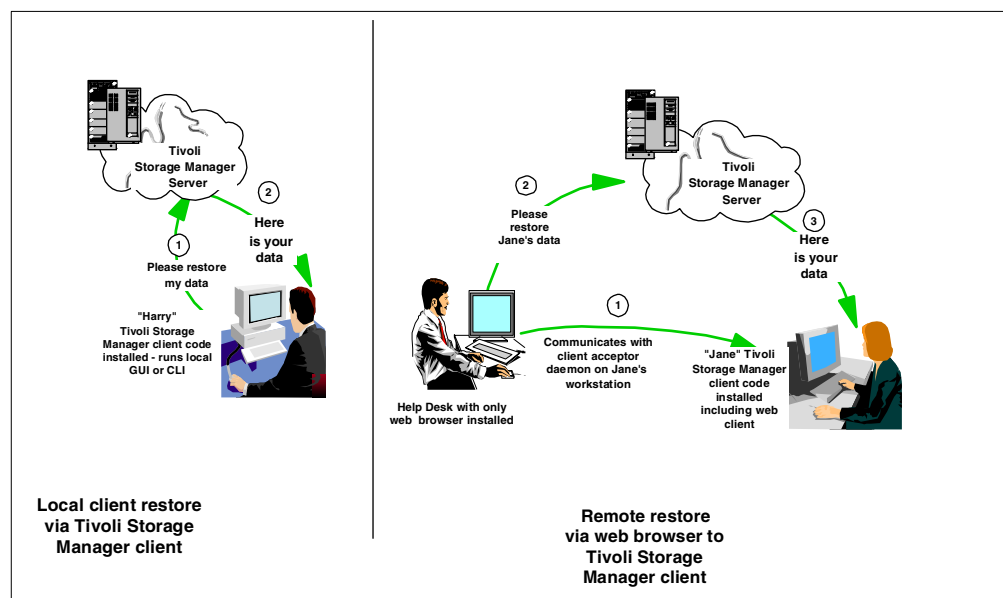


Figure 22. Difference between a local and web-based restore

10.1.1 Backup/archive client code

The backup/archive client is the software piece that is installed on each machine that you need to backup. After installing the client code and customizing how it should interact with the Tivoli Storage Manager server, you have a working machine ready to send and receive data.

Tivoli Storage Manager client code fixes and enhancements are released on a regular basis. The fixes are available from IBM through the Internet or CD-ROM. See 3.1, "Latest code updates" on page 47 for further information.

For further assistance on code installation procedures, see the Tivoli Storage Manager manual *Tivoli Storage Manager Installing the Clients*, SH26-4102, as well as the appropriate platform *Using the Backup-Archive Client* manuals (listed in D.3, "Tivoli Storage Management Publications" on page 304).

10.1.2 Web backup/archive client

The web backup/archive client main window is shown in Figure 23. Its functions are described in the following sections.

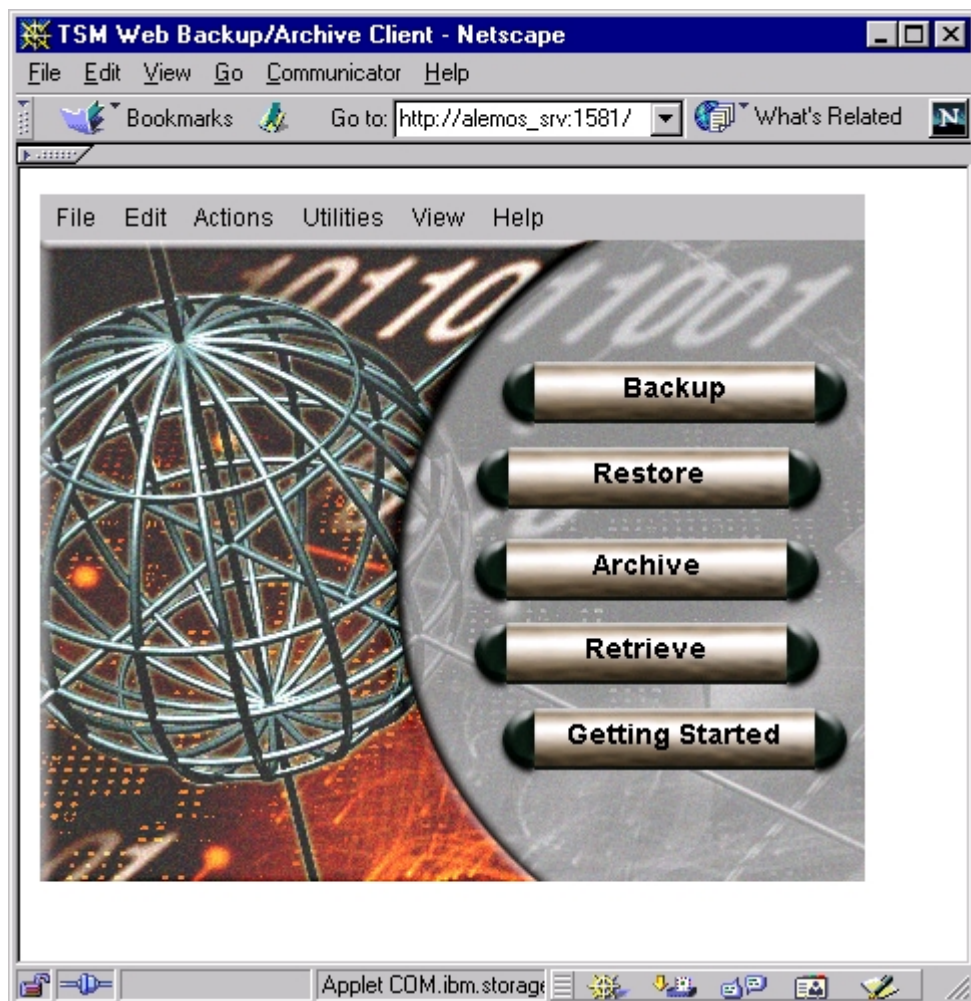


Figure 23. Web backup/archive client main window

The web client consists of new client components that must be installed and configured before the client can be accessed with a Web browser.

The web client is installed with the backup/archive client package. The web client consists of two processes on the client workstation: the client acceptor and the remote client agent.

The client acceptor is an HTTP daemon that serves the web client Java applet to the web browser. The name of the executable is DSMCAD. On AIX and other UNIX clients, it should be run as a daemon. For Windows NT it is installed and run as a service. For Windows 95 and 98 it must be executed in a DOS window. For Netware it is an NLM (NetWare Loadable Module) that should be loaded as part of the NetWare startup.

The remote client agent performs the client functions initiated with the web client interface. The name of the executable is DSMAGENT. The agent does not have to be running all the time. The acceptor daemon starts the agent when client functions are initiated through the web client interface.

10.2 Customization

The following describes the different settings for the backup/archive client and the web backup/archive client.

10.2.1 Backup/archive client

This section summarizes the steps you must follow to customize your client installation:

- Setting environment variables
- Defining client options files
- Defining include-exclude lists

For a more detailed explanation see *Tivoli Storage Manager Installing the Clients*, SH26-4102.

10.2.1.1 Environment variables

The following are the variables that you must update and set in your client environment:

- **PATH:** This is the default search path that the operating system uses to locate executable files. Set this to include the fully qualified paths of the Tivoli Storage Manager client directories.
- **DSM_CONFIG:** Tivoli Storage Manager uses this environment variable to locate the client options file `dsm.opt`. It points to the client user options file for users who create their own personalized options file.
- **DSM_DIR:** Tivoli Storage Manager uses this environment variable to locate all other client files. It points to the executable files DSMTCA and DSMSTAT, the resource files, and the `dsm.sys` file (on UNIX only).
- **DSM_LOG:** This points to the directory where you want the `dsmerror.log` file to reside. The error log file contains information about any errors that occur during processing. The error log is intended to help you diagnose severe errors.

10.2.1.2 Options file

Tivoli Storage Manager includes options that control processing for user sessions. For example, you can use options to inform Tivoli Storage Manager which communication method to use, or what format to use for dates.

We provide recommended options files for the various client platforms, as the layout of the supplied options file is not easy to understand. Section B.3, “Client options files” on page 297 contains our client option files.

The backup/archive client has at least one configuration file, which is divided into the following parts:

- Communication options
- Operational options
- Site-dependent options
- Include-exclude options

Some of the configuration steps that you perform from the client side may need complementary configuration on the Tivoli Storage Manager server side.

On all non-UNIX platforms, all options reside in the client options file `dsm.opt` (Preferences file for Macintosh) which resides by default in the Tivoli Storage Manager client directory.

On UNIX, you can set options in three different files:

- Client system options file (`dsm.sys`): In the client system options file, a root user sets options that are required to establish communication with a Tivoli Storage Manager server, and options that authorize users on your workstation to use Tivoli Storage Manager services. A root user can also set options that affect backup and archive processing, and options that affect scheduled services.
- Client user options file (`dsm.opt`): In the default client user options file, a root user can set options that determine which Tivoli Storage Manager server your client node contacts, and that specify the formats to use for date, time, and numbers. A root user can also set options that affect backup, archive, restore, and retrieve processing. In addition, users can also create their own personalized client user options file if they want to use different options. Users can overwrite an option contained in a client user options file by entering a different value for the option with an appropriate Tivoli Storage Manager command.
- Include-exclude options file: In the include-exclude options file, a root user can set options to exclude specific files from backup services. In addition, a root user can set options to associate specific files with different management classes. This file can be called any legal filename and be located in any directory, provided that you reference its full pathname in the `dsm.opt` file.

We provide recommended client options files for various client platforms, because the layout of the supplied options file is not easy to understand. Section B.3, “Client options files” on page 297 contains our client option files. We assume that TCP/IP is the network protocol. Most options can be centrally set using client option sets. We also provide examples in B.1.2, “Define client option sets” on page 286 for various client platforms.

For more detailed information about client options, refer to the following Tivoli Storage Manager manuals:

- *Tivoli Storage Manager for UNIX Using the Backup-Archive Clients*, SH26-4105
- *Tivoli Storage Manager for NetWare Using the Backup-Archive Client*, SH26-4100
- *Tivoli Storage Manager for Windows Using the Backup-Archive Client*, SH26-4101

Communication options

Any Tivoli Storage Manager client needs some basic communication settings to send or receive data. You can choose any protocols that are supported in Tivoli Storage Manager to use in your environment, provided that both client and server are properly configured for them. Depending on the type of Tivoli Storage Manager server, you may need extra help from the system administrator to have all the necessary communication definitions in place. The server may have multiple communication protocols configured. If that is the case, you may choose any of those to start communicating with the Tivoli Storage Manager server.

The following is a list of communication options you need to set:

- **COMMMETHOD**: The COMMMETHOD option specifies the communication method you are using to provide connectivity for client-server communication. The Tivoli Storage Manager server configuration file (`dsmserve.opt`) must have specified the same communication method with a valid port address, so that the server can accept client requests.
- **TCPSERVERADDRESS**: This option specifies the TCP/IP address for a Tivoli Storage Manager server. This is either the IP address or the hostname of the Tivoli Storage Manager server. If you choose to use the hostname, make sure that there is a name resolution service (for example, DNS, local HOSTS file or WINS) in place for the TCP/IP configuration in the client machine. Otherwise, you may not be able to connect to the Tivoli Storage Manager server. Bear in mind that if you use name resolution to figure out a server address, you will depend on the name resolution protocol to be up and running. This may not be practicable in your environment. If you believe that name resolution may not be always available during your backup windows, consider using the IP addresses instead. Although it gives you less flexibility, it will work despite any name resolution service problem.
- **TCPPORT**: The TCPPORT option specifies a TCP/IP port address for a Tivoli Storage Manager server. By default, this value is 1500.
- **TCPWINDOWSIZE**: This option specifies the size, in kilobytes, of the TCP/IP sliding window for your client node. This setting is highly operating-system specific. You must only use the allowed values for your TCP/IP implementation.
- **TCPBUFFSIZE**: This option specifies the size, in kilobytes, of the Tivoli Storage Manager internal TCP/IP communication buffer. This setting is highly operating-system specific. You must only use the allowed values for your TCP/IP implementation.
- **NODENAME**: This option assigns a new name to your client node if you do not want to use the default. The default is the hostname of your client machine, which we recommend using.

Operational options

The following list gives some examples of client options for backup, restore, and scheduling services.

- **PASSWORDACCESS:** This option specifies whether you want your Tivoli Storage Manager password automatically generated or always manually prompted for. We highly recommend setting **PASSWORDACCESS** to **GENERATE**.
- **REPLACE:** The **REPLACE** option specifies what you want Tivoli Storage Manager to do when it restores files that already exist on your workstation. This option applies to the `restore` and `retrieve` commands only.
- **SUBDIR:** The **SUBDIR** option specifies whether you want Tivoli Storage Manager to include subdirectories of named directories. This option applies, for example, to `selective`, `restore`, `archive`, and `retrieve`.
- **TAPEPROMPT:** This option specifies whether to wait for a tape to mount if it is required for a backup, archive, restore, or retrieve process, or to prompt you for a choice.
- **SCHEDMODE:** The **SCHEDMODE** option specifies whether you want to use the client-polling mode (your client node periodically asks the Tivoli Storage Manager server for scheduled work), or the server-prompted mode (the server contacts your client node when it is time to start a scheduled operation). All communication methods can use the client polling mode, but only TCP/IP can use the server prompted mode.

Site dependent options

If your locale is not the USA, then we recommend that you adjust the date, number, and time format options and also the language option in your client options file to support your requirements.

Include-exclude options

The include-exclude options must be placed as part of the client options file `dsm.opt` on non-UNIX platforms, or in a separate include-exclude file on UNIX. On UNIX, you need to add the name of the file in your client system options file `dsm.sys` with the keyword **INCLEXCL** as shown as follows:

```
INCLEXCL /tsm/server1/inclexcl.file
```

The **INCLUDE** option specifies files within a broad group of excluded files that you want to include for backup services. You also use this option to assign a management class either to specific files or to all files to which you have not already assigned a specific management class, and for which you do not want Tivoli Storage Manager to use the default management class.

The **EXCLUDE** option excludes files from backup services. When you back up files, any files you exclude are not considered for backup. For example, we recommend you exclude the Tivoli Storage Manager installation directory.

Include-exclude options are checked from the bottom up to the top of the list until a match is found. If a match is found, the processing stops and checks whether the option is **INCLUDE** or **EXCLUDE**. If the option is **INCLUDE**, the file is backed up, using the assigned management class. If the option is **EXCLUDE**, the file is not backed up.

See B.1.2, “Define client option sets” on page 286 for examples of include-exclude lists.

10.2.2 Web backup/archive client

Three client options are important for the web client. The PASSWORDACCESS option must be set to GENERATE and the password must be generated by first running a backup/archive client session. The remote client agent establishes connection to Tivoli Storage Manager in the same manner as the backup/archive GUI and command line clients. It requires the generated client password to authenticate a client session when the web client is used.

The acceptor daemon listens on a TCP/IP port for incoming connections from an administrator’s web browser. By default it listens on port 1581. This default port can be overridden with the HTTPPORT client option.

The REVOKEREMOTEACCESS option has two possible values: NONE, which is the default, and ACCESS. If the option is set to NONE, any administrator userid with either client access or client owner authority can perform client operations. If the option is set to ACCESS, administrator userids with only client access authority are prevented from performing remote client operations. A pop-up message displayed in the web browser indicates that the administrator userid being used has insufficient authority. This option does not prevent administrators with client owner or higher authorities from performing client functions. Table 33 summarizes the web client options.

Table 33. Web client options

	Option
PASSWORDACCESS	GENERATE
HTTPPORT	1581
REVOKEREMOTEACCESS	NONE

10.3 Interfaces

Tivoli Storage Manager client interfaces let you submit and receive information from the Tivoli Storage Manager server. In this section we describe how you can start, use, and stop the two available client interfaces: the command line and the GUI. Table 34 lists all available backup/archive client interfaces by platform.

Table 34. Backup-Archive client interfaces by platform

Platform	Command Line Interface	GUI	Web backup/archive client
Apple Macintosh		x	
Data General DG/UX	x	x	
Compaq Tru64 (formerly Digital UNIX)	x	x	x
Hewlett-Packard HP/UX	x	x	x

Platform	Command Line Interface	GUI	Web backup/archive client
IBM AIX	x	x	x
IBM OS/2	x	x	x
IBM OS/390 UNIX System Services	x	x V2 only	x
Linux	x	x	x
Microsoft Windows 95, 98, and NT/Intel	x	x	x
Microsoft Windows NT DEC Alpha	x	x	x
NCR UNIX SVR4	x	x	
NEC EWS-UX/V	x	x	
Novell NetWare	x		x
NUMA-Q (formerly Sequent)	x	x	x
SCO UNIX	x	x	s
Siemens Nixdorf Reliant UNIX	x	x	
Silicon Graphics IRIX UNIX	x	x	x
Sun Solaris	x	x	x

There are minor differences between the backup/archive client code among the platforms. For example, in Windows NT you have specific options to handle the NT Registry information, which are not found in any other platform. Despite those specific options, all commands are the same except for the file space specification for each platform (in UNIX this uses the format `/usr/tivoli`, and in Windows, `D:\newdir`). Although the examples in this section are based on a UNIX machine, we also give examples of Windows commands.

10.3.1 Command line

The command line interface is a character mode interface, and therefore is well suited for those users who want to type the commands. You may also consider using it when you cannot access the GUI or when you want to create scripts for automating a backup process. You may use it for backup, restore, archive, retrieve operations and to start the Tivoli Storage Manager scheduler.

10.3.1.1 Starting a session

To start the command line interface, you must execute the backup/archive client program by typing `DSMC` at the operating system prompt.

```

root@ / > dsmc

Tivoli Storage Manager
Command Line Backup Client Interface - Version 3, Release 7, Level 1.0
(C) Copyright IBM Corporation, 1990, 1999, All Rights Reserved.

tsm>

```

You then get the prompt `tsm>` where you can type the backup/archive client commands. When issuing the first command, you may be prompted for your password before the command is executed. This occurs if you are not using automatic password handling. If you set up your client to use automatic password handling, by setting the `PASSWORDACCESS` option to `GENERATE`, the command you issue will be executed directly.

Once a command is executed you always return back to the `tsm>` prompt. This interactive method is also called *loop* mode.

Another way to start a client session is the so-called *batch* mode. You simply type `DSMC` together with the actual command that you want to execute. In this mode, the backup/archive client processes the command and returns to the operating system prompt. This is the recommended way to use the `DSMC` command when creating scripts for automation. You can create each `DSMC` command in turn (as well as any other operating system command you want to include in the script) and save them in a batch file for processing later.

Help is available from the `tsm>` command prompt by typing `help` and then selecting a number from the list.

```

root@ / > dsmc

Tivoli Storage Manager
Command Line Backup Client Interface - Version 3, Release 7, Level 1.0
(C) Copyright IBM Corporation, 1990, 1999, All Rights Reserved.

tsm> help
The following help topics are available.
Enter the number of the desired help topic or 'q' to quit,
'd' to scroll down, 'u' to scroll up.

    0 - Using Commands
    1 - ARCHIVE
    2 - BACKUP IMAGE
    3 - CANCEL RESTORE
    4 - DELETE ACCESS
    5 - DELETE ARCHIVE
    6 - DELETE FILESPACE
    7 - HELP
    8 - INCREMENTAL
    9 - LOOP
   10 - MACRO
   11 - QUERY ACCESS
   12 - QUERY ARCHIVE
   13 - QUERY BACKUP
   14 - QUERY BACKUPSET
   15 - QUERY FILESPACE
   16 - QUERY IMAGE
   17 - QUERY MGMTCLASS

```

10.3.1.2 Running backup operations

To start backing up a file, directory, or the whole machine, you use the `incremental` command. This command backs up all new or changed files in your

default client domain, or in the file systems you specify that are not excluded from backup operations.

Another backup command is the `selective` command. During a selective backup, Tivoli Storage Manager sends copies of the files to the server even if they have not changed since the last backup. This might result in having more than one copy of the same file on the server. If that occurs, you might not have as many different file copy versions on the server as you intended. Your version limit might consist of identical files. To avoid that, use the `incremental` command to back up only changed and new files. However you could consider using a periodic `selective` backup operation to regroup all the files one machine onto one tape. This might be helpful if you are not collocating data, or as an alternative to the backup set functionality. Remember the backup set functionality would usually perform this task more efficiently as it is done entirely within the server without requiring re-sending of any client data to the server.

On selected UNIX platforms you also have the `backup image` command for raw logical volume or complete file system backup.

Table 35 gives you some examples on how to perform a backup operation using a UNIX backup/archive client.

Table 35. Backup command examples: UNIX

If you want to perform this then this is the client command
Incremental backup of the client domain	<code>dsmc incremental</code>
Incremental backup of a filesystem (/home)	<code>dsmc incremental /home</code>
Incremental backup of all files in a directory (/home/lucia)	<code>dsmc incremental /home/lucia/</code>
Incremental backup of all files in a directory (/home/lucia) and all its subdirectories	<code>dsmc incremental /home/lucia/ -subdir=yes</code>
Selective backup of a filesystem (/home) and all its subdirectories	<code>dsmc selective /home/ -subdir=yes</code>
Selective backup of all files in a directory (/home/lucia) and all its subdirectories	<code>dsmc selective /home/lucia/ -subdir=yes</code>
Point-in-time, full Logical Volume backup for a filesystem (/bruno)	<code>dsmc backup image /bruno</code>
Logical Volume backup for a RAW device (/dev/test)	<code>dsmc backup image /dev/test</code>

Table 36 gives you some examples on how to perform a backup operation using a Windows backup/archive client.

Table 36. Backup command examples: Windows

If you want to perform this then this is the client command
Incremental backup of the client domain	<code>dsmc incremental</code>
Incremental backup of the C: volume only	<code>dsmc incremental c:</code>
Incremental backup of all files in a directory (C:\work)	<code>dsmc incremental c:\work*</code>

If you want to perform this then this is the client command
Incremental backup of all files in a directory (C:\work) and all its subdirectories	dsmc incremental c:\work* -subdir=yes
Selective backup of a directory (C:\work)	dsmc selective c:\work*
Selective backup of all files in a directory (C:\work) and all its subdirectories	dsmc selective c:\work* -subdir=yes

Include/Exclude Processing:

Especially on Windows platforms, you need to be aware of the difference between the pattern (*.*) — double wildcard and (*) — single wildcard. In an INCLUDE/EXCLUDE statement, these patterns are processed differently. The specification INCLUDE C:\WORK* will match all files in the WORK directory. The specification INCLUDE C:\WORK*. * will match files in the WORK directory only if they include the period or dot. For the backup commands however, the patterns are considered the same. For example the two commands DSMC INCREMENTAL C:\WORK* and DSMC INCREMENTAL C:\WORK*. * both match all files.

Therefore if you put INCLUDE C:\WORK*. * in your INCLUDE/EXCLUDE list and then issue the command DSMC INCREMENTAL C:\WORK* you will not see any files without the dot being backed up. This is because they do not match the INCLUDE rule (assuming an EXCLUDE statement excludes everything else).

We recommend that you configure all include/exclude options with single wildcards, (for example INCLUDE C:\WORK*, EXCLUDE D:\TEMP*) to avoid potential hard-to-diagnose problems, unless you have a specific requirement for files with dots in their names.

The next example illustrates the relationship between wildcards in the INCLUDE/EXCLUDE statements and backup commands.

Exclude rules preventing some files from being backed up

Suppose we have two files (c:\work\myfile.doc and c:\work\myfile) and we configure an exclude rule to avoid some of the files, as below:

include/exclude settings in dsm.opt:

...

exclude c:\...*

include c:\work*. *

...

The DSMC INCR command will never backup the file c:\work\myfile, because it is not caught by the INCLUDE statement and therefore is excluded under the previous rule. It does not matter if you use *.* or * on the INCR command:

```

C:\PROGRA~1\Tivoli\TSM\baclient>dsmc incr c:\work\*. *

Tivoli Storage Manager
Command Line Backup Client Interface - Version 3, Release 7, Level 2.0 n
(C) Copyright IBM Corporation, 1990, 1999, All Rights Reserved.

Node Name: DEVNODE1
Session established with server TSM010: AIX-RS/6000
  Server Version 3, Release 7, Level 2.0
  Server date/time: 02/15/2000 15:31:47  Last access: 02/15/2000 15:30:44

Incremental backup of volume 'C:\WORK\*. *'
Normal File-->          7  \\a78axfwy\c$\work\myfile.doc [Sent]
Successful incremental backup of '\\a78axfwy\c$\work\*. *'

```

```

C:\PROGRA~1\Tivoli\TSM\baclient>dsmc incr c:\work\*

Tivoli Storage Manager
Command Line Backup Client Interface - Version 3, Release 7, Level 2.0 n
(C) Copyright IBM Corporation, 1990, 1999, All Rights Reserved.

Node Name: DEVNODE1
Session established with server TSM010: AIX-RS/6000
  Server Version 3, Release 7, Level 2.0
  Server date/time: 02/15/2000 15:33:28  Last access: 02/15/2000 15:33:11

Incremental backup of volume 'C:\WORK\*'
Normal File-->          7  \\a78axfwy\c$\work\myfile.doc [Sent]
Successful incremental backup of '\\a78axfwy\c$\work\*'

```

The same will happen for the corresponding DSMC SELECTIVE commands.

UNIX examples

The following is an incremental backup operation and its command output:

```
tsm> incremental /homeusr/ -subdir=yes
Node Name: PUTNEY
Session established with server TSM010: AIX-RS/6000
  Server Version 3, Release 7, Level 1.0
  Server date/time: 02/10/00 17:06:28 Last access: 02/10/00 17:05:23

Incremental backup of volume '/homeusr/'
Directory--> 512 /homeusr [Sent]
Directory--> 512 /homeusr/bruno [Sent]
Directory--> 512 /homeusr/claudia [Sent]
Directory--> 512 /homeusr/daniela [Sent]
Directory--> 512 /homeusr/ecastro [Sent]
Directory--> 512 /homeusr/lala [Sent]
Directory--> 512 /homeusr/lucia [Sent]
Directory--> 512 /homeusr/pp [Sent]
Normal File--> 6 /homeusr/daniela/finalreport [Sent]
Normal File--> 11 /homeusr/daniela/project_y2k [Sent]
Normal File--> 15 /homeusr/lala/cubano [Sent]
Normal File--> 2,523 /homeusr/lucia/db2setup.log [Sent]
Normal File--> 593 /homeusr/lucia/dsmerror.log [Sent]
Normal File--> 173 /homeusr/lucia/dtappint.log [Sent]
Normal File--> 409 /homeusr/lucia/ibmnpm2Pze7a.log [Sent]
Normal File--> 409 /homeusr/lucia/ibmnpm921_qa.log [Sent]
Normal File--> 409 /homeusr/lucia/ibmnpmBAu.ya.log [Sent]
Normal File--> 408 /homeusr/lucia/ibmnpmHUFILMa.log [Sent]
Normal File--> 409 /homeusr/lucia/ibmnpmYMyKUa.log [Sent]
Normal File--> 409 /homeusr/lucia/ibmnpmJlb7a.log [Sent]
Normal File--> 409 /homeusr/lucia/ibmnpmfzuQya.log [Sent]
Normal File--> 1,021 /homeusr/lucia/inst.log [Sent]
Normal File--> 1,227 /homeusr/lucia/listener.log [Sent]
Normal File--> 0 /homeusr/lucia/wpreinst.log [Sent]
Successful incremental backup of '/homeusr/*'

Total number of objects inspected: 24
Total number of objects backed up: 24
Total number of objects updated: 0
Total number of objects rebound: 0
Total number of objects deleted: 0
Total number of objects failed: 0
Total number of bytes transferred: 8.57 KB
Data transfer time: 0.00 sec
Network data transfer rate: 18,351.84 KB/sec
Aggregate data transfer rate: 2.79 KB/sec
Objects compressed by: 0%
Elapsed processing time: 00:00:03
```

This is an example of a full Logical Volume backup operation:

```

root@ / > dsmc backup image /bruno

Tivoli Storage Manager
Command Line Backup Client Interface - Version 3, Release 7, Level 2.0 n
(C) Copyright IBM Corporation, 1990, 1999, All Rights Reserved.

Node Name: PUTNEY
Session established with server TSM010: AIX-RS/6000
  Server Version 3, Release 7, Level 2.0
  Server date/time: 02/16/00 11:54:28 Last access: 02/16/00 11:00:25

Backup Image Function Invoked.

Volume --> 16,777,216 /bruno [Sent]
Selective Backup processing of '/bruno' finished without failure.

Total number of objects inspected:      1
Total number of objects backed up:     1
Total number of objects updated:       0
Total number of objects rebound:      0
Total number of objects deleted:       0
Total number of objects expired:       0
Total number of objects failed:        0
Total number of bytes transferred:    15.25 MB
Data transfer time:                    9.00 sec
Network data transfer rate:            1,735.10 KB/sec
Aggregate data transfer rate:          1,411.56 KB/sec
Objects compressed by:                 0%
Elapsed processing time:                00:00:11

```

AIX JFS and raw volumes

The backup image command knows the difference between a raw logical volume (with no filesystem) and a logical volume with a JFS (journalled filesystem) defined to it. You cannot use the backup image function to try to backup a defined JFS filesystem using its raw device specification. In the previous example, if you try to run a `dsmc backup image /dev/bruno`, the command will fail because this is a JFS filesystem. On the other hand, you will be able to perform a `dsmc backup image /dev/test`, provided it is a raw logical volume, with no JFS filesystem definitions in place.

The next example shows a successful raw image backup (`/dev/test`) and a failure (`/dev/bruno`) when trying to use the backup image function on a defined JFS filesystem:

```

root@ / > lsfs /test
lsfs: 0506-915 No record matching /test was found in /etc/filesystems.

root@ / > dsmc backup image /dev/test
Tivoli Storage Manager
Command Line Backup Client Interface - Version 3, Release 7, Level 2.0 n
(C) Copyright IBM Corporation, 1990, 1999, All Rights Reserved.

Node Name: PUTNEY
Session established with server TSM010: AIX-RS/6000
  Server Version 3, Release 7, Level 2.0
  Server date/time: 02/16/00 14:19:34 Last access: 02/16/00 14:18:40

Backup Image Function Invoked.

Volume --> 16,777,216 /dev/test [Sent]
Selective Backup processing of '/dev/test' finished without failure.

Total number of objects inspected:      1
Total number of objects backed up:     1
Total number of objects updated:       0
Total number of objects rebound:      0
Total number of objects deleted:       0
Total number of objects expired:       0
Total number of objects failed:        0
Total number of bytes transferred:     15.00 MB
Data transfer time:                    9.00 sec
Network data transfer rate:            1,706.66 KB/sec
Aggregate data transfer rate:          1,391.67 KB/sec
Objects compressed by:                 0%
Elapsed processing time:                00:00:11

root@ / > lsfs /bruno
Name          Nodename    Mount Pt          VFS   Size   Options   Auto Accounting
/dev/bruno    --         /bruno            jfs   32768  rw        yes no

root@ / > dsmc backup image /dev/bruno
Tivoli Storage Manager
Command Line Backup Client Interface - Version 3, Release 7, Level 2.0 n
(C) Copyright IBM Corporation, 1990, 1999, All Rights Reserved.

Node Name: PUTNEY
Session established with server TSM010: AIX-RS/6000
  Server Version 3, Release 7, Level 2.0
  Server date/time: 02/16/00 14:20:05 Last access: 02/16/00 14:19:42

Backup Image Function Invoked.

ANS1228E Sending of object '/dev/bruno' failed
ANS1063E Invalid path specification

```

10.3.1.3 Running restore operations

To restore a file, a directory, or even the whole machine, you need to know two things: *what* you want to restore (file name, directory), and optionally *from when* (point-in-time), if you want to restore a file other than the most recent one.

Important: You do not need to know where the data actually is. When you request a file, Tivoli Storage Manager looks into its database and mounts the desired volume. The `restore` command obtains copies of backup versions of your files from a Tivoli Storage Manager server.

To restore files, specify the directories or selected files either at the command line or by clicking down through the directory structure displayed in the GUI. Both the GUI and the command line interface have search functions in case you don't know the exact filenames or directories needed for restore. You can also select the files from a list using the PICK option with the command line `restore` command, as shown in the following example:

```

tsm> restore /homeusr/* -subdir=yes -pick
Restore function invoked.

TSM Scrollable PICK Window - Restore

# Backup Date/Time File Size A/I File
-----
1. 02/10/00 17:06:28 512 A /homeusr/bruno
2. 02/10/00 17:06:28 512 A /homeusr/claudia
3. 02/10/00 17:06:28 512 A /homeusr/daniela
4. 02/10/00 17:06:28 512 A /homeusr/ecastro
5. 02/10/00 17:06:28 512 A /homeusr/lala
6. 02/10/00 17:06:28 512 A /homeusr/lucia
7. 02/10/00 17:06:28 512 A /homeusr/pp
8. 02/10/00 17:06:28 6 A /homeusr/daniela/finalreport
x 9. 02/10/00 17:06:28 11 A /homeusr/daniela/project_y2k
x 10. 02/10/00 17:06:28 15 A /homeusr/lala/cubano
11. 02/10/00 17:06:28 2.46 KB A /homeusr/lucia/db2setup.log
12. 02/10/00 17:06:28 593 A /homeusr/lucia/dsmerror.log
13. 02/10/00 17:06:28 173 A /homeusr/lucia/dtappint.log
14. 02/10/00 17:06:28 409 A /homeusr/lucia/ibmmpm2Pze7a.
0-----10-----20-----30-----40-----50-----60-----7
<U>=Up <D>=Down <T>=Top <B>=Bottom <R#>=Right <L#>=Left
<G#>=Goto Line # <#>=Toggle Entry <+>=Select All <->=Deselect All
<#:#+>=Select A Range <#:#->=Deselect A Range <O>=Ok <C>=Cancel

```

The options at the bottom of the screen allow you to select all or some of the file entries for restore.

You can restore files to the directory from which you backed them up, or you can restore files to a different location. Depending on what you want to restore, you may need to use the PRESERVEPATH option, which specifies how much of the source path you want to preserve and append to the destination path specified with the restore operation.

A file copy can be in one of three states: active, inactive, or expired. An active file copy is the most current copy of the file, an inactive file copy is a previous copy of the file (or any copy of a file which no longer exists on the client), and an expired file copy is a copy to be removed from the Tivoli Storage Manager server. Only active versions are presented for restore unless you use the INACTIVE or LATEST options: The INACTIVE option instructs Tivoli Storage Manager to display and restore an inactive backup if an active one is not available, and the LATEST option restores the most recent backup version of a file, even if that backup is inactive. You cannot restore an expired file version.

Table 37 gives you some examples on how to perform a restore operation using a UNIX backup/archive client:

Table 37. Restore command examples: UNIX

If you want to perform this then this is the client command
Restore a single file (/home/myfile1)	dsmc restore /home/myfile1
Restore a directory and all files from that level (/home)	dsmc restore "/home/"
Restore all files in a directory (/home) with all its subdirectories	dsmc restore "/home/" -subdir=yes
Restore a full directory (/home to /temp) with all subdirectories and write full path on destination	dsmc restore /home/myfile1 /temp -preservepath=complete -subdir=yes

If you want to perform this then this is the client command
Restore a Logical Volume image to an existing mounted filesystem (/bruno)	dsmc restore image /bruno
Restore a RAW device to an existing allocated area (/dev/test)	dsmc restore image /dev/test
Restore a backupset (armando_monthlyBS.12345) from the server	dsmc restore backupset armando_monthlyBS.12345
Restore a backupset using local tape unit (/dev/rmt0)	dsmc restore backupset /dev/rmt0 -location=tape
Restore a subset of files from a tape backupset	dsmc restore backupset /dev/mt0 /homeusr/*.log -location=tape

UNIX Users:

Only a UNIX root user can perform the backup and restore image operations.

