

Disaster Recovery Strategies with Tivoli Storage Management

Keeping your TSM server and clients
safe from disaster

How and why to build a disaster
recovery plan

Testing your disaster
recovery plan



Charlotte Brooks
Matthew Bedernjak
Igor Juran
John Merryman



International Technical Support Organization

**Disaster Recovery Strategies with Tivoli Storage
Management**

November 2002

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

Second Edition (November 2002)

This edition applies to IBM Tivoli Storage Manager Version 5, Release 1.

© Copyright International Business Machines Corporation 2002. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	xi
Tables	xv
Examples	xvii
Notices	xxi
Trademarks	xxii
Preface	xxiii
The team that wrote this redbook	xxiii
Become a published author	xxv
Comments welcome	xxvi
Part 1. Disaster Recovery Planning	1
Chapter 1. Introduction to TSM and Disaster Recovery	3
1.1 Overview of this book	4
1.2 Objective: Business Continuity	5
1.3 Disaster Recovery Planning terms and definitions	9
1.4 IBM Tivoli Storage Manager overview	11
1.4.1 TSM platform support	11
1.4.2 How TSM works	13
1.4.3 TSM Server administration	14
1.4.4 TSM Backup/Archive Client	15
1.4.5 TSM backup and archive concepts	17
1.4.6 Tivoli Data Protection for Applications modules	17
1.4.7 TSM security concepts	18
1.5 TSM and Disaster Recovery	18
Chapter 2. The tiers of Disaster Recovery and TSM	21
2.1 Seven tiers of Disaster Recovery	22
2.1.1 Tier 0 - Do nothing, no off-site data	23
2.1.2 Tier 1 - Offsite vaulting (PTAM)	23
2.1.3 Tier 2 - Offsite vaulting with a hot site (PTAM + hot site)	24
2.1.4 Tier 3 - Electronic vaulting	25
2.1.5 Tier 4 - Electronic vaulting to hot site (active secondary site)	26
2.1.6 Tier 5 - Two-site, two-phase commit	27
2.1.7 Tier 6 - Zero data loss	28

2.2	Costs related to DR solutions	29
2.3	Strategies to achieve tiers of DR using TSM.	31
Chapter 3. Disaster Recovery Planning		37
3.1	Demand for Comprehensive DR Planning	38
3.1.1	Business requirements for data recovery and availability	38
3.1.2	Increasing availability requirements	39
3.1.3	Storage growth trends.	39
3.1.4	Storage management challenges	39
3.1.5	Networking growth and availability	41
3.1.6	Capacity planning trends.	43
3.2	Business Continuity, Disaster Recovery, and TSM.	43
3.2.1	Business Continuity Planning	44
3.2.2	Disaster Recovery Planning	44
3.3	Disaster Recovery Planning and TSM planning	48
3.3.1	How BIA and RTO relate to TSM policy	48
3.3.2	Dependencies between systems	49
3.3.3	Data classification procedures	50
3.3.4	Relating DR Planning to TSM planning.	52
3.4	DRP development	54
3.4.1	Background information	55
3.4.2	Notification/activation phase	56
3.4.3	Recovery phase	58
3.4.4	Reconstitution phase.	59
3.4.5	Design appendices	60
3.5	Related planning considerations	60
3.5.1	Insourcing versus outsourcing	60
3.5.2	Offsite tape vaulting	62
Chapter 4. Disaster Recovery Plan testing and maintenance.		63
4.1	How to provide Disaster Recovery testing	64
4.1.1	How is DR testing different from a real disaster	65
4.1.2	Partial activity testing or overall plan testing	67
4.1.3	Timeframe for DRP testing	67
4.1.4	Example of a simple testing plan scenario	68
4.1.5	Example of a more complex test scenario	69
4.2	When to provide Disaster Recovery testing	71
4.2.1	How to select the best time for testing	71
4.3	DRP maintenance procedures	71
4.3.1	Maintenance policy	73
4.3.2	Reference information.	76
Chapter 5. Planning for data center availability		79
5.1	Infrastructure, the first line of defense	80

5.2	Data center planning	81
5.2.1	Space and growth considerations	81
5.2.2	Fundamental elements of data center design	81
5.2.3	Preventative controls	82
5.2.4	Onsite parts maintenance locker	83
5.2.5	Data center facilities security	83
5.2.6	Disaster Recovery Planning and infrastructure assessment	83
5.3	Distributed systems architecture considerations	84
5.3.1	Server architecture	84
5.3.2	Server operating system clustering technologies	84
5.3.3	Capacity planning	85
5.3.4	System software	85
5.3.5	Configuration management	85
5.3.6	Problem and change management	85
5.4	Disk storage architecture considerations	86
5.4.1	Advanced storage subsystem functionality	87
5.5	Network architecture considerations	87
5.5.1	Network architecture	87
5.5.2	Network documentation	89
5.5.3	Network management	89
5.5.4	Remote access considerations	90
5.5.5	WAN considerations	90
5.5.6	Network security considerations	90
5.6	Hotsite, warmsite, coldsite considerations	90
5.6.1	Alternate site cost analysis	93
5.6.2	Partnering strategies for alternate sites	93
	Chapter 6. Disaster Recovery and TSM	95
6.1	Mapping general concepts to TSM	96
6.1.1	TSM policy drivers	97
6.1.2	TSM software architecture review	98
6.1.3	TSM policy configuration overview	99
6.2	Mapping DR requirements to TSM configurations	104
6.2.1	A functional example	104
6.3	Creating TSM disk pools for recovery objectives	105
6.3.1	Disk storage pool sizing	107
6.4	TSM tape storage pools	107
6.4.1	Tape device performance considerations	107
6.4.2	Tape library sizing fundamentals	109
6.4.3	Versioning and tape storage pool sizing	110
6.5	LAN and SAN design	111
6.5.1	LAN considerations	111
6.5.2	SAN considerations	112

6.5.3	TSM server LAN, WAN, and SAN connections	112
6.5.4	Remote device connections	112
6.6	TSM server and database policy considerations	113
6.7	Recovery scheduling	113
Chapter 7.	TSM tools and building blocks for Disaster Recovery	115
7.1	Introduction	116
7.2	The TSM server, database, and log	117
7.2.1	TSM database page shadowing	118
7.2.2	TSM database and recovery log mirroring	119
7.3	TSM storage pools	120
7.4	Requirements for TSM server Disaster Recovery	121
7.5	TSM backup methods and supported topologies	124
7.5.1	Client backup and restore operations	124
7.5.2	Traditional LAN and WAN backup topology	127
7.5.3	SAN (LAN-free) backup topology	128
7.5.4	Server-free backup	129
7.5.5	Split-mirror/point-in-time copy backup using SAN	130
7.5.6	NAS backup and restore	131
7.5.7	Image backup	133
7.6	TSM Disaster Recovery Manager (DRM)	134
7.6.1	DRM and TSM clients	137
7.7	TSM server-to-server communications	138
7.7.1	Server-to-server communication	139
7.7.2	Server-to-server virtual volumes	140
7.7.3	Using server-to-server virtual volumes for Disaster Recovery	142
7.7.4	Considerations for server-to-server virtual volumes	143
7.8	TSM and high availability clustering	144
7.8.1	TSM and High Availability Cluster Multi-Processing (HACMP)	144
7.8.2	TSM backup-archive and HSM client support with HACMP	146
7.8.3	TSM and Microsoft Cluster Server (MSCS)	148
7.8.4	TSM backup-archive client support with MSCS	150
7.9	TSM and remote disk replication	151
7.10	TSM and tape vaulting	153
7.10.1	Electronic tape vaulting	153
7.11	Remote disk mirroring and tape vaulting solutions	155
7.11.1	Collocation considerations for offsite vaulting	156
7.11.2	Reclamation considerations for offsite vaulting	157
Part 2.	Implementation procedures and strategies	161
Chapter 8.	IBM Tivoli Storage Manager and DRM	163
8.1	Introduction	164
8.2	Best practices for hardware design	164

8.2.1	TSM server platform selection from DR point of view	164
8.2.2	Disk storage devices	168
8.2.3	Tape storage devices	169
8.2.4	Storage Area Network (SAN)	171
8.2.5	Common mistakes in backup/recovery hardware planning.	172
8.3	Best practices for TSM Server implementation	173
8.3.1	TSM database and recovery log	173
8.3.2	TSM server database management	175
8.3.3	TSM storage pools	178
8.4	Best practices for environment management	180
8.4.1	Cold standby	181
8.4.2	Warm standby	181
8.4.3	Hot standby	183
8.4.4	Configuration considerations.	186
8.5	Proactive hints/tips for management availability	187
8.5.1	Failover	188
8.5.2	Geographic failover.	189
8.6	DRM — what it is and what it is not.	190
8.7	Example of DRM execution.	192
8.7.1	DRM setup	193
8.7.2	Daily operations	199
8.7.3	Server restore setup	207
Chapter 9. Overview of bare metal recovery techniques.		223
9.1	Bare metal recovery steps.	225
9.1.1	Manual bootstrap BMR	225
9.1.2	Operating system image bootstrap BMR	226
9.1.3	Live media bootstrap BMR	227
9.1.4	Automated BMR suite	228
Chapter 10. Solaris client bare metal recovery.		229
10.1	Sun Solaris Version 7 recovery	230
10.1.1	Product Overview	230
10.1.2	Backup	230
10.2	Pre disaster preparation	231
10.2.1	TSM backups	231
10.2.2	Operating environment of the Solaris client	231
10.2.3	Creating ufsdump backup	233
10.2.4	Restore root file system using ufsrestore command.	235
Chapter 11. AIX client bare metal recovery.		239
11.1	AIX system backup overview	240
11.1.1	AIX root volume group overview	240
11.2	Using the mkysb for bare metal restore.	241

11.2.1	Exclude files for mksysb backup	242
11.2.2	Saving additional volume group definitions	243
11.2.3	Classic mksysb to tape	245
11.2.4	Mksysb to CD-ROM or DVD	246
11.2.5	The use of TSM in addition to mksysb procedures	247
11.2.6	Bare metal restore using mksysb media	248
11.3	Using NIM for bare metal restore	248
11.3.1	Disaster Recovery using NIM and TSM	249
11.3.2	Basic NIM Setup	250
11.3.3	AIX bare metal restore using NIM	261
11.3.4	NIM administration	263
11.4	SysBack overview	264
11.4.1	Network Boot installations using SysBack	264
Chapter 12. Windows 2000 client bare metal recovery		267
12.1	Windows 2000 client bare metal restore using TSM	268
12.1.1	Collecting client machine information for Disaster Recovery	268
12.1.2	Collect partition and logical volume information with diskmap	271
12.1.3	Storing system information for DRM access	271
12.1.4	Inserting client machine information into DRM	272
12.1.5	Using machchar.vbs to insert machine reports into DRM	273
12.1.6	Backing up Windows 2000 with TSM for BMR	276
12.1.7	Basic Windows 2000 BMR procedure with TSM	279
12.2	Using third party products for BMR with TSM	285
12.2.1	Cristie	285
12.2.2	Ultrabac	286
12.2.3	VERITAS BMR	286
Chapter 13. Linux client bare metal recovery		289
13.1	Linux Red Hat bare metal recovery and TSM	290
13.1.1	Collecting client machine information for Disaster Recovery	290
13.1.2	Saving client machine information for DRM	294
13.1.3	Backing up your Linux system	295
13.1.4	Preparing a BMR bootable diskette	296
13.2	Basic Linux BMR procedure with TSM	297
Chapter 14. Putting it all together		307
14.1	TSM implementation scenarios	308
14.1.1	Storage networking considerations	308
14.1.2	Remote disk copy and mirroring considerations	310
14.1.3	Tier 6 mirrored site implementation with TSM	311
14.1.4	Tier 5 hot site with TSM electronic vaulting	313
14.1.5	Tier 5 hot site with clustered TSM and electronic vaulting	314
14.1.6	Tier 4 warm site TSM implementation with electronic vaulting	315

14.1.7 Tier 3 warmsite TSM implementation	317
14.1.8 Dual production sites using TSM	318
14.1.9 Remote TSM server at warmsite, with manual offsite vaulting	319
14.2 Case studies	320
14.2.1 Case study 1 - background	320
14.2.2 Case study 2 - background	324
14.3 TSM and DR from beginning to end - summary	325
14.4 Application and database backup Redbooks	326
Part 3. Appendices	329
Appendix A. DR and Business Impact Analysis Planning Templates	331
Disaster Recovery Plan Template	332
A.1 Introduction	333
A.1.1 Purpose	333
A.1.2 Applicability	333
A.1.3 Scope	334
A.1.4 References/requirements	335
A.1.5 Record of changes	336
A.2 Concept of operations	336
A.2.1 System description and architecture	336
A.2.2 Line of succession	337
A.2.3 Responsibilities	337
A.3 Notification and Activation phase	337
A.3.1 Damage assessment procedures	338
A.3.2 Alternate assessment procedures	338
A.3.3 Criteria for activating the DR plan	339
A.4 Recovery operations	339
A.5 Return to normal operations	340
A.5.1 Concurrent processing	341
A.5.2 Plan deactivation	341
A.6 Plan appendices	341
Business Impact Analysis Template	343
Appendix B. Windows BMR configuration scripts	347
Insert machine characteristics into DRM	348
Reducing msinfo32 output using a VBScript	351
Appendix C. DRM plan output	353
Appendix D. Additional material	387
Locating the Web material	387
Using the Web material	387
How to use the Web material	388

Abbreviations and acronyms	389
Related publications	391
IBM Redbooks	391
Other resources	392
Referenced Web sites	392
How to get IBM Redbooks	393
IBM Redbooks collections	393
Index	395

Figures

1-1	Planned versus unplanned outages for IT operations	5
1-2	Types of disasters	6
1-3	Potential cost factors for downtime	7
1-4	Average cost of downtime for various US industries	8
1-5	TSM client, server, application, and database support	12
1-6	Data movement with TSM and the TSM storage hierarchy	13
1-7	Sample TSM administration screenshot	15
1-8	The TSM Version 5 client interface	16
2-1	Summary of Disaster Recovery tiers (SHARE)	22
2-2	Tier 0 - Do nothing, no offsite data.	23
2-3	Tier 1 - Offsite vaulting (PTAM)	24
2-4	Tier 2 - Offsite vaulting with a hot site (PTAM + hot site).	25
2-5	Tier 3 - Offsite electronic vaulting	26
2-6	Tier 4 - Electronic vaulting with hot site (active secondary site)	27
2-7	Tier 5 - Two-site, two-phase commit	28
2-8	Tier 6 - Zero data loss (advanced coupled systems)	29
2-9	Seven tiers of Disaster Recovery solutions	30
3-1	Networking technologies and transfer rates.	42
3-2	Example of the Business Impact Analysis process	46
3-3	Disaster Recovery Planning process	48
3-4	The relationship between BIA, RTO, and TSM planning.	49
3-5	Basic example of enterprise data dependencies	50
3-6	Sample Business Impact Analysis Worksheet.	53
3-7	Disaster Recovery Planning processes and plan content	55
5-1	Example of power systems design for availability	82
5-2	Redundant network architecture	88
5-3	Storage Area Network architecture for data availability	89
5-4	Cost versus recovery time for various alternate site architectures.	91
6-1	Planning, TSM policy creation, and infrastructure design	97
6-2	TSM data movement and policy subjects	99
6-3	TSM policy elements	100
6-4	Relation of RTO and RPO requirements to management classes	101
6-5	Policy, storage pools, device classes, and devices	103
6-6	Management class relation to disk storage pools for critical data	106
6-7	Relating RTO, management classes and device classes	108
6-8	File versions, management class grouping, and storage pools	110
6-9	Sample recovery schedule using resource scheduling	114
7-1	TSM components.	116

7-2	TSM server, database, and recovery log overview	117
7-3	TSM database page shadowing	118
7-4	TSM database and recovery log mirroring.	119
7-5	TSM storage pools.	120
7-6	TSM server Disaster Recovery requirements	122
7-7	TSM LAN and WAN backup	128
7-8	TSM LAN-free backup	129
7-9	TSM Server-free backup	130
7-10	TSM split-mirror/point-in-time copy backup	131
7-11	IBM TotalStorage NAS backup	132
7-12	TSM and NDMP backup	133
7-13	TSM image backup	134
7-14	Tivoli DRM functions	136
7-15	TSM server-to-server communications	139
7-16	Server-to-server virtual volumes	141
7-17	HACMP and TSM server high availability configuration	145
7-18	HACMP and TSM client high availability configuration	147
7-19	MSCS and TSM server high availability configuration	149
7-20	MSCS and TSM client high availability configuration.	151
7-21	TSM and remote disk replication	152
7-22	TSM and electronic tape vaulting.	154
8-1	Available TSM server platforms	165
8-2	TSM database with two database volumes and triple volume mirroring	174
8-3	TSM recovery log volume, triple mirrored	174
8-4	Type of TSM database backup	176
8-5	Define disk storage pool with migration delay	179
8-6	Hot standby TSM server and TSM server	183
8-7	Hot standby TSM server	185
8-8	Hot instance on TSM server	186
8-9	DRM lab setup	192
8-10	DRM media states and lifecycle.	199
8-11	Daily operations - primary pools backup and TSM database backup .	205
8-12	Disaster Recovery Plan generation	206
8-13	Generating instruction, command, and cmd input files from DR plan .	208
8-14	TSM server restore	209
9-1	Manual bootstrap BMR	225
9-2	Operating system image bootstrap BMR.	226
9-3	Live media bootstrap BMR.	227
11-1	Typical AIX file system for system data.	241
11-2	Layout of a mksysb tape	242
11-3	AIX client dataflow with NIM and TSM.	249
11-4	Bare metal restore of AIX clients using NIM and TSM.	250
12-1	Using DRM INSERT MACHINE via TSM Web-based interface.	273

12-2	DRM machine information output	274
12-3	Backup Selection of Local Domain and System Objects	279
12-4	Windows 2000 System Partition Restore Options	283
13-1	Linux full system TSM backup	296
14-1	Storage network extension technologies	309
14-2	Tier 6 TSM implementation, mirrored production sites	312
14-3	Tier 5 hotsite with TSM electronic vaulting	313
14-4	Clustered TSM Servers, with site to site electronic vaulting	314
14-5	Asynchronous TSM DB mirroring and copy storage pool mirroring . . .	316
14-6	Electronic vaulting of TSM DRM, manual DB, and copy pool vaulting	317
14-7	Dual production sites, electronic TSM DB and copy pool vaulting. . . .	318
14-8	Remote TSM Server at warm site, plus manual offsite vaulting	319
14-9	Starting situation at XYZ Bank	321
14-10	TSM data mirrored on backup location	322
14-11	Recommended solution for Bank - final situation.	323
14-12	Summary of TSM and Disaster Recovery considerations	326

Tables

2-1	TSM strategies to achieve Disaster Recovery tiers	32
3-1	Data classification descriptions	51
4-1	Differences between DR testing and real disaster for Tiers 1-3	65
4-2	Differences between DR testing and real disaster for Tiers 4-6	66
4-3	Materials to restore NT server	68
4-4	Preparation steps	69
4-5	Disaster Recovery testing scenario	70
4-6	Record of DR executed steps	70
4-7	Approval Board members list	73
4-8	Document Owner	73
4-9	Task Owners	74
4-10	Plan of audits	76
4-11	Change history	76
4-12	Release protocol	77
4-13	Distribution list	77
5-1	Alternate site decision criteria	91
5-2	Alternate site cost analysis template	93
6-1	The basic tape drive aggregate performance equation	108
6-2	Network topologies and relative data transfer rates	111
7-1	Summary of client backup and restore operations	125
7-2	Extended electronic vaulting technologies	155
8-1	TSM server platform consideration	166
8-2	Classes of tape drive by CNT	171
8-3	Comparison of point-in-time and roll-forward recovery	177
8-4	DRM principal features	191
8-5	Review of the TSM device configuration	213
10-1	Commands for copying files and files systems	230
11-1	Mksysb, NIM, and SysBack characteristics	240
12-1	Basic Windows 2000 BMR Procedure with TSM	280
13-1	Linux system directories backed up as tar files	295
13-2	Basic Linux Red Hat BMR Procedure with TSM	297
A-1	Record of changes	336
A-2	Preliminary system information	343
A-3	Identify system resources	344
A-4	Identify critical roles	344
A-5	Link critical roles to critical resources	344
A-6	Identify outage impacts and allowable outage times	345
A-7	Prioritize resource recovery	345

Examples

7-1	DMR Plan File Example Stanza	137
8-1	Register DRM license	193
8-2	Define copy storage pool	194
8-3	Query storage pool	194
8-4	Define DRM plan prefix	195
8-5	Set DRM instruction prefix	196
8-6	Set DRM replacement volume names postfix	196
8-7	Command SET DRMCHECKLABEL	196
8-8	Specify primary storage pools to be managed by DRM.	196
8-9	Specify primary storage pools to be managed by DRM.	197
8-10	Set the DRM couriername	197
8-11	Set the database backup series expiration value	197
8-12	Set whether to process FILE device class media	197
8-13	Set NOTMOUNTABLE location name	198
8-14	Set expiration period for recovery plan files.	198
8-15	Set DRM vault name	198
8-16	Command QUERY DRMSTATUS	198
8-17	Backup of primary disk storage pool	200
8-18	Backup of primary tape storage pool	200
8-19	TSM database backup.	201
8-20	Query the DR media	202
8-21	Check and eject DR media	202
8-22	Eject DR media from the library	203
8-23	Sign over the DR media to the courier.	204
8-24	Confirm DR media arrives at vault	204
8-25	Display offsite media status	204
8-26	Prepare the recovery plan	206
8-27	Break out DRP on Windows platforms.	209
8-28	Break out DRP on AIX platform	209
8-29	Break out DRP on Sun Solaris platform.	209
8-30	Generating instruction, command and cmd files	210
8-31	Volumes required for TSM server recovery	211
8-32	File PRIMARY.VOLUMES.REPLACEMENT.CREATE.CMD.	212
8-33	File PRIMARY.VOLUMES.REPLACEMENT.MAC	213
8-34	DEVICE.CONFIGURATION.FILE	214
8-35	Invoke RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE	215
8-36	The start of TSM server after database recovery	216
8-37	Commands UPDATE PATH	218

8-38	Invoke RECOVERY.SCRIPT.NORMAL.MODE	219
10-1	Start system in single user mode	233
10-2	File-system check command fsck	233
10-3	Back up of root file system	234
10-4	Boot from CD to single user mode	235
10-5	Format partition using newfs command	235
10-6	Checking partition using fsck command	236
10-7	Mount partition for restore	236
10-8	Restore root partition using ufsrestore	236
10-9	Remove restoresymtable	237
10-10	Umount file system	237
10-11	Reinstallation of boot sector	237
11-1	Sample /etc/exclude.rootvg file	243
11-2	The vg_recovery script	243
11-3	Restoring a volume group using the restvg command	244
11-4	Mksysb interface on AIX	245
11-5	Sample dsm.opt file exclude statements	247
11-6	Basic NIM Master setup	251
11-7	NIM Client configuration	253
11-8	List standalone NIM clients on NIM master	253
11-9	Client mksysb backup definition to the NIM master	254
11-10	Defining the Exclude_files resource on the NIM Master	255
11-11	Sample view of NIM machine characteristics	256
11-12	Sample etc/bootptab NIM client entry on NIM master	257
11-13	Sample client.info file output	257
11-14	Sample screen for installing the BOS for Standalone clients	258
11-15	Sample output for NIM client verification	260
12-1	How to run msinfo32	269
12-2	Example of msinfo32 report output	269
12-3	Batch file for saving machine information and starting TSM client	270
12-4	Example diskmap output	271
12-5	Setting client system information report access in TSM	272
12-6	Using TSM command line to insert machine information	273
12-7	Running machchar.vbs to start machine information collection	273
12-8	Running the TSM macro to insert machine date into DRM	274
12-9	Querying DRM for Client Machine Information	274
12-10	Incorporating machine information in the DRM Plan file	275
12-11	Machine information seen in DRM Plan file	275
12-12	Include-exclude modification to backup system objects	278
12-13	System Object Verification using QUERY SYSTEMOBJECT	279
13-1	Using /proc/version to view operating system information	291
13-2	Using /proc/pci to view PCI device information	291
13-3	Using fdisk to output system partition information	292

13-4	viewing mounted file systems in /etc/fstab	293
13-5	Using df to output filesystem size, percent used, and mount points . .	293
13-6	Using ifconfig to display network card settings	294
13-7	Viewing system network settings in the /etc/sysconfig/network file . .	294
13-8	Using the tar command to package system directories	295
13-9	Creating a bootable rescue diskette using tomsrtbt	297
13-10	Deleting the Linux partition table	298
13-11	tomsrtbt root shell login	299
13-12	Floppy drive mount, driver install, zip drive mount commands	300
13-13	Using fdisk to repartition the root disk to pre-disaster state	300
13-14	Creating a file system and /root directory	303
13-15	Restoring our system files to the /root directory	303
13-16	Replacing the '/' directory and restoring the lilo boot block	304
13-17	Performing a command line install of a TSM Backup/Archive Client . .	304
13-18	Performing a command line restore of files systems and data	305
B-1	Example of VBScript command to insert machine characteristics	348
B-2	VBScript to Reduce Output from MSINFO32 Report	351
C-1	Example of DRP generated by TSM server	353

Notices

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

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Trademarks

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

AIX®	IBM eServer™	RMF™
AIX 5L™	IMS™	S/390®
DB2®	Informix®	SANergy™
DFS™	NetView®	Sequent®
DPI®	OS/2®	SP™
e Strategy™	OS/390®	Tivoli®
Enterprise Storage Server™	OS/400®	TotalStorage™
ESCON®	Perform™	Wave®
FICON™	pSeries™	WebSphere®
FlashCopy®	Redbooks™	xSeries™
IBM®	Redbooks(logo)™ 	z/OS™

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

Domino™	Lotus Notes®
Lotus®	Notes®

The following terms are trademarks of other companies:

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

C-bus is a trademark of Corollary, Inc. in the United States, other countries, or both.

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

SET, SET Secure Electronic Transaction, 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.

Preface

Disasters, by their very nature, cannot be predicted, in either their intensity, timing, or effects. However, all enterprises can and should prepare for whatever might happen in order to protect themselves against loss of data or, worse, their entire business. It is too late to start preparing after a disaster occurs. This IBM Redbook will help you protect against a disaster — taking you step by step through the planning stages, with templates for sample documents. It explores the role that IBM Tivoli Storage Manager plays in disaster protection and recovery, from both the client and server side. Plus, it describes basic sample procedures for bare metal recovery of some popular operating systems, such as Windows 2000, AIX, Solaris, and Linux.

This redbook is organized into two parts. Part 1 presents the general Disaster Recovery Planning process. It shows the relationship (and close interconnection) of Business Continuity Planning with Disaster Recovery Planning. It also describes how you might set up a Disaster Recovery Plan test. Various general techniques and strategies for protecting your enterprise are presented. Part 2 focuses on the practical, such as how to use IBM Tivoli Disaster Recovery Manager to create an auditable and easily executed recovery plan for your Tivoli Storage Manager server. It also shows approaches for bare metal recovery on different client systems.

This book is written for any computing professional who is concerned about protecting their data and enterprise from disaster. It assumes you have basic knowledge of storage technologies and products, in particular, IBM Tivoli Storage Manager.

The team that wrote this redbook

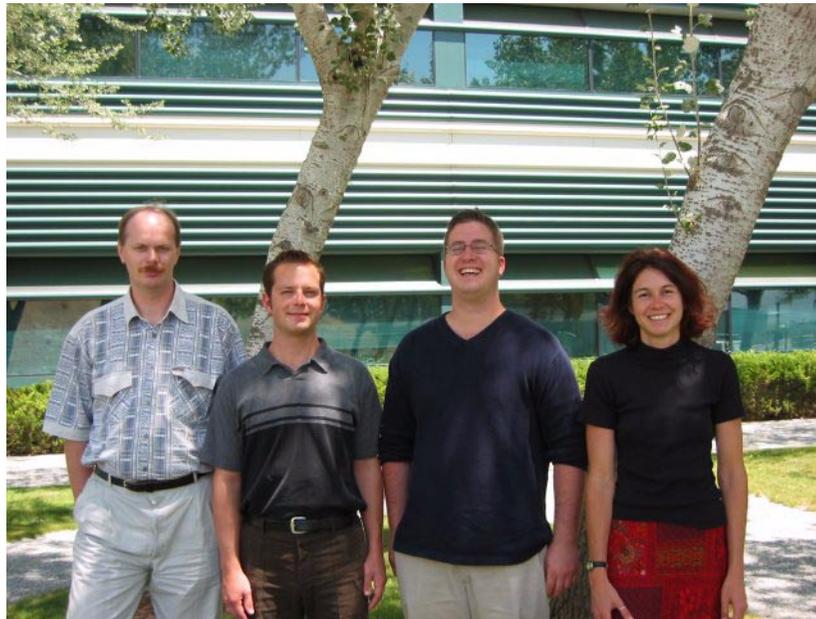
This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

Charlotte Brooks is a Project Leader for Tivoli Storage Management and Open Tape Solutions at the International Technical Support Organization, San Jose Center. She has 12 years of experience with IBM in the fields of RISC System/6000 and Storage. She has written eight redbooks, and has developed and taught IBM classes on all areas of storage management. Before joining the ITSO in 2000, she was the Technical Support Manager for Tivoli Storage Manager in the Asia Pacific Region.

Matthew Bedernjak is an IT Specialist in Toronto, Canada. He has over 4 years of experience in UNIX (IBM @server pSeries Server and AIX) and IBM Storage in Technical Sales Support. He holds a bachelor's degree in Civil Engineering from the University of Toronto, and is currently completing a master's degree in Mechanical and Industrial Engineering. His areas of expertise include AIX and pSeries systems, tape storage systems, Storage Area Networks, and disaster recovery. He has written extensively on disaster recovery planning and architecture.

Igor Juran is an IT Specialist at the IBM International Storage Support Centre/CEE located in Bratislava, Slovakia. He has 16 years experience in the IT field. He holds a degree in Computer Engineering from Slovak Technical University. His areas of expertise include Tivoli Storage Manager implementation on Windows and UNIX platforms, project management, storage planning and hardware configuration, and the development of disaster recovery plans.

John Merryman is a Data Storage Management Team Lead for Partners HealthCare System in Boston. He holds a bachelor's degree in Geology from the University of North Carolina at Chapel Hill and has 7 years of experience in the Enterprise Storage industry. His areas of expertise include UNIX and enterprise storage architecture, Storage Area Networking, storage management, project management, and disaster recovery planning. He has written extensively on disaster recovery planning and emerging technologies in the storage industry.



The team: (left to right) Igor, John, Matt, Charlotte

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

Jon Tate
International Technical Support Organization, San Jose Center

Emma Jacobs, Yvonne Lyon
International Technical Support Organization, San Jose Center

Mike Dile, Jim Smith
IBM Tivoli Storage Manager Development, San Jose

Dan Thompson
IBM Tivoli, Dallas

Jeff Barckley
IBM Software Group, San Jose

Tony Rynan
IBM Global Services, Australia

Ernie Swanson
Cisco Systems

Karl Evert, Jonathan Jeans, Terry Louis
CNT

Nate White
UltraBac Software

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

- ▶ Use the online **Contact us** review redbook form found at:

ibm.com/redbooks

- ▶ Send your comments in an Internet note to:

redbook@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. QXXE Building 80-E2
650 Harry Road
San Jose, California 95120-6099



Part 1

Disaster Recovery Planning

In this part we overview the Disaster Recovery Planning process, including basic definitions, the SHARE Disaster Recovery tiers, and an introduction to Business Impact Analysis and Business Continuity Planning. We describe how to test and maintain a Disaster Recovery Plan, and factors to consider when setting up a data center to maximize availability. Finally, we concentrate on IBM Tivoli Storage Manager, how to relate the planning process to specific product capabilities, and some of the basic Tivoli Storage Manager tools for protecting the client and server.



Introduction to TSM and Disaster Recovery

This chapter provides an introduction to this redbook. It discusses the purpose and intended audience, common terms and definitions, and structure of the following chapters. As a basis for a disaster recovery discussion, we introduce key considerations for business continuity planning. We introduce IBM Tivoli Storage Manager (TSM) and specific features and strategies TSM offers for disaster recovery.

Terminology convention: Throughout this redbook we refer to Disaster Recovery, Disaster Recovery Planning, which is the planning process, and also to the Disaster Recovery Plan, which is the plan to be used. For brevity, we will shorten these terms in this manner:

- ▶ Disaster Recovery (DR)
- ▶ Disaster Recovery Planning (DR Planning)
- ▶ Disaster Recovery Plan (DRP)

1.1 Overview of this book

The purpose of this book is to provide a comprehensive guide for developing and executing disaster recovery strategies using Tivoli Storage Management (TSM). We hope that our readers will be able to use this information to gain insights and develop specific strategies for Disaster Recovery in their own business environment.

This book covers a wide range of topics, from high level business disaster recovery planning, to specific TSM Disaster Recovery functions and to operating system bare metal restore. Therefore, we believe this book can provide value to those in a wide range of roles, including: customers developing DR strategies, IT Managers, IT Specialists, TSM Administrators, Disaster Recovery Consultants, Sales Specialists, and other related professionals.

This book is organized into the following parts:

Part 1 - Disaster Recovery Planning and TSM Building Blocks

Business Continuity Planning and TSM DR Overview, Tiers of Disaster Recovery and TSM, Disaster Recovery Planning, Disaster Recover Plan Testing and Maintenance, Planning for Data Center Availability, Disaster Recovery and TSM, and TSM Tools and Building Blocks for Disaster Recovery.

Part 2 - Implementation Procedures and Strategies

IBM Tivoli Storage Manager Server and Disaster Recovery Manager (DRM), Windows 2000 Bare Metal Recovery, Linux Redhat Bare Metal Recovery, AIX Bare Metal Recovery, and Sun Solaris Bare Metal Recovery, Typical Implementation Scenarios, Cases Studies, TSM and DR Summary, and Application and Database Backup References.

Part 3 - Appendices

Disaster Recovery and Business Impact Analysis Templates, Windows BMR Configuration Scripts, Sample DRM Plan.

This book focuses on those TSM concepts specific to Disaster Recovery and assumes that the reader has experience with general TSM concepts. Therefore, we do not intend to provide a basic introduction to TSM. However, we provide references to other TSM Redbooks and manuals where appropriate.

1.2 Objective: Business Continuity

Business Continuity Planning (BCP) is an enterprise wide planning process which creates detailed procedures to be used in the case of a large unplanned outage or disaster. Maintaining continuity of business processes is the overall objective. Disaster Recovery Planning (DR Planning) is a logical subset of the BCP process, which focuses on continuity of IT operations.

From an IT-centric perspective, outages are classified as planned or unplanned disruptions to operations. Figure 1-1 shows the types of outages commonly experienced in enterprise computing environments.

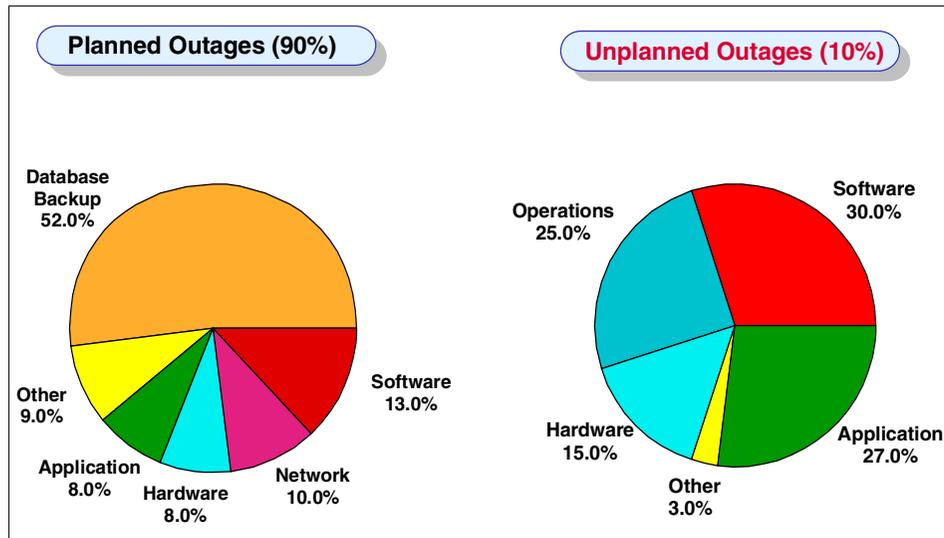


Figure 1-1 Planned versus unplanned outages for IT operations

An unplanned IT outage can equate to a disaster, depending on the scope and severity of the problem. Many Disaster Recovery plans focus solely on risks within the data center. We stress the importance of looking beyond the data center operations by implementing the BCP process in addition to traditional IT Disaster Recovery Planning. Beyond the data center, IT operations face a variety of risks shown in Figure 1-2.

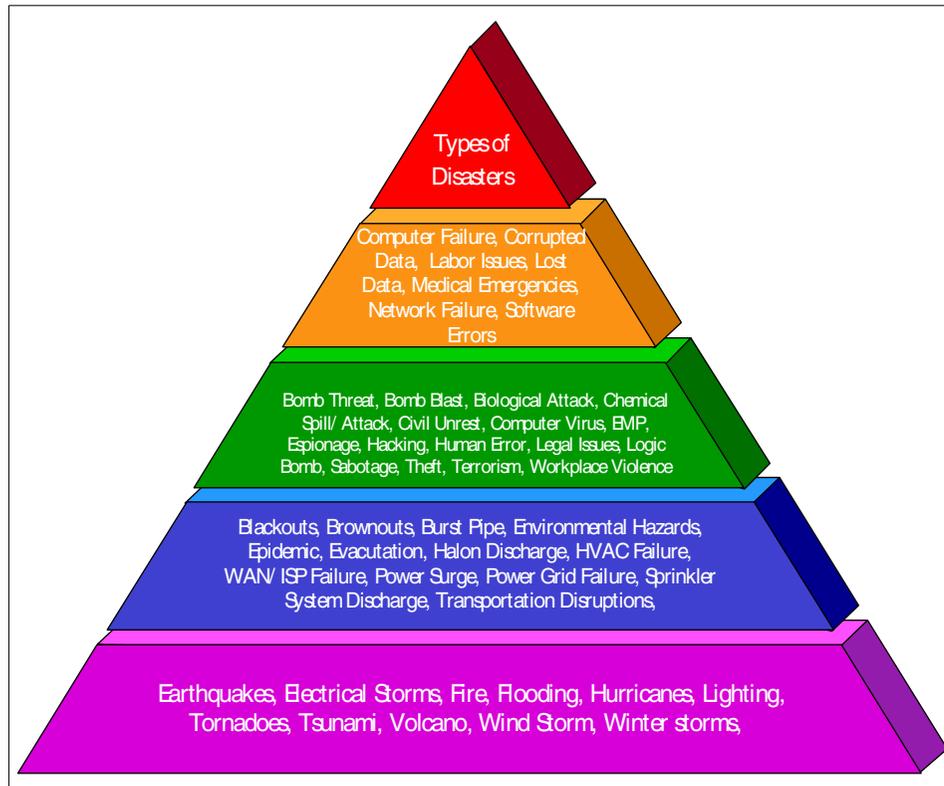


Figure 1-2 Types of disasters

An important part of preparing for a disaster is understanding the type of risks your organization faces. The top level of Figure 1-2 lists some of the common IT failures that disrupt operations. The other risk types are grouped into malicious behavior, infrastructure related, and natural disaster categories. Nearly every organization in the world faces feasible risks from many if not most of these levels.

Business continuity is achieved through rigorous planning, strategy, and process development. Part of the BCP planning process quantifies the critical business processes, the cost of downtime, and the risks an organization faces. The risks help to justify the means by which an organization builds availability, disaster tolerance, and disaster recovery capability into the IT infrastructure. The supporting IT infrastructure is closely associated with the Disaster Recovery Planning process.

Understanding risks and the associated cost of downtime for your business is a critical element of the planning process. Lost revenue is only a portion of the

comprehensive loss which could be sustained during an unplanned outage. Figure 1-3 lists the types of direct and indirect losses linked to IT downtime.

- **Employee Costs**
 - ▶ Employee and contractor idle time
 - ▶ Salaries paid to staff unable to undertake billable work
- **Direct Fiscal Losses**
 - ▶ Lost revenues
 - ▶ Delays in enterprise accounting
 - ▶ Loss of revenue for existing service contracts (customer SLA failure)
 - ▶ Lost ability to respond to contract opportunities
 - ▶ Loss of interest on overnight balances
 - ▶ Cost of interest on lost cash flow
- **Long Term Losses**
 - ▶ Penalties from failure to provide tax and annual reports
 - ▶ Loan rate fluctuations based on market valuation
 - ▶ Loss of control over debtors
 - ▶ Loss of credit control and increased bad debt.
 - ▶ Delayed profitability for new products and services
 - ▶ Brand image recovery
 - ▶ Loss of share value
 - ▶ Lost market share
- **Recovery Site Costs**
 - ▶ Cost of replacement of buildings and plant
 - ▶ Cost of replacing infrastructure and equipment
 - ▶ Cost of replacing software
 - ▶ Cost of DR contract activations
 - ▶ Cost of third party and contractor support

Figure 1-3 Potential cost factors for downtime

These losses are quantified and documented during the Business Impact Analysis (BIA) phase of the BCP process. Critical processes are identified and analyzed at a business level to determine the actual cost of downtime for each process. The enterprise cost of downtime varies from industry to industry, but in general the costs can be staggering. Figure 1-4 shows the cost of IT downtime across many US industries. These costs impose rigorous demands for data availability on the enterprise.

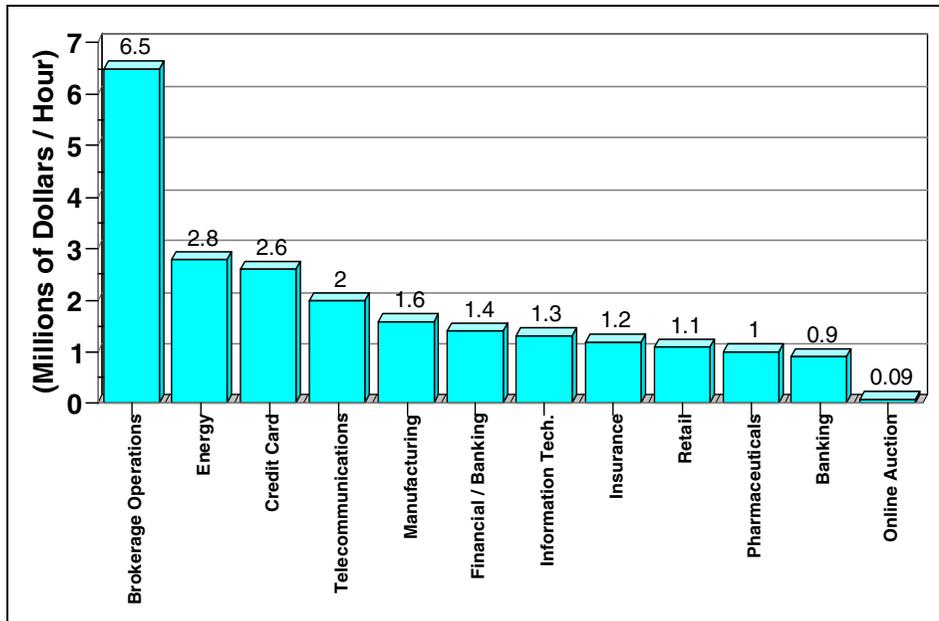


Figure 1-4 Average cost of downtime for various US industries

It doesn't take an independent analyst to realize that the costs associated with creating and assuring availability for the enterprise rise dramatically as you approach the requirement for 100% availability. The real challenge is defining the balance between the relative cost of downtime and the cost of maintaining availability for critical business processes. The following chapters about Disaster Recovery and TSM discuss planning methods to maximize availability for the enterprise.

Consider this scenario. You are an executive director in charge of IT operations for a large manufacturing outfit. In the early morning of a business day, you receive a phone call from a confused third shift network administrator, saying there was a network intrusion, a few key systems were attacked, and application data seems to be corrupted "everywhere". He is waiting for instructions and guidance. Four thousand full shift employees will begin arriving for critical business-day processes in three hours.

What steps do you follow to assess and remedy this situation? Does your staff have a detailed plan to follow? Have you designed the IT infrastructure to recover from this kind of outage in a timely fashion? Read the rest of this book to learn about strategies to deal with these questions and help you sleep better at night.

1.3 Disaster Recovery Planning terms and definitions

This section provides definitions of the most important and commonly used terms as used throughout this book. It is important to understand these terms, as many of them may seem similar and familiar, but are in fact often confused and describe slightly different concepts.

Business Continuity

Business continuity describes the processes and procedures an organization puts in place to ensure that essential functions can continue during and after a disaster. Business Continuity Planning seeks to prevent interruption of mission-critical services, and to re-establish full functioning as swiftly and smoothly as possible.

Business Impact Analysis (BIA)

A business impact analysis is performed to determine the impacts associated with disruptions to specific functions or assets in a firm — these include operating impact, financial impact, and legal or regulatory impact. For example, if billing, receivable, and collections business functions be crippled by inaccessibility of information, cash flow to the business will suffer. Additional risks are that lost customers will never return, the business' credit rating may suffer, and significant costs may be incurred for hiring temporary help. Lost revenues, additional costs to recover, fines and penalties, overtime, application and hardware, lost good will, and delayed collection of funds could be the business impact of a disaster.