Table 38 gives you some examples on how to perform a restore operation using a Windows backup-archive client:

Table 38. Restore command examples: Windows

If you want to perform this then this is the client command
Restore a single file (C:\work\myfile1)	dsmc restore c:\work\myfile1
Restore a directory and all files from that level (C:\work)	dsmc restore c:\work\
Restore all files from a directory (C:\work) with all its subdirectories	dsmc restore c:\work\ -subdir=yes
Restore a full subdirectory (C:\work to D:\temp) with all subdirectories, and write full path on destination	dsmc restore c:\work\ d:\temp -preservepath=complete -subdir=yes

The following is an example of a restore operation and its command output. The restored data does not exist on the client workstation before doing the restore, so we do not have a problem with collision options.

```
tsm> restore /homeusr/* -subdir=yes
Node Name: PUTNEY
Session established with server TSM010: AIX-RS/6000
Server Version 3, Release 7, Level 1.0
Server date/time: 02/10/00 17:13:46 Last access: 02/10/00 17:13:17

Restore function invoked.

ANS1247I Waiting for files from the server...
Restoring          512 /homeusr [Done]
Restoring          512 /homeusr/bruno [Done]
Restoring          512 /homeusr/claudia [Done]
Restoring          512 /homeusr/daniela [Done]
Restoring          512 /homeusr/ecastro [Done]
Restoring          512 /homeusr/lala [Done]
Restoring          512 /homeusr/lucia [Done]
Restoring          512 /homeusr/pp [Done]
Restoring           6 /homeusr/daniela/finalreport [Done]
Restoring          11 /homeusr/daniela/project_y2k [Done]
Restoring          15 /homeusr/lala/cubano [Done]
Restoring       2,523 /homeusr/lucia/db2setup.log [Done]
Restoring          593 /homeusr/lucia/dsmerror.log [Done]
Restoring          173 /homeusr/lucia/dtappint.log [Done]
Restoring          409 /homeusr/lucia/ibmnpm2Pze7a.log [Done]
Restoring          409 /homeusr/lucia/ibmnpm921_qa.log [Done]
Restoring          409 /homeusr/lucia/ibmnpmBAu_ ya.log [Done]
Restoring          408 /homeusr/lucia/ibmnpmHUFILMa.log [Done]
Restoring          409 /homeusr/lucia/ibmnpmYMyKUa.log [Done]
Restoring          409 /homeusr/lucia/ibmnpmeJlb7a.log [Done]
Restoring          409 /homeusr/lucia/ibmnpmfzuQya.log [Done]
Restoring          1,021 /homeusr/lucia/inst.log [Done]
Restoring          1,227 /homeusr/lucia/listener.log [Done]
Restoring           0 /homeusr/lucia/wpreinst.log [Done]

Restore processing finished.

Total number of objects restored:      24
Total number of objects failed:        0
Total number of bytes transferred:    8.57 KB
Data transfer time:                    0.00 sec
Network data transfer rate:            4,942.51 KB/sec
Aggregate data transfer rate:         14.17 KB/sec
Elapsed processing time:                00:00:00
```

The next example shows a partial restore from a local backupset which is on tape. In this case, some of the files already exist so we had the option to replace the existing copy, skip restoring that object or terminate the restore option. Note that a partial restore from a backupset is available only using the command line interface. If you use the GUI then you can only restore the complete backupset.

```

root@ /r1 > dsms restore backupset /dev/mt0 /homeusr/*.log -subdir=yes -location=tape

Tivoli Storage Manager
Command Line Backup Client Interface - Version 3, Release 7, Level 2.0 n
(C) Copyright IBM Corporation, 1990, 1999, All Rights Reserved.

Restore function invoked.

Node Name: PUTNEY

--- User Action is Required ---
File '/homeusr/lucia/dsmerror.log' exists

Select an appropriate action
  1. Replace this object
  2. Replace all objects that already exist
  3. Skip this object
  4. Skip all objects that already exist
  A. Abort this operation
Action [1,2,3,4,A] : 2
Restoring          593 /homeusr/lucia/dsmerror.log [Done]
Restoring          2,523 /homeusr/lucia/db2setup.log [Done]
Restoring           173 /homeusr/lucia/dtappint.log [Done]
Restoring           409 /homeusr/lucia/ibmnpm2Pze7a.log [Done]
Restoring           409 /homeusr/lucia/ibmnpm921_qa.log [Done]
Restoring           409 /homeusr/lucia/ibmnpmBAu.ya.log [Done]
Restoring           408 /homeusr/lucia/ibmnpmHUFILMa.log [Done]
Restoring           409 /homeusr/lucia/ibmnpmYMyKUa.log [Done]
Restoring           409 /homeusr/lucia/ibmnpmeJlb7a.log [Done]
Restoring           409 /homeusr/lucia/ibmnpmfzuQya.log [Done]
Restoring          1,021 /homeusr/lucia/inst.log [Done]
Restoring          1,227 /homeusr/lucia/listener.log [Done]
Restoring           0 /homeusr/lucia/wpreinst.log [Done]

Restore processing finished.

Total number of objects restored:      13
Total number of objects failed:        0
Total number of bytes transferred:    9.06 KB
Data transfer time:                    0.00 sec
Network data transfer rate:            12,719.96 KB/sec
Aggregate data transfer rate:          0.00 KB/sec
Elapsed processing time:                00:16:59
root@ / >

```

Here is an example of restoring from a backupset which is available on the Tivoli Storage Manager server. We first query the server to find out the name of the backupsets which have already been generated.

```

root@ /r1 > dsmc
Tivoli Storage Manager
Command Line Backup Client Interface - Version 3, Release 7, Level 2.0 n
(C) Copyright IBM Corporation, 1990, 1999, All Rights Reserved.

tsm> query backupset
Node Name: PUTNEY
Session established with server TSM010: AIX-RS/6000
Server Version 3, Release 7, Level 2.0
Server date/time: 02/16/00 15:23:37 Last access: 02/16/00 15:23:02

Backup Set Name          Generation Date  Retention Description
-----
1 ARMANDO_MONTHLYBS.40481 02/16/00 15:20:57 365 No Description

tsm> restore backupset armando_monthlybs.40481
Restore function invoked.

ANS1247I Waiting for files from the server...
Restoring          512 /r1/ [Done]
Restoring          512 /r1/lost+found [Done]
Restoring           81 /r1/KINDU_VICTORIA_TDP [Done]
Restoring       19,342 /r1/dsierror.log [Done]
Restoring    3,791,565 /r1/dsm [Done]
Restoring     867,493 /r1/dsm.jar [Done]
Restoring          925 /r1/dsm.opt [Done]
Restoring          925 /r1/dsm.opt.smp [Done]
Restoring         1,768 /r1/dsm.sys [Done]
Restoring         1,115 /r1/dsm.sys.smp [Done]
Restoring    1,436,927 /r1/dsmadmc [Done]
Restoring    1,925,002 /r1/dsmagent [Done]
Restoring    1,992,796 /r1/dsmc [Done]
Restoring    1,444,747 /r1/dsmcad [Done]
Restoring     282,115 /r1/dsmcliios.driv [Done]
Restoring     4,567 /r1/dsmerror.log [Done]
Restoring     1,860 /r1/dsmsched.log [Done]
Restoring    14,399 /r1/dsmsna.driv [Done]
Restoring    18,787 /r1/dsmstat [Done]
Restoring    1,368,068 /r1/dsmtca [Done]
Restoring          962 /r1/dsmwebcl.log [Done]
Restoring     537,622 /r1/mkkfe [Done]
Restoring     22,304 /r1/secnls.cat [Done]

Restore processing finished.

Total number of objects restored:      23
Total number of objects failed:        0
Total number of bytes transferred:    13.11 MB
Data transfer time:                    25.75 sec
Network data transfer rate:            521.62 KB/sec
Aggregate data transfer rate:          107.73 KB/sec
Elapsed processing time:                00:02:04

```


10.3.1.4 Running archive operations

The `archive` command archives a single file, selected files, or all files in a directory and its subdirectories on a Tivoli Storage Manager server. You use this command to save files that you want to keep in their present condition for a long period of time, and to control them by expiration date only. There is no version controlling for archives. You can also use archives to release storage space on your workstation and delete files as you archive them using the `DELETEFILES` option. Typically, you may use archives to save information that has either a legal requirement (for example, account information, annual reports, billing information, or annual customer reports); or an internal audit requirement (for example, application logs, user activity information, or employee files).

When you use the `archive` command with the `ARCHMC` option, you can select any available management class in the active policy set that has an archive copygroup. If you omit this option then the default management class is used. This enables you to select the retention period (according to the retention period in days specified in the archive copygroup) for all the data you are archiving. The output of each archive operation is called an archive package and is by default given a description of the time and date when the archive was made. You can override the default description by specifying your own text with the `DESCRIPTION` option. You can later use the assigned archive package descriptions to add additional files to an existing package. The package description can also be used to query for existing packages or to retrieve files from the package, without needing to know exactly what files were contained in it. The description is case-sensitive.

Table 39 gives you some examples on how to perform an archive operation using a UNIX backup/archive client:

Table 39. Archive command examples: UNIX

If you want to perform this then this is the client command
Archive all files from a directory (/home)	<code>dsmc archive /home/ -subdir=yes</code>
Archive files from different directories (/home and /temp)	<code>dsmc archive /home/ /temp/ -subdir=yes</code>
Archive files from a directory (/home) to a management class called Y1 (which is one year retention)	<code>dsmc archive /home/ -archmc=y1</code>
Create a new archive package called "DANCLA-FILES" from a directory (/home)	<code>dsmc archive /home/ -subdir=yes -description="DANCLA-FILES"</code>
Add files from a directory (/temp) to an archive package called "DANCLA-FILES"	<code>dsmc archive /temp/ -subdir=yes -description="DANCLA-FILES"</code>

Table 40 gives you some examples of how to perform an archive operation using a Windows backup/archive client:

Table 40. Archive command examples: Windows

If you want to perform this then this is the client command
Archive all files from a directory (C:\work)	<code>dsmc archive c:\work\ -subdir=yes</code>
Archive files from different directories (C:\work and C:\temp)	<code>dsmc archive c:\work\ c:\temp\ -subdir=yes</code>

If you want to perform this then this is the client command
Archive files from a directory (C:\work) to a management class called Y1 (which is one year retention)	dsmc archive c:\work\ -archmc=y1
Create a new package called "DANCLA-FILES" from a directory (C:\work)	dsmc archive c:\work\ -subdir=yes -description="DANCLA-FILES"
Add files from a directory (C:\temp) to an archive package called "DANCLA-FILES"	dsmc archive c:\temp\ -subdir=yes -description="DANCLA-FILES"

The following is an example of an archive operation and its command output:

```

tsm> archive /homeusr/* -subdir=yes -description="Year End 2000" -archmc=y1

Archive function invoked.

Directory-->          512 /homeusr [Sent]
Directory-->          512 /homeusr/bruno [Sent]
Directory-->          512 /homeusr/claudia [Sent]
Directory-->          512 /homeusr/daniela [Sent]
Directory-->          512 /homeusr/ecastro [Sent]
Directory-->          512 /homeusr/lala [Sent]
Directory-->          512 /homeusr/lucia [Sent]
Directory-->          512 /homeusr/pp [Sent]
Normal File-->         6 /homeusr/daniela/finalreport [Sent]
Normal File-->        11 /homeusr/daniela/project_y2k [Sent]
Normal File-->        15 /homeusr/lala/cubano [Sent]
Normal File-->       2,523 /homeusr/lucia/db2setup.log [Sent]
Normal File-->        593 /homeusr/lucia/dsmerror.log [Sent]
Normal File-->        173 /homeusr/lucia/dtappint.log [Sent]
Normal File-->        409 /homeusr/lucia/ibmmpm2Pze7a.log [Sent]
Normal File-->        409 /homeusr/lucia/ibmmpm921_qa.log [Sent]
Normal File-->        409 /homeusr/lucia/ibmmpmBAu.ya.log [Sent]
Normal File-->        408 /homeusr/lucia/ibmmpmHfLlMa.log [Sent]
Normal File-->        409 /homeusr/lucia/ibmmpmYMyKUa.log [Sent]
Normal File-->        409 /homeusr/lucia/ibmmpmeJlb7a.log [Sent]
Normal File-->        409 /homeusr/lucia/ibmmpmfzuQya.log [Sent]
Normal File-->       1,021 /homeusr/lucia/inst.log [Sent]
Normal File-->       1,227 /homeusr/lucia/listener.log [Sent]
Normal File-->         0 /homeusr/lucia/wpreinst.log [Sent]
Archive processing of '//homeusr/*' finished without failure.

Total number of objects inspected:      31
Total number of objects archived:      24
Total number of objects updated:        0
Total number of objects rebound:       0
Total number of objects deleted:        0
Total number of objects failed:         0
Total number of bytes transferred:     8.57 KB
Data transfer time:                     0.00 sec
Network data transfer rate:             19,611.69 KB/sec
Aggregate data transfer rate:           2.80 KB/sec
Objects compressed by:                  0%
Elapsed processing time:                 00:00:03

```

10.3.1.5 Running retrieve operations

The `retrieve` command obtains copies of archived files from the Tivoli Storage Manager server. You can specify either selected files or whole directories to retrieve archived files. The `DESCRIPTION` field allows you to search for the descriptions assigned to the files when they were archived.

Table 41 gives you some examples on how to perform an retrieve operation using a UNIX backup/archive client:

Table 41. Retrieve command examples: UNIX

If you want to perform this then this is the client command
Retrieve a single file (/home/myfile1)	<code>dsmc retrieve /home/myfile1</code>
Retrieve files from a directory (/home)	<code>dsmc retrieve /home</code>
Retrieve files from a directory (/home) with a description "project1998" and save them into another location (/temp)	<code>dsmc retrieve /home/ /temp/ -subdir=yes -description="project1998"</code>

Table 42 gives you some examples on how to perform a retrieve operation using a Windows backup/archive client:

Table 42. Retrieve command examples: Windows

If you want to perform this then this is the client command
Retrieve a single file (C:\work\myfile1)	<code>dsmc retrieve c:\work\myfile1</code>
Retrieve files from a directory (C:\work)	<code>dsmc retrieve c:\work\ -subdir=yes</code>
Retrieve files from a directory (C:\work) with a description "project1998" and save them into another location (C:\temp)	<code>dsmc retrieve c:\work\ c:\temp\ -subdir=yes -description="project1998"</code>

Note:

1. The description field is case-sensitive.
2. The same restore rules applies for the retrieve operations, when using the PRESERVEPATH option.

The following is an example of a retrieve operation and its command output. The retrieved data does not exist on the client workstation before the retrieve operation.

```

tsm> retrieve /homeusr/* -subdir=yes -descr="Year End 2000"
Node Name: PUTNEY
Session established with server TSM010: AIX-RS/6000
Server Version 3, Release 7, Level 1.0
Server date/time: 02/10/00 17:32:46 Last access: 02/10/00 17:30:10

Retrieve function invoked.

Retrieving          512 /homeusr/bruno [Done]
Retrieving          512 /homeusr/claudia [Done]
Retrieving          512 /homeusr/daniela [Done]
Retrieving          512 /homeusr/ecastro [Done]
Retrieving          512 /homeusr/lala [Done]
Retrieving          512 /homeusr [Done]
Retrieving          512 /homeusr/lucia [Done]
Retrieving           0 /homeusr/lucia/wpreinst.log [Done]
Retrieving          512 /homeusr/pp [Done]
Retrieving           6 /homeusr/daniela/finalreport [Done]
Retrieving          11 /homeusr/daniela/project_y2k [Done]
Retrieving          15 /homeusr/lala/cubano [Done]
Retrieving         2,523 /homeusr/lucia/db2setup.log [Done]
Retrieving          593 /homeusr/lucia/dsmerror.log [Done]
Retrieving          173 /homeusr/lucia/dtappint.log [Done]
Retrieving          409 /homeusr/lucia/ibmmpm2Pze7a.log [Done]
Retrieving          409 /homeusr/lucia/ibmmpm921_qa.log [Done]
Retrieving          409 /homeusr/lucia/ibmmpmBAu.ya.log [Done]
Retrieving          408 /homeusr/lucia/ibmmpmHUF1Ma.log [Done]
Retrieving          409 /homeusr/lucia/ibmmpmMYMyKUa.log [Done]
Retrieving          409 /homeusr/lucia/ibmmpmJ1b7a.log [Done]
Retrieving          409 /homeusr/lucia/ibmmpmfzuQya.log [Done]
Retrieving          1,021 /homeusr/lucia/inst.log [Done]
Retrieving          1,227 /homeusr/lucia/listener.log [Done]

Retrieve processing finished.

Total number of objects retrieved:      24
Total number of objects failed:         0
Total number of bytes transferred:     8.57 KB
Data transfer time:                     0.00 sec
Network data transfer rate:             5,038.39 KB/sec
Aggregate data transfer rate:           15.90 KB/sec
Elapsed processing time:                 00:00:00

```

10.3.1.6 Stopping a session

If the Tivoli Storage Manager backup/archive client is running in loop mode and resting at the `tsm>` prompt, then you have to enter `quit` to end the session. This terminates the connection with the Tivoli Storage Manager server and returns to the calling program, which is normally the operating system prompt.

```

tsm> quit
root@ / >

```

Note:

If you want to terminate a running backup/archive session, you can cancel the operation by pressing CTRL-C in the session window or ask the Tivoli Storage Manager administrator to issue the `cancel session` command. If you choose to cancel it by using CTRL-C while there is a restore operation in progress, the Tivoli Storage Manager administrator needs to cancel the restartable restore by issuing the `cancel restore` command from the Tivoli Storage Manager server console or from an administrative client.

10.3.2 GUI

Most end-users will find the user-friendly interface of the GUI more attractive and intuitive than the command-line. You may use it for all Tivoli Storage Manager client operations except for scheduled operations.

10.3.2.1 Starting a session

To start the GUI, you must execute the backup/archive client program by typing DSM in the operating system prompt. The GUI has a slightly different appearance on Windows and UNIX machines. Although we will use the Windows GUI in the next examples to show how to run client operations, keep in mind that the same functions apply for the UNIX version as well.

UNIX Users:

Some UNIX clients (currently AIX, HP-UX and Solaris) have an additional option for RAW devices in the Backup and Restore options if you have installed the image backup component. You will also notice two additional options in the backup selection criteria, such as "*Image backup*" and "*Incremental image (date only)*".

Figure 24 shows the main screen of the backup/archive client GUI interface on Windows:



Figure 24. Backup/archive client GUI: Windows

Figure 25 shows the main screen of the backup/archive client GUI interface on UNIX:

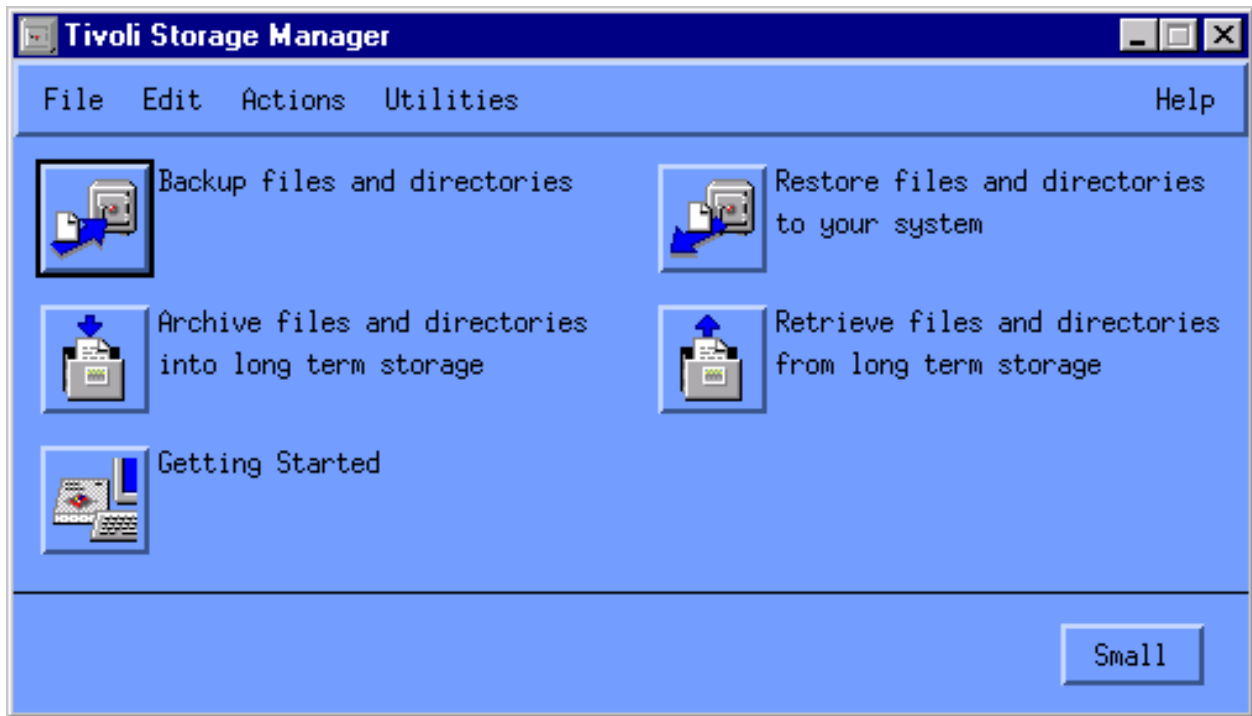


Figure 25. Backup/archive client GUI: UNIX

Depending on the features of your UNIX machine, for example, if it has the Common Desktop Environment (CDE) installed, you can add an option in the workplace screen so that it is easier to start Tivoli Storage Manager.

10.3.2.2 Running Backup Operations

To open the Backup window, click on **Backup** in the main window as shown in Figure 24. Figure 26 shows an example of a Backup window with some files selected.

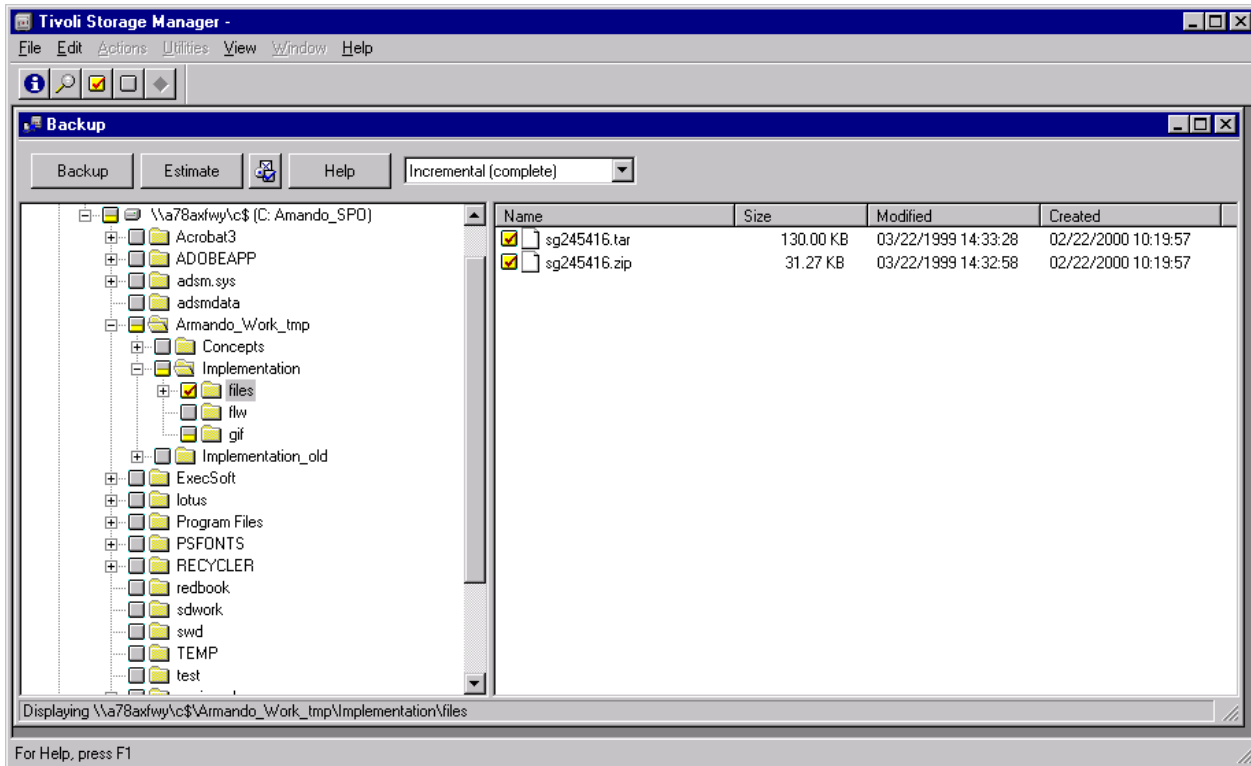


Figure 26. GUI backup window

To back up some of your files:

1. In the Backup window, expand the directory tree. Click on the plus sign (+) next to a directory or drive.
2. Click the selection boxes next to the drives or directories you want to back up.
3. Be sure that the drop-down list near the top of the Backup window displays the Incremental (Complete) choice.
4. Click on **Backup**. The Backup Status window appears.
5. You can view the progress of the backup in the Backup Status window. When it is finished, you get a completion dialog box.

10.3.2.3 Running restore operations

To open the Restore window, click on **Restore** in the main window as shown in Figure 24 on page 194. Figure 27 shows an example of a Restore window with some files selected. Figure 28 on page 200 and Figure 29 on page 201 show two backupset restore examples (one for a server-based restore and one for a local, LAN-free situation).

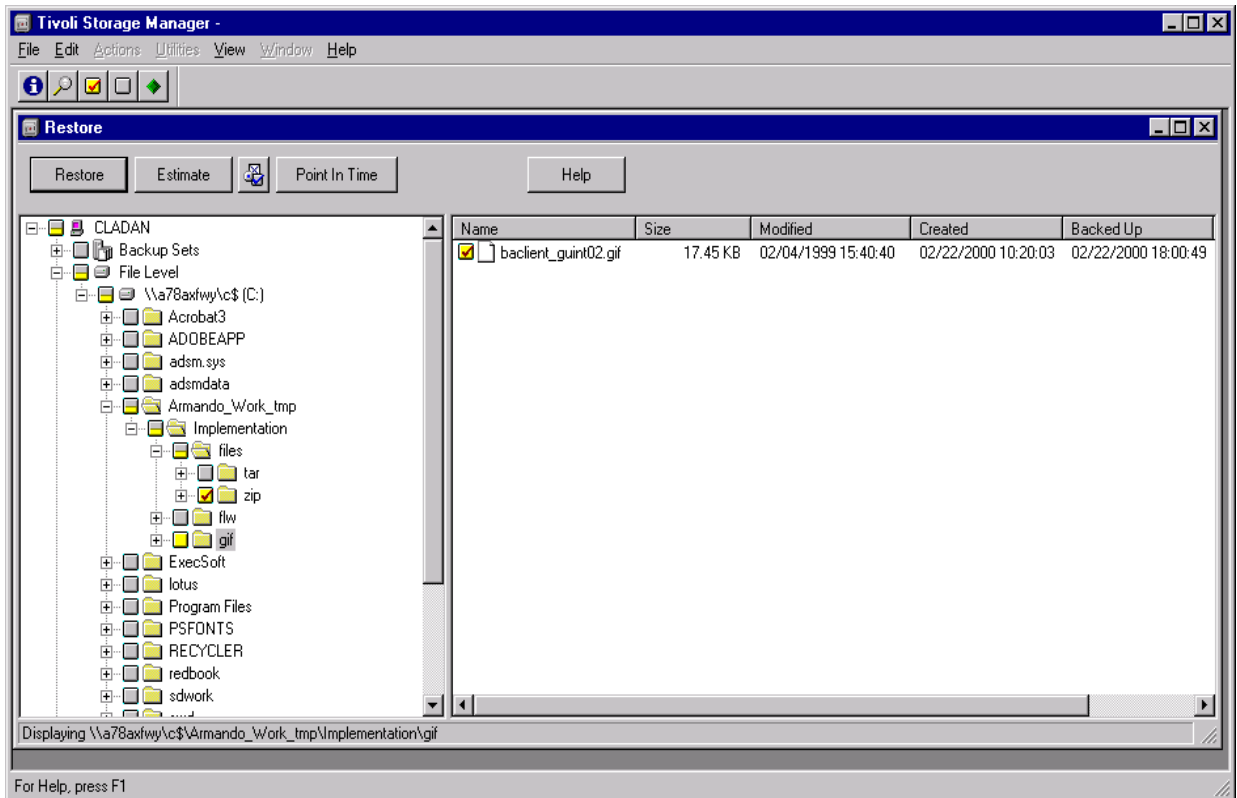


Figure 27. GUI restore window

The Restore window displays the directories and files that you previously backed up to the Tivoli Storage Manager server. Use this window to select the objects that you want to restore. You can also request Tivoli Storage Manager to run an estimate for the restore or change processing options.

There are several restore operations that you can perform:

- Restore backup versions from directory tree
- Restore backup versions by name
- Restore backup versions by filtering the directory tree
- Restore someone else's backup versions
- Restore inactive backup versions
- Run a point-in-time restore
- Work with restartable restore sessions
- Restore a backupset
- Restore a Logical Volume (UNIX only)

Restore backup versions from directory tree

1. In the Restore window, expand the directory tree. Click the plus sign (+) next to a directory or drive.
2. Click the selection boxes next to the files or directories that you want to restore.

3. Click **Restore**. The Restore Destination window appears.
4. Fill in the Restore Destination window, then click **Restore** to start the restore.
5. You can view the progress of the restore in the Restore Status window. When it is finished, you get a completion dialog box.

Restore backup versions by name

1. In the Restore window, expand the directory tree. Click the plus sign (+) next to a directory or drive.
2. Highlight the drive or directory you want to work with.
3. Click the Search tool on the tool bar.
4. Enter your search criteria in the Find Files window, then click **Search**.
5. When a list of files appears in the Matching Files window, click the selection boxes next to the files you want to restore. Your selections appear in the Restore by Tree window.
6. Click **Restore**. The Restore Destination window appears.
7. Fill in the Restore Destination window, then click **Restore** to start the restore.
8. You can view the progress of the restore in the Restore Status window.

Restore backup versions by filtering the directory tree

1. In the Restore window, expand the directory tree. Click the plus sign (+) next to a directory or drive.
2. Highlight the drive or directory you want to work with.
3. Click the Search tool on the tool bar.
4. Enter your criteria in the Find Files window, then click the Filter button. Close the window. The directory tree displays only those directories that contain the files that match your criteria.
5. Click the selection boxes next to the files or filtered directories that you want to restore.
6. Click **Restore**. The Restore Destination window appears.
7. Fill in the Restore Destination window, then click **Restore** to start the restore.
8. You can view the progress of the restore in the Restore Status window.

Restore someone else's backup versions

1. From the main window, access the other user's stored data
 1. Select the Access Another User choice from the Utilities menu. The Access Another User window appears.
 2. Type the user's node name in the Node Name field.
 3. Click the Set button. You can display another user's stored data only if the other user has authorized you to do so in the Authorization List window.
2. Restore the other user's backup versions using one of the following methods:
 - Restoring backup versions by name
 - Restoring backup versions from the directory tree
 - Restoring backup versions by filtering the directory tree

Restore inactive backup versions

1. In the Restore window, select the Display Active/Inactive Files choice from the View menu. The inactive file icon appears next to file names of inactive backup versions in the file list.
2. Expand the directory tree. Click the plus sign (+) next to a directory or drive.
3. Click the selection boxes next to the directories or files that you want to restore.
4. Click the Restore button. The Restore Destination window appears.
5. Fill in the Restore Destination window, then click **Restore** to start the restore.
6. You can view the progress of the restore in the Restore Status window.

Run a point-in-time restore

1. In the Restore window, click on **Point in Time Restore**.
2. Fill in the Point in Time Restore window and click **OK**. The point in time that you specified appears in the Point in Time display field in the Restore window.
3. Display the objects you want to restore. To display the objects, you can search for an object by name, filter the directory tree, or work with the directories in the directory tree.
4. Click the selection boxes next to the objects you want to restore.
5. Click **Restore**. The Restore Destination window appears.
6. Fill in the Restore Destination window, then click **Restore** to start the restore.
7. You can view the progress of the restore in the Restore Status window.

Work with restartable restore sessions

To restart a restore session that was interrupted:

1. From the main window, select the Restartable Restores choice from the Actions menu. If there are any restartable restore sessions available, the Restartable Restores window appears.
2. Select a restore session that you want to restart from the list.
3. Click **Restart**.

To delete a restore session that was interrupted:

1. From the main window, select the Restartable Restores choice from the Actions menu. If there are any restartable restore sessions available, the Restartable Restores window appears.
2. Select an object that you want to delete from the list.
3. Click **Delete**.

Restoring from a backupset

The backupset is a collection of files that can be either stored on the server or locally. If the backupset files are stored on the server side, then the required volumes will be mounted and all files from that backupset will be transferred through the LAN to your local client machine. On the other hand, you could use the unique features of the backupset to either restore the files from a local device (that is, 8MM tape in UNIX) or you could use the FILE device class to make it possible to restore them in any client machine (respecting cross-platform dependencies).

To restore a backupset that is on the server:

1. From the main menu, select **Restore (Figure 28)**
2. Select **Backup Sets —> Server**
3. Select the desired backupset to restore (expand the inner levels to find the backupsets)

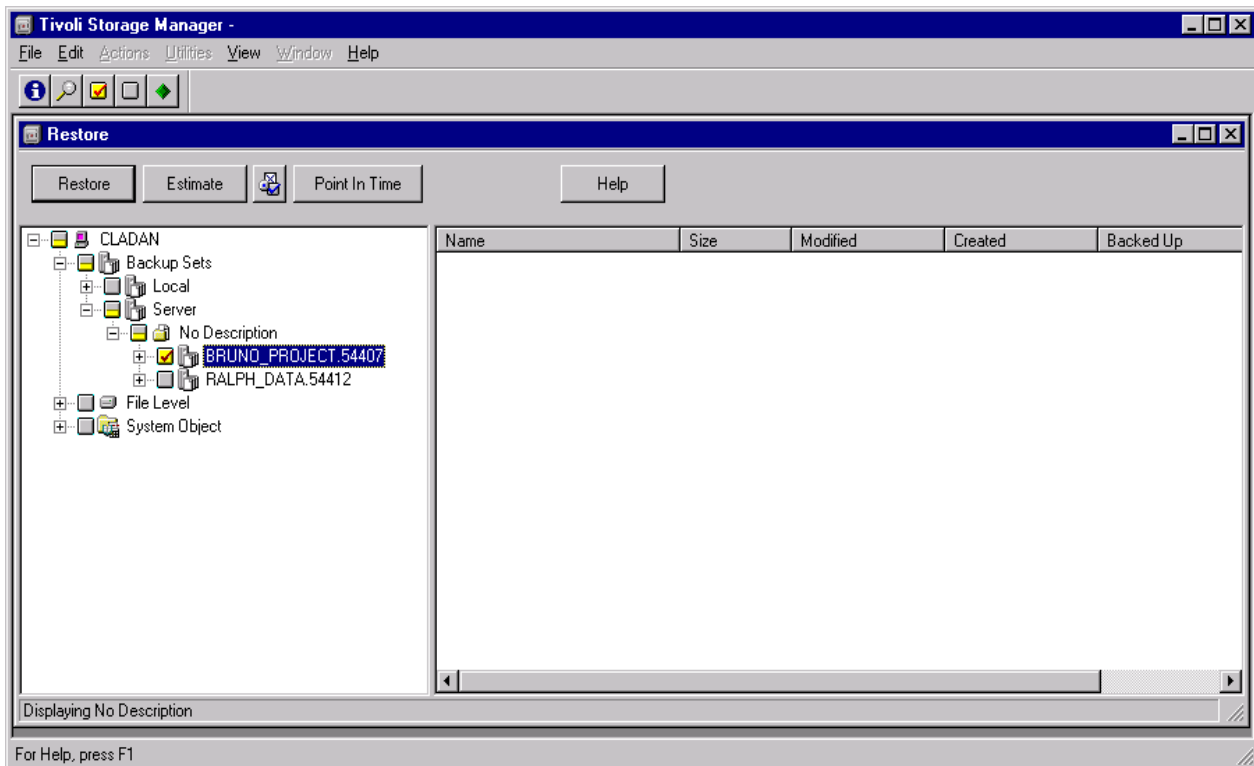


Figure 28. GUI backupset restore window from server data

To restore from a local backupset

1. From the main menu, select **Restore (Figure 29)**
2. Select **Backup Sets —> Local**
3. Specify where the backupset data is. Depending on the client features, you might be able to restore from tape (for example, a UNIX tape as /dev/rmt0 or a Windows NT tape as \\.\Tape0) or by using a FILE format.
4. Select the desired backupset to restore (expand the inner levels to find the backupsets)

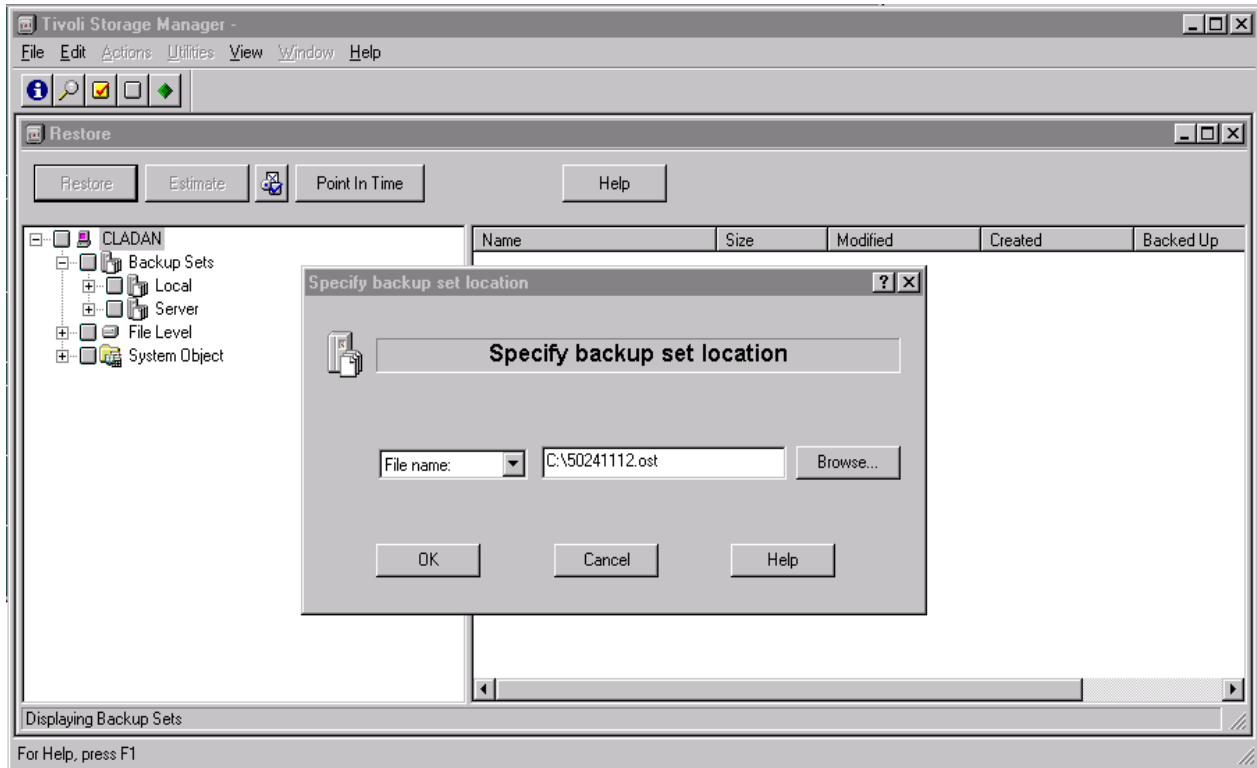


Figure 29. GUI for a local backupset restore

10.3.2.4 Running archive operations

To open the Archive window, click **Archive** in the main window as shown in Figure 24 on page 194. Figure 30 shows an example of an Archive window with some files selected.

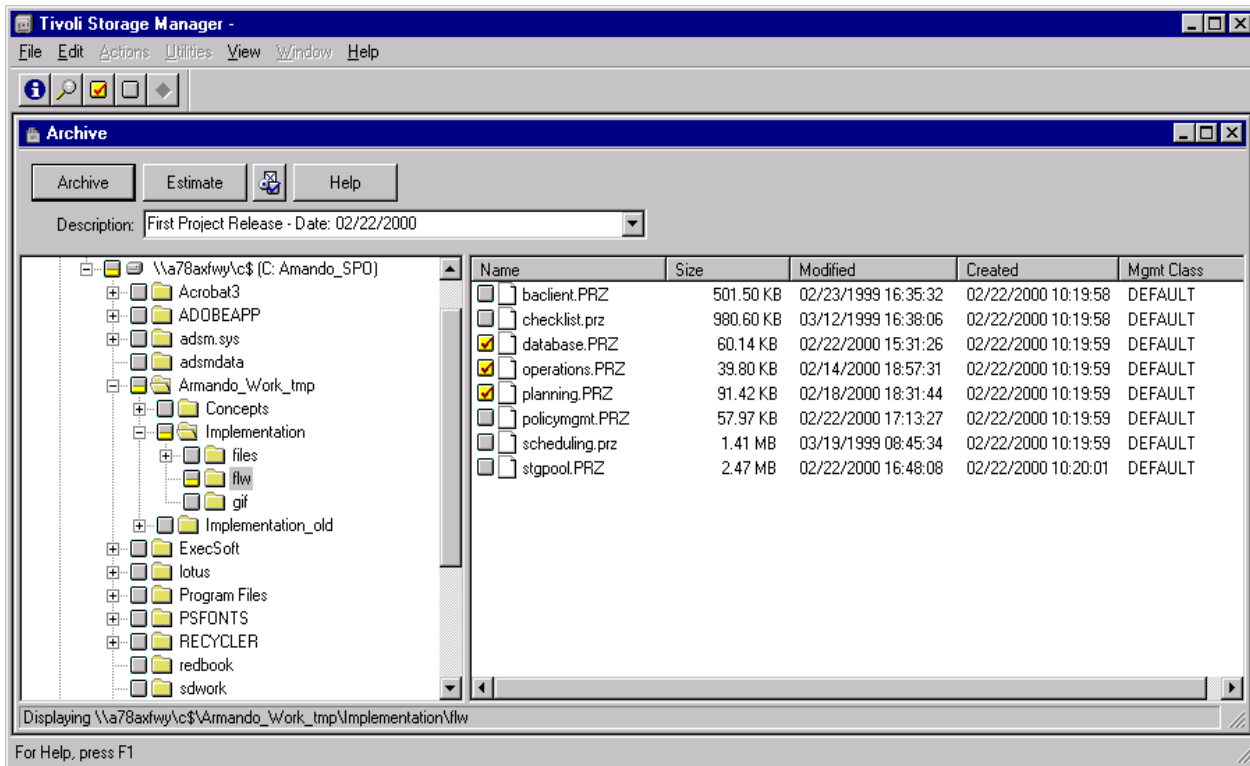


Figure 30. GUI archive window

The Archive window displays the directories and files on your workstation. Use this window to select the directories and files you want to archive.

Tivoli Storage Manager requires an archive description for all archived files. An archive description identifies data through a meaningful description that you can use later to identify files and directories. You can enter as many as 255 characters to describe your archived data. If a description is not entered, Tivoli Storage Manager assigns a default archive description.

When you select the archive function from the backup/archive GUI, Tivoli Storage Manager displays a list of all previously used archive descriptions. You can use these displayed archive descriptions on future archives, either to edit so as to use a standard format, or as is, in which case the files will be added to the package which already exists.