Risk Analysis

A risk analysis identifies important functions and assets that are critical to a firm's operations, then subsequently establishes the probability of a disruption to those functions and assets. Once the risk is established, objectives and strategies to eliminate avoidable risks and minimize impacts of unavoidable risks can be set. A list of critical business functions and assets should first be compiled and prioritized. Following this, determine the probability of specific threats to business functions and assets. For example, a certain type of failure may occur once in 10 years. From a risk analysis, a set objectives and strategies to prevent, mitigate, and recover from disruptive threats should be developed.

Disaster Recovery Plan (DRP)

The DRP is an IT-focused plan that is designed to restore operability of the target systems, applications, or computer facility at an alternate site after an emergency. A DRP addresses major site disruptions that require site relocation. The DRP applies to major, usually catastrophic, events that deny access to the normal facility for an extended period. Typically, Disaster Recovery Planning involves an analysis of business processes and continuity needs; it may also include a significant focus on disaster prevention.

Disaster Tolerance

Disaster tolerance defines an environment's ability to withstand major disruptions to systems and related business processes. Disaster tolerance at various levels should be built into an environment and can take the form of hardware redundancy, high availability/clustering solutions, multiple data centers, eliminating single points of failure, and distance solutions.

DR hotsite

A DR hotsite is a data center facility with sufficient hardware, communications interfaces and environmentally controlled space capable of providing relatively immediate backup data processing support.

DR warmsite

A DR warmsite is a data center or office facility which is partially equipped with hardware, communications interfaces, electricity and environmental conditioning capable of providing backup operating support.

DR coldsite

A DR coldsite is one or more data center or office space facilities equipped with sufficient pre-qualified environmental conditioning, electrical connectivity, communications access, configurable space and access to accommodate the installation and operation of equipment by critical staff required to resume business operations.

Bare Metal Recovery

A bare metal recovery describes the process of restoring a complete system, including system and boot partitions, system settings, applications, and data to their original state at some point prior to a disaster.

High Availability

High availability describes a system's ability to continue processing and functioning for a certain period of time — normally a very high percentage of time, for example 99.999%. High availability can be implemented in your IT infrastructure by reducing any single points-of-failure (SPOF), using redundant components. Similarly, clustering and coupling applications between two or more systems can provide a highly available computing environment.

Recovery Time Objective (RTO)

The Recovery Time Objective is the time needed to recover from a disaster or, saying it another way, how long you can afford to be without your systems.

Recovery Point Objective (RPO)

Recovery Point Objective describes the age of the data you want the ability to restore in the event of a disaster. For example, if your RPO is six hours, you want to be able to restore systems back to the state they were in, as of no longer than

six hours ago. To achieve this, you need to be making backups or other data copies at least every six hours. Any data created or modified inside your recovery point objective will be either lost or must be recreated during a recovery. If your RPO is that no data is lost, synchronous remote copy solutions are your only choice.

Network Recovery Objective (NRO)

Network Recovery Objective indicates the time required to recover or fail over network operations. Keep in mind that systems level recovery is not fully complete if customers cannot access the application services via network connections. Therefore, the NRO includes the time required to bring online alternate communication links, reconfigure routers and name servers (DNS) and alter client system parameters for alternative TCP/IP addresses. Comprehensive network failover planning is of equal importance to data recovery in a Disaster Recovery scenario.

1.4 IBM Tivoli Storage Manager overview

IBM Tivoli Storage Manager (TSM) is a storage management application built for the enterprise. TSM provides an enterprise solution for data protection, disaster recovery, space management, and record retention. TSM facilitates flexible and scalable storage management policies to support complicated business needs for storage management and disaster recovery. Most importantly, TSM automates storage management tasks by eliminating labor and cost intensive manual procedures for backup, archive, and recovery.

1.4.1 TSM platform support

TSM protects and manages data on more than 30 operating platforms. The TSM server application is supported on over 10 platforms, and it supports hundreds of disk, tape, and optical storage devices. The TSM server software provides built-in device drivers for directly connecting more than 300 different device types from every major manufacturer. All common LAN, WAN, and SAN infrastructures are also supported by TSM. Figure 1-5 summarizes TSM platform support.

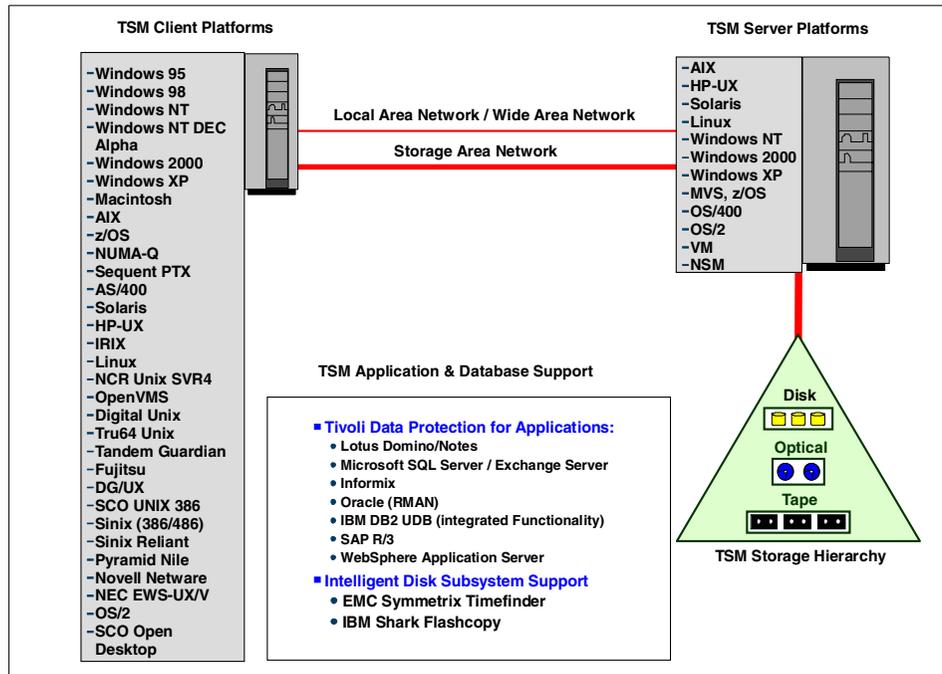


Figure 1-5 TSM client, server, application, and database support

Tivoli Storage Manager provides data protection, disaster recovery, and storage management functionality for the enterprise. TSM storage management services include:

- ▶ **Operational Backup and Restore of Data:** The backup process creates a copy of the data to protect against the operational loss or destruction of file or application data. The customer defines how often to back up (frequency) and how many copies (versions) to hold. The restore process places the backup copy of the data back onto the designated system or workstation.
- ▶ **Disaster Recovery:** By the creation of multiple copies of enterprise data, TSM supports the implementation of site to site recovery operations. Disaster Recovery with TSM includes moving data to offsite locations, rebuilding or initializing TSM infrastructure, and reloading data to clients in an acceptable time frame. Many such scenarios are discussed in later chapters.
- ▶ **Vital Record Retention, Archive, and Retrieval:** The archive process creates a copy of a file or a set of files for long term storage. Files can remain on the local storage media or can be deleted. The customer controls how long (retention period) an archive copy is to be retained. The retrieval process locates the copies within the archival storage and places them back into a customer-designated system.

- ▶ **Hierarchical Space Management:** This process provides the automatic and transparent movement of operational data from the user system disk space to a central storage repository. If the user needs to access this data, it is dynamically and transparently restored to the client storage.

1.4.2 How TSM works

Tivoli Storage Manager is implemented as a client-server software application. The TSM server software component coordinates the movement of data from TSM Backup/Archive clients across the network or SAN to a centrally managed storage hierarchy. The classic TSM hierarchy includes disk, tape, and in some cases optical devices for data storage. Figure 1-6 shows the general movement of data in a TSM environment.

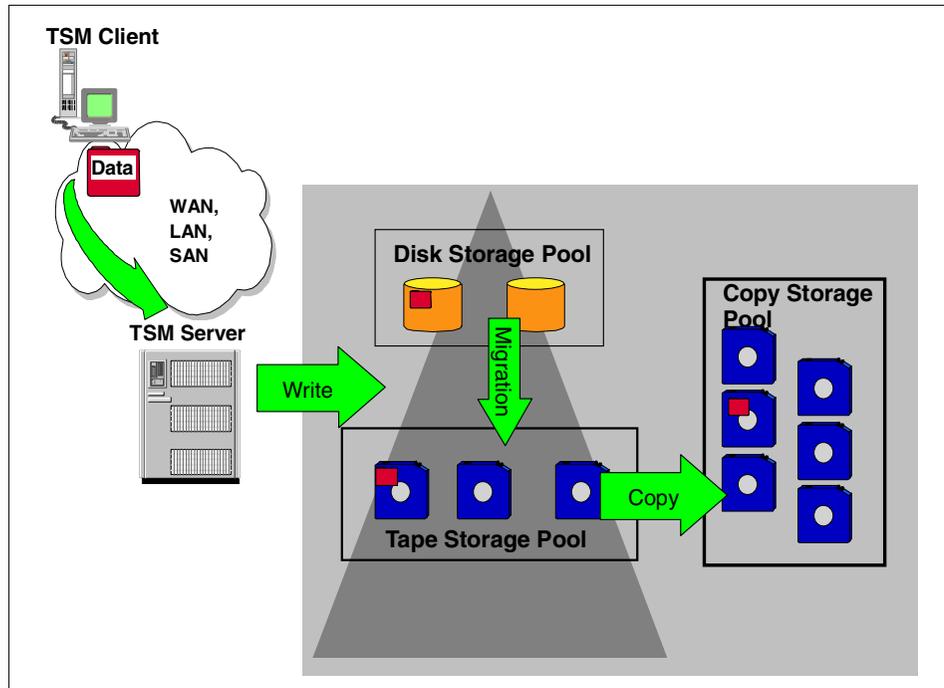


Figure 1-6 Data movement with TSM and the TSM storage hierarchy

TSM client data is moved via SAN or LAN connections to the TSM server, written directly to disk or tape primary storage pool (and optionally simultaneously to a copy storage pool in TSM 5.1), migrated to other storage primary storage pools, and copied as many times as necessary to additional copy storage pools.

TSM manages and stores all metadata about policies, operations, locations of data, and TSM component definitions. It uses an internal relational database as the repository for all its data. This component of TSM server architecture makes TSM extremely scalable for large implementations. The TSM database requires minimal database administration.

1.4.3 TSM Server administration

The TSM Server can be administered via a command-line interface or via the Web browser GUI. Both interfaces can be invoked either locally or remotely (with a TCP/IP connection to the TSM Server). The Web browser interface also features a panel for entering in text-commands. TSM commands can also be scripted for automation of routine administrative processes. Figure 1-7 shows the TSM Web browser administrative interface.

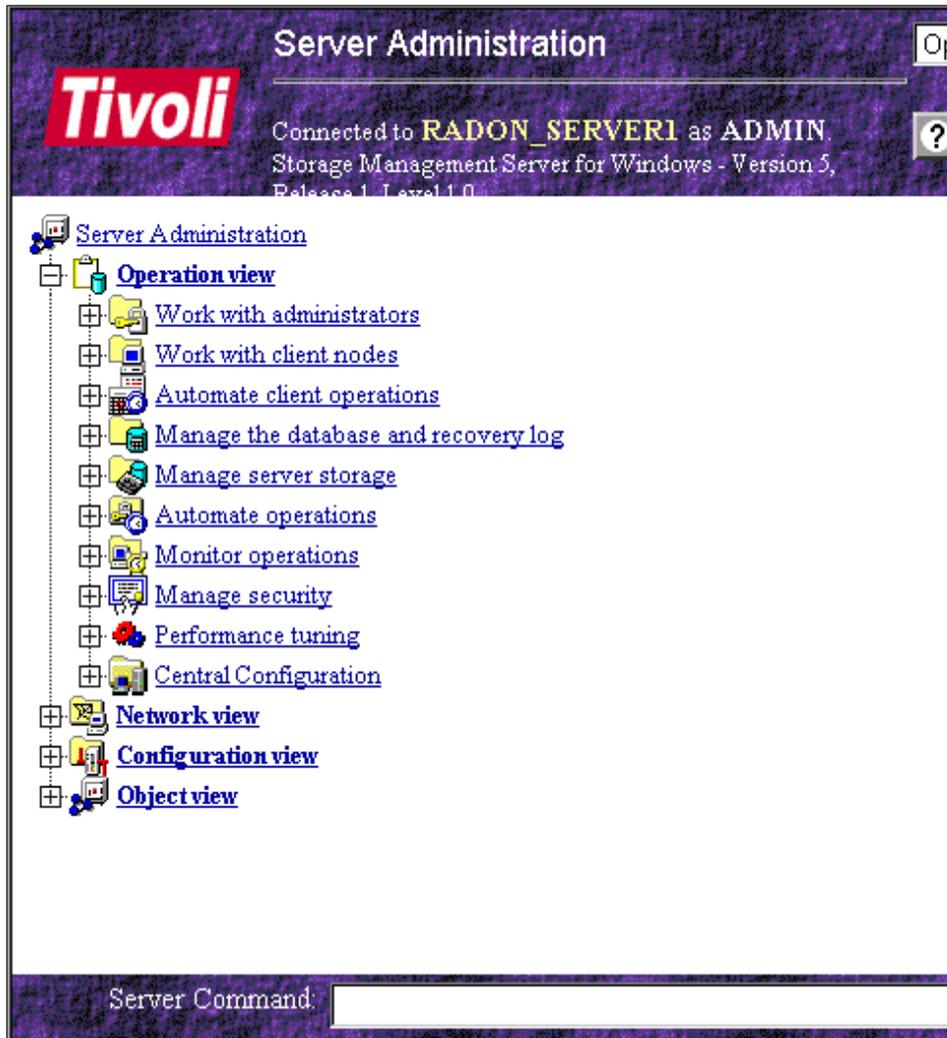


Figure 1-7 Sample TSM administration screenshot

For multi-server TSM environments, TSM enterprise administration allows one TSM administrator to centrally configure, maintain, route commands, and log events for multiple TSM server environments. This feature lowers management cost and complexity significantly in large multi-TSM server environments.

1.4.4 TSM Backup/Archive Client

The TSM Backup/Archive Client interface provides full backup and restore functionality in a user-friendly format. Users can select any combination of full,

individual directory or individual file backups, as well as backup sets and image backups, via the interface. TSM end users can also use the interface to restore files, without requiring administrator intervention. The Backup/Archive Client also provides the archive and retrieve function. Depending on the client platform, the Backup/Archive Client may be available as a command-line, graphical or Web based interface. Figure 1-8 shows the Windows TSM GUI Client interface.

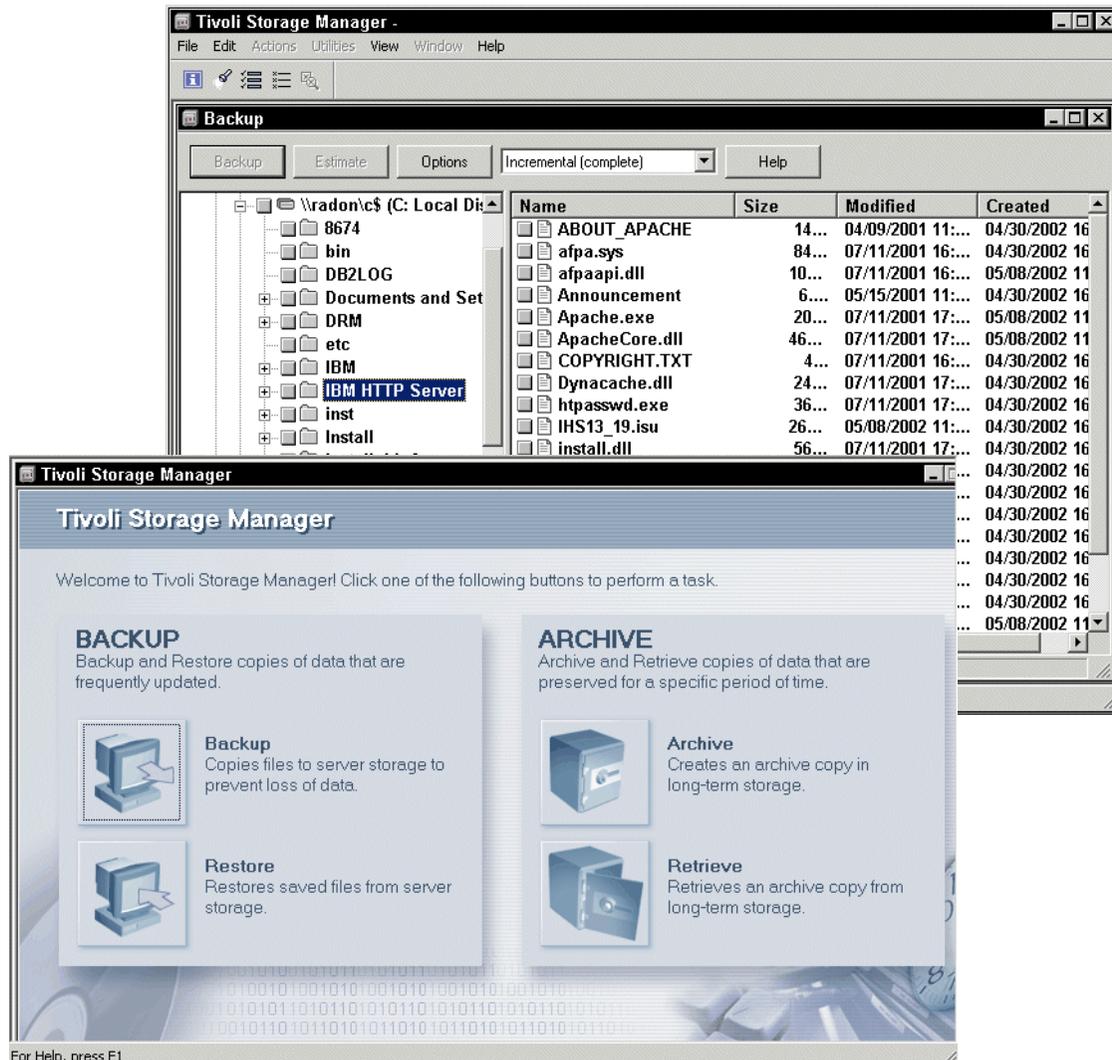


Figure 1-8 The TSM Version 5 client interface

TSM implements the patented *progressive backup methodology* and *adaptive subfile backup technology*. Most of the Backup/Archive Clients are able to exploit

the multi-threading capabilities of modern operating systems. TSM supports parallel backup and recovery processes to allow expedient movement of data to and from the TSM client.

1.4.5 TSM backup and archive concepts

Backup, in Tivoli Storage Manager terms, means creating a copy of a data object to be used for recovery. A TSM data object can be a file, a part of a file, a directory or a user defined data object like a database table. The backup version of this data object is stored separately in the Tivoli Storage Manager server storage hierarchy. TSM policy tools allow great flexibility for the way data is managed for each client. Backup frequency, retention, and copy policies are easily implemented on the TSM client. The variety of backup types supported by TSM are explained in detail in 7.5.1, “Client backup and restore operations” on page 124.

In addition to data backup, *archive* copies of data can also be created using TSM. Archive creates an additional copy of data and stores it for a specific amount of time — known as the retention period. TSM archives are not expired until the retention period is past, even if the original files are deleted from the client system.

Therefore, the difference between *backup* and *archive* is that backup creates and controls multiple backup versions that are directly attached to the original file; whereas archive creates an additional file that is retained for a specific period of time.

Progressive incremental backups

One of the many advantages of TSM is the *progressive incremental backup methodology*. After the first full backup, TSM then operates with incremental backups only. Also known as *incremental forever*, progressive incremental means only those files that have changed since the last backup will be backed up. Incremental backup by date is also available. This methodology reduces network and storage resource consumption and lowers the overall cost of storage management. Tivoli Storage Manager’s file level progressive backup methodology is far superior to other traditional backup methods such as Full+Incremental or Full+Differential, because progressive incremental backups are never redundant.

1.4.6 Tivoli Data Protection for Applications modules

Tivoli Data Protection for Applications modules are separate program products which connect business applications to the Tivoli Storage Manager

environment. These applications allow application specific storage management controls for backup and restore operations.

Oracle, Informix, Lotus Notes, Lotus Domino, Microsoft Exchange, Microsoft SQL Server, SAP R/3, and WebSphere Application Server each have their own storage management interface or TDP application which integrates with the TSM data management API in each TSM data protection application. DB2 from IBM integrates the TSM API directly, without requiring a separately purchased TDP product. Some of the TSM data protection applications leverage IBM and EMC *intelligent disk subsystem* advanced copy functions such as FlashCopy and TimeFinder. This functionality bridges TSM and high-availability storage infrastructures to maximize application availability.

1.4.7 TSM security concepts

Security is a vital aspect for enterprise storage management. Data must be protected, available, and secure. From the moment data is backed up from the client, TSM provides a secure storage management environment. TSM is the only interface to your backup and archive data. In later chapters we discuss methods to safeguard the TSM server and storage environment for Disaster Recovery.

Before a communication session between the TSM Client and the TSM Server begins, an authentication *handshaking* process occurs with authentication tickets and a *mutual suspicion algorithm*. The TSM security protocol is modeled after the Kerberos network authentication protocol, which is a highly respected method for secure signon cryptography. The client uses its password as part of an encryption key, and does not send the password over the network. Each session key is unique, so replaying a session stream will not result in a signon to the TSM server. This significantly lowers the chance of a TSM session being hijacked by an outside user.

To heighten security for TSM sessions, data sent to the TSM server during backup and archive operations can be encrypted with standard DES 56-bit encryption. For WAN implementations of TSM across public networks, data encryption compliments and completes data security for TSM.

1.5 TSM and Disaster Recovery

TSM protects enterprise data. TSM also has many tools, functions, and potential strategies that can be used specifically for Disaster Recovery. This long list of features and strategies that specifically addresses Disaster Recovery needs are the topic of this book. Note, while TSM does provide many tools for defense against a disaster, a good understanding of individual business continuity needs,

careful planning, testing, and use of appropriate hardware technologies are also critical to a successful Disaster Recovery strategy. To this end, we also discuss these in this part of the book.

A list of TSM tools and strategies for protection against disasters and for recovering in the event of disasters is given here:

- ▶ Database and recovery log mirroring
- ▶ Database page shadowing
- ▶ Storage pool manipulation for disaster protection
- ▶ Varied client backup operations
- ▶ Varied backup methods and topologies
- ▶ TSM Disaster Recovery Manager (DRM)
- ▶ TSM server-to-server communications
- ▶ TSM server-to-server virtual volumes
- ▶ TSM and high availability clustering
- ▶ TSM and remote disk replication
- ▶ TSM traditional and electronic tape vaulting.
- ▶ TSM and system bare metal restore integration

DRM overview

Tivoli Storage Manager delivers Tivoli Disaster Recovery Manager (DRM) as part of its Extended Edition. DRM offers various options to configure, control and automatically generate a Disaster Recovery Plan containing the information, scripts, and procedures needed to automate restoration of the TSM Server and helps ensure quick recovery of client data after a disaster. It also manages and tracks the media on which TSM data is stored, whether on site, in-transit, or in a vault, so that data can be easily located if disaster strikes. It generates scripts which assist in documenting IT systems and recovery procedures, as well as providing automated steps to rebuild the TSM server.

TSM APIs

The TSM APIs are used for Tivoli's own TDP products (see 1.4.6, "Tivoli Data Protection for Applications modules" on page 17), but they are also published and documented. This allows ISVs to adapt their solutions to integrate with TSM to extend its functionality. In particular, various vendors have used the APIs to provide bare metal recovery solutions for various platforms. Among the vendors exploiting these APIs for Disaster Recovery include Cristie, UltraBac Software, and VERITAS Bare Metal Restore. More information on these companies and

products is provided in 12.2, “Using third party products for BMR with TSM” on page 285.



The tiers of Disaster Recovery and TSM

This chapter introduces and describes the seven tiers of Disaster Recovery solutions established by the SHARE User Group. It describes each of the tiers in the context of recovery time, cost, and tier characteristics.

Once the tiers have been described we outline specific TSM functions and strategies that can be used to achieve the various tiers.

An understanding of the seven tiers of Disaster Recovery solutions lays the foundation to help you determine what level of Disaster Recovery solution you currently have and what level you may want or need to achieve in the future.

2.1 Seven tiers of Disaster Recovery

Understanding DR strategies and solutions can be very complex. To help categorize the various solutions and their characteristics (for example costs, recovery time capabilities, recovery point capabilities) definitions of the various levels and required components can be defined. The idea behind such a classification would be to help those concerned with DR determine:

- ▶ What kind of solution they have?
- ▶ What kind of solution they require?
- ▶ What it would require to meet greater DR objectives?

In 1992, the SHARE user group in the United States, in combination with IBM, defined a set of DR tier levels. This was done to address the need to properly describe and quantify various different methodologies for successful mission-critical computer systems DR implementations. Accordingly, within the IT Business Continuity industry, the tier concept continues to be used, and is very useful for describing today's DR capabilities. The tiers' definitions are designed so that emerging DR technologies can also be applied. These tiers are summarized in Figure 2-1.

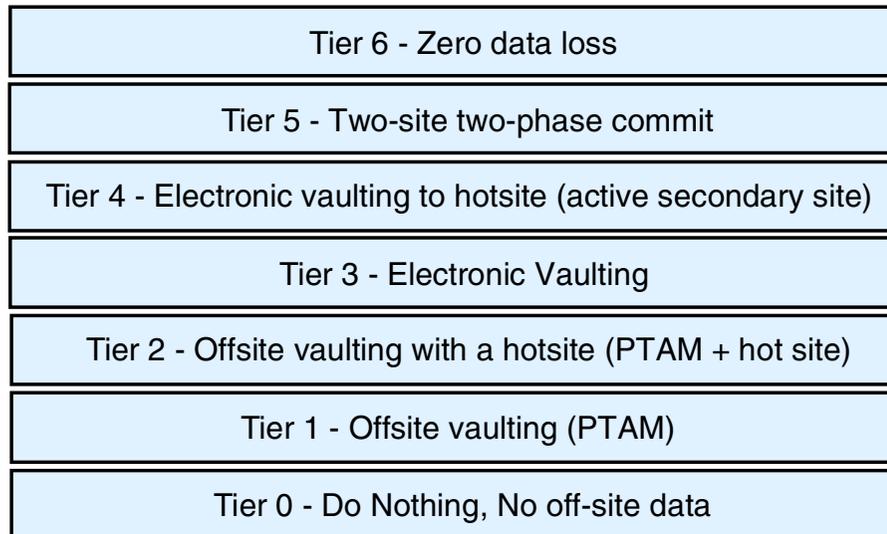


Figure 2-1 Summary of Disaster Recovery tiers (SHARE)

The following sections provide an overview of each of the tiers, describing their characteristics and associated costs. Typical recovery times (based on industry experience and the capabilities of the recovery strategy) are also noted. The purpose is to introduce these tiers for those not familiar with them, and then later directly link these recovery tiers with TSM DR strategies.

2.1.1 Tier 0 - Do nothing, no off-site data

Tier 0 is defined as a single site data center environment having no requirements to backup data or implement a Disaster Recovery Plan. See Figure 2-2 for an illustration.

On this tier, there is no saved information, no documentation, no backup hardware, and no contingency plan. There is therefore no DR capability at all. In our experience, some customers still reside in this tier. For example, while some customers actively make backups of their data, these backups are left onsite in the same computer room, or occasionally are not removed from the site due to lack of a rigorous vaulting procedure. A customer data center residing on this tier is exposed to a disaster from which they may never recover their business data!

Note: The typical length of recovery time in this instance is unpredictable. In many cases complete recovery of applications, systems, and data is never restored.

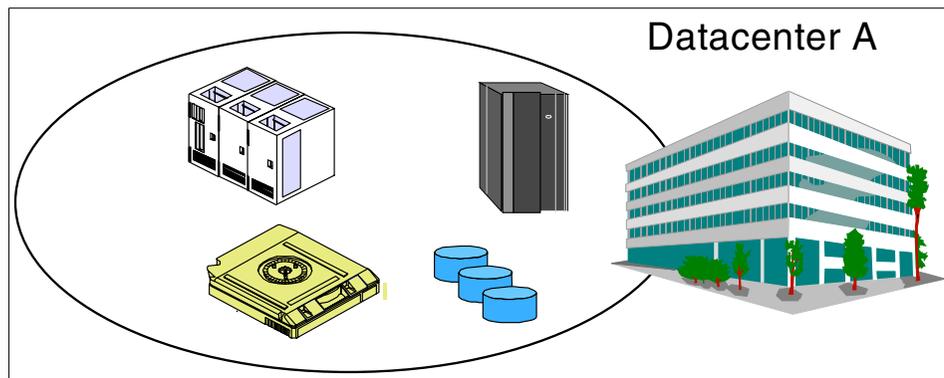


Figure 2-2 Tier 0 - Do nothing, no offsite data

2.1.2 Tier 1 - Offsite vaulting (PTAM)

A Tier 1 installation is defined as having a DRP, backs up and stores its data at an offsite storage facility and has determined some recovery requirements. As shown in Figure 2-3, backups are being taken which are being stored at an offsite storage facility. This environment may also have established a backup platform, although it does not have a site at which to restore its data, nor the necessary hardware on which to restore the data, for example, compatible tape devices.

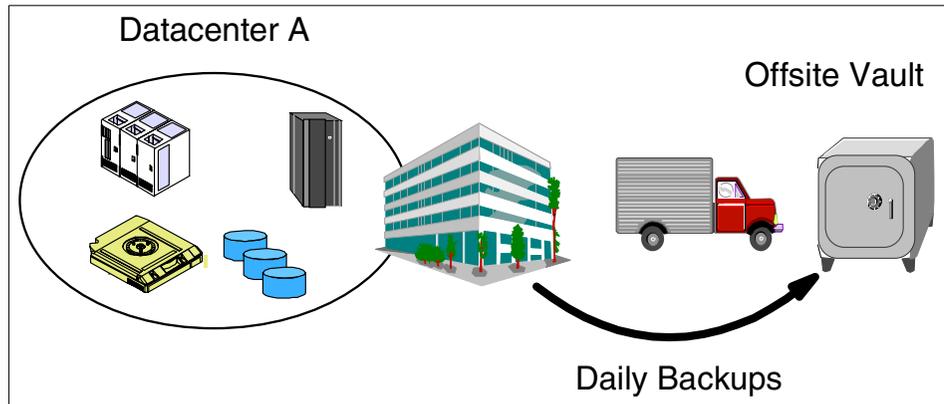


Figure 2-3 Tier 1 - Offsite vaulting (PTAM)

Because vaulting and retrieval of data is typically handled by couriers, this tier is described as the Pickup Truck Access Method (PTAM). PTAM is a method used by many sites, as this is a relatively inexpensive option. It can, however, be difficult to manage, that is, it is difficult to know exactly where the data is at any point. There is probably only selectively saved data. Certain requirements have been determined and documented in a contingency plan and there is optional backup hardware and a backup facility available.

Recovery is dependent on when hardware can be supplied, or possibly when a building for the new infrastructure can be located and prepared.

While some customers reside on this tier and are seemingly capable of recovering in the event of a disaster, one factor that is sometimes overlooked is the recovery time objective (RTO). For example, while it may be possible to eventually recover data, it may take several days or weeks. An outage of business data for this period of time can have an impact on business operations that lasts several months or even years (if not permanently).

Note: The typical length of time for recovery is normally more than a week.

2.1.3 Tier 2 - Offsite vaulting with a hot site (PTAM + hot site)

Tier 2 encompasses all requirements of Tier 1 (offsite vaulting and recovery planning) plus it includes a hot site. The hot site has sufficient hardware and a network infrastructure able to support the installation's critical processing requirements. Processing is considered critical if it must be supported on hardware existing at the time of the disaster. As shown in Figure 2-4, backups are being taken and they are being stored at an offsite storage facility. There is

also a hot site available and the backups can be transported there from the offsite storage facility in the event of a disaster.

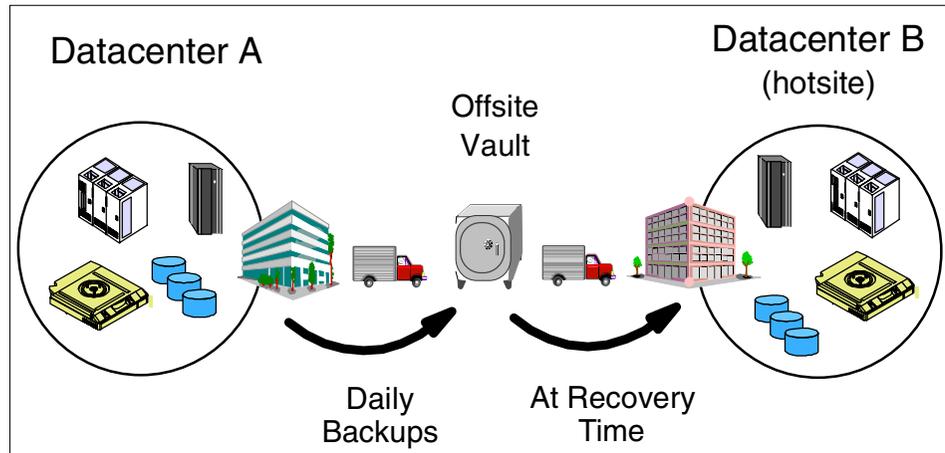


Figure 2-4 Tier 2 - Offsite vaulting with a hot site (PTAM + hot site)

Tier 2 installations rely on a courier (PTAM) to get data to an offsite storage facility. In the event of a disaster, the data at the offsite storage facility is moved to the hot site and restored onto the backup hardware provided. Moving to a hot site increases the cost but reduces the recovery time significantly. The key to the hot site is that appropriate hardware to recover the data (for example, a compatible tape device) is present and operational.

Note: The typical length of time for recovery is normally more than a day.

2.1.4 Tier 3 - Electronic vaulting

Tier 3 encompasses all the components of Tier 2 (offsite backups, disaster recovery plan, hot site) and, in addition, supports electronic vaulting of some subset of the critical data. Electronic vaulting consists of electronically transmitting and creating backups at a secure facility, moving business-critical data offsite faster and more frequently than traditional data backup processes allow. The receiving hardware must be physically separated from the primary site and the data stored for recovery should there be a disaster at the primary site. As shown in Figure 2-5, backups are being taken and they are then being stored at an offsite storage facility. There is also a hot site available and the backups can be transported there from the offsite storage facility. There is also electronic vaulting of critical data occurring between the primary site and the hot site.

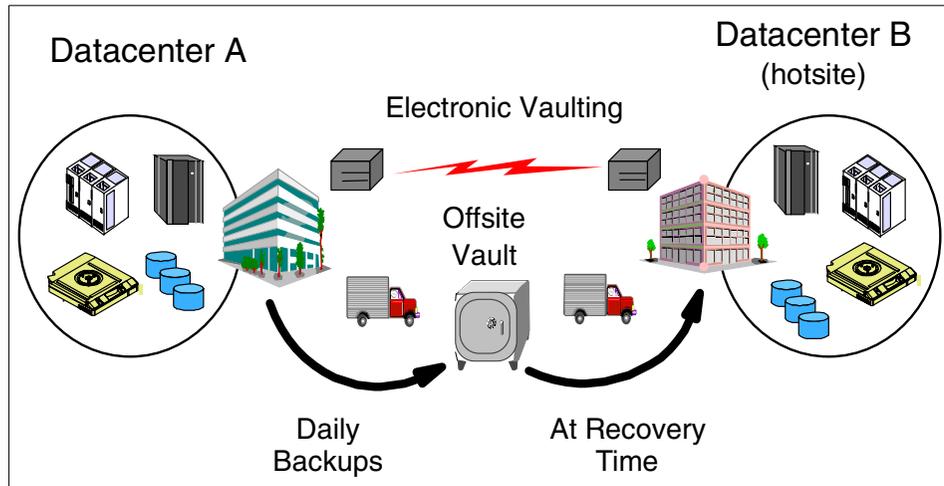


Figure 2-5 Tier 3 - Offsite electronic vaulting

The hotsite is kept running permanently, thereby increasing the cost. As the critical data is already being stored at the hotsite, the recovery time is once again significantly reduced. Often, the hotsite is a second data center operated by the same firm or a Storage Service Provider.

Note: The typical length of time for recovery is normally about one day.

2.1.5 Tier 4 - Electronic vaulting to hotsite (active secondary site)

Tier 4 is defined as using two data centers with electronic vaulting between both sites and introduces the requirements of active management of the data being stored at the recovery site. This is managed by a processor at the recovery site and can support bi-directional recovery. The receiving hardware must be physically separated from the primary platform. As shown in Figure 2-6, backups are being taken and they are being stored at an offsite storage facility. There is also a hotsite available and the backups can be transported there from the offsite storage facility. There is also continuous transmission of data or connection between the primary site and the hot site, supported by high bandwidth connections.

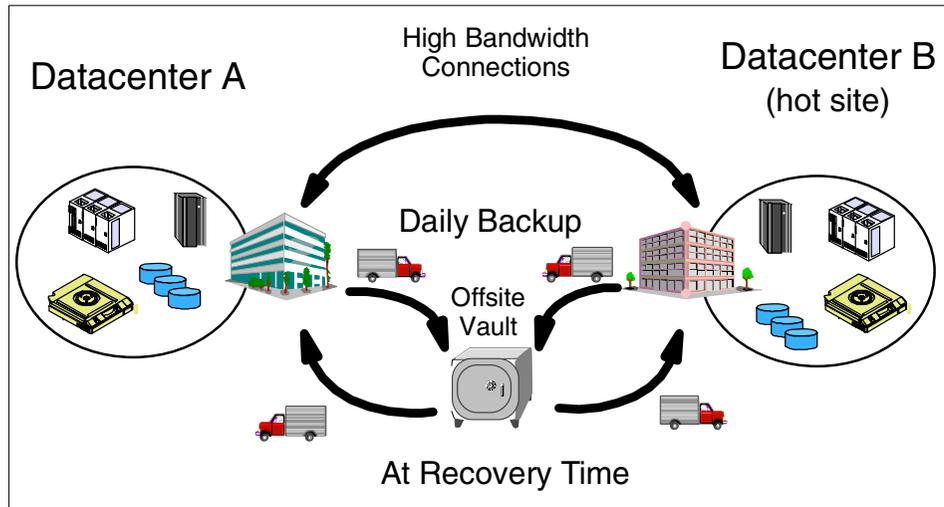


Figure 2-6 Tier 4 - Electronic vaulting with hot site (active secondary site)

In this scenario, the workload may be shared between the two sites. There is a continuous transmission of data between the two sites with copies of critical data available at both sites. Any other non-critical data still needs to be recovered from the offsite vault via courier in the event of a disaster.

Note: The typical length of time for recovery is usually up to one day.

2.1.6 Tier 5 - Two-site, two-phase commit

Tier 5 encompasses all the requirements of Tier 4 (offsite backups, disaster recovery plan, electronic vaulting, and active secondary site), and in addition, will maintain selected data in image status (updates will be applied to both the local and the remote copies of the database within a single-commit scope). Tier 5 requires that both the primary and secondary platforms' data be updated before the update request is considered successful. As shown in Figure 2-7, the two sites are synchronized utilizing a high-bandwidth connection between the primary site and the hot site.

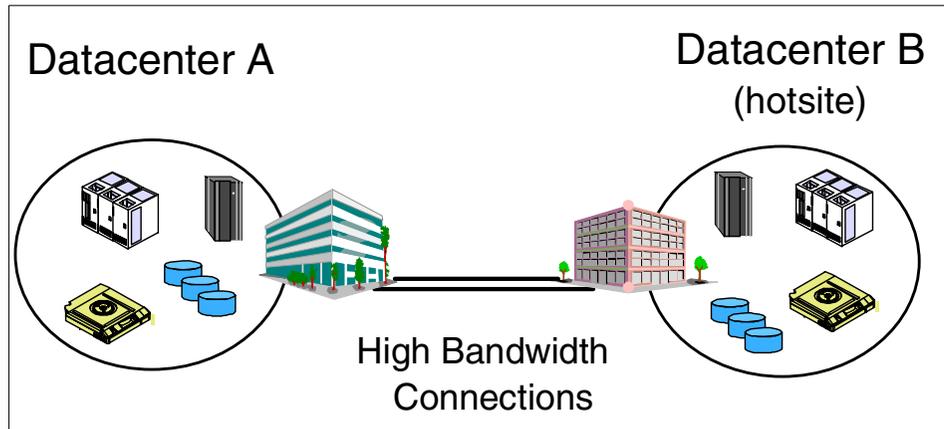


Figure 2-7 Tier 5 - Two-site, two-phase commit

Tier 5 also requires partially or fully dedicated hardware on the secondary platform with the ability to automatically transfer the workload over to the secondary platform. We now have a scenario where the data between the two sites is synchronized by remote two-phase commit. The critical data and applications are therefore present at both sites and only the in-flight data is lost during a disaster. With a minimum amount of data to recover and reconnection of the network to implement, recovery time is reduced significantly.

Note: The typical length of time for recovery is usually less than 12 hours.

2.1.7 Tier 6 - Zero data loss

Tier 6 encompasses zero loss of data and immediate and automatic transfer to the secondary platform. Data is considered lost if a transaction has commenced (for example, a user hits the Enter key to initiate an update), but the request has not been satisfied. Tier 6 is the ultimate level of Disaster Recovery. Local and remote copies of all data are updated and dual online storage is utilized with a full network switching capability. As shown in Figure 2-8 the two sites are fully synchronized utilizing a high-bandwidth connection between the primary site and the hotsite. The two systems are advanced coupled, allowing an automated switchover from one site to the other when required.

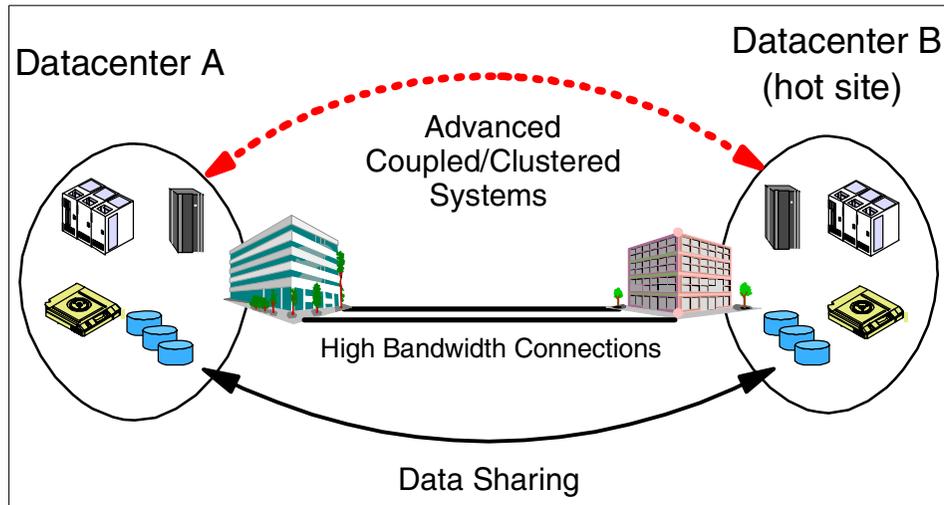


Figure 2-8 Tier 6 - Zero data loss (advanced coupled systems)

This is the most expensive Disaster Recovery solution as it requires coupling or clustering applications, additional hardware to support data replication, and high bandwidth connections over extended distances. However, it also offers the speediest recovery by far.

Note: The typical length of time for recovery is normally a few minutes.

2.2 Costs related to DR solutions

As Disaster Recovery objectives move towards the higher tiers, fixed costs, implementation costs, ongoing network costs, and maintenance costs generally grow exponentially. For example, higher tier DR solutions may require:

- ▶ A secondary site with redundant operational equipment, media, software licensing, and staff
- ▶ High bandwidth connections over long distances
- ▶ Additional backup software modules to support advanced features
- ▶ Coupling or clustering management software
- ▶ Hardware that supports or provides point-in-time volume copies or remote data replication

Figure 2-9 illustrates the relationship between the tiers of disaster recovery solutions, recovery time, and cost.

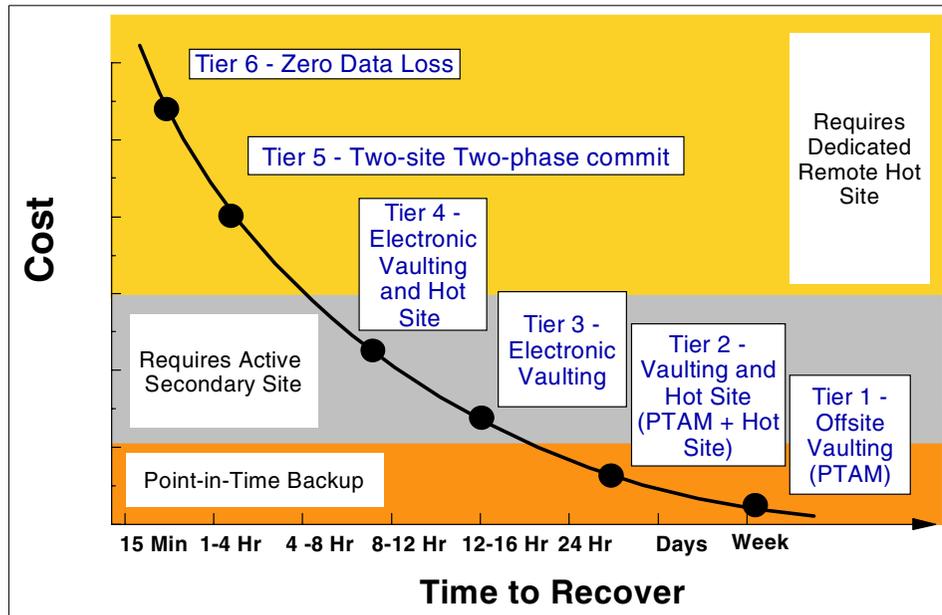


Figure 2-9 Seven tiers of Disaster Recovery solutions

Many of the tiers described define the ability to recover your data. The distinction between the tiers are how quickly you need to recover your data (RTO), how quickly you need to recover the services provided by your environment, and how much data you cannot afford to lose (RPO). Therefore, a recovery solution should be chosen based on your business' unique recovery criteria versus how much it will cost the company in lost revenue due to being down and unable to continue normal business processing. The shorter the time period required to recover the data to continue business processing, the higher the cost. Almost always, the longer a company is unable to process transactions the more expensive the outage is going to be for the company.