You can archive specific files or entire directories from the directory tree. You can assign a unique description for each group of files you archive. This group of files is called an archive package. Archiving files and directories in such a way makes retrieval easier. You can archive and retrieve files or directories as one entity, retrieve individual files, add files to an existing package using the archive description, or delete files from a package.

To archive your files:

1. In the Archive window, expand the directory tree. Click the plus sign (+) next to a directory or drive.
2. Click the selection boxes next to the files or directories that you want to archive.

3. Type a description, accept the default description, or select an existing description for your archive package in the Description combination box.
4. Click **Archive**. The Archive Status window appears as shown in Figure 30.
5. You can view the progress of the archive in the Archive Status window. When it is finished, you get a completion dialog box.

10.3.2.5 Running retrieve operations

To open the Retrieve window, click on **Retrieve** in the main window as shown in Figure 24 on page 194. Figure 31 shows an example of a Retrieve window with some files selected.

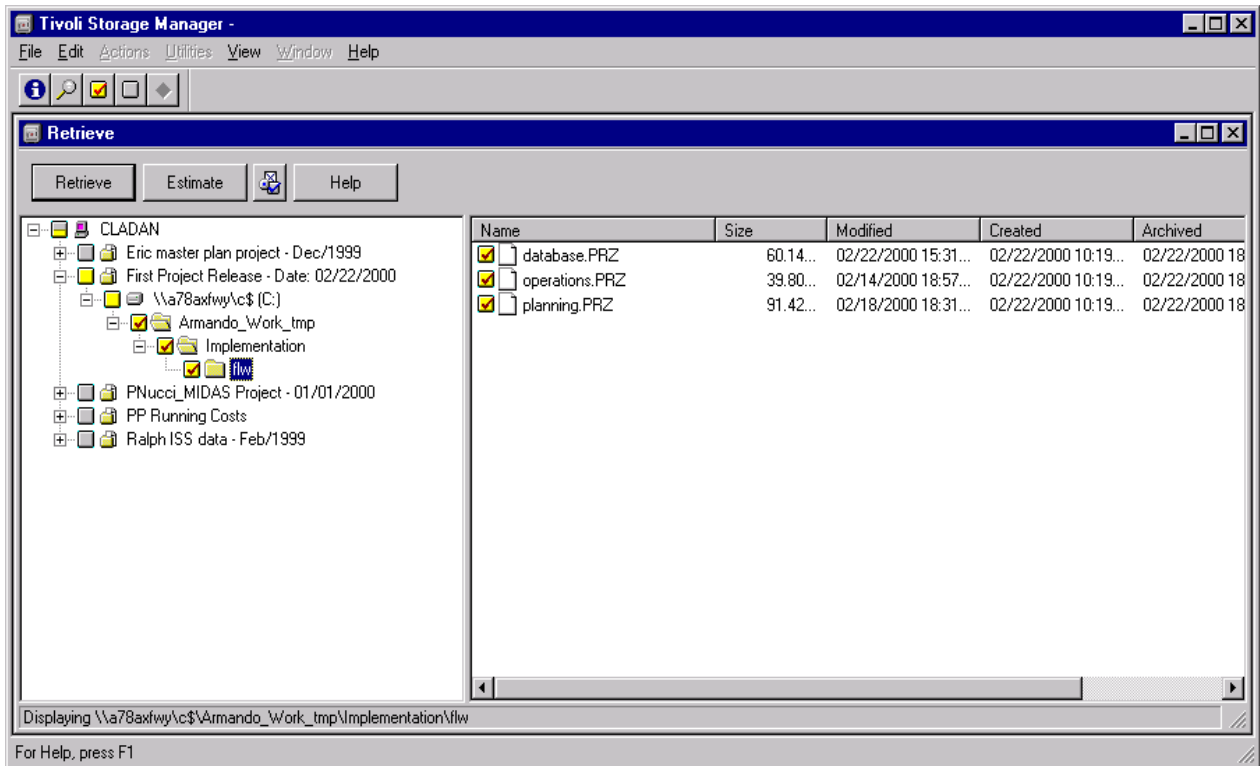


Figure 31. GUI retrieve window

The Retrieve window displays the packages that you archived to the Tivoli Storage Manager server. Use this window to select the objects you want to retrieve.

To retrieve your archive copies:

1. In the Retrieve window, expand the directory tree. Click the plus sign (+) next to an object you want to expand. If you are connected to the Tivoli Storage Manager server, the objects on the tree are grouped by archive package description.
2. Click the selection boxes next to the objects that you want to retrieve in the tree view or file view.
3. Click **Retrieve**. The Retrieve Destination window appears.
4. Fill in the Retrieve Destination window, then click **Retrieve** to start the operation.

5. You can view the progress of the retrieve in the Retrieve Status window. When it is finished, you get a completion dialog box.

You can also retrieve another user's archive copies to your local drives if the other user has authorized you to do so.

To retrieve another user's archive copies:

1. From the main window, access the other user's stored data
 1. Select the **Access Another User** choice from the Utilities menu. The **Access Another User** window appears.
 2. Type the user's node name in the **Node Name** field.
 3. Click **Set**.
2. Retrieve the other user's archive copies using the previous method explained.

10.3.2.6 Stopping a session

To stop a GUI session, you must select **File** → **Exit** from the pull-down menu in the main window.

10.4 Client scheduler

A Tivoli Storage Manager administrator can create schedules for the clients so that they perform defined tasks automatically on a regular basis. For example, you can define schedules to automatically run an incremental file backup at the end of each day, or to archive a particular directory every Friday. This procedure, known as central scheduling, is a cooperative effort between the server and the client node. You associate clients with one or more schedules that are part of the policy domain maintained in the Tivoli Storage Manager database. You use the central scheduler on the server, and you start the client scheduler on the workstation. Once you start the client scheduler, further intervention is not necessary.

Note:

The schedule start time is based on the server clock, not the client clock. This makes it easier for you to manage schedule start times for clients in different time zones.

With client scheduling, you can also:

- Display information about available schedules.
- Display information about work that the schedule has completed.
- Modify scheduling options in the client options file.

Note:

Install the Tivoli Storage Manager command-line client and ensure that you have the client to server communication working correctly before you start the client scheduler. You cannot start the client scheduler from the Tivoli Storage Manager GUI.

10.4.1 Starting the client scheduler

You can start the Tivoli Storage Manager client scheduler either manually or with each system start. Tivoli Storage Manager does not recognize changes to the client options files while the client scheduler is running. If you make changes to these files while the client scheduler is running, and you want Tivoli Storage Manager to use the new values immediately, stop the client scheduler and restart it.

UNIX Users:

The client scheduler can only be started by a Tivoli Storage Manager-authorized user. This is a user who has administrative authority for the Tivoli Storage Manager client on a workstation. This user changes passwords, performs open registrations, and deletes filespaces.

10.4.1.1 Manual start

You may perform a manual startup for the Tivoli Storage Manager scheduler when you accidentally stopped it, when you want to re-start it to pick up a changed client option, or when you want to force it to run immediately, for example during testing of a new function.

To start the client scheduler on your client node and connect to the server scheduler, enter the `dsmc schedule` command from the operating system prompt. You cannot start the scheduler in loop mode. You will be prompted to type in the password unless you have set `PASSWORDACCESS GENERATE` option set in client option file (`dsm.opt`). The scheduling requirement is one of the strongest reasons for recommending the `PASSWORDACCESS GENERATE` option. If you don't specify this you will always be prompted for a password unless you include it in a script file. Either possibility is undesirable - there is probably no-one around to enter the password and entering it in a script in plain text exposes a security risk.

```
root@ / > dsmc schedule
Tivoli Storage Manager
Command Line Backup Client Interface - Version 3, Release 7, Level 1.0
(C) Copyright IBM Corporation, 1990, 1999, All Rights Reserved.

Querying server for next scheduled event.
Node Name: PUTNEY
Session established with server TSM010: AIX-RS/6000
  Server Version 3, Release 7, Level 1.0
  Server date/time: 02/11/00 14:42:19 Last access: 02/11/00 14:42:00

Next operation scheduled:
-----
Schedule Name:      INCR_DAILY
Action:             Incremental
Objects:
Options:
Server Window Start: 14:31:05 on 02/12/00
-----
Waiting to be contacted by the server.
```

When you start the client scheduler, it runs continuously until you close the window, therefore ending the process, or log off your system.

UNIX

To run the schedule command in background and to keep the client scheduler running even if you log off your UNIX system, enter the following command from the UNIX shell:

```
root@ / > nohup dsmc schedule 2> /dev/null &
```

Windows

You may create a batch file, for example, `sched.bat`, to perform the client scheduler start, so that there will be no need to type the Tivoli Storage Manager command. The following commands should accomplish this:

```
rem
rem BATCH file for Tivoli Storage Manager client scheduler start
rem
cd c:\Program Files\Tivoli\TSM\baclient
set DSM_CONFIG=c:\Program Files\Tivoli\TSM\baclient\dsm.opt
set DSM_DIR=c:\Program Files\Tivoli\TSM\baclient
dsmc schedule -password=secret
```

Figure 32. Batch file for scheduler start on windows

Note in this case how the password had to be entered into the file because `PASSWORDACCESS GENERATE` was not used in the options file. You can also specify the `dsmc schedule` command in your startup folder so that every time you start Windows, the client scheduler starts.

10.4.1.2 Automatic start

We normally recommend automatically starting the Tivoli Storage Manager scheduler client. This is to ensure that the server can contact the client to execute regular schedules, especially if you have an enterprise policy to have Tivoli Storage Manager machines backed up at some point in the day.

If you rely on manual startup, this can be easily forgotten if the owner of the client system has to reboot it for some reason, resulting in a missed backup schedule.

AIX

If you want to start the client scheduler when you start your operating system, set the `PASSWORDACCESS` option parameter to `GENERATE` in your client system options file, and include the following command in the `/etc/inittab` file:

```
tsm::once:/usr/bin/dsmc sched > /dev/null 2>&1 # TSM scheduler
```

The scheduler starts the next time the system is rebooted. On other UNIX variants enter a similar command in the startup files.

Windows NT

You can register the Tivoli Storage Manager scheduler as a service under Windows NT using the `DSMCUTIL` program, so that it becomes part of the operating system startup procedure.

The following is an example of how to manually install the client scheduler named "TSM Scheduler" for node CLADAN on a Windows NT machine. For further information on the syntax of this utility program, check the `dsmcutil.txt` file in the Tivoli Storage Manager backup/archive client directory:

```

C:\Program Files\Tivoli\TSM\baclient> dsmcutil install /name:"TSM Scheduler"
/optfile:"c:\program files\tivoli\tsm\baclient\dsm.opt"
/clientdir:"c:\program files\tivoli\tsm\baclient" /node:CLADAN /password:secret
/eventlog:yes /autostart:yes

TSM Windows NT Client Service Configuration Utility
Command Line Interface Version 3.00.a
Last Updated Feb  2 2000
TSM Api Version 3.7.2

Command: Install TSM Client Service
Machine: A78AXFWY (Local Machine)

Installing TSM Client Service:

      Machine      : A78AXFWY
      Service Name  : TSM Scheduler
      Client Directory : c:\program files\tivoli\tsm\baclient
      Automatic Start : Yes
      Logon Account  : LocalSystem

The service was successfully installed.

Creating Registry Keys ...

Updated registry value 'ImagePath' .
Updated registry value 'EventMessageFile' .
Updated registry value 'TypesSupported' .
Updated registry value 'OptionsFile' .
Updated registry value 'EventLogging' .
Updated registry value 'ClientNodeName' .
Updated registry value 'ADSMClientKey' .
Updated registry value 'TSM Scheduler' .
Authenticating TSM password for node CLADAN ...

Connecting to TSM Server via client options file 'c:\program files\tivoli\tsm\baclient\dsm.o
Password authentication successful.

The registry password for TSM node CLADAN has been updated.

Starting the 'TSM Scheduler' service .....

The service was successfully started.

Generating registry password ...

```

There are a couple of other commands that you can issue in the DSMCUTIL utility, such as *install cad*, *install remoteagent*, *query*, *showpw*, *remove*, *update*, *list* and *updatepw*. For further reference, please see *Tivoli Storage Manager Installing the Clients*, SH26-4102. The next example shows a query execution:

```

C:\PROGRA~1\Tivoli\TSM\baclient> dsmcutil query /name:"tsm scheduler"

TSM Windows NT Client Service Configuration Utility
Command Line Interface Version 3.00.a
Last Updated Feb  2 2000
TSM Api Verison 3.7.2

Command: Query TSM Client Service Parameters
Machine: A78AXFWY (Local Machine)

Connecting to service 'tsm scheduler' ...

Service Configuration/Status:

Service Name      : tsm scheduler
Logon Account     : LocalSystem
Start Type        : Auto
Current Status    : Started

TSM Client Service Registry Settings:

Client Service Type = Client Scheduler Service
Client Directory    = c:\program files\tivoli\tsm\baclient\dsmcsvc.exe
Options file        = c:\program files\tivoli\tsm\baclient\dsm.opt
Event Logging       = YES
TSM Client Node     = CLADAN
Comm Protocol       = (value not currently set)
Server              = (value not currently set)
Server Port         = (value not currently set)
Schedule Log        = (value not currently set)
Error Log           = (value not currently set)
MSCS Enabled Node   = (value not currently set)
Cluster group name  = (value not currently set)

```

10.4.2 Stopping the scheduler

If you want to terminate a running client scheduler, you can cancel the operation by pressing CTRL-C in the session window, or you can ask the Tivoli Storage Manager administrator to issue the `cancel session` command (The administrator may need to cancel all multi-threaded sessions for the client). If you choose to cancel it by using CTRL-C while there is a restore operation in progress, the Tivoli Storage Manager administrator needs to cancel the restartable restore by issuing the `cancel restore` command from the Tivoli Storage Manager server console or from an administrative client.

On Windows NT, you can terminate the client scheduler by stopping its service. To change the status of the scheduler service, you have to go to **Settings** → **Control Panel** → **Services**, select the Tivoli Storage Manager client scheduler service and stop it.

10.5 Web client usage

The Tivoli Storage Manager web backup/archive client is a Java client interface that enables authorized users to perform remote operations on a backup/archive client system.

10.5.1 Remote client GUI

The web client is a Java applet that provides a remote client GUI for Tivoli Storage Manager backup/archive clients. Authorized users can access the web client remotely using a Web browser. The graphic is the initial client hub window displayed in the Web browser when the web client is accessed.

The web client can be used from any workstation running a Web browser with support for Java 1.1.6 or higher. Netscape 4.06 has the required Java support. Netscape 4.03, 4.04, and 4.05 requires a JDK upgrade for this support. This upgrade can be obtained from the Netscape homepage. Microsoft's Internet Explorer 4.01 has the required Java support. If the browser does not have the correct level of Java support, the client hub window is displayed as a grey box.

10.5.2 Client functions

The web client supports most of the backup/archive client functions. It can be used to perform all file-based functions such as backup, restore, archive, and retrieve operations and also supports the backupset features. For UNIX-based clients, the web interface also supports Logical Volume and RAW operations. For file-based restore operations, both active and inactive files can be restored, and the point-in-time restore function is supported. The functions not currently supported include cross client restores and the file find and search functions.

The Windows NT web client can back up and restore the NT registry. The NetWare web client can backup and restore the NetWare 3.x bindery and NetWare 4.x Directory Services (NDS) and, for the first time, provides a NetWare client GUI.

The web client is another client interface. It can be used only to perform backup/archive client operations. If used from a Web browser on a remote workstation, with a user ID with client access privilege (see 10.5.3) no local access to backup or archive data is provided on that remote workstation. Data cannot be restored or retrieved locally; it can only be restored to or retrieved to the client workstation that owns the data. A Web browser can also be used locally on the client workstation to invoke the web client as an alternative interface to the backup/archive client command line interface or GUI.

10.5.3 Access authorization

The web client is an interface in addition to the client's existing GUI and command line interface. Use of the web client interface is authenticated whenever backup, restore, archive, or retrieve web client functions are performed. Authentication of the web client interface is separate, and independent, from authentication between the client node and the server.

An administrator userid is required to use the web client. This administrator userid, and associated password, is used to authenticate that the user has sufficient authority to perform remote client functions. Two administrative authorities are provided to enable this authentication:

- Client owner
- Client access

These authorities can be used to enable usage of the web client interface for backup/archive client owners and helpdesk personnel. A user ID with client access privilege can only restore data back to its original system. A user ID with client owner privilege can restore data to another system if required.

The web backup/archive client connects to a Tivoli Storage Manager client via the Web to present an interface to remotely perform the functions of the backup/archive client on that client. Using a Web browser means that you actually do not have to have any Tivoli Storage Manager code on the Web browser workstation.

This interface is very useful because it provides all the file-based, backupset and UNIX Logical Volume functions remotely and without authorization problems.

10.5.4 Starting the web client

The web client consists of two processes on the client workstation: the client acceptor and remote client agent.

10.5.4.1 Client acceptor

The client acceptor is an HTTP daemon that serves the web client Java applet to the Web browser. The name of the executable is DSMCAD.

On AIX and other UNIX clients, it should be run as a daemon.

To have the web client automatically available when the system is started, you have to edit the /etc/inittab file (on AIX - use the equivalent startup file on other UNIX variants) and add the following entry:

```
dsmcad:2:once:/usr/tivoli/tsm/client/ba/bin/dsmcad >/dev/console 2>&1 #Start  
the TSM WEB agent
```

For Windows NT it is installed and run as a service. By default, the installation procedure registers the Web Client Acceptor as a manual service. Therefore, if you want to make it an automatic service, you have to go to **Settings —> Control Panel —> Services**, and change the startup setting of the Web Client Acceptor to automatic.

For Windows 95 and 98, the DSMCAD program must be executed in a DOS window.

For NetWare, DSMCAD is an NLM that should be loaded as part of the NetWare startup.

10.5.4.2 Remote client agent

The remote client agent performs the client functions initiated with the web client interface. The name of the executable is DSMAGENT. The agent does not have

to be running all the time. The acceptor daemon starts the agent when client functions are initiated through the web client interface.

10.5.4.3 Web client

To access the web backup/archive client, start your Web browser. In the Location or Address field, enter the URL:

`http://<machinename>:1581`

Where, <machinename> is the TCP/IP DNS name or dotted IP address of your Tivoli Storage Manager client machine and 1581 is the TCP/IP port number set by the HTTPPORT option.

Figure 33 shows you the main web client screen.

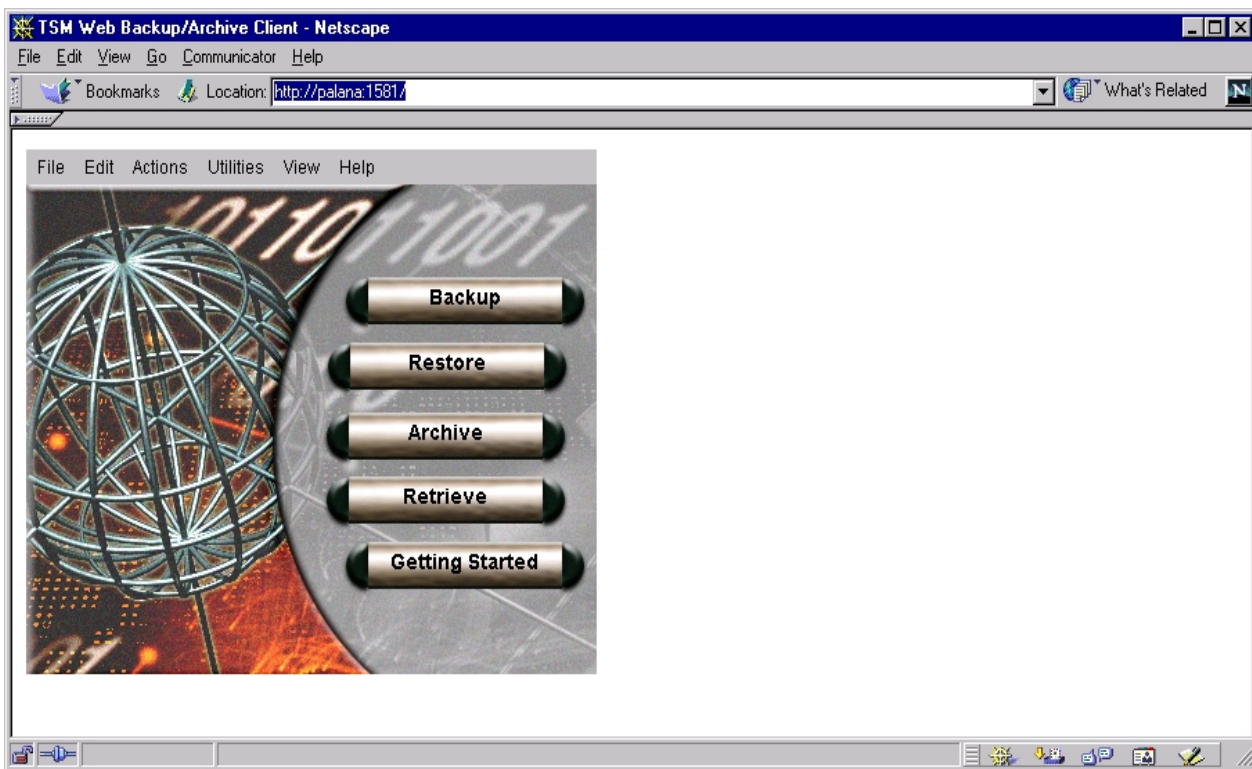


Figure 33. Web backup/archive client main window

When you try to connect to a client by using the web client interface, Tivoli Storage Manager validates the administrative ID you supply. You are then granted access to the web client interface, the client is connected to the Tivoli Storage Manager and the main window displays. Figure 34 shows a login dialog box for a web client session.



Figure 34. Web client login dialog box

10.5.5 Stopping the remote agent services

In general, there is no need to stop the remote agent services manually.

On Windows NT, if you wish to disable administrators from accessing the web client, you can prevent the service from starting by disabling it. To do that, you have to go to **Settings** —> **Control Panel** —> **Services** and then change the startup settings to disabled.

In a UNIX environment, you can stop the DSMCAD process by using the `KILL` command:

```
root@ / > ps -ef | grep dsmcad
root 95890 97892  3 13:15:08 pts/6  0:00 grep dsmcad
root 96492    1  0 13:04:54   -  0:00 ./dsmcad

root@ / > kill -15 96492
root@ / >
root@ / > ps -ef | grep dsmcad
root 95894 97892  3 13:15:36 pts/6  0:00 grep dsmcad
```

Chapter 11. Scheduling

Scheduling is at the heart of efficient Tivoli Storage Manager operations. Having a comprehensive and integrated set of schedules means that the server runs with very little requirement for intervention during normal operations.

Tivoli Storage Manager scheduling is divided into two different categories: administrative and client scheduling. The two categories differ in two key areas:

- Execution location: An administrative schedule performs some action on the Tivoli Storage Manager server, while the client schedule can only execute on the Tivoli Storage Manager client.
- Domain privilege: An administrative schedule can only be managed by an administrator with system privilege, while the client schedule can be managed by an administrator with policy privileges in only one domain. The granularity provided by this feature can be very useful when distributing management control across a large enterprise.

For both types of schedules, there are four key pieces of information:

- A command that needs to be executed
- When the command executes
- How long the command takes to complete
- How often the command needs to be repeated

The command that you run may be an incremental backup (client schedule) or a storage pool migration (administrative schedule) which should be run every day at a particular time. You also have to estimate how long the command will run so that you can synchronize your schedules and balance the load on the server.

This chapter provides scheduling rules of thumb and recommendations that will give you a good base configuration.

11.1 The wheel of life

Figure 35 shows the recommended series of operations that should occur in a Tivoli Storage Manager environment and the sequence in which those operations should occur. The circle represents a clock and the tick marks indicate the hours of the day. The actual start time and duration of the various operations are subject to your scheduling requirements, but not the sequence.

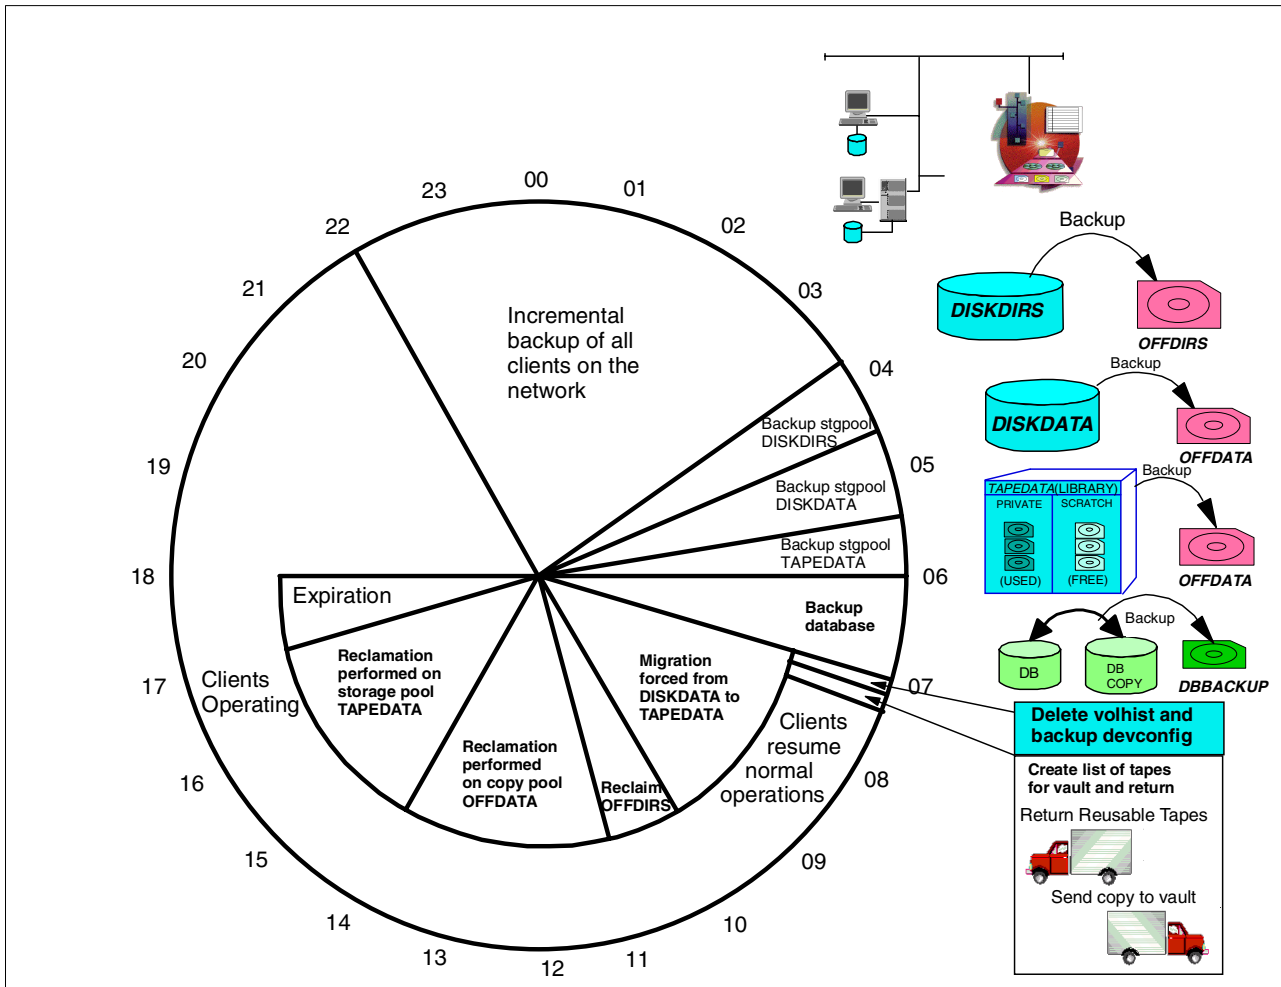


Figure 35. Scheduling of operations

11.2 Administrative schedules

An administrative schedule is a directive to trigger some sort of action on the Tivoli Storage Manager server. It consists of a command or sequence of commands and some details on when the actions should happen. Any actions that are used on a regular basis to manage the Tivoli Storage Manager environment should be defined as an administrative schedule.

In 1.2, “Redbook support material” on page 2, we show how to load a predefined macro into Tivoli Storage Manager. The macro `mac.schedules`, which we provide to define schedules for our redbook environment, is shown in B.1.4, “Define schedules” on page 289.

The following command illustrates the syntax for an administrative schedule. This command defines a schedule named `BACKUP_DATABASE` that specifies a full database backup, and that will start at 7 a.m. every morning.

```

tsm: TSM010> define schedule backup_database type=administrative \
cont> cmd="backup db devclass=coffsite type=full" starttime=07:00 active=yes \
cont> description="Database backup"
ANR2577I Schedule BACKUP_DATABASE defined.

```

We have assembled a series of administrative schedules that should help you set up a Tivoli Storage Manager environment that will minimize user intervention while providing a high level of data availability.

The first step in setting up your administrative schedules is knowing which commands to run. Table 43 shows our recommendations for commands that should be scheduled and where to find more information on them. The table also includes our estimates on various other factors based on a typical implementation. Client schedules are not considered administrative, but they are included to emphasize that all of the other scheduling in your environment should hinge on your client schedules. This is a recommended set of schedules with all of the specifics required to execute in your environment if you follow the recommendations made in this redbook.

Table 43. Recommended schedules

#	Function	Task	Duration (Hrs)	Task Dependency	Refer to:
1	Backup Client Data to Tivoli Storage Manager Server	Schedule Client Backup Windows	6.0	Site Requirements	11.3 on page 224
2	Copy Backup Data to Offsite Storage Media	Backup Storage Pool DISKDIRS	0.5	Task 1 Complete	11.2.1 on page 218
3		Backup Storage Pool DISKDATA	1.0	Task 1 Complete	
4		Backup Storage Pool TAPEDATA	1.0	Task 1 Complete	
5		Backup Database	0.5	Task 2,3,4 Complete	
6	Volume History File Management	Delete Volume History Data for Database Backup Volumes	0.25	Task 5 Complete	11.2.2 on page 220
7	Prepare Disk Storage Pools for Next Backup Window	Migrate data from DISKDATA to TAPEDATA	3.0	Task 5 Complete	11.2.3 on page 221
8	Defragment Tape Volumes in Sequential Storage Pools	Reclaim volumes from OFFDIRS	1.0	None	11.2.4 on page 222
9		Reclaim volumes from OFFDATA	3.0	None	
10		Reclaim volumes from TAPEDATA	3.0	None	
11	Database Space Management	Expire Inventory	1.0	None	11.2.5

#	Function	Task	Duration (Hrs)	Task Dependency	Refer to:
12	Enable Statistics	Audit Licenses	0.25	None	11.2.6

11.2.1 Defining an offsite backup schedule

When the client backups are complete, all of the data that has changed in the primary storage pools should be copied to a copy storage pool. The copy storage pool tapes and the database backup tapes should be taken to another site, called the offsite location. These tapes are referred to collectively as the offsite backups, and will be used to restore your environment in the event of a catastrophic failure or disaster.

11.2.1.1 Flow of events

The order of execution of these tasks should be:

1. Check that there are no running sessions with the clients. When you back up the primary storage pools and database, you want to make sure that you are capturing an up to date copy. If there is still a client backing up its data, you should wait until it is done before starting the offsite backups.
2. Back up the primary storage pools to their copy storage pool. See 5.9.3, “Back up storage pool” on page 112 for more information.
 - a. Back up the DISKDIRS storage pool to the OFFDIRS storage pool.
 - b. Back up the DISKDATA storage pool to the OFFDATA storage pool.
 - c. Back up the TAPEDATA storage pool to the OFFDATA storage pool.
3. Back up the database. See 4.8, “Database backup” on page 72 for more information.

Since all of these commands are going to take a variable amount of time to execute, and they should be run in sequence, it makes sense to use a server script to run them. This will guarantee that each command will complete its tasks before another command is started.

11.2.1.2 Defining the server script

Figure 36 displays the commands that we would type into an administrative command line interface to accomplish the offsite backup tasks. We can use a server script to run these commands and check return codes to determine if we can proceed with the remaining commands.

```
select * from sessions where upper(session_type)='NODE'
backup stgpool diskdirs offdirs wait=yes
backup stgpool diskdata offdata wait=yes
backup stgpool tapedata offdata wait=yes
backup db devclass=coffsite type=full scratch=yes
```

Figure 36. Administrative commands to back up storage pools and database

Figure 37 shows the finished script for our recommended configuration. In 1.2, “Redbook support material” on page 2, we show how to load a predefined macro into Tivoli Storage Manager. The macro `mac.scripts`, which we provide to define a script to backup storage pools and the database to an offsite location, is shown in B.1.5, “Define server scripts” on page 290.

```

/* Script file for automation of the offsite storage tasks */

/* Make sure that any previous copies of the script are removed */
del script offsite_backups

/* Backup all primary storage pools followed by the database */
def script offsite_backups desc='Backup all primary storage pools &
database'
upd script offsite_backups /* -----*/
upd script offsite_backups /* Script Name: OFFSITE_BACKUPS */
upd script offsite_backups /* Description: Backup all primary storage */
upd script offsite_backups /* pools followed by the */
upd script offsite_backups /* database. If there are */
upd script offsite_backups /* active node sessions, don't*/
upd script offsite_backups /* run now, but reschedule to */
upd script offsite_backups /* check and run again in 20 */
upd script offsite_backups /* minutes. */
upd script offsite_backups /* The command will backup */
upd script offsite_backups /* three storage pools called */
upd script offsite_backups /* diskdirs, diskdata and */
upd script offsite_backups /* tapdata to backup copypools */
upd script offsite_backups /* called offdata and offdirs. */
upd script offsite_backups /* Example: run offsite_backups */
upd script offsite_backups /* -----*/
upd script offsite_backups ' select * from sessions where -'
upd script offsite_backups " upper(session_type)='NODE'"
upd script offsite_backups ' if (rc_ok) goto reschedule'
upd script offsite_backups ' ba stg diskdirs offdirs wait=yes'
upd script offsite_backups ' ba stg diskdata offdata wait=yes'
upd script offsite_backups ' ba stg tapedata offdata wait=yes'
upd script offsite_backups ' ba db dev=coffsite type=full scratch=yes'
upd script offsite_backups ' exit'
upd script offsite_backups ' reschedule:'
upd script offsite_backups ' del sch retry_offsite_backups type=a'
upd script offsite_backups ' def sch retry_offsite_backups t=a cmd="run
offsite_backups -'
upd script offsite_backups ' " active=yes starttime=NOW+0:20 peru=o'

```

Figure 37. Script file for offsite backups

The name of the script is completely arbitrary, but it should be representative of the function performed by the script. You should check that the script name you choose is not already in use. This can be accomplished by using the following command:

```

tsm: TSM010> select * from script_names where NAME='REDBOOK_OFFSITE'
ANR2034E SELECT: No match found using this criteria.
ANS8001I Return code 11.

```

Once the script is loaded into the server, you need to define a schedule that will run this script for you every day at 4 a.m. The macro `mac.schedules`, which we provide to define recommended schedules in our redbook environment, is shown in B.1.4, “Define schedules” on page 289 and includes the following schedule definition command:

```

tsm: TSM010> define schedule redbook_offsite type=administrative \
cont> cmd="run REDBOOK OFFSITE" starttime=04:00 \
cont> description="Backup primary storage pools & database"
ANR2577I Schedule REDBOOK_OFFSITE defined.

tsm: TSM010> query schedule redbook* type=administrative

```

*	Schedule Name	Start Date/Time	Duration	Period	Day
-	REDBOOK_OFFSITE	02/12/00 04:00:00	1 H	1 D	Any

11.2.2 Defining the volume history schedule

Every volume that is used by Tivoli Storage Manager, including the volumes used for the server database backups, is tracked within the server database. You can access this information while the server is up with the `query volhistory` command. The volume history information is very important because it tells you which volume holds your most recent database backup. In the event of loss or corruption of the server database, you need to know this in order to be able to restore your database automatically. However if the database is not available, you cannot retrieve the information from there. Therefore, a separate copy of the volume history information is periodically copied out to a text file that you can specify with the `VOLUMEHISTORY` option in the `dsmserv.opt` file. We recommend that you maintain at least two copies of the volume history file, in case one becomes unusable.

There are two recommended schedules that affect the volume history file: `DELETE_VOLHIST` and `BACKUP_VOLHIST`.

11.2.2.1 DELETE_VOLHIST

When you back up the database, the previous database backups become obsolete and should be returned to scratch volume status for reuse. The way to accomplish this is to delete the volume history record of this volume using the `delete volhistory` command. We recommend that you delete any database backups that are older than five days, since we back up the database every day. To define a delete volume history schedule, we define a schedule called `DELETE_VOLHIST` specifying a type of `TYPE=DBBACKUP`. The backslash (`\`) indicates to Tivoli Storage Manager to continue the command on the next line. For example:


```
tsm: TSM010> define schedule delete_volhist type=administrative \
cont> cmd="delete volhistory type=dibackup todate=today-5" starttime=07:00 \
cont> description="Delete volume history information for database backups"
```

```
ANR2577I Schedule DELETE_VOLHIST defined.
```

```
tsm: TSM010> query schedule delete* type=administrative
```

*	Schedule Name	Start Date/Time	Duration	Period	Day
-	DELETE_VOLHIST	02/12/00 07:00:00	1 H	1 D	Any

11.2.2.2 BACKUP_VOLHIST

Some operations, like the `delete volhistory` command, will update the volume history in the server database, but will not trigger the server to update the external volume history files. To ensure that the volume history data is always current, you should schedule a daily backup of the volume history database using the `backup volhistory` command. Use these commands to schedule a daily backup of the volume history from the database to the external volume history files:

```
tsm: TSM010> define schedule backup_volhist type=administrative \
cont> cmd="backup volhistory" starttime=07:05 \
cont> description="Backup volume history file"
ANR2577I Schedule BACKUP_VOLHIST defined.
```

```
tsm: TSM010> query schedule backup_volh* type=administrative
```

*	Schedule Name	Start Date/Time	Duration	Period	Day
-	BACKUP_VOLHIST	02/12/00 07:05:00	1 H	1 D	Any

11.2.3 Defining a migration schedule

Migration of your primary storage pools during client backup can slow down the client backup sessions and impact the scheduling of other server processes. We recommend that you avoid migration during client backup by scheduling primary storage pool migration outside of the client backup window. This will only prevent migration during client backup if the storage pools are correctly sized to hold the daily backup data from your clients.

To schedule a migration of a storage pool, you change the storage pool high and low migration thresholds to zero using the `update stgpool` command. This will force the storage pools to migrate immediately to their next level of storage pool. The migration process will continue until the storage pool is empty or until you reset the migration thresholds to a non-zero value. It is important not to leave the migration thresholds at zero after they are finished migrating. If more data is sent to these storage pools while they are set to zero, the data will be migrated immediately, defeating the purpose of having the storage pool.

Two schedules are required, one to start the migration and one to stop it. To define the migration schedules to start and stop storage pool migration, do the following:

```

tsm: TSM010> define schedule migration_start type=administrative \
cont> cmd="update stgpool diskdata hi=0 lo=0" starttime=07:00 \
cont> description="Start migration on DISKDATA storage pool"
ANR2577I Schedule MIGRATION_START defined.

tsm: TSM010> define schedule migration_stop type=administrative \
cont> cmd="update stgpool diskdata hi=70 lo=30" starttime=10:00 \
cont> desc="Stop migration on DISKDATA storage pool"
ANR2577I Schedule MIGRATION_STOP defined.

tsm: TSM010>query schedule migration* type=administrative

```

*	Schedule Name	Start Date/Time	Duration	Period	Day
-	MIGRATION_START	02/12/00 07:00:00	1 H	1 D	Any
	MIGRATION_STOP	02/12/00 10:00:00	1 H	1 D	Any

11.2.4 Defining a reclamation schedule

Tapes in the sequential access storage pools will eventually get fragmented due to file expiration and node removal. Over time, this can cause many tapes to have very little valid data on them. A reclamation process will take all of this data and consolidate it onto a smaller number of tapes thus returning many tapes to scratch status for reuse. We recommend controlling when the reclamation process runs by using schedules which set and reset the reclamation threshold for the storage pools. Outside of the scheduled reclamation periods, the RECLAIM parameter for the storage pool should remain at 100, disabling reclamation from running. Controlling reclamation is desirable because it demands heavy use of the tape drives and might interfere with normal backup or restore operations if allowed to run ad hoc.