Therefore, it is important to understand that the cost of a solution must be in reasonable proportion to the business value of IT. You do not want to spend more money on a DR solution than the financial loss you would suffer from a disaster. The financial loss you may suffer from a disaster should be determined from past experiences with disruptions, a Business Impact Analysis (BIA), a risk assessment, and careful planning. There is no doubt that your data is very important and downtime can be very costly. In many cases, such as financial institutions, companies are required to implement very high levels of disaster protection and recoverability.

What we are recommending is a determination of the best solution for your environment, based on a careful assessment of your business, planning, and preparation.

2.3 Strategies to achieve tiers of DR using TSM

IBM Tivoli Storage Manager (TSM) can be used to help provide all the tiers of DR described in the previous section. A variety of functions provided by TSM can be used to meet DR objectives. These TSM functions are the subject of the remainder of this book. A short summary of these strategies includes:

- ▶ Use of Disaster Recovery Manager (DRM) which can automate the TSM server recovery process and manage offsite volumes.
- ▶ Vaulting of TSM database, recovery log, volume history information, device configuration information, DRP file (if using DRM) and copy pools for storage at an offsite location.
- ▶ Use of TSM server-to-server communications to enable enterprise configuration (multiple TSM servers), enterprise event logging and monitoring, and command routing.
- ▶ TSM servers installed at multiple locations, optionally setup as peer to peer servers (that is, each server able to recover at the alternate site).
- ▶ Use of TSM virtual volumes over TCP/IP connection to allow storage of TSM entities (TSM database backups, recovery log backups, and primary and copy storage pools, DRM plan files) on remote target servers.
- ▶ Use of high bandwidth connections and data replication technology (such as IBM PPRC, EMC SRDF) to support asynchronous/synchronous data replication of TSM databases backups, recovery log backups, TSM database and recovery log mirrors, and storage pools.
- ▶ Use of remote electronic tape vaulting of TSM database and recovery log backups, primary or copy storage pools. Extended distances can be achieved by using distance technologies, for example, extended SAN, DWDM, IP/WAN channel extenders.

Table 2-1 provides a summary of strategies and techniques that can be used with TSM to achieve the various tiers of disaster recovery.

Table 2-1 TSM strategies to achieve Disaster Recovery tiers

Tier #	Description	TSM integrated strategies
0	Do nothing, No offsite backups	<p>No Offsite Strategy.</p> <p>Normal TSM-based onsite backups.</p> <p>With site disaster there will be no ability to recover except for rebuild of environment.</p>
1	<p>Offsite Vaulting</p> <p>Also Known as Pickup Truck Access Method (PTAM)</p>	<p>Storage pool vaulting with TSM server environment (fully integrated).</p> <p>Requires a DRP and careful management of offsite volumes. Consider use of Disaster Recovery Manager (DRM) which can automate the TSM server recovery process and manage offsite volumes.</p> <p>Strategy includes vaulting of TSM database, recovery log, volume history information, device configuration information, DRP file (if using DRM) and copy pools for storage at an offsite location.</p>
2	<p>Offsite Vaulting with a hot site</p> <p>Also Known as PTAM + hot site</p>	<p>TSM server installed at both locations.</p> <p>Storage pool vaulting with TSM server environment (fully integrated).</p> <p>Requires a DRP and careful management of offsite volumes. Consider use of Disaster Recovery Manager (DRM) which can automate the TSM server recovery process and manage offsite volumes.</p> <p>Strategy includes vaulting of TSM database, recovery log, volume history information, device configuration information, DRP file (if using DRM) and copy pools for storage at an offsite location.</p> <p>Consider use of TSM server-to-server communications to enable enterprise configuration (multiple TSM servers), enterprise event logging and monitoring, and command routing.</p>

Tier #	Description	TSM integrated strategies
3	Electronic Vaulting	<p>TSM server installed at both locations.</p> <p>Requires a DRP and careful management of offsite volumes. Consider use of Disaster Recovery Manager (DRM) which can automate the TSM server recovery process and manage offsite volumes.</p> <p>Consider use of TSM virtual volumes over TCP/IP connection to allow storage of TSM entities (TSM database backups, recovery log backups, and primary and copy storage pools, DRM plan files) on remote target servers, which store as archive files on a the target server.</p> <p>Strategy can include vaulting of TSM database, recovery log, volume history information, device configuration information, DRP file (if using DRM) and copy pools for storage at an offsite location.</p> <p>Consider use of TSM server-to-server communications to enable enterprise configuration (multiple TSM servers), enterprise event logging and monitoring, and command routing.</p>

Tier #	Description	TSM integrated strategies
4	Electronic Vaulting to hot site (active secondary site)	<p>TSM servers installed at both locations, optionally setup as peer to peer servers (that is, each server able to recover at the alternate site).</p> <p>Requires a DRP and careful management of offsite volumes. Consider use of Disaster Recovery Manager (DRM) which can automate the TSM server recovery process and manage offsite volumes.</p> <p>Strategy may include use of TSM virtual volumes over TCP/IP connection to allow storage of TSM entities (TSM database backups, recovery log backups, and primary and copy storage pools, DRM plan files) on remote target servers, which store as archive files on a target server.</p> <p>High bandwidth connections and data replication technology (for example, IBM PPRC, EMC SRDF) will support asynchronous data replication of TSM databases backups and recovery log backups. TSM storage pools with critical data can be replicated as well. Extended distances can be achieved by using distance technologies, for example, extended SAN, DWDM, IP/WAN Channel Extenders.</p> <p>Solution may include remote electronic tape vaulting of TSM database and recovery log backups, primary or copy storage pools. Extended distances can be achieved by using distance technologies, for example, extended SAN, DWDM, IP/WAN Channel Extenders.</p> <p>Strategy includes vaulting of TSM database, recovery log, volume history information, device configuration information, DRP file (if using DRM) and copy pools for storage at an offsite location.</p> <p>Consider use of TSM server-to-server communications to enable enterprise configuration (multiple TSM servers), enterprise event logging and monitoring, and command routing.</p>

Tier #	Description	TSM integrated strategies
5	Two-site, Two-phase Commit	<p>TSM servers installed at both locations, optionally setup as peer to peer servers (that is, each server able to recover at the alternate site).</p> <p>High bandwidth connections and data replication technology (for example, IBM PPRC, EMC SRDF) will support synchronous data replication of TSM databases backups and recovery log backups or mirrors. TSM storage pools with critical data can be replicated as well. Extended distances can be achieved by using distance technologies, for example, extended SAN, DWDM, IP/WAN Channel Extenders.</p> <p>Solution may include remote electronic tape vaulting of TSM database and recovery log backups, primary or copy storage pools. Extended distances can be achieved by using distance technologies, for example, extended SAN, DWDM, IP/WAN Channel Extenders.</p> <p>Consider use of TSM server-to-server communications to enable enterprise configuration (multiple TSM servers), enterprise event logging and monitoring, and command routing.</p>
6	Zero Data Loss	<p>TSM servers installed at both locations, optionally setup as peer to peer servers (that is, each server able to recover at the alternate site).</p> <p>Requires dual active data centers, high availability application (for example, HACMP, HAGEO, MSCS) to support hot failover of server(s) from one data center to the other. Use of clustering/High Availability solutions to failover TSM server environment.</p> <p>High bandwidth connections and data replication technology (for example, IBM PPRC, EMC SRDF) will support synchronous data replication of TSM databases backups and recovery log backups or mirrors. TSM storage pools with critical data can be replicated as well. Extended distances can be achieved by using distance technologies, for example, extended SAN, DWDM, IP/WAN Channel Extenders.</p> <p>Solution may include remote electronic tape vaulting of TSM database and recovery log backups, primary or copy storage pools. Extended distances can be achieved by using distance technologies, for example, extended SAN, DWDM, IP/WAN Channel Extenders.</p> <p>Consider use of TSM server-to-server communications to enable enterprise configuration (multiple TSM servers), enterprise event logging and monitoring, and command routing.</p>



Disaster Recovery Planning

In this chapter, we discuss the demand, concepts, and methodologies for Disaster Recovery Planning (DR Planning) as it relates to IBM Tivoli Storage Manager. In particular, we discuss:

- ▶ The demand for comprehensive Disaster Recovery Planning
- ▶ Disaster Recovery Planning overview
- ▶ Disaster Recovery Planning processes and procedures
- ▶ DRP methodologies, templates for plan creation

3.1 Demand for Comprehensive DR Planning

Worldwide, businesses continually increase their dependence on IT systems for routine business processes. The business processes which directly rely on information systems and the supporting IT infrastructure often require high levels of availability and recovery in the case of an unplanned outage. As a result, the process of business continuity planning must intimately relate business processes to the traditional process of IT disaster recovery. Here we explore the interrelation of IBM Tivoli Storage Manager (TSM) to the processes of Business Continuity Planning and Disaster Recovery.

3.1.1 Business requirements for data recovery and availability

Industries and governments are becoming increasingly accountable for how data is managed, protected, and secured. Policies and regulations vary from industry to industry, and the overall landscape of technical requirements continues to grow in complexity.

The financial industry traditionally leads in terms of stringent regulations for data protection, security, and contingency planning. While some countries or regions still require hardcopy contingency copies of financial data, others are quickly migrating towards a completely electronic format for data. Increasing reliance on electronic data, forms, and processes underscores the importance of integration of enterprise storage management into the Disaster Recovery Planning process.

Recent legislative trends also are driving government organizations and health care providers to meet similar requirements for business continuity and disaster preparedness. In the US, the Health Insurance Portability and Accountability Act (HIPAA) requires the entire health care industry to securely manage and protect patient data through contingency planning and security measures.

Legal systems and regulatory groups also have varying definitions for court-admissible electronic records. In many cases, WORM optical media is the only acceptable format for non-tampered data for court proceedings. Moving data to optical media is not complex, however developing enterprise policies and systems which support legal and technical requirements such as these is increasingly challenging.

Often, entire data center operations have grown through ad-hoc planning processes, with little or no metrics for storage management and enterprise disaster recovery. These trends, along with corporate mergers, acquisitions, consolidations, and globally distributed IT operations have created a myriad of scenarios where no single solution can simply solve the storage management challenge.

3.1.2 Increasing availability requirements

Overall, the increasing dependence on e-mail, electronic messaging, IP services, and cross platform e-commerce applications is changing the way businesses use and rely on information systems. For many companies, the use of e-mail has eclipsed that of voice for corporate communications, customer care, and vendor interactions. This dependency has created much more rigorous demands on operations to ensure availability of services and functional continuity plans. In some cases, data corruption has brought entire corporate e-mail services to a halt due to two to three day recovery and rebuild times.

Application services now pull data from dozens of data sources and the web of consequent dependencies for data continues to increase in complexity. The proliferation of complex Web based applications, messaging, and data management applications is changing the paradigm for what backup/restore has traditionally meant to a business. Now, strategic planning and systems management design is necessary to meet business requirements for availability and recovery.

3.1.3 Storage growth trends

Year by year, storage requirements in enterprise operating environments continue to increase. Typical industry growth rates for disk storage range from 30% to 100% each year.

The ubiquity of relational databases, e-mail systems, and rich media dependent systems (scanned documents, high-quality images, audio, video) all contribute to the growth of storage in data processing and customer environments. Emerging technologies (image recognition, advanced image analysis, wireless applications, smart card, and so on) will only increase the demand for open, scalable, manageable, high performing, and relatively sophisticated storage systems.

3.1.4 Storage management challenges

Mainframe operating environments typically access large pooled arrays of disk and tape, which are managed by a central storage management application. Mainframe storage costs originally prevented decentralization and drove the need for highly efficient use of disk space. In addition to extensive use of tape for backup, these environments often also employ hierarchical storage management (HSM) applications, where infrequently used files are transferred to tape archives to free up disk space. A key aspect of mainframe environments which enables these approaches was the ability to logically partition (LPAR) an operating system environment. A logically partitioned system allows multiple instances of

an operating system to essentially share locally connected hardware and network resources.

The traditional storage techniques for distributed open systems environments included locally attached and assigned disk and tape resources for data storage, backup, and recovery operations. Initially, the operating and capital costs of distributed system environments seemed lower than traditional mainframe environments. This continues to be a debatable issue. Demands for system performance and scalability continue to advance the development of hardware and operating systems technology in the open systems market.

As a result, the original paradigm of locally attached storage resources is shifting rapidly towards a centralized model for disk storage, tape storage, and storage management, not unlike those used in mainframe environments for decades. One of the main reasons for this shift is the prohibitive cost of, and human resource required for, decentralized storage management. Therefore, increasing the amount of disk storage to alleviate the immediate capacity problem only compounds the storage management problem. Storage management costs are also rapidly outstripping the costs of disk and tape storage.

The general movement towards pooled/shared resources for storage also underscores the storage management problems encountered in many open systems environments. It is not uncommon to find a wild array of operating systems, disk devices, tape devices, tape formats, network formats, and storage management tools in place in medium to large data center environments. Through years of ad hoc planning and various technology fads, many large customers now face a daunting challenge for consolidating and simplifying storage management practices.

In the mid-late 1990's many open systems customers adopted a disk-mirroring strategy to solve the business continuity storage problem. The general trend included consolidating storage into large shared storage arrays and mirroring the data locally and/or remotely to similar storage arrays, to ensure continuity of operations from a storage centric perspective. While this style of implementation provides advantages for data backup, primary disk performance, and one level of contingency for application data, it falls short of providing complete strategic business continuity value.

One of the primary limitations of this approach is the fact that only one logical copy of data exists on the primary, secondary, or tertiary mirror at any given time. If the data is corrupted for any of a host of reasons, which is likely in an actual disaster (unstable network connection, abrupt application shutdown, hardware/operating system failure), the mirrored copies could also be rendered useless for recovering operations. In subsequent chapters, we will discuss the importance of policy driven systems and procedures, which in some cases do incorporate these technologies for systems availability.

The selection of storage management tools and devices that comply fully with technology standards and strategic market directions has never been more important. An enterprise wide storage management application is an essential component for developing continuity of operations and enterprise recovery. In Chapter 7, “TSM tools and building blocks for Disaster Recovery” on page 115 we discuss, in detail, how business continuity policy requirements map to technology implementations using TSM.

3.1.5 Networking growth and availability

TCP/IP based networks constitute the majority of open systems networking environments today. Up until the late 1990s, most local area network (LAN) based backup/recovery operations functioned over TCP/IP networks along with production network traffic. As storage management requirements increased along with data volume, many customers deployed locally attached storage management solutions and dedicated TCP/IP networks for backup/recovery operations in order to remedy TCP/IP bottle necking and congestion from backup operations.

The emergence of the Gigabit Ethernet standard (formalized in 1998) supports data transfer rates of 1 Gigabit per second (Gbps). The Gigabit standard opened the door to network attached storage (NAS) devices, which are increasingly popular in file serving and departmental environments. Gigabit Ethernet is also frequently deployed in dedicated networks for backup and recovery operations.

Today, Storage Area Network (SAN) technologies are being quickly adopted to enhance performance, scalability, and flexibility of shared storage resources. The SAN is a dedicated infrastructure for storage I/O, based on widely adopted industry standards for hardware components, Fibre Channel protocol (FCP), and ubiquitous SCSI standards. SANs allow hosts to access storage resources over a network as if they were locally attached. Current Fibre Channel devices support data transfer rates of 2 Gbps and in the future will support up to 10 Gbps rates per connection. Like network connections, multiple Fibre Channel connections can be established between hosts and devices, allowing highly scalable and reliable methods for data transfer. To date, approximately 50% of enterprise scale organizations have deployed a SAN in a production capacity.

Fibre optic distance technologies open a new paradigm of data movement over long distances. Dense wavelength division multiplexing (DWDM) is a technology that allows multiple streams and protocols of data to be combined on one or several long distance fibre optic connections. This means that IP and FCP traffic can be consolidated and routed at high speeds across long distances. Up to 200 billion bits per second (200 Gbps) can be delivered over a single optical fibre. In theoretical terms, this means 100 TB of data could be moved to a remote site in approximately 68 minutes over a single fibre/DWDM connection.

This is why we consider these technologies to be *paradigm shifting*. These technologies are particularly important for designing disaster recovery TSM architectures for multiple site backup/recovery operations. Figure 3-1 lists the range of networking technologies available for implementation, the theoretical maximum performance and the number of minutes taken to transfer varying amounts of data, from 10 GB to 100 TB.

The numbers are calculated as follows. Assume a 1 Gbps Ethernet link. Dividing by 8 to show the number of GBps gives .125. Therefore, 10 GB would be transferred in 80 seconds, or 1.33 minutes. Note that no network actually achieves its theoretical performance — real rates of between 60% and 80% usually are seen more in production networks.

Network Technology	Protocol	Bandwidth	Time Required to Transfer Data Volume (Minutes)				
LAN Network	Protocol	(Gbps)	10 GB	100 GB	1 TB	10 TB	100 TB
10 Mbps Ethernet LAN	TCP/IP	0.01	133.33	1333.33	13333	133333	1333333
16 Mbps Token Ring	TCP/IP	0.016	83.33	833.33	8333	83333	833333
100 Mbps Ethernet LAN	TCP/IP	0.1	13.33	133.33	1333	13333	133333
Gigabit Ethernet	TCP/IP	1	1.33	13.33	133	1333	13333
Storage Area Network	FCP (SCSI-3)	2	0.67	6.67	67	667	6667
SAN (Planned Capacity)	FCP (SCSI-3)	10	0.13	1.33	13	133	1333
WAN/ Optical Network	Protocol	(Gbps)	10 GB	100 GB	1 TB	10 TB	100 TB
T1	TCP/IP	0.0015	888.89	8888.89	88889	888889	8888889
T3	TCP/IP	0.0447	29.83	298.28	2983	29828	298285
OC-1 / STS1	ATM	0.0518	25.74	257.40	2574	25740	257400
OC-3 / STS3	ATM	0.1552	8.59	85.91	859	8591	85911
OC-12 / STS12	ATM	0.6221	2.14	21.43	214	2143	21433
OC-48	ATM	2.488	0.54	5.36	54	536	5359
OC-96	ATM	4.976	0.27	2.68	27	268	2680
OC-192	ATM	10	0.13	1.33	13	133	1333
OC-255	ATM	13.21	0.10	1.01	10	101	1009
SAN + DWDM	FCP+DWDM	200	0.01	0.07	1	7	67

Figure 3-1 Networking technologies and transfer rates

An important factor in this data is that each set of calculations is based on a single element or connection per network type. These technologies can be implemented in multiples. For most networking technologies, distance, signal loss, hop count, and line conditions can degrade performance capacities.

The use of optical fibre is quickly evolving for long-distance data communications. *Dark fibre*, or dormant fibre available for purchase or lease, ranges in availability from region to region. While fibre is abundantly available through competitive providers in New York City or Berlin, other geographies may lack the service altogether. Costs fluctuate greatly too, ranging from \$1000+USD/month/fibre in remote continental locations, to \$50USD/month/fibre in certain European markets. Customers also can choose to create wavelength management services in house, or out source these services to a variety of Metropolitan Area Network (MAN) providers, who specialize in wavelength management services and provisioning. We acknowledge that the use of dark fibre and DWDM is an emerging and dynamic market, and we see current and future potential for using these technologies for site to site movement of data for disaster recovery.

In later chapters we will explore the integration of these technologies into TSM architectural planning scenarios.

3.1.6 Capacity planning trends

Continuity of operations depends on having access to the data along with the ability to recover the data in the event of a disaster. As storage consumption and growth continues, businesses depend more and more on accurate forecasting and capacity planning. Poor planning or lack of planning for storage often results in surprising halts to operations (unplanned outages), due to storage or network resource over-consumption. Some platforms and technologies support dynamic expansion of resources, while others do not.

TSM planning and policy development, if properly done, will make this valuable data accessible at the host and enterprise level. Understanding the enterprise landscape, and creating policies to achieve business continuity, disaster recovery, and performance goals can all be achieved through a methodical planning process, which we discuss in greater detail in Chapter 7, “TSM tools and building blocks for Disaster Recovery” on page 115.

3.2 Business Continuity, Disaster Recovery, and TSM

Business requirements for availability and continuity of IT operations is driving Business Continuity Planning (BCP) and Disaster Recovery Planning (DR Planning) to new levels of interdependence. Storage management processes and infrastructure impact not only the recovery times for operations but also the overall continuity of operations. Backup procedures often affect availability and restore procedures depend on carefully documented and planned procedures.

TSM planning and implementation procedures relate to both the BCP and DRP process.

3.2.1 Business Continuity Planning

Business Continuity Planning (BCP) is an enterprise wide planning process which creates detailed procedures to be used in the case of a disaster. Business Continuity Plans take into account processes, people, facilities, systems, and external elements which could alter the ability of a organization to function in a variety of disaster scenarios. The overall BCP objective aims to maintain critical business processes in the event of an unplanned outage.

The scope of a Business Continuity Plan will most certainly involve the IT infrastructure. Most Business Continuity professionals view Disaster Recovery Planning as an IT-centric logical subset of the Business Continuity Planning process.

The focus of this redbook is to emphasize critical elements of BCP and DRP as they relate to the use of TSM in an enterprise environment. For more comprehensive background information on Business Continuity Planning, please refer to *IBM TotalStorage Solutions for Disaster Recovery*, SG24-6457.

3.2.2 Disaster Recovery Planning

Disaster Recovery Planning (DR Planning) functions as a logical subset to the Business Continuity Planning (BCP) process. Traditionally, IT operation's teams manage the DRP process to ensure continuity of operations in the event of a wide variety of disaster scenarios.

Since system availability requirements continue to increase, organizations must now view DRP as an ongoing practice, not unlike other core business processes. Traditionally, IT operations handles DRP, while BCP functions as a closely coupled process. The published DRP is typically an IT focused plan, which is designed to provide continuity of operations for applications, databases, system, networks, telephony, staff, and supporting infrastructure (power, cooling, space).

Generally, the planning process follows a method based approach to analyzing business requirements, business impacts, risk, and the IT infrastructure to establish a formal action plan.

There are several effective methods and approaches to delivering a well designed plan, but most planning procedures follow the general format described here:

1. Project Initiation and Team Selection

Project Initiation begins with executive/management leadership, the selection of a project leader, planning team selection, and the implementation team selection. From the onset of the planning process, information collection methods (forms, templates, questionnaires, and so on) are established along with clear expectations of individual roles in the project.

2. Business Process Analysis

The Business Process Analysis entails the collection of information about business processes, their technology infrastructure supports/dependencies, and techniques/strategies used by similar organizations to mitigate risk. In this stage, high level business processes, regulatory requirements, and enterprise wide availability requirements are defined and documented.

3. Risk Analysis/BIA

The Risk Analysis phase entails the processing of collected data to arrive at DRP targets and objectives. Basically, the outcome of this analysis is an understanding of what systems or processes are at risk and what resources are required for recovery within an acceptable time frame. The risk analysis phase is a critical step in the collection of data for storage management policy creation.

The Business Impact Analysis (BIA), illustrated in Figure 3-2, is typically performed by a Business Continuity Planning team. The BIA enables the disaster recovery group to fully characterize the system requirements, processes, and interdependencies. It is important that the DRP team have access to this information for subsequent planning and policy creation phases.

The Business Impact Analysis (BIA) identifies the financial and non-financial value of a business process and the supporting application environments. This information can be used to classify systems and create policies and strategic plans to support the business requirements for availability and reliability. The BIA is a critical element for balancing the cost of supporting infrastructure versus cost of downtime.

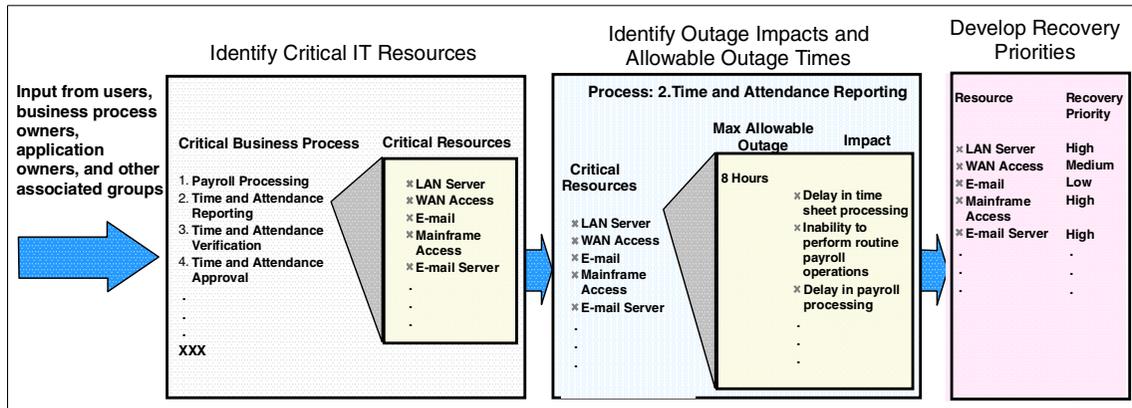


Figure 3-2 Example of the Business Impact Analysis process

Another important aspect of the risk analysis phase is general risk assessment. Not only does the definition of credible risks help to create realistic contingency planning scenarios contained in the DRP, it also helps an organization create strategic policy and infrastructure decisions which will help them to avoid disasters of an operational nature. These are intrinsic benefits of DR and storage management planning.

4. BIA/RPO/RTO Analysis

The BIA/RPO/RTO analysis stage focuses on organizing business requirements in such a way that systems and data can be classified by relative importance. The systems classification is translated into technical classifications and policies. While these elements are often included in the risk analysis phase, we separate them to stress the importance of strategic planning and classification of systems.

In addition to the BIA data, the following elements must be fully understood for each system:

Recovery Point Objective (RPO) defines the amount of data that you can afford to recreate during a recovery, by determining the most recent point in time for data recovery.

Recovery Time Objective (RTO) is the time needed to recover from a disaster, or how long the business can survive without the systems.

Network Recovery Objective (NRO) details the time to recover or failover network operations. Keep in mind that the systems level recovery is not fully complete if customers cannot access the application services via network connections. For instance, in a WAN environment where data processing services are transitioned from site A to the recovery site B, numerous problems can exist outside of the local network configuration, including DNS

updates, DNS configuration problems, IP name conflicts, subnetting issues, and ISP failures. Comprehensive network failover planning is of equal importance to data recovery in a Disaster Recovery scenario.

5. Data Protection/Policy Creation

The Data Protection/Policy Creation phase is a linchpin for effective Disaster Recovery Planning with a storage management focus. At this stage, business process, risk, BIA, RTO, RPO, and NRO data are collectively synthesized into actual policy, which is later implemented in TSM. The policies created directly impact the TSM project scope, architectural design, and capacity planning.

In general, policy creation is a formal procedure to classify and safeguard digital information assets.

6. Recovery Plans

The Recovery Plans phase involves the formulation of strategies for replacing systems, networks, and end users in the wake of an unplanned interruption event. Every organization differs, so there is no standard approach to effectively create a DRP.

We propose that customers use a data centric approach to create policies for TSM data management and DRP. This type of plan is based on tiered systems classifications and supporting infrastructure. Systems are classified according to relative importance to the organization (BIA) and the time needed to recover the systems (RTO). The RPO translates directly to the frequency of backups and offsite vaulting policies. With a basic method driven approach, the enterprise would have recovery operations based on business logic, clear procedural documentation, and a clear plan for recovery of operations. Section 6.2, “Mapping DR requirements to TSM configurations” on page 104 discusses how this process maps to a TSM implementation.

7. Training and Testing

Training and Plan Testing validate the disaster recovery strategies that have been developed and provide a vehicle for identifying plan deficits, so they can be resolved. Testing also improves skills sets and can provide an opportunity for cross training within an organization. Chapter 4, “Disaster Recovery Plan testing and maintenance” on page 63, describes the DR testing process, in detail.

8. Change Management

Change management provides a mechanism for keeping the plan up to date with changing business and IT parameters. Change management provides feedback that is used to keep the disaster recovery capability current at all times. Section 4.3, “DRP maintenance procedures” on page 71, describes the DRP change management process, in detail.

The general Disaster Recovery Planning process is illustrated in Figure 3-3.

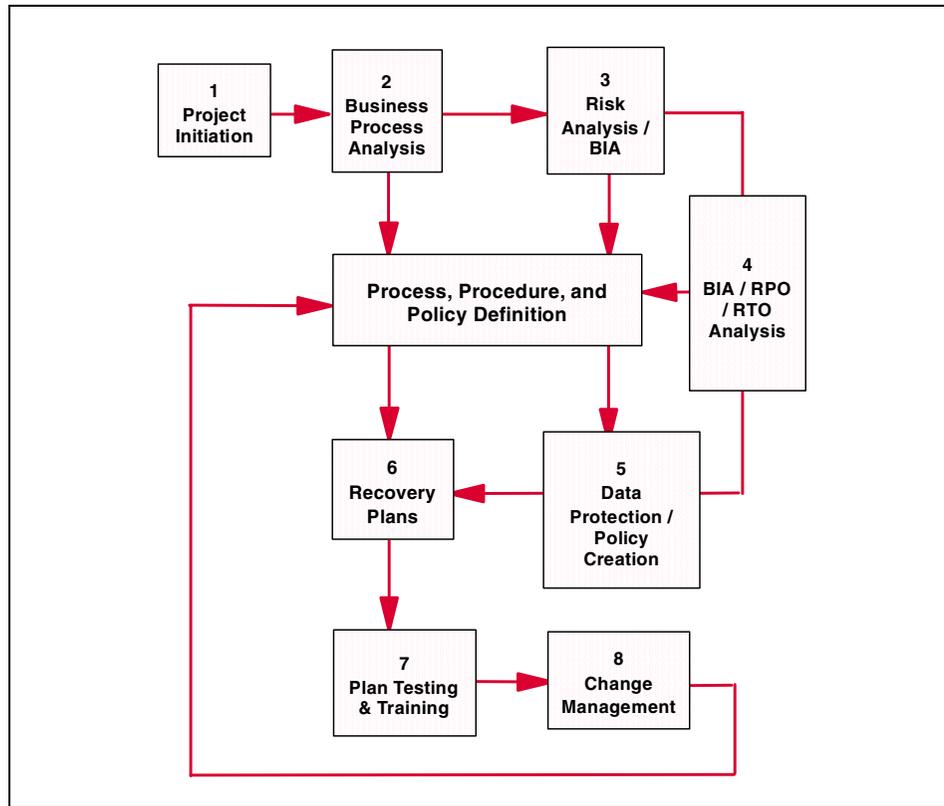


Figure 3-3 Disaster Recovery Planning process

3.3 Disaster Recovery Planning and TSM planning

Disaster Recovery Planning and storage management planning function as separate processes in many environments. We see a synergy between these processes and illustrate the relationships in the following sections.

3.3.1 How BIA and RTO relate to TSM policy

As illustrated in Figure 3-4, the traditional output from Business Continuity Planning provides detailed information about critical systems, their supporting systems, the value of each system to the business, Risk Analysis, and the Recovery Time Objective for each system. At various levels, these concepts relate to storage management planning and TSM planning. Since BCP, DRP, and TSM policy definitions accommodate a wide variety of business and technical

needs, a data centric planning approach is necessary for a well defined set of policies.

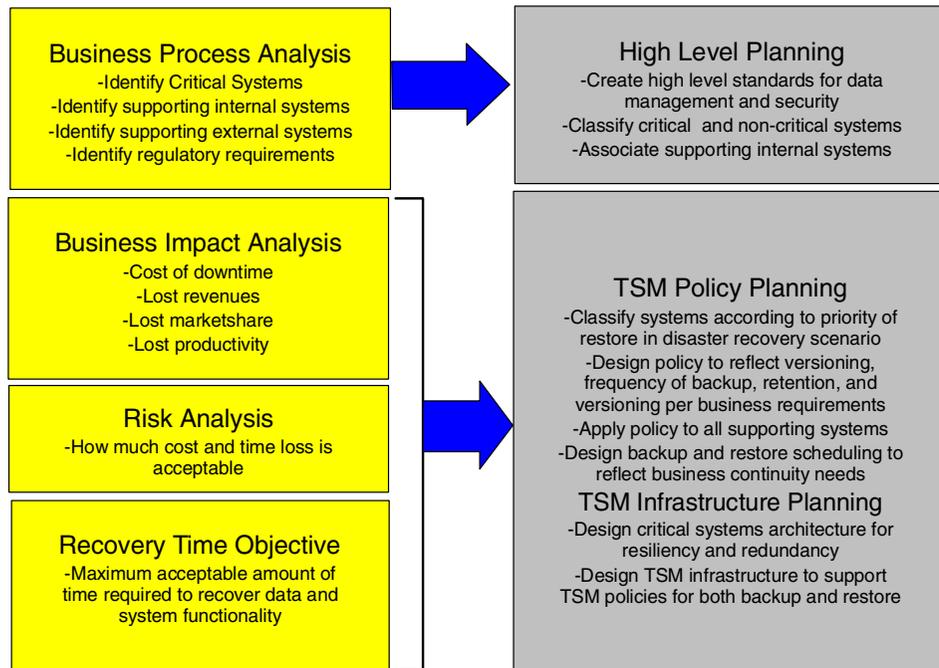


Figure 3-4 The relationship between BIA, RTO, and TSM planning

The Business Process Analysis identifies critical processes for the Business Impact Analysis (BIA). From the BIA, an application environment will usually be assigned a cost per hour value, which directly impacts continuity plans and the Recovery Time Objective (RTO). Essentially, the value of the data and the access to the data directly correlates to policies and infrastructure decisions for backup and recovery.

3.3.2 Dependencies between systems

Once a critical business process is identified, the systems from which it obtains data must also be identified and classified as part of the *supporting infrastructure*. As illustrated in Figure 3-5, a critical application environment may have multiple sources for data, in addition to core customer input. A central application may source data from other applications, messaging systems, EDI systems, FTP updates, e-mail processes, batch updates from legacy systems, and Web-based customer interfaces. Understanding the flow of data within the

enterprise is critical to designing an effective storage management infrastructure for Disaster Recovery.

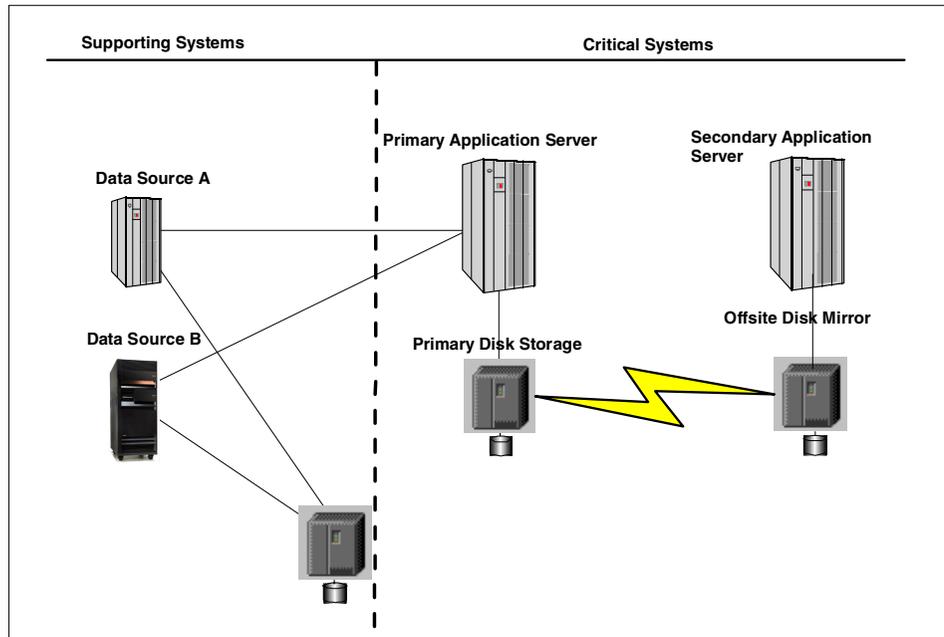


Figure 3-5 Basic example of enterprise data dependencies

Once critical and supporting systems are identified, general policies for storage management can be applied. The rationale behind this approach stems from the principle that an application environment is only as effective as its weakest component. If a supporting system which sends critical data to a critical system fails, the overall critical business process may be compromised. These situations are easily overlooked in enterprise planning efforts due to the increasing complexity of data processing environments. An additional analysis of data within each system can then be used to classify data based on restore priority.

3.3.3 Data classification procedures

Data classification is an important component of DRP, storage management planning, and TSM policy planning. Within every operating environment, various types of data exist. While progressive (incremental) backup methods can allow frequent backups with minimal resource consumption, effective restore operations depend on the efficient use of system resources. System data can be classified into four tiers, as shown in Table 3-1.

Table 3-1 Data classification descriptions

Data Classification	Description
Critical	Application data which is critical for business processes to provide a minimum acceptable level of service in the event of a disaster. Also, data which must be available for regulatory audits. Some examples are customer and company financial data in the financial industry or electronic patient records in the health care industry.
Important	Application data needed for standard business processes, which is impossible or extremely expensive to recreate. Data which is not needed for minimal critical operations, but has significant operating value (secret and classified data included).
Semi-Important	Application data which is needed for normal operational procedures, but is possible and cost effective to recreate from original data sources at minimal to moderate costs. An example may be switch call detail record data which often resides on remote cellular tele-communication switches for 3-5 days.
Non-Critical	Non-secret or non-classified data which can easily be recreated from original source data. An example could be data mart query logs/tables, which could be easily regenerated once the original data mart database is restored.

In a DR event, efficiently restoring an enterprise environment depends on methodical planning and classifications of systems, data, and resources. By classifying data, restore procedures (as documented in the DRP and implemented in TSM policy) will follow an orderly series of events to restore the mission-critical servers, applications and data first. From an enterprise view or a single system view, this means that critical application data will be treated with appropriate priority during DR operations. This approach also limits resource consumption and accelerates recovery procedures.

3.3.4 Relating DR Planning to TSM planning

Disaster Recovery Planning and TSM planning are both customer specific and methodology driven processes. Through the BIA process of classifying systems, determining dependencies, and classifying data, the information should then be organized into a format which complements both the DRP and the TSM policy documentation. Starting with the business impact analysis, systems can be classified and grouped according to business value. Appendix A, “DR and Business Impact Analysis Planning Templates” on page 331 includes a general outline for the BIA process. Once the BIA is completed for critical business processes, this information can be organized into a planning worksheet as shown in Figure 3-6.

In Figure 3-6, the RTO and NRO are usually the same, because customer access to the data is just as important as the system restoration. The scope of network recovery procedures can center on critical systems environments and spread to the enterprise, depending on network architectures in place. The RPO depends strictly on the cost of recreating application data and can vary from system to system. In general, this kind of planning exercise aids in the organization of system priorities, storage management policies, and DR plans.

Business Impact Analysis Sample Worksheet							
		(Usually measured in hours)					
<u>Application</u>	<u>Cost Per Hour</u>	<u>RTO</u>	<u>NRO</u>	<u>RPO</u>	<u>System Resources</u>	<u>Recovery Priority</u>	<u>Outage Impact</u>
Application A	200000	2	2	1	Server / OS	A	\$400,000
(Dependencies: Application G, X)					Network / SAN	B	
					Disk Resources	C	
					TSM Client Software	D	
					Application Data "A"	E	
					Application Data "B"	F	
Application G	30000	2	2	5	Server / OS	A	\$60,000
					Network / SAN	B	
					Disk Resources	C	
					TSM Client Software	D	
					Application Data "A"	E	
Application X	65000	2	2	3	Server / OS	A	\$130,000
					Network / SAN	B	
					Disk Resources	C	
					TSM Client Software	D	
					Application Data "A"	E	
Application B	35000	3	3	2	Server / OS	B	\$105,000
(Dependencies: Application Z)					Network / SAN	C	
					Disk Resources	D	
					TSM Client Software	E	
					Application Data "A"	F	
					Application Data "B"	G	
					Application Data "B"	H	
Application Z	10000	3	3	2	Server / OS	B	\$30,000
					Network / SAN	C	
					Disk Resources	D	
					TSM Client Software	E	
					Application Data "A"	F	
					Application Data "B"	G	
					Application Data "B"	H	

Figure 3-6 Sample Business Impact Analysis Worksheet

3.4 DRP development

The objective of a Disaster Recovery Plan is to coordinate the recovery of the IT infrastructure through plans, procedures, people, and technical assets. Recovery plans can include alternate sites, equipment, and even staff.

A DRP targets an audience which includes, but is not limited to:

- ▶ Executive management
- ▶ Operations management
- ▶ IT security teams
- ▶ System Architects and Engineers
- ▶ System Administrators
- ▶ Customers who use IT services
- ▶ Application/Database Administrators

Since the audience includes a wide array of talent, the plan language needs to be concise and accessible to technical and non-technical readers. A well written plan provides a roadmap to IT recovery for current or replacement IT staff. After all, non-standard staff could be partly or fully responsible for IT recovery in a disaster scenario.

Since every IT environment is unique, a DRP must be built from a thorough and site specific planning process. A balance must be struck between technical detail and plan flexibility, to ensure a concise, functional, and scalable plan for the enterprise. A DRP outlines team roles, responsibilities, and specific procedures for restoring an environment during an unplanned outage. A DRP can be used on several scales, ranging from system specific outages, to partial site outages, to massive site failures. A well designed DRP supports business continuity for a variety of outage situations.

A DRP follows a basic format, as illustrated in Figure 3-7.

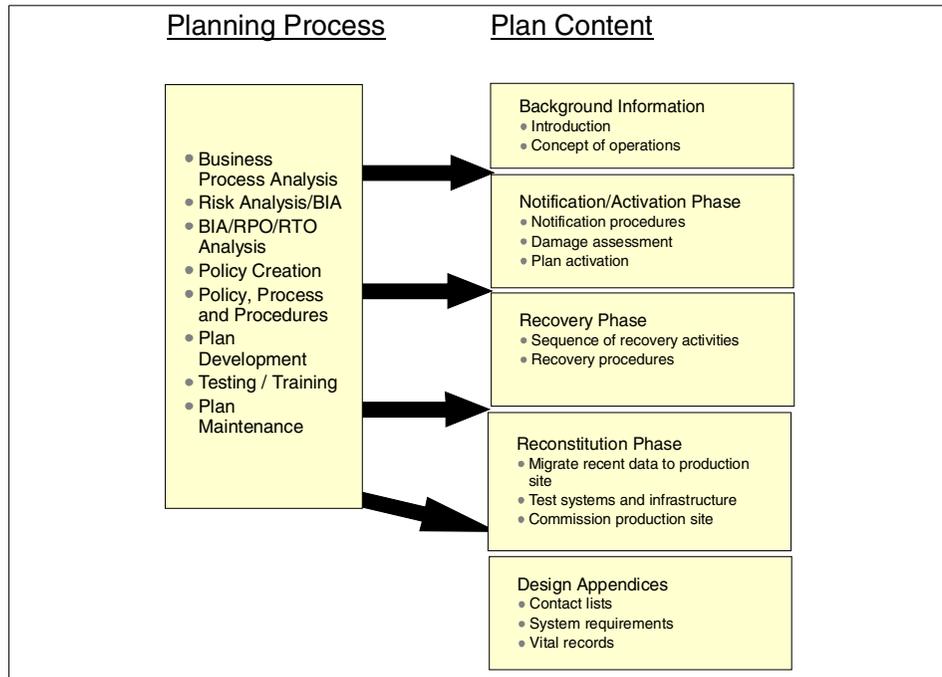


Figure 3-7 Disaster Recovery Planning processes and plan content

The DRP organizes the data from the BCP and DR Planning processes into an action plan for IT recovery. The five core sections of the Disaster Recovery Plan are described in detail in the following sections. We have also included a sample IT Disaster Recovery Plan template, which is located in “Disaster Recovery Plan Template” on page 332.

3.4.1 Background information

The background information section explains the purpose and scope of the DRP and how it relates to the operational environment. The background information subsections can include an *introduction* and *concept of operations*, as specified in the following descriptions.

Introduction

An introduction has a purpose, scope, and commitment statement:

- ▶ **Purpose:** An explanation of the DRP development, operational needs for availability and recovery, and the overall plan objective.
- ▶ **Scope:** Identifies what the plan covers and does not cover. The plan design and intent is related to specific situations and scenarios as addressed in this

section. Also, key assumptions are outlined, such as availability of staff, facilities, external vendors, among other variables.

- ▶ **Commitment statement:** The DRP should be introduced as a living and dynamic process and procedure. Specific elements for plan maintenance and testing are covered extensively in Chapter 4, “Disaster Recovery Plan testing and maintenance” on page 63.