To start a reclamation process on a sequential storage pool, use the `update stgpool` command to set the RECLAIM parameter for the storage pool to a value between 50 and 99. This value represents the percentage of reclaimable space on the volume which is the opposite of the volume utilization. That is to say, if the output from a `query volume` command shows a volume as 20% utilized, then its reclaimable space is 80%. Reclamation stops when all eligible volumes have been reclaimed or the reclaim parameter is set to 100.

Each sequential storage pool should have a schedule to start and stop reclamation. We have three sequential storage pools in our recommended configuration. So we need six schedules to start and stop the reclamations. To set up the reclamation schedules, do the following:

```

tsm: TSM010> define schedule RECLAIM_OFFDIRS_START type=administrative \
cont> cmd="update stgpool offdirs reclaim=75" starttime=10:00 \
cont> description="Start reclaim on the OFFDIRS storage pool"
ANR2577I Schedule RECLAIM_OFFDIRS_START defined.

tsm: TSM010> define schedule RECLAIM_OFFDIRS_STOP type=administrative \
cont> cmd="update stgpool offdirs reclaim=100" starttime=11:00 \
cont> description="Stop reclaim on the OFFDIRS storage pool"
ANR2577I Schedule RECLAIM_OFFDIRS_STOP defined.

tsm: TSM010> define schedule RECLAIM_OFFDATA_START type=administrative \
cont> cmd="update stgpool offdata reclaim=75" starttime=11:00 \
cont> description="Start reclaim on the OFFDATA storage pool"
ANR2577I Schedule RECLAIM_OFFDATA_START defined.

tsm: TSM010> define schedule RECLAIM_OFFDATA_STOP type=administrative \
cont> cmd="update stgpool offdata reclaim=100" starttime=14:00 \
cont> description="Stop reclaim on the OFFDATA storage pool"
ANR2577I Schedule RECLAIM_OFFDATA_STOP defined.

tsm: TSM010> define schedule RECLAIM_TAPEDATA_START type=administrative \
cont> cmd="update stgpool tapedata reclaim=75" starttime=14:00 \
cont> description="Start reclaim on the TAPEDATA storage pool"
ANR2577I Schedule RECLAIM_TAPEDATA_START defined.

tsm: TSM010> define schedule RECLAIM_TAPEDATA_STOP type=administrative \
cont> cmd="update stgpool tapedata reclaim=100" starttime=17:00 \
cont> description="Stop reclaim on the TAPEDATA storage pool"
ANR2577I Schedule RECLAIM_TAPEDATA_STOP defined.

tsm: TSM010> query schedule reclaim* type=administrative

```

*	Schedule Name	Start Date/Time	Duration	Period	Day
-	RECLAIM_OFFDATA- START	02/12/00 11:00:00	1 H	1 D	Any
	RECLAIM_OFFDATA- STOP	02/12/00 14:00:00	1 H	1 D	Any
	RECLAIM_OFFDIRS- START	02/12/00 10:00:00	1 H	1 D	Any
	RECLAIM_OFFDIRS- STOP	02/12/00 11:00:00	1 H	1 D	Any
	RECLAIM_TAPEDAT- A START	02/12/00 14:00:00	1 H	1 D	Any
	RECLAIM_TAPEDAT- A STOP	02/12/00 17:00:00	1 H	1 D	Any

11.2.5 Defining an expiration schedule

Managing the amount of space used by your database will be very important as your environment grows. The client backup process handles half of the task by marking expired data references while backing up client data. The second half of the task is handled by the `expire inventory` command. This command will remove the marked entries from the database.

The Tivoli Storage Manager server configuration file defaults to a 24 hour cycle for expiration. This would result in daily expiration processing beginning 24 hours from when you start the server. This command can be computationally intensive and it is best scheduled during a quiet server time to minimize its impact on server operations. Therefore, we recommend scheduling the `expire inventory` command daily. This assumes you have disabled the default expiration execution as recommended in 6.5, “Enforcing your policy” on page 141.

```

tsm: TSM010> define schedule expire_inventory type=administrative \
cont> cmd="expire inventory" startt=17:00 \
cont> description="Inventory expiration"
ANR2577I Schedule EXPIRE_INVENTORY defined.

```

```

tsm: TSM010> query schedule expire* type=administrative

```

*	Schedule Name	Start Date/Time	Duration	Period	Day
-	EXPIRE_INVENTORY	02/12/00 17:00:00	1 H	1 D	Any

11.2.6 Defining a licensing audit schedule

It is helpful to be able to check statistics for the Tivoli Storage Manager server usage by using the `query auditoccupancy` command. This command returns some valuable information on client data usage in the Tivoli Storage Manager environment.

```

tsm: TSM010> query auditoccupancy
License information as of last audit on 02/12/00 at 14:21:03.

```

Node Name	Backup Storage Used (MB)	Archive Storage Used (MB)	Space-Managed Storage Used (MB)	Total Storage Used (MB)
DEVNODE1	5	11	0	16
PAGOPAGO	0	0	0	0
FIJI	0	0	0	0
PUTNEY	67	0	0	67
CLADAN	0	0	0	0

The data returned from this command is updated when the `audit licenses` command is run on the Tivoli Storage Manager server. We recommend running this command daily using the administrative schedule below:

```

tsm: TSM010> define schedule audit_license type=administrative \
cont> cmd="audit licenses" starttTime=00:00 \
cont> description="Audit licenses"
ANR2577I Schedule AUDIT_LICENSE defined.

```

```

tsm: TSM010> query schedule audit* type=administrative

```

*	Schedule Name	Start Date/Time	Duration	Period	Day
-	AUDIT_LICENSE	02/12/00 00:00:00	1 H	1 D	Any

11.3 Client schedules

A client schedule is a directive to trigger an action on a group of Tivoli Storage Manager client machines. It is different from an administrative schedule in that it specifies that an action be performed on the Tivoli Storage Manager client. The client scheduling system consists of a server portion and a client portion. The server part is integrated into the Tivoli Storage Manager process and is responsible for defining the schedule parameters and associating nodes with the schedule. The client scheduler is a separate process on the Tivoli Storage Manager client and provides communication between the server and client. See 10.4, "Client scheduler" on page 204 for further information on the client portion of the client scheduling system.

A Tivoli Storage Manager client machine with an active scheduler component can be scheduled to perform any of the following actions:

- Backup (incremental/selective)
- Archive
- Restore
- Retrieve
- Operating system command
- Macro
- Image backup/restore (UNIX)

The most common use of client schedules is to implement an automatic incremental backup process for a group of machines.

At this point in the configuration of the Tivoli Storage Manager system, you will have installed the Tivoli Storage Manager client code on your clients and defined them to a policy domain. You will also have activated the scheduler service on the client so that it is ready to perform the actions that we schedule for it. The remaining activities are concentrated on the server and its scheduling facility. To schedule an action on the client you will need to do the following:

1. Define a client schedule
2. Associate a client with a schedule
3. Verify the client schedules

11.3.1 Defining a client schedule

In 1.2, “Redbook support material” on page 2, we show how to load a predefined macro into Tivoli Storage Manager. The macro `mac.schedules`, which we provide to define schedules for our redbook environment, is shown in B.1.4, “Define schedules” on page 289.

The `define schedule` command defines a schedule to the Tivoli Storage Manager server. It does not contain any reference to a client; that operation will be handled by the `define association` command. To define a client schedule, you will need the following information:

- Policy domain name: The schedule will be created within this policy domain and can only be associated with nodes within the same domain. This will allow you to manage which administrators have control of these schedules by limiting access to the domain that owns them.
- Schedule name: This can be any name you choose, up to 30 characters.
- Start time: The time of day that you want the schedule to trigger your action.
- Duration: The scheduler can only initiate the schedule within the time period specified with the duration parameters.

The following command illustrates the syntax for a client schedule. This command defines a schedule named `BACKUP_NIGHTLY` in the `ABCCORP` policy domain. This schedule will trigger an incremental backup (the default action for a client schedule) at 10 p.m. every night. If there is a problem contacting the client, the scheduler will keep trying until 1 a.m.

```

tsm: TSM010> define schedule server backup_nightly starttime=22:00 duration=3 \
cont> durunits=hours description="Nightly backup schedule for SERVER domain"
ANR2500I Schedule BACKUP_NIGHTLY defined in policy domain SERVER.

```

Client schedules are very dependent on site requirements, but we have some recommendations that are applicable for a wide range of environments.

Our recommended policy configuration includes two policy domains, so we must define a client backup schedule for each. The parameters for each schedule are very similar except for the name which is chosen to reflect the domain. The schedules start at 10 p.m. and will continue attempting to start incremental backups until 1 a.m. The actual commands and the output look like this:

```

tsm: TSM010> define schedule server server_nightly starttime=22:00 duration=3 \
cont> durunits=hours description="Nightly backup schedule for SERVER domain"
ANR2500I Schedule SERVER_NIGHTLY defined in policy domain SERVER.

tsm: TSM010> define schedule workstn workstn_nightly starttime=22:00 duration=3 \
cont> durunits=hours description="Nightly backup schedule for WORKSTN domain"
ANR2500I Schedule WORKSTN_NIGHTLY defined in policy domain WORKSTN.

tsm: TSM010> query schedule

```

Domain	* Schedule Name	Action	Start Date/Time	Duration	Period	Day
SERVER	SERVER_NIGHTLY	Inc Bk	02/12/00 22:00:00	3 H	1 D	Any
WORKSTN	WORKSTN_NIGHTLY	Inc Bk	02/12/00 22:00:00	3 H	1 D	Any

11.3.2 Associating a client with a schedule

Once the schedule is defined, we need to specify which nodes belong to it using the `define association` command. We must associate nodes to a schedule within their own domain. The nodes must be registered to the server before running this command, but they do not necessarily have to be in contact with the server. The actual commands and the output look like this:

```

tsm: TSM010> define association server server_nightly putney
ANR2510I Node PUTNEY associated with schedule SERVER_NIGHTLY in policy
domain SERVER.

tsm: TSM010> define association workstn workstn_nightly cocos
ANR2510I Node COCOS associated with schedule WORKSTN_NIGHTLY in policy domain
WORKSTN.

tsm: TSM010> query association

```

```

Policy Domain Name: SERVER
Schedule Name: SERVER_NIGHTLY
Associated Nodes: PUTNEY

Policy Domain Name: WORKSTN
Schedule Name: WORKSTN_NIGHTLY
Associated Nodes: COCOS

```

11.3.3 Verifying the client schedules

The schedules can be checked out by using the `query event` command. This will show you if the nodes and schedules are all set up correctly. The actual command and its output look like this:

```
tsm: TSM010> query event * *
```

Scheduled Start	Actual Start	Schedule Name	Node Name	Status
02/12/00	22:00:00	SERVER_NIGHT- LY	PUINEY	Future
02/12/00	22:00:00	WORKSTIN_NIGH- TLY	COCOS	Future

Chapter 12. Operations

Tivoli Storage Manager, like any other software system, has to be maintained and operated. This is to check that the tasks which are implemented are still working. To accomplish this, you may need to consider how many people will be dedicated to Tivoli Storage Manager. Depending on the size of your installation, this can range from a single person to a complete team. The operations team is responsible for monitoring, correcting or working together with technical support to isolate any non-specific error condition. Basically, you need the following logical roles among your current or future Tivoli Storage Manager operations team. Depending on the actual number of people in the team, some or all of these roles may be combined into the same actual person:

- **Operator:** Handles Tivoli Storage Manager daily tasks and special server procedures.
- **Technical Support:** Installs new client code, upgrades server code, activates trace functions, deals with hardware configurations, answers end-user questions and handles error conditions.
- **Storage Administrator:** Performs sanity checks in the environment, evaluates capacity planning issues, and manages storage pools.
- **System Administrator:** Manages Tivoli Storage Manager policies, defines new nodes and management classes, and reports error conditions to technical support.

Tivoli Storage Manager provides a centralized point for backup, restore, archive, and retrieve operations, as well as backupsets and logical volume operations. All of these functions require monitoring and special operational procedures. As a rule of thumb, consider having at least two people responsible for Tivoli Storage Manager, so that each one can handle the workload in case of a problem (this way you can also avoid concentrating Tivoli Storage Manager information in only one person). All the personnel directly involved in Tivoli Storage Manager administration—the Tivoli Storage Manager operations team—need to divide the Tivoli Storage Manager workload into some basic activities:

- **Server procedures:** Start and stop the Tivoli Storage Manager server.
- **Daily operations:** Check database free space, check activity log, check volumes in and out, send tapes to offsite, request to bring tapes onsite, check Tivoli Storage Manager devices, back up Tivoli Storage Manager database (or monitor its execution), and label new tapes.
- **Error or critical conditions:** Bring an offsite volume for restore, create another database or recovery log, deal with storage pool shortages, mirror write failures, communication errors, device problems, media errors, offsite tape not found, full or partial disaster recovery and any other non-specific problems.

12.1 Server procedures

In this section we show how to start and stop a Tivoli Storage Manager server. To a great extent, Tivoli Storage Manager is a software program that runs as any other application in your system. This means that you can start the server and let it run continuously and uninterrupted for days, weeks, or even longer if desired.

12.1.1 Starting the Tivoli Storage Manager server

The Tivoli Storage Manager server can run as either a foreground or a background process. The foreground method allows you to have an active open console, with all system messages displayed. From this server console, you are able to issue all administrative commands (depending on the authority which is available for the SERVER_CONSOLE administrator ID — see 7.1.2, “Default environment” on page 143). The background method allows you to close the active window without stopping the server process. This is the recommended way if you are planning to have the Tivoli Storage Manager server automatically started or stopped.

This section explains how to start a Tivoli Storage Manager server (we assume that it is not running yet). Otherwise, you must stop it first and issue the commands we explain. You can have the Tivoli Storage Manager server automatically started, in which case you need to configure it as a background process. If you wish to have it started manually, then the Tivoli Storage Manager server needs to be configured as a foreground process. For further details on how to start the Tivoli Storage Manager server, refer to the Tivoli Storage Manager Quick Start manual which comes with the product software.

12.1.1.1 Active console mode

To start a Tivoli Storage Manager server in foreground, you can start the Tivoli Storage Manager server by simply executing the DSMSESV program in the Tivoli Storage Manager system directory. If you run this way, you will not be able to close the window without halting the server. All output messages will be redirected to standard output.

The next screen shows you how to manually start the Tivoli Storage Manager server on AIX in active console mode. This server console session is established using a predefined administrator ID SERVER_CONSOLE. All Tivoli Storage Manager system messages are displayed on the console. The Tivoli Storage Manager server is ready to receive and send data to the client machines.

```

root@ /usr/tivoli/tsm/server/bin > dmserv
ANR7800I DSMSESV generated at 12:24:44 on Feb  4 2000.

Tivoli Storage Manager for AIX-RS/6000
Version 3, Release 7, Level 2.0

Licensed Materials - Property of IBM

5697-TSM (C) Copyright IBM Corporation 1999. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

ANR0900I Processing options file dmserv.opt.
ANR0990I Server restart-recovery in progress.
ANR0200I Recovery log assigned capacity is 48 megabytes.
ANR0201I Database assigned capacity is 52 megabytes.
ANR0306I Recovery log volume mount in progress.
ANR0353I Recovery log analysis pass in progress.
ANR0354I Recovery log redo pass in progress.
ANR0355I Recovery log undo pass in progress.
ANR0352I Transaction recovery complete.
ANR2283I Database backup trigger defined, but is disabled.
ANR2100I Activity log process has started.
ANR1305I Disk volume /tsm/stgpool/dirs01 varied online.
ANR1305I Disk volume /tsm/stgpool/file01 varied online.
ANR2803I License manager started.
ANR2560I Schedule manager started.
ANR8200I TCP/IP driver ready for connection with clients on port 1500.
ANR8190I HTTP driver ready for connection with clients on port 1580.
ANR0993I Server initialization complete.
ANR0984I Process 1 for EXPIRATION started in the BACKGROUND at 09:14:35.
ANR0811I Inventory client file expiration started as process 1.
ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli is now ready for use.
ANR0812I Inventory file expiration process 1 completed: examined 10 objects,
deleting 0 backup objects, 0 archive objects, 0 DB backup volumes, and 0
recovery plan files. 0 errors were encountered.
ANR0985I Process 1 for EXPIRATION running in the BACKGROUND completed with
completion state SUCCESS at 09:14:35.
TSM:TSM010>
ANR2835I Server is licensed for 102 clients.
ANR2861I Server is licensed to support NETWORK connections.
ANR2856I Server is licensed to support Open Systems Environment clients.
ANR2855I Server is licensed to support Tivoli Space Manager clients.
ANR2860I Server is licensed to support Tivoli Disaster Recovery Manager.
ANR2880I Server is licensed to support Tivoli Data Protection for Oracle clients.
ANR2881I Server is licensed to support Tivoli Data Protection for MS SQL Server clients.
ANR2882I Server is licensed to support Tivoli Data Protection for MS Exchange clients.
ANR2883I Server is licensed to support Tivoli Data Protection for Lotus Notes clients.
ANR2884I Server is licensed to support Tivoli Data Protection for Lotus Domino clients.
ANR2885I Server is licensed to support Tivoli Data Protection for Infomix clients.
ANR2827I Server is licensed to support Tivoli Data Protection for SAP R/3 clients.
ANR2869I Server is licensed for Advanced Device Support.
ANR8439I SCSI library L3570 is ready for operations.

Storage Management Server for AIX-RS/6000 - Version 3, Release 7, Level 2.0

TSM:TSM010>

```

When you do this, you should check the startup messages — either as they appear on the screen or by querying the server activity log. You will soon learn what messages indicate an error-free startup, and what messages are cause for concern. For example, if one of the mirrored database or log volumes is not able to be brought online, usually because of a disk problem, it will trigger a startup error message. If you note this when the server starts you can start investigating and fixing the problem straight away.

On Windows NT, you also can start the server in active console mode from the Tivoli Storage Manager server utilities program. In this method, you will not be able to close your desktop, since Tivoli Storage Manager is running in foreground

mode. Figure 38 shows the main Tivoli Storage Manager server utilities screen and the Launch option to start Tivoli Storage Manager in console mode.

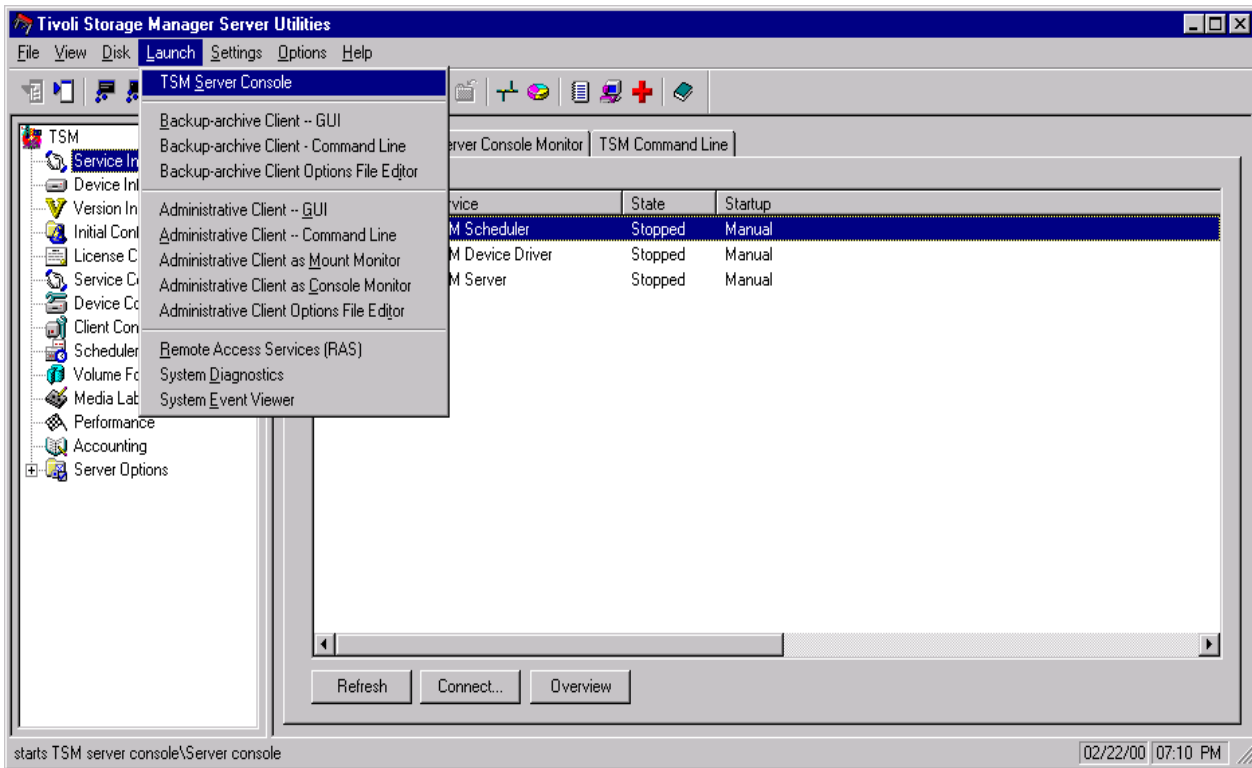


Figure 38. Launch server console in foreground

By using the launch option, the Tivoli Storage Manager console stays open as long as your desktop (login session) stays active. This is the recommended way if you plan to have a Tivoli Storage Manager operation staff near the machine, or if you want to have all Tivoli Storage Manager system activity displayed in the server, so that you can quickly check what is running.

12.1.1.2 Background mode

You can start the Tivoli Storage Manager server in background mode either manually or with each system start.

Manual start

On AIX, to manually start the Tivoli Storage Manager server in background mode, you have to run the script `rc.admserv`, which is a part of the Tivoli Storage Manager server software.

```
root@ /usr/tivoli/tsm/server/bin >nohup rc.admserv >/dev/console 2>&1 &
```

On Windows NT, the Tivoli Storage Manager server is installed as a manual service. To start it, you have to go to **Settings** —> **Control Panel** —> **Services**, select the Tivoli Storage Manager server service, and start it.

Automatic start

On AIX, after the Tivoli Storage Manager server installation, the system is configured to start the Tivoli Storage Manager server automatically during the

system startup. The system initialization file /etc/inittab is edited. The following entry is added to the file:

```
autosrvr:2:once:/usr/tivoli/tsm/server/bin/rc.admserv >/dev/console 2>&1
#Start the Tivoli Storage Manager server
```

Solaris and HP-UX also use the /etc/inittab for automatic startup of processes, however you may prefer to put the commands into another of the possible startup files. We recommend you refer to your operating system documentation and discuss with the system administrator to decide on the best configuration.

On Windows NT, the Tivoli Storage Manager server is installed as a manual service. To change its status to automatic, you have to go to **Settings** → **Control Panel** → **Services**, select the Tivoli Storage Manager server service and change its startup setting to automatic.

The Tivoli Storage Manager server will be started automatically the next time the system is rebooted. You can then check that the server is running by trying to connect to the server from an administrative client or a backup-archive client, or by checking whether the Tivoli Storage Manager server process (UNIX) or service (Windows NT) is running.

Administrative console session

If you start the Tivoli Storage Manager server in the background, we highly recommend running an additional administrative client session in console mode on the administrative workstation. This session displays all Tivoli Storage Manager messages, however, it is not an active session, that is, you cannot issue administrative commands here. This way you can easily monitor your Tivoli Storage Manager system. The following command shows how to start an administrative session in console mode.

```
root@ / > dsmadm -console -id=sysadmin -password=island
Tivoli Storage Manager
Command Line Administrative Interface - Version 3, Release 7, Level 2.0
(C) Copyright IBM Corporation, 1990, 1999, All Rights Reserved.

Session established with server TSM010: AIX-RS/6000
Server Version 3, Release 7, Level 2.0
Server date/time: 02/14/00 09:24:06 Last access: 02/12/00 15:56:52
```

Note:

We recommend that you create a specific analyst administrator for that purpose, so that it can only query the server.

12.1.2 Stopping the Tivoli Storage Manager server

There are two ways of stopping a Tivoli Storage Manager server:

- Immediate stop: This brings the server down immediately. All system processes and client activity are interrupted.
- Drain system activity and stop: This is the recommended way, because you can understand what processes were in progress and what you may need to restart later on. Performing basic pre-checks and then stopping gives you (and the operations team) an idea of the workload that may need to be restarted. It

is also a more orderly and therefore cleaner shutdown than the immediate stop.

12.1.2.1 Stopping the Tivoli Storage Manager server immediately

You can halt the server without warning if an unplanned operating system problem requires the server to be stopped.

On Windows NT, you can terminate a server session by stopping the service. To change the status of the server service, you have to go to **Settings** → **Control Panel** → **Services**, select the Tivoli Storage Manager server service and stop it.

When you halt the server, all processes are abruptly stopped and client sessions are canceled, even if they are not completed. Any in-progress transactions are rolled back when the server is restarted. When the server is halted, neither administrator activity nor client operations are possible. We recommend that you halt the server only after current administrative and client node sessions have completed or canceled.

You must use the `halt` command to shutdown the Tivoli Storage Manager server:

```
TSM:TSM010> halt
ANR2017I Administrator SERVER_CONSOLE issued command: HALT
ANR7835I Server thread 1 terminated in response to server shutdown.
ANR7835I Server thread 14 terminated in response to server shutdown.
....
ANR7835I Server thread 41 terminated in response to server shutdown.
ANR7835I Server thread 43 terminated in response to server shutdown.
ANR7835I Server thread 44 terminated in response to server shutdown.
ANR0991I Server shutdown complete.
```

Note:

The QUIESCE option with the `halt` command is recommended only if you plan to do a database dump by using the `DSMSERV DUMPDB` command immediately after halting. Because Tivoli Storage Manager supports online database backup (`backup db` command), the `DSMSERV DUMPDB` command should rarely, if ever, be needed.

12.1.2.2 Draining system activity before stopping the server

To shut down the server without severely impacting administrative and client node activity with the server, we recommend that you take some steps before actually issuing the `halt` command to the server:

1. Disable the server: This prevents new client node sessions from starting. Disabling new client sessions prevents users from establishing client node sessions with the server. This command does not affect current client sessions in progress or system processes like migration and reclamation. To disable client node access to the server, use the `disable sessions` command:

```
tsm: TSM010> disable sessions
ANR2097I Server now disabled for backup/archive client access.
```

When you disable client sessions from the server, administrators can still access it, and current client node activity completes unless the user logs off or you cancel the client node session.

You can issue the `query status` command to determine if the server is enabled or disabled.

2. Query for session information: This is to identify any existing administrative and client node sessions.

When an administrator or client accesses Tivoli Storage Manager, either an administrative or client node session is established with the server. Each session is assigned a unique session number. Note that due to the multi-threading features of Tivoli Storage Manager client, node PUTNEY has many active sessions with the server. To request information about client sessions, enter the `query session` command:

```
tsm: TSM010> query session
```

Sess Number	Comm. Method	Sess State	Wait Time	Bytes Sent	Bytes Recvd	Sess Type	Platform	Client Name
1	Tcp/Ip	Run	0 S	617	168	Admin	AIX	SYSADMIN
2	Tcp/Ip	IdleW	2.6 M	1.2 K	268	Node	WinNT	DEVNODE1
5	Tcp/Ip	IdleW	5 S	502.1 K	551	Node	AIX	PUTNEY
6	Tcp/Ip	RecvW	0 S	428	63.2 M	Node	AIX	PUTNEY
7	Tcp/Ip	RecvW	0 S	406	40.1 M	Node	AIX	PUTNEY
9	Tcp/Ip	RecvW	0 S	400	19.4 M	Node	AIX	PUTNEY
11	Tcp/Ip	RecvW	0 S	391	15.7 M	Node	AIX	PUTNEY
13	Tcp/Ip	RecvW	0 S	375	3.6 M	Node	AIX	PUTNEY
14	Tcp/Ip	Run	0 S	346	8.0 M	Node	AIX	PUTNEY

3. Notify users: This is to notify any existing administrative and client node sessions that you plan to shut down the server. Tivoli Storage Manager does not provide a network notification facility; you must use external means to notify users.
4. Cancel any existing administrative or client node sessions: To cancel a session, you must identify it by its session number. You can identify the session number by issuing the `query session` command. For example, in the screen above, one session number is 11 (client PUTNEY). You can cancel that session by entering the `cancel session` command:

```
tsm: TSM010> cancel session 11
ANR0490I Canceling session 11 for node PUTNEY (AIX) .
```

Note that one client may still have other active sessions running. If you want to cancel just those client sessions, you will have to cancel them one by one.

On the other hand, if you want to cancel all backup-archive client sessions, enter the `cancel session all` command:

```

TSM:TSM010> cancel session all
ANR0490I Canceling session 2 for node DEVNODE1 (WinNT) .
ANR0490I Canceling session 5 for node PUTNEY (AIX) .
ANR0490I Canceling session 6 for node PUTNEY (AIX) .
ANR0490I Canceling session 7 for node PUTNEY (AIX) .
ANR0490I Canceling session 9 for node PUTNEY (AIX) .
ANR0490I Canceling session 11 for node PUTNEY (AIX) .
ANR0490I Canceling session 13 for node PUTNEY (AIX) .
ANR0490I Canceling session 14 for node PUTNEY (AIX) .

ANR0483W Session 2 for node DEVNODE1 (WinNT) terminated - forced by administrator.
ANR0483W Session 5 for node PUTNEY (AIX) terminated - forced by administrator.
ANR0483W Session 6 for node PUTNEY (AIX) terminated - forced by administrator.
ANR0483W Session 7 for node PUTNEY (AIX) terminated - forced by administrator.
ANR0483W Session 9 for node PUTNEY (AIX) terminated - forced by administrator.
ANR0483W Session 11 for node PUTNEY (AIX) terminated - forced by administrator.
ANR0483W Session 13 for node PUTNEY (AIX) terminated - forced by administrator.
ANR0483W Session 14 for node PUTNEY (AIX) terminated - forced by administrator.

```

If an operation, such as a backup or an archive process, is interrupted when you cancel the session, Tivoli Storage Manager rolls back the results of the current transaction. That is, any changes that are not yet committed to the database are undone. If necessary, the cancellation process may be delayed. For example, while a client restore session is in a restartable state, the file space is locked and no files can be moved from sequential volumes. This prevents the data from being migrated, moved, reclaimed, or backed up by another operation. These sessions will automatically expire when the specified restore interval has passed. For further details, see the `query restore` and `cancel restore` commands.

5. Find out if any other processes are running, such as server migration or inventory expiration, by using the `query process` command. If a database backup process is running, allow it to complete before halting the server. If other types of processes are running, cancel them by using the `cancel process` command.
6. Halt the server and all server operations by using the `halt` command as explained in 12.1.2.1, “Stopping the Tivoli Storage Manager server immediately” on page 234.

12.2 Event monitoring

One of the features that is integrated to Tivoli Storage Manager is a complete set of messages that can be enabled, disabled, stored, or even forwarded to other Tivoli Storage Manager servers, user programs, systems management consoles like T/EC from Tivoli, or to SNMP management applications. These components are known as EVENT RECEIVERS. You can enable the event receivers which are available and appropriate to your environment. Since event logging has a full set of functions and features that are beyond of the scope of this book, we only cover some basic settings.

Tivoli Storage Manager divides the range and the source of the messages by prefixes. For example, prefix ANR is for server messages, and most of prefix ANS messages are from backup/archive client sessions. Details of all these messages are explained in *Tivoli Storage Manager Messages*, GC35-0382.

You can enable or disable messages for any particular receiver with the exception that you cannot disable server events from being sent to the server activity log.

Logging of events to both the Tivoli Storage Manager server console and the Tivoli Storage Manager activity log begins automatically at server startup. Messages can appear on the server console, the administrative client, an operator terminal, the administrative graphical user interface, the backup/archive client, or the space-management client.

Tivoli Storage Manager provides an activity log to help the administrator track server activity and monitor the system. The activity log contains messages generated by the server and is stored in the database. Any messages sent to the server console are stored in the activity log. Examples of the types of messages stored in the activity log include:

- When client sessions start or end
- When migration starts or ends
- When backed up files are expired from data storage
- Any output generated from background processes

You can check if event logging is enabled for a receiver by using the `query status` command and checking the Active Receivers field:

```

tsm: TSM010> query status
Storage Management Server for AIX-RS/6000 - Version 3, Release 7, Level 2.0

Server Name: TSM010
Server host name or IP address:
Server TCP/IP port number:
Server URL:
Crossdefine: Off
Server Password Set: No
Server Installation Date/Time: 01/24/00 16:44:02
Server Restart Date/Time: 02/14/00 09:32:27
Authentication: On
Password Expiration Period: 90 Day(s)
Invalid Sign-on Attempt Limit: 3
Minimum Password Length: 6
WEB Admin Authentication Time-out (minutes): 0
Registration: Closed
Availability: Enabled
Accounting: On
Activity Log Retention Period: 10 Day(s)
Activity Summary Retention Period: 30 Day(s)
License Audit Period: 30 Day(s)
Last License Audit: 02/12/00 14:21:03
Server License Compliance: Valid
Central Scheduler: Active
Maximum Sessions: 25
Maximum Scheduled Sessions: 12
Event Record Retention Period: 10 Day(s)
Client Action Duration: 5 Day(s)
Schedule Randomization Percentage: 25
Query Schedule Period: Client's Choice
Maximum Command Retries: Client's Choice
Retry Period: Client's Choice
Scheduling Modes: Any
Log Mode: Normal
Database Backup Trigger: Disabled
BufPoolSize: 2,048 K
Active Receivers: CONSOLE ACTLOG
Configuration manager?: Off
Refresh interval: 60
Last refresh date/time:

```

We recommend that you run the `begin eventlogging` command to make sure that all defined receivers are active. Use this command to begin logging events to one or more receivers. Event logging automatically begins when the server is started for the console and activity log, and for any receivers that are started automatically based on entries in the server options file. A receiver for which event logging has begun is an active receiver. The next screen shows you how to activate all defined receivers:

```

tsm: TSM010> begin eventlogging

ANR1825I Event logging active for the CONSOLE receiver.
ANR1825I Event logging active for the ACTLOG receiver.

```

We recommend that you enable nodename logging, so that you can also monitor client activity centrally in the server. Note that enabling client events to the activity log will increase the Tivoli Storage Manager database utilization. You can set a retention period for the log records by using the `set actlogretention` command. To enable a specific receiver (CONSOLE, ACTLOG, EVENTSERVER, FILE, FILETEXT, SNMP, TIVOLI, USEREXIT) enter the `enable events` command:

```
tsm: TSM010> enable events actlog all nodename=*  
ANR1844I ENABLE EVENTS command processed.
```

Windows NT Users:

On Windows NT, you also have a NTEVENTLOG receiver. This enables you to send events to the application eventlog, so that you can use the Windows Event Viewer to monitor a Tivoli Storage Manager system. In this case, we recommend that you issue the `enable events` command for the error and severe category.

12.3 Daily sanity checks

This section shows you how to perform some basic, yet important, system validations on your Tivoli Storage Manager environment. These checks assume that you do not have any kind of automation in place, so that you (or the Tivoli Storage Manager operations team) may have to control some of the server and client aspects to have it running properly.

All commands are issued in the administrative command line interface. This makes the monitoring process easier and can be further enhanced to accommodate any automation facility you may desire. We recommend that you start using all of the commands on a daily basis, then once you feel confident of your environment settings, you can use any or all of them on a weekly basis. You should balance the level of information that you believe is necessary as you start running the commands.

You should run monitoring commands as part of the daily operations check, as a sync-point for all Tivoli Storage Manager personnel, especially if your company has different working hours and different working teams.

Tivoli Storage Manager has a rich set of commands that allow you to monitor and create your own reports. The `query` command set is the base for most information from the server. Typing `help query` or `help query <command>` provides you with further details on all possible combinations.

For more sophisticated reporting requirements, there is a module available for the Tivoli Decision Support product called Tivoli Decision Support for Storage Management Analysis. This provides complex text and graphical reporting and analysis for one or more Tivoli Storage Manager servers. This product is outside the scope of this book, but more information can be found in *Tivoli Storage Manager Reporting*, SG24-6109.

12.3.1 Database and recovery log

The following is a list of important Tivoli Storage Manager queries for the database and recovery log in alphabetical order.

12.3.1.1 Display database information

To display information about your current database occupancy and how much space is still available for extension, use the `query db` command:

```

tsm: TSM010> query db

Available Assigned Maximum Maximum Page Total Used Pct Max.
Space Capacity Extension Reduction Size Usable Pages Util Pct
(MB) (MB) (MB) (MB) (bytes) Pages
-----
52 52 0 44 4,096 13,312 949 7.1 7.1

tsm: TSM010> query db format=detail

Available Space (MB): 52
Assigned Capacity (MB): 52
Maximum Extension (MB): 0
Maximum Reduction (MB): 44
Page Size (bytes): 4,096
Total Usable Pages: 13,312
Used Pages: 949
Pct Util: 7.1
Max. Pct Util: 7.1
Physical Volumes: 10
Buffer Pool Pages: 512
Total Buffer Requests: 74,591
Cache Hit Pct.: 99.59
Cache Wait Pct.: 0.00
Backup in Progress?: No
Type of Backup In Progress:
Incrementals Since Last Full: 0
Changed Since Last Backup (MB): 3.71
Percentage Changed: 100.11
Last Complete Backup Date/Time: 02/10/00 13:30:55

```

12.3.1.2 Display database volumes information

Use the `query dbvolume` command to display information on one or more database volumes, including available, allocated, and free space on the volume. This command displays information about the specified database volume and any database volume copies.

```

tsm: TSM010> query dbvolume

Volume Name      Copy   Volume Name      Copy   Volume Name      Copy
(Copy 1)         Status (Copy 2)         Status (Copy 3)         Status
-----
/tsm/database/p- Sync'd /tsm/database/c- Sync'd      Undef-
rimary/file01   opy/file01      ined

```

12.3.1.3 Display information on the recovery log

Use the `query log` command to display allocation information about the recovery log, including utilization, expansion, and reduction abilities.

```

tsm: TSM010> query log

Available Assigned Maximum Maximum Page Total Used Pct Max.
Space Capacity Extension Reduction Size Usable Pages Util Pct
(MB) (MB) (MB) (MB) (bytes) Pages Pages Util Pct
-----
48 48 0 44 4,096 11,776 89 0.8 5.0

tsm: TSM010> query log format=detail

Available Space (MB): 48
Assigned Capacity (MB): 48
Maximum Extension (MB): 0
Maximum Reduction (MB): 44
Page Size (bytes): 4,096
Total Usable Pages: 11,776
Used Pages: 89
Pct Util: 0.8
Max. Pct Util: 5.0
Physical Volumes: 2
Log Pool Pages: 128
Log Pool Pct. Util: 1.81
Log Pool Pct. Wait: 0.00
Cumulative Consumption (MB): 73.22
Consumption Reset Date/Time: 01/24/00 16:44:11

```

12.3.1.4 Information on log volumes

Use the `query logvolume` command to display information on one or more recovery log volumes, including available, allocated, and free space on the volume. This command displays information about the specified recovery log volume and any recovery log volume copies. The next screen shows you the output for a UNIX server.

```

tsm: TSM010> query logvolume

Volume Name      Copy   Volume Name      Copy   Volume Name      Copy
(Copy 1)         Status (Copy 2)         Status (Copy 3)         Status
-----
/tsm/log/primar- Sync'd /tsm/log/copy/f- Sync'd          Undef-
y/file01         ile01          ined

```

12.3.1.5 Display sequential volume history information

Use the `query volhistory` command to display sequential volume history information that has been collected by the server.

```
tsm: TSM010> query volhist begindate=02/10/2000
```

```
      Date/Time: 02/10/00  10:04:43
      Volume Type: STGDELETE
      Backup Series:
      Backup Operation:
      Volume Seq:
      Device Class: CFILE
      Volume Name: /tsm/stgpool/00000005.BFS
      Volume Location:
      Command:
```

```
      Date/Time: 02/10/00  10:05:43
      Volume Type: STGDELETE
      Backup Series:
      Backup Operation:
      Volume Seq:
      Device Class: C3570
      Volume Name: ALF007
      Volume Location:
      Command:
```

12.3.2 Data storage

The following is a list of important Tivoli Storage Manager queries for the data storage in alphabetical order.