Concept of operations

The concept of operations provides a clear description of what infrastructure exists, how the DR/operations teams are organized, and how the recovery plan functions through various disaster scenarios. The concept of operations subsections can include:

- ▶ **Environment overview:** Provides a high-level overview of IT operations with written and graphic explanations of where systems reside, where production and recovery facilities reside, and where disaster recovery staff work in normal and recovery operations.
- ▶ **System descriptions:** Provides an architectural view and description of critical infrastructure (including servers, storage devices, backup/recovery systems, networks, Storage Area Networks, firewalls, and telecommunications/ISP connections) and a general written description of “how things work” in IT operations.
- ▶ **Recovery scenarios:** Describes how the plan functions in a variety of unplanned outage events, ranging from entire site loss to single systems failure. This section is an extremely critical component for establishing expectations for how the recovery plan will function in a variety of scenarios. There is no way to plan perfectly for a disaster, however understanding the plan design and testing the plan in a variety of scenarios is the best way to achieve full disaster preparedness.
- ▶ **Responsibilities:** Outlines team member roles and responsibilities. Usually an organization chart shows the team hierarchy to establish rules for decision making, succession, escalation and replacement in the event of staff loss. Greater detail is then assigned to specific roles (typically organized by position name instead of personal name) for Disaster Recovery Plan activation.

3.4.2 Notification/activation phase

The notification/activation phase describes the communications and decision making process during or immediately after a disaster. The components of this section include notification procedures, damage assessment, and plan activation.

Notification procedures

Notification procedures provide communication plans in the event of a disaster. First the damage assessment team is notified of an event, and depending on the initial analysis, additional DR team members are notified of plan activation and the overall scope of recovery operations. Multiple notification procedures (cell phone, land-line, pager, e-mail, text messaging, manual procedures) should be built into the overall plan. There are no guarantees that any particular communication methods will function in the event of a given disaster. A detailed call list and organization chart is critical to the notification process. Notification procedures may include injury/casualty reporting, incident reporting, damage estimates, response/recovery procedures, stock market notifications, contact information for family members (such as 800 numbers or Web site information broadcasted via various media), instructions for relocation, and instructions for DR operations teams.

Damage assessment

The damage assessment is typically performed by the damage assessment team and follows a critical action plan to determine the nature of and extent of site damage. Assuming the assessment team is not at risk of injury, the following elements are generally included in the assessment scope:

- ▶ Cause of disruption
- ▶ Potential risks for additional disruptions or damage
- ▶ Scope of damage and disruption
- ▶ Physical infrastructure assessment (including structural integrity, power, cooling, heating, ventilation, fire-suppression, telecommunications, and HVAC)
- ▶ Functional status of equipment (fully functional, partially functional, nonfunctional)
- ▶ Type of damage to IT equipment and media (including water damage, fire and heat, physical impact, electrical surge, electromagnetic pulse, and so on)

The damage assessment impacts directly the extent to which the DRP is implemented. Scenarios range from system specific outages (such as, hardware failure, virus, or hacking) to site wide disaster situations, which would require the full implementation of the Disaster Recovery Plan. Depending on the scope of damage and the consequent recovery operations, the appropriate teams and resources are notified according to the documented procedures.

Plan activation

The plan activation depends on the plan activation criteria, which is a set of organization-specific metrics for decision making in an unplanned outage event.

Variables for plan activation may include, but are not limited to staff safety, damage to facilities, extent of systems damage, extent of damage to critical infrastructure/systems, duration of disruption, and the likelihood of additional disruptions. The activation of the DRP is generally made by the Disaster Recovery team lead.

3.4.3 Recovery phase

Once the DRP is activated, various recovery teams and procedures take effect. Manual processing measures are typically activated, while IT systems are recovered according to the extent of plan activation. Backup system architecture and planning plays a critical role in the recovery phase as systems are rebuilt, applications are restored, data is recovered, and systems are put back online into production. Data recovery operations must work synergistically with overall DR operations. At the completion of the recovery phase, systems will be functioning to the extent determined in the plan. Beyond the critical recovery phases, less critical recovery operations may ensue.

Two critical phases comprise the recovery phase: sequence of recovery activities and recovery procedures.

Sequence of recovery activities

The BIA, systems classifications, and data classifications play an integral role in defining the sequential order of recovery events. In an enterprise environment, many individuals and teams will be working in parallel to recover infrastructure and systems, so the logical planning and coordination of events, resources, and team communications is essential.

Critical systems and data are recovered based on priority classifications, however infrastructure and resources to support the systems must be recovered in a logical sequence to support the overall recovery. High level procedural documentation can outline critical steps for replacing infrastructure (people, power, cooling, hardware, software, operating system, network, storage), which would then support focused data recovery operations. The Network Recovery Objective (NRO) must relate closely to the steps which depend on network resources (LAN communications, operating system bare metal restore, data recovery, application services, and so on). Project management software tools can be used to model and plan parallel events and resource utilization for a large scale recovery operation.

If systems are being recovered at an alternate site, some or all of the following components must be either already available at, or delivered to the recovery site: backup tapes, hardware, software, software licenses, recovery plans, staff, and even food/water supplies. Such activities are site and plan specific, but careful planning and preparation simplifies the movement of resources from one site to

another. Remember that in a disaster event, even the basic modes of transportation could be compromised.

Third party vendor procedures should also be considered in the sequence of recovery events. Plans for facilities, additional hardware, additional staff, external networks, and other various components must complement the overall recovery sequence.

Recovery procedures

Recovery procedures provide detailed procedures to restore systems and supporting infrastructure components. Assuming general facilities and infrastructure is in place or restored, recovery procedures generally target specific recovery team members and address the following broad guidelines:

- ▶ Installing hardware components
- ▶ Recovery or re-install of operating system images/backups
- ▶ Configuring network resources
- ▶ Restoring/configuring system data
- ▶ Restoring application software
- ▶ Restoring application data
- ▶ Testing system functionality and security controls
- ▶ Connecting system to production network
- ▶ Testing
- ▶ Administering and monitoring replacement systems

Once all systems are restored, tested, and functional, the Disaster Recovery team transitions back to normal operational procedures. This will continue while further decisions are made, about if or when a roll-back to the original or replacement site is necessary.

3.4.4 Reconstitution phase

Depending on the event and the extent of damage to a primary facility, systems will either be restored to the primary facility, or transitioned to a replacement primary facility for standard operations. TSM functions well in a large site recovery and reconstitution phase, because the most recent versions of backed up data created in a secondary site environment can be migrated and recovered to the original site environment using standard TSM restore procedures.

Systems reconstitution generally involves these steps:

- ▶ Audit infrastructure support functions (power, water, security, telecommunications, environmental controls, equipment, supplies).
- ▶ Install system hardware, software, firmware.
- ▶ Configure and test LAN, WAN, and SAN configurations.

- ▶ Test system operations.
- ▶ Restore all current operational data to production systems.
- ▶ Test, validate, and start primary production systems.
- ▶ Shutdown and terminate contingency operations.

Once the infrastructure, systems, application, network, and customer interfaces are tested for production, the primary or replacement site reverts to a state of normal operations, and the original DRP and contingency measures are maintained.

3.4.5 Design appendices

All systems specific recovery procedures, BIA documents, systems and data classification worksheets, TSM policy documentation, transportation plans, contact lists, vendor contact information, vendor SLA's, equipment configuration inventories, and infrastructure recovery plans can be easily organized or referenced in the plan appendices.

3.5 Related planning considerations

For every enterprise organization, Disaster Recovery Planning will depend on external factors for success. Components of a Disaster Recovery Plan may be out sourced or in some way linked to external organizations for an additional level of contingency.

3.5.1 Insourcing versus outsourcing

The fundamental responsibility for business continuity lies within any organization. Various components of a Business Continuity Plan, including the Disaster Recovery Plan, can be out sourced to third-party consultants or companies. Companies such as IBM and Sunguard provide DR services, ranging from equipment replacement contracts, mobile site services, facilities contracts, Business Continuity Planning consulting services, to full disaster recovery infrastructure and consulting services. A trade-off exists between insourcing and outsourcing Disaster Recovery Planning responsibilities.

On one hand, insourcing BCP/DRP can bind the most detailed knowledge of business processes and supporting infrastructure to an in-house plan, because the full time employees who are ultimately responsible for the plan generation, maintenance, and testing will have years of experience within the environment. Since much of the planning process is based on the discovery of existing process and resources, this strategy offers a distinct advantage over outsourcing

strategies. Many organizations employ specialized software planning tools (such as Strohl Systems LDPRS application) to create the Business Continuity Plan. Factors weighing against insourcing might include lack of technical depth in specialized technologies, the cost of additional resources and assets, and the lack of focus and commitment that sometimes plagues in-house programs for business continuity.

Outsourcing, on the other hand, offers customers the ability to off load several tasks to specialized service providers. Offsite tape vaulting, facilities, workspace, telephone, and hardware contracts comprise the majority of disaster recovery services available today. Risk assessments, threat assessments, and business process analysis studies can also be out sourced components of a Business Continuity Plan.

Outsourcing companies sometimes provide business continuity planning tools and specialists, along with the core service and contract offerings. Outsource providers help organizations navigate the complex BCP/DRP process, however some caution should be taken in placing too much reliance in a product or service-contract-centric plan provided by an external organization. Many companies are now realizing that a facilities/hardware/tape vaulting contract alone will no longer meet business requirements for IT availability and recovery. Developing a storage management strategy which closely relates to systems management, availability, and the BCP/DRP processes is essential in today's enterprise.

If contracting for the DR site with a commercial vendor, adequate testing time, work space, security requirements, hardware requirements, telecommunications requirements, support services, and recovery days (how long the organization can occupy the space during the recovery period) must be negotiated and clearly stated in the contract.

Customers should be aware that multiple organizations may contract with a vendor for the same alternate site; as a result, the site may be unable to accommodate all of the customers if a disaster affects enough of those customers simultaneously. The vendor's policy on how this situation should be addressed and how priority status is determined should be negotiated.

Depending on your organization, one strategy or a mix of both insourcing and outsourcing may provide the best level of protection in a DRP. Before pursuing either route, the organization strategy, an internal skills assessment, and a cost benefit analysis for both options should be taken into consideration.

3.5.2 Offsite tape vaulting

As a standard, backup and archive tapes should be taken offsite (either physically or electronically) for secure and safe storage. Customers have a choice to develop in-house or outsourced methods for tape vaulting and storage. Either way, several factors must be taken into consideration for secure off-site vaulting.

- ▶ **Proximity:** Tapes must be stored far enough away from the production site to avoid exposure to risk. For instance, a data center that sits near an area prone to seismic activity or tsunamis, should of necessity store backup tapes in a less vulnerable geographic location.
- ▶ **Access time:** Tapes must be accessible within a time frame which supports RTO objectives for critical systems. In a medium-to-large scale disaster scenario, many variables (including traffic congestion, martial law, and lack of air or road transportation services) could significantly hamper the simple movement of tapes from one location to another. The location of offsite tape storage is a strategic element for recovery.
- ▶ **Environment and Security:** A physical offsite storage location should provide data center class environmental and security controls. Data should be treated as a high security asset in a well protected and conditioned environment. Considerations include temperature, humidity, access control lists (and backup lists), check-in/check-out procedures, fire and flood prevention, and power management.

In defiance of logic, there are dozens of high-cost backup/recovery systems, which send tapes offsite in the trunk of a regular motor vehicle. Tape vaulting is often an afterthought and a cost-saving opportunity, however we stress the importance of designing or using a well built tape vaulting service. TSM architecture supports offsite volume management and a rigorous schedule for offsite tape movement.



Disaster Recovery Plan testing and maintenance

Once a Disaster Recovery Plan is developed, it is important to establish and maintain the functionality and validity of the plan. This chapter discusses methodologies for Disaster Recovery testing, frequency and approaches for testing and how to maintain the Disaster Recovery Plan.

Creating a Disaster Recovery Plan is not enough — it must be carefully and regularly tested to ensure that it works and that the people responsible for executing it know how to follow it.

The overall business structure of an enterprise will remain relatively stable over a period of time. A Disaster Recovery Plan is a vital element for an enterprise to describe how the continuity of the business processes will be preserved in case of a disaster. The technical details and the human resources of a business requirement typically change more frequently. An update process for the Disaster Recovery Plan is necessary, so its functionality and effectiveness is preserved.

4.1 How to provide Disaster Recovery testing

Regular testing is vital for maintaining an effective Disaster Recovery Plan (DRP). Therefore the following steps are required:

1. Develop a test schedule with a pre-defined set of test scenarios. This ensures that each element of the Disaster Recovery Plan will be tested to confirm the accuracy of individual recovery procedures and the overall effectiveness of the plan.
2. Fully document each activity during the particular recovery test. This will include all steps taken, any errors found, timings recorded. It is important to have a plan which not only reaches the Recovery Point Objective but also the Recovery Time Objective. If a particular restore operation takes a longer time than planned, or, if a procedure does not work as expected, an action to fix this is required.
3. Review the results of each test and initiate necessary corrections to the DRP. Each test (whether successful or unsuccessful) is a great source of information. Test results and lessons learned should be documented and reviewed by all participating technical staff. Information collected during the test itself and the post-test reviews that improve plan effectiveness must be incorporated into the DRP.

Ensure that all technical staff dealing with IT production take part in regular testing. You must consider the possibility that some or all of the key technical personnel may not be available during an actual disaster (for example, because of vacation, injury, or other personal circumstance). Ideally, the testing should be done by staff not specifically experienced in the platforms involved. This will expose any “missing steps” in procedures which an expert would automatically know how to do. We recommend that you have backups identified and assigned for all the critical members of the DR team.

The details of the testing plan will depend on the tier level of your DR solution. If you do not have a hot site and only Tier 1 level of recovery, the testing plan is very simple. You have base disaster recovery capability, you keep your vital data offsite and you will establish the recovery on appropriate hardware. On the other hand if you have invested in a dedicated remote hot site (Tier 6 level) the testing and planning will be much more complex.

Remember, your success in being able to recover in a real disaster depends on the quality of your DR Plan, and your capability to execute it. Testing is really the key to validating and enhancing your Disaster Recovery Plan, as well as giving your Disaster Recovery team the experience and confidence to execute it.

4.1.1 How is DR testing different from a real disaster

Obviously, testing a DRP is different to executing it in a real disaster situation. A disaster test is scheduled and the start and end times are known. Because of this, applications can be shut down in an orderly manner before starting the test. There will be a definite roll-back (return to the original equipment) at the end of the test period. There are few uncontrolled factors occurring during a DRP test. By contrast, a real disaster could occur at any time, and typically is accompanied by some (or possibly many) factors outside of the control of the DRP. The roll-back may be to the same site (if it was not destroyed in the disaster), or to a totally new site with all new equipment. The time taken for this could be indeterminate.

Nevertheless, an attempt should be made as far as possible to emulate the conditions of a real disaster when testing the DRP. For comparing DR testing with a real disaster we will consider two cases. The first case, shown in Table 4-1, covers Tiers 1 through 3 and the second case (Table 4-2) is for Tiers 4 through 6, where mirroring is available.

Table 4-1 Differences between DR testing and real disaster for Tiers 1-3

DR testing	Real disaster
<ul style="list-style-type: none">▶ Databases and applications are backed up as of the most current status, after restore on backup hardware daily operation can continue.	<ul style="list-style-type: none">▶ After restore of databases, administrator has to apply logs to bring the databases to the most current state after a disaster.
<ul style="list-style-type: none">▶ No lost transactions.	<ul style="list-style-type: none">▶ Users must check for last finished transactions and eventually manually insert lost transactions.
<ul style="list-style-type: none">▶ All IT staff is available — each recovery procedure can be done by responsible person (administrator), who is knowledgeable about the procedure.	<ul style="list-style-type: none">▶ Because not all IT staff may be available, some of the activities have to be done by less experienced people (for example, a Windows administrator has to restore a UNIX machine).
<ul style="list-style-type: none">▶ People under less stress than a real disaster.	<ul style="list-style-type: none">▶ People under stress can make more mistakes.

DR testing	Real disaster
<ul style="list-style-type: none"> ▶ Usually for testing purposes, it is sufficient just to complete the restore of data and applications — it is not necessary to stay on the backup site and plan roll-back. 	<ul style="list-style-type: none"> ▶ Roll-back plan needed when replacement hardware for the primary site is available.

Table 4-2 Differences between DR testing and real disaster for Tiers 4-6

DR testing	Real disaster
<ul style="list-style-type: none"> ▶ Systems, databases, and applications are shutdown correctly and in the correct order. ▶ You don't have to pay attention to database integrity. 	<ul style="list-style-type: none"> ▶ Database and application administrators must run a check for database and application integrity (while databases and applications are mirrored using hardware remote mirroring techniques).
<ul style="list-style-type: none"> ▶ No lost transactions. 	<ul style="list-style-type: none"> ▶ Users must check for last finished transactions and eventually manually insert lost transactions.
<ul style="list-style-type: none"> ▶ All IT staff is available — each recovery procedure can be done by responsible person (administrator), who is knowledgeable about the procedure. 	<ul style="list-style-type: none"> ▶ Because not all IT staff may be available, some of the activities have to be done by less experienced people (for example, a Windows administrator has to restore UNIX machine).
<ul style="list-style-type: none"> ▶ People under less stress than a real disaster. 	<ul style="list-style-type: none"> ▶ People under stress can make more mistakes.
<ul style="list-style-type: none"> ▶ It is recommended to stay for some time on the backup site, but roll-back time is set. 	<ul style="list-style-type: none"> ▶ If disaster strikes, there may be no definite information on when the roll-back time will be, and operations on the backup site can run in some cases for months.
<ul style="list-style-type: none"> ▶ Roll-back process is similar to the DR test, but in reverse (backup site restoring to the primary). This gives the opportunity to test the DRP twice with different staffing. 	<ul style="list-style-type: none"> ▶ If the primary site has been destroyed by the disaster, a new roll-back plan should be prepared, because the new primary site may be different from the old one.

4.1.2 Partial activity testing or overall plan testing

Basically you can test the whole Disaster Recovery Plan or test it part by part. Providing a regular testing of the whole Disaster Recovery Plan could be difficult especially in the case of higher tier levels. It can jeopardize daily operations and be costly and time consuming. Testing part by part is more easily planned and minimizes adverse side effects, however it may not cover all the interrelationships between components of the plan, and will not give information on the total time for recovery required.

Some core procedures contained in the DRP are probably performed regularly (in non-disaster situations) by expert IT administrators. Examples include restoring an AIX system from a **mksysb** tape, or re-loading Windows 2000 from the operating system CD. However all of these procedures are still required to be documented properly in the DRP. Even an expert, when under stress, can make mistakes, and even the experts may not be available at the crucial time.

The DRP should provide for two recovery cases. In the first case, the original (primary) site is completely destroyed. You will have to continue operations at the backup site for some time while the replacement equipment is ordered, installed and commissioned. In this case, the eventual switch to the new (replacement) site will usually occur over a period of time, and could be phased. However, some disasters may only be temporary — for example, if physical, electrical, or communications access is denied to your primary site for a period of time. In this case, once services were restored, you would want to roll-back (return) to the primary site as soon as possible. This procedure would be equivalent to the original DR in terms of timing, but in this case returning to the primary from the secondary. When you perform the test of the overall DRP you should consider how the return to normal operation will be achieved. The development of complex Disaster Recovery testing plans on higher tier levels can be difficult and can have a significant effect on daily operations. We suggest that from Tier 4 and higher to consult with a specialized Disaster Recovery service provider.

4.1.3 Timeframe for DRP testing

Many customers run their DR test over a weekend. They stop operations at the primary site, restore backups at the recovery site, run some tests, then switch back to the primary. While this will validate the actual recovery procedures, it does not ensure that normal operations are possible at the recovery site. If possible, when you are planning a test of the overall Disaster Recovery Plan, it is a good practice to stay on the hot site for some time — up to one week. This time is generally long enough to check issues such as connectivity, computing and network performance, to see if the hot site has satisfactory capacity for a typical production workload. Of course, running a production workload on the secondary site has issues of how to synchronize data updates back to the primary site when

roll-back occurs, since the ensuing transactions will have updated the data. Therefore, this practice usually is performed by customers with remote mirroring capability between the primary and recovery site (for example, PPRC). This means that all transactions performed at the recovery site will be reflected back at the primary.

Customers have had good DR test experiences whereby they switch to the backup site one weekend, remain there for one week, then roll-back to the original site on the following weekend. In this way, the customer stays on the backup site for five working days. Disaster Recovery Plan testing should be done annually, because IT equipment, and applications are changing very rapidly. Also IT staff volatility mandates regular education, re-education and practical training.

4.1.4 Example of a simple testing plan scenario

The DRP test will consist of executing a number of procedures and tests. Each should be individually documented, showing the steps and material required, with space to record the results of the test execution.

This example shows how a simple individual test scenario can be documented. Note that it makes reference to the enterprise's own procedures manuals.

Scope of test scenario

The scope gives the task, estimated time, and essential steps.

Task: Restore the NT application server, *SV002*, after a total hardware crash. The operating system (system disk) and the application will be restored. This scenario assumes that the application data is located on a disk array which was not destroyed. The system partition was backed up by an IBM TSM server the previous night, and the backup should be consistent.

Estimated time: 60 minutes.

The materials shown in Table 4-3 are required.

Table 4-3 Materials to restore NT server

Material	Quantity	Note
Bootable CD-ROM for Windows NT server version 4.0	1	
CD with Windows NT SP6	1	
IBM Tivoli Storage Manager Desktop Clients CD	1	

Material	Quantity	Note
Record of network settings for SV002	1	
Administrator password for SV002	1	
Appropriate hardware	1	

Essential steps: Follow the internal procedure for restore of Windows NT server. This procedure is described in the internal *Systems Restore Procedures* on page xxx (Note: of enterprise’s own procedures manual).

The procedure for each test scenario must be provided in writing — containing all the required steps, supporting information and suggestions for improvement (for example, after a test run).

4.1.5 Example of a more complex test scenario

Suppose we have an operational environment, which consists of a clustered database server, two application servers, and a TSM server with tape library. The online application data is located on a disk array, which is mirrored to a remote (backup) location. The backup data created by TSM is moved offsite on the tapes. At the remote location we have a disk array with online mirrored data, backup database server, and two application servers — one of these also acts as a TSM server.

Task: Restore operations at the remote site after the primary site was totally destroyed by disaster.

Preparation steps: Before a disaster is declared some preparation steps should be done so that the DR testing will not jeopardize operations. These are shown in Table 4-4.

Table 4-4 Preparation steps

Time	Task	Who
10:00pm Friday	Check if weekly processing finished.	Operator on duty.
05:00am Saturday	Check if weekly TSM backup finished and all offsite tape was moved to backup location.	Operator on duty.
05:30am Saturday	Shutdown of application servers.	Operator on duty.
05:45am Saturday	Shutdown of database server.	Operator on duty.

Time	Task	Who
06:00am Saturday	Declaration of "Disaster" - primary site was stricken by disaster - time t_0	Operator on duty.

Essential steps: Next, Table 4-5 shows an example of the DR testing scenario.

Table 4-5 Disaster Recovery testing scenario

Time	Task	Who	Note
t_0	Start escalation procedure (Disaster Recovery Plan page xx).	Operator on duty.	
$t_0 + 30$ min.	Manager on duty arrives at primary site (Disaster Recovery Plan page xx - first steps).	Manager on duty.	
$t_0 + 45$ min.	Decision to move to backup site (Disaster Recovery Plan page xx).	Manager on duty.	
$t_0 + 120$ min.	Start of recovery tasks (Disaster Recovery Plan page xx).	Manager on duty. DR teams.	
$t_0 +$	Further site-specific tasks.	DR teams.	

Review: All steps done in the DR testing procedure should be recorded. This record will become basic document for post-testing review. An example of a format which could be used is shown in Table 4-6.

Table 4-6 Record of DR executed steps

Planned start time	Planned finish time	Actual start time	Actual finish time	Task	Comments

4.2 When to provide Disaster Recovery testing

One of the most frequent questions is how often testing has to be done. A general recommendation for testing the entire Disaster Recovery Plan is at least once a year.

Partial test activities, such as a recovery of individual file servers, application servers or database servers can be done a few times per year. Every member of the technical staff should at least be able to do these tasks using the relevant part of the DRP. The opportunity for these kinds of activities often arises spontaneously in normal IT operations. For example, if there is a normal business requirement to reinstall a server, this is a good opportunity to test the readiness of the technical staff. Before reinstalling it, try to get time to do a test recovery. An untested procedure will almost certainly fail when performed in a real disaster situation.

Utilize every opportunity for testing. Investment in testing is never a waste of time or money. A DRP is like car insurance. You pay money to the insurance company (which is a financial commitment) but in the case of an accident you expect that all damages will be covered by the insurance company. If you didn't think the insurance company could pay the damages, you would not pay for their coverage. Similarly, if your DRP does not work, it needs to be fixed — however you will not know if it works, and what is wrong with it, until you test it.

4.2.1 How to select the best time for testing

There is no right answer to this question for all companies and circumstances. Some companies schedule DRP testing for long weekends. Other customers try to predict a time when business activities are at a lower volume, for example, in the summer months. We recommend that you discuss the timing with all the participating departments so that all agree — if a single time can't be found which pleases everybody, then management has to make the decision.

4.3 DRP maintenance procedures

A process of change management for the DRP must be in place to:

- ▶ Determine whether any planned technical and organizational change will either meet the standards of the existing DRP or necessitate a change to it.
- ▶ Ensure that each single function will report immediately any changes which could impact the DRP.
- ▶ Regularly revise the DRP in accordance with the actual recovery policy.

A procedure must exist to regulate the following issues:

- ▶ Ensure that the documentation of the DRP is valid at all times (version control of document).
- ▶ Ensure that all people listed on the distribution list have received a copy of the last valid update.
- ▶ Ensure that old, no longer valid versions of the documentation have been destroyed.

If involved personnel either have no access to the current DRP, or have instead, an obsolete version of it, this could seriously jeopardize a company's ability to correctly respond to a disaster.

This section describes how the maintenance of the DRP could be performed. The procedures given here are examples only. The actual procedures must be discussed and agreed to internally and modified in order to meet individual specific requirements.

The actual implementation of these procedures is the customer's responsibility. Generally, however, the following disciplines must be covered by the maintenance procedure:

- ▶ Delegating and specifying responsibility for the DRP document as a whole or for individual chapters
- ▶ Maintaining the distribution list for updated copies of the DRP
- ▶ Establishing standards for correct identification of the document. Identifying fields such as security classification, date of change, date of printing, date of release, name of owner should appear on the title page and probably on each individual page as well
- ▶ Securing access to the data file, that is, specifying who is allowed to modify the document
- ▶ Ensuring correct handling of the document in order to maintain security and availability
- ▶ Putting in place a release procedure for the document after modifications
- ▶ Ensuring correct document control (valid documents distributed, obsolete documents collected)
- ▶ Designating a contact point in case of errors, required modifications and updates
- ▶ Defining a regular process of document revision
- ▶ Establishing a *scratch* procedure for obsolete versions of the document

4.3.1 Maintenance policy

The purpose of this policy is to ensure the validity, availability and confidentiality of the DRP.

This policy defines the responsibilities regarding maintenance and audits. In addition, it describes the distribution process, the audits process, the release process, and the obsolescence process.

Responsibilities - Approval Board

The Approval Board is responsible for officially signing off the performed changes and updates and ultimately approves the release of the change or update.

The list of members of the Approval Board should be recorded with at least the information shown in Table 4-7.

Table 4-7 Approval Board members list

Name	Function	Department

Responsibilities - Document Owner

The Document Owner is responsible for the overall administration of the Disaster Recovery Plan.

The following person is defined as the Document Owner in Table 4-8.

Table 4-8 Document Owner

Name	Function	Department

Responsibilities - Task Owners

The Task Owners are responsible for the correctness of specific parts or chapters of the DRP. All Task Owners should be recorded as shown in Table 4-9.

If the chapters are stored in separate files, also record the name of each file in the appropriate column.

Table 4-9 Task Owners

Chapter			TASK OWNER	
No.	Title	Filename	Name	Department

Duties and tasks

The Task Owners are responsible for forwarding the file or files for which they are responsible to the Document Owner. The Document Owner collects all the individual parts or file from the Task Owners. The Document Owner composes the total document by merging the components into a single document.

At regular intervals (for example, every six months) the Document Owner contacts the Task Owners to see if changes are necessary. In case of updates, the Document Owner contacts the Approval Board.

The Approval Board is responsible for the official sign off of all changes and finally approves the release of the changes or updates made.

After receiving the release from the Approval Board, the Document Owner issues a new Version Number. The Document Owner records the changes in the Change History. The Document Owner arranges for the printout of the amended DRP and initiates the distribution of the DRP to the Distribution List.

As this document contains vital customer information, the Document Owner has to choose a secure distribution channel. The Document Owner requests the return of all obsolete versions of the DRP.

The Document Owner collects the obsolete versions of the DRP and initiates destruction of these documents.

Version control

To ensure clear identification of each page of the DRP, each page must contain the following data:

- ▶ Classification
- ▶ Version Number
- ▶ Release Date
- ▶ Total number of pages of the document

Obsolescence

Scrapping of obsolete documents must be according to the general scrapping procedure for internal customer documents.

The following additional instructions apply:

- ▶ Obsolete versions of the document must be returned to the Document Owner.
- ▶ The Document Owner verifies that all obsolete versions of the document have been returned.
- ▶ The Document Owner initiates secure scrapping of all returned documents.
- ▶ Scrapping can be performed by use of standard paper shredder equipment or other secure techniques, which prevent documents from being accessible to non-eligible personnel.
- ▶ If an external vendor performs the scrapping, an appropriate certificate signed by this vendor must be requested.

Audits

If no changes or updates have been initiated within the last 12 months, the Approval Board initiates a review or audit of the DRP in order to ensure the correctness of the document. The review must include the correctness of the technical procedures and that all contact names, telephone numbers and addresses are valid.

General issues

The source-files of the DRP must be protected against uncontrolled modification and unauthorized access. The hard copies of the DRP should not be made available to unauthorized persons.

The document must be available to authorized staff at any time in order to ensure immediate activation of the DRP when necessary.

The members of the Management Team, who are responsible for the recovery process should have access to the valid DRP copy at all times, for example, by

keeping a copy at home. This allows them to have the required information available immediately after the notification of a disaster incident.

4.3.2 Reference information

This section gives examples of how to collect and format information to be used for the administration of DRP maintenance procedures.

Plan of audits

Table 4-10 records the audit process.

Table 4-10 Plan of audits

Date Audit planned	Date Audit performed	Audit requested by	Signature

Change History

Table 4-11 can be used to document the changes performed.

Table 4-11 Change history

Release Date	Version	Changes	Performed	Approved

Release Protocol

Table 4-12 may be used to keep track of version releases.

Table 4-12 Release protocol

Version Number (new):			
Release Date			
Approved by:	Function	Signature	Date
Brief description of changes / updates:			
Sign off by Document Owner			

Distribution

The DRP will be distributed to the distribution list shown in Table 4-13.

Table 4-13 Distribution list

Name	Department

Cover letter

The following letter could be used as a cover for the distribution of the Disaster Recovery Plan.

Date

To distribution-list:

Subject: Update of the Disaster Recovery Plan

Attached find the following update of the Disaster Recovery Plan:

Version:

Release Date:

Please replace the complete handbook with this new version and return the obsolete document to my attention by (insert date).

Signature



Planning for data center availability

In this chapter we discuss planning strategies and metrics for designing fault tolerant infrastructure, including the following topics:

- ▶ Infrastructure, the first line of defense
- ▶ Data center availability and planning fundamentals
- ▶ Distributed systems considerations
- ▶ Storage architecture
- ▶ Network architecture (LAN, WAN, SAN)
- ▶ Hotsite, warmsite and coldsite planning considerations

5.1 Infrastructure, the first line of defense

In a data processing environment, application services depend on every fundamental building block of the data center. A systemic view of a data center shows the relationships and dependencies of every basic component. Every system depends directly on a long list of services. Designing a world class data center from the ground up follows a design process focused on intrinsic properties for security, reliability, and availability of services. Services can range from power, to cooling, to security, to data storage, to networking.

TSM provides storage management services for backup, recovery, and disaster recovery. TSM functionality depends directly on the existence of a well designed and managed environment.

In many cases, large IT infrastructures have grown at tremendous rates, with few, if any design standards for technology. The resulting environments sometimes contain 20+ platform types, 5 to 6 database platforms, and a heterogeneous mix of applications and development tools from different vendors. Usually development teams select application and database technology, while operations teams influence hardware platform decisions. Everyone has their personal technology preference, and through time and employee turnover this naturally tends to lead to increasing entropy in the data center environment.

Since storage management usually crosses so many departmental boundaries, storage infrastructures are often completely unplanned and enterprise data recovery systems are more often distributed than centralized. Beyond the physical infrastructure, enterprise data management policy is a rare occurrence. Storage and storage management continues to move into the limelight for data center operational challenges.

In the following sections, we discuss the importance of strategic planning and the importance of mapping design metrics to availability requirements. The role of IT architecture planning has become central to strategic planning for the enterprise. To create a manageable enterprise environment, organizations must incorporate strategic architecture planning into data center procedures. Enterprise hardware and software standards, architectural design metrics, and the establishment of business requirement driven policy all make a tremendous difference in creating a manageable enterprise environment.

We encourage enterprise deployments of storage management software tools. However designing functional infrastructure to support operational requirements for availability should come first.

5.2 Data center planning

During the late 1990's e-commerce boom, several companies specialized in providing data center outsourcing facilities for high-growth startup companies. These data centers were built around design fundamentals which incorporated redundancy, resiliency, and high security concepts. The selling point for this high tech real-estate was the fact that companies could *trust* their hardware operations to a world-class data center specialist. Any organization can adopt the fundamental elements of data center design which created this market for world class data center space.

5.2.1 Space and growth considerations

One of the most challenging elements of data center planning includes space planning. Many organizations fill up space at 2 or 3 times the planned rate of space consumption. Technology density (such as, servers and storage) continues to improve, but space planning should be the first element of design. Data center space should accommodate 5-10 years of growth, depending on the business type and scale. The cost of building extra space typically pales in comparison to data center relocation costs and risks.

5.2.2 Fundamental elements of data center design

From a structural standpoint, the data center must be designed to withstand many kinds of disruptions, ranging from building evacuations to regional seismic events. The Risk Analysis phase of the Disaster Recovery Planning process, described in 3.2.2, "Disaster Recovery Planning" on page 44, helps organizations understand the credible risks to the environment. Location of the data center is extremely important too, since building access can be controlled by external entities. This could hinder or complement recovery and security procedures for the data center.

Design fundamentals for a data center components center on scalability, redundancy, and resiliency. As an example, the power infrastructure must provide redundancy and scalability without disruption. Every power management device (transformers, UPS, systems, and so on) must be built with redundancy in mind, just like high availability systems architecture. Some organizations take measures to source power from separate power grids and suppliers, so as to reduce points of infrastructure failure, back to multiple power generation sources. An example of infrastructure redundancy is shown in Figure 5-1.

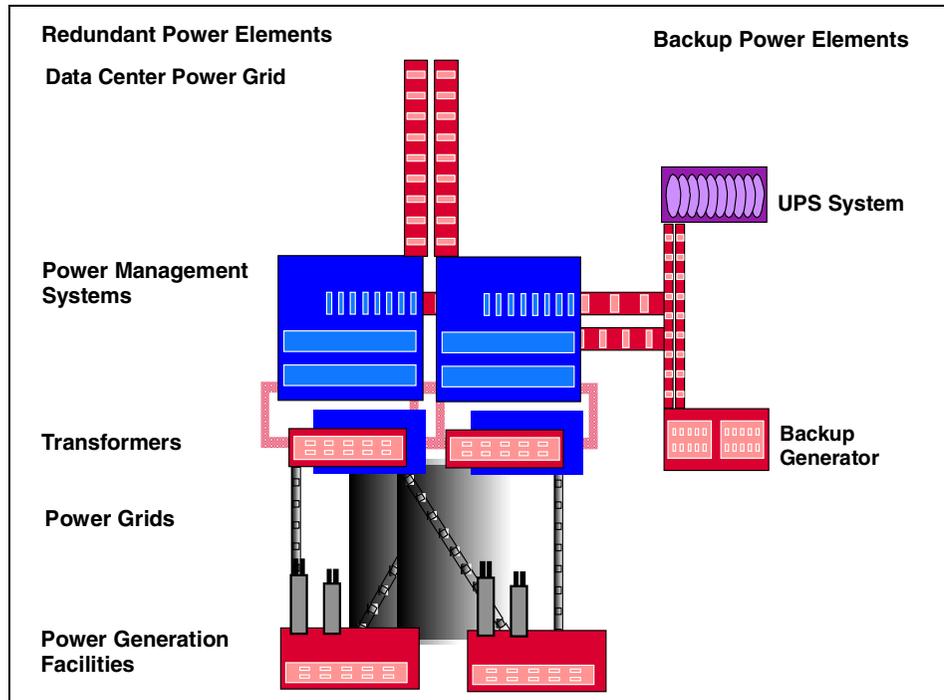


Figure 5-1 Example of power systems design for availability

Similar design methods can be applied to other infrastructure components, including cooling, HVAC, halon, fire prevention, networks, telecommunications, and fibre optics. Support in the design or retro-fitting process can be obtained through professional services groups, including the data center planning specialists within IBM Global Services.

5.2.3 Preventative controls

An additional element of data center design includes the use of preventative controls. Redundancy and protection come at a cost that must balance with risks identified in the BIA process. Frequently found preventative controls include:

- ▶ **Uninterruptible power supplies (UPS)** to provide short-term backup power to all systems components (including environmental and safety systems)
- ▶ **Petroleum powered generators** to provide long-term backup power
- ▶ **Air conditioning systems** with adequate excess capacity
- ▶ **Fire suppression systems**
- ▶ **Fire and smoke detectors**

- ▶ **Water detectors**
- ▶ **Plastic tarpaulins to protect equipment from water damage**

Preventative controls must be documented and integrated into the overall DRP. General awareness of how these measures are used is very important for personnel (for example, appropriate response if a fire alarm sounds), and can be instilled via drills, documentation provided to employees and so on. This will ensure they know what to do in the case of a real disaster.

5.2.4 Onsite parts maintenance locker

A simple and effective way to fortify site redundancy is to initiate a program for onsite parts inventory. Components which frequently fail (network cables, connectors, disk drives, tapes, HBAs, and so on) can be stocked in the data center for fast access in the event of component level failures. Some organizations compile and analyze component failures for mean time between failure (MTBF) data, to help cost-effectively prepare for these basic kinds of service interruptions.

5.2.5 Data center facilities security

Security controls for data center access are also extremely important. Creating access controls and procedures helps to fortify operations against malicious human behavior. Scenarios range from card-key access doors to armed and fully supervised one person entry/exit chambers, which include visual recognition, armed guards, and two-stage points of clearance. At a minimum, data center access should be controlled and monitored.

Site locations can also be kept virtually private from public awareness by the use of unmarked buildings and data center locations. Some government and energy installations, for instance, limit the number of people who have knowledge of data center locations and access procedures. However paranoid, these measures provide an excellent level of protection and security.

5.2.6 Disaster Recovery Planning and infrastructure assessment

We realize that the majority of organizations rarely have the opportunity to design a data center from the ground up. Best practices for infrastructure and data center design are applied retroactively at best, and in most cases affect only new systems. The DRP process provides incredible value for mapping logical and physical environments and identifying contingencies within the infrastructure. We encourage readers to approach DRP with this frame of reference.

5.3 Distributed systems architecture considerations

Every organization faces the challenge of standardizing platform architectures to help streamline operating efficiencies in an open systems environment. Analysts, such as Gartner Group, have suggested the identification of dual strategic platform vendors, to help control the proliferation of various operating systems and hardware devices in heterogeneous distributed systems environments while still encouraging competition and best pricing for acquisitions. The dual vendor strategy is worth consideration for long term strategic planning efforts.

5.3.1 Server architecture

Servers should be designed to provide maximum redundancy and scalability. Standards for platform architecture should include redundant power supplies, redundant cooling devices, hot-swappable PCI devices, and hot-swappable internal disk devices. The internal system design should also eliminate most, if not all, single points of failure. Predictive failure analysis capabilities for CPU and memory should also be available from the operating system. For any organization, these architecture standards should be designed, documented, and followed for all production systems.

5.3.2 Server operating system clustering technologies

Several software applications allow multiple servers to be clustered together for high availability. Clustering software generally monitors systems resources and devices through protocol driven “heartbeat” messages. Depending on events and definitions, the absence of system resources triggers the failover of an application environment from one server to another. HACMP (AIX), Microsoft Cluster Server (Windows), MC Serviceguard (HP-UX), and Sun Cluster software (Solaris) provide high availability cluster services for their respective hardware platforms.

Clustering software adds a layer of complexity to system administration and application administration. Changes made in a production cluster can accidentally trigger a failover event. We suggest high availability clustering for mission-critical environments that have development or testing platforms available to test significant systems modifications. A clustered server environment should be infrequently modified during production.

For AIX environments, the HACMP application allows multiple servers to be clustered together for high availability. HACMP clusters support the TSM application so that if the primary system fails, the server will be brought back up on another system in the cluster. Several HACMP start/stop scripts are provided with the TSM server fileset. These scripts need to be customized to suit the local

environment. More information on using HACMP to cluster AIX servers is provided in 7.8.1, “TSM and High Availability Cluster Multi-Processing (HACMP)” on page 144.

Microsoft cluster server is also supported with the TSM application running on Windows server platforms. Clustering techniques for TSM in the Windows environment are discussed in *Using TSM in a Clustered Windows NT Environment*, SG24-5742, and in 7.8.3, “TSM and Microsoft Cluster Server (MSCS)” on page 148.

5.3.3 Capacity planning

Capacity planning and growth forecasting also affects metrics for systems architecture. Depending on the platform environment, several classes of machine type and size can be applied to specific performance and growth requirements.

5.3.4 System software

System software, such as operating systems, patches, upgrades, security patches, and systems utilities should be standardized for all platforms in a production environment. This software should be documented and considered a functional component of systems level recovery. Copies of all systems software and licensing information should be cataloged and stored in a secure offsite location.

5.3.5 Configuration management

An overall objective for infrastructure planning is configuration standardization across the enterprise. Using configuration and platform standards not only affects efficiencies for systems management, but it also directly impacts the ability of an organization to recover systems in the event of a disaster. Detailed records for system configurations should be integrated into the DRP, as an additional guide for internal staff and vendors. Disaster Recovery Manager with TSM provides a mechanism for this.

5.3.6 Problem and change management

Problem and change management procedures should be established and documented as part of an enterprise server environment. For every production system, a problem and change management procedure and log should be defined and consistently used. Copies of these logs should also be stored offsite and possibly at the Disaster Recovery location.

5.4 Disk storage architecture considerations

Like server architecture, disk storage devices should incorporate design capabilities to limit single points of failure. At a minimum, disk devices should include redundant power supplies, redundant cooling devices, hot-swappable adapters, and hot-swappable internal disk drives. The internal system design (that is, the system backplane) should include no single points of failure.

Advanced storage subsystems tend to be either cache-centric or controller centric. Cache-centric disk subsystems provide generous amounts of memory-cache for I/O operations, generally improving I/O performance and dramatically increasing system cost. Controller-centric disk subsystems offer front-end processing capabilities, with less cache capabilities and generally lower costs. Disk sub-system cost can be evaluated in terms of price per MB for raw storage volumes. The analysis should then include equivalent functionality and overall management cost between each subsystem.

Today, disk subsystems provide some level of RAID (Redundant Array Of Independent Disk) protection. In general, RAID algorithms spread data across multiple disks, to lower the odds of disk failure losing data. Hardware and software level RAID protection is available from most open systems vendors, and should be considered a standard for any disk environment in production.

RAID subsystems typically use three redundancy techniques:

- ▶ **Mirroring:** When data is mirrored, the system writes data simultaneously to separate hard drives or drive arrays. If a disk drive fails, the system can still read data from the mirrored logical volume.
- ▶ **Parity:** Parity is a technique used to determine whether or not data has been lost or overwritten. In RAID implementations, disk drives can be dedicated or partially used for data parity. Essentially, data redundancy is shared across multiple volumes with parity.
- ▶ **Striping:** Striping distributes data across multiple physical drives in a RAID array. A logical group of data is split into block or byte sized increments, and written sequentially across multiple drives. Striping alone is known to improve disk I/O performance, because more drives can access a logical sequence of data at one time. However by itself, it does not provide redundancy.

Several varieties of RAID are available for data protection, however the most commonly used RAID solutions use RAID-1, RAID-3, and RAID 5, described here:

- ▶ **RAID-1** Otherwise known as disk mirroring, RAID-1 is generally the most expensive and fault tolerant method for protecting data. Some RAID-1 implementations also use striping (RAID-0) to improve performance, with the effective result being a RAID 0+1 solution. Mirroring can be either positive or

negative on performance, depending on the application types and other operating system characteristics, and striping generally improves performance. Cost is typically a limiting factor for large scale RAID-1 implementations.

- ▶ **RAID-3** Large file applications sometimes benefit from the performance attributes of a RAID-3 environment. RAID-3 stripes data at a byte level across multiple drive, with one drive dedicated to parity data.
- ▶ **RAID-5** The most commonly implemented RAID technique, RAID-5 combines parity with striping across multiple drives in a disk array. This means that data and parity data is spread across multiple volumes. Hot-spare drives can also be built into a RAID-5 array, to provide an additional level of fault tolerance. While performance sometimes degrades due to parity I/O operations, front end caching in modern disk subsystem helps speed the performance of large scale RAID-5 storage environments.

At a minimum, data should be stored in a RAID storage environment. This provides a basic level of protection to guard data from component level failures in a disk environment.

5.4.1 Advanced storage subsystem functionality

Within and between disk subsystems, several additional capacities exist for data protection and data management. Within subsystems, data volumes can be instantly copied, using techniques such as IBM ESS FlashCopy or EMC TimeFinder. Data can also be mirrored synchronously or asynchronously between subsystems remotely and locally using IBM PPRC or EMC SRDF. TSM also integrates policy and tools with many of such advanced copy services, to help relate BIA policy needs to systems architecture.

Techniques to use these services along with TSM are discussed, in detail, in 7.9, “TSM and remote disk replication” on page 151.

5.5 Network architecture considerations

Enterprise design and contingency standards can be applied to LAN, SAN, and WAN environments alike.

5.5.1 Network architecture

From an architectural view, single points of failure should be avoided for critical systems. Redundant backbones, switches, routers, cables, and adapters should

be in place for critical systems identified in the BIA. For recovery procedures, network failover plans should focus on critical systems and TSM architecture.

Network architecture relates to both data and network service availability. Figure 5-2 illustrates multiple levels of network redundancy to ensure availability for LAN and WAN type network services.

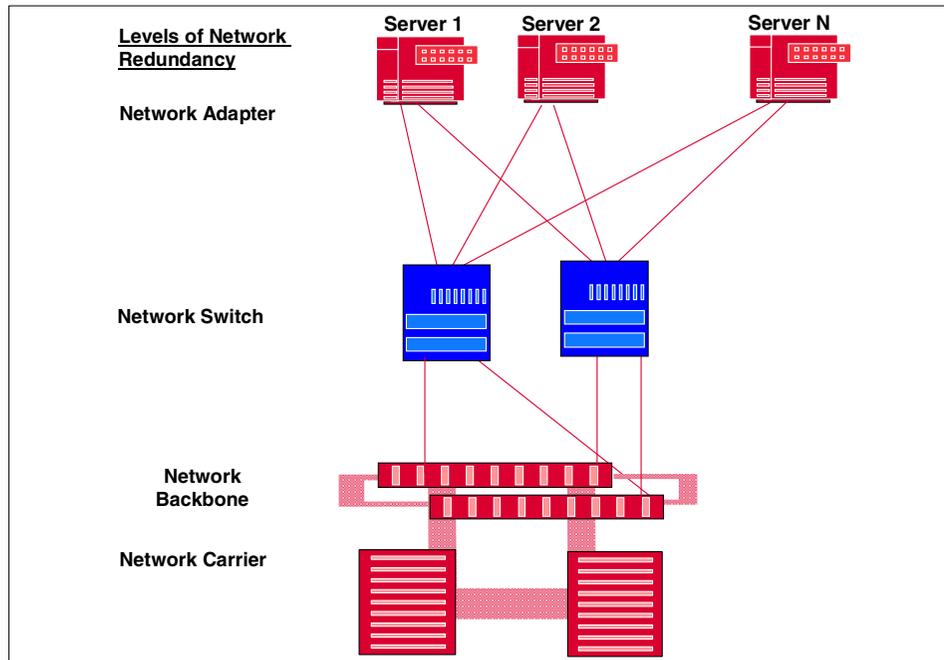


Figure 5-2 Redundant network architecture

A similar design methodology applies to Storage Area Networking (SAN) and data availability, as illustrated in Figure 5-3. This figure shows a typical highly-available SAN, with duplicate HBAs, links and FC switches.

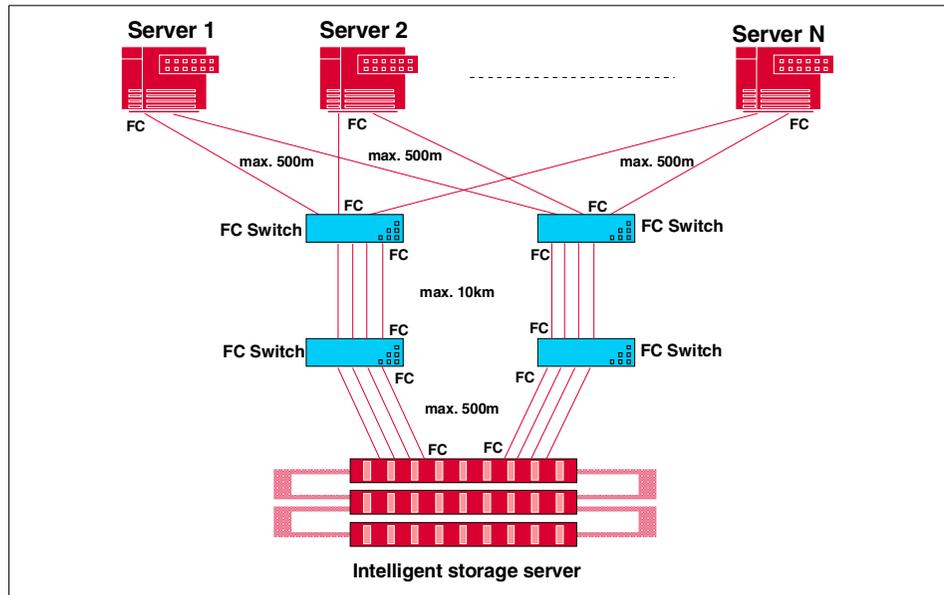


Figure 5-3 Storage Area Network architecture for data availability

The distance components in a Storage Area Network, such as IP extenders, multimode (longwave) fibre, or DWDM devices follow the same design metrics found in a SAN for high availability. As long as single points of failure are eliminated from any point to point connection, transaction integrity and data availability is best insured.

5.5.2 Network documentation

The physical and logical network diagrams should accurately reflect the production environment. A physical diagram should display the physical layout of the facility and connected devices. The logical diagram should present the network devices and the connected nodes. Detailed information device names, addresses, DNS, and subnets (or zones) should be documented as well.

5.5.3 Network management

All network types should also be monitored and managed from an enterprise view. Problem management and resolution in a large environment depends not only on skilled resources, but also intelligent monitoring tools for network topology and event management. Various network tools, such as Tivoli NetView, HP OpenView, and CA Unicenter provide this functionality. Tools for SAN management with similar capacities are emerging as well.

5.5.4 Remote access considerations

Remote access should also be securely configured for all network components, so that systems can be administered from a remote location. In a Disaster Recovery scenario, remote access to systems might determine the ability of the DR team to work with normal production systems. Some threats to systems environments, such as biological or chemical terrorism events, might limit the ability for onsite work, while systems could maintain functionality in a primary production facility.

5.5.5 WAN considerations

WAN system configurations should also be well documented. Any vendor or service provider can be a single point of failure for an enterprise, especially for e-commerce or IP communications-dependent businesses. Now that voice and data services rely more on IP services, identifying and eliminating contingencies with network service providers is of ultimate importance. The power planning methods described in 5.2.2, “Fundamental elements of data center design” on page 81 can be similarly applied to WAN providers. The overall WAN architecture can be designed to eliminate multiple layers of potential failure, ranging from switches to service providers.

5.5.6 Network security considerations

Security services for all networks must be an enterprise priority. Firewall device and software configurations should be well documented and secured offsite. IP addresses should be treated as confidential information. All network security measures included in production network operations also must be included in the Disaster Recovery Plan. The plan should document the communications providers (and their SLAs) involved in all components of network security.

5.6 Hotsite, warmsite, coldsite considerations

Depending on the BIA and high-level strategies for enterprise availability, an organization may choose to use mirrored, hot, warm, or coldsite facilities for disaster recovery. Some measure of facilities, equipment, and data contingency must be in place for an enterprise DRP. Facilities planning is usually a function of business continuity planning, and typically includes the data center and a fully functional employee workspace.

Alternate site selection prepares an organization for a disaster event which renders the production site unusable or irreplaceable for short to long-term periods of time. The overall risk and cost of downtime for IT operations directly

determines the kind of alternate site architecture. Figure 5-4 illustrates the cost and recovery time benefit for the different kinds of alternate site architectures.

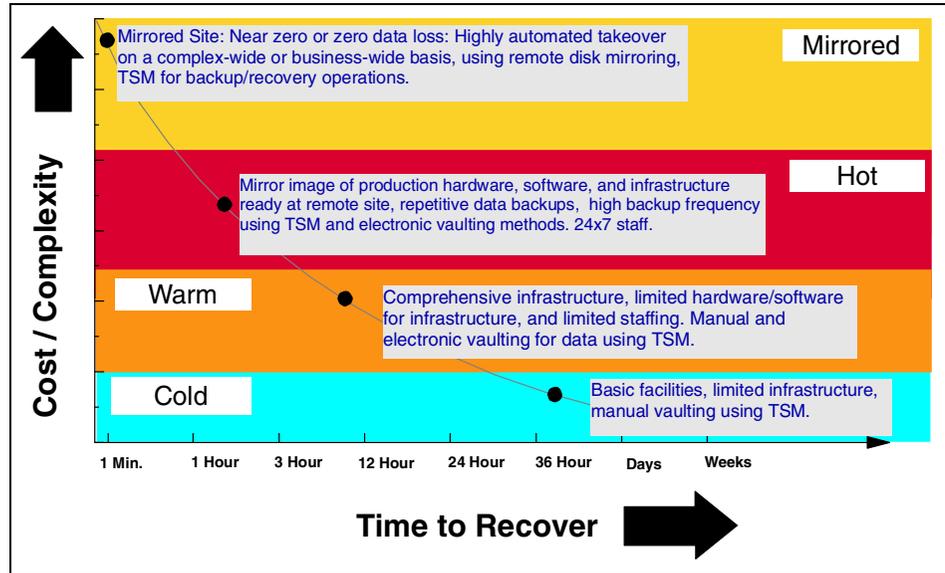


Figure 5-4 Cost versus recovery time for various alternate site architectures

There are obvious cost and recovery time differences among the options. The mirrored site is the most expensive choice, because it ensures virtually 100 percent availability. Coldsites are the least expensive to maintain; however, they require substantial time to acquire and install necessary equipment. Partially equipped sites, such as warm sites, fall in the middle of the spectrum. Table 5-1 summarizes the criteria that can be employed to determine which type of alternate site meets the organization's business continuity and BIA requirements.

Table 5-1 Alternate site decision criteria

Site Type	Capital Costs	Hardware/ Software	Networking/ Communications	Setup Time
Mirrored	High	Full	Full	Minimal
Hot	Medium/High	Full	Full	Hours
Warm	Medium	Partial	Partial or Full	Days
Cold	Low	None	None	Days/Weeks

We will discuss four main types of alternate sites, which include the use of TSM for storage management and disaster recovery:

- ▶ **Mirrored Sites:** Mirrored sites are fully redundant facilities with full, real-time information mirroring. Mirrored sites are identical to the production site and provide the highest level of availability, because data is written and stored synchronously at both sites. Mirrored sites are built and maintained at approximately twice the operating costs of the production data center facilities.

In a mirrored site, TSM provides backup, restore, and archive operations and adds additional layers of protection against data corruption. Site to site mirroring is the best way to provide 100% availability, but data corruption and cyber-threats pose a credible risk to even fully redundant operations. Maintaining functional versions of data in a mirrored site is key to ensuring high availability.

- ▶ **Hotsites:** Hotsites are equipped with fully functional and prepared servers, storage, network, and software systems. Hotsites are staffed for 24x7 operations, and in some cases share a portion of the production workload as a measure to justify ongoing capital and human costs. Hotsite costs vary according to the scale of operations and recovery requirements.

TSM provides continuous backup and availability of production data to a hot-site operation. Electronic vaulting techniques might be used to move production data backups from site to site. Once the DRP is activated, hotsite personnel immediately begin the restore process in the recovery environment. In some instances, production data backups may be routinely restored to hotsite systems to minimize the overall time to restore in the event of a disaster. Specific requirements and procedures for data availability can be developed from RTO data for each system.

- ▶ **Warm sites:** Warm sites provide infrastructure and equipment to restore critical business functions. Generally, all power and environmental services are in place, along with hardware, software, and network components needed for operations. The site is maintained in an operational status with minimal staff.

TSM provides the capacity to vault data to a warm site through manual or automated procedures. Production data backups can be vaulted to a warm site TSM environment and restored to recovery systems in the event of a disaster.

- ▶ **Coldsites:** Coldsites typically consist of a facility with adequate space and infrastructure (electric power, telecommunications, environmental controls) to support the IT data processing environment. The site does not contain IT equipment or office equipment needed for operations. Usually, coldsite facilities are leased through third-party service providers, and the majority of equipment is insured through service level agreements with hardware

vendors or disaster recovery hardware vendors. Organizations such as Sunguard and IBM Business Continuity and Recovery Services specialize in designing these kinds of service contracts for enterprise organizations.

For coldsite operations, TSM can provide offsite copies of data to be stored at or near the cold site. TSMs automated capability to manage offsite tape volumes provides a cost effective and automated method for coldsite data management.

5.6.1 Alternate site cost analysis

Several variables factor into the cost analysis for an alternate site. Facilities can be either owned or leased through commercial vendors. Either way, facilities costs represent a continual capital cost for any alternative site plan. Infrastructure, telecommunications, and networking costs are usually grouped with facilities costs. Hardware and software is secured via full acquisition, hardware DR vendor contracts, or a mix of the two elements. Travel, labor, and testing costs are site and event dependent, but must factor into the overall operational cost model. Other miscellaneous variables may also affect the cost models. Table 5-2 provides a basic template for alternate site cost analysis and comparison.

Table 5-2 Alternate site cost analysis template

Site Type	Facilities Costs	Hardware/ Software	Travel / Logistics	Labor / Contracting	Maintenance/ Testing
Mirrored					
Hot					
Warm					
Cold					

As site costs and architectures are evaluated, primary production site security, management, operational, and technical controls must map adequately to the alternative site design.

5.6.2 Partnering strategies for alternate sites

Two or more organizations with similar or identical IT configurations and backup technologies may enter a formal agreement to serve as alternate sites for each other or to enter a joint contract for an alternate site. Enterprise customers with multiple production sites may also design alternative site policies between locations.

This type of site is set up via a reciprocal agreement or memorandum of understanding. A reciprocal agreement should be entered into carefully because each site must be able to support the other, in addition to its own workload, in the event of a disaster.

Written agreements for the specific recovery alternatives selected should be prepared, including the following special considerations:

- ▶ Contract duration, extension of service, and termination conditions
- ▶ Testing and maintenance procedures
- ▶ Shared and unshared capital costs
- ▶ Security procedures
- ▶ Change and problem management procedures
- ▶ Hours of operation, availability, and personnel requirements
- ▶ Specific equipment required for recovery
- ▶ Disaster Recovery Plan notification procedures
- ▶ Guarantee of compatibility and service
- ▶ Priorities for system recovery in the event of a multi-site disaster

Multiple party commitments for alternative site provisioning should carefully map business requirements for availability to a joint DR Planning process. Identification of mutual risks is key to developing a multi-party Disaster Recovery Plan and infrastructure. For both recovery sites, the DR sequence for systems from both organizations needs to be prioritized from a joint perspective. Testing should be conducted at the partnering sites to evaluate the extra processing thresholds, compatible system and backup configurations, sufficient telecommunications and network connections, and compatible security measures, in addition to the functionality of the recovery strategy.



Disaster Recovery and TSM

In this chapter we discuss the integration of IBM Tivoli Storage Manager policy with Disaster Recovery Planning data. The Disaster Recovery Planning and Business Continuity Planning processes yield a tremendous amount of procedural and policy related data about how an organization should react in the event of a disaster. Translating this information into storage management policy and infrastructure requires a thorough understanding of TSM and the basic methodology for mapping business continuity requirements to the technical infrastructure.

This chapter is not intended to be a TSM primer, or a detailed discussion of TSM system design. This information is contained in *Tivoli Storage Management Concepts*, SG24-4877, and *Getting Started with Tivoli Storage Manager: Implementation Guide*, SG24-5416.

6.1 Mapping general concepts to TSM

DRP functionality depends directly on the policy, procedures, people, and technical infrastructure supporting recovery operations. If any element which supports recovery operations becomes a bottleneck, business recovery objectives are compromised. This is why we introduce a data centric methodology for creating TSM policy and infrastructure. This approach relates business continuity requirements, to TSM policy, to the supporting TSM infrastructure resources.

Throughout the following sections, we stress the importance of carefully relating policy decisions to technical infrastructure design. TSM is a storage management application, and the most critical overall function is data recovery for the enterprise. Backup performance often receives primary focus, while recovery procedures and performance are an afterthought. The technical planning procedures we recommend make recovery performance the *primary* objective, since the time available for recovery will almost always be less than that for backup, and keeping in mind that backups are done at leisure, while recoveries are done under pressure. Backup-centric planning also excludes many complex variables experienced in Disaster Recovery scenarios.

As we discussed in 3.3, “Disaster Recovery Planning and TSM planning” on page 48, the BIA and DRP processes yield a great deal of data about critical processes, recovery time objectives, system and process dependencies, classification of systems and data, and general policies for data backup and versioning. In this chapter, we discuss how these specifics can be mapped to TSM technical policy.

At an enterprise level, TSM policy must meet overall business requirements for data availability, data security, and data retention. Enterprise policy standards can be established and applied to all systems during the policy planning process.

At the systems level, RTO and RPO requirements vary across the enterprise. Systems classifications and data classifications typically delineate the groups of systems and data along with their respective RTO/RPO requirements. Data classification schemes add a necessary layer of sophistication to policy generation, to effectively streamline backup and recovery operations. Specific variables for backup type, backup frequency, number of backup copies, archive variables also can be mapped to these groups of systems and data.

Figure 6-1 shows the overall relationship between planning, policy, and infrastructure design processes, including TSM definitions which we will discuss in more detail in this chapter.

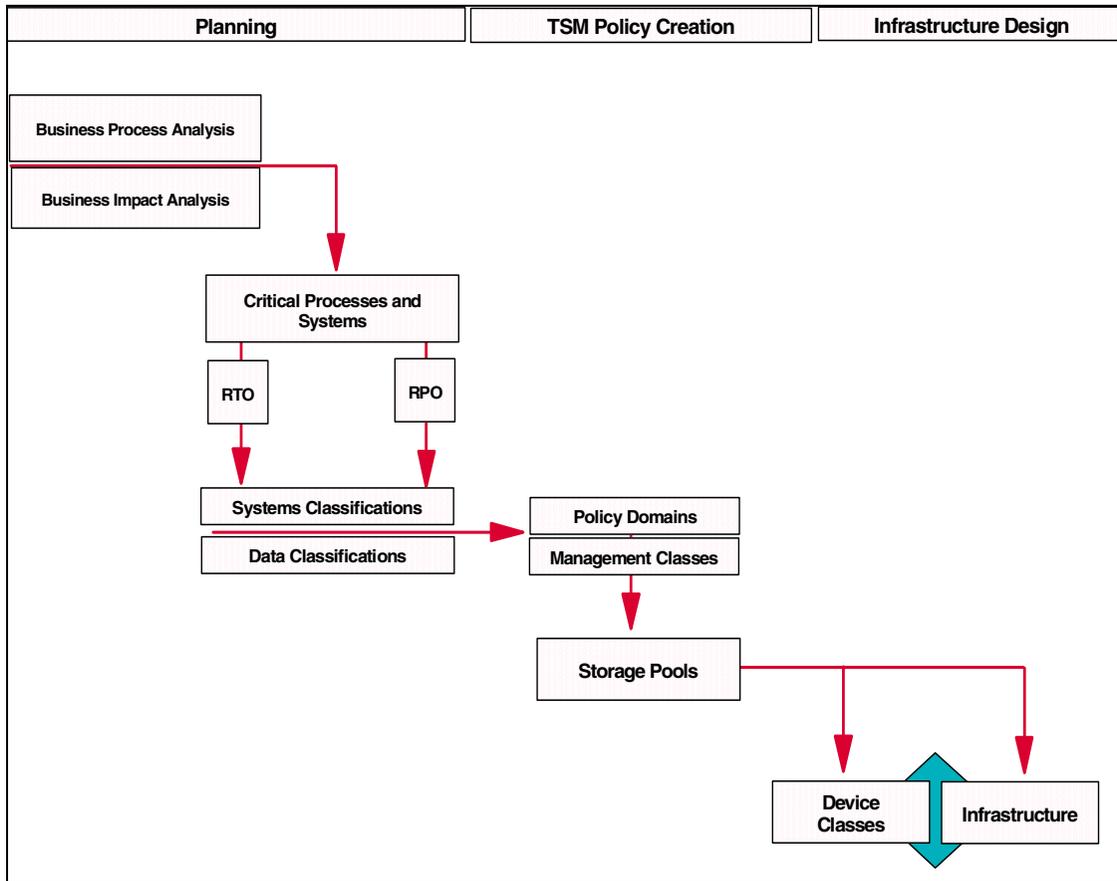


Figure 6-1 Planning, TSM policy creation, and infrastructure design

Organization of this data in a usable format is critical to the policy planning procedure for critical systems. Methodical policy planning ensures that objectives for data availability and recoverability will be met at both enterprise and systems levels.

6.1.1 TSM policy drivers

The elements of the BIA/DRP process which directly affect TSM policy creation range from enterprise wide requirements to systems specific requirements for data management and data availability. Organizing this data into a functional format presents a tremendous challenge to both Disaster Recovery planners and storage administrators. Starting from the top allows us to simplify the translation of business requirement to TSM specifics.

Enterprise level TSM policy drivers include:

- ▶ Regulatory and performance requirements (data availability, data retention, data security)
- ▶ Business Impact Analysis (critical systems, processes, dependencies)

System Level TSM policy drivers include:

- ▶ Recovery Time Objective (acceptable time to restore systems)
- ▶ Recovery Point Objective (how much data you can afford to recreate)
- ▶ Data classification (the relative importance of data)
- ▶ Versioning and retention requirements (application/business specific)
- ▶ Archive requirements
- ▶ On site and offsite contingency requirements

We recognize that every organization faces unique requirements and resources for enterprise Disaster Recovery, however these fundamental requirements consistently map to the TSM policy creation process.

6.1.2 TSM software architecture review

TSM software follows traditional client/server architecture principles. The TSM server coordinates the movement of data between clients, primary storage pools (disk and tape), disk copy pools, and tape copy pools. The server also controls several specific variables for each primary or copy storage pool, such as reclamation, collocation and migration thresholds. Figure 6-2 shows the general flow of data with TSM and related policy subjects.

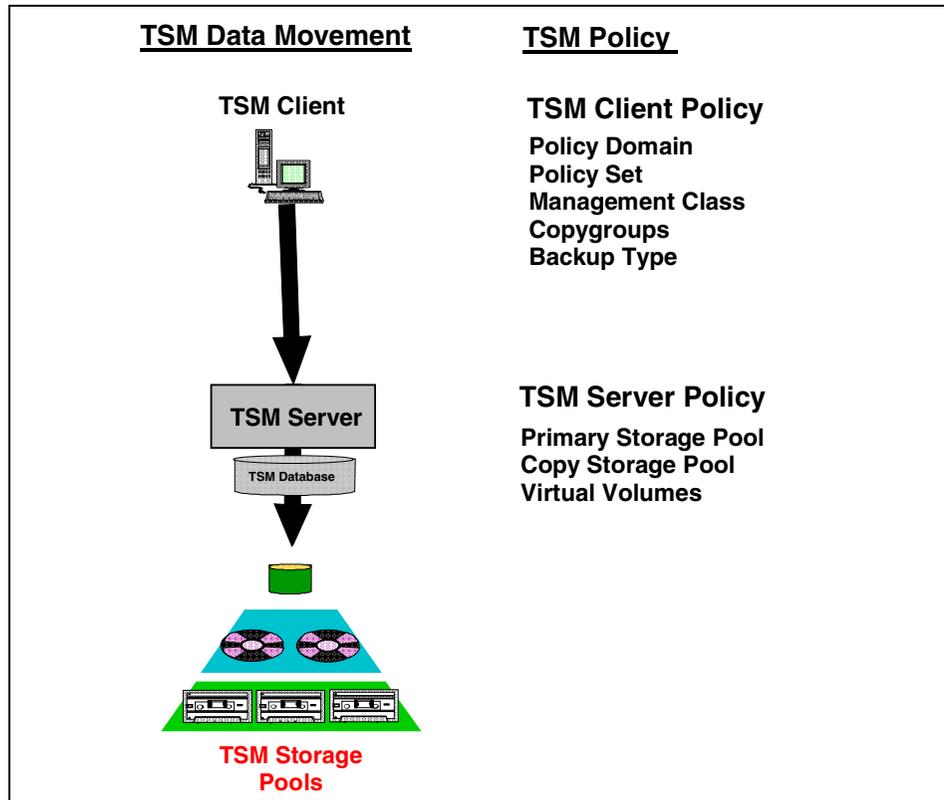


Figure 6-2 TSM data movement and policy subjects

TSM capabilities allow a great deal of flexibility, control, and granularity for policy creation. Due to the abundant choices that exist for policy creation, we can make general suggestions which can then be implemented in a variety of ways. In an enterprise environment, flexibility and choice in policy creation directly affects resource utilization, systems architecture, and overall feasibility of restore procedures.

6.1.3 TSM policy configuration overview

TSM policies are generated for both TSM clients and for the TSM server. Both policy specifications are distinct, yet interrelated to the overall performance and functionality of a TSM environment. The bottom line is that critical data needs to be mapped to TSM system resources, which will provide satisfactory performance for recovery.

TSM client policy

TSM client policy maps backup type, backup, archive, versioning, retention, and location variables to specific data. These policies are organized and defined in policy sets, management classes, and copy groups, as illustrated in Figure 6-3.

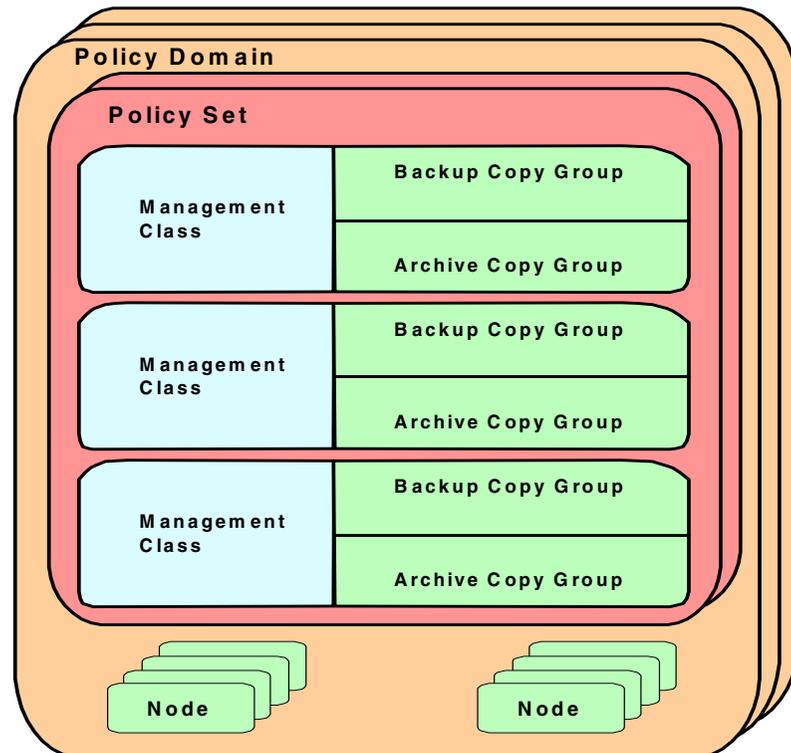


Figure 6-3 TSM policy elements

► Policy domains and policy sets

A policy domain is a policy object that contains policy sets, management classes, and copy groups that are used by a group of client nodes. Policy domains typically contain only one active policy set. A policy set contains a group of management class definitions. The policy domain is used typically to group together client nodes with common requirements or characteristics, for example, common operating systems such as Windows or UNIX servers. A similar approach can be applied to disaster recovery requirements.

For instance, critical systems can be organized into specific policy domains based on similarities in RTO requirements. Or, groupings could originate from departmental associations and RTO requirements. Other situations might warrant grouping systems by both platform and RTO requirements. The possibilities are numerous.

► **Management class**

A management class contains specific rules and is the link between rules and the client data. The rules are contained in constructs called copy groups, one for backup data and one for archive data. The management class allows files, directories, or specified groups of files to be linked or *bound* to how the server will manage them. Since many management classes can be built within a policy set, specific groups of client data can be organized into particular management classes. For example, a system’s data might be grouped into tiers of criticality, as described in 3.3.3, “Data classification procedures” on page 50, and mapped specifically to management classes for restore priority.

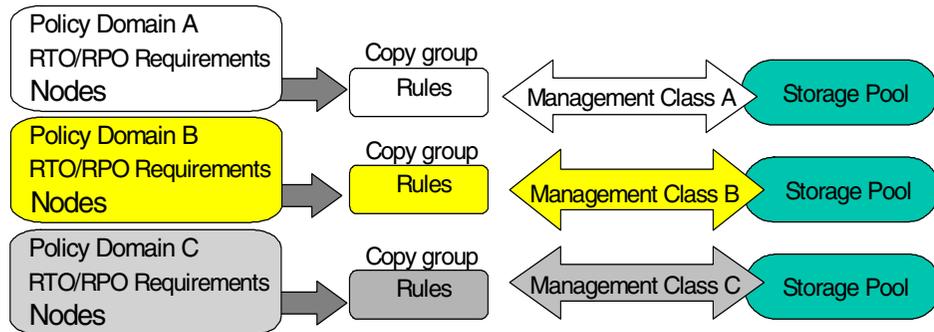


Figure 6-4 Relation of RTO and RPO requirements to management classes

Figure 6-4 shows how systems and data, based on criticality, can be bound to management classes to support specific recovery objectives. Within each management class, specific copy group rules are applied.

► **Copy Group**

Copy groups belong to management classes. Copy groups control the generations, storage destination, and expiration of backup and archive files. The copy group destination parameter specifies a valid primary storage pool to hold the backup or archive data. For performance and capacity considerations, the association of copy groups to specific storage pool sizing and planning procedures is extremely important.

► **Backup Methods**

Another critical variable for client policy is the type of backup method to be performed. Full, progressive, adaptive subfile, backup sets, or image backups can be performed for different sets of data.

TSM server policy

TSM server policies include definitions for primary storage pools and copy storage pools. Storage pools can reside on locally or SAN attached disk or tape,

or can be remote or offsite. Within each pool, specific variables can be set for migration thresholds and synchronous writing to additional storage pools. Other copy pool specific variables include collocation, reclamation, and virtual volume management.

► **Primary storage pools**

A typical TSM configuration includes a primary disk pool and primary tape pool for data backup. Typically, client data fills up the disk pool until thresholds are met, and then the data is migrated to the primary tape pool. TSM V5.1 allows data to be written to multiple storage pools in parallel (simultaneous writes to copy storage pools during client backup or archive). With LAN-free and server-free backups, data can move to disk or directly to tape without passing through the TSM server.

Multiple primary storage pools and copy pools can be designed within one TSM environment. These storage pools (which should be viewed as critical resources) are logically associated with specific management classes and copy groups. Every storage pool resource (on either disk or tape) should be viewed as an architectural element with its own capacity and performance attributes, because the storage resource performance is a key component for restore performance. Simply creating two large storage pools will rarely provide the level of control needed in a large data recovery event.

► **Copy Storage Pools**

Copy storage pools contain duplicates of data residing in primary storage pools. Tape volumes in a copy storage pool are typically sent manually or electronically vaulted to a disaster recovery site. Data is copied to copy storage pools incrementally. Administrative schedules can be used to automate the movement of data from storage pools to copy pools via the **BACKUP STGPPOOL** command.

Collocation and reclamation parameters affect the restore performance of tape storage pool data and should be considered for both primary and copy storage pools.

► **Storage pool migration**

Within primary storage pool definitions, migration threshold values are used to automate the movement of data between storage pools. When a primary storage pool fills up to a maximum percent utilization, it will write data to the next storage pool in the hierarchy until it reaches a low threshold. The cycle continues when the pool fills up again.

Storage pools, device classes, and devices

Storage pool policy directly impacts the availability of data for client restore operations. Disk storage pools usually provide the fastest restore capability, while tape storage pools provide cheaper, usually higher capacity, storage and

potentially higher access times. Several elements for tape resource planning are outlined in 6.4, “TSM tape storage pools” on page 107.

Figure 6-5 illustrates the connection between storage pools, device classes, and devices. Each copy group in a management class has a particular storage pool defined to it where data bound to that management class will be stored. The device class maps the particular physical devices (disk, drives and libraries) defined to TSM to the storage pools.

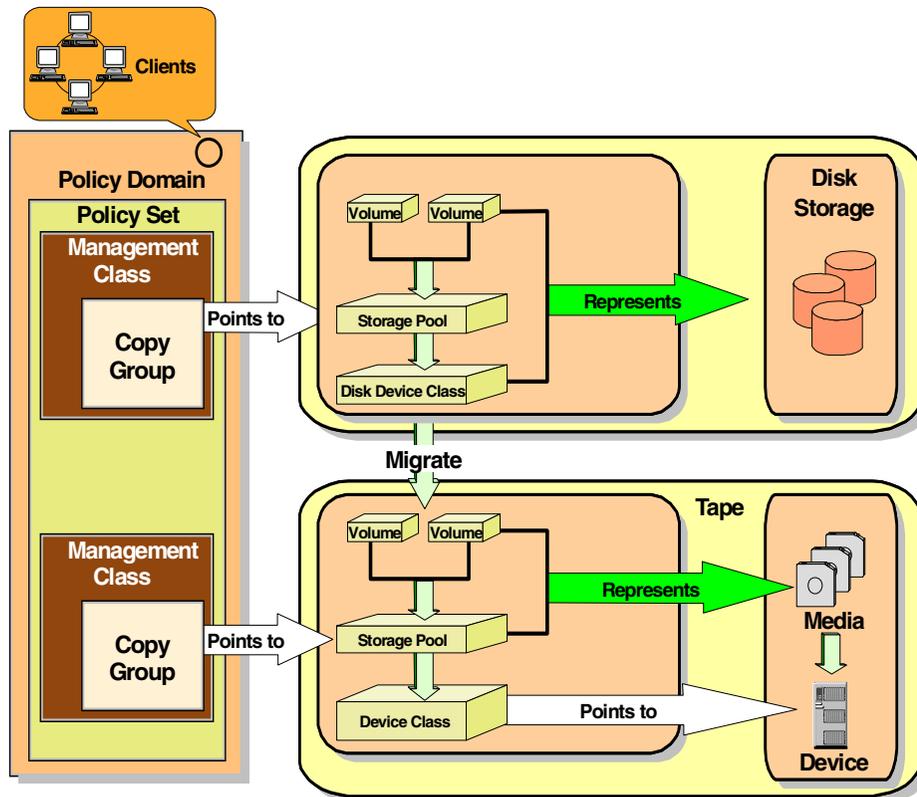


Figure 6-5 Policy, storage pools, device classes, and devices

Storage pools logically map to device classes, which can be designed to meet enterprise level requirements for performance and capacity, or more granular performance requirements based on specific RTO objectives for critical data.

6.2 Mapping DR requirements to TSM configurations

Once storage management policies are developed and critical data is grouped into specific policy sets and management classes, detailed planning procedures for server, disk, tape, and network resources should bridge policy requirements and technical architecture supporting the TSM application.

Every component in the TSM architecture must be designed to support the most rigorous RTO requirements in the enterprise.

TSM architectural design should not only factor in current requirements, but also future requirements for growth and change. A one to three year horizon for capacity planning should be built into any TSM sizing event. A three to five year technology plan should also be defined for TSM resources, including platform, disk, SAN, tape media, and tape devices. Strategic architectural planning and growth metrics are key.

At a high level, the overall data volume and data type (number and type of files) shape the basic metric for TSM solutions design. Policies for versioning, retention, data location, backup frequency, and archiving all directly impact the overall amount of data managed by TSM. The amount of data and how it is managed then affects the following elements of a TSM environment:

- ▶ TSM server (CPU, Memory, I/O)
- ▶ TSM database and recovery log volumes and layout
- ▶ TSM disk pool architecture and performance
- ▶ TSM tape pool architecture and performance
- ▶ TCP/IP and SAN configurations
- ▶ Tape media type

The most effective way to accomplish these goals is to methodically classify critical systems into policy domains, associate their data to management classes, and design storage pools and policies which functionally meet RTO performance requirements for each management class or group of management classes. For the most critical systems, consider dedicating storage pools exclusively to their management classes. Many such micro-environments can then exist within one logical TSM application environment.

6.2.1 A functional example

A good example of how management classes map to RTO requirements is the use of the DIRMC option in a client options file. Essentially, the DIRMC is a management class created specifically for directory data. All of the directories from the associated clients will be bound to the directory management class. Since directory data is restored before the actual files, this data can be located on a primary disk pool or a disk copy pool (in addition to tape pools), for fast

access and restore times. The DIRMC concept shows the relation of restore priority to management class policy. Note that simple directory structures are stored directly in the TSM database, without requiring storage pool space. More complex directory pool structures (which include ACLs for example) do get stored in a storage pool.

While this is a simple example, the same methodology can be applied to critical systems and critical data in an enterprise environment. Critical data can be grouped together and associated with storage pool resources and network connections which meet RTO requirements for recovery.

Once data policy requirements have been defined, the physical devices required (disk, tape drives and libraries) can be then selected to meet general performance and capacity requirements.

6.3 Creating TSM disk pools for recovery objectives

The disk device class is used for TSM disk storage pools. The disk device class is predefined by TSM and cannot be modified. However, virtually any number of disk volumes on any number of physical disk systems can be defined to TSM and subsequently used for disk storage pools. This means that disk storage pools can easily be designed to provide the required level of performance and availability (for example, by using locally or remotely mirrored disk systems). Disk storage pools can be designed specifically to meet performance and capacity requirements for critical management classes.

Figure 6-6 illustrates how management class definitions bind specific data to disk storage pools contained within the disk device class.

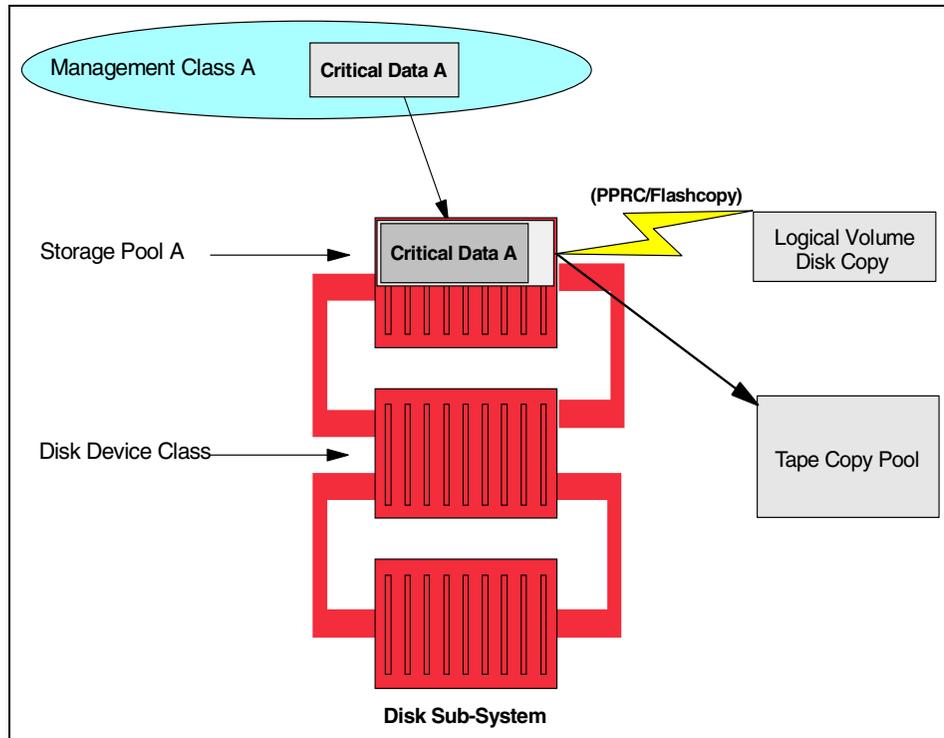


Figure 6-6 Management class relation to disk storage pools for critical data

While critical data remains in the onsite disk storage pool for fast access and restore, logical volume copies can be made using subsystem specific copy functions (such as FlashCopy or SRDF) and additional copies of the data can be made to tape copy storage pools. In this way, TSM strategically positions critical data for highly efficient disk based restores if available, while maintaining additional disk or tape copies in local and remote locations.

Within a disk storage pool definition, the MIGDELAY parameter is used to define retention periods for data within the storage pool. If designing individual disk storage pools to meet rigorous RTO requirements, this variable adds an excellent level of control for storage pool retention. MIGDELAY time settings should match backup frequencies for this kind of data.

High-end disk subsystems also offer the ability to create local or remote mirrored copies of data contained in disk storage pools. Advanced copy functions include ESS Flashcopy for local volumes and ESS PPRC for site to site volume replication or mirroring. These advanced copy functions can complement contingency requirements for mission critical management class data which needs to be stored on disk for availability reasons. These methods are discussed

in detail in 7.9, “TSM and remote disk replication” on page 151, and in *Introduction to SAN Distance Solutions*, SG24-6408.

6.3.1 Disk storage pool sizing

The disk pool will contain backup data that will eventually migrate to tape. Disk storage pool capacity planning must take into effect retention policies for disk storage pools and the daily volume of incremental and full backups.

Performance and capacity considerations must be made when sizing the disk device and related disk storage pools. We recommend the use of scalable and high performance disk subsystems for disk storage pools. The disk subsystem I/O performance and connectivity must support enterprise level demands for RTO.

6.4 TSM tape storage pools

Tape storage pools are typically used within TSM for both primary and copy storage pools. Tape is an ideal medium for backup and archive data storage because of its cost/MB ratio, ease of transporting offsite and high read/write performance. Tape devices are defined to TSM via library and drive definitions. Each physical library (of whatever tape technology) is associated with or mapped to a tape device class definition. The device class definition informs TSM about the type of drive being used, for example, format and capacity. Tape drives within a large tape library can be logically grouped to meet performance requirements for different groups of data.

6.4.1 Tape device performance considerations

Different tape technologies have performance characteristics, which can be quantitatively related to RTO requirements for any given volume of data. To meet RTO performance requirements for management classes, tape libraries are sized and assigned to tape storage pools. This mapping enables TSM to know the device characteristics of the storage pool media, and how to access it when data is written to or accessed from a storage pool. Figure 6-7 illustrates the fundamental reasoning behind designing tape storage pools to meet recovery time objectives for different classes of data.

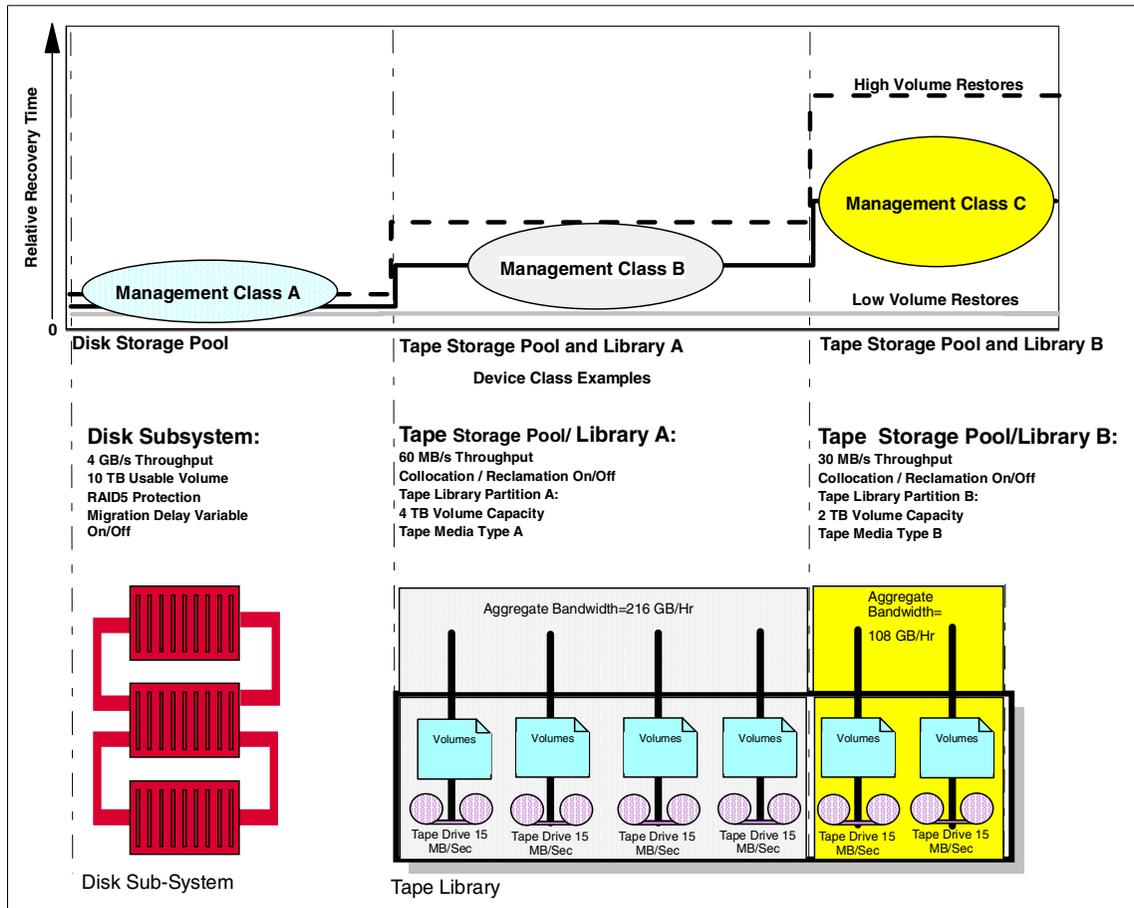


Figure 6-7 Relating RTO, management classes and device classes

As shown in Figure 6-7, tape drive performance can be estimated in terms of individual and aggregate throughput. Performance benchmarks for uncompressed data I/O are available from tape drive technology vendors. If multiple drives are installed in a library, the aggregate throughput capacity can be calculated and directly compared to RTO requirements for data volume versus time, as shown in Table 6-1.