12.3.2.1 Client node storage utilization

Use the `query auditoccupancy` command to display information about the client node server storage utilization. The displayed information is current as of the last license audit processed by the Tivoli Storage Manager server. You can use this information to determine if and where client node storage utilization needs to be balanced. This information can also assist with billing clients for storage usage.

```
tsm: TSM010> query auditoccupancy
License information as of last audit on 02/14/00 at 10:43:37.
```

Node Name	Backup Storage Used (MB)	Archive Storage Used (MB)	Space-Managed Storage Used (MB)	Total Storage Used (MB)
DEVNODE1	5	11	0	16
PAGOPAGO	0	0	0	0
FLJI	0	0	0	0
PUTNEY	269	0	0	269
CLADAN	0	0	0	0
COCOS	0	0	0	0

12.3.2.2 Display drive information

Use the `query drive` command to display information about a drive located in a server-attached library. You may use this command to validate if all drives are online. The next screen shows you the output for a UNIX server.

```
tsm: TSM010> query drive
```

Library Name	Drive Name	Device Type	Device	ON LINE
L3570	RMT1	3570	/dev/rmt1	Yes
L3570	RMT2	3570	/dev/rmt2	Yes
M8MM	RMT0	8MM	/dev/mt0	Yes

12.3.2.3 Query a library volume

Use the `query libvolume` command to display information about one or more volumes that have been previously checked into an automated library for use by the Tivoli Storage Manager server.

```
tsm: TSM010> query libvolume
```

Library Name	Volume Name	Status	Owner	Last Use	Home Element
L3570	07D9A3	Scratch			43
L3570	080CED	Private		Data	47
L3570	081DF3	Private		BackupSet	39
L3570	083573	Scratch			40

12.3.2.4 Number of scratch tapes

You can use an SQL query to display the number of scratch tapes available in the library. You want to be sure not to run out of scratch tapes or backup operations may fail.

```
tsm: TSM010> select count(*) as NUM_SCRATCH_TAPES from LIBVOLUMES
where STATUS ='Scratch'
```

NUM_SCRATCH_TAPES
1

12.3.2.5 Sequential access storage pool media

Use the `query media` command to display information about the sequential access primary and copy storage pool volumes moved by the `move media` command.

```
tsm: TSM010> query media * stg=offdata
```

Volume Name	State	Location	Automated LibName
086094	Mountable in library		L3570
CLA001	Mountable in library		L3570

12.3.2.6 Display file space information by storage pool

Use the `query occupancy` command to display information on where a client's file spaces are stored and how much space they occupy. This command can be used to display information about file spaces that are stored in primary or copy storage pools.

```
tsm: TSM010> query occupancy
```

Node Name	Type	Filespace Name	Storage Pool Name	Number of Files	Physical Space Occupied (MB)	Logical Space Occupied (MB)
DEVNODE1	Bkup	\\a78axfw-y\c\$	TAPEDATA	18	0.13	0.13
DEVNODE1	Bkup	\\bonnie\c\$	TAPEDATA	1	5.08	5.08
DEVNODE1	Arch	\\a78axfw-y\c\$	TAPEDATA	155	10.83	10.83
PUTNEY	Bkup	/	TAPEDATA	1,921	43.41	43.41
PUTNEY	Bkup	/usr	TAPEDATA	1,687	173.89	173.89
PUTNEY	Bkup	/r1	TAPEDATA	22	29.11	29.11
PUTNEY	Bkup	/var	TAPEDATA	280	4.59	4.59
PUTNEY	Bkup	/home	TAPEDATA	98	18.48	18.48
PUTNEY	Arch	/	TAPEDATA	15	0.01	0.01

12.3.2.7 Query one or more pending mount requests

Use the `query request` command to display information about one or more pending mount requests.

```
tsm: TSM010> query req
```

```
ANR8352I Requests outstanding:
ANR8323I 001: Insert 3570 volume ALF261 R/W into entry/exit port of library
L3570 within 60 minute(s); issue 'REPLY' along with the request ID when
ready.
```

12.3.2.8 Query one or more storage pools

Use the `query stgpool` command to display information about one or more storage pools.

```
tsm: TSM010> query stgpool
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKDATA	DISK	200.0	0.0	0.0	0	0	TAPEDATA
DISKDIRS	DISK	64.0	0.0	0.0	100	70	
DISKFILE	CFILE	0.0	0.0	0.0	100	70	
NONE	DISK	0.0	0.0	0.0	90	70	
OFFDATA	COFFSITE	0.0	0.0				
OFFDIRS	COFFSITE	0.0	0.0				
TAPEDATA	C3570	50,000,000	0.0	0.0	90	70	

12.3.2.9 Query one or more storage pool volumes

Use the `query volume` command to display information about one or more storage pool volumes. This command displays information about volumes defined to Tivoli Storage Manager. The next screen shows you the output for a UNIX server.


```

tsm: TSM010> query volume

Volume Name          Storage      Device      Estimated   Pct   Volume
                    Pool Name   Class Name  Capacity    Util  Status
                    -----
/tsm/stgpool/dirs01  DISKDIRS   DISK        64.0        0.0  On-Line
/tsm/stgpool/file01  DISKDATA   DISK        200.0       0.0  On-Line
080CED               TAPEDATA   C3570      5,000.0     5.7  Filling
086094               OFFDATA    COFFSITE    0.0         0.0  Pending
CLA001               OFFDATA    COFFSITE    0.0         0.0  Pending

```

12.3.3 Client-server activity

The following is a list of important Tivoli Storage Manager queries for client-server activities in alphabetical order.

12.3.3.1 Search activity log for messages

One of the basic commands to access the activity log is the `query actlog` command. Use the `query actlog` command to search the server activity log for messages selected by a range of time or date, message number, or string expression. If you do not specify any parameters for this command, all messages that were generated in the last hour are displayed.

```

tsm: TSM010> query actlog

Date/Time          Message
-----
02/14/00  09:58:47  ANR2017I Administrator SYSADMIN issued command: HELP
enable events
02/14/00  10:00:26  ANR2017I Administrator SYSADMIN issued command: ENABLE
EVENTS actlog all nodename=*
02/14/00  10:00:26  ANR1844I ENABLE EVENTS command processed.
02/14/00  10:01:11  ANR8325I Dismounting volume 080CED - 15 minute mount
retention expired.
02/14/00  10:01:20  ANR8336I Verifying label of 3570 volume 080CED in drive
RMT1 (/dev/rmt1).
02/14/00  10:01:46  ANR8468I 3570 volume 080CED dismounted from drive RMT1
(/dev/rmt1) in library L3570.
02/14/00  10:03:14  ANR0407I Session 27 started for administrator ADMIN
(WinNT) (Tcp/Ip 9.1.151.106(1490)).
02/14/00  10:03:14  ANR0423W Session 27 for administrator ADMIN (WinNT)
refused - administrator name not registered.
02/14/00  10:03:14  ANR0405I Session 27 ended for administrator ADMIN (WinNT).
02/14/00  10:15:56  ANR0482W Session 1 for node SYSADMIN (AIX) terminated -
idle for more than 15 minutes.
02/14/00  10:16:47  ANR0407I Session 28 started for administrator SYSADMIN
(AIX) (Tcp/Ip 9.1.150.113(33241)).
02/14/00  10:54:27  ANR2017I Administrator SYSADMIN issued command: LABEL libv
l3570 ALF261
02/14/00  10:54:27  ANR0984I Process 11 for LABEL LIBVOLUME started in the
BACKGROUND at 10:54:27.
02/14/00  10:54:27  ANR8799I LABEL LIBVOLUME: Operation for library L3570
started as process 11.
02/14/00  10:54:27  ANR8323I 001: Insert 3570 volume ALF261 R/W into
entry/exit port of library L3570 within 60 minute(s);
issue 'REPLY' along with the request ID when ready.

```

To search for warning messages, you would use the query:

```
tsm: TSM010> query actlog search="ANR????W"
```

Date/Time	Message
02/14/00 10:03:14	ANR0423W Session 27 for administrator ADMIN (WinNT) refused - administrator name not registered.
02/14/00 10:15:56	ANR0482W Session 1 for node SYSADMIN (AIX) terminated - idle for more than 15 minutes.

12.3.3.2 Query scheduled and completed events

Use the `query event` command to display scheduled and completed events. This command takes two forms, depending on whether the query applies to scheduled client operations or scheduled administrative commands.

Each scheduled client operation and administrative command is called an event. The server tracks each scheduled event and records the results of each function in the database. An event record is created whenever processing of a scheduled command is started or missed.

The `query event` command may display information about an event for which there is no event record. For example, if you request information about a scheduled event in the future, the `query event` command displays the status of *Future* for the event even though no event record has been created for this event. In addition, the `query event` command displays the status for past events whose event records have already been removed from the database. In this case, the server returns the status of *Uncertain*, because the actual event status cannot be determined without the event records. The following shows the command and its output to query scheduled client events:

```
tsm: TSM010> query event * *
```

Scheduled Start	Actual Start	Schedule Name	Node Name	Status
02/14/00 14:31:05		INCR_DAILY	PAGOPAGO	Future
02/14/00 14:31:05		INCR_DAILY	FIJI	Future
02/14/00 14:31:05		INCR_DAILY	PUTNEY	Future
02/14/00 14:32:27		CMD_ORACLE_C- OLDBKP	PUTNEY	Future
02/14/00 22:00:00		SERVER_NIGHT- LY	PUTNEY	Future
02/14/00 22:00:00		WORKSTN_NIGH- TLY	COCOS	Future

The following shows the command and its output to query scheduled administrative events:

```
tsm: TSM010> query event * type=administrative
```

Scheduled Start	Actual Start	Schedule Name	Status
02/14/00 07:00:00	02/14/00 07:00:12	BACKUP_DATAB-ASE	Failed
02/14/00 11:00:00	02/14/00 11:14:29	RECLAIM_OFFD-ATA_START	Completed
02/14/00 11:00:00	02/14/00 11:14:29	RECLAIM_OFFD-IRS_STOP	Completed
02/14/00 14:00:00		RECLAIM_OFFD-ATA_STOP	Future
02/14/00 14:00:00		RECLAIM_TAPE-DATA_START	Future
02/14/00 17:00:00		EXPIRE_INVEN-TORY	Future
02/14/00 17:00:00		RECLAIM_TAPE-DATA_STOP	Future

12.3.3.3 Query one or more server processes

Use the `query process` command to display information about one or more active background processes.

```
tsm: TSM010> query proc
```

Process Number	Process Description	Status
11	LABEL LIEVOLUME	ANR8804I Labelling volume ALF261 in library L3570.
12	Database Backup	Full backup: 0 pages of 1041 backed up. Waiting for mount of scratch volume (39 seconds).

12.3.3.4 Query restartable restore sessions

Use the `query restore` command to display information about the restartable restore sessions recorded in the server database. Certain restore operations invoke a special protocol with Tivoli Storage Manager. These special restores are called restartable restore sessions. It is desirable to know which restores are using this protocol for two reasons:

1. Restartable restore sessions that fail for some reason (network outage, client failure, or server outage) can be restarted from the last committed transaction where the restore operation left off. This can save valuable time because these restore operations do not have to return to the beginning of the operation. The `query restore` command displays the restartable restore sessions.
2. Restartable restore sessions lock the filespace, and do not allow files to be moved off of sequential volumes. The `query restore` command displays the restartable restore sessions and their associated filespace. The `cancel restore` command can be used to cancel a restartable restore session.

```
tsm: TSM010> query restore
```

Sess Number	Restore State	Elapsed Minutes	Node Name	Filespace Name
-1	Restartable	0	PUINEY	/r1

12.3.3.5 Query one or more client sessions

Use the `query session` command to display information on one or more administrative and client node sessions.

```
tsm: TSM010> query session
```

Sess Number	Conn. Method	Sess State	Wait Time	Bytes Sent	Bytes Recvd	Sess Type	Platform	Client Name
24	Tcp/Ip	Run	0 S	401	17.8 M	Node	AIX	PUTNEY
36	Tcp/Ip	Run	0 S	7.1 K	1.0 K	Admin	AIX	SYSADMIN
39	Tcp/Ip	IdleW	2.0 M	9.0 K	617	Node	WinNT	DEVNODE1

12.4 Tape management

This section shows you how to use Tivoli Storage Manager tapes, how to label them, and how to manage the many possible stages a tape can have. Although we use the term *tape* for simplicity, this section applies equally to optical media.

Tape is a vital component of the storage hierarchy. When you start using Tivoli Storage Manager, you have to create storage pools to hold all your data. After a storage pool is defined, you also need to define volumes to the storage pool, so that you have space to store the data. Tivoli Storage Manager allows you to use and reuse removable media to store data. One of the key concepts for Tivoli Storage Manager is the distinction between onsite data and offsite copy. For our discussion, we assume that all data that is available in the company is onsite data. All data that is remotely stored in another location is offsite data. Figure 39 shows you onsite and offsite tapes.

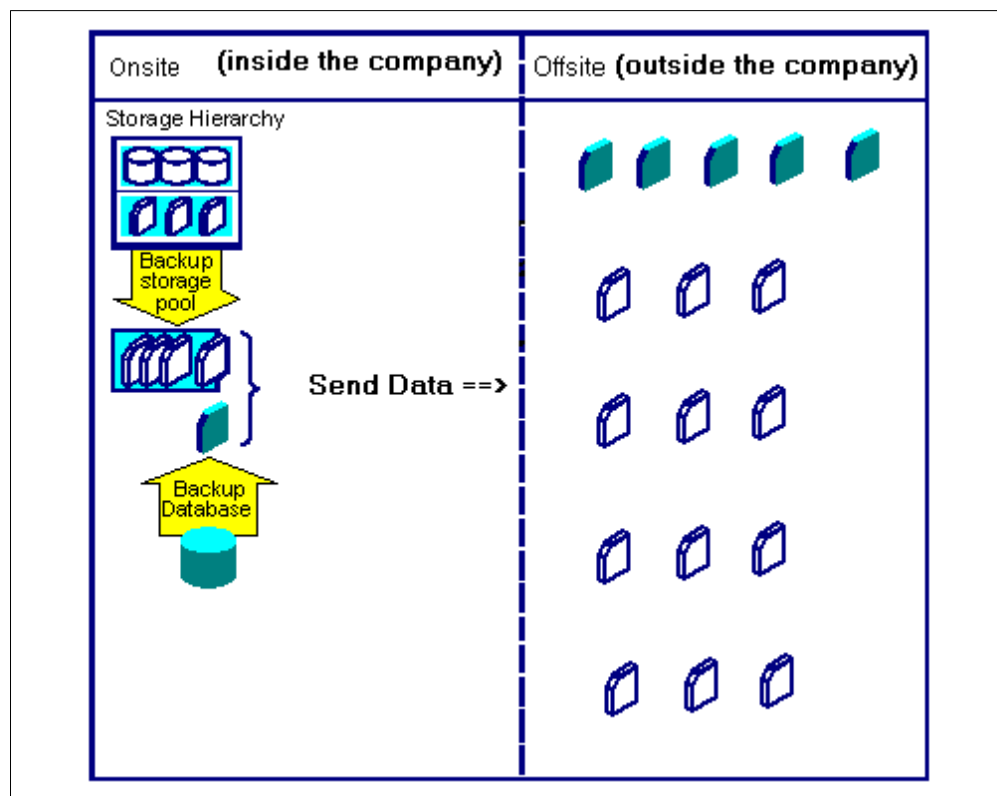


Figure 39. Onsite and offsite distinction

All tapes used for onsite data are copied to different tapes to be sent to an offsite location. Tivoli Storage Manager keeps track of all volumes (onsite and offsite) so that you do not have to care about which volumes are in the company or in the offsite location. Tivoli Storage Manager can handle both onsite and offsite tapes. Figure 40 shows how data is generically processed on tape from the moment the data is backed up until it is sent to the offsite location.

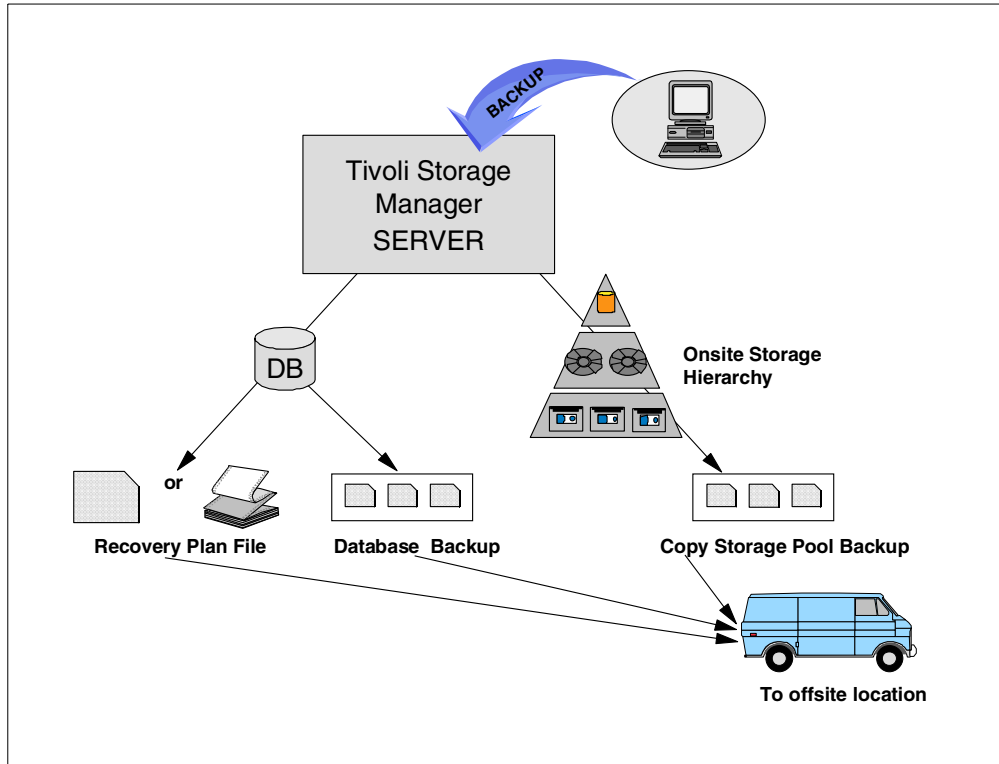


Figure 40. Tivoli Storage Manager tape processing

Providing that all your primary storage pools are being backed up, then under normal conditions, the offsite tapes contain exactly the same versions of valid data that you have onsite. This means that onsite and offsite volumes mirror the same data, correct to the time of the last storage pool backup operation. This does not mean that you have the same amount of tapes, but simply shows that what you have “inside” is the same that you have “outside” (regardless of the number of tapes).

Depending on the environment you are working in (manual devices or automated libraries), you may have different operational procedures. For example, if you have a 3494 tape library, the CHECKIN and CHECKOUT operations are part of the tape movement, with no manual intervention required. Otherwise, if you have a single I/O slot in a library (or even a manual device), you need to REPLY to the mount/unmount requests before executing the next command.

As you can see in Figure 41, the steps that you must accomplish to store data on tape are:

1. Label tape (create volume)
2. Add tape volume to Tivoli Storage Manager inventory
3. Use the tapes (Tivoli Storage Manager will do this)
4. Create offsite copies of your data
5. Send offsite volumes to safe location
6. Manage expired data (reclamation)
7. Receive free offsite volumes no longer needed (reclaimed)

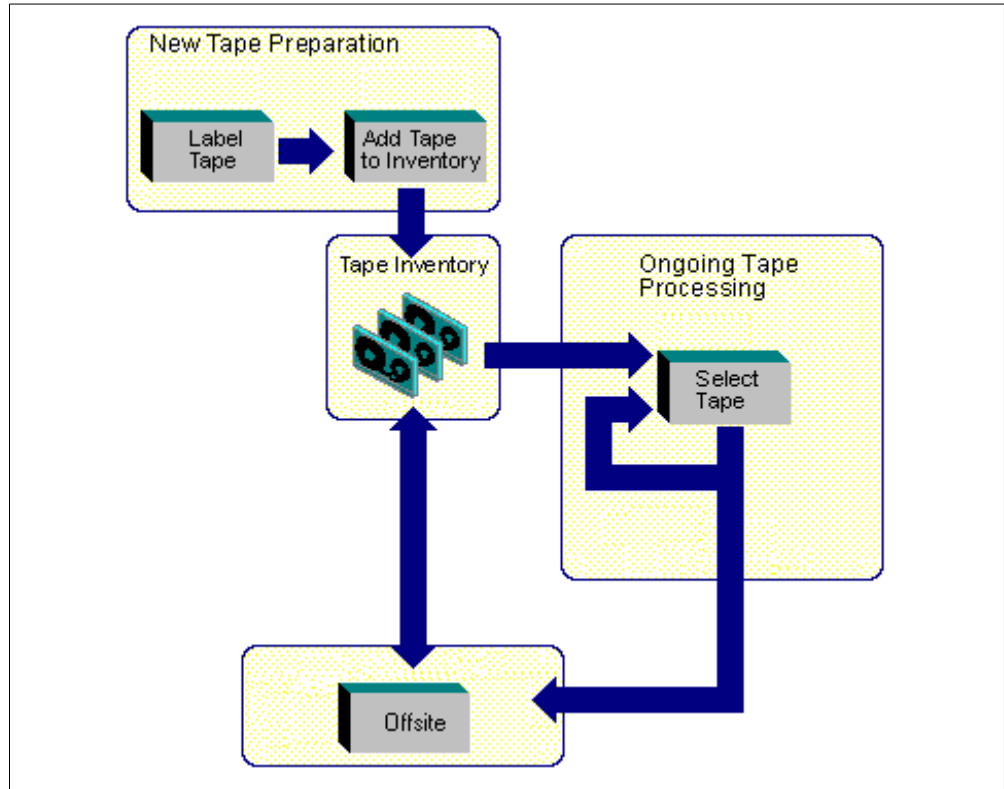


Figure 41. How Tivoli Storage Manager tapes are processed

Although, it is not in the scope of this book to show you how to use Tivoli Disaster Recovery Manager, it is a key component to successfully handle many tape management procedures. We recommend that you further evaluate Tivoli Disaster Recovery Manager as an automated approach to handling tape movement and server disaster recovery planning. As an example, Figure 42 shows a typical tape management with enhanced control of the actual state of the tapes:

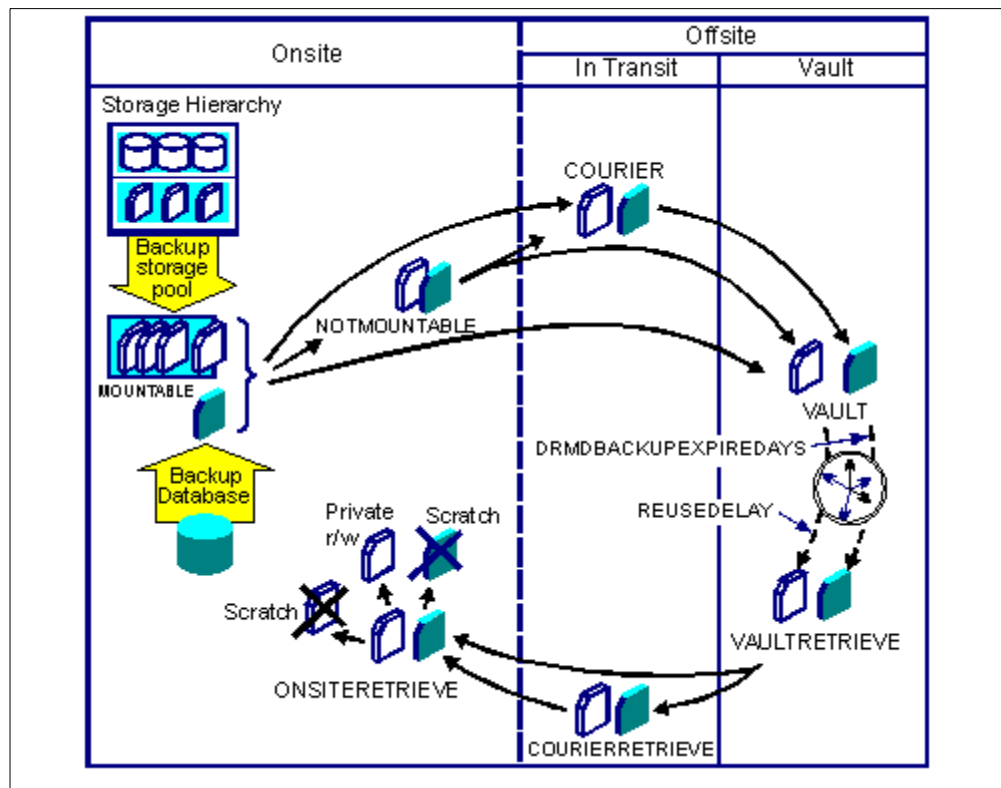


Figure 42. Tivoli Storage Manager and tape lifecycle

12.4.1 Labeling tapes

To start using new tapes, you need to first prepare removable media by using the `label libvolume` command. If you have an automated library (like the example we show), you can load the library with the new volumes and label them all with one single command. Depending on the library that you have (3494, SCSI, or MANUAL) the command syntax may be slightly different. For further details, type:

```
help label libvolume
```

Remember to use unique label names. If your tapes come already physically labeled with external codes, we recommend to use the same string for the Tivoli Storage Manager tape label. If your tape library has a barcode reader, the process of labeling pre-barcode tapes will take significantly less time and be less prone to error.

To label newly mounted tapes (scratch tapes) in a SCSI library, which has no barcode reader, use the following command:


```

tsm: TSM010> label libv l3570 search=yes labelsource=prompt checkin=scratch
ANS8003I Process number 14 started.

tsm: TSM010> query request
ANR8809I 005: Please provide the label name for the volume in slot element 37
of library L3570 by issuing REPLY n LABEL=xxx within 60 minutes, where n is
the re quest ID and xxx is the desired label name.

tsm: TSM010> reply 5 label=AVOL01
ANR8499I Command accepted.

tsm: TSM010> query actlog search=AVOL01

```

Date/Time	Message
02/14/00 13:56:29	ANR2017I Administrator SYSADMIN issued command: REPLY 5 Label=AVOL01
02/14/00 13:56:45	ANR8810I Volume AVOL01 has been labeled in library L3570.
02/14/00 13:56:58	ANR8427I CHECKIN LIBVOLUME for volume AVOL01 in library L3570 completed successfully.
02/14/00 13:57:09	ANR2017I Administrator SYSADMIN issued command: QUERY ACTLOG search=AVOL01

If you have a library of type MANUAL (for example, a single standalone 8mm drive named m8mm) and you want to label a volume as AVOL01, use the following command:

```

tsm: TSM010> label libvolume m8mm avol01 overwrite=yes
ANS8003I Process number 17 started.

tsm: TSM010> query process

```

Process Number	Process Description	Status
17	LABEL LIBVOLUME	ANR8804I Labelling volume AVOL01 in library M8MM.

```

tsm: TSM010> query actlog search=avol01

```

Date/Time	Message
02/14/00 14:07:25	ANR2017I Administrator SYSADMIN issued command: LABEL libvolume m8mm avol01 overwrite=yes
02/14/00 14:07:25	ANR8326I 007: Mount 8MM volume AVOL01 R/W in drive RMT0 (/dev/mt0) of library M8MM within 60 minutes.
02/14/00 14:09:41	ANR8372I 007: Remove 8MM volume AVOL01 from drive RMT0 (/dev/mt0) of library M8MM.
02/14/00 14:09:41	ANR8800I LABEL LIBVOLUME for volume AVOL01 in library M8MM completed successfully.
02/14/00 14:09:57	ANR2017I Administrator SYSADMIN issued command: QUERY ACTLOG search=avol01

Note:

Although Tivoli Storage Manager allows you to use a volume identification longer than six characters, we strongly recommend that you use up to six alphanumeric characters for the label (also known as VOLSER or Volume ID). This should be compatible with other ANSI styled label systems.

12.4.2 Onsite tape management to offsite

If you think of onsite tape management as the source of all data, then all commands that you run from this point on are to send data to an offsite location, which can be thought as a destination. The commands explained in this section show how to move data from onsite to offsite. Note that Tivoli Disaster Recovery Manager automatically performs most of the tape movement operations described here, reducing the need for generating, testing and maintaining separate scripts.

For manually mounted devices, Tivoli Storage Manager sends messages to request that volumes be mounted when they are needed. When Tivoli Storage Manager needs to write to a tape, it first tries to write to a previously used tape in the storage pool with enough room for new data (within the restrictions of collocation, if enabled). If that is not possible, Tivoli Storage Manager asks to mount a SCRATCH (unused) tape for use. In manual libraries, there is no inventory - that is, the command `query libvolume` does not apply. Therefore there is no concept of checking volumes in and out.

For devices in automated libraries (such as a tape autochanger), Tivoli Storage Manager interacts with the library to mount volumes, but sends messages when the library needs attention from an operator. Tivoli Storage Manager also tracks the inventory of media in each automated library. You need to manage new tapes that you have labeled by the `label libvolume` command or they are reclaimed volumes from offsite.

The steps that you need to perform to send data offsite are:

1. Backup your vital data to the copy storage pools
2. List all volumes that must be sent to offsite
3. Check out volumes from the library
4. Update the volume location to offsite

12.4.2.1 Backup your vital data

Use the `backup stgpool` command to create copies of files that reside in a primary storage pool, and store them in a copy storage pool. File versions are only copied if they do not already exist in the copy storage pool — if a file is already present in the specified copy storage pool, then it is not recopied. However, if a copy of a physical file already exists in the copy storage pool, but the copy is marked damaged, a new copy will be created, provided that the primary physical file is not also marked damaged. The following command backs up onsite data from the DISKDATA and DISKDIRS storage pool to the OFFDATA and OFFDIRS storage pools respectively. This task can also be automated using a server script as shown in 11.2.1, “Defining an offsite backup schedule” on page 218.

```
TSM:TSM010> backup stgpool diskdata offdata
ANR2017I Administrator SERVER_CONSOLE issued command: BACKUP STGPOOL diskdata
offdata
ANR0984I Process 19 for BACKUP STORAGE POOL started in the BACKGROUND at
14:21:33.
ANR2110I BACKUP STGPOOL started as process 19.
ANR1210I Backup of primary storage pool DISKDATA to copy storage pool OFFDATA
started as process 19.
ANR8337I 3570 volume 0838B5 mounted in drive RMT1 (/dev/rmt1).
ANR1340I Scratch volume 0838B5 is now defined in storage pool OFFDATA.
ANR1212I Backup process 19 ended for storage pool DISKDATA.
ANR0986I Process 19 for BACKUP STORAGE POOL running in the BACKGROUND processed
120 items for a total of 10,502,144 bytes with a completion state of SUCCESS at
14:21:58.
ANR1214I Backup of primary storage pool DISKDATA to copy storage pool OFFDATA
has ended. Files Backed Up: 120, Bytes Backed Up: 10502144, Unreadable Files:
0, Unreadable Bytes: 0.

tsm: TSM010> query volume * stg=offdata
```

Volume Name	Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Volume Status
0838B5	OFFDATA	COFFSITE	5,000.0	0.2	Filling

```
TSM:TSM010> backup stgpool diskdirs offdirs
ANR2017I Administrator SERVER_CONSOLE issued command: BACKUP STGPOOL diskdirs offdirs
ANR0984I Process 3 for BACKUP STORAGE POOL started in the BACKGROUND at 15:16:58.
ANR2110I BACKUP STGPOOL started as process 3.
ANR1210I Backup of primary storage pool DISKDIRS to copy storage pool OFFDIRS
started as process 3.
ANR8337I 3570 volume 083964 mounted in drive RMT2 (/dev/rmt2).
ANR1340I Scratch volume 083964 is now defined in storage pool OFFDIRS.
ANR1212I Backup process 3 ended for storage pool DISKDIRS.
ANR0986I Process 3 for BACKUP STORAGE POOL running in the BACKGROUND processed
104 items for a total of 106,496 bytes with a completion state of SUCCESS at
15:17:21.
ANR1214I Backup of primary storage pool DISKDIRS to copy storage pool OFFDIRS
has ended. Files Backed Up: 104, Bytes Backed Up: 106496, Unreadable Files: 0,
Unreadable Bytes: 0.
```

```
TSM:TSM010> query volume * stg=offdirs
```

Volume Name	Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Volume Status
083964	OFFDIRS	COFFSITE	5,000.0	0.0	Filling

12.4.2.2 List volumes to send to offsite

You use the `query volume` command to list all offsite tapes that need to go to the offsite location. In our example, we assume that an offsite tape is any tape that belongs to the copy storage called OFFDATA and OFFDIRS which is still available in the library, and therefore, must be moved to a separate safe location (a vault). By using the ACCESS and STATUS parameters, we can distinguish between volumes newly added to the offsite pool, which therefore need to be taken offsite, and volumes already safely offsite.

```

tsm: TSM010> query volume * access=readwrite,readonly status=full,filling \
cont> stgpool=offdata

Volume Name          Storage      Device      Estimated   Pct   Volume
Pool Name            Class Name  Capacity    Util   Status
-----
0838B5               OFFDATA     COFFSITE    5,000.0    0.2   Filling

tsm: TSM010> query volume * access=readwrite,readonly status=full,filling \
cont> stgpool=offdirs

Volume Name          Storage      Device      Estimated   Pct   Volume
Pool Name            Class Name  Capacity    Util   Status
-----
083964               OFFDIRS     COFFSITE    5,000.0    0.0   Filling

```

As you can see from the example, the `query volume` listed two volumes (0838B5 and 083964). The offsite movement operation must take into account two types of volumes: FULL and FILLING. You must issue the command for both status, so that you move all needed volumes.

We recommend that you keep a record of these volumes for tracking purposes.

12.4.2.3 Checking out volumes from the library

This command is only used if you have an automated tape library. Use the `checkout libvolume` commands to remove Tivoli Storage Manager control of a storage volume located in a library. You must use the same volume listing that you got from the previous `query volume` section. In our case, the volumes are 0838B5 and 083964:

```

TSM:TSM010> checkout libv l3570 0838b5

ANR2017I Administrator SERVER_CONSOLE issued command: CHECKOUT libv l3570 0838b5
ANR0984I Process 5 for CHECKOUT LIBVOLUME started in the BACKGROUND at 15:26:12.
ANR8434I CHECKOUT LIBVOLUME: Operation for volume 0838B5 in library L3570
started as process 5.
ANR8336I Verifying label of 3570 volume 0838B5 in drive RMT2 (/dev/rmt2).
ANR8307I 001: Remove 3570 volume 0838B5 from slot with element number 49 of
library L3570; issue 'REPLY' along with the request ID when ready.

TSM:TSM010> reply 1

ANR8499I Command accepted.
ANR8438I CHECKOUT LIBVOLUME for volume 0838B5 in library L3570 completed
successfully.
ANR0985I Process 5 for CHECKOUT LIBVOLUME running in the BACKGROUND completed
with completion state SUCCESS at 15:26:46.

TSM:TSM010> checkout libv l3570 083964

ANR2017I Administrator SERVER_CONSOLE issued command: CHECKOUT libv l3570 083964
ANR0984I Process 6 for CHECKOUT LIBVOLUME started in the BACKGROUND at 15:29:02.
ANR8434I CHECKOUT LIBVOLUME: Operation for volume 083964 in library L3570
started as process 6.
ANR8336I Verifying label of 3570 volume 083964 in drive RMT2 (/dev/rmt2).
ANR8307I 002: Remove 3570 volume 083964 from slot with element number 41 of
library L3570; issue 'REPLY' along with the request ID when ready.
ANR0403I Session 41 ended for node PUTNEY (AIX).

TSM:TSM010> reply 2

ANR2017I Administrator SERVER_CONSOLE issued command: REPLY 2
ANR8499I Command accepted.
ANR8438I CHECKOUT LIBVOLUME for volume 083964 in library L3570 completed
successfully.
ANR0985I Process 6 for CHECKOUT LIBVOLUME running in the BACKGROUND completed
with completion state SUCCESS at 15:29:38.

```

12.4.2.4 Update volume location to offsite

Before you update the volume information to the offsite location, you have to make sure that the desired volumes have been physically sent and delivered to the offsite location. To avoid a situation in which the volumes are requested for mounts while in transition to the offsite location, you can set the volumes to unavailable while in transition, by using the `update volume` command with `STATUS` set to `UNAVAILABLE`.

You can change the offsite location name to one suitable for your site. Once in the offsite location, Tivoli Storage Manager never asks for volumes with a status of offsite to be mounted.

The next example assumes that you have the volumes 0838B5 and 083964 from the previous `backup stgpool` commands and the volumes have been delivered to the offsite location. Now you want to update their location to "ITSO vault":

```

tsm: TSM010> update volume * access=offsite location="ITSO vault" \
cont> whereaccess=readwrite,readonly wherestg=offdata wherestatus=full,filling
ANR2207I Volume 0838B5 updated.

tsm: TSM010> query volume 0838b5 format=detail

          Volume Name: 0838B5
          Storage Pool Name: OFFDATA
          Device Class Name: COFFSITE
          Estimated Capacity (MB): 5,000.0
          Pct Util: 0.2
          Volume Status: Filling
          Access: Offsite
          Pct. Reclaimable Space: 99.8
          Scratch Volume?: Yes
          In Error State?: No
          Number of Writable Sides: 1
          Number of Times Mounted: 1
          Write Pass Number: 1
          Approx. Date Last Written: 02/14/00 14:21:58
          Approx. Date Last Read: 02/14/00 14:21:58
          Date Became Pending:
          Number of Write Errors: 0
          Number of Read Errors: 0
          Volume Location: ITSO vault
          Last Update by (administrator): ARMANDO
          Last Update Date/Time: 02/14/00 15:39:23

tsm: TSM010> update volume * access=offsite location="ITSO vault" \
cont> whereaccess=readwrite,readonly wherestg=offdirs wherestatus=full,filling
ANR2207I Volume 083964 updated.

tsm: TSM010> query vol 083964 format=detail

          Volume Name: 083964
          Storage Pool Name: OFFDIRS
          Device Class Name: COFFSITE
          Estimated Capacity (MB): 5,000.0
          Pct Util: 0.0
          Volume Status: Filling
          Access: Offsite
          Pct. Reclaimable Space: 100.0
          Scratch Volume?: Yes
          In Error State?: No
          Number of Writable Sides: 1
          Number of Times Mounted: 1
          Write Pass Number: 1
          Approx. Date Last Written: 02/14/00 15:17:21
          Approx. Date Last Read: 02/14/00 15:17:21
          Date Became Pending:
          Number of Write Errors: 0
          Number of Read Errors: 0
          Volume Location: ITSO vault
          Last Update by (administrator): ARMANDO
          Last Update Date/Time: 02/14/00 15:42:31

```

12.4.3 Offsite tape management to onsite

In this section we show how to move offsite volumes to the onsite location. These offsite volumes either have no active data on them (status of EMPTY) or are being returned to recreate damaged data in a primary storage pool.

Offsite tapes can be seen as your second good copy of critical data. As you perform backup and archive operations, all data is saved in storage pools, which are the basic unit for onsite operations. When you perform a `backup stgpool` command, Tivoli Storage Manager verifies which data has not previously been sent to offsite yet and then performs the copy operation of any new data. This means the operation is incremental and hence minimizes the amount of time required to execute it.

As backup and archive data age and expire, the offsite tapes contain less and less active data. Offsite tape reclamation consolidates these tape volumes, replacing the partially filled tapes with full and empty tapes. The newly emptied tapes remain in a *pending* state until after the reuse delay period has expired. At this time the status of the tape changes to EMPTY. Empty tapes can be brought back onsite and returned to the scratch pool.

The steps that you need to perform to receive tapes from offsite are the following:

1. Identify the volumes to be brought onsite.
2. Update the volume location.
3. Move the tape volumes from offsite to onsite.
4. Check in storage volumes into a library.

12.4.3.1 Identify volumes to bring onsite

You must first check which volumes are available to return onsite. The following `query volume` command shows two volumes (0838B5 and 083964) that are now available to be brought back:

```
tsm: TSM010> query volume * access=offsite status=empty
```

Volume Name	Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Volume Status
0838B5	OFFDATA	COFFSITE	0.0	0.0	Empty
083964	OFFDIRS	COFFSITE	0.0	0.0	Empty

We recommend that you keep a record of these volumes for tracking purposes and use it as a shipping list for your offsite location.

The tape volumes should now be moved from the offsite location to your onsite location.

12.4.3.2 Updating volume locations to onsite

When the previously identified tape volumes have been returned onsite, you must update the location and access mode, so that the volume can be reused. In the next example, both volumes are deleted from the storage pool location and are now considered scratch volumes:

```
tsm: TSM010> update volume * access=readwrite location="" wherestg=offdata \
cont> wherestatus=empty whereaccess=offsite
ANR2208I Volume 0838B5 deleted from storage pool OFFDATA.

tsm: TSM010> update volume * access=readwrite location="" wherestg=offdirs \
cont> wherestatus=empty whereaccess=offsite
ANR2208I Volume 083964 deleted from storage pool OFFDIRS.
```

12.4.3.3 Checking in storage volumes into a library

This command may only be used if you have an automated tape library. Use the `checkin libvolume` command to add a sequential access storage volume to Tivoli Storage Manager's inventory for an automated library. This command informs the Tivoli Storage Manager server that a volume is available for use. The server does

not use any volumes that physically reside in an automated library until a volume has been checked in via the `checkin libvolume` command. The volumes are checked in as scratch, because they are empty.