Table 6-1 The basic tape drive aggregate performance equation

Drive I/O	# of Drives	Aggregate I/O	RTO (GB/hour)
15 MB/s	4	216 GB/hour	190

These measurements assume *perfect* performance, that is, they do not take account of tape seek and mount time, the number of volumes required for a set of data, and the physical distribution of data across one or multiple volumes. Since these variables are unpredictable, it is difficult to quantitatively factor them into preliminary real-world tape library performance estimates. Therefore, we advise you to reduce specified performance rates, say by 20 percent. Once a TSM environment is running in a live customer environment with real data, more accurate performance measurements and sizing metrics can be developed.

Collocation and reclamation settings also significantly impact tape restore performance, and must be taken into consideration for each tape storage pool definition. For each storage pool and associated management class data, these variables must be carefully considered for restore performance. Section 7.11.1, “Collocation considerations for offsite vaulting” on page 156 and 7.11.2, “Reclamation considerations for offsite vaulting” on page 157 discuss these factors in greater detail.

6.4.2 Tape library sizing fundamentals

At an enterprise level, tape library sizing must take into consideration current backup and archive volume, data characteristics, and high level performance requirements. High level tape library sizing methods are discussed in *Getting Started with Tivoli Storage Manager: Implementation Guide*, SG24-5416. The same capacity planning methods can be applied at a lower level against specific groups of management classes and device classes as needed.

Another important element of tape library selection involves strategic planning for tape media and devices. Enterprise tape considerations should include:

- ▶ Tape library hardware strategy for technology, performance, and scalability
- ▶ Tape format/density and tape media road map
- ▶ Hardware compression (on or off)
- ▶ Software or TSM compression standards for client data
- ▶ Tape drive connectivity (SAN or SCSI)

For DR technical planning, we suggest the consolidation of tape resources into a single tape media and library format. Alternate site tape infrastructure must match production site tape infrastructure. Otherwise, the tape media type could become a risk to recovery operations. We recognize that many environments host a variety of tape and media types, however we stress the importance of strategic planning for DR capability and operational efficiencies. Mixed tape media types limit the functionality of collocation and reclamation, and can make offsite tape management procedures overly complex.

An important factor for tape media management is the volume per cartridge, which inevitably increases with time and innovation. If tape pools contain mixed

capacity volumes, advanced tape management functions such as collocation or reclamation become less efficient. Currently, tape cartridge technology does not contain header information to describe the cartridge capacity, so all tapes within a tape library are treated as volumetric equals. A principal way to control tape volume efficiency is to allocate specific cartridge types to device classes and associated library partitions. New device classes with higher capacity cartridges can later be introduced. The migration process, or the **MOVE NODEDATA** or **MOVE DATA** commands can then be used to move data from one storage pool to another, where the storage pools use different device classes.

6.4.3 Versioning and tape storage pool sizing

The number and type of backup versions directly impacts tape storage pool sizing for capacity. Figure 6-8 illustrates a sample matrix which can be used to visualize and organize data volumes. Using this kind of matrix allows RTO requirements (through device class associations to storage pools), to be related to management classes, versioning policies, and the overall amount of volume required for each storage pool.

Number of Versions	Active + 30 Inactive Versions	Management Class - A4	Management Class - B4	Management Class - C4	Management Class - D4
	Active + 14 Inactive Versions	Management Class - A3	Management Class - B3	Management Class - C3	Management Class - D3
	Active + 7 Inactive Versions	Management Class - A2	Management Class - B2	Management Class - C2	Management Class - D2
	Active + 1 Inactive Version	Management Class - A1	Management Class - B1	Management Class - C1	Management Class - D1
		Disk Storage Pool	Tape Storage Pool A	Tape Storage Pool B	Tape Storage Pool C
		Storage Pool Type			

Figure 6-8 File versions, management class grouping, and storage pools

Once data is grouped by criticality into management classes, versioning and retention policies are applied. These policies impact directly the amount of storage volume required. Collectively, this overall volume of data must be known, in addition to growth forecasts, to accurately size the storage pools (both local and remote).

6.5 LAN and SAN design

Data can be restored from storage devices to TSM clients through LAN or SAN connections. For LAN-based TSM clients, data typically travels from either disk or tape, to the TSM server, then to the client. For LAN-free clients, data travels from tape or disk directly to the client over a Fibre Channel connection. Between clients and the TSM environment, the LAN or SAN infrastructure must be designed to support RTO requirements.

Table 6-2 shows the basic metrics for network topologies and data transfer rates. We divide each base rate by eight to convert bits to bytes. For example, 10Mbps is .125 Mbps. We then multiply by 3600 (seconds in an hour), and divide by 1000 and 1 million respectively to give the rates for GB/hour and TB/hour. These metrics do not take into consideration protocol overhead, and assume perfect disk, tape and application performance. Therefore, it is more typical to expect and achieve actual throughput numbers of between 60 and 80% of these numbers in real environments. LAN and SAN network connections must support RTO requirements for each critical TSM client.

Table 6-2 Network topologies and relative data transfer rates

Network topology	MB/second	GB/hour	TB/hour
10 Mb/sec Ethernet	1.25	4.5	.0045
100 Mb/sec Ethernet	12.5	45	.045
Gigabit Ethernet	125	450	.45
1 Gb SAN	125	450	.45
2 Gb SAN	250	900	.9

6.5.1 LAN considerations

Bottlenecks in production network environments are not uncommon. When restore traffic mixes with production network traffic, the network itself can compromise data recovery efforts. Backups are usually scheduled during off-peak times, but restore traffic is less predictable.

We suggest the implementation of dedicated LAN networks for enterprise storage management operations. Even a dedicated Gigabit ethernet network could become a bottleneck in a multi-system restore scenario. Duplexed network connections to the TSM server and redundantly designed network architectures can help to limit bottleneck issues.

6.5.2 SAN considerations

For large data management requirements (ranging from 10's of GB to multiple TB's), we encourage the use of SAN infrastructure for TSM clients. Fibre channel protocol (FCP) introduces relatively low protocol stack overhead for client systems, and eliminates the reliance on a single network for backup and recovery. LAN-Free backups over Fibre Channel also fail-over to LAN infrastructure in the unlikely event of a SAN disruption. Performance oriented design for the SAN infrastructure ensures maximum restore performance.

6.5.3 TSM server LAN, WAN, and SAN connections

Last but not least, the TSM server must host adequate connections to support RTO requirements for client data restore operations. Connectivity to networks, storage pools, and remote storage pools must also be properly designed and tested. Starting from the tape storage pool, performance measurements can be estimated to ensure the TSM architecture is free of bottlenecks.

As an example, disk storage device connections to TSM must support I/O requirements for the backup and migration traffic depicted in Figure 6-8. In this case, the amount of data moving in and out of the disk storage pool is 175 GB/hour at the peak hour of 6 a.m. A minimum of two 1 Gb Fibre Channel connections would provide the disk device sufficient connectivity to the TSM server. This presumes the disk system is also capable of sustaining that throughput. Doubling the Fibre Channel connections to four would improve both performance and redundancy.

6.5.4 Remote device connections

Design principles for long distance data transfer using DWDM or IP follow the same guidelines as SAN or LAN architecture. Single points of failure must be eliminated through redundant components and connections. Performance measurements must be made to support enterprise RTO requirements in the event of a major disaster recovery effort.

For remote operations, such as copy storage pool routines from a local storage pool to a remote tape copy pool, distance network architecture needs to support DR requirements to restore large volumes of critical data. If distance network

topologies only factor incremental backup traffic, the network architecture will not meet RTO requirements for full restores of critical systems.

Defining restore priorities and procedures is a necessary step in designing site to site network architecture for recovery operations. Critical system RTO requirements (in terms of GB or TB per hour) can then be mapped to bandwidth requirements for long distance architectures.

6.6 TSM server and database policy considerations

From a policy perspective, availability and performance of the TSM server is critical to enterprise disaster recovery. The server must be designed to sustain I/O requirements for all connected TSM clients and devices. In enterprise environments, multiple TSM servers can be implemented and centrally managed to help balance these requirements. Section 8.2, “Best practices for hardware design” on page 164 discusses TSM server and database sizing in detail.

The TSM database can (and should) be protected on many levels. The locally attached disk storage must support mirroring or RAID-5 protection for the TSM database and recovery logs. The TSM database must be backed up as frequently as the most frequently backed up data in the enterprise. Specifically, the RTO and RPO for TSM must be equally or more aggressive than the most mission critical system in the enterprise.

For instance, if a mission critical file system is backed up every hour, the same backup policy must be applied to the TSM database. Large TSM environments often take frequent incremental backups and daily full backups of the TSM server database. Since TSM client data can only be accessed through the TSM application, the TSM database must be routinely backed up and safeguarded.

TSM policy should also be documented in the DRP appendix. Detailed information should be stored in an offsite location, along with current backup copies of the TSM database, the TSM DRM output, and hard copies of the TSM DRM plan.

6.7 Recovery scheduling

Backup operations are almost always automated by using TSM scheduling facilities, however recovery operations almost always occur ad hoc and are manually initiated by either the user or the administrator initiated commands. For DR of enterprise data, recovery priorities and the sequence of data recovery events must be well defined in the DR plan.

Assuming the TSM systems design supports RTO requirements for critical data, a logical sequence of events must still be planned to recover critical application data. System classifications and data classifications should be analyzed to determine the most logical recovery sequence for your environment.

If specific storage or network resources are used to recover multiple critical systems data in parallel, the recovery plan must accommodate a logical sequence of events to avoid resource contention. Without recovery scheduling, tape drives or network resources can be overwhelmed and effectively cause bottlenecks for large recovery operations. Other considerations, such as the availability of tape volumes in a recovery site library play a critical role in recovery efficiency. Figure 6-9 shows an example resource utilization and recovery sequence plan for a TSM environment which contains multiple tape libraries and storage pools.

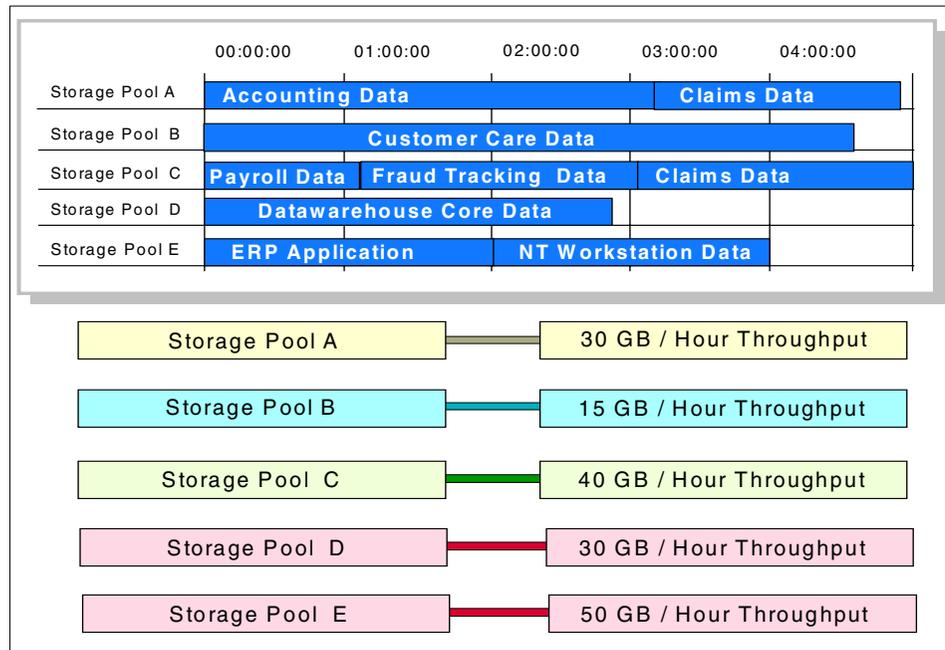


Figure 6-9 Sample recovery schedule using resource scheduling

Project Management tools, such as Microsoft Project, can be used to model resource utilization in a large TSM environment. Traditionally these tools factor employee time and costs versus specific tasks on a timeline. Instead, TSM resource utilization can be factored against RTO time requirements using these software tools. Parallel recovery events can also be displayed on Gantt charts to help identify a complex sequence of events involving multiple systems and resources.



TSM tools and building blocks for Disaster Recovery

The goal of this chapter is to introduce DR tools and building blocks available with IBM Tivoli Storage Manager (TSM). Such tools include: database page shadowing, database and recovery log mirroring, storage pool manipulation, client backup operations, backup methods and topologies, TSM Disaster Recovery Manager (DRM), TSM server-to-server communications, TSM server-to-server virtual volumes, TSM and high availability clustering, TSM and remote disk replication, and TSM traditional and electronic tape vaulting.

For additional references describing these TSM components we recommend the following redbooks and manuals:

- ▶ *Tivoli Storage Manager Version 5.1 Technical Guide*, SG24-6554
- ▶ *Getting Started with Tivoli Storage Manager: Implementation Guide*, SG24-5416
- ▶ The *Administrator's Guide* for your TSM server platform. See “Related publications” on page 391 for platform-specific titles.
- ▶ The *IBM Tivoli Storage Manager Client Backup-Archive Guide* for your TSM platform. See “Related publications” on page 391 for platform-specific titles.

7.1 Introduction

TSM protects enterprise data. This section will provide an overview of concepts and functions in TSM that could help you protect your environment and recover from a disaster. These building blocks for DR discussed here and include: database and recovery log mirroring, database page shadowing, storage pool manipulation, the variety of backup methods and topologies, TSM Disaster Recovery Manager (DRM), TSM server-to-server communications, TSM server-to-server virtual volumes, TSM and high availability clustering, TSM and remote disk replication, and TSM traditional and electronic tape vaulting. We have assumed for the purpose of this section that the reader has some knowledge of the following major TSM components as shown in Figure 7-1.

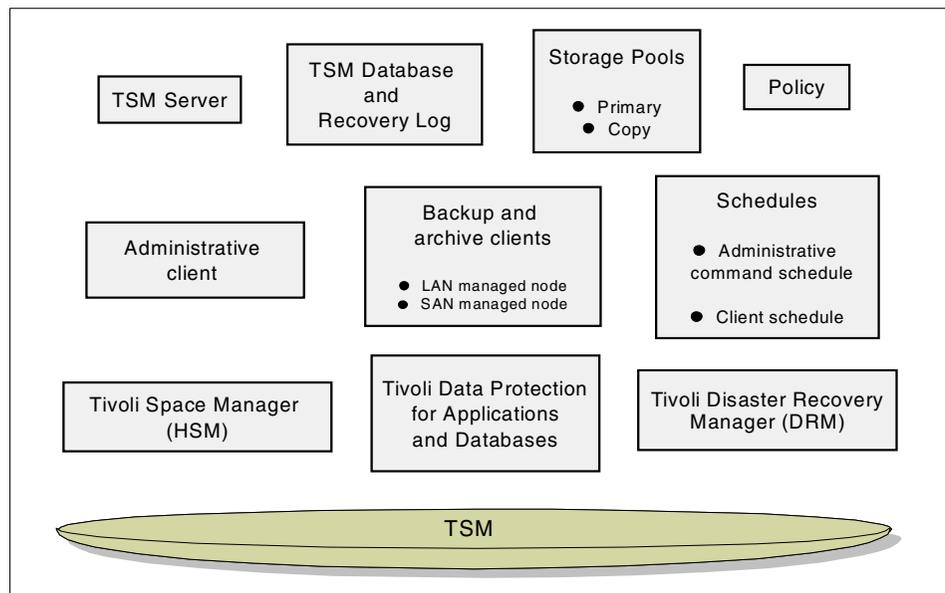


Figure 7-1 TSM components

A comprehensive discussion of each of these items could also be found in the books referred to in the introduction to this chapter.

7.2 The TSM server, database, and log

A Tivoli Storage Manager environment begins with a TSM server and its associated database and recovery log. These components provide the following functions:

TSM server: The TSM server is the program that provides backup, archive, restore, retrieve, space management (HSM), and administrative services to client systems (also known as nodes). There can be one or many TSM servers in your environment to meet data protection needs or balance resources. The TSM server contains and uses its own dedicated database and recovery log.

TSM database: A collection of information (or metadata) about all objects managed by the server, including policy management objects, users and administrators, client nodes, and client data. Because the TSM database is often considered the heart of the TSM server (and can be recovered in the event of a TSM server failure) a variety of methods exist to protect and recover it.

TSM recovery log: A log of updates that are about to be written to the database. The server uses the recovery log as a scratch pad for the database, recording information about client and server transactions while the actions are being performed. The log can be used to recover from system and media failures.

Figure 7-2 shows these components.

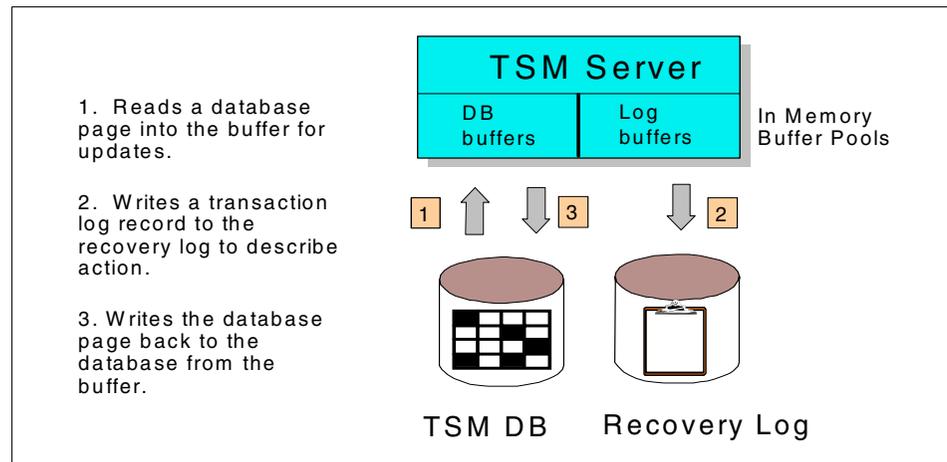


Figure 7-2 TSM server, database, and recovery log overview

When a transaction occurs on the TSM server or between it and a TSM client, the TSM server updates (reads and writes to) the TSM database and recovery log as required. Backing up files from a client node to the server and storing them

in a storage pool is an example of a transaction. When a transaction occurs, the server:

1. Reads a database page in the database buffer and updates it.
2. Writes a transaction log record to the recovery log describing the action occurring and associates it with the database page in case the page needs to be rolled back during recovery.
3. Writes the database page to the database, releasing it from the buffer pool.

7.2.1 TSM database page shadowing

The TSM database page shadowing process enables a mirrored write of the last batch of pages written to the server database. See Figure 7-3. Database page shadowing can be explicitly enabled and disabled through the **DBPAGESHADOW Yes | No** option in TSM. With **DBPAGESHADOW=Yes** the server mirrors the latest batch of pages written to the server database. In this way if an outage occur that affects both mirrored volumes, the server can recover pages that have been partially written.

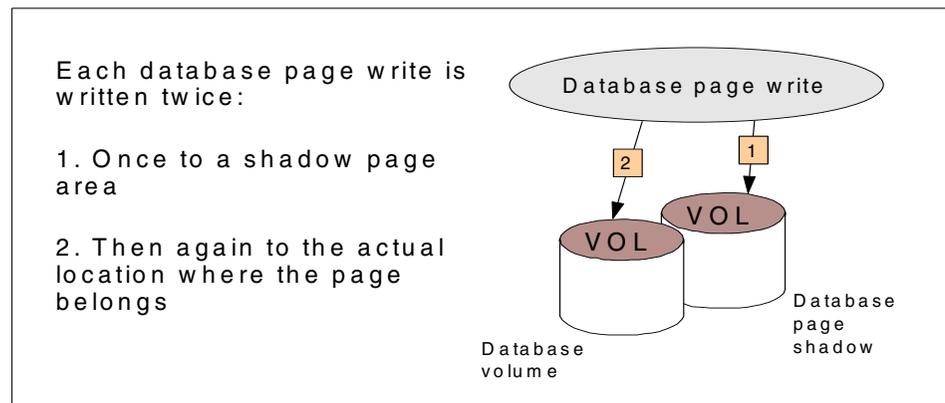


Figure 7-3 TSM database page shadowing

During server startup the pages in the shadow area are compared with those in the real location in the database to determine if any have been partially written. If partially written pages are detected in the shadow page area, processing simply continues as before, that is, the real pages are rebuilt and written during transaction recovery. If the pages in the shadow are intact, but one or more pages in their real location are partially written, the pages in the shadow page area are copied over the real page addresses.

Important: Database page shadowing is not a replacement for database mirroring.

7.2.2 TSM database and recovery log mirroring

The database contains information about the client data in your storage pools. The recovery log contains records of changes to the database. If you lose the recovery log, you lose the changes that have been made since the last database backup. If you lose the database, you lose all your client data.

You can prevent the loss of the database or recovery log due to a hardware failure on a single drive, by mirroring them on separate physical drives. Mirroring simultaneously writes the same data to multiple volumes as shown in Figure 7-4. However, mirroring does not protect against a disaster or a hardware failure that affects multiple drives or causes the loss of the entire system. While Tivoli Storage Manager is running, you can dynamically start or stop mirroring and change the capacity of the database. TSM provides 2-way or 3-way mirroring.

Database and recovery log mirroring provides the following benefits:

- ▶ Protection against database and recovery log media failures.
- ▶ Uninterrupted operations if a database or recovery log volume fails.
- ▶ Avoidance of costly database recoveries.

Mirroring however comes at the following costs:

- ▶ Mirroring doubles (or triples) the required disk for those volumes that are mirrored.
- ▶ Mirroring can potentially affect performance.

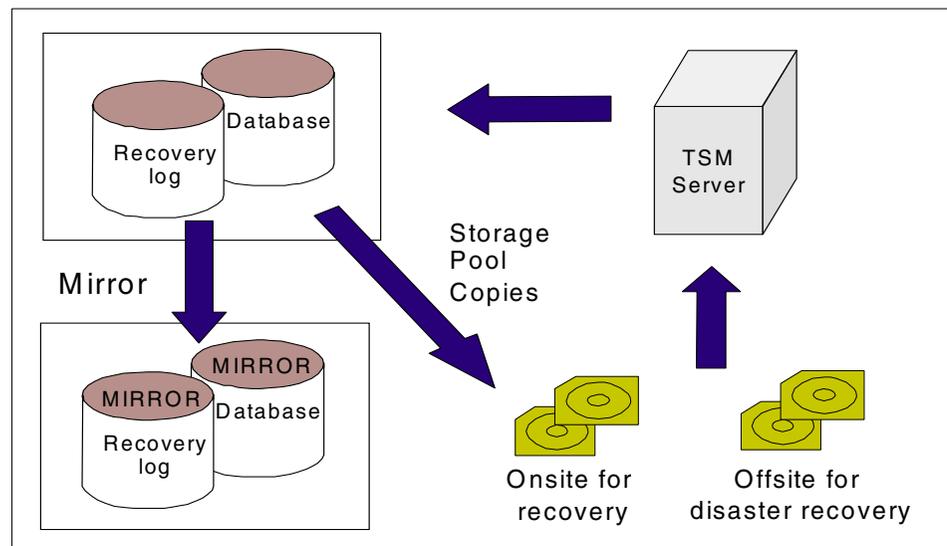


Figure 7-4 TSM database and recovery log mirroring

Mirroring can be crucial in the recovery process. Consider the following scenario. Because of a sudden power outage, a partial page write occurs. The recovery log is corrupted and not completely readable. Without mirroring, recovery operations cannot complete when the server is restarted. 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 data.

7.3 TSM storage pools

The TSM server stores a client's backed up and archived files in data storage pools, as shown in Figure 7-5. Data storage pools can be configured on direct access storage (hard disks), optical media, and sequential tape media. These data storage pools are defined within a TSM storage pool hierarchy. Files are placed by the TSM server on different storage pools according to the desired storage management policy. They can then be automatically moved to other devices or other states (for example, offsite) to satisfy business requirements for space utilization, performance, and recovery. A storage pool should be configured to match user requirements for data with the physical characteristics of storage devices. For example, if a user has a requirement for immediate access to their data on the TSM server then a disk storage pool should be considered. These primary storage pools can also be protected by backing them up to copy storage pools.

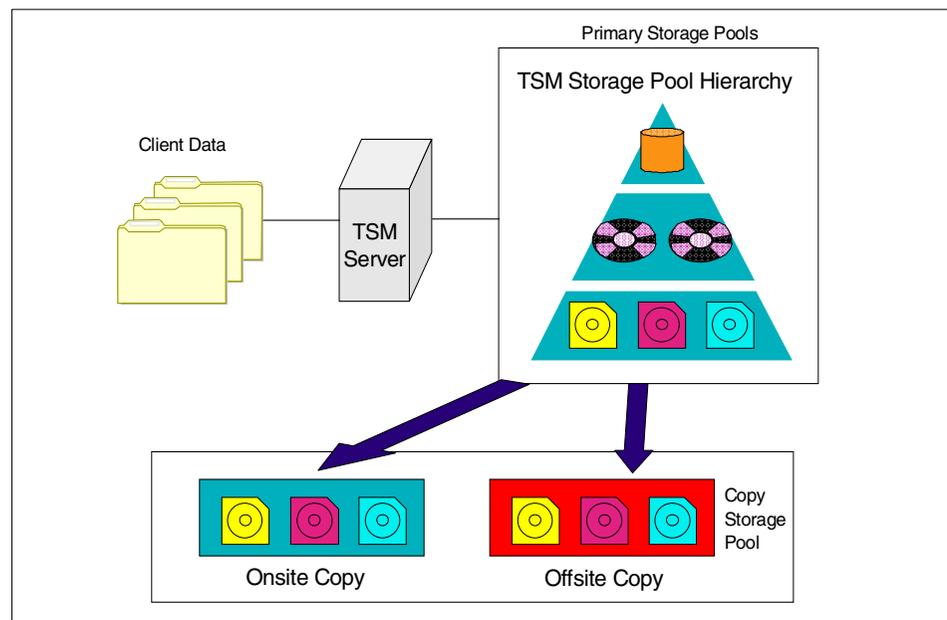


Figure 7-5 TSM storage pools

A primary storage pool can use random access storage (DISK device class) or sequential access storage (for example, tape, optical or FILE device classes). Data which comes from the TSM client to the server, whether backed up, archived or HSM migrated, is always stored in a primary storage pool.

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 (using the **COPY STGPPOOL** command) to another storage pool defined as a copy. The copy storage pool contains all current versions of all files, active and inactive, exactly as they appear in the primary storage pool. A copy storage pool provides protection 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 which is lost or found to be defective. When a client attempts to restore a file which was on this volume, the server will automatically use a copy storage pool volume (if available onsite) containing that file to transparently 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 sequential access storage (for example, tape, optical or FILE device classes), or copy storage pools can also be created remotely on another Tivoli Storage Manager server, therefore providing electronic vaulting.

Tip: TSM supports an extensive list tape drives, autoloaders, libraries and optical devices. For a full list of supported devices please refer to the following Web site:

http://www.tivoli.com/support/storage_mgr/requirements.html

7.4 Requirements for TSM server Disaster Recovery

In the event of a system failure or site disaster, TSM is designed with several means to restore the TSM server. These include: server recovery using database and storage pool backups, restoring your server using mirrored volumes, restoring storage pool volumes, auditing a storage pool volume, and correcting damaged files. In any case, to ensure the ability to restore your TSM server you should have the following components stored offsite.

- ▶ A full database backup
- ▶ Any incremental database backups between the last full backup and the point-in-time to which you are recovering
- ▶ Copy storage pool volumes
- ▶ On removable media: CD, tape, ZIP drive or diskette, or as printouts:
 - Server options file

- Volume history file
- Device configuration file with the applicable device information (library, drive, and device class definitions)
- Database and recovery log setup (the output from detailed queries of your database and recovery log volumes)

Database backups

TSM can perform full and incremental database backups while the server is running and available to clients. The backup media can then be stored onsite or offsite and can be used to recover the database up to the point of the backup. You can run full or incremental backups as often as needed to ensure that the database can be restored to an acceptable point-in-time (Figure 7-6).

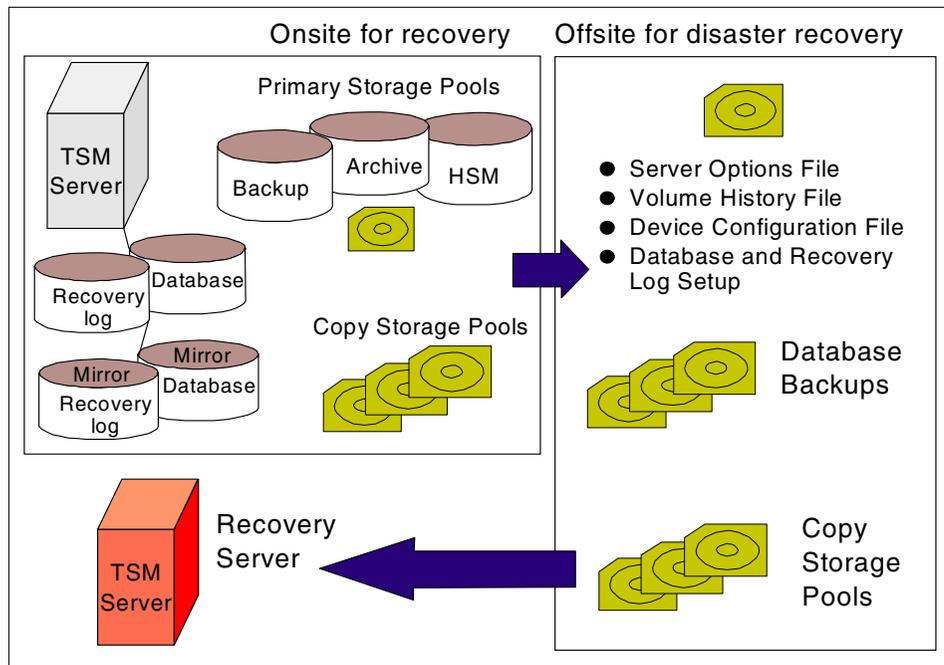


Figure 7-6 TSM server Disaster Recovery requirements

A snapshot database backup can also provide disaster protection for the TSM database. A snapshot backup is a full database backup that does not interrupt the current full and incremental backup series. Snapshot database tapes can then be taken offsite for recovery purposes and therefore kept separate from the normal full and incremental backup tapes.

You can provide even more complete protection if you specify rollforward mode for the TSM database recovery log. With rollforward mode and an intact recovery log, you can recover the database up to its most current state (the point at which the database was lost). For the fastest recovery time and greatest availability of the database, mirror both the database and recovery log, and periodically back up the database.

Volume history

Every volume that is used by TSM for storage pools and server database backups, is tracked within the server database. This information is very important because it indicates which volume holds the most recent server database backup. Volume history information is stored in the database, but during a database restore, it is not available from there. To perform a restore, therefore, the server must get the information from the volume history file. The volume history file can be maintained as a text file by specifying its name and location with the **VOLUMEHISTORY** option in the `dsmserv.opt` file. It is very important to save the volume history file regularly with the **BACKUP VOLHISTORY** command. You can specify multiple volume history files by repeating the **VOLUMEHISTORY** stanza in the server options file. If you use DRM, then it will automatically save a copy of the volume history file in its Disaster Recovery Plan file.

Device configuration

The device configuration file contains information required to read backup data. This information includes devices class definitions, library definitions, drive definitions, server definitions. This information is stored in the database, but during a database restore, it is not available from there. To perform a restore, therefore, the server must get the information from the device configuration file. The device configuration file can be maintained as a text file by specifying its name and location with the **DEVCONFIG** option in the `dsmserv.opt` file. It is very important to save your device configuration file regularly with the **BACKUP DEVCONFIG** command. You can specify multiple device configuration files by repeating the **DEVCONFIG** stanza in the server options file. If you use DRM, then it will automatically save a copy of the device configuration file in its Disaster Recovery Plan file.

When device information is updated in the database, it is also updated in the device configuration file. The device information must match the devices configured on the system where the restore will be performed. You may have to edit those commands in an existing file so that they match, for example if you have a single manual tape drive rather than a library at the recovery location.

Procedures and recommendations for TSM server and client recovery are discussed in more detail in:

- ▶ Chapter 8, “IBM Tivoli Storage Manager and DRM” on page 163
- ▶ Chapter 10, “Solaris client bare metal recovery” on page 229
- ▶ Chapter 11, “AIX client bare metal recovery” on page 239
- ▶ Chapter 12, “Windows 2000 client bare metal recovery” on page 267
- ▶ Chapter 13, “Linux client bare metal recovery” on page 289.

7.5 TSM backup methods and supported topologies

This section provides an overview of TSM backup methods, operations, and supported topologies. For example, a variety of client backup operations are available that can provide different levels of backup efficiency (such as adaptive subfile backup) or restore (for example, image restore). Similarly a variety of supported backup methods and topologies exist in a TSM environment, including TSM support for LAN-free and server-free backups.

This information is provided to help you determine which approach you may want to consider as part of your Disaster Recovery strategy. It is not intended to exhaustively discuss all backup scenarios.

7.5.1 Client backup and restore operations

It is important to understand TSM options for client backup and restore operations. It is also important to understand the characteristics of each of these operations because each method may have an effect on backup and restore efficiency, retention periods, portability, CPU utilization, connection time, and network utilization. Table 7-1 describes the various client backup and restore operations supported with TSM and provides a description, usage, and restore options. The standard backup method that Tivoli Storage Manager uses is called progressive incremental backup. It is a unique and efficient method for backup.

Table 7-1 Summary of client backup and restore operations

Type of backup operation	Description	Usage	Restore options
Progressive incremental backup	<p>The standard method of backup used by the Tivoli Storage Manager backup/archive client. After the first, full backup of a client system, incremental backups are done. Incremental backup by date is also available.</p> <p>No additional full backups of a client are required after the first backup.</p>	<p>Helps ensure complete, effective, policy-based backup of data. Eliminates the need to retransmit backup data that has not been changed during successive backup operations.</p>	<p>The user can restore just the version of the file that is needed (depending on the retention parameters). Tivoli Storage Manager does not need to restore a base file followed by incremental backups. This means reduced time and fewer tape mounts, as well as less data transmitted over the network.</p>
Selective backup	<p>Backup of files that are selected by the user, regardless of whether the files have changed since the last backup.</p>	<p>Allows users to protect a subset of their data independent of the normal incremental backup process.</p>	<p>The user can restore just the version of the file that is needed. TSM does not need to restore a base file followed by incremental one. This means reduced time, fewer tape mounts, and less data over the network.</p>
Adaptive subfile backup	<p>A method that backs up only the parts of a file that have changed since the last backup. The server stores the base file and subsequent subfiles (the changed parts) that depend on the base file. The process works with either the standard progressive incremental backup or with selective backup.</p>	<p>Maintains backups of data while minimizing connect time and data transmission for the backup of mobile and remote users.</p> <p>Applicable to clients on Windows systems.</p>	<p>The base file plus a maximum of one subfile is restored to the client.</p>

Type of backup operation	Description	Usage	Restore options
Journal-based backup	Aids all types of backups (progressive incremental backup, selective backup, adaptive subfile backup) by basing the backups on a list of changed files. The list is maintained on the client by the journal engine service of the Tivoli Storage Manager backup/archive client.	Reduces the amount of time required for backup. The files eligible for backup are known before the backup operation begins. Applicable to clients on Windows NT and Windows 2000 systems.	Journal-based backup has no effect on how files are restored; this depends on the type of backup performed.
Image backup	Full volume backup. Nondisruptive, online backup is possible for Windows 2000 and Linux clients by using the Tivoli Storage Manager snapshot function.	Allows backup of an entire file system or raw volume as a single object. Can be selected by backup-archive clients on UNIX and Windows systems. Used by Windows clients that are using server-free data movement.	The entire image is restored.
Image backup with differential backups	Full volume backup, which can be followed by subsequent differential backups.	Used only for the image backups of NAS file servers, performed by using Tivoli Data Protection for NDMP.	The full image backup plus a maximum of one differential backup are restored.
Backup using hardware snapshot capabilities	A method of backup that exploits the capabilities of IBM Enterprise Storage Server FlashCopy and EMC TimeFinder to make copies of volumes used by database servers. The Tivoli Data Protection for applications and databases use the volume copies to back up the database volumes.	Implements high-efficiency backup and recovery of business-critical applications while virtually eliminating backup-related downtime or user disruption on the database server.	See Section 7.5.5, "Split-mirror/point-in-time copy backup using SAN" on page 130 for details.

Type of backup operation	Description	Usage	Restore options
Archive	The process creates a copy of files and stores them for a specific time.	Use for maintaining copies of vital records for legal or historical purposes. If you frequently create archives for the same data, consider using instant archive (backup sets) instead. Frequent archive operations can create a large amount of metadata in the server database resulting in increased database growth and decreased performance of expiration server operations.	The selected version of the file is retrieved on request.
Instant archive	The process creates a backup set of the most recent versions of the files for the client, using files already in server storage from earlier backup operations.	Use when portability of the recovery media or rapid recovery of a backup-archive client is important. Also use for efficient archiving.	The files are restored directly from the backup set. The backup set resides on media that can be mounted on the client system, for example, CD, tape drive, file system. The TSM server does not have to be contacted for the restore process, so the network and TSM server are not used.

7.5.2 Traditional LAN and WAN backup topology

In a traditional LAN and WAN environment the Tivoli Storage Manager backup and archive client or application would read data from locally attached disks and send it over the LAN to the Tivoli Storage Manager backup server as shown in Figure 7-7. The server receives the data then writes it out to its storage pool — tape, disk, or optical media — based on predefined policies and server configuration. Data is read and written by both the Tivoli Storage Manager client and Tivoli Storage Manager server machines. In addition, control information is also sent over the LAN to the Tivoli Storage Manager server.

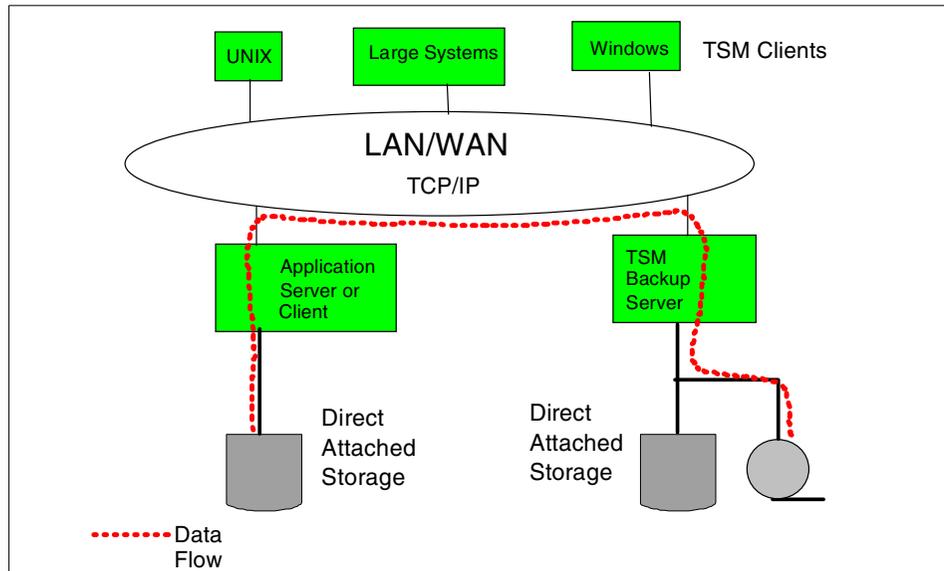


Figure 7-7 TSM LAN and WAN backup

7.5.3 SAN (LAN-free) backup topology

SAN technology provides an alternative path for data movement between the TSM client and the server. Shared storage resources (disk, tape) are accessible to both the client and the server through the Storage Area Network. Data movement is off-loaded from the LAN and from the server processor and allows for greater scalability. LAN-free backups decrease the load on the LAN by introducing a Storage Agent. The Storage Agent can be thought of as a small TSM server (without a database or recovery log) which is installed and run on the TSM client machine. The Storage Agent handles the communication with the TSM server over the LAN but sends the data directly to SAN attached tape devices, relieving the TSM server from the actual I/O transfer. A LAN-free backup environment is shown in Figure 7-8.

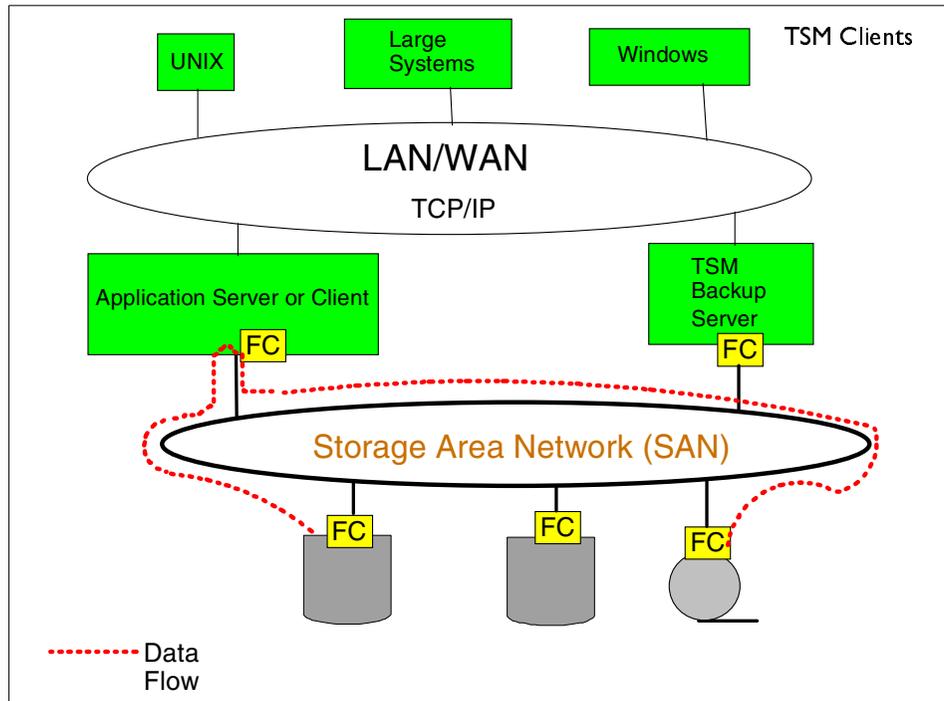


Figure 7-8 TSM LAN-free backup

7.5.4 Server-free backup

Server-free backup/restore capability is introduced into IBM Tivoli Storage Manager (TSM) in Version 5. In a server-free backup environment, data is copied directly from the SAN attached Tivoli Storage Manager client disk to the SAN attached tape drive via the SAN Data Gateway data mover as shown in Figure 7-9. The Storage Agent used in LAN-free backups is not used. The data movement is actually done by a SAN Data Gateway (SDG) or similar device on the SAN. Therefore both TSM client and server machines do not have to read and write the data at all. The TSM server sends commands to the SDG device to tell it which blocks to move from which SAN attached disk to which SAN attached tape device. The data is actually copied rather than moved from one location to another. This provides a way to back up and restore large volumes of data between client-owned disks and storage devices using a method that considerably reduces overhead on the TSM server and the client. Only volume images and not individual files can be moved by server-free data movement. The data is transferred block by block rather than by doing file I/O. Both raw and NTFS volumes can be backed up using server-free data movement.

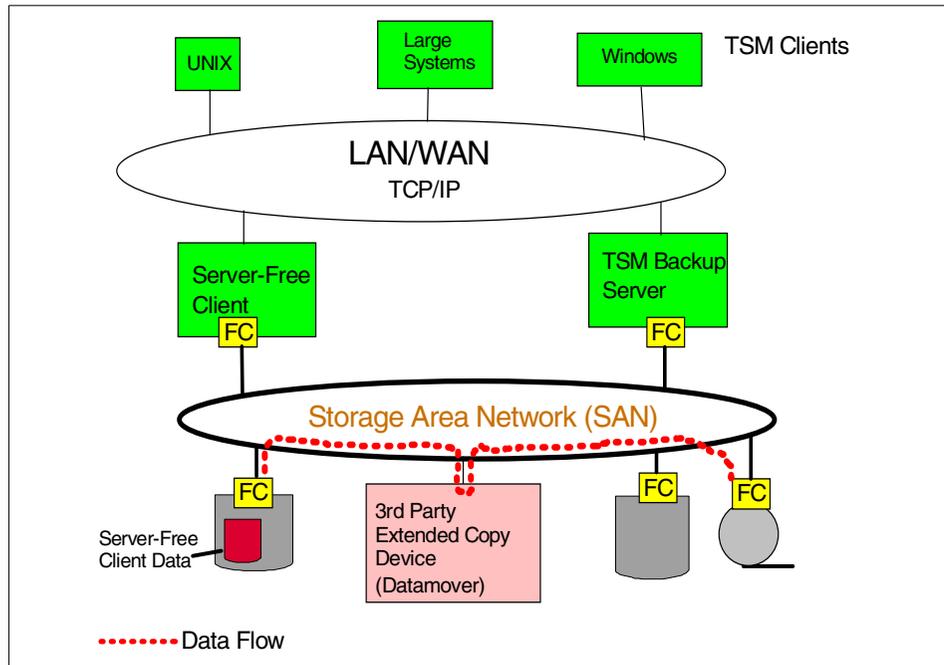


Figure 7-9 TSM Server-free backup

Data that has been backed up using server-free data movement can be restored over a server-free path, over a LAN-free path, or over the LAN itself. The impact on application servers is now minimized with Server-Free Data Movement. It reduces both TSM client and server CPU utilization. The use of a SCSI-3 extended copy command causes data to be transferred directly between devices over the SAN or SCSI bus. The data mover device must support the SCSI-3 EXTENDED COPY command, which conforms to the ANSI T10 SPC-2 standard. The data mover device can be anywhere in the SAN, but it has to be able to address the LUNs for both the disk and tape devices it is moving data between.

7.5.5 Split-mirror/point-in-time copy backup using SAN

A split-mirror/point-in-time backup occurs when a copy volume generated by Operating System mirroring or a hardware assisted instant copy function (as found on many of today's high-end storage systems) is backed up to a TSM server, as shown in Figure 7-10. Such a backup method virtually eliminates the backup-related performance impact on the production host. This approach is facilitated and automated with the IBM Tivoli Storage Manager for Hardware component by integrating IBM's Enterprise Storage Server FlashCopy function with IBM Tivoli Storage Manager and its database protection capabilities for Oracle, R/3, and DB2. Likewise, EMC's Symmetrix TimeFinder integrates with

IBM Tivoli Storage Manager and its protection capabilities for Oracle and R/3. This *Copy-Backup* procedure adds value to storage and backup procedures, because it helps ensure that essential applications can continue to run 24 x 7 with minimal backup-related impact.

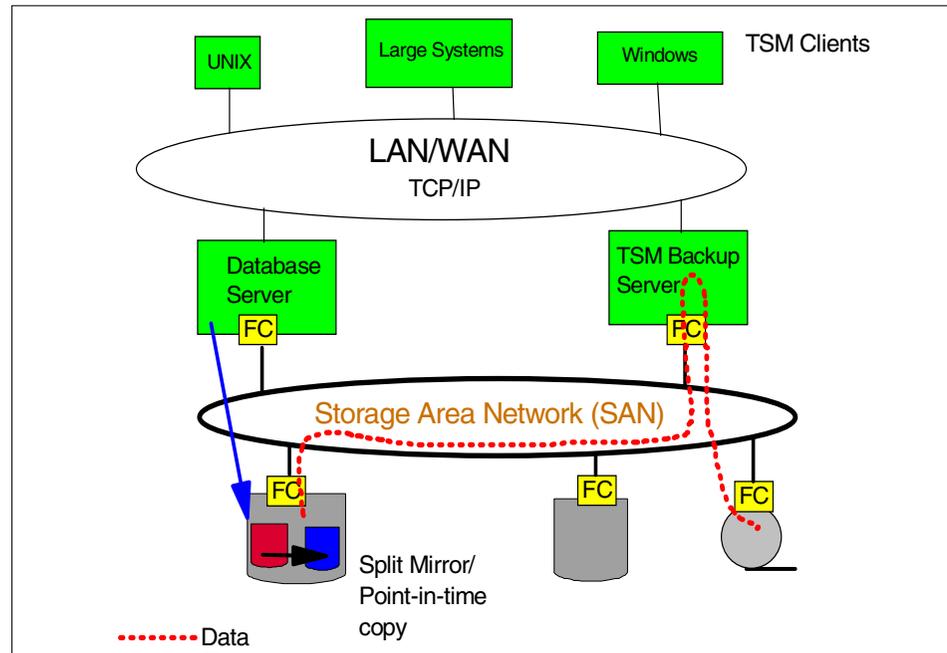


Figure 7-10 TSM split-mirror/point-in-time copy backup

7.5.6 NAS backup and restore

The IBM Network Attached Storage products have been pre-installed with the TSM Client which enables an existing or planned TSM server to back up the data in the NAS system. This backup client is designed to provide file level and sub-file level backup and restore functionality. A TSM environment integrated with NAS systems can be used to manage a broad range of data storage, recovery, and availability functions across the entire computing infrastructure. Based on the TSM server's configuration, the final destination of the NAS appliance's backup may either be located in the TSM server's disk storage or an attached tape subsystem. Persistent images must be created, using the pre-loaded Persistent Storage Manager (PSM) software before activating this backup function. Automated scheduling to back up these PSM images can then be configured in the TSM server, see Figure 7-11. Just like a standard TSM client an *option* file is created and stored on the NAS system. For more information on backing up IBM NAS appliances, see *IBM TotalStorage NAS Backup and Recovery Solutions*, SG24-6831.

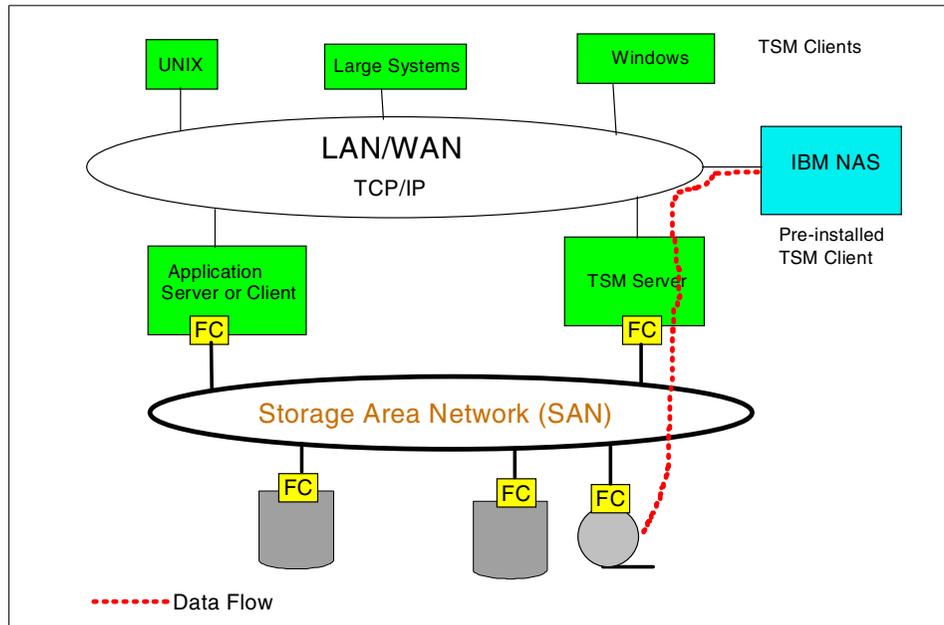


Figure 7-11 IBM TotalStorage NAS backup

The Tivoli Data Protection for NDMP product is now available with Tivoli Storage Manager (TSM) Version 5 Extended Edition. It provides backup and recovery support on TSM servers for NAS file servers from Network Appliances. NAS file servers often require a unique approach to providing backup and recovery services, because these file servers are not typically intended to run third-party software. The NAS file server does not require installation of Tivoli Storage Manager software. Instead, the TSM server uses NDMP to connect to the NAS file server to initiate, control, and monitor a file system backup or restore operation, as shown in Figure 7-12. Tivoli Data Protection for NDMP utilizes the Network Data Management Protocol (NDMP) to communicate with and provide backup and recovery services for NetApps NAS file servers. NDMP is an industry-standard protocol that allows a network storage-management application to control the backup and recovery of an NDMP-compliant file server without installing third-party software on that server. The implementation of the NDMP server protocol enables the NAS file servers to be backup-ready and enables higher-performance backup to tape devices without moving the data over the LAN.

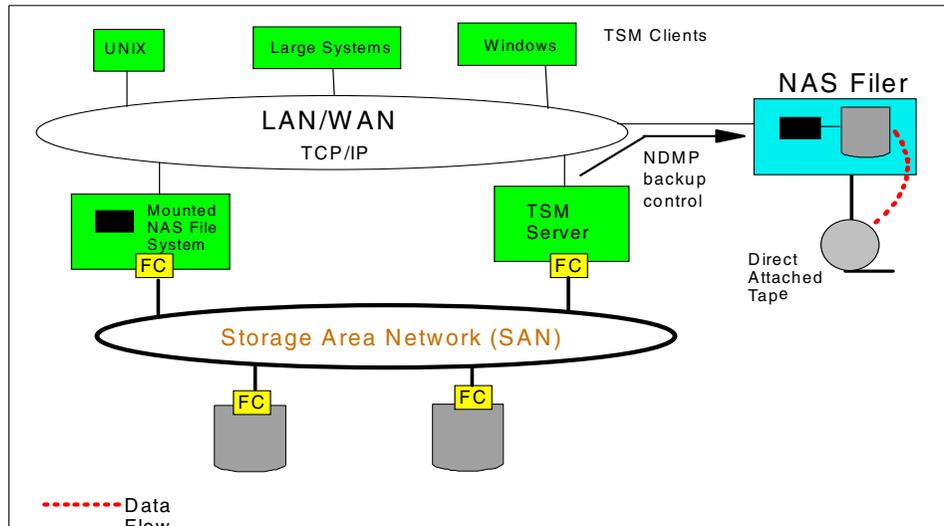


Figure 7-12 TSM and NDMP backup

7.5.7 Image backup

An image backup is a block by block copy of data from the TSM client to the TSM backup server. One important function of an image restore is to accelerate recovery in a Disaster Recovery scenario. Image backup is available at the time of writing on AIX, HP, Sun, Linux, and Windows client platforms. Specific image backup requirements for each platform should be reviewed in their associated administration documents found at:

http://www.tivoli.com/support/public/Prodman/public_manuals/td/TD_PROD_LIST.html

With image backup, the TSM server does not track individual files in the file system image. File system images are tracked as individual objects and the management class policy will be applied to the file system image as a whole. An image backup provides the following benefits.

- ▶ Image backup can provide a quicker backup and restore than a file-by-file backup as there is no overhead involved in creating individual files.
- ▶ Conserves resources on the server during backups since only one entry is required for the image.
- ▶ Provides a point-in-time picture of your file system, which may be useful if your enterprise needs to recall that information.
- ▶ Restores a corrupt file system or raw logical volume. Data is restored to the same state it was when the last logical volume backup was performed.

Figure 7-13 illustrates the process for image backup with TSM.

On the Windows 2000 client platform a Logical Volume Storage Agent (LVSA) has been introduced which is capable of taking a snapshot of the volume while it is online. Optionally only occupied blocks can be copied. If the snapshot option is used (rather than static) then any blocks which change during the backup process are first kept unaltered in an Original Block File. In this way the client is able to send a consistent image of the volume as it was at the start of the snapshot process to the TSM server.

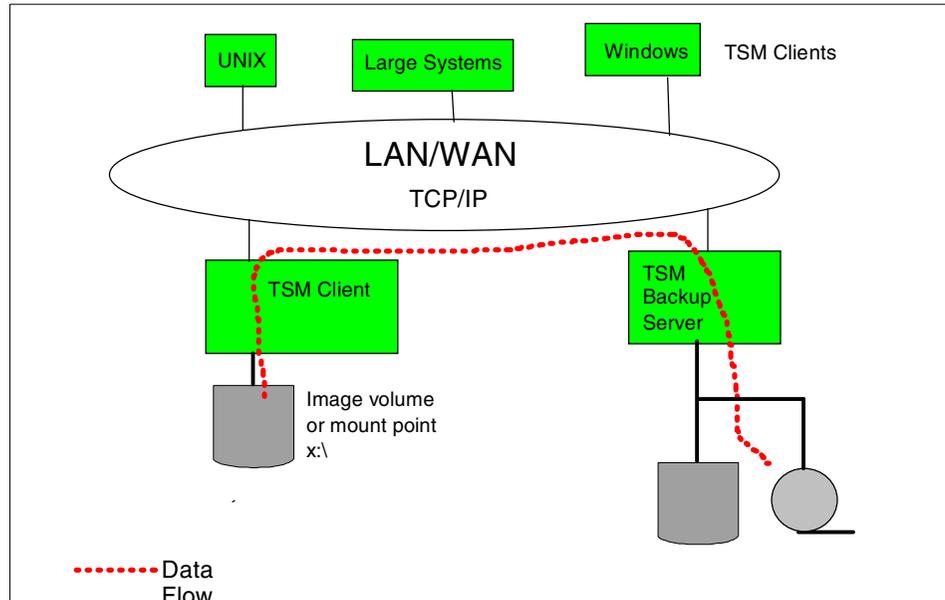


Figure 7-13 TSM image backup

This section has provided an overview of many of the common backup methods supported by TSM. These methods can be integrated with DR strategies being considered in your environment.

7.6 TSM Disaster Recovery Manager (DRM)

Disaster Recovery Manager (DRM) is a component of TSM that is now included with IBM Tivoli Storage Manager Version 5 Extended Edition. Tivoli Disaster Recovery Manager provides Disaster Recovery for the TSM server and assists in Disaster Recovery for clients.

Disaster Recovery Manager offers various options to configure, control and automatically generate a Disaster Recovery Plan file containing the information, scripts, and procedures needed to automate restoration and help ensure quick recovery of data after a disaster. It also manages and tracks the media on which the data is stored, whether onsite, in-transit, or in a vault, so that required data can be easily located if disaster strikes. The scripts can help document the basic Information Technology recovery strategy, the steps to rebuild core systems, as well as the critical machines that must be recovered.

One of the key features of Tivoli Storage Manager and Tivoli Disaster Recovery Manager is the ability to track media in all states that it could possibly be, such as onsite, in transit or in a vault. Because of the criticality of data in the production environment, controls are needed to make sure that all previously backed up data can be found and restored in a reasonable amount of time.

Tivoli Disaster Recovery Manager functions helps maintain business continuance by:

- ▶ Establishing and helping to automate a thorough *server* Disaster Recovery Plan — clients can then subsequently restore their data from the server if required.
- ▶ Ensuring that customer-provided information is available in the same plan.
- ▶ Automating vital recovery steps to bring the TSM server and backup environment back to normal operation.
- ▶ Managing and indentifying offsite media needed for recovery.
- ▶ Tracking and reporting systems destroyed, in the event of a disaster.
- ▶ Storing client configuration information and assigning client recovery priorities.

With DRM you can recover at an alternate site, on replacement computer hardware, recover using different hardware configuration at the recovery site, and with people who are not familiar with the applications. You can also use the DRP for audits to certify the recoverability of the server. The DRP can be easily recreated daily so that it is up to date. The main functions of DRM are illustrated in Figure 7-14.

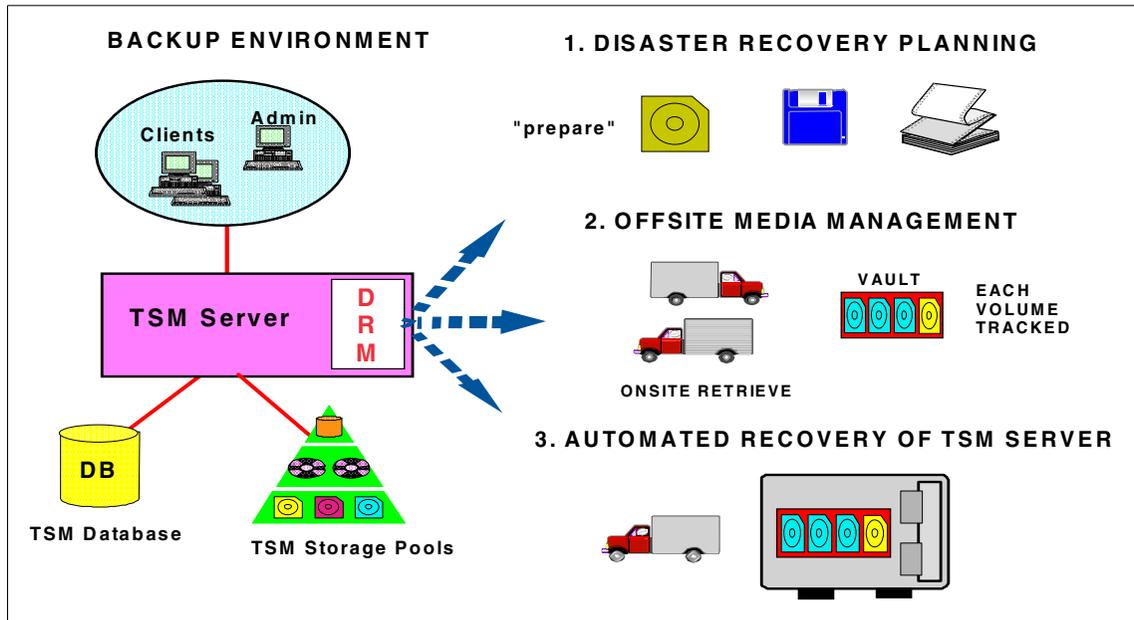


Figure 7-14 Tivoli DRM functions

During a real disaster, errors commonly encountered include:

- ▶ The DRP was not tested
- ▶ A skilled technical team performed the testing, who "filled-in" missing steps
- ▶ The recovery plan is out of date
- ▶ Disk volume definitions for the recovery site are not known
- ▶ Location of recovery tapes is not known
- ▶ It is not known which tapes are to be applied first.

DRM will help answer questions, such as:

- ▶ Where is the current server configuration information located?
- ▶ What are the current server database volumes?
- ▶ What is my recovery sequence?
- ▶ Is my recovery plan current, is this guaranteed?
- ▶ What was the client and server machine configuration like?
- ▶ Who should be contacted in a disaster?
- ▶ Where is the recovery media located?
- ▶ Can I restore my environment to any point in time?

During recovery from a disaster, DRM automates the following procedures to restore the TSM servers:

- ▶ Restores TSM server's key option files
- ▶ Copies files from alternate locations to production locations
- ▶ Initializes TSM database and log volumes
- ▶ Matches sizes and locations of TSM database and log volumes
- ▶ Launches current DB restore automatically
- ▶ Tracks media needed and availability
- ▶ Registers TSM server features installed
- ▶ Returns server state to a valid license configuration
- ▶ Updates TSM volume catalog information
- ▶ Marks volume information for recovery, that is, is it destroyed or not?
- ▶ Rebuilds TSM hierarchical storage configuration
- ▶ Re-creates customer backup environment

DRM uses the **PREPARE** command to generate a plan file that will contain critical information needed for recovery. Information in the plan is arranged in stanzas, these can be considered to be somewhat like headers. For example the stanza `PLANFILE.DESCRPTION` shown in Example 7-1 provides summary information about the plan file as a whole.

Example 7-1 DMR Plan File Example Stanza

```
begin PLANFILE.DESCRPTION

Recovery Plan for Server RADON_SERVER1
Created by DRM PREPARE on 07/26/2002 18:30:58
DRM PLANPREFIX C:\DRM\PLAN\RADON
Storage Management Server for Windows - Version 5, Release 1, Level 1.0

end PLANFILE.DESCRPTION
```

A detailed description, recovery scenario, and recovery plan built with DRM is given in Chapter 8, “IBM Tivoli Storage Manager and DRM” on page 163. Also, recommendations and examples of using DRM to store client machine information in the DRM plan file for use during a client disaster recovery are given in Chapter 12, “Windows 2000 client bare metal recovery” on page 267 and Chapter 13, “Linux client bare metal recovery” on page 289.

In summary, DRM will systematically re-build the storage management server environment and ensure current application data for the entire enterprise is available for recovery. This is all possible from a single scripted command, automatically.

7.6.1 DRM and TSM clients

DRM is designed to automate the recovery of your TSM server(s), rather than automate client recovery. The recovery of the TSM server(s) is most critical

because it is from the server that clients can continue to backup their data and recover if required. For the clients, DRM allows the machine information needed to help recover the TSM clients to be stored in the TSM server database. The type of information stored includes:

- ▶ TSM client machine location, machine characteristics, and recovery instructions.
- ▶ Business priorities associated with the TSM client machines.
- ▶ Description, location, and volume/diskette/CD labels of TSM client boot media.

With this information stored in the DRM plan file you can then use appropriate operating system CDs or tape images to perform a bare metal recovery (BMR) of your client system. Finally, the TSM client on that machine can be reinstalled which allows for restoration of client data from the TSM server. Strategies for bare metal recovery and recovery of data on various client platforms are discussed in detail in:

- ▶ Chapter 10, “Solaris client bare metal recovery” on page 229
- ▶ Chapter 11, “AIX client bare metal recovery” on page 239
- ▶ Chapter 12, “Windows 2000 client bare metal recovery” on page 267
- ▶ Chapter 13, “Linux client bare metal recovery” on page 289

7.7 TSM server-to-server communications

With a solution that includes multiple TSM servers, you can use server-to-server communications and virtual volumes to enhance management and improve disaster recoverability. TSM server-to-server communications provide the capability to:

- ▶ Configure and manage multiple servers with enterprise administration.
- ▶ Distribute a consistent configuration for TSM servers through a configuration manager to managed servers. By having consistent configurations, you can simplify the management of a large number of servers and clients.
- ▶ perform tasks simultaneously on multiple servers by using command routing, enterprise logon and enterprise console.
- ▶ Send server and client events to another server for logging.
- ▶ Monitor events of many servers and clients from a single server.
- ▶ Store data on another server using virtual volumes.

This section provides an overview of server-to-server communication and virtual volumes.

7.7.1 Server-to-server communication

The server-to-server communication function in TSM provides the ability to enable enterprise configuration, enterprise event logging and monitoring, command routing (Figure 7-15).

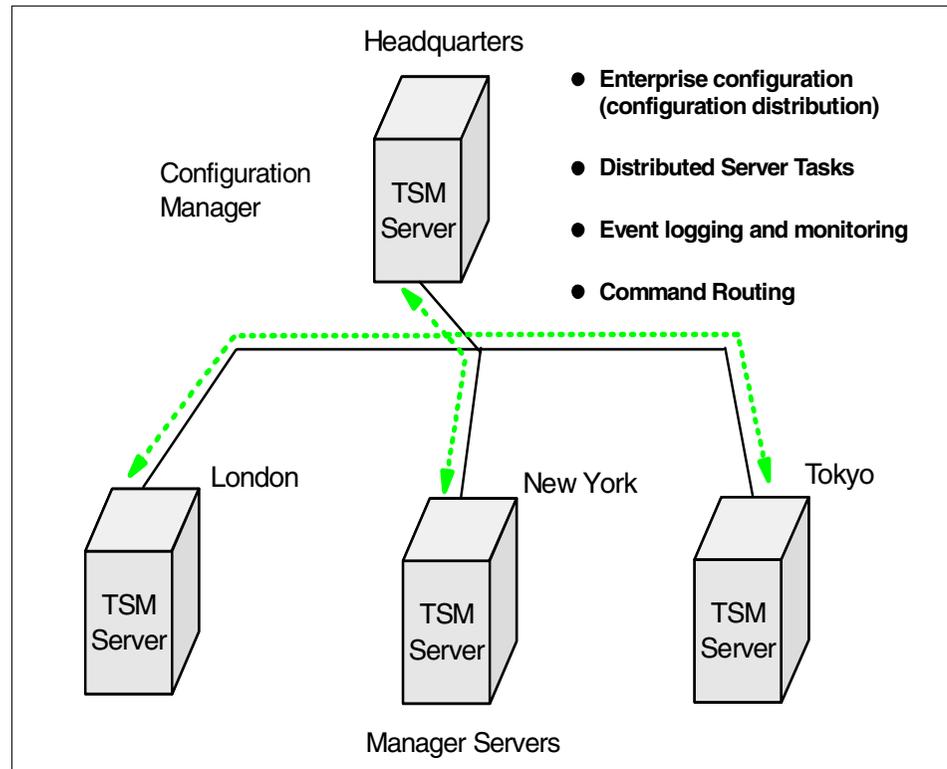


Figure 7-15 TSM server-to-server communications

The enterprise configuration functions of the Tivoli Storage Manager make it easier to consistently set up and manage a network of TSM servers. You set up configurations on one server and distribute the configurations to the other servers. You can make changes to configurations and have the changes automatically distributed. To use enterprise configuration, you first select the TSM server that is to act as the *configuration manager*. You may want to dedicate a new server for this purpose. At the configuration manager, you define the details of the server configurations that you want to distribute. For example:

- ▶ You set up backup and archive policies and client option sets.
- ▶ You designate one or more administrators to have access to the servers, and control their authority levels.

- ▶ You define the servers that you want the configuration manager to manage or communicate with, and you set up communications among the servers.

On each server that is to receive the configuration information, you identify the server as a *managed server* by defining a *subscription* to one or more profiles owned by the configuration manager.

When you connect to the configuration manager via a Web browser, you are presented with the enterprise console. From the enterprise console you can perform tasks on the configuration manager and on one or more of the managed servers. You can also connect to another server to perform tasks directly on that server. As long as you are registered with the same administrator ID and password, you can do this work on many servers without having to log on each time.

From the command line of the administrative Web interface or from the command-line administrative client, you can also route commands to other servers. The other servers must be defined to the server to which you are connected. You must also be registered on the other servers as an administrator with the administrative authority that is required for the command. Command routing enables an administrator to send commands for processing to one or more servers at the same time. The output is collected and displayed at the server that issued the routed commands. A system administrator can configure and monitor many different servers from a central server by using command routing. To make routing commands easier, you can define a server group that has servers as members. Commands that you route to a server group are sent to all servers in the group.

One or more servers can send server events and events from their own clients to another server for logging. The sending server receives the enabled events and routes them to a designated event server. This is done by a receiver that Tivoli Storage Manager provides. At the event server, an administrator can enable one or more receivers for the events being routed from other servers.

7.7.2 Server-to-server virtual volumes

TSM server-to-server virtual volume support lets a server (a *source* server) use another TSM server (a *target* server) as the destination device class for certain TSM entities. The data is stored as *virtual volumes*, which are actually archive files on a target server. The target server has no knowledge of the contents or type of objects which are stored in the virtual volumes. The only detailed information about the virtual volume contents is contained in the source server's database. Virtual volumes can be any of the following:

- ▶ Database backups
- ▶ Storage pool backups (copy storage pools)
- ▶ Data that is backed, archived or space managed from client nodes (primary storage pools)
- ▶ Any data that can be moved by the **EXPORT** and **IMPORT** commands
- ▶ DRM plan files (output of the DRM **PREPARE** command)

The source server is a client of the target server, and the data for the source server is managed only by the source server. In other words, the source server controls the expiration and deletion of the files that comprise the virtual volumes on the target server. At the target server, the virtual volumes from the source server are seen as archive data. The relationship between the source and target TSM servers is illustrated in Figure 7-16.

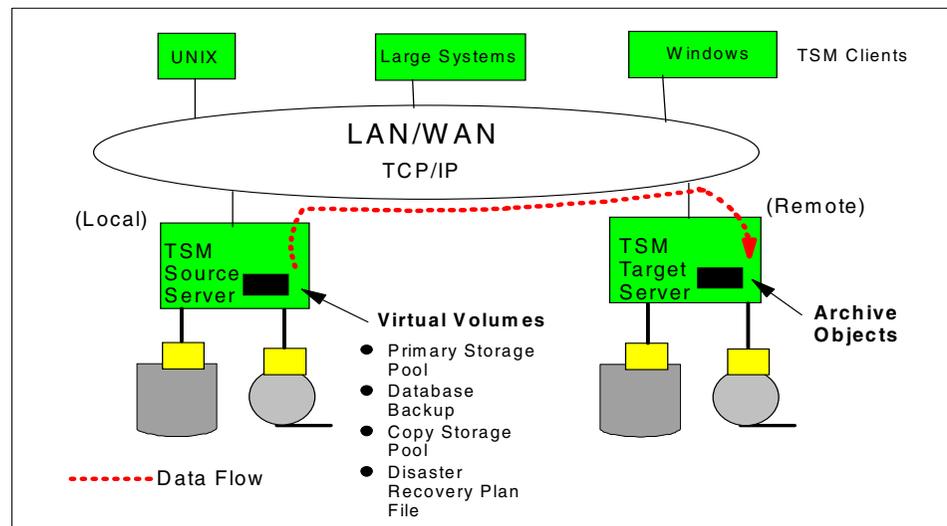


Figure 7-16 Server-to-server virtual volumes

The source server is registered as a client node (of TYPE=SERVER) at the target server and is assigned to a policy domain. The archive copy group of the default management class of that domain specifies the storage pool for the data from the source server. TCP/IP is the only communications method that can be used for server-to-server virtual volumes.

To use virtual volumes, the source server needs to define one or more device classes of TYPE=SERVER, using the **DEFINE DEVCLASS** command. The device class definition indicates on which remote or target server the volumes will be

created. Having defined the device class, it can be used for primary or copy storage pools, database backups and other virtual volume functions.

All data destined for virtual volumes is sent to the target server, using virtual volumes rather than direct attached storage devices. For example, if a client is backing up data which is bound to a backup copy group which is using a virtual volume primary storage pool, this data will be sent to the target server. If a client needs to restore the data, the source server gets the data back from the target server. A TSM client can always perform the same granularity of restore, retrieve or recall operation, whether the data is stored on a local TSM server, or on a target server using server-to-server communication. That is, remote storage pools (using server-to-server communication) are transparent to the client. The only requirement is that the TCP/IP communication link between the source and target server must be working correctly.

Note: Source server objects such as database and storage pool backups are stored on the target server as *archived* data. Therefore, the target server cannot directly restore these objects in the event of a disaster at the source server site. In the event of a disaster at the source server site, the source server should be re-installed (likely at an alternate location) and then objects originally stored on the target server can be restored over the network, using the same server-to-server communication.

Using virtual volumes can benefit you in the following ways:

- ▶ The source server can use the target server as an electronic vault. This removes the need to physically ship media offsite, or retrieve it physically in a disaster.
- ▶ Smaller TSM source servers can use the storage pools and tape devices of larger TSM servers.
- ▶ For incremental database backups, it can decrease wasted space on volumes and under use of high-end tape drives.

The TSM Disaster Recovery Manager (now included with TSM Extended Edition Version 5) provides full support of server-to-server virtual volumes for database and storage pool backup. This function is required on a source server but not on the target server.

7.7.3 Using server-to-server virtual volumes for Disaster Recovery

Server-to-server virtual volumes can be used as part of strategies for site disaster protection and DR.

For site disaster protection, server-to-server virtual volumes (with Tivoli Disaster Recovery Manager) can be used to back up storage pools from one server to another and back up the TSM database of one server to another. This form of virtual vaulting can occur on a peer-to-peer basis, between any number of servers, and at different data centers or different sites. For example, two servers can back up each other in an equal level relationship.

For disaster recovery, server-to-server virtual volumes can be used to store the Disaster Recovery Plan file remotely. In this strategy, the source server creates the Disaster Recovery Manager plan files, then stores the files on a remote target server. You can display information about recovery plan files from the server that created the files (the source server) or from the server on which the files are stored (the target server). To display plan file information on the target server you can issue the **QUERY RPFIL** command, specifying the source server node name that prepared the plan. You can easily display a list of all recovery plan files that have been saved on a target server.

7.7.4 Considerations for server-to-server virtual volumes

When using server-to-server virtual volumes consideration must be given to network bandwidth, the location of TSM database and storage pools, and data flow duplication. As mentioned earlier, TCP/IP is the only supported communications protocol. The following should be considered when implementing server-to-server virtual volumes.

- ▶ If you use virtual volumes for database backups, you might have the following situation: SERVER_A backs up its database to SERVER_B in the same location, and SERVER_B backs up its database to SERVER_A. If this is the only way databases are backed up, and a disaster strikes at that location, you may have no backups available from which to restore your databases.
- ▶ Moving large amounts of data between the servers may slow down your communications significantly, depending on the network bandwidth and availability. In general a TSM server-to-server communications operation will achieve slower throughput than a typical client backup operation. You should carefully test to make sure sufficient performance is provided.
- ▶ Storage space limitations on the target server affect the amount of data that you can store on that server. Make sure that you include virtual volume requirements in your planning for each server.
- ▶ Duplicate data flow over the network can occur when one storage pool is backed up to another storage pool and both primary and copy storage pools are using virtual volumes on the same target server. In this situation, data will be moved from a primary storage pool volume (on the target server), back to the source server, then to a copy storage pool volume on the target server.

This is clearly more time consuming than a simple copy storage pool operation without using server-to-server communication.

7.8 TSM and high availability clustering

What exactly is a cluster? In simple terms, a cluster is a group of computers that, together, provide a set of resources to a client. A cluster consists of a minimum of two and in theory, up to any number, of systems. The key point of clustering is that the client has no knowledge of the underlying physical hardware of the cluster. This means that the client is isolated and protected from changes to the physical hardware, which brings a number of benefits.

Perhaps the most important of these benefits is high availability. Resources on clustered servers act as highly available versions of unclustered resources. If a node (an individual computer) in the cluster is unavailable or too busy to respond to a request for a resource, the request is transparently passed to another node capable of processing it. Clients are therefore unaware of the exact locations of the resources they are using. For example, a client can request the use of an application without being concerned about either where the application resides or which physical server is processing the request. The user simply gains access to the application in a timely and reliable manner.

Another benefit is scalability. If you need to add users or applications to your system and want performance to be maintained at existing levels, additional systems can be incorporated into the cluster. A typical example would be a Web site that shows rapid growth in the number of demands for Web pages from browser clients. Running the site on a cluster would allow the growth in demand to be easily accommodated by adding servers to the cluster as needed.

Whether you will configure your system to include clusters depends on your business needs. A cluster can provide system level high availability to ensure a TSM server or client can continue normal backup and restore processes without significant disruption for users and administrators. In addition to assuring the right type of hardware and the applicable software, varying failover patterns between cluster nodes exist and play different roles, for example, hot-standby versus concurrent cluster operation. This section overviews TSM server and client clustering support and configuration using IBM High Availability Cluster Multi-Processing (HACMP) and Microsoft Cluster Server (MSCS).

7.8.1 TSM and High Availability Cluster Multi-Processing (HACMP)

A TSM server can use IBM's High Availability Cluster Multiprocessing (HACMP) software for high availability. HACMP provides the leading AIX-based clustering solution, which allows automatic system recovery on system failure detection. By

using HACMP together with TSM, you can ensure server availability. HACMP offers local or campus disaster survivability with real-time automated failover and reintegration within distance limitations. In an HACMP environment, TCP/IP is the communications method used to support the checking of status and availability of be the production and failover server, also commonly referred to as the *heartbeat* connection.

HACMP detects system failures and manages failover to a recovery processor with a minimal loss of end-user time. You can set up a TSM server on a system in an HACMP cluster so that, if the system fails, the TSM server will be brought back up on another system in the cluster. In both failover and fallback, it appears that the TSM server has crashed or halted and was then restarted. Any transactions that were in progress at the time of the failover or fallback are rolled back, and all completed transactions are still complete. Tivoli Storage Manager clients see this as a communications failure and try to re-establish their connections as shown in Figure 7-17.

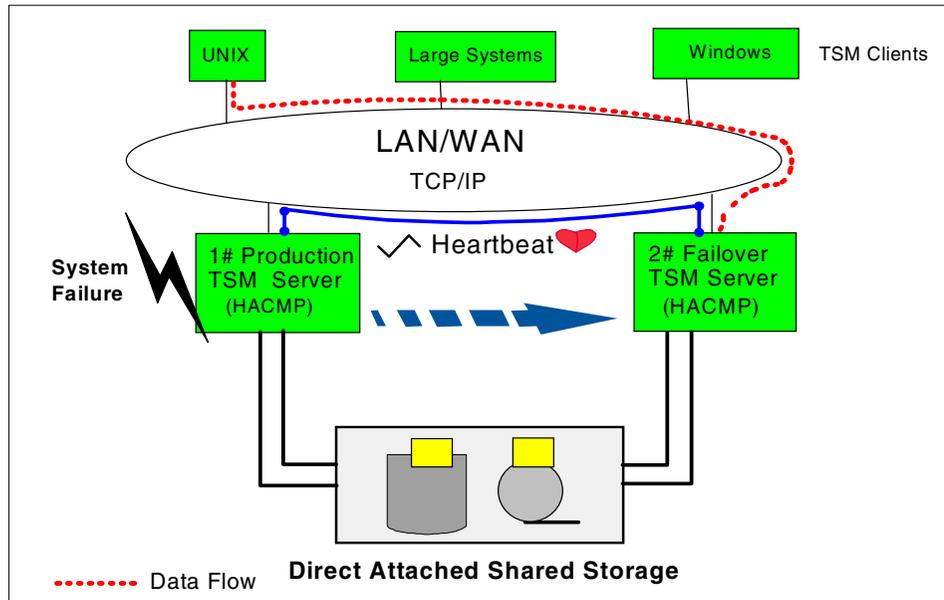


Figure 7-17 HACMP and TSM server high availability configuration

HACMP support for the TSM server on AIX has been officially supported since V4.2. With V5.1 of TSM, several HACMP scripts are provided with the TSM server fileset. These scripts can then be customized to suit the local environment. When failover occurs, HACMP calls the TSM *startserver* script on the standby node. The script verifies the devices, breaking the SCSI reserves, and starts the server. On fallback, the *stopserver* script runs on the standby node, which causes the TSM server to halt. Then the *startserver* script runs on the

production node. HACMP handles taking over the TCP/IP address and mounting the shared file systems on the standby node or production node, as appropriate. By default, the *startserver* script will not start the Tivoli Storage Manager server unless all the devices in the VerifyDevice statements can be made available. However, you can modify the *startserver* script to start the server even if no devices can be made available.

Both failover and fallback act as if a Tivoli Storage Manager server has crashed or halted and was then restarted. Any transactions that were in-flight at the time are rolled back, and all completed transactions are still complete. Tivoli Storage Manager clients see this as a communications failure and try to re-establish connection based on their COMMRESTARTDURATION and COMMRESTARTINTERVAL settings. The backup-archive client can usually restart from the last committed transaction. The clients and agents will behave as they normally do if the server was halted and restarted while they were connected. The only difference is that the server is physically restarted on different hardware.

For a detailed discussion on supported environments, prerequisites, install, setup, and testing of a HACMP and TSM server failover environment see *Tivoli Storage Manager Version 5.1 Technical Guide*, SG24-6554, and *Tivoli Storage Manager for AIX Quick Start Version 5.1*, GC32-0770.

7.8.2 TSM backup-archive and HSM client support with HACMP

As of TSM Version 5.1, the backup-archive client itself (including the administrator, backup/archive, HSM and API pieces) is supported for use in an HACMP cluster environment. This configuration allows scheduled Tivoli Storage Manager client operations to continue processing in the event of a system failure on a redundant clustered failover server. See Figure 7-18 for an illustration of how this works.

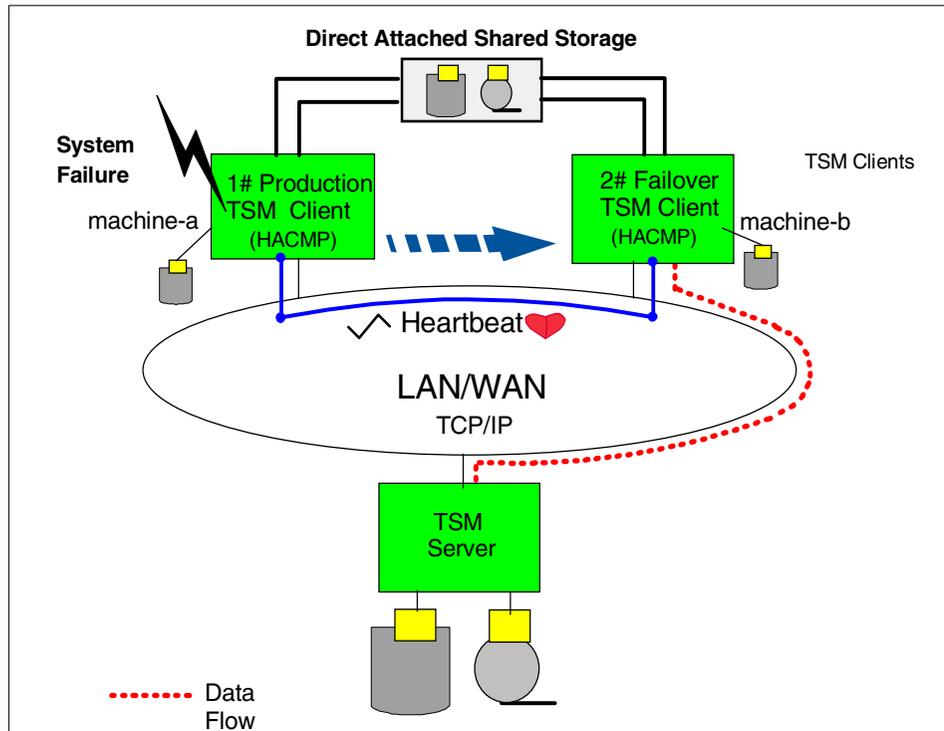


Figure 7-18 HACMP and TSM client high availability configuration

The CLUSTERNODE option in the AIX client `dsm.sys` file determines if you want the TSM client to back up cluster resources and participate in cluster failover for high availability.

If a scheduled incremental backup of a clustered volume is running on machine-a and a system failure causes a failover to machine-b, machine-b then reconnects to the server. If the reconnection occurs within the start window for that event, the scheduled command is restarted. This scheduled incremental backup will reexamine files sent to the server before the failover. The backup will then “catch up” to where it terminated before the failover situation.

If a failover occurs during a user initiated (that is, non-scheduled) client session, the TSM client starts on the node that is handling the takeover. This allows it to process scheduled events and provide Web client access. You can install the TSM client locally on each node of an HACMP environment. You can also install and configure the TSM Scheduler Service for each cluster node to manage all local disks and each cluster group containing physical disk resources.

HACMP support for Hierarchical Storage Management (HSM) clients on AIX provides support for HACMP failover on AIX so that HSM managed filesystems can continue to operate in the case of an HACMP node failover and fallback.

For a detailed discussion on supported environments, prerequisites, install, setup, and testing of a HACMP and TSM client failover environment see *Tivoli Storage Manager Version 5.1 Technical Guide*, SG24-6554, and *Tivoli Storage Manager for UNIX Backup-Archive Clients Installation and user's Guide Version 5.1*, GC32-0789.

7.8.3 TSM and Microsoft Cluster Server (MSCS)

Tivoli Storage Manager is a cluster-aware application and can be configured in a MSCS high availability environment. The administrator uses the MSCS Cluster Administrator interface and Tivoli Storage Manager to designate cluster arrangements and define the failover pattern. The systems are connected to the same disk subsystem and provide a high-availability solution that minimizes or eliminates many potential sources of downtime. Microsoft Cluster Server (MSCS) is software that helps configure, monitor, and control applications and hardware components that are deployed on a Windows cluster. When you use cluster configurations, you enhance the availability of your servers. Clustering allows you to join two Windows servers, or nodes, using a shared disk subsystem. This provides the nodes with the ability to share data, which provides high server availability.

For example, in the MSCS failover environment shown in Figure 7-19, a clustered TSM server called TSMSERVER1 runs on node A and a clustered TSM server called TSMSERVER2 runs on node B. Clients connect to the TSM server TSMSERVER1 and the TSM server TSMSERVER2 without knowing which node currently hosts their server. The MSCS concept of a virtual server ensures that the server's location is transparent to client applications. To the client, it appears that the TSM server is running on a virtual server called TSMSERVER1. When one of the software or hardware resources fails, failover occurs. Resources (for example, applications, disks, or an IP address) migrate from the failed node to the remaining node. The remaining node takes over the TSM server resource group, restarts the TSM service, and provides access to administrators and clients. If node A fails, node B assumes the role of running TSMSERVER1. To a client, it is exactly as if node A were turned off and immediately turned back on again.

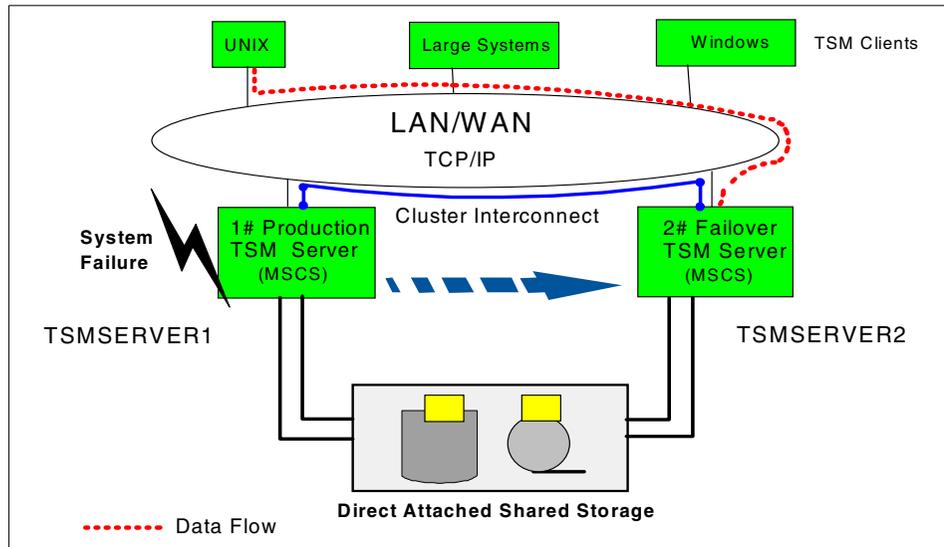


Figure 7-19 MSCS and TSM server high availability configuration

Clients experience the loss of all connections to TMSERVER1 and all active transactions are rolled back to the client. Clients must reconnect to TMSERVER1 after this occurs, which is normally handled as an automatic attempt to reconnect by the TSM client. The location of TMSERVER1 is transparent to the client. A node can host physical or logical units, referred to as resources. Administrators organize these cluster resources into functional units called groups and assign these groups to individual nodes. If a node fails, the server cluster transfers the groups that were being hosted by the node to other nodes in the cluster. This transfer process is called failover. The reverse process, failback, occurs when the failed node becomes active again and the groups that were failed over to the other nodes are transferred back to the original node.