```
TSM:TSM010> checkin libvolume l3570 search=yes status=scratch
ANR2017I Administrator SERVER_CONSOLE issued command: CHECKIN libvolume l3570
search=yes status=scratch
ANR0984I Process 12 for CHECKIN LIBVOLUME started in the BACKGROUND at 15:57:36.
ANR8422I CHECKIN LIBVOLUME: Operation for library L3570 started as process 12.
ANR8430I Volume 0838B5 has been checked into library L3570.
ANR8430I Volume 083964 has been checked into library L3570.
ANR8431I CHECKIN LIBVOLUME process completed for library L3570; 2 volume(s)
found.
ANR0985I Process 12 for CHECKIN LIBVOLUME running in the BACKGROUND completed
with completion state SUCCESS at 15:58:29.
```

12.4.4 Reclaiming offsite tapes

Offsite storage pool volumes expire their data in the same way as onsite volumes. The expiration process run by the Tivoli Storage Manager updates expired file version for both onsite and offsite volumes. You do not need to mount the offsite volume for this — it is done entirely as a server database update operation.

Offsite tapes also become emptier over time and need to be reclaimed. We need the reclamation process so that the size of our offsite storage pool does not grow indefinitely. We already know that reclaiming onsite volumes requires mounting the volume to be reclaimed, and copying its valid data onto a new tape, thereby freeing the empty tape for reuse. It would be very inconvenient if reclaiming offsite volumes also required them to be mounted. Fortunately, this is not the case. When an offsite volume is eligible for reclamation, Tivoli Storage Manager will determine the ONSITE volumes, which contain the same valid data remaining on the offsite tape. It mounts the onsite tape and copies those files onto a new or existing volume belonging to the copy storage pool. This volume is then eligible to be transported offsite and the old offsite volume now switches to PENDING and then the EMPTY state. We can then bring it back onsite for reuse.

Tivoli Storage Manager will only ask to bring the old (unused offsite) tape after you send the new (and now true) tape to offsite.

12.4.5 Database backup management

The management of Tivoli Storage Manager database backups is a critical task in the management of your environment. Backups of the database should be done every day, and the copies should be moved offsite to provide disaster recovery.

The steps involved in the movement of backups from onsite to offsite, and their return, are similar to those for storage pool data. Tivoli Disaster Recovery Manager simplifies the management of onsite-offsite tape movement for both storage pool and database data. We highly recommend using Tivoli Disaster Recovery Manager.

12.4.5.1 Moving database backups offsite

The steps that you need to perform to send database backups offsite are:

1. Back up your database.

2. Identify the volumes to be taken offsite.
3. Remove volumes from the library.
4. Move the volumes offsite.

Back up your database

You can perform a backup of the database at any time, by using the `backup db` command. For examples, refer to 4.8, “Database backup” on page 72.

This task can also be automated using a server script as shown in 11.2.1, “Defining an offsite backup schedule” on page 218.

Database backup volumes are tracked in a different manner to those of storage pools. You can use the `query volhistory` command to identify tape volumes containing database backups. Alternately, you can use the activity log messages to identify those tapes. Assuming that the database backup has been done, the command to display the volumes is:

```

tsm: TSM010> query volhist type=dbbackup begindate=today

      Date/Time: 02/16/00  16:03:24
      Volume Type: BACKUPFULL
      Backup Series: 17
      Backup Operation: 0
      Volume Seq: 1
      Device Class: C3570
      Volume Name: 083983
      Volume Location:
      Command:

```

The volumes you require are those identified in the database backup with the latest time stamp.

Remove backups from library

This step only applies if your backup tapes are in a library.

You can use the `checkout libvolume` command to remove the database backup tapes from the library. This is the same process as that for storage pool data.

```

TSM:TSM010> checkout libv l3570 083983
ANR2017I Administrator SYSADMIN issued command: CHECKOUT libv l3570 083983
ANR0984I Process 17 for CHECKOUT LIBVOLUME started in the BACKGROUND at 16:06:33.
ANR8434I CHECKOUT LIBVOLUME: Operation for volume 083983 in library L3570
started as process 17.
ANR8336I Verifying label of 3570 volume 083983 in drive RMT2 (/dev/rmt2).
ANR8322I 005: Remove 3570 volume 083983 from entry/exit port of library L3570;
issue 'REPLY' along with the request ID when ready.
TSM:TSM010> reply 5
ANR2017I Administrator SERVER_CONSOLE issued command: REPLY 5
ANR8499I Command accepted.
ANR8438I CHECKOUT LIBVOLUME for volume 083983 in library L3570 completed
successfully.
ANR0985I Process 17 for CHECKOUT LIBVOLUME running in the BACKGROUND completed
with completion state SUCCESS at 16:07:10.

```

Move volumes offsite

You should move the backup volumes to your offsite location. We recommend that you track the tapes in a movement log.

12.4.5.2 Reusing backup tapes

You should keep a minimum of three days of database backup tapes. If you want to allow for holiday weekends, then a larger number, such as four or five, may be more appropriate. Our redbook environment uses five tapes.

The steps that you need to perform to reuse database backup tapes are:

1. Identify the volumes to be brought onsite.
2. Move the tape volumes from offsite to onsite.
3. Update the volume to scratch.
4. Check in storage volumes into a library.

Identify the volumes to be brought onsite

Database backups remain active until their entry is deleted from the volume history. However, once it has been deleted, Tivoli Storage Manager has no knowledge of the volume at all.

As there is no way to display just the oldest database backup volumes, we use the `query volhistory` command to identify all tape volumes containing database backups older than a specific number of days. The volumes you require are those with the oldest timestamp. The following example uses nine days.

```
tsm: TSM010> query volhistory type=dbbackup enddate=today-9

      Date/Time: 02/16/00 16:03:24
      Volume Type: BACKUPFULL
      Backup Series: 17
      Backup Operation: 0
      Volume Seq: 1
      Device Class: C3570
      Volume Name: 083983
      Volume Location:
      Command:
```

Move the tape volumes from offsite to onsite

You should move the backup volumes from your offsite location to your site. We recommend that you track the tapes in a movement log.

Update the volume to scratch

You use the `delete volhistory` command to return the database backup volumes to scratch. As there is no way to just delete the oldest database backup volumes, you must delete backups older than a number of days. The following example uses five days.

```
tsm: TSM010> delete volhistory type=dbbackup todate=today-5

Do you wish to proceed? (Yes/No) y
ANR2467I DELETE VOLHISTORY: 1 sequential volume history entries were
successfully deleted.
```

Check in storage volumes into a library

This step only applies if your backup tapes are in a library.

You can use the `checkin libvolume` command to put the database backup tapes into the library. This is the same process as that for storage pool data.

```

TSM:TSM010> checkin libvolume l3570 search=yes status=scratch
ANR2017I Administrator SERVER_CONSOLE issued command: CHECKIN libvolume l3570
search=yes status=scratch
ANR0984I Process 29 for CHECKIN LIBVOLUME started in the BACKGROUND at 16:35:12.
ANR8422I CHECKIN LIBVOLUME: Operation for library L3570 started as process 29.
ANR8430I Volume 083983 has been checked into library L3570.
ANR8431I CHECKIN LIBVOLUME process completed for library L3570; 1 volume(s) found.
ANR0985I Process 29 for CHECKIN LIBVOLUME running in the BACKGROUND completed
with completion state SUCCESS at 16:35:39.

```

12.5 Error conditions

Tivoli Storage Manager can write error information to both the Tivoli Storage Manager activity log or the operating system log. Some errors may not be related to Tivoli Storage Manager. Thus, it may detect a system failure which may be related to hardware problems, or even a software component reporting error conditions.

12.5.1 Tivoli Storage Manager errors

Error analysis is a key point to your Tivoli Storage Manager environment. If for any reason Tivoli Storage Manager is not able to write data to a tape, it may retry the operation (using hardware built-in recovery features) or it may fail due to a non-recoverable error. Tivoli Storage Manager will notify all error conditions to the console, the Tivoli Storage Manager activity log, or the operating system log. For example, you can check that Tivoli Storage Manager has an error condition logged by using the `query actlog` command with the `SEARCH` parameter:

```

TSM:TSM010> query actlog begindate=01/25/2000 search=*error*
ANR2017I Administrator SERVER_CONSOLE issued command: QUERY ACTLOG
begindate=01/25/2000 search=*error*

```

Date/Time	Message
...	
01/25/00 00:50:03	ANR8302E I/O error on drive RMT2 (/dev/zmt2) (OP=LOCATE, CC=0, KEY=03, ASC=31, ASCQ=00, SENSE=70.00.03.00.00.00.00.58.00.00.00.00.31.00.FE.0A.36-.30.40.51.00.09.01.31.08.0E.B8.50.00.00.2B.00.00.31.89.0-3.18.33.36.00.06.33.3D.00.06.00.00.00.DD.39.00.33.42.04.-00.00.03.FF.00.00.DD.73.06.00, Description=An undetermined error has occurred). Refer to Appendix B in the 'Messages' manual for recommended action.
01/25/00 00:50:03	ANR8302E I/O error on drive RMT2 (/dev/zmt2) (OP=LOCATE, CC=0, KEY=03, ASC=31, ASCQ=00, SENSE=70.00.03.00.00.00.00.58.00.00.00.00.31.00.FF.0A.3A-.5E.40.51.00.01.01.31.08.0E.10.40.00.00.2B.00.00.33.3D.0-0.00.00.00.00.00.00.00.00.00.00.00.00.00.00.DD.39.00.33.42.04.-00.00.03.FF.00.00.DD.73.06.00, Description=An undetermined error has occurred). Refer to Appendix B in the 'Messages' manual for recommended action.
....	
01/28/00 14:15:30	ANR1414W Volume SP0552 access mode is "read-only" due to previous write error.

As you can see from the previous example, there are two library errors and one volume error. The volume error may be because of a true hardware problem or a

media error. In this case, you may try to use another tape to isolate the problem or even use another tape unit. You might realize that it was just a matter of cleaning the unit.

12.5.2 Machine errors

Although Tivoli Storage Manager can handle many error conditions, it cannot resolve all possible errors. Therefore, when an unexpected error condition is met, Tivoli Storage Manager notifies the operating system event logging so that you may further isolate the problem. Note that not all error messages are true problems. Sometimes, they are simply information messages (for example, if a 3590 tape needs cleaning, then the 3494 library manager may mount the tape and create an information record indicating that a CLEAN operation was performed).

12.5.2.1 AIX errpt

You can locate Tivoli Storage Manager errors on the AIX error log by using the operating system command `errpt`. You also can use the system management interface tool (SMIT) fast path to run the `errpt` command. To use the SMIT fast path, enter `smit errpt`.

Note that the information shown in the `errpt` command may not be all from Tivoli Storage Manager. In this example, the LFTDD error is not related to Tivoli Storage Manager. The `lmcpd` and `rmt1` errors may be a Tivoli Storage Manager problem or a hardware problem. You and/or the technical staff must investigate the conditions when the problem occurs and take the required actions.

```
root@ / > errpt | pg
IDENTIFIER  TIMESTAMP  T C RESOURCE_NAME  DESCRIPTION
0F78A011    0215052600 T H rmt1           RECOVERY LOGIC INITIATED BY DEVICE
C6ACA566    0214122700 U S lmcpd         MESSAGE REDIRECTED FROM SYSLOG
C6ACA566    0214121300 U S lmcpd         MESSAGE REDIRECTED FROM SYSLOG
C6ACA566    0214083800 U S lmcpd         MESSAGE REDIRECTED FROM SYSLOG
D1A1AE6F    0212153100 I H rmt1           TAPE SIM/MIM RECORD
E507DCF9    0212152400 I H rmt1           TAPE DRIVE NEEDS CLEANING
D1A1AE6F    0212152400 I H rmt1           TAPE SIM/MIM RECORD
E85C5C4C    0212083000 P S LFTDD          SOFTWARE PROGRAM ERROR
E85C5C4C    0212083000 P S LFTDD          SOFTWARE PROGRAM ERROR
0BA49C99    0212082600 T H scsi0        SCSI BUS ERROR
2BFA76F6    0212075000 T S SYSPROC      SYSTEM SHUTDOWN BY USER
9DBCDFDEE   0212083000 T O errdemon     ERROR LOGGING TURNED ON
192AC071    0212074800 T O errdemon     ERROR LOGGING TURNED OFF
8BDD5B42    0212062600 T H pdisk17      DISK OPERATION ERROR
```

You can use the `errpt -a` command to look at the error details, as shown:

```
root@ / > errpt -a

LABEL:SYSLOG
IDENTIFIER:C6ACA566

Date/Time:      Fri Feb 11 02:12:42
Sequence Number: 41826
Machine Id:     000D08394C00
Node Id:        sm010
Class:          S
Type:           UNKN
Resource Name:  lmcpsd

Description
MESSAGE REDIRECTED FROM SYSLOG

User Causes
OPERATOR REDIRECTED SYSLOG MESSAGES TO ERROR LOG

Recommended Actions
REVIEW DETAILED DATA

Detail Data
SYSLOG MESSAGE
ERROR on 3494a, volume SP0380, ERA 83 Library Drive Exception
```

12.5.2.2 Windows NT event viewer

Tivoli Storage Manager logs error, information, and warning conditions to the Windows NT application log. You can browse it by using the Event Viewer utility which comes with Windows NT. Figure 43 shows you an example of a Windows NT error screen with some Tivoli Storage Manager errors. Note that not all of these may be true error conditions. You must evaluate if it is actually a problem or simply a misconfiguration.

Date	Time	Source	Category	Event	User	Computer
2/16/00	10:58:01 AM	AdsmClientService	None	4097	SYSTEM	A78AXFWY
2/16/00	10:53:59 AM	AdsmClientService	None	4097	SYSTEM	A78AXFWY
2/16/00	10:53:56 AM	AdsmClientService	None	4103	SYSTEM	A78AXFWY
2/16/00	10:53:56 AM	AdsmClientService	None	4097	SYSTEM	A78AXFWY
2/15/00	4:50:52 PM	IBM Personal Comrn	None	3008	N/A	A78AXFWY
2/15/00	2:36:31 PM	IBM Personal Comrn	None	3008	N/A	A78AXFWY
2/15/00	2:33:42 PM	AdsmClientService	None	4099	SYSTEM	A78AXFWY
2/15/00	2:33:42 PM	AdsmClientService	None	4099	SYSTEM	A78AXFWY
2/15/00	2:33:42 PM	AdsmClientService	None	4097	SYSTEM	A78AXFWY
2/15/00	2:33:42 PM	AdsmClientService	None	4099	SYSTEM	A78AXFWY
2/15/00	2:33:42 PM	AdsmClientService	None	4099	SYSTEM	A78AXFWY
2/15/00	2:33:42 PM	AdsmClientService	None	4097	SYSTEM	A78AXFWY
2/15/00	1:57:53 PM	IBM Personal Comrn	None	3008	N/A	A78AXFWY
2/15/00	1:55:54 PM	IBM Personal Comrn	None	3008	N/A	A78AXFWY
2/15/00	1:54:29 PM	AdsmClientService	None	4097	SYSTEM	A78AXFWY
2/15/00	1:54:27 PM	AdsmClientService	None	4097	SYSTEM	A78AXFWY
2/15/00	1:54:16 PM	ADSMServer	None	3	N/A	A78AXFWY
2/15/00	11:35:13 AM	ADSMServer	None	27	N/A	A78AXFWY
2/15/00	11:35:13 AM	ADSMServer	None	27	N/A	A78AXFWY
2/15/00	11:35:13 AM	ADSMServer	None	6	N/A	A78AXFWY
2/15/00	11:35:13 AM	ADSMServer	None	27	N/A	A78AXFWY
2/15/00	11:35:13 AM	ADSMServer	None	27	N/A	A78AXFWY
2/15/00	11:35:13 AM	ADSMServer	None	6	N/A	A78AXFWY
2/15/00	11:35:13 AM	ADSMServer	None	27	N/A	A78AXFWY
2/15/00	11:35:13 AM	ADSMServer	None	27	N/A	A78AXFWY
2/15/00	11:35:13 AM	ADSMServer	None	6	N/A	A78AXFWY
2/15/00	11:35:13 AM	ADSMServer	None	4	N/A	A78AXFWY
2/15/00	9:38:29 AM	ADSMServer	None	4	N/A	A78AXFWY
2/15/00	9:38:29 AM	ADSMServer	None	2	N/A	A78AXFWY
2/14/00	10:09:13 AM	ADSMServer	None	4	N/A	A78AXFWY
2/14/00	10:09:08 AM	ADSMServer	(2000)	21	N/A	A78AXFWY
2/14/00	10:09:00 AM	ADSMServer	(110)	21	N/A	A78AXFWY
2/14/00	10:09:00 AM	ADSMServer	(110)	21	N/A	A78AXFWY
2/14/00	10:08:58 AM	ADSMServer	(110)	21	N/A	A78AXFWY
2/14/00	10:08:58 AM	ADSMServer	(110)	21	N/A	A78AXFWY
2/14/00	10:08:58 AM	ADSMServer	(110)	21	N/A	A78AXFWY

Figure 43. Tivoli Storage Manager error entries in Windows Event Viewer

In this example, we selected for further analysis, one of the red events (indicated by the small Stop sign), which is an error condition. Figure 44 shows us that a severe error occurred (most probably, an attempt to start Tivoli Storage Manager when it was already running).

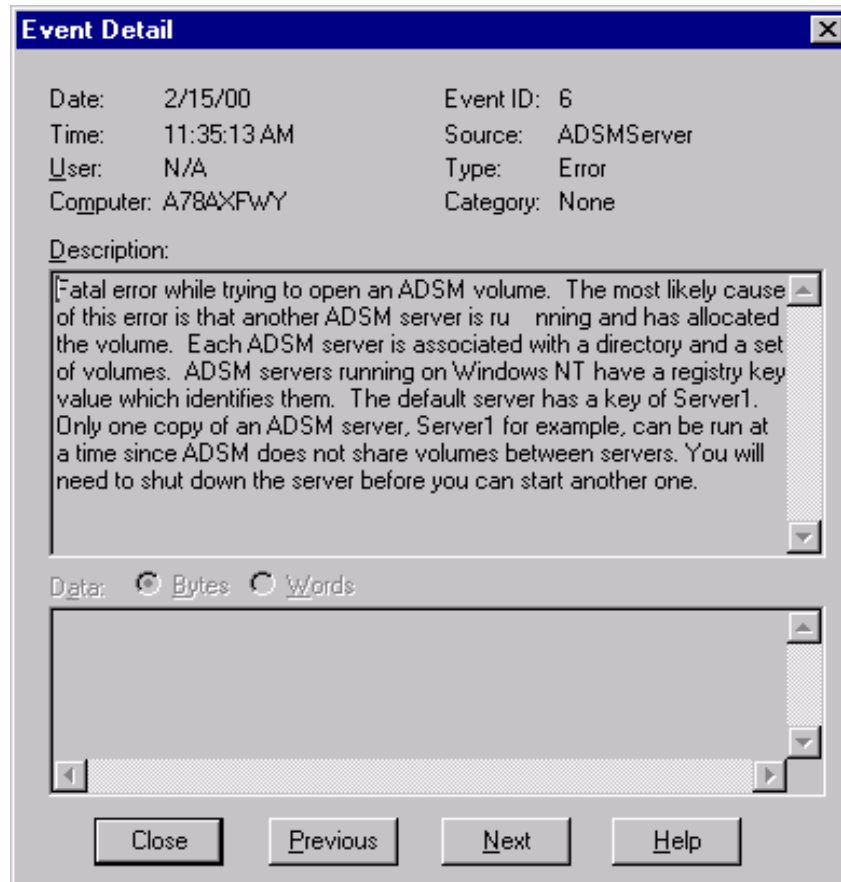


Figure 44. Detailed event information of a Tivoli Storage Manager error

Chapter 13. Performance considerations

This chapter focuses on some parameters that should be tuned to get the maximum performance during backup, archive, restore and retrieve operations. The first section is a discussion of the parameters applicable on the Tivoli Storage Manager server that are self-tuned, as well as the general performance improvements that you should consider for your environment. The next section covers the client node parameters.

Remember that a Tivoli Storage Manager implementation is typically quite complex, covering a range of variables which could affect performance. In this section we discuss only Tivoli Storage Manager-specific parameters but bear in mind that these must be taken in the context of the entire environment. We will discuss the wider performance picture in 13.3, “How to measure performance” on page 276.

13.1 Tivoli Storage Manager server

This section is a discussion of the most common tuning parameters applicable to all Tivoli Storage Manager servers. There are other tuning parameters that can also be considered, but due to the scope limitations of this guide, we will not be able to discuss all of them.

13.1.1 Self tuning parameters

This section covers the Tivoli Storage Manager self-tuning parameters that the server uses to get optimum performance. These are adaptive, auto-configured settings that require little or no intervention from the administrator.

13.1.1.1 BUFPOOLSIZE

Cache storage is provided by the database buffer pool size, which allows frequently used database pages to remain cached in memory. This means faster access times and better performance because the pages can be read and updated from memory instead of requiring I/O operations to external storage. The larger the database buffer pool, the more space there is to keep pages in memory. However since the buffer pool is implemented in real memory, its size must be balanced against total memory requirements for the whole system and any other workloads.

You can measure the effectiveness of your database buffer pool by monitoring the cache hit percentage. This number tells you how often a database page request could be satisfied from the cache as opposed to requiring a disk access. An optimal setting for the database buffer pool is one in which the cache hit percentage is greater than or equal to 98%. To check the cache hit percentage, use the `query db` command with `FORMAT=DETAILED`.

```

tsm: TSM010> query db format=detail

    Available Space (MB): 100
    Assigned Capacity (MB): 100
    Maximum Extension (MB): 0
    Maximum Reduction (MB): 88
    Page Size (bytes): 4,096
    Total Usable Pages: 25,600
    Used Pages: 391
    Pct Util: 1.5
    Max. Pct Util: 8.3
    Physical Volumes: 15
    Buffer Pool Pages: 512
    Total Buffer Requests: 1,525,110
    Cache Hit Pct.: 99.79
    Cache Wait Pct.: 0.00
    Backup in Progress?: No
    Type of Backup In Progress:
    Incrementals Since Last Full: 0
    Changed Since Last Backup (MB): 1.26
    Percentage Changed: 82.61
    Last Complete Backup Date/Time: 02/17/00 16:10:38

```

Increasing the BUFPOOLSIZ parameter can improve the performance of many Tivoli Storage Manager server functions, such as multi-client backup, storage pool migration, storage pool backup, expiration processing, and move data. If the cache hit percentage is lower than 98%, increase the size of the BUFPOOLSIZ parameter in the server options file. For most servers, we recommend starting with a value of 32768, which equals 8192 database pages. If you have enough memory, increase in 1MB increments. A cache hit percentage greater than 98% is an indication that the proper BUFPOOLSIZ has been reached. However, continuing to raise BUFPOOLSIZ beyond that level may still be helpful. While increasing BUFPOOLSIZ, care must be taken not to cause paging in the virtual memory system. Monitor system memory usage to check for any increased paging after the BUFPOOLSIZ change.

Tivoli Adaptive Algorithm

Tivoli Storage Manager can use the SELFTUNEBUFPOOLSIZ YES parameter in the server options file to activate self-tuning. When you activate this option, the server will check the cache hit ratio each time inventory expiration is run. If the value is less than 98%, then it will increase the BUFPOOLSIZ by a predictable factor related to the actual cache hit ratio achieved. An upper limit for BUFPOOLSIZ with self-tuning is imposed of 10% of the real memory for AIX, HP UX, Sun Solaris and Windows NT, and 50% of the region size in OS/390. If you still need a larger BUFPOOLSIZ and it is acceptable in your environment, you will have to set the value manually.

13.1.1.2 MOVEBATCHSIZE, MOVESIZETHRESH and TXNGROUPMAX

The MOVEBATCHSIZE and MOVESIZETHRESH options help to tune the performance of the server processes that involve the movement of data between storage media. These processes include storage pool backup and restore, migration, reclamation, and move data.

These options specify the number of files that are to be moved and grouped together in a batch, within the same server transaction.

The number of client files moved for each server database transaction during a server storage pool backup or restore, migration, reclamation, or move data operation will be determined by the number and size of the files in the batch. If the number of files in the batch equals the MOVEBATCHSIZE before the cumulative size of the files becomes greater than the MOVESIZETHRESH, then the MOVEBATCHSIZE is used to determine the number of files moved or copied in the transaction. If the cumulative size of files being gathered for a move or copy operation exceeds the MOVESIZETHRESH value before the number of files becomes equivalent to the MOVEBATCHSIZE, then the MOVESIZETHRESH value is used to determine the number of files moved or copied in the transaction.

Note that when the MOVEBATCHSIZE or MOVESIZETHRESH parameters are increased from their default values, the server will require more space in the recovery log. The recovery log may require an allocation space two or more times larger than a recovery log size which uses the defaults. In addition, the server requires a longer initialization time at startup. The impact of a larger recovery log size will be felt while running the server with the logmode set to NORMAL (the default value). If you choose to increase these values for performance reasons, be sure to monitor recovery log usage during the first few storage pool backup/restore, migration, reclamation, or move data executions to ensure sufficient recovery log space is available.

TXGROUPMAX specifies the number of files that are transferred as a group between a client and the server between transaction commit points. You may improve the performance of the backup/archive client BACKUP, ARCHIVE, RESTORE, and RETRIEVE commands by using a larger value for this option.

Tivoli Adaptive Algorithm

Tivoli Storage Manager can be configured to auto configure and monitor these parameters, so that it can increase its ranges for better throughput.

When the SELFTUNETXNSIZE parameter is set to YES in the server options file, the MOVEBATCHSIZE and MOVESIZETHRESH server options will both be set to their maximum possible values to optimize server throughput. Also, the server will monitor and change the value of TXNGROUPMAX to optimize client-server throughput.

13.1.2 General performance parameters

This parameters can affect system performance and can be manually tuned according to customer requirements.

Note:

The following parameters must be set in the server options file. You need to restart the Tivoli Storage Manager server for the changes to take effect.

13.1.2.1 EXPINTERVAL

The Tivoli Storage Manager server runs automatic inventory expiration, and this option specifies the interval in hours for this process. Inventory expiration removes client backup and archive file copies from the server.

Backup and archive copy groups can specify the criteria that make copies of files eligible for deletion from data storage. However, even when a file becomes eligible for deletion, the file is only marked for deletion in the server database. It is not physically deleted from the database and storage pools until expiration processing occurs. If expiration processing does not occur periodically, storage pool space is not reclaimed from expired client files, and the Tivoli Storage Manager server requires increased disk storage space.

Expiration processing is very CPU intensive. If possible, it should be run when other Tivoli Storage Manager processes are not occurring. To enable this, either schedule expiration once per day, or set EXPINTERVAL to 0 and manually start the process with the `expire inventory` command at the server.

Tivoli Adaptive Algorithm

If SELFTUNEBUFPOOLSIZE is set to YES, `expire inventory` processing will also trigger the buffer pool performance monitoring algorithm to check if a performance improvement could be gained (based on previous server execution and available real memory).

13.1.2.2 LOGPOOLSIZE

This parameter specifies the size of the recovery log buffer pool size in kilobytes. A large recovery log buffer pool may increase the rate by which recovery log transactions are committed to the database, but it also requires more memory. The recovery log buffer pool is used to hold new transaction records until they can be written to the recovery log. The size of the recovery log buffer pool can affect the frequency in which the server forces records to the recovery log. To determine if LOGPOOLSIZE should be increased, monitor the value of Log Pool Percentage Wait. To check the wait percentage, use the `query log` command with `FORMAT=DETAIL`.

```
tsm: TSM010> query log format=detail

    Available Space (MB): 48
    Assigned Capacity (MB): 48
    Maximum Extension (MB): 0
    Maximum Reduction (MB): 44
      Page Size (bytes): 4,096
    Total Usable Pages: 11,776
      Used Pages: 179
      Pct Util: 1.5
    Max. Pct Util: 3.4
    Physical Volumes: 3
      Log Pool Pages: 128
    Log Pool Pct. Util: 0.26
    Log Pool Pct. Wait: 0.00
    Cumulative Consumption (MB): 116.57
    Consumption Reset Date/Time: 01/24/00 16:44:11
```

If the value is greater than zero, increase the value of LOGPOOLSIZE. As the size of the recovery log buffer pool is increased, remember to monitor system memory usage.

13.1.2.3 MAXSESSION

The MAXSESSION parameter specifies the maximum number of simultaneous client sessions that can connect with the Tivoli Storage Manager server. The default value is 25 client sessions. The minimum value is two client sessions. The maximum value is limited only by available virtual memory or communication resources. This parameter specifies the maximum number of simultaneous client sessions that can connect with the Tivoli Storage Manager Server. By limiting the number of clients, server performance can be improved, but the availability of Tivoli Storage Manager services to the clients is reduced.

Due to the multi-threading features of Tivoli Storage Manager clients, you should consider the total number of concurrent nodes, times the number of concurrent sessions that the server must handle. If you need 10 concurrent nodes to perform backup operations and each of them has a multi-threading capability of (resource utilization) of 4, then you must set MAXSESSION to at least 40 (10 x 4).

13.1.2.4 MAXSCHEDESESSIONS

The MAXSCHEDESESSIONS parameter specifies the number of sessions that the server can use for processing scheduled operations. It is specified as a percentage of the total number of server sessions available (in the MAXSESSION parameter). You should set this to less than 100, otherwise you could be in a situation where all the available sessions are being used for running client scheduled backups and a critical restore cannot run because there are no spare session slots.

You can display the current value of MAXSCHEDESESSIONS with the `query status` command and change it to another value by using the `set maxschedsessions` command.

13.1.2.5 USELARGEBUFFER

This parameter increases communication and device I/O buffers. Both the client-server communication buffer and disk device I/O buffers have been increased from 32 KB to 256 KB. The communication I/O buffer is used during data transfer with a client session, such as a backup session. The disk I/O buffer is used when data is read from or written to a disk storage pool.

Significant improvement in data transfer operations and CPU usage has been observed when the USELARGEBUFFER feature is enabled. Increasing the buffer sizes allows client-server communications and disk I/O to be more efficiently performed than environments without this feature enabled. Reads and writes are quicker and server resources are better utilized. By reducing CPU utilization, more clients can concurrently be serviced, improving overall system performance. With Tivoli Storage Manager these benefits are complemented with server aggregation. Aggregation groups smaller, logical client files into fewer but larger physical files at the server level. Larger files are better able to take advantage of the larger buffers.

The USELARGEBUFFER option is enabled by default.

13.1.2.6 TCPWINDOWSIZE

This option specifies the size of the TCP sliding window in kilobytes. The TCPWINDOWSIZE option overrides the operating system's TCP send and receive spaces. In AIX for instance, these parameters are named `tcp_sendspace` and `tcp_recvspace`. The TCPWINDOWSIZE option specifies the size of the TCP sliding window for all clients and all but MVS servers. A larger window size can improve communication performance, but uses more memory. It enables multiple frames to be sent before an acknowledgment is obtained from the receiver. If long transmission delays are being observed, increasing the TCPWINDOWSIZE may improve throughput.

The size of the TCP/IP buffer is used when sending or receiving data. The window size used in a session is the smaller of the server and client window sizes. Larger window sizes use additional memory but may improve performance.

13.2 Client node

This section focuses on some tuning parameters on the Tivoli Storage Manager client to obtain maximum performance when using Tivoli Storage Manager. These parameters can be changed or added to the client options file.

13.2.1 COMPRESSION

The compression option compresses files before you send them to the Tivoli Storage Manager server. Compressing your files decreases the amount of data storage that is required to store backup versions and archive copies of your files. It can, however, affect Tivoli Storage Manager throughput.

Client data compression will save storage space, network capacity and server cycles. However, compression may have an adverse effect on throughput. If compression is on, throughputs can be significantly slower, depending on client processor speed, than if it were disabled. Compression may be beneficial for a fast processor on a slow network, but probably not for a slow processor on a fast network.

By default, files continue compressing even if the operation is causing the file size to increase. This can be the case if you are backing up already compressed files, such as ZIP archives. To prevent continued compression if the filesize grows, and to send the file again without compression, use the `COMPRESSALWAYS` option set to `NO`.

The client `COMPRESSION` option is only in effect if your Tivoli Storage Manager administrator specifies that your client node may choose its own compression. When a node is registered, using the `register node` command, there is a `COMPRESSION` parameter which allows the server to enforce compression, enforce non-compression, or allow the client to choose for itself. The default value is that the client can choose by setting the `COMPRESSION` value in the client options file. If the administrator wants to force a client to use compression, regardless of the contents of the client options file, then they use the `update node` command with the parameter `COMPRESSION` set to `YES`.

Some other general rules for using Tivoli Storage Manager compression:

- If you are backing up directly to tape, and the tape drive supports its own compaction, use the tape drive compaction. If you have already compressed

the data at the client-level, then the tape drive compaction will not be able to reduce the data further.

- Do not use Tivoli Storage Manager compression if a client currently has built-in file compression support. Tivoli Storage Manager compression on these clients will not yield additional reduction in the amount of data backed up to the server. The following platforms have built-in file compression: NetWare 4.x, Windows NT, and Windows 95.

13.2.2 COMPRESSALWAYS

The COMPRESSALWAYS option controls what occurs when a file grows during compression. You can continue compressing, or send the object again as is if it grows during compression. This option is used with the COMPRESSION option.

13.2.3 DIRMC

This option will have a big impact during restore. The DIRMC option specifies the management class you want Tivoli Storage Manager to use for backing up directories. The main benefit of managing directories separately from their associated files is faster file restore throughput from tape. The directories should be stored in a disk storage pool. The disk storage pool should not require a lot of space, since directories are typically very small.

In a restore operation, Tivoli Storage Manager first restores the directories and then the files. This order of events can have a negative impact on restore from tape. For example, if you have slow tape storage pools, and the directories are stored on tape with the files, a large amount of time can be spent searching and mounting tapes to restore the directories. This time is saved with the DIRMC option by storing the directories in a disk storage pool.

13.2.4 QUIET

The QUIET option keeps messages from being written to the screen during Tivoli Storage Manager backups. By default, Tivoli Storage Manager displays information about each file it backs up. To prevent this, use the QUIET option. However, messages and summary information are still written to the log files.

Two main benefits exist in using the QUIET option:

- For tape backup, the first transaction group of data is always resent. To avoid this, use the QUIET option to reduce retransmissions at the client.
- If you are using the client scheduler to schedule Tivoli Storage Manager backups, using the QUIET option dramatically reduces disk I/O overhead to the schedule log and improves Tivoli Storage Manager throughput.

13.2.5 LARGECOMMBUFFERS

The LARGECOMMBUFFERS option specifies whether the client will use increased buffers to transfer large amounts of data between the client and the server. You can disable this option when your machine is running low on memory.

13.2.6 TCPBUFFSIZE

The TCPBUFFSIZE option specifies the size of the internal TCP communication buffer that is used to transfer data between the client node and the server. A large buffer can improve communication performance, but requires more memory.

13.2.7 TCPWINDOWSIZE

This option specifies the size of the TCP/IP sliding window in kilobytes. The TCPWINDOWSIZE option overrides the operating system's TCP send and receive spaces.

The TCPWINDOWSIZE option specifies the size of the TCP sliding window for all clients and all but MVS servers. A larger window size can improve communication performance, but uses more memory. It enables multiple frames to be sent before an acknowledgment is obtained from the receiver. If long transmission delays are being observed, increasing the TCPWINDOWSIZE may improve throughput.

13.2.8 TXNBYTELIMIT

Specifies the batch size, in kilobytes, for Tivoli Storage Manager server transactions. The TXNBYTELIMIT option is used in conjunction with TXNGROUPMAX server option. This option reduces the number of server transactions by increasing the amount of data within any one transaction. Therefore, the amount of overhead during backup, restore, archive, and retrieve caused by database commits is reduced.

There are several items to consider when setting this parameter:

- Increasing the amount of data per transaction will increase recovery log requirements on the server. Check the log and log pool space to ensure there is enough space. Also note that a larger log may result in longer server start-up times.
- Increasing the amount of data per transaction may result in more data being retransmitted if a retry occurs. This may negatively affect performance.
- The benefits of changing this parameter are dependent on configuration and workload characteristics. In particular, this parameter benefits tape storage pool backup more than disk storage pool backup, especially if many small files are in the workload.

When setting the size of transactions, consider setting a smaller size if you are suffering from many resends due to files changing during backup when using static, shared static, or shared dynamic. This would apply to static as well as to shared, because when the client realizes a file has changed during a backup and decides to not send it, the file that is, it would still have to re-send the other files in that transaction.

13.3 How to measure performance

In measuring the performance of the Tivoli Storage Manager solution, there is a need for the above mentioned parameters to be fine tuned and the system to be restarted. Fine tuning Tivoli Storage Manager does not end with the backup, as the restore window is as important as the backup window. Some parameters might improve backup at the expense of restore and vice versa. Or they might improve both operations but require more system resources. You need to determine the best balance in your environment, depending on the frequency and criticality of each backup and restore operation.

In most cases where the backup parameters are different from the restore parameters, the backup and restore windows can be maximized.

To get the maximum performance of the system, a series of tests must be performed and recorded. It is only in this manner that the throughput can be tabulated. Tests should be conducted both for the backup and restore.

Tivoli Storage Manager performance can be influenced by various tuning parameters. Tuning these functions for good performance requires diligence and expertise on the part of the installer. The number of parameters that may be set within Tivoli Storage Manager are quite small; it is the tuning of the client, server, and network options that can become very complex.

Performance tuning for a single platform function is quite complex, but due to years of experience, it is generally well understood. However, Tivoli Storage Manager functions in the client-server realm typically supports many operating systems, works across networks, and accepts different communication protocols. Consequently, there are many more factors that affect performance. All of the factors below can affect Tivoli Storage Manager performance significantly:

- Client type
- Client speed
- Client activity
- Communication protocol type
- Communication protocol tuning
- Communication controller hardware
- Network activity
- Network speed
- Network reliability
- Server type
- Server speed
- Server activity
- Size and number of files
- Final output repository type (disk, tape, optical)

Clearly, with so many combinations and possibilities, it is not feasible to discuss all possible combinations of these parameters within the scope of this guide. However, this should be taken as advice when faced with performance questions, not to limit yourself to considering only those few Tivoli Storage Manager tunable parameters.

Appendix A. Planning and sizing worksheets

The following is a collection of worksheets we introduce and recommend to use in Chapter 1, “Tivoli Storage Manager implementation checklists” on page 1.

The redbook support material is available in softcopy on the internet from the redbooks Web server. Point your Web browser to:

<ftp://www.redbooks.ibm.com/redbooks/SG245416>

Alternatively, you can get to the same Web page at:

<http://www.redbooks.ibm.com>

Select Additional Materials and click the suggested link (or follow the instructions given, since the Web pages change frequently!), and then select SG245416/.

Table 44. Client requirements worksheet

	Client 1	Client 2	Client 3	Client 4
Client name				
Contact information				
Operating system				
Total storage available (GB)				
Total storage used (GB)				
GB changed per backup				
Number of files backed up				
Data compression				
Backup window times				
Backup number of hours				
Required recovery time frame				
Tivoli Storage Manager restore time frame				
GB copied per archive				
Number of files archived				
Number of archives kept				
Archive frequency				
Archive window times				
Archive number of hours				
Number of image backups				
Image backup frequency				
Number of backupsets				
Backupset frequency				

	Client 1	Client 2	Client 3	Client 4
Policy domain				
Client option set				

Table 45. Storage policy requirements worksheet

	Example 1	Example 2	Example 3
Group name			
Number of backup versions			
Backup file retention period			
Number of deleted versions			
Last deleted file version retention period			
Archive retention period			
Offsite copies			
Onsite collocation			
Offsite collocation			
Image backup retention			
Backupset retention			

Table 46. Database worksheet

Database volume	Filename (Primary)	Size (MB)	Filename (Copy)	Size (MB)
	Total		Total	

Table 47. Recovery log worksheet

Log Volume	Filename (Primary)	Size (MB)	Filename (Copy)	Size (MB)

Log Volume	Filename (Primary)	Size (MB)	Filename (Copy)	Size (MB)
	Total		Total	

Table 48. Primary storage pool worksheet

Filename	Size (MB)
Total	

Table 49. Device configuration and volume history worksheet

Name	Size (MB)
Total	

Table 50. Total Tivoli Storage Manager disk required worksheet

	Size (MB)
Tivoli Storage Manager code (dependent on platform)	
Tivoli Storage Manager database	
Tivoli Storage Manager recovery log	
Primary storage pools	
Device configuration table and volume history table	

	Size (MB)
Other (RAID, Operating system)	
Total	

Table 51. Tape drive configuration worksheet

	Option
Library model	
Number of drives	
Drive model	
Number of onsite tape volumes	
Number of offsite tape volumes	
Number of database volumes	
Number of scratch tapes	
Number of backupset tape volume	
Total tape volumes required	

Table 52. Administrator IDs worksheet

Functions	Tivoli Storage Manager ID	Authority
Server console	SERVER_CONSOLE	System

Table 53. License requirements worksheet

	Required
Server type	
Client connections	
Network connections	
Open Systems Environment clients	
Space management	

	Required
Tivoli Disaster Recovery Manager	
Advanced device support	
Tivoli Data Protection Features	
TDP for Oracle	
TDP for MS SQL	
TDP for MS Exchange	
TDP for Lotus Notes	
TDP for Lotus Domino	
TDP for Informix	
TDP for SAP R/3	

Appendix B. Redbook support material — macros and scripts

This appendix contains listings of support material for this redbook. See section 1.2, “Redbook support material” on page 2 for considerations regarding this material.

The redbook support material is available in softcopy on the Internet from the redbooks Web server. Point your Web browser to:

<ftp://www.redbooks.ibm.com/redbooks/SG245416>

Alternatively, you can get to the same Web page at:

<http://www.redbooks.ibm.com>

Select Additional Materials and click the suggested link (or follow the instructions given, since the Web pages change frequently!), and then select SG245416.

B.1 Macros

We have provided macros to help you implement your Tivoli Storage Manager environment. The names and values within those macros reflect the recommendations for our redbook environment. You may change them to suit your environment as required.

B.1.1 Define administrators

The following is a Tivoli Storage Manager macro which contains the administrative commands to define our administrators. The name of the macro is `mac.admins`.

```
/*=====*/
/* Redbook Macro - Define Administrators (mac.admins) */
/* Getting Started with Tivoli Storage Manager - Implementation */
/* SG24-5416-01 */
/* */
/* These scripts are supplied to help you implement a */
/* Tivoli Storage Manager */
/* environment. The names and values reflect the recommendations */
/* for our Redbook environment. You may change them to suit your */
/* environment as required. */
/* */
/* This file is designed to be run as an Administrative Command */
/* Line macro. You should use the -itemcommit parameter when you */
/* start the admin session, otherwise the macro could fail!! */
/* */
/*=====*/

/*-----*/
/* System Administrators */
/*-----*/
register admin sysadmin sysadmin contact='System Administrator'
grant authority sysadmin classes=system

/*-----*/
/* System Support Administrators */
/*-----*/
register admin support support contact='System Support'
grant authority support classes=system
```

```

/*-----*/
/* System Reporting Administrators */
/*-----*/
register admin reporter reporter contact='System Reporting'

/*-----*/
/* Client Administrators */
/*-----*/
register admin helpdesk helpdesk contact='Client Administrator'
/* grant authority helpdesk classes=node node=* */
/* You first have to register client nodes for this command to succeed! */

/*-----*/
/* Query all administrators */
/*-----*/
query admin

```

B.1.2 Define client option sets

The following is a Tivoli Storage Manager macro which contains the administrative commands to define our client option sets. The name of the macro is `mac.optionsets`.

```

/*=====*/
/* Redbook Support Materials File - (mac.optionsets) */
/* Getting Started with Tivoli Storage Manager - Implementation */
/* (SG24-5416-01) */
/* */
/* These scripts are supplied to help you implement a */
/* Tivoli Storage Manager */
/* environment. The names and values reflect the recommendations */
/* for our Redbook environment. You may change them to suit your */
/* environment as required. */
/* */
/* This file is designed to be run as an Administrative Command */
/* Line macro. You should use the -itemcommit parameter when you */
/* start the admin session, otherwise the macro could fail!! */
/* */
/*=====*/

/*--Special Note-----*/
/* Remove the comments from the following delete cloptset */
/* commands if you want to rerun this macro. */
/*-----*/
/* delete cloptset redbook */
/* delete cloptset aix */
/* delete cloptset netware */
/* delete cloptset windows */

/* Base Option Set */
define cloptset redbook description="Redbook Base Set"
define clientopt redbook changingretries 4
define clientopt redbook compressalways yes
define clientopt redbook compression off
define clientopt redbook dimc directory
define clientopt redbook domain all-local
define clientopt redbook maxcmdretries 2
define clientopt redbook memoryefficientbackup no
define clientopt redbook quiet
define clientopt redbook retryperiod 20
define clientopt redbook runasservice yes
define clientopt redbook schedmode prompted
define clientopt redbook scrolllines 20
define clientopt redbook scrollprompt no
define clientopt redbook subdir no
define clientopt redbook tapeprompt no
define clientopt redbook txbytelimit 25600

```

```

/*-----*/
/* AIX Options Set */
/*-----*/
copy cloptset redbook aix
update cloptset aix description="AIX Clients"
define clientopt aix inclexcl "exclude /unix/" seq=1
define clientopt aix inclexcl "exclude.dir /unix/" seq=2
define clientopt aix inclexcl "exclude /.../core" seq=3
define clientopt aix inclexcl "exclude /tmp/.../*" seq=4
define clientopt aix inclexcl "include /.../dsmwebcl.log special" seq=5
define clientopt aix inclexcl "include /.../dmsched.log special" seq=6
define clientopt aix inclexcl "include /.../dsmerror.log special" seq=7
define clientopt aix inclexcl "include.image /.../* image" seq=8

/*-----*/
/* Netware Options Set */
/*-----*/
copy cloptset redbook netware
update cloptset netware description="Netware Clients"
define clientopt netware inclexcl "exclude sys:vol$log.err" seq=1
define clientopt netware inclexcl "exclude sys:tts$log.err" seq=2
define clientopt netware inclexcl "exclude sys:system/sys$log.err" seq=3
define clientopt netware inclexcl "exclude sys:system/events.log" seq=4
define clientopt netware inclexcl "exclude sys:system/secaudit.log" seq=5
define clientopt netware inclexcl "exclude sys:system/system.log" seq=6
define clientopt netware inclexcl "exclude sys:system/cmaste.dba" seq=7
define clientopt netware inclexcl "exclude sys:system/btrieve.trn" seq=8
define clientopt netware inclexcl "exclude sys:system/tsa/tsa$temp.*" seq=9
define clientopt netware inclexcl "include */.../dsmwebcl.log special" seq=10
define clientopt netware inclexcl "include */.../dmsched.log special" seq=11
define clientopt netware inclexcl "include */.../dsmerror.log special" seq=12

/*-----*/
/* Windows Options Set */
/*-----*/
copy cloptset redbook windows
update cloptset windows description="Windows Clients"
define clientopt windows inclexcl "exclude *:\...\pagefile.sys" seq=1
define clientopt windows inclexcl "exclude *:\...\netlogon.chg" seq=2
define clientopt windows inclexcl "exclude *:\...\system32\config\...*" seq=3
define clientopt windows inclexcl "exclude *:\...\ntuser.dat" seq=4
define clientopt windows inclexcl "exclude *:\...\ntuser.dat.log" seq=5
define clientopt windows inclexcl "exclude *:\...\temp\...*" seq=6
define clientopt windows inclexcl "exclude *:\...\cache\*" seq=7
define clientopt windows inclexcl "exclude *:\...\recycler\*" seq=8
define clientopt windows inclexcl "exclude *:\...\Temporary Internet Files\*" seq=9
define clientopt windows inclexcl "exclude *:\microsoft uam volume\...*" seq=10
define clientopt windows inclexcl "exclude *:\ibmio.com" seq=11
define clientopt windows inclexcl "exclude *:\ibmdos.com" seq=12
define clientopt windows inclexcl "exclude *:\msdos.sys" seq=13
define clientopt windows inclexcl "exclude *:\io.sys" seq=14
define clientopt windows inclexcl "include */.../dsmwebcl.log special" seq=15
define clientopt windows inclexcl "include */.../dmsched.log special" seq=16
define clientopt windows inclexcl "include */.../dsmerror.log special" seq=17

/*-----*/
/* Query all client option sets */
/*-----*/
select * from cloptsets

```

B.1.3 Define policy structure

The following is a Tivoli Storage Manager macro which contains the administrative commands to delete the default storage pools. The name of the macro is `mac.policy`.

```

/*=====*/
/* Redbook Support Materials File - (mac.policy) */
/* Getting Started with Tivoli Storage Manager - Implementation */
/* (SG24-5416-01) */
/* */
/* These scripts are supplied to help you implement a */
/* Tivoli Storage Manager */
/* environment. The names and values reflect the recommendations */
/* for our Redbook environment. You may change them to suit your */
/* environment as required. */
/* */
/* This file is designed to be run as an Administrative Command */
/* Line macro. You should use the -itemcommit parameter when you */
/* start the admin session, otherwise the macro could fail!! */
/* */
/*=====*/

/*-----*/
/* Policy Domains */
/*-----*/
define domain server description="Server nodes" backretention=100 \
archretention=365
define domain workstn description="Workstation nodes" backretention=100 \
archretention=365

/*-----*/
/* Policy Sets */
/*-----*/
define policyset server server description="Server nodes"
define policyset workstn workstn description="Workstation nodes"

/*-----*/
/* Management Classes */
/*-----*/
define mgmtclass server server data migdestination=NONE \
description="Default management class for server domain"
assign defmgmtclass server server data
define mgmtclass server server directory migdestination=NONE \
description="Directory management class for server domain"
define mgmtclass server server special migdestination=NONE \
description="Special management class for server domain"
define mgmtclass server server image migdestination=NONE \
description="Image management class for server domain"
define mgmtclass workstn workstn data migdestination=NONE \
description="Default management class for workstn domain"
assign defmgmtclass workstn workstn data
define mgmtclass workstn workstn directory migdestination=NONE \
description="Directory management class for workstn domain"
define mgmtclass workstn workstn special migdestination=NONE \
description="Special management class for workstn domain"
define mgmtclass workstn workstn image migdestination=NONE \
description="Image management class for workstn domain"

```

```

/*-----*/
/* Copy Groups */
/*-----*/
define copygroup server server data type=Backup destination=DISKDATA \
frequency=1 verexists=3 verdeleted=1 retextra=100 retonly=365 \
mode=modified serialization=shrstatic
define copygroup server server directory type=Backup destination=DISKDIRS \
frequency=1 verexists=nolimit verdeleted=1 retextra=100 retonly=365 \
mode=modified serialization=shrstatic
define copygroup server server special type=Backup destination=DISKDATA \
frequency=1 verexists=3 verdeleted=1 retextra=100 retonly=365 \
mode=modified serialization=shrdynamic
define copygroup server server image type=backup destination=diskdata \
frequency=1 verexists=3 verdeleted=1 retextra=120 retonly=365 \
mode=modified serialization=static
define copygroup workstn workstn data type=Backup destination=DISKDATA \
frequency=1 verexists=2 verdeleted=1 retextra=30 retonly=100 \
mode=modified serialization=shrstatic
define copygroup workstn workstn directory type=Backup destination=DISKDIRS \
frequency=1 verexists=nolimit verdeleted=1 retextra=30 retonly=100 \
mode=modified serialization=shrstatic
define copygroup workstn workstn special type=Backup destination=DISKDATA \
frequency=1 verexists=2 verdeleted=1 retextra=30 retonly=100 \
mode=modified serialization=shrdynamic
define copygroup workstn workstn image type=backup destination=diskdata \
frequency=1 verexists=3 verdeleted=1 retextra=120 retonly=365 \
mode=modified serialization=static
define copygroup server server data type=Archive destination=DISKDATA \
retver=365 serialization=shrstatic
define copygroup workstn workstn data type=Archive destination=DISKDATA \
retver=100 serialization=shrstatic
/*-----*/
/* Validate policy sets */
/*-----*/
validate policyset server server
validate policyset workstn workstn

```

B.1.4 Define schedules

The following is a Tivoli Storage Manager macro which contains the administrative commands to define our schedules. The name of the macro is `mac.schedules`.

```

/*=====*/
/* Redbook Support Materials File - (mac.schedules) */
/* Getting Started with Tivoli Storage Manager - Implementation */
/* (SG24-5416-01) */
/* */
/* These scripts are supplied to help you implement a */
/* Tivoli Storage Manager */
/* environment. The names and values reflect the recommendations */
/* for our Redbook environment. You may change them to suit your */
/* environment as required. */
/* */
/* This file is designed to be run as an Administrative Command */
/* Line macro. You should use the -itemcommit parameter when you */
/* start the admin session, otherwise the macro could fail!! */
/* */
/*=====*/

/*-----*/
/* Offsite Backup Schedule */
/*-----*/
define schedule redbook_offsite type=admin cmd="run redbook_offsite" \
description="Backup all data for offsite storage" starttime=04:00 active=yes

/*-----*/
/* Volume History File Management */
/*-----*/
define schedule delete_volhist type=admin \
cmd="delete volhistory type=dbbackup todate=today-5" \
description="Delete volume history information for database backups" \
starttime=07:00 active=yes
define schedule backup_volhist type=admin cmd="backup volhistory" \
description="Backup volume history file" starttime=07:05 active=yes

```

```

/*-----*/
/* Disk Migration */
/*-----*/
define schedule MIGRATION_START type=admin cmd="update stgpool diskdata hi=0 lo=0"
description="Start migration on DISKDATA storage pool" starttime=07:00 active=yes
define schedule MIGRATION_STOP type=admin cmd="update stgpool diskdata hi=70 lo=30"
description="Stop migration on DISKDATA storage pool" starttime=10:00 active=yes

/*-----*/
/* Tape Space Reclamation */
/*-----*/
define schedule RECLAIM_OFFDIRS_START type=admin \
description="Start reclaim on the OFFDIRS storage pool" \
cmd="update stg offdirs rec=75" starttime=10:00 active=yes
define schedule RECLAIM_OFFDIRS_STOP type=admin \
description="Stop reclaim on the OFFDIRS storage pool" \
cmd="update stg offdirs rec=100" starttime=11:00 active=yes
define schedule RECLAIM_OFFDATA_START type=admin \
description="Start reclaim on the OFFDATA storage pool" \
cmd="update stg offdata rec=75" starttime=11:00 active=yes
define schedule RECLAIM_OFFDATA_STOP type=admin \
description="Stop reclaim on the OFFDATA storage pool" \
cmd="update stg offdata rec=100" starttime=14:00 active=yes
define schedule RECLAIM_TAPEDATA_START type=admin \
description="Start reclaim on the TAPEDATA storage pool" \
cmd="update stg tapedata rec=75" starttime=14:00 active=yes
define schedule RECLAIM_TAPEDATA_STOP type=admin \
description="Stop reclaim on the TAPEDATA storage pool" \
cmd="update stg tapedata rec=100" starttime=17:00 active=yes

/*-----*/
/* Inventory Expiration */
/*-----*/
define schedule EXPIRE_INVENTORY type=admin description="Inventory expiration" \
cmd="expire inventory" starttime=17:00 active=yes

/*-----*/
/* Audit Licences */
/*-----*/
define schedule AUDIT_LICENSE type=admin description="Audit licenses" \
cmd="audit licenses" starttime=00:00 active=yes

/*-----*/
/* Client Schedules */
/*-----*/
define schedule server server nightly starttime=22:00 duration=3 durunits=hours \
description="Nightly backup schedule for SERVER domain"
define schedule workstn workstn nightly starttime=22:00 duration=3 durunits=hours \
description="Nightly backup schedule for WORKSTN domain"

/*-----*/
/* Query all schedules */
/*-----*/
query schedule
query schedule type=admin

```

B.1.5 Define server scripts

The following is a Tivoli Storage Manager macro which contains the administrative commands to define our server scripts. The name of the macro is `mac.scripts`.

```

/*=====*/
/* Redbook Support Materials File - (mac.scripts) */
/* Getting Started with Tivoli Storage Manager - Implementation */
/* (SG24-5416-01) */
/* */
/* These scripts are supplied to help you implement a */
/* Tivoli Storage Manager */
/* environment. The names and values reflect the recommendations */
/* for our Redbook environment. You may change them to suit your */
/* environment as required. */
/* */
/* This file is designed to be run as an Administrative Command */
/* Line macro. You should use the -itemcommit parameter when you */
/* start the admin session, otherwise the macro could fail!! */
/* */
/*=====*/

/*-----*/
/* Define Offsite Storage Process */
/*-----*/

delete script redbook_offsite
define script redbook_offsite description="Backup all data for offsite storage"
update script redbook_offsite /*-----*/
update script redbook_offsite /* Script Name: redbook_offsite */
update script redbook_offsite /* Description: Backup data for offsite storage */
update script redbook_offsite /*-----*/
update script redbook_offsite /* If there are active node sessions, then */
update script redbook_offsite /* reschedule this script to run again in 20 */
update script redbook_offsite /* 20 minutes. If there are no active node */
update script redbook_offsite /* sessions, then backup all the onsite storage */
update script redbook_offsite /* pools and the Tivoli Storage Manager database*/
update script redbook_offsite /*-----*/

update script redbook_offsite "select * from sessions where -"
update script redbook_offsite "      upper(session_type)='NODE'"
update script redbook_offsite "if (rc_ok) goto reschedule"
update script redbook_offsite "backup stgpool diskdirs offdirs wait=yes"
update script redbook_offsite "backup stgpool diskdata offdata wait=yes"
update script redbook_offsite "backup stgpool tapedata offdata wait=yes"
update script redbook_offsite "backup db devclass=coffsite type=full scratch=yes"
update script redbook_offsite "exit"
update script redbook_offsite "reschedule:"
update script redbook_offsite "delete schedule redbook_offsite_retry type=admin"
update script redbook_offsite "define schedule redbook_offsite_retry type=admin -"
update script redbook_offsite "      cmd='run redbook_offsite' active=yes -"
update script redbook_offsite "      starttime=NOW+0:20 perunits=onetime"

/*-----*/
/* Query all redbook scripts */
/*-----*/
query script redbook*

```

B.1.6 Create storage pools

The following is a Tivoli Storage Manager macro which contains the administrative commands to create storage pools. The name of the macro is `mac.stgcreate`.

```

/*=====*/
/* Redbook Support Materials File - (mac.stgcreate) */
/* Getting Started with Tivoli Storage Manager - Implementation */
/* (SG24-5416-01) */
/* */
/* These scripts are supplied to help you implement a */
/* Tivoli Storage Manager */
/* environment. Names for device classes used should be entered */
/* between the << >> symbols and these << >> symbols removed. */
/* The MAXScratch is set purposefully at a high value to avoid */
/* any misleading out of space messages. You may set it to a more */
/* meaningful value to better reflect the %util value of your */
/* pool when q stgpool is run. */
/* */
/* This file is designed to be run as an Administrative Command */
/* Line macro. You should use the -itemcommit parameter when you */
/* start the admin session, otherwise the macro could fail!! */
/* */
/*=====*/

define stgpool tapedata <<library device class name>> highmig=100 \
maxscratch=10000 collocate=no reclaim=100 reusedelay=1
define stgpool diskdata disk nextstgpool=tapedata highmig=70 lowmig=30 \
cache=no
define stgpool diskdirs disk highmig=100
define stgpool offdirs <<offsite library device class name>> pooltype=copy \
reclaim=100 maxscratch=10000 reusedelay=5
define stgpool offdata <<offsite library device class name>> pooltype=copy \
reclaim=100 maxscratch=10000 reusedelay=5
define stgpool none disk

```

B.1.7 Delete default storage pools

The following is a Tivoli Storage Manager macro which contains the administrative commands to delete the default storage pools. The name of the macro is `mac.stgdelete`.

```

/*=====*/
/* Redbook Support Materials File - (mac.stgdelete) */
/* Getting Started with Tivoli Storage Manager - Implementation */
/* (SG24-5416-01) */
/* */
/* These scripts are supplied to help you implement a */
/* Tivoli Storage Manager */
/* environment. The names and values reflect the recommendations */
/* for our Redbook environment. You may change them to suit your */
/* environment as required. */
/* */
/* This file is designed to be run as an Administrative Command */
/* Line macro. You should use the -itemcommit parameter when you */
/* start the admin session, otherwise the macro could fail!! */
/* */
/*=====*/

delete stgpool backuppool
delete stgpool archivepool
delete stgpool spacemgpool
delete stgpool diskpool

```

B.2 Server options files

We have created server option files for the AIX, MVS, and Windows NT server environments. Our environment assumes that TCP/IP is the network protocol, Web access is enabled for administrators, and basic performance tuning values are specified.

All possible server options for each platform have been specified in the respective files. Server options have been grouped into the same categories as

they appear in the administrator reference manuals for ease of reference. Those categories are:

- Communications
- Automated Cartridge System Library
- Client-server communication processing
- Site dependent
- Database and recovery log
- Group and transfer data
- Messages
- Event logging
- Miscellaneous

Within each category, the options are ordered alphabetically. Options which are not active are preceded by an "*". All other options are active. Inactive options are those that are either not applicable for our TCP/IP communication protocol or require further setup such as those for Tivoli event monitoring.

B.2.1 AIX

The following shows our server options file for AIX. This file can also be used as the basis for a server options file for the HP/UX and Solaris platforms.

Tivoli Storage Manager HP/UX and Solaris servers only support TCP/IP as a network protocol. HP/UX does not support the STK Automated Cartridge System Library Software or the ENABLE3590LIBRARY parameters.

```
=====
* Tivoli Server Options File - AIX Version - (dmserv.aix)
* Getting Started with Tivoli Storage Manager - Implementation
* SG24-5416-01
=====

*-----*
*Communications *
*-----*
COMMETHOD HTTP
COMMETHOD SHAREDMM
COMMETHOD TCPIP
HTTPPORT 1580
HTTPSPORT 1543
*IPXBUFFERSIZE
*IPXSOCKET
*LANADAPTER
*LLNAME
*NETBIOSBUFFERSIZE
*NETBIOSNAME
*NETBIOSSESSION
SHMPORT 1510
*SNMPHEARTBEATINTERVAL
*SNMPMESSAGECATEGORY
*SNMPSUBAGENT
*SNMPSUBAGENTHOST
*SNMPSUBAGENTPORT
TCPBUFSIZE 32
TCPNODELAY YES
TCPPORT 1500
TCPWINDOWSIZE 2048
*TPNAME
*TPNPROFILENAME
```

```

*-----*
*Automated Cartridge System Library Software *
*-----*
*ACSACCESSID
*ACSLOCKDRIVE
*ACSQUICKINIT
*ACSTIMBOUTX

*-----*
* Client-Server *
*-----*
COMMTIMEOUT 60
DISABLESCHEDS YES
IDLETIMEOUT 15
MAXSESSIONS 40
USELARGEBUFFERS YES

*-----*
* Site Dependent *
*-----*
DATEFORMAT 2
LANGUAGE AMENG
NUMBERFORMAT 1
TIMEFORMAT 1

*-----*
* Database & Recovery Log *
*-----*
BUFFPOOLSIZE 32768
LOGPOOLSIZE 2048
SELTUNEBUFFPOOLSIZE YES
MIRRORREAD DB NORMAL
MIRRORREAD LOG NORMAL
MIRRORWRITE DB SEQUENTIAL
MIRRORWRITE LOG SEQUENTIAL

*-----*
* Group *
*-----*
MOVEBATCHSIZE 500
MOVESIZETHRESH 256
TXNGROUPMAX 256

*-----*
* Message Options *
*-----*
EXPQUELET YES
MESSAGEFORMAT 1
MSGINTERVAL 1
STATUSMSGCNT 10

*-----*
* Event Logging *
*-----*
EVENTSERVER YES
FILEEXIT NO
USEREXIT NO
*TECBEGINEVENTLOGGING
*TECHOSTNAME
*TECPORT

*-----*
* Miscellaneous *
*-----*
DEVCONFIG /tsm/files/primary/devconfig
DEVCONFIG /tsm/files/copy/devconfig
ENABLE3590LIBRARY NO
EXPINTERVAL 0
*NOAUDITSTORAGE
*NOPREEMPT
RESTORE INTERVAL 1440
VOLUMEHISTORY /tsm/files/primary/volumehistory
VOLUMEHISTORY /tsm/files/copy/volumehistory

```

B.2.2 MVS

The following shows our server options file for MVS.

```

*=====*
* Tivoli Server Options File - MVS Version - (dmserv.mvs) *
* Getting Started with Tivoli Storage Manager - Implementation *
* SG24-5416-01 *
*=====*

*-----*
* Communications *
*-----*
*CLIOPORT
*HTTPICSPORT
*HTTPTCSPORT 1580
*ICSSNAME
*ICSPORT
*IUCV
*LLNAME
*TCPBUFSIZE 32
*TCPNAME TCPIP
*TCPPORT 1500
*TPNPROFILENAME

*-----*
* Client-Server *
*-----*
*COMMITIMEOUT 60
*DISABLESCHEDS YES
*IDLETIMEOUT 15
*MAXSESSIONS 40
*USELARGEBUFFERS YES

*-----*
* Site Dependent *
*-----*
*DATEFORMAT 2
*LANGUAGE AMENG
*NUMBERFORMAT 1
*TIMEFORMAT 1

*-----*
* Database & Recovery Log *
*-----*
*BUFPoolsize 16384
*LOGPoolsize 2048
*SELFTUNEBUFPOLSIZE YES
*MIRRORRead DB Normal
*MIRRORRead LOG Normal
*MIRRORWrite DB Sequential
*MIRRORWrite LOG Sequential

*-----*
* Group *
*-----*
*MOVEBatchsize 500
*MOVESizethresh 256
*TXNGroupmax 256

*-----*
* Message Options *
*-----*
*EXPQUIET YES
*MESsageformat 1
*MSGHIGHLIGHT 0
*MSGSUPPRESS 0
*ROUTECode 11

*-----*
* Event Logging *
*-----*
*EVENTSERVer YES
*FILEEXIT NO
*USEREXIT NO
*TECBeginEventLogging
*TECHostname
*TECPort

```

```

*-----*
* Miscellaneous *
*-----*
DELetionexit ARCTVEXT
DEVCONFIG 'TSM.SERVER.DEVCONFG'
DEVCONFIG 'TSM.SERVER.DEVCONFG.ALT'
EXPINterval 0
*HPDTENABLE
*HPDTSTORAGE
*HPDTBUFFER
*HPDTEXPAND
LICENSE CLIENTS 1
*NOAUDITSTORAGE
*NOPREEMPT
RESTORE INTERVAL 1440
VOLUMEHistory 'TSM.SERVER.VOLHIST'
VOLUMEHistory 'TSM.SERVER.VOLHIST.ALT'

```

B.2.3 Windows NT

The following shows our server options file for Windows NT.

```

=====
* Tivoli Storage Manager V3 Redbook Server Options File - Windows NT Version *
* Getting Started with Tivoli Storage Manager - Implementation (SG24-5416) *
=====

*-----*
*Communications *
*-----*
COMMMETHOD HTTP
COMMMETHOD NAMEDPIPE
COMMMETHOD TCPIP
HTTPSPORT 1580
HTTPSPORT 1543
*IPXBUFFER SIZE
*IPXSOCKET
*LANADAPTER
NAMEDPIPENAME \\.\PIPE\ADSMPIPE
*NETBIOSBUFFER SIZE
*NETBIOSNAME
*NETBIOSSESSION
*SNMPHEARTBEATINTERVAL
*SNMPMESSAGECATEGORY
*SNMPSUBAGENT
*SNMPSUBAGENTHOST
*SNMPSUBAGENTPORT
TCPBUF SIZE 32
TCPNODELAY YES
TCPPORT 1500
TCPWINDOW SIZE 2048

*-----*
* Client-Server *
*-----*
COMMTIMEOUT 60
DISABLESCHEDS YES
IDLETIMEOUT 15
MAXSESSIONS 40
USELARGEBUFFERS YES

*-----*
* Site Dependent *
*-----*
DATEFORMAT 2
LANGUAGE AMENG
NUMBERFORMAT 1
TIMEFORMAT 1

```

```

*-----*
* Database & Recovery Log *
*-----*
BUFFPOOLSIZE 16384
SELFTUNEBUFFPOOLSIZE Yes
LOGPOOLSIZE 2048
MIRRORREAD DB NORMAL
MIRRORREAD LOG NORMAL
MIRRORWRITE DB SEQUENTIAL
MIRRORWRITE LOG SEQUENTIAL

*-----*
* Group *
*-----*
MOVEBATCHSIZE 500
MOVESTZETHRESH 256
TXNGROUPMAX 256

*-----*
* Message Options *
*-----*
EXPQUIET YES
MESSAGEFORMAT 1
MSGINTERVAL 1
STATUSMSGCNT 10

*-----*
* Event Logging *
*-----*
EVENTSERVER YES
FILEEXIT NO
USEREXIT NO
*TECBEGINEVENTLOGGING
*TECHOSTNAME
*TECPORT

*-----*
* Miscellaneous *
*-----*
DEVCONFIG c:\tsm\files\primary\devconfig
DEVCONFIG d:\tsm\files\copy\devconfig
EXPINTERVAL 0
*NOAUDITSTORAGE
*NOPREEMPT
RESTORE INTERVAL 1440
VOLUMEHISTORY c:\tsm\files\primary\volumehistory
VOLUMEHISTORY d:\tsm\files\copy\volumehistory

```

B.3 Client options files

We have created client option files for the NetWare and Windows 95/98/NT environments. Our environment assumes that TCP/IP is the network protocol, Web access is enabled for administrators, and basic performance tuning values are specified. You must replace fields surrounded by angle brackets (<< and >>) with the specific values for your site.

B.3.1 AIX

B.3.1.1 Client system options file

The following shows our client system options file for AIX. This file can also be used as the basis for a client system options file for the HP/UX and Solaris platforms.

```

=====
* Redbook Client System Options File (dsm.sys) - AIX Version *
* Getting Started with Tivoli Storage Manager - Implementation *
* SG24-5416-01 *
=====

*-----*
* Communications *
*-----*
SERVERNAME <<TSM-Server-Name>>
COMMETHOD TCPIP
TCPBUFFSIZE 32
TCPNODELAY Yes
TCPPOINT 1500
TCPSEVERADDRESS <<TSM-Server-Address>>
TCPWINDOWSIZE 58
LARGECOMMBUFFERS Yes

*-----*
* Operations *
*-----*
*ERRORLOGNAME
ERRORLOGRETENTION 10,D
PASSWORDACCESS GENERATE
SCHEDLOGNAME dsm sched.log
SCHEDLOGRETENTION 10,D

*-----*
* Site Dependent *
*-----*
***These site dependent options are supported on all UNIX
*DATEFORMAT 1
*LANGUAGE AMENG
*NUMBERFORMAT 1
*TIMEFORMAT 1

*-----*
* Include/Exclude *
*-----*
*None - specified via the Client Option Set (AIX) instead

```

B.3.1.2 Client user options file

The following shows our client user options file for AIX. This file can also be used as the basis for a client user options file for the HP/UX and Solaris platforms.

```

=====
* Redbook Client User Options File (dsm.opt) - AIX Version *
* Getting Started with Tivoli Storage Manager - Implementation *
* SG24-5416-01 *
=====

*-----*
* Communications *
*-----*
SERVERNAME <<TSM-Server-Name>>

```

B.3.2 NetWare

The following shows our client options file for NetWare.

```

=====
* Redbook Client Options File - Netware Version (dsm.opt) *
* Getting Started with Tivoli Storage Manager - Implementation *
* SG24-5416-01 *
=====

```

```

*-----*
* Identification *
*-----*
NODENAME <<Netware-Server-Name>>

*-----*
* Communications *
*-----*
COMMETHOD TCPIP
TCPBUFFSIZE 63
TCPNODELAY Yes
TCPPOINT 1500
TCPSEVERADDRESS <<TSM-Server-Name>>
TCPWINDOWSIZE 32

*-----*
* Operations *
*-----*
*ERRORLOGNAME
ERRORLOGRETENTION 10,D
PASSWORDACCESS GENERATE
REPLACE PROMPT
SCHEDLOGNAME dsmsched.log
SCHEDLOGRETENTION 10,D

*-----*
* Site Dependent *
*-----*
DATEFORMAT 1
LANGUAGE AMENG
NUMBERFORMAT 1
TIMEFORMAT 1

*-----*
* Include/Exclude *
*-----*
*None - specified via the Client Option Set (NETWARE) instead

```

B.3.3 Windows 95/98/NT

The following shows our client options file for Windows 95/98/NT.

```

=====
* Redbook Client Options File - Windows Version (dsm.opt) *
* Getting Started with Tivoli Storage Manager - Implementation *
* SG24-5416-01 *
=====

*-----*
* Communications *
*-----*
COMMETHOD TCPIP
TCPBUFFSIZE 63
TCPNODELAY Yes
TCPPOINT 1500
TCPSEVERADDRESS <<TSM-Server-Name>>
TCPWINDOWSIZE 32

*-----*
* Operations *
*-----*
*ERRORLOGNAME
ERRORLOGRETENTION 10,D
PASSWORDACCESS GENERATE
REPLACE PROMPT
SCHEDLOGNAME dsmsched.log
SCHEDLOGRETENTION 10,D

```

```
*-----*
* Site Dependent *
*-----*
DATEFORMAT 1
LANGUAGE AMENG
NUMBERFORMAT 1
TIMEFORMAT 1

*-----*
* Include/Exclude *
*-----*
*None - specified via the Client Option Set (WINDOWS) instead
```

Appendix C. Special notices

This publication is intended to help IBM customers, Business Partners, consultants, and staff implement Tivoli Storage Manager. The information in this publication is not intended as the specification of any programming interfaces that are provided by Tivoli. See the PUBLICATIONS section of the IBM Programming Announcement for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter 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 the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

This document contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples contain 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.

You can reproduce a page in this document as a transparency, if that page has the copyright notice on it. The copyright notice must appear on each page being reproduced.

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

IBM ®	Tivoli Storage Manager
AIX	AS/400
MVS/ESA	NetView
OS/2	OS/390
OS/400	

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and/or other countries licensed exclusively through X/Open Company Limited.

SET and the SET logo are trademarks owned by SET Secure Electronic Transaction LLC.

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

Appendix D. Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

D.1 International Technical Support Organization publications

For information on ordering these ITSO publications, see “How to get IBM Redbooks” on page 307.

Book Title	Publication Number
General Topics	
Tivoli Storage Management Concepts	SG24-4877
Tivoli Storage Manager Version 3.7: Technical Guide	SG24-5477
Using ADSM Hierarchical Storage Management	SG24-4631
ADSM Client Disaster Recovery: Bare Metal Restore	SG24-4880
Planning for IBM Remote Copy	SG24-2595
Specific Server Books	
A Practical Guide to Implementing Tivoli Storage Manager on AS/400	SG24-5472
Windows NT Backup and Recovery with ADSM	SG24-2231
Using TSM in a Clustered NT Environment	SG24-5742
ADSM Server-to-Server Implementation and Operation	SG24-5244
ADSM Server for Windows NT Configuration and Recovery Examples	SG24-4878
ADSTAR Distributed Storage Manager/6000 on 9076 SP2	GG24-4499
Tivoli Storage Manager with Other Products	
A Practical Guide to Network Storage Manager	SG24-2242
Using ADSM to Back Up Databases	SG24-4335
Tivoli Storage Management and Lotus Domino R5	SG24-5247
Using Tivoli Storage Manager to Back Up Lotus Notes	SG24-4534
Using ADSM to Back Up and Recover Microsoft Exchange Server	SG24-5266
Backup, Recovery, and Availability with DB2 Parallel Edition on RISC/6000	SG24-4695
ADSM Operation and Management with TME10	SG24-2214
Using Tivoli Data Protection for Workgroups	SG24-5490

D.2 Redbooks on CD-ROMs

Redbooks are also available on the following CD-ROMs. Click the CD-ROMs button at <http://www.redbooks.ibm.com/> for information about all the CD-ROMs offered, updates and formats.

CD-ROM Title	Collection Kit Number
System/390 Redbooks Collection	SK2T-2177
Networking and Systems Management Redbooks Collection	SK2T-6022
Transaction Processing and Data Management Redbooks Collection	SK2T-8038
Lotus Redbooks Collection	SK2T-8039
Tivoli Redbooks Collection	SK2T-8044
AS/400 Redbooks Collection	SK2T-2849
Netfinity Hardware and Software Redbooks Collection	SK2T-8046
RS/6000 Redbooks Collection (BkMgr Format)	SK2T-8040
RS/6000 Redbooks Collection (PDF Format)	SK2T-8043
Application Development Redbooks Collection	SK2T-8037
IBM Enterprise Storage and Systems Management Solutions	SK3T-3694

D.3 Tivoli Storage Management Publications

Book Title	Publication Number
Tivoli Storage Manager for AIX Quick Start	GC35-0367
Tivoli Storage Manager for AIX Administrator's Guide	GC35-0368
Tivoli Storage Manager for AIX Administrator's Reference	GC35-0369
Tivoli Storage Manager for HP-UX Quick Start	GC35-0370
Tivoli Storage Manager for HP-UX Administrator's Guide	GC35-0371
Tivoli Storage Manager for HP-UX Administrator's Reference	GC35-0372
Tivoli Storage Manager for Sun Solaris Quick Start	GC35-0373
Tivoli Storage Manager for Sun Solaris Administrator's Guide	GC35-0374
Tivoli Storage Manager for Sun Solaris Administrator's Reference	GC35-0375
Tivoli Storage Manager for MVS and OS/390 Quick Start	GC35-0376
Tivoli Storage Manager for MVS and OS/390 Administrator's Guide	GC35-0377
Tivoli Storage Manager for MVS and OS/390 Administrator's Reference	GC35-0378
Tivoli Storage Manager for MVS and OS/390 Quick Start	GC35-0376
Tivoli Storage Manager for Windows NT Quick Start	GC35-0379
Tivoli Storage Manager for Windows NT Administrator's Guide	GC35-0380
Tivoli Storage Manager for Windows NT Administrator's Reference	GC35-0381
Tivoli Storage Manager Messages	GC35-0382
Tivoli Storage Manager for AS/400 Quick Start	GC35-0317
Tivoli Storage Manager for AS/400 Administrator's Guide	GC35-0315
Tivoli Storage Manager for AS/400 Administrator's Reference	GC35-0316
Tivoli ADSM V3R1 VM Quick Start	GC35-0351
Tivoli ADSM V3R1 VM Administrator's Guide	GC35-0352
Tivoli ADSM V3R1 VM Administrator's Reference	GC35-0353
Tivoli ADSM Messages	GC35-0271
ADSM V2R1 and V2R6 Messages	SH35-0133
Tivoli Storage Manager for UNIX Using the Backup-Archive Client	SH26-4105
Tivoli Storage Manager for Windows Using the Backup-Archive Client	SH26-4101
Tivoli Storage Manager for NetWare Using the Backup-Archive Client	SH26-4100

Book Title	Publication Number
Tivoli Storage Manager: Installing the Clients	SH26-4102
Tivoli Storage Manager Using the Application Programming Interface	SH26-4107
Tivoli Storage Manager AFS/DFS Backup Clients	SH26-4106
Tivoli Storage Manager Trace Facility Guide	SH26-4104
Tivoli Storage Manager Reference Cards for the Backup-Archive Clients	SX26-6021
Tivoli Data Protection for Informix: Installation and User's Guide	SH26-4095
Tivoli Data Protection for Lotus Domino for UNIX: Installation and User's Guide	SH26-4088
Tivoli Data Protection for Lotus Domino for Windows NT: Installation and User's Guide	GC26-7320
ADSMConnect Agent for Lotus Notes on Windows NT Installation and User's Guide	SH26-4065
ADSMConnect Agent for Lotus Notes on AIX Installation and User's Guide	SH26-4067
Tivoli Data Protection for Microsoft Exchange Server Installation and User's Guide	SH26-4110
Tivoli Data Protection Microsoft SQL Server Installation and User's Guide	SH26-4111
Tivoli Data Protection for Oracle for UNIX Installation and User's Guide	SH26-4112
Tivoli Data Protection for Oracle on Windows NT Installation and User's Guide	SH26-4113
Tivoli Data Protection for SAP R/3 Installation and User's Guide	SC33-6388

How to get IBM Redbooks

This section explains how both customers and IBM employees can find out about IBM Redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs by fax or e-mail is also provided.

- **Redbooks Web Site** ibm.com/redbooks

Search for, view, download, or order hardcopy/CD-ROM Redbooks from the Redbooks Web site. Also read redpieces and download additional materials (code samples or diskette/CD-ROM images) from this Redbooks site.

Redpieces are Redbooks in progress; not all Redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

- **E-mail Orders**

Send orders by e-mail including information from the IBM Redbooks fax order form to:

	e-mail address
In United States or Canada	pubscan@us.ibm.com
Outside North America	Contact information is in the "How to Order" section at this site: http://www.elink.ibm.ibm.com/pbl/pbl

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	Country coordinator phone number is in the "How to Order" section at this site: http://www.elink.ibm.ibm.com/pbl/pbl

- **Fax Orders**

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	Fax phone number is in the "How to Order" section at this site: http://www.elink.ibm.ibm.com/pbl/pbl

This information was current at the time of publication, but is continually subject to change. The latest information may be found at the Redbooks Web site.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and Redbooks by accessing the IBM Intranet Web site at <http://w3.itso.ibm.com/> and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may access MyNews at <http://w3.ibm.com/> for redbook, residency, and workshop announcements.

Index

A

- access control lists
 - See ACLs
- ACLs
 - backup of directory information
- ACTIVATE POLICYSET command 141
- ADMIN, admin ID 144
- administrative client
 - ACTIVATE POLICYSET command 141
 - ADMIN ID 144
 - ASSIGN DEFMGMTCLASS command 128
 - AUDIT LICENSES command 160, 224
 - AUDIT VOLUME command 111
 - BACKUP DB command 72, 261
 - BACKUP STGPOOL command 112, 254, 257
 - BACKUP VOLHISTORY command 221
 - batch mode option of command line interface 164
 - BEGIN EVENTLOGGING command 238
 - CANCEL PROCESS command 236
 - CANCEL RESTORE command 236, 247
 - CANCEL SESSION command 209, 235
 - CHECKIN LIBVOLUME command 114, 260, 262
 - CHECKOUT LIBVOLUME command 114, 256, 261
 - client option set
 - See client option set
 - command line interface 163
 - console session 233
 - CONTACT option 143
 - COPY CLOPTSET command 155
 - DEFINE ASSOCIATION command 225, 226
 - DEFINE CLIENTOPT command 155
 - DEFINE CLOPTSET command 155
 - DEFINE COPYGROUP command 131, 132
 - DEFINE DBBACKUPTRIGGER command 63
 - DEFINE DBCOPY command 68
 - DEFINE DBVOLUME command 53, 58
 - DEFINE DOMAIN command 125
 - DEFINE LOGCOPY command 69
 - DEFINE LOGVOLUME command 53, 60
 - DEFINE MGMTCLASS command 128
 - DEFINE POLICYSET command 126
 - DEFINE SCHEDULE command 225
 - DEFINE SPACETRIGGER command 64
 - DEFINE VOLUME command 53, 109
 - DELETE CLIENTOPT command 155
 - DELETE CLOPTSET command 155
 - DELETE DBVOLUME command 70
 - DELETE DEVCLASS command 117
 - DELETE DOMAIN command 138
 - DELETE FILESPACE command 138, 151
 - DELETE LIBRARY command 117
 - DELETE LOGVOLUME command 71
 - DELETE STGPOOL command 110, 116
 - DELETE VOLHISTORY command 76, 220, 262
 - deleting client node filespaces 151
 - DISABLE SESSIONS command 234
 - DISMOUNT VOLUME command 111, 117
 - dsmadm program 163
 - ENABLE EVENTS command 238
 - EXPIRE INVENTORY command 141, 223, 272
 - EXTEND DB command 70, 74
 - EXTEND LOG command 71, 74
 - GRANT AUTHORITY command 143
 - granting node authority to administrator 150
 - HALT command 234
 - HELPDESK ID 145
 - installation 162
 - ITEMCOMMIT parameter 53
 - LABEL LIBVOLUME command 55, 99, 115, 252
 - loop mode option of command line interface 164
 - MACRO command 52
 - managing and controlling resources 161
 - MOVE DATA command 119
 - MOVE MEDIA command 243
 - privileges to manage resources 143
 - QUERY ACTLOG command 245, 263
 - QUERY ADMIN command 143, 146
 - QUERY AUDITOCCUPANCY command 224, 242
 - QUERY CONTENT command 119
 - QUERY COPYGROUP command 139
 - QUERY DB command 239, 269
 - QUERY DBBACKUPTRIGGER command 63
 - QUERY DBVOLUME command 67, 240
 - QUERY DEVCLASS command 73
 - QUERY DRIVE command 242
 - QUERY EVENT command 246
 - QUERY FILESPACE command 151
 - QUERY LIBRARY command 73
 - QUERY LIBVOLUME command 243
 - QUERY LICENSE command 159
 - QUERY LOG command 240, 272
 - QUERY LOGVOLUME command 68, 241
 - QUERY MEDIA command 243
 - QUERY NODE command 149
 - QUERY OCCUPANCY command 120, 243
 - QUERY OPTIONS command 48
 - QUERY PROCESS command 112, 152, 236, 247
 - QUERY REQUEST command 108, 114, 244
 - QUERY RESTORE command 236, 247
 - QUERY SESSION command 248
 - QUERY STATUS command 51, 62, 235, 237
 - QUERY STGPOOL command 116, 244
 - QUERY VOLHISTORY command 241, 261, 262
 - QUERY VOLUME command 222, 244, 255, 259
 - QUIESCE option 234
 - REDUCE DB command 75
 - REDUCE LOG command 76
 - REGISTER ADMIN command 143
 - REGISTER LICENSE command 157
 - REGISTER NODE command 148, 149
 - registering admin IDs 144
 - registering client nodes 149
 - REMOVE ADMIN command 147
 - REMOVE NODE command 138
 - removing client nodes 152

- RENAME NODE command 149
- RENAME STGPOOL command 120
- REPLY command 115
- RESET LOGCONSUMPTION command 62
- scheduling backup of volume history 221
- scheduling deletion of volume history 220
- scheduling inventory expiration 223
- scheduling license audit 224
- scheduling offsite backup 218
- scheduling operations 216
- scheduling reclamation 222
- scheduling storage pool migration 221
- SELECT command 52, 121, 154
- SERVER_CONSOLE ID 143, 230
- SET ACCOUNTING command 52
- SET ACTLOGRETENTION command 238
- SET AUTHENTICATION command 51
- SET INPWLENGTH command 51
- SET INVALIDPWLIMIT command 51
- SET LICENSEAUDITPERIOD command 160
- SET LOGMODE command 61
- SET PASSEXP command 51
- SET WEBAUTHTIMEOUT command 51
- TCPSERVERNAME option 163
- TSO interface on MVS 161
- UPDATE ADMIN command 147
- UPDATE NODE command 139, 149, 154
- UPDATE SPACETRIGGER command 64, 65
- UPDATE STGPOOL command 221, 222
- UPDATE VOLUME command 257
- VALIDATE POLICYSET command 141
- verifying schedule events 246
- Web administrative interface
 - See Web administrative interface
- administrator
 - See administrative client
- ARCHIVE command 189
- archive copy group
 - ARCHMC option 189
 - assigning files to archive copy group 189
 - defining a copy group 132
 - DESTINATION option 132
 - RETVR option 132
 - SERIALIZATION option 132
 - TYPE option 132
- archive description 189
- ARCHMC
 - ARCHIVE command 189
 - client option 127
- ARCHRET option, DEFINE DOMAIN command 125
- ASSIGN DEFMGMTCLASS command 128
- assigning files to specific management class 174, 189
- AUDIT LICENSES command 160, 224
- AUDIT VOLUME command 111

B

- BACKRET option, DEFINE DOMAIN command 125
- backup copy group
 - active file 128
 - assigning files to backup copy group 174
 - defining a copy group 128
 - DESTINATION option 129
 - expired file 128
 - FREQUENCY option 130
 - inactive file 128
 - INCLUDE option 174
 - MODE option 130
 - retention period 129
 - RETEXTRA option 129
 - REONLY option 129
 - SERIALIZATION option 131
 - TYPE option 129
 - VERDELETED option 129
 - VEREXISTS option 129
- BACKUP DB command 72, 261
- BACKUP IMAGE command 178
- BACKUP STGPOOL command 112, 254, 257
- BACKUP VOLHISTORY command 221
- backup-archive client
 - ARCHIVE command 189
 - archive description 189
 - archiving files to server 189
 - ARCHMC option 127
 - assigning files to archive copy group 189
 - assigning files to backup copy group 174
 - associating client with schedule 226
 - automatic start of scheduler 207
 - backing up files to server 178
 - BACKUP IMAGE command 178
 - batch mode option of command line interface 177
 - central scheduling 204
 - changing node password 151
 - CHANGINGRETRIES option 131
 - CLIENT ID 149
 - client option set
 - See client option set
 - client system options file 172
 - client user options file 172
 - command line interface 176
 - COMMETHOD option 173
 - communication options 173
 - COMPRESSALWAYS option 274
 - COMPRESSION option 274
 - deleting files during archive operation 189
 - DIRMC option 127, 275
 - DOMAIN option 148
 - dsm program 193
 - dsm.opt file 172
 - dsm.sys file 172
 - DSM_CONFIG variable 171
 - DSM_DIR variable 171
 - DSM_LOG variable 171
 - dsmc program 176
 - DSMCUTIL program 207
 - DSMTCA, client communication agent 171
 - EXCLUDE option 174
 - granting node authority to administrator 150
 - INCLEXCL option 174
 - INCLUDE option 127, 174
 - include-exclude list 127

- INCREMENTAL command 178
- LARGECOMMBUFFERS option 275
- loop mode option of command line interface 177
- manual start of scheduler 205
- NODENAME option 173
- operational options 174
- options file 172
- PASSWORDACCESS option 174, 175
- point-in-time restore 199
- Preferences file for Macintosh 172
- QUIET option 275
- QUIT command 192
- registering client nodes 149
- removing client nodes 152
- REPLACE option 174
- restartable restore sessions 199
- RESTORE command 183
- restoring backed up files to client machine 183
- restoring inactive file copies 184
- restoring most recent file copies 184
- RETRIEVE command 190
- retrieving archived files to client machine 190
- REVOKEREMOTEACCESS option
 - See Web client
- SCHEDMODE option 174
- SCHEDULE command 205
- scheduling client operation 224
- SELECTIVE command 178
- site dependent options 174
- SUBDIR option 174
- TAPEPROMPT option 174
- TCPBUFFSIZE option 173, 275
- TCPPORT option 173
- TCPSERVERADDRESS option 173
- TCPWINDOWSIZE option 173, 276
- TXNBYTELIMIT option 276
- TXNGROUPMAX option 276
- USERID option 148
- verifying schedule events 227
- banner frame, Web administrative interface 167
- batch mode of command line interface 164, 177
- BEGIN EVENTLOGGING command 238
- BUFPOOLSIZ, server option 50, 269

C

- CANCEL PROCESS command 236
- CANCEL RESTORE command 236, 247
- CANCEL SESSION command 209, 235
- CHANGINGRETRIES, client option 131
- Chargeback 45
- CHECKIN LIBVOLUME command 114, 260, 262
- CHECKLABEL option, DISMOUNT VOL command 111
- checklists
 - client implementation, client tasks 8
 - client implementation, server tasks 8
 - daily operations 9
 - planning 5
 - server implementation 6
- CHECKOUT LIBVOLUME command 114, 256, 261
- client acceptor, Web client 171, 211

- client access authority
 - See Web client
- client node
 - See backup-archive client
- client option set
 - adding options to set 155
 - associating client node 154
 - defining set 155
 - deleting set 155
 - easing administration of client nodes by centralized management 153
 - include-exclude statements 153
 - protecting option settings against overwriting 153
 - removing option from set 155
- client owner authority
 - See Web client
- CLIENT, client node ID 149
- command line client
 - Web administrative interface 165, 167
- commands, administrator
 - ACTIVATE POLICYSET 141
 - ASSIGN DEFMGMTCLASS 128
 - AUDIT LICENSES 160, 224
 - AUDIT VOLUME 111
 - BACKUP DB 72, 261
 - BACKUP STGPOOL 112, 254, 257
 - BACKUP VOLHISTORY 221
 - BEGIN EVENTLOGGING 238
 - CANCEL PROCESS 236
 - CANCEL RESTORE 236, 247
 - CANCEL SESSION 209, 235
 - CHECKIN LIBVOLUME 114, 260, 262
 - CHECKOUT LIBVOLUME 114, 256, 261
 - COPY CLOPTSET 155
 - DEFINE ASSOCIATION 225, 226
 - DEFINE CLIENTOPT 155
 - DEFINE CLOPTSET 155
 - DEFINE COPYGROUP 131, 132
 - DEFINE DBBACKUPTRIGGER 63
 - DEFINE DBCOPY 68
 - DEFINE DBVOLUME 53, 58
 - DEFINE DOMAIN 125
 - DEFINE LOGCOPY 69
 - DEFINE LOGVOLUME 53, 60
 - DEFINE MGMTCLASS 128
 - DEFINE POLICYSET 126
 - DEFINE SCHEDULE 225
 - DEFINE SPACETRIGGER 64
 - DEFINE VOLUME 53, 109
 - DELETE CLIENTOPT 155
 - DELETE CLOPTSET 155
 - DELETE DBVOLUME 70
 - DELETE DEVCLASS 117
 - DELETE DOMAIN 138
 - DELETE FILESPACE 138, 151
 - DELETE LIBRARY 117
 - DELETE LOGVOLUME 71
 - DELETE STGPOOL 110, 116
 - DELETE VOLHISTORY 76, 220, 262
 - DISABLE SESSIONS 234

DISMOUNT VOLUME 111, 117
 ENABLE EVENTS 238
 EXPIRE INVENTORY 141, 223, 272
 EXTEND DB 70, 74
 EXTEND LOG 71, 74
 GRANT AUTHORITY 143
 HALT 234
 LABEL LIBVOLUME 55, 99, 115, 252
 MACRO 52
 MOVE DATA 119
 MOVE MEDIA 243
 QUERY ACTLOG 245, 263
 QUERY ADMIN 143, 146
 QUERY AUDITOCCUPANCY 224, 242
 QUERY CONTENT 119
 QUERY COPYGROUP 139
 QUERY DB 239, 269
 QUERY DBBACKUPTRIGGER 63
 QUERY DBVOLUME 67, 240
 QUERY DEVCLASS 73
 QUERY DRIVE 242
 QUERY EVENT 246
 QUERY FILESPACE 151
 QUERY LIBRARY 73
 QUERY LIBVOLUME 243
 QUERY LICENSE 159
 QUERY LOG 240, 272
 QUERY LOGVOLUME 68, 241
 QUERY MEDIA 243
 QUERY NODE 149
 QUERY OCCUPANCY 120, 243
 QUERY OPTIONS 48
 QUERY PROCESS 112, 152, 236, 247
 QUERY REQUEST 108, 114, 244
 QUERY RESTORE 236, 247
 QUERY SESSION 248
 QUERY STATUS 51, 62, 235, 237
 QUERY STGPOOL 116, 244
 QUERY VOLHISTORY 241, 261, 262
 QUERY VOLUME 222, 244, 255, 259
 REDUCE DB 75
 REDUCE LOG 76
 REGISTER ADMIN 143
 REGISTER LICENSE 157
 REGISTER NODE 148, 149
 REMOVE ADMIN 147
 REMOVE NODE 138
 RENAME NODE 149
 RENAME STGPOOL 120
 REPLY 115
 RESET LOGCONSUMPTION 62
 SELECT 52, 121, 154
 SET ACCOUNTING 52
 SET ACTLOGRETENTION 238
 SET AUTHENTICATION 51
 SET INPWLENGTH 51
 SET INVALIDPWLIMIT 51
 SET LICENSEAUDITPERIOD 160
 SET LOGMODE 61
 SET PASSEXP 51

SET WEBAUTHTIMEOUT 51
 UPDATE ADMIN 147
 UPDATE NODE 139, 149, 154
 UPDATE SPACETRIGGER 64, 65
 UPDATE STGPOOL 221, 222
 UPDATE VOLUME 257
 VALIDATE POLICYSET 141
 commands, client
 ARCHIVE 189
 BACKUP IMAGE 178
 INCREMENTAL 178
 QUIT 192
 RESTORE 183
 RETRIEVE 190
 SCHEDULE 205
 SELECTIVE 178
 COMMMETHOD
 client option 173
 server option 162
 COMPRESSALWAYS, client option 274
 COMPRESSION option, UPDATE NODE command 274
 configuration view, Web administrative interface 167
 CONTACT option, DEFINE ADMIN command 143
 COPIED option, QUERY CONTENT command 119
 COPY CLOPTSET command 155

D

data storage policy, management of backup data 123
 database
 defining server volumes 54, 58
 extending capacity 70
 incremental and full backup 72
 managing backups 260
 MAXIMUMSIZE option 64
 mirroring volumes 66
 moving backups offsite 260
 reducing capacity 75
 removing volumes 70
 synchronization of server volumes 68
 triggering automatic backup 63
 triggering automatic expansion 64
 DATEFORMAT
 client option 163
 server option 50
 DEFINE ASSOCIATION command 225, 226
 DEFINE CLIENTOPT command 155
 DEFINE CLOPTSET command 155
 DEFINE COPYGROUP command 131, 132
 DEFINE DBBACKUPTRIGGER command 63
 DEFINE DBCOPY command 68
 DEFINE DBVOLUME command 53, 58
 DEFINE DOMAIN command 125
 DEFINE LOGCOPY command 69
 DEFINE LOGVOLUME command 53, 60
 DEFINE MGMTCLASS command 128
 DEFINE POLICYSET command 126
 DEFINE SCHEDULE command 225
 DEFINE SPACETRIGGER command 64
 DEFINE VOLUME command 53, 109
 DELETE CLIENTOPT command 155

DELETE CLOPTSET command 155
 DELETE DBVOLUME command 70
 DELETE DEVCLASS command 117
 DELETE DOMAIN command 138
 DELETE FILESPACE command 138, 151
 DELETE LIBRARY command 117
 DELETE LOGVOLUME command 71
 DELETE STGPOOL command 110, 116
 DELETE VOLHISTORY command 76, 220, 262
 DELETEDFILES option, ARCHIVE command 189
 DESCRIPTION option, ARCHIVE command 189
 DESTINATION option, DEFINE COPYGROUP command 129
 detail frame, Web administrative interface 167
 DEVCONFIG, server option 50
 device class
 8MM type 93
 defining device class 93
 deleting device class 117
 DISK 94
 FILE type 94
 logical component of storage device 91
 MOUNTRETENTION option 93
 devices
 physical and logical components 90
 random access 90
 sequential storage 90
 DIRMC, client option 127, 275
 DISABLE SESSIONS command 234
 DISMOUNT VOLUME command 111, 117
 domain
 See policy domain
 DOMAIN option, REGISTER NODE command 148
 drive
 defining for automatic library 101
 defining for manual library 92
 physical component of storage device 91
 DRM 46, 260
 dsm, backup-archive GUI program 193
 dsm.opt, client user options file 172
 dsm.sys, client system options file 172
 DSM_CONFIG, client variable 171
 DSM_DIR, client variable 171
 DSM_LOG, client variable 171
 dsmadm, administrative client program 163
 DSMAGENT, remote client agent 171, 211
 dsmsvc, backup-archive command line interface 176
 DSMCAD, client acceptor 171, 211
 DSMCUTIL, Windows NT client service program 207
 DSMFMT, server utility 53
 DSMLABEL, server utility 55
 dsmserv, server program 52, 230
 dsmserv.opt, server options file 48
 DSMTCA, client communication agent 171

E

ENABLE EVENTS command 238
 ENABLE3590LIBRARY, server option 50
 enterprise console
 see Web administrative interface

environment variables for client 171
 event monitoring 236
 event receiver 238
 event viewer, Web administrative interface 165, 167
 EXCLUDE, client option 174
 excluding files from backup processing 174
 EXPINTERVAL, server option 50, 142, 272
 EXPIRE INVENTORY command 141, 223, 272
 export-import function 46
 EXTEND DB command 70, 74
 EXTEND LOG command 71, 74

F

file state
 active 128, 184
 expired 128, 184
 inactive 128, 184
 FORMATSIZES, volume option 53, 58, 60
 FREQUENCY option, DEFINE COPYGROUP command 130

G

GRANT AUTHORITY command 143

H

HALT command 234
 HELPDESK, admin ID 145, 150
 HTTPPORT
 client option 175
 server option 162

I

INACTIVE option, RESTORE command 184
 INCLEXCL, client option 174
 INCLUDE, client option 127, 174
 include-exclude list 127, 172
 including excluded files to backup processing 174
 INCREMENTAL command 178
 Internet Explorer, supported levels 164, 210
 ITEMCOMMIT, admin parameter 53

J

Java applets 164
 Java requirements
 Web administrative interface 164
 Web client 210

L

LABEL LIBVOLUME command 55, 99, 115, 252
 LANGUAGE, client option 163
 LARGECOMMBUFFERS, client option 275
 LATEST option, RESTORE command 184
 library
 LIBTYPE option 92
 physical component of storage device 91
 steps to delete a library 115

- library, automatic
 - defining to Tivoli Storage Manager 101
- library, manual
 - defining to Tivoli Storage Manager 92
- LIBTYPE option, DEFINE LIBRARY command 92
- licensing
 - base server license
 - compliance with license terms 159, 160
 - enrollment certificate files 157
 - NOAUDITSTORAGE server option 160
 - saving in nodelock file 159
 - scheduling audit 224
- logical volume backup 178
- LOGPOOLSIZE, server option 50, 272
- loop mode of command line interface 164, 177
- lsdev, AIX command 93

M

- MACRO command 52
- macros
 - loading macro into Tivoli Storage Manager 4
 - mac.admins 144
 - mac.optionsets 154
 - mac.policy 133
 - mac.schedules 216, 220, 225
 - mac.scripts 219
 - mac.stgcreate 104
 - mac.stgdelete 110
- management class
 - ARCHMC option 127
 - binding backup data 127
 - defining a management class 126
 - DIRMC option 127
 - INCLUDE option 127
 - MIGDESTINATION option 135
- MAXIMUMSIZE option, DEFINE SPACETRIGGER command 64
- MAXSCHEDESSESSIONS 273
- MAXSCRATCH option, DEFINE STGPOOL command 95
- MAXSESSION, server option 273
- MIGDESTINATION option, DEFINE MGMTCLASS command 135
- migration between storage pools 118
- mirroring of volumes 66
- MIRRORWRITE, server option 50
- MODE option, DEFINE COPYGROUP command 130
- modified files during processing, serialization 131, 132
- MOUNTRETENTION option, DEFINE DEVICECLASS command 93
- MOVE DATA command 119
- MOVE MEDIA command 243
- MOVEBATCHSIZE, server option 50, 270
- MOVESIZETHRESH, server option 50, 270

N

- Netscape, supported levels 164, 210
- network view, Web administrative interface 167
- NOAUDITSTORAGE, server option 160
- nodelock, license file 159

- NODENAME, client option 173
- NUMBERFORMAT
 - client option 163
 - server option 50

O

- object view, Web administrative interface 167
- operation view, Web administrative interface 167

P

- PASSWORDACCESS, client option 174, 175
- passwords
 - automatic handling 174
 - Web client 175
- PATH, client variable 171
- PICK option, RESTORE command 183
- point-in-time restore 199
- policy domain
 - ARCHRET option 125
 - BACKRET option 125
 - defining a policy domain 125
 - enforcing policy 141
 - moving node to another domain 139
 - retention grace period 133
 - steps to remove default 138
- policy set
 - activating policy sets 126, 141
 - defining policy sets 126
 - validating policy sets 141
- PRESERVEPATH option
 - RESTORE command 184
 - RETRIEVE command 191
- problem determination 46, 263

Q

- QUERY ACTLOG command 245, 263
- QUERY ADMIN command 143, 146
- QUERY AUDITOCUPANCY command 224, 242
- QUERY CONTENT command 119
- QUERY COPYGROUP command 139
- QUERY DB command 239, 269
- QUERY DBBACKUPTRIGGER command 63
- QUERY DBVOLUME command 67, 240
- QUERY DEVCLASS command 73
- QUERY DRIVE command 242
- QUERY EVENT command 246
- QUERY FILESPACE command 151
- QUERY LIBRARY command 73
- QUERY LIBVOLUME command 243
- QUERY LICENSE command 159
- QUERY LOG command 240, 272
- QUERY LOGVOLUME command 68, 241
- QUERY MEDIA command 243
- QUERY NODE command 149
- QUERY OCCUPANCY command 120, 243
- QUERY OPTIONS command 48
- QUERY PROCESS command 112, 152, 236, 247
- QUERY REQUEST command 108, 114, 244

QUERY RESTORE command 236, 247
 QUERY SESSION command 248
 QUERY STATUS command 51, 62, 235, 237
 QUERY STGPOOL command 116, 244
 QUERY VOLHISTORY command 241, 261, 262
 QUERY VOLUME command 222, 244, 255, 259
 QUIESCE option, HALT command 234
 QUIET, client option 275
 QUIT command 192

R

RAID
 RECLAIM option
 UPDATE STGPOOL command 222
 reclamation
 scheduling 222
 recovery log
 defining server volumes 54, 59
 extending capacity 71
 mirroring volumes 66
 reducing capacity 76
 removing volumes 71
 ROLLFORWARD option 61
 saving records 61
 synchronization of volumes 69
 triggering automatic expansion 65
 redbook support material
 checklists
 See checklists 6
 creating storage pools 104
 defining administrator 144
 defining client option sets 154
 defining policy structure 133
 defining schedules 216, 220, 225
 defining server scripts for offsite backup 219
 deleting default storage pools 110
 downloading from redbooks Web server 3, 11, 279, 285
 loading macro into Tivoli Storage Manager 4
 planning worksheets
 See worksheets 11
 redbook Tivoli Storage Manager environment 2
 REDUCE DB command 75
 REDUCE LOG command 76
 redundant array of independent disks
 See RAID
 REGISTER ADMIN command 143
 REGISTER LICENSE command 157
 REGISTER NODE command 148, 149
 remote client agent, Web client 171, 211
 REMOVE ADMIN command 147
 REMOVE NODE command 138
 RENAME NODE command 149
 RENAME STGPOOL command 120
 REPLACE, client option 174
 REPLY command 115
 RESET LOGCONSUMPTION command 62
 restartable restore sessions 199
 RESTORE command 183
 retention grace period 133

RETEXTRA option, DEFINE COPYGROUP command 129
 RETONLY option, DEFINE COPYGROUP command 129
 RETRIEVE command 190
 RETVER option, DEFINE COPYGROUP command 132
 REVOKEREMOTEACCESS, client option
 See Web client
 ROLLFORWARD option, SET LOGMODE command 61

S

SCHEDMODE, client option 174
 SCHEDULE command 205
 scripts, server 46, 52, 218
 SEARCH option, QUERY ACTLOG command 263
 SELECT command 52, 121, 154
 SELECTIVE command 178
 SELFTUNEBUFPOOLSIZE 50
 SERIALIZATION option, DEFINE COPYGROUP command 131
 server
 active console mode 230
 background mode 232
 BUFPOOLSIZE option 50, 269
 COMMMETHOD option 162
 communication options 49
 database
 See database
 DATEFORMAT option 50
 DEVCONFIG option 50
 DSMFMT utility 53
 DSMLABEL utility 55
 dsmserv program 52, 230
 dsmserv.opt, options file 48
 ENABLE3590LIBRARY option 50
 event monitoring 236
 event receiver 238
 EXPINTERVAL option 50, 142, 272
 export-import function 46
 HTTPPORT option 162
 installation 48
 LANGUAGE option 50
 licensing
 See licensing
 loading scripts 52
 LOGPOOLSIZE option 50, 272
 MAXSESSION option 273
 MIRRORWRITE option 50
 MOVEBATCHSIZE option 50, 270
 MOVESIZETHRESH option 50, 270
 NOAUDITSTORAGE option 160
 NUMBERFORMAT option 50
 performance options 50
 recovery log
 See recovery log
 sample scripts 52
 scripts 218
 security related settings 51
 site dependent options 49
 starting the program 230
 stopping the program 233

- storage pool
 - See storage pool
- TCPWINDOWSIZE option 50, 274
- TIMEFORMAT option 50
- TXNGROUPMAX option 50
- USELARGEBUFFER option 273
- VOLUMEHISTORY option 50, 220
- volumes
 - See volumes
- SERVER_CONSOLE, admin ID 143, 230
- SET ACCOUNTING command 52
- SET ACTLOGRETENTION command 238
- SET AUTHENTICATION command 51
- SET INPWLENGTH command 51
- SET INVALIDPWLIMIT command 51
- SET LICENSEAUDITPERIOD command 160
- SET LOGMODE command 61
- SET PASSEXP command 51
- SET WEBAUTHTIMEOUT command 51
- SQL 46, 121
- STATUS option, UPDATE VOLUME command 257
- storage pool
 - checkout of tape volumes 114
 - collection of storage pool volumes 79
 - defining copy 106
 - defining disk volumes 100, 109
 - defining for device class of DISK 95
 - defining for manual library 95
 - defining primary 105
 - defining server volumes 54
 - defining volumes for automatic library 107
 - defining volumes for manual library 99
 - deleting a storage pool 116
 - labelling and checkin of tape volumes 107
 - logical component of storage device 91
 - MAXSCRATCH option 95
 - migration 118
 - moving data within storage pool 119
 - NEXTSTGPOOL option 95
 - primary volumes 80
 - RECLAIM option 222
 - RECLAIMSTGPOOL option 97
 - relabelling of tape volumes 115
 - removing default 110
 - scheduling migration 221
 - threshold parameter for migration 118
- structured query language
 - See SQL
- SUBDIR, client option 174

T

- TAPEPROMPT, client option 174
- tapes
 - See volumes
- TCPBUFFSIZE, client option 173, 275
- TCPPORT, client option 173
- TCPSERVERADDRESS, client option 173
- TCPSERVERNAME, client option 163
- TCPWINDOWSIZE
 - client option 173, 276

- server option 50, 274
- TIMEFORMAT
 - client option 163
 - server option 50
- Tivoli Storage Manager homepage
 - software fixes 47
- tree frame, Web administrative interface 167
- trustee directory assignments, NetWare
 - backup of directory information 110
- TSO, administrative interface on MVS 161
- TXNBYTELIMIT, client option 276
- TXNGROUPMAX, server option 50, 276
- TYPE option, DEFINE COPYGROUP command 129

U

- UPDATE ADMIN command 147
- UPDATE NODE command 139, 149, 154
- UPDATE SPACETRIGGER command 64, 65
- UPDATE STGPOOL command 221, 222
- UPDATE VOLUME command 257
- USELARGEBUFFER, server option 273
- USERID option, REGISTER NODE command 148

V

- VALIDATE POLICYSET command 141
- VERDELETED option, DEFINE COPYGROUP command 129
- VEREXISTS option, DEFINE COPYGROUP command 129
- VOLUMEHISTORY, server option 50, 220
- volumes
 - checking out from library 256
 - CHECKLABEL option 111
 - checkout of tapes 114
 - COPIED option 119
 - defining disk storage pool volumes 109
 - defining for automatic library 107
 - defining storage pool volumes for manual library 99
 - FORMATSIZE option 53, 58, 60
 - formatting server volumes 53
 - history
 - See administrative client
 - labelling and checkin of tapes 107
 - listing to send offsite 255
 - listing to send onsite 259
 - logical component of storage device 91
 - managing tapes between onsite and offsite 250
 - mirroring 66
 - mount request of tapes 108, 114
 - moving data within storage pool 119
 - moving tapes from offsite to onsite 258
 - moving tapes from onsite to offsite 254
 - reclamation, offsite 260
 - relabelling of tapes 115
 - scheduling reclamation 222
 - STATUS option 257

W

WAIT option, running processes in foreground 152

Web administrative interface

- accessing the Web 165
- banner frame 167
- browser requirements 164
- command line 165, 167
- COMMETHOD server option 162
- configuration view 167
- detail frame 167
- enterprise console 164
- event viewer 165, 167
- HTTPPORT server option 162
- Internet Explorer 164
- Java applets 164
- network view 167
- object view 167
- operation view 167
- software requirements 162
- tree frame 167
- WEBAUTHTIMEOUT option

Web client

- administrative authorities 211
- client acceptor 171, 211
- client access authority 211
- client functions 210
- client owner authority 211
- configuration 171
- DSMAGENT, remote client agent 171, 211
- DSMCAD, client acceptor 171, 211
- enterprise console 164
- HELPDESK admin ID 145
- HTTPPORT option 175
- Java applet 210
- PASSWORDACCESS option 175
- passwords 175
- remote client agent 171, 211
- REVOKEREMOTEACCESS client option 175
- userid 211

WEBAUTHTIMEOUT option

See Web administrative interface

wheel of live, recommended scheduling of operations 215

Windows NT configuration wizards 55

WINDOWS NT event viewer 265

worksheets

- administrator IDs 43
- client requirement 13
- database 31
- device configuration and volume history 36
- license requirements 43
- primary storage pool 35
- recovery log 33
- storage policy requirements 19
- tape drive configuration 37
- total Tivoli Storage Manager disk required 36

IBM Redbooks review

Your feedback is valued by the Redbook authors. In particular we are interested in situations where a Redbook "made the difference" in a task or problem you encountered. Using one of the following methods, **please review the Redbook, addressing value, subject matter, structure, depth and quality as appropriate.**

- Use the online **Contact us** review redbook form found at ibm.com/redbooks
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Document Number	SG24-5416-01
Redbook Title	Getting Started with Tivoli Storage Manager: Implementation Guide
Review	
What other subjects would you like to see IBM Redbooks address?	
Please rate your overall satisfaction:	<input type="radio"/> Very Good <input type="radio"/> Good <input type="radio"/> Average <input type="radio"/> Poor
Please identify yourself as belonging to one of the following groups:	<input type="radio"/> Customer <input type="radio"/> Business Partner <input type="radio"/> Solution Developer <input type="radio"/> IBM, Lotus or Tivoli Employee <input type="radio"/> None of the above
Your email address: The data you provide here may be used to provide you with information from IBM or our business partners about our products, services or activities.	<input type="checkbox"/> Please do not use the information collected here for future marketing or promotional contacts or other communications beyond the scope of this transaction.
Questions about IBM's privacy policy?	The following link explains how we protect your personal information. ibm.com/privacy/yourprivacy/



Redbooks

Getting Started with Tivoli Storage Manager: Implementation Guide

(0.5" spine)
0.5" <-> 0.875"
250 <-> 459 pages



Getting Started with Tivoli Storage Manager:

Implementation Guide



Redbooks

How to plan and set up your Tivoli Storage Manager server from scratch

Includes worksheets, scripts, and macros to make your job easier

Tips for both new and experienced administrators

This redbook provides material describing how to implement and operate Tivoli Storage Manager for people who have a general understanding of Tivoli Storage Manager but have no practical experience of Tivoli Storage Manager implementations.

Support materials in the form of example option files and macros are included where appropriate, and checklists lead you step-by-step through an implementation of the basic functions of Tivoli Storage Manager:

- Planning and sizing
- Product implementation
- Customization
- Operation
- Performance and tuning

This practical guide is intended for the following audience:

- New Business Partners and other implementers aiming for Tivoli Storage Manager Certification.
- System administrators, new to Tivoli Storage Manager, who are asked to commence a basic Tivoli Storage Manager implementation for the very first time
- Tivoli Storage Manager administrators who want to learn more about the basic Tivoli Storage Manager components and their implementation

**INTERNATIONAL
TECHNICAL
SUPPORT
ORGANIZATION**

**BUILDING TECHNICAL
INFORMATION BASED ON
PRACTICAL EXPERIENCE**

IBM Redbooks are developed by IBM's 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**

SG24-5416-01

ISBN 0-7384-1907-9