Two failover configurations are supported with MSCS and TSM: *active/passive* and *active/active*. In the *active/passive* configuration you create one instance of a TSM server that can run on either node. One system runs actively as the production TSM server, while the other system sits passively as an online (hot) backup. In the *active/active* configuration the cluster runs two independent instances of a TSM server, one on each server. In the event of a system failure, the server on the failed instance transfers to the surviving instance, so that it is running both instances. Even if both instances are running on the same physical server, users believe they are accessing a separate server.

Clusters consist of many components such as nodes, cluster objects, Microsoft Cluster Server (MSCS) virtual servers, and even the hardware and software. If any one of these components is missing, the cluster cannot work. MSCS

requires each TSM server instance to have a private set of disk resources. Although nodes can share disk resources, only one node can actively control a disk at a time. TCP/IP is used as the communications method in a MSCS environment with TSM.

MSCS does not support the failover of tape devices. However, TSM can handle this type of a failover pattern with the correct set up. TSM uses a shared SCSI bus for the tape devices. Each node (two only) involved in the tape failover must contain an additional SCSI adapter card. The tape devices (library and drives) are connected to the shared bus. When failover occurs, the TSM server issues a SCSI bus reset during initialization. In a failover situation, the bus reset is expected to clear any SCSI bus reserves held on the tape devices. This allows the TSM server to acquire the devices after the failover.

For a detailed discussion on supported environments, prerequisites, install, setup, and testing of a MSCS and TSM server failover environment, see *Tivoli Storage Manager for Windows Administrator's Guide Version 5.1*, GC32-0782, and *Tivoli Storage Manager Version 3.7.3 and 4.1: Technical Guide*, SG24-6116.

7.8.4 TSM backup-archive client support with MSCS

The TSM client is supported within a MSCS cluster environment. This configuration allows scheduled Tivoli Storage Manager client operations to continue processing in the event of a system failure on a redundant clustered failover server as shown in Figure 7-20.

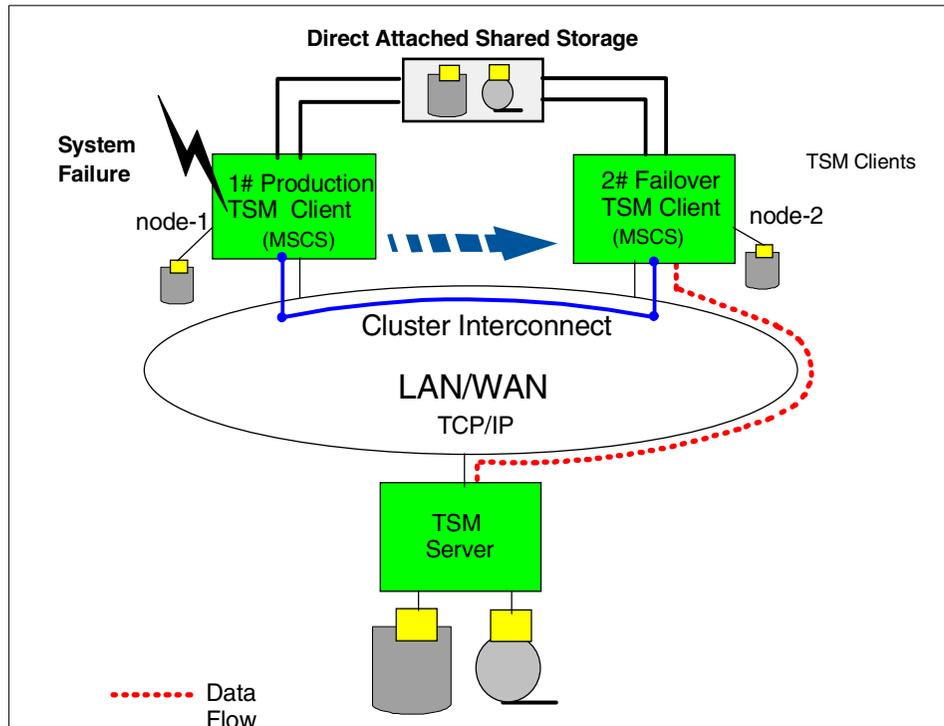


Figure 7-20 MSCS and TSM client high availability configuration

In this example, the cluster contains two nodes: node-1 and node-2, and two cluster groups containing physical disk resources. In this case, an instance of the TSM Backup-Archive Scheduler Service should be installed for each node node-1, node-2, and physical disk resources. This ensures that proper resources are available to the Backup-Archive client when disks move (or fail) between cluster nodes. The CLUSTERNODE option in the client option file ensures that TSM manages backup data logically, regardless of which cluster node backs up a cluster disk resource.

For a detailed discussion on supported environments, prerequisites, install, setup, and testing of a MSCS and TSM client failover environment see *Tivoli Storage Manager for Windows Backup-Archive Clients Installation and User's Guide Version 5.1*, GC32-0788.

7.9 TSM and remote disk replication

Using TSM with remote disk replication solutions can provide full disaster tolerance and among the highest of levels of RTO and RPO, that is, Tier 6.

Hardware disk replication technologies, such as IBM's Peer-to-Peer Remote Copy (PPRC) or EMC's SRDF can be used to provide real-time (synchronous or asynchronous) mirroring of the Tivoli Storage Manager database, log, or storage pools to a remote site, as illustrated in Figure 7-21. In the event of a disaster, the target volume at Datacenter #2 can be suspended and the volume pair terminated which will return the target volume to the simplex state where the secondary volume is now accessible to the host on the remote site. This online copy of the TSM database at the remote site can be used to resume the TSM environment. Later the volumes could be re-synched to re-establish a remote mirror pair.

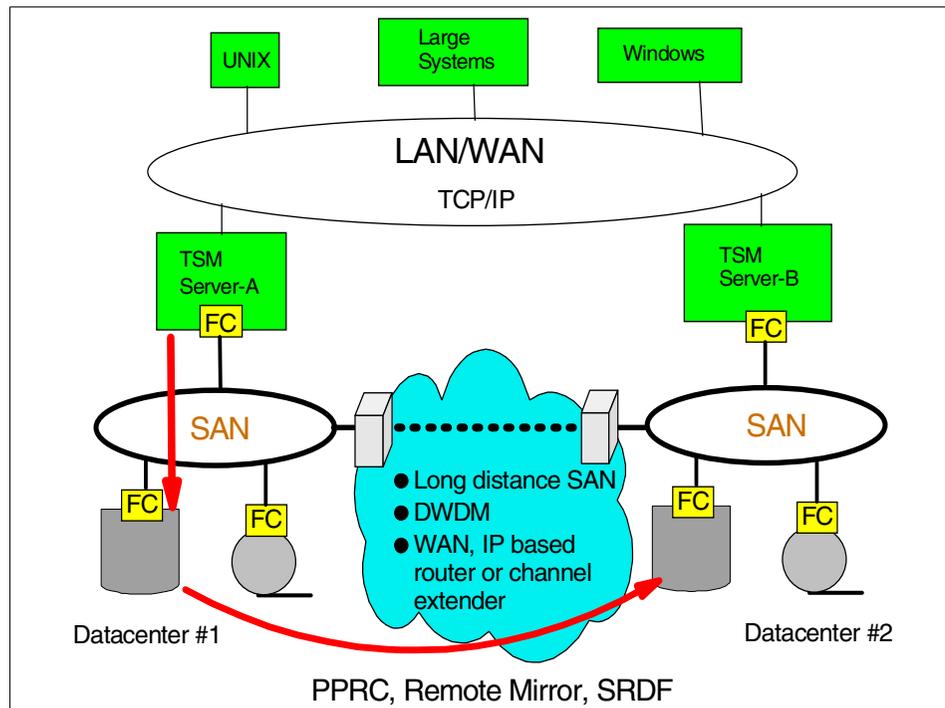


Figure 7-21 TSM and remote disk replication

If remote disk replication is used for mirroring of the TSM database and storage pools, the TSM server could be recovered very quickly without any loss of client data. A peer-to-peer configuration could be used to balance the load of TSM services in the enterprise and provide data protection and rapid recovery for a failure at either site. Various configurations with remote disk replication can exist. For example, if only the TSM database and logs are mirrored remotely, recovery of client data can begin once TSM copy pools are restored from tape. Electronic tape vaulting, (see Figure 7-22), can be used along with remote disk replication to improve recovery time and recovery point objectives. Optionally, data storage

pools can also be replicated in real-time to further improve recovery time and recovery point objectives. Failover to a DR site while using remote disk replication and electronic vaulting enables services to be restored in the shortest time possible. The decision on what to mirror (database or database and storage pools) depends on factors such as cost of extra disk and communication hardware, availability of remote links and ability to set up additional procedures to support this.

Attention: TSM database updates may be broken into multiple I/Os at the operating system level. Therefore, it is necessary to replicate *all* Tivoli Storage Manager-managed mirrors of the Tivoli Storage Manager database to the remote site. Replicating all database mirror copies will protect against a broken chain of I/Os.

7.10 TSM and tape vaulting

Traditionally, disaster recovery plans include daily offsite tape backups that are picked up from the local site and transported via a courier to a secure facility, often a tape vaulting service provider. Vaulting of tapes at offsite locations can provide a secure means to protect data in the event of a disaster at the primary site. To recover from a disaster, you must know the location of offsite recovery media. DRM helps you to determine which volumes to move offsite and back onsite and tracks the location of the volumes. With tape vaulting you can back up primary storage pools to a copy storage pool and then send the copy storage pool volumes offsite. You can track these copy storage pool volumes by changing their access mode to offsite, and updating the volume history to identify their location. If an offsite volume becomes expired, the server does not immediately return the volume to the scratch pool. The delay prevents the empty volumes from being deleted from the database, making it easier to determine which volumes should be returned to the onsite location. DRM handles all of this automatically.

7.10.1 Electronic tape vaulting

Using TSM with electronic tape vaulting provides additional data protection capabilities, with backups made to remote tape drives over communication links. Electronic vaulting can enable shorter recovery times and reduced data loss should the server be damaged. An electronic tape vaulting solution combined with TSM is fundamental to achieving Tier 3 and above RPO and RTOs, that is, less than 24 hours, as outlined in Chapter 2, “The tiers of Disaster Recovery and TSM” on page 21. With electronic tape vaulting the TSM server will have an alternate location to store primary and copy storage pools as though they are directly attached. The TSM server can first write a copy of disk storage pool data

to tape pools at the remote site (Datacenter #2), then the data can be migrated to the tape storage pools at the primary site (Datacenter #1). See Figure 7-22.

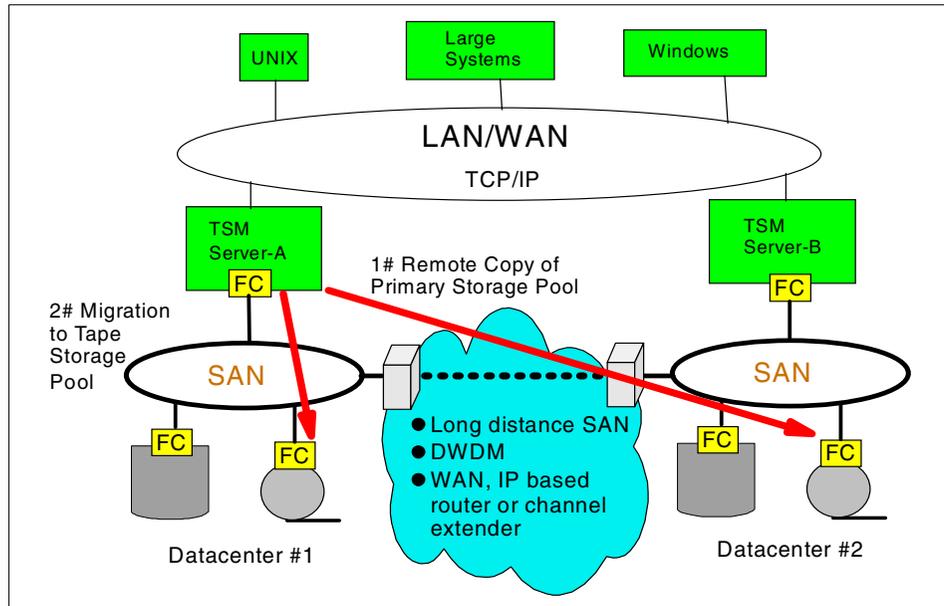


Figure 7-22 TSM and electronic tape vaulting

Depending on your configuration (and whether or not remote disk replication is being used in conjunction with electronic tape vaulting) you may choose to backup the TSM database and configuration files by this method. This ensures a copy of the data is stored at both sites and that the TSM server can rapidly recover at the remote site. If remote disk replication is used for mirroring of the TSM database and storage pools, the TSM server could be recovered very quickly without any loss of client data. A peer-to-peer configuration could be used to balance the load of TSM services in the enterprise and provide data protection and rapid recovery for a failure at either site. Some advantages for using electronic tape vaulting with TSM include:

- ▶ Critical data can be frequently and rapidly vaulted to remote offsite locations.
- ▶ In the event of a disaster, up-to-date data can be restored at a hot site, therefore improving recovery time and recovery point objectives.
- ▶ Eliminates physical tape handling which could result in: damaged tapes, lost tapes, tapes delayed in transit, or data that is sabotaged. Increased reliability.
- ▶ Eliminates costs associated with couriers and offsite vaulting vendors.
- ▶ Government offsite vaulting regulations are satisfied.

- ▶ Lowers cost of downtime and storage management.
- ▶ Increases company control and data security.
- ▶ Peer solutions eliminate or reduce costs associated with hot-site service providers.

7.11 Remote disk mirroring and tape vaulting solutions

A variety of technologies exist for remote disk mirroring and electronic tape vaulting, these include. These include:

- ▶ Long distance SANs
- ▶ Dense Wavelength Division Multiplexing (DWDM)
- ▶ Fibre extenders
- ▶ WAN based channel extension using telco and IP protocols
- ▶ Newly emerging NAS and iSCSI gateways

Table 7-2 summarizes some of these. The use of these various technologies also may depend on a particular vendor's replication or vaulting solution. For example, the IBMs TotalStorage Enterprise Storage Server uses PPRC to achieve data replication. PPRC is supported via ESCON links which can be further extended via DWDM or WAN channel extension.

Table 7-2 Extended electronic vaulting technologies

Electronic vaulting technology	Commonly supported distances between sites	Common product vendors	Relative technology costs
Long Distance SAN (shortwave/longwave Fibre Channel)	up to 11 km	Brocade (OEM) Mcddata (OEM) INRANGE (OEM)	Low
DWDM (MAN) and fibre extenders	up to 180 km	Cisco Nortel Ciena Finisar CNT	Medium
WAN, IP Based Routers and Channel Extension	up to 1000's km	CNT Cisco (OEM)	Low to High (depending on solution)

SANs were designed to overcome the distance limitations of other storage channel protocols, such as SCSI. Longwave laser GBICs available on most SAN hubs, switches and directors enable a transmission distance of up to 10 kilometers (11 kilometers when including switch to host connections) when used with 9 micron diameter single-mode optical fibre. Shortwave GBICs use

multi-mode fibre and is the ideal choice for shorter distance (less than 500m from transmitter to receiver or vice versa).

DWDM is a way to open up the conventional optical fibre bandwidth by breaking it up into many channels, each at a different optical wavelength (a different color of light). Each wavelength can carry a signal at any bit rate less than an upper limit defined by the electronics, typically up to several gigabits per second. DWDMs are implemented in areas that have dark fibre available through telcos and service providers. The DWDM is deployed as part of the physical layer. It is therefore independent of protocol, simply passing signal information in the format it is received. Examples of the protocols it can support are ATM, Gigabit Ethernet, ESCON, FICON and Fibre Channel.

WAN and IP based channel extenders typically use telecommunication lines for data transfer and therefore enable application and recovery sites to be located longer distances apart. The use of WAN and IP channel extenders provides the separation for disaster recovery purposes and avoids some of the barriers imposed when customers do not have a “right of way” to lay their fibre cable. WAN and IP channel extenders generally compress the data before sending it over the transport network, however the compression ratio needs to be determined based on the application characteristics and the distance.

Network attached storage (NAS) and iSCSI solutions are beginning to offer low cost IP based storage, such as the IBM TotalStorage IP 200i. Copies of TSM storage pools and the TSM database can be stored at a remote site using IP based storage to offer a low cost implementation while utilizing existing infrastructure. Configurations can include TSM clients attached to iSCSI based data backing up to a TSM server or TSM servers using iSCSI based storage as storage pools.

For a detailed overview of technologies, products, costs and best practices with distance solutions we recommend you review *Introduction to SAN Distance Solutions*, SG24-6408.

7.11.1 Collocation considerations for offsite vaulting

With collocation, large numbers of files belonging to a client node can be restored, retrieved and recalled more quickly. However, using collocation on copy storage pools requires special consideration. Primary and copy storage pools perform different recovery roles. Normally you use primary storage pools to recover data to clients directly. You use copy storage pools to recover data to the primary storage pools. In a disaster where both clients and the server are lost, the copy storage pool volumes will probably be used directly to recover clients. The types of recovery scenarios that concern you the most will help you to determine whether to use collocation on your copy storage pools.

You may also want to consider that collocation on copy storage pools will result in more partially filled volumes and probably increased offsite reclamation activity. Collocation typically results in a partially filled sequential volume for each client or client file space. This may be acceptable for primary storage pools because these partially filled volumes remain available and can be filled during the next migration process. However, for copy storage pools this may be unacceptable because the storage pool backups are usually made to be taken offsite immediately. If you use collocation for copy storage pools, you will have to decide between:

- ▶ Taking more partially filled volumes offsite, thereby increasing the reclamation activity when the reclamation threshold is lowered or reached.

Or,

- ▶ Leaving these partially filled volumes onsite until they fill and risk not having an offsite copy of the data on these volumes.

With collocation disabled for a copy storage pool, typically there will be only a few partially filled volumes after storage pool backups to the copy storage pool are complete. Consider carefully before using collocation for copy storage pools. Even if you use collocation for your primary storage pools, you may want to disable collocation for copy storage pools. Or, you may want to restrict collocation on copy storage pools to certain critical clients, as identified by the Business Impact Analysis.

7.11.2 Reclamation considerations for offsite vaulting

Space on a sequential volume becomes reclaimable as files expire or are deleted from the volume. For example, files become obsolete because of aging or limits on the number of versions of a file. In reclamation processing, the TSM server rewrites files on the volume being reclaimed to other volumes in the storage pool, making the reclaimed volume available for reuse.

When an offsite volume is reclaimed, the files on the volume are rewritten to another copy storage pool volume which is onsite. The TSM server copies valid files contained on the offsite volumes being reclaimed, from the original files in the primary storage pools. In this way, the server can reclaim offsite copy storage pool volumes without having to recall and mount these volumes. Logically, these files are moved back to the onsite location. The new volume should be moved offsite as soon as possible. However the files have not been physically deleted from the original offsite volume. In the event of a disaster occurring before the newly written copy storage pool volume has been taken offsite, these files could still be recovered from the offsite volume, provided that it has not already been reused and the database backup that you use for recovery references the files on the offsite volume. The server reclaims an offsite volume which has reached the reclamation threshold as follows:

1. The server determines which files on the volume are still valid.
2. The server obtains these valid files from a primary storage pool, or if necessary, from an onsite volume of a copy storage pool.
3. The server writes the files to one or more volumes in the copy storage pool and updates the database. If a file is an aggregate file with unused space, the unused space is removed during this process.
4. A message is issued indicating that the offsite volume was reclaimed.
5. The newly written volumes are then marked to be sent offsite, and after this has occurred, the reclaimed volume can be returned to an onsite scratch pool.

Volumes with the access value of offsite are eligible for reclamation if the amount of empty space on a volume exceeds the reclamation threshold for the copy storage pool. The default reclamation threshold for copy storage pools is 100%, which means that reclamation is not performed.

If you plan to make daily storage pool backups to a copy storage pool, then mark all new volumes in the copy storage pool as offsite and send them to the offsite storage location. This strategy works well with one consideration — if you are using automatic reclamation (the reclamation threshold is less than 100%). Each day's storage pool backups will create a number of new copy storage pool volumes, the last one being only partially filled. If the percentage of empty space on this partially filled volume is higher than the reclaim percentage, this volume becomes eligible for reclamation as soon as you mark it offsite. The reclamation process would cause a new volume to be created with the same files on it. The volume you take offsite would then be empty according to the TSM database. If you do not recognize what is happening, you could perpetuate this process by marking the new partially filled volume offsite.

If you send copy storage pool volumes offsite, we recommend that you control copy storage pool reclamation by using the default value of 100. This turns reclamation off for the copy storage pool. You can start reclamation processing at desired times by changing the reclamation threshold for the storage pool.

Depending on your data expiration patterns, you may not need to do reclamation of offsite volumes each day. You may choose to perform offsite reclamation on a less frequent basis. For example, suppose you send copy storage pool volumes to and from your offsite storage location once a week. You can run reclamation for the copy storage pool weekly, so that as offsite volumes become empty they are sent back for reuse.

When you do perform reclamation for offsite volumes, the following sequence is recommended:

1. Back up your primary storage pools to copy storage pools.
2. Turn on reclamation for copy storage pools by lowering the reclamation threshold below 100%.
3. When reclamation processing completes, turn off reclamation for copy storage pools by raising the reclamation threshold to 100%.
4. Mark any newly created copy storage pool volumes as offsite and then move them to the offsite location.

This sequence ensures that the files on the new copy storage pool volumes are sent offsite, and are not inadvertently kept onsite because of reclamation.

Attention: If collocation is enabled and reclamation occurs, the server tries to reclaim the files for each client node or client file space onto a minimal number of volumes.



Part 2

Implementation procedures and strategies

Having covered planning strategies, this part of the book provides practical procedures for protecting and restoring TSM servers and clients. Disaster Recovery Manager (DRM) is discussed in depth, including how to set it up, how to maintain the plan, and how to recover the TSM server using the plan. Next, procedures for bare metal recovery on popular operating systems are described, including Windows 2000, Linux, Solaris and AIX. Finally, we draw together many of the threads in this book to present some complete disaster protection implementation ideas and case studies.



IBM Tivoli Storage Manager and DRM

This chapter discusses the following topics:

- ▶ Best practices for hardware design
- ▶ Best practices for TSM server implementation
- ▶ Proactive hints and tips for management and availability
- ▶ Best practices for environment management
- ▶ Disaster Recovery Manager (DRM) — what it is and what it isn't
- ▶ Detailed setup steps for DRM

8.1 Introduction

The exploding growth of corporate data combined with falling storage prices has created both an opportunity and a challenge for Information Technology managers. More affordable storage technology allows IT managers to purchase additional storage devices and store rapidly increasing volumes of corporate data. Yet managing these expanding storage networks is becoming a complex, resource-intensive task. In many organizations, storage management is executed without strategy, reducing cost-effectiveness and efficiency.

Applying the discipline of storage management, combined with the appropriate technology and a well-crafted set of Storage Management Best Practices, can provide significant business value, helping enterprises to increase revenues and decrease costs.

8.2 Best practices for hardware design

IBM Tivoli Storage Manager was developed from the ground up to be a comprehensive Storage Management solution. IBM Tivoli Storage Manager is more than just a distributed backup product: it is a distributed data backup, recovery, and Storage Management solution. Unlike many other storage products, TSM is built on a common, highly portable code base that is consistent across all Tivoli Storage Manager server platforms. This common code base allows TSM to manage computer environments ranging from a single database server, to thousands of desktop clients and everything in between.

8.2.1 TSM server platform selection from DR point of view

IBM Tivoli Storage Manager server and client software is available on many different operating system platforms and can exploit different communications protocols. IBM Tivoli Storage Manager server runs on nine different operating platforms plus the IBM Network Storage Manager (NSM). IBM Tivoli Storage Manager server V5.1 is available for these server platforms, shown in Figure 8-1:

- ▶ AIX
- ▶ Sun Solaris
- ▶ HP-UX
- ▶ Windows NT
- ▶ Windows 2000
- ▶ OS/390
- ▶ z/OS
- ▶ OS/400 (using PASE, expected availability, October 11, 2002)
- ▶ Linux (expected availability, end 2002)



Figure 8-1 Available TSM server platforms

How to choose one TSM server platform over another? Sometimes there will not be a choice because of system availability, historical reasons, or corporate policy. However if a new server is to be commissioned, then there are certain factors to consider. With only minor differences, a TSM server provides the same functionality on every platform at the same version and release level. From the TSM client perspective, there is no major difference between any TSM server — they will provide the same services, even for a bare metal recovery. However there are some differences dictated by operating system platform which affect overall costs, administration, and operations. Every platform has different procedures for bare metal recovery or re-installation of the operating system. Some systems are more scalable or more versatile in terms of peripheral support. If it is desired to use clustered TSM systems or a hot backup system, the question of cost can be important. Table 8-1 summarizes these considerations.

Table 8-1 TSM server platform consideration

	Windows NT/W2K	UNIX	OS/400	z/OS
Operating system installation and/or restore	Simple to Medium	Simple (especially AIX)	Complex	Very complex
Cost	Low	Medium	Medium to High	High
Operation	Simple	Medium to Complex	Complex	Very complex
Supported devices	Many	Many	Limited	Limited
Capacity	Medium	Medium to High	Medium to High	Medium to Very High

Let's consider each of these factors in further detail.

Operating system installation and/or restore and operation

If a TSM server is completely destroyed, and there is no hot or warm standby (that is, pre-installed/configured system), re-installation and re-configuration of the operating system on the server will be the first step to site recovery. Since the TSM server must be up first, before any client restores can take place, commissioning of the server is on the critical path. Therefore, we can speed up the overall process of DR by choosing a TSM server platform which is easy to install or recover.

Installation consists of the hardware platform installation and operating system installation and/or restore. Installing from the CD media will take some time and require certain post-installation configuration steps to bring the TSM server back to working order. This step can be expedited by using available operating system utilities or third-party packages. For example, a TSM server using Windows could be more quickly restored using one of the many third-party "ghosting" applications. AIX provides the `mksysb` command for complete backup and restore of the root volume group. The "cleaner" that the TSM server is (that is minimum of extra applications), the easier it will be to recover from scratch. For this reason, we recommend that the TSM server be dedicated for its purpose so that if a new system needs to be rebuilt, this will only require operating system re-install, TSM server install, and then use of DRM to re-build the TSM server.

The operation of the TSM server itself is fairly consistent across all platforms — the administrative commands and interfaces are the same. The largest differences are in attaching and installing devices, and in managing disk volumes for database and log volumes, and disk storage pools. If there are already

specific operating system administration skills in the enterprise, these will help to influence the choice of the TSM server platform.

Cost and capacity

Cost is a very relative parameter in an IT solution. A company's most valuable asset is its data, therefore the first thing to consider when designing a DR solution is the Business Impact Analysis. This will help to understand how much the loss or unavailability of data will cost the enterprise. The real cost of computer hardware for a TSM server represents only a portion of the total investment required to protect data in an actual IT environment. A much larger fraction is taken up by storage software, administration costs and of course services and education.

The TSM server is architected to manage a virtually unlimited number of clients and amount of data. Various operating system and system platforms however, have different capacities in regard to the CPU power it can deliver to TSM server, the number and size of devices it can attach, and the throughput it can deliver. A TSM server can be installed on a small Windows NT laptop, but this is clearly not a useful production environment. For many companies, the amount of data to be protected is doubling every 18 months. To avoid the disruption of frequent migrations of the TSM server from one system to another, or even upgrading storage devices, it is a good practice to plan for the future today and install capacity for growth. This means choosing a server platform with ample slots, so that more peripherals can be attached, it has a large maximum memory capacity, ability to add processors, and so on. If a server migration becomes necessary, moving to a larger computer of the same operating system type is relatively simple although it will require an outage. It could even be treated as part of DR testing. Moving between different server platforms will require use of the TSM **IMPORT** and **EXPORT** commands which is time-consuming as well as labor and resource intensive. Therefore, the decision to change platforms should not be taken lightly.

Supported devices

We have already frequently mentioned the enormous growth in the amount of enterprise data. Storage devices and technologies are being continually upgraded to keep abreast with this growth. There are a wide variety of supported devices available on the Windows and UNIX platforms, including disk drives, disk arrays, tape drives, optical drives and automated tape libraries. OS/400 and z/OS tend to have a more restricted range. For a list of supported devices on each TSM server platform, see the Web site:

http://www.tivoli.com/support/storage_mgr/devices/all.html

It is also important to carefully check support for SAN-attached devices for using functions such as tape library sharing, LAN-free backup and server-free backup.

Recommendations

Choose a TSM server platform where there are suitable skills existing in the company (or where these skills are readily obtainable). Consider ease of operating system recovery in terms of being able to deploy a replacement TSM server as quickly as possible. Ensure that the server system chosen will be able to attach enough storage devices to meet your capacity requirements into the future. Consider how the operating system supports clustering and high availability and any fault-tolerant technology built into the hardware.

In general, most new sites will choose either Windows or one of the UNIX platforms, as these are somewhat easier to administer, and skilled staff are readily available. Companies with significant OS/400 investment and skill will be more comfortable with an OS/400 platform. If you have skilled staff in either Windows or particular UNIX-variants, choose those platforms. For sites with little or no commitment and skills to other platforms, choose AIX. It is easy to install or restore, and robust enough to handle most implementations while remaining relatively cost effective. AIX-based TSM servers are very well supported by IBM in the operating system, TSM and DR Planning. The AIX platform scales from small to very large TSM systems.

Remember that the TSM server is CPU-intensive, I/O-intensive, and memory-intensive. The amount of data managed is almost invariably going to grow in the first year as new opportunities to use TSM in the enterprise are uncovered. In the majority of cases where a TSM server installation started on hardware shared with other applications, it has been moved to a dedicated machine within the first year. Our recommendation is to dedicate hardware specifically to the TSM server, equip it with enough memory, CPUs, and the appropriate number and type of I/O channels. A typical recommendation is at least 512 MB memory on any platform, Gigabit Ethernet card(s), and at least two SCSI or FC SAN adapters (depending on your storage devices). Make sure that the TSM server hardware has enough free slots for more adapters.

8.2.2 Disk storage devices

Tivoli Storage Manager requires a disk to operate: it needs a disk to hold the database, logs, and nearly always the primary storage pools. 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 TSM implementations grow substantially. Choose a disk model that meets your present estimated needs and has ample room for expansion. Multiple I/O paths and hot swappable components should also be considered, both for performance and availability.

To protect data on a disk-type storage device, you should consider some form of RAID, which can be implemented at either a hardware or software level. Protecting data stored on a disk is a subject in itself, and below we just touch on some of the possibilities you may want to look into further. The best solution will depend on many factors:

- ▶ Server hardware and hardware RAID solutions available
- ▶ Operating system type and software RAID solutions available
- ▶ Price, performance, and redundancy level of the solution

A hardware RAID solution can be beneficial in minimizing the performance penalty on the server, when implementing RAID. It can also provide disks and controllers that can be hot-swapped in the event of failure. Dual pathing, power redundancy, and disk hot-spare features are other possibilities. A software RAID solution can 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 fibre, which all have advantages and disadvantages.

In general IT trends, usage of direct attached storages (SCSI, SSA) is dropping and use of SAN attached storage and NAS is increasing. Analysts forecast that this trend will continue. SAN attached storage offers more flexibility and scalability in using storage resources. Trends for the future are to virtualize disk storage, making storage independent of the server type, and providing enterprise-wide data sharing.

From the perspective of high availability and disaster recoverability we recommend that you use disk arrays with hardware implemented RAID. There are many types of disk arrays on the market, which use the common RAID technologies of RAID 1, 3, 5 and 0-1. Another important feature for selecting disk technology is the ability to create instant and local mirrors of the data (for example, using Snapshot or FlashCopy functions). Even if you don't initially plan to build a DR solution at Tier Level 4 or above, having these functions available means you will be able to go in that direction in the future.

8.2.3 Tape storage devices

Most TSM systems use tape as the ultimate destination for storing their data. There are a variety of tape subsystems available for most TSM server platforms. 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 TSM implementations, because tape is generally faster and more convenient to work with. You should choose technology and methods which will minimize, if not eliminate completely, the use of manual TSM procedures for identifying and moving volumes from onsite to the

offsite location and back. DRM can be used to automate and manage these procedures.

The main reasons why we suggest the biggest, fastest, and most automated tape library are:

- ▶ Constant growth in quantity of data. Forecasted data capacity is reached often in half the projected time period. Retention requirements for data can also increase, requiring more media storage.
- ▶ As data volumes increase, backup windows are shortening, and the availability of dedicated backup windows (data offline) can even be zero in today's 24x7, mission-critical environments.
- ▶ Automation helps reduce overall costs, increases reliability (less human error) and makes the system as a whole faster.

Automation of course introduces additional technology costs, as well as requiring new operating processes and education. On the upside, automation reduces or eliminates the need for humans to handle tapes, leading to fewer errors and significant reductions in tape mount delays. A combination of fully automated tape libraries with SAN connection between primary and backup site can fully automate the process of offsite tape creation and transfer. TSM can operate an automated tape library in the backup location connected through a long-distance SAN connection as well as using a locally attached library. This represents a higher service level, since copies of the data are automatically made at the remote location and simultaneously the data is available onsite in the primary storage pool.

Creating copy storage pools for offsite storage requires at least two drives in a library. With multiple drives, housekeeping operations such as reclamation are also much easier. The TSM backup client and many of the API applications (for example, DR backup for SAP, RMAN for Oracle) are able to automatically and simultaneously stream the data to multiple drives, providing efficient utilization of multi-drive libraries. For better data and system availability, we therefore recommend usually at least three drives in a library. This provides better availability in the event of failure of one drive and allows more concurrent TSM operations to occur. Detailed TSM planning will determine the optimum number required.

There is a large variety of different tape technologies available on the market, each has different price points, speed, reliability, capacity, and product form factors (single drives, libraries, automated libraries). Table 8-2 shows some of the most popular types broken down into various classifications.

Table 8-2 Classes of tape drive by CNT

	Type I (Up to 5 MB/sec.)	Type II (5-10 MB/sec.)	Type III (10-15 MB/sec.)	Type IV (over 15 MB/sec.)
Enterprise Class (Class 3)	N/A	IBM 3590B STK 9840	IBM 3590E	IBM 3590H
Network Class (Class 2)	Mammoth-1	DLT 7000 AIT-1 Mammoth-2 IBM 3575	SuperDLT LTO-1 AIT-2	LTO-2 AIT-3
Server Class (Class 1)	DAT (4mm, 8mm)	N/A	N/A	N/A

With large amounts of data, the 4mm, 8mm, and QIC tapes are a good less suitable fit and they do not have large capacity. They are also becoming increasingly less common. DLT is a quite common tape system, and may store large amounts of data, but also has a number of restrictions. These tape formats should only be recommended for smaller systems.

LTO drives and tape libraries (IBM 358x product line) are a good choice if there is no requirement to attach to a mainframe. This tape technology offers high performance and reliability and can connect to multiple operating systems for library sharing. IBM 3584 LTO library provides coexistence with DLT drives, therefore protecting the present investment in DLT technology.

The IBM 3590 class type system is a high-end system usable in any environment, including mainframe. The 3590 technology offers high speed, capacity and reliability.

8.2.4 Storage Area Network (SAN)

Before SANs, each server required its own dedicated storage. The explosion of servers caused serious storage management problems, because of the proliferation of storage devices and the inability to scale and share data.

SANs have altered the nature of distributed systems storage. SANs create new methods of attaching storage to processors by repositioning storage devices onto a shared network and by providing new, improved data access methods to this storage. These new connectivity and access options promise improvements in both data availability and performance. One operational area benefiting from the development of SAN technology is data backup/restore and archive/retrieve. For backup/restore operations, SANs allow:

- ▶ Sharing and dynamic allocation of shared storage resources, such as tape drives within a library.
- ▶ LAN-free backup: Off loading data movement from LAN to the SAN during backup/restore operations.
- ▶ Server-free data movement: Not only off loading backup/restore data movement from the LAN, but also greatly reducing the application server processor data movement from the backup server through the use of a third-party copy agent.

TSM has progressively included SAN exploitation in the product, starting with tape library sharing, first made available in V3.7. IBM TSM Extended Edition V5.1 now includes the base tape library sharing, plus LAN-free backup to both tape and disk (using Tivoli SANergy) and server-free backup. LAN-free and server-free backup allows the data traffic to be off-loaded to the SAN, instead of moving the data through file servers over the LAN. This reduces LAN and server overhead, minimizing the impact of backups on the efficiency of the network.

8.2.5 Common mistakes in backup/recovery hardware planning

Here are some statements which are sometimes made by those planning a backup/recovery environment. They are usually driven by the desire to save money — but at what overall cost?

- ▶ “This server is a bit old, but it fits the minimal CPU/memory requirements for a TSM server.”

Do you really mean that the TSM server is less important than your other application servers? To provide high quality backup and restore services which meet the enterprises RTOs, do you think that an old machine, which unusable for other applications, is enough?

- ▶ “We would like to use a small tape library. The overnight backup window is eight hours which is sufficient to perform the incremental backup.”

The backup window and backup time is not so important as restore. Restore is not incremental and if disaster strikes you will need to restore all your critical data as soon as possible. How long will a restore from a small library take?

- ▶ “A single drive library, or a library with a simple autoloader is enough for our environment.”

The library with one drive, or a drive with a simple autoloader does not allow TSM to make copy storage pools. It also greatly complicates the reclamation operation, since it requires additional disk space and data movement. Using this kind of library might save money in the short term, but will require much more manual (and mistake-prone) intervention.

As much attention should be paid to the backup server as to other production servers. After all, the protection of the production server depends on it. This means having the right storage capacity, performance and component quality.

8.3 Best practices for TSM Server implementation

Here we discuss some specifics for setting up your TSM server for availability and recoverability.

8.3.1 TSM database and recovery log

The database is the most important part of the TSM server. It stores and manages all the information about client files residing in storage pools and provides for the other server components. The TSM database is not a relational database but provides an SQL interface for performing queries. An administrator can access this function using the TSM administrative command line or Web browser interfaces by entering SQL `SELECT` statements. The TSM database contains the following information:

- ▶ Metadata information (information about stored files)
- ▶ Information about registered client nodes
- ▶ Access control information for administrative clients
- ▶ Policies and schedules
- ▶ Activity log and event records
- ▶ Data storage inventory
- ▶ Information about TSM volumes

The recovery log is used by the server to keep a record of all changes to the database. When a change occurs, the recovery log is updated with some transaction information before the database is updated. This enables uncommitted transactions to be rolled back during recovery so the database remains consistent.

The TSM database consists of one or more database volumes on a disk. Spreading database volumes over multiple physical volumes may improve performance, because this allows the logical volume manager to spread the I/O load over more volumes. Database volumes can also be mirrored within TSM (2-way or 3-way). As well as providing higher availability (continued operation even if a disk containing a database volume fails, providing mirrors are created on separate disks), mirroring can improve performance, because the logical volume manager can perform read operations to the least busy volume in a mirror set. Figure 8-2 shows a TSM database configured with 3-way mirroring on the Windows platform. Each mirror copy is on a separate physical volume.

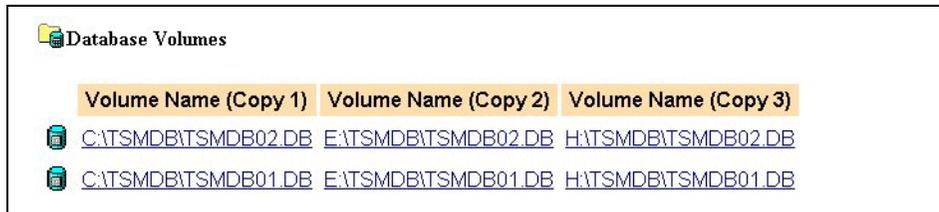


Figure 8-2 TSM database with two database volumes and triple volume mirroring

The TSM recovery log also consists of one or more recovery log volumes on a disk. Access to the recovery log is predominately write-oriented, with the writes and the few reads clustered together. The writes are done in moving cursor format which does not lend itself to multiple volume organization. Therefore, fewer recovery log volumes are normally required. Mirroring of the recovery log is highly recommended, even if database mirroring is not done (usually for reasons of the cost of disks). Recovery log volumes can also be mirrored within TSM (2-way or 3-way). Figure 8-3 shows a TSM recovery log configured with 3-way mirroring on the Windows platform. Each mirror copy is on a separate physical volume.

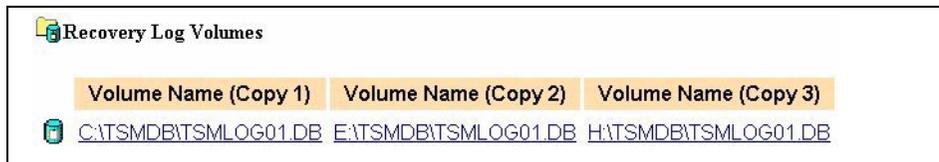


Figure 8-3 TSM recovery log volume, triple mirrored

The TSM server can use one of two log modes for recovery log: NORMAL or ROLLFORWARD. The log mode determines how log TSM saves records in the recovery log and the kind of database recovery you can use. When log mode is set to NORMAL, TSM saves only those records needed to restore the database to the point of the last backup. TSM deletes any unnecessary records from the recovery log. Changes made to the database since the last backup cannot be recovered. In NORMAL log mode, you may need less space for the recovery log, because TSM does not keep records already committed to the database. In the ROLLFORWARD log mode TSM saves all recovery log records that contain changes made to the database since the last time it was backed up. Log records are deleted only after a successful database backup. In this mode the database can be restored to the most current state (rollforward recovery) after loading the most current database backup series. Using ROLLFORWARD log mode may require a significant amount of space to record all activity.

Recommendations for TSM database and recovery log design

- ▶ For the performance and availability reasons define one or two mirror copies of each database and recovery log volume.
- ▶ Place these volumes on separate disks to eliminate impact of single disk failure and allow simultaneous access to different parts of the database and recovery log.
- ▶ It can be good practice to place one set of database and recovery log volumes on local disk(s) and a mirror copy to an external (SAN or SCSI attached disk array). This protects the database and recovery log volumes from media failure, and in the case of disaster, even if the TSM server is destroyed, an external (especially remote) disk array can survive and the mirrored volumes can be used to restore the TSM server. Considering that the disk array can be remotely mirrored using technology like PPRC, this can significantly improve the restore process in case of disaster.
- ▶ To use disk space for the database efficiently, allocate a few large disk volumes on disk drives rather than many small disk volumes. This avoids losing space to overhead processing and makes the system more manageable.
- ▶ For improved reliability, use the most reliable file system of your TSM server platform for allocating database and recovery log volumes. Example: in Windows environment use NTFS file system, rather than FAT or FAT32 file system. If you use the most reliable file system, you can take advantage of operating system ability to recover from problems that can occur during I/O to a disk.
- ▶ If using NORMAL log mode, a recovery log size of 10% of the database size should be sufficient. If ROLLFORWARD mode is used, the space for the log will be greater, depending on other factors like the frequency of DB backup, and amount of backed up files. The selection of the log mode depends on the required RPO. For disaster recovery to the most current state we recommend using ROLLFORWARD log mode.

8.3.2 TSM server database management

TSM database and recovery log volume mirroring protect the TSM server in case of individual media failure. For DR (and possible DB corruption problems), you also need regular TSM database backups. TSM can perform full and incremental database backups to sequential storage pools while the server is running and available to clients. The recovery log mode does not affect the backup type. TSM provide three type of database backup, shown in Figure 8-4:

- ▶ Full backup
- ▶ Full plus incremental backup
- ▶ Snapshot backup

Backup database Help

Device Class: CLASS1

Type: INCREMENTAL FULL DBSNAPSHOT

Volume names:

Use scratch volumes ? YES NO

Wait for completion ? YES NO

Step 1 of 1 Finish Cancel

Figure 8-4 Type of TSM database backup

A full database backup is required when:

- ▶ The database has never been backed up.
- ▶ Maximum incremental backups (32) are reached.
- ▶ Database has been extended or reduced.
- ▶ Recovery log mode changed to ROLLFORWARD.

A full backup takes longer to run than an incremental backup, because it copies the entire database. However, with the full backup you can more quickly recover your TSM server because only one set of volumes needs to be loaded to restore the entire database. Incremental backup takes less time to run because it copies only database pages that have changed since the last time the database was backed up. However, this increases the restore time because you must first restore the last full backup followed by all the incremental backups.

Snapshot backup is a full backup that does not interrupt the full plus incremental backup series. Snapshot backups are usually shipped offsite (with the assistance of DRM) while regular full and incremental backups are kept onsite.

TSM provides three types of TSM database recovery:

- ▶ Point-in-time recovery — restore full backup and all incremental backups
- ▶ Rollforward recovery to most recent state — restore full backup and all incremental backups plus rollforward the recovery logs
- ▶ Single database volume recovery

Table 8-3 shows a comparison of point-in-time recovery and rollforward recovery.

Table 8-3 Comparison of point-in-time and roll-forward recovery

Point-in-time recovery	Roll-forward recovery
Smaller size for recovery log.	Database can be restored to most current state.
After restore no audit required.	No auditing of database and storage pools required after database restore.
Cannot restore to most current state.	Single database volume can be restored
Single database volume cannot be restored.	Larger recovery log.
Auditing of storage pools required after restore.	Mirrored recovery log consumes additional disk space.
NORMAL or ROLLFORWARD logmode	ROLLFORWARD logmode only

Recommendations for TSM database management for DR

- ▶ To restore the TSM server after a disaster you will need also these four files: dsmserv.dsk, dsmserv.opt, devcnfg.out, and volhist.out. Make copies of this files at least after each full backup, and keep them off-site. This ensures your ability to recover the TSM server smoothly. DRM automatically keeps copies of these files.
- ▶ A combination of backup techniques, using full database backup and ROLLFORWARD log mode can significantly help shorten the restore time to the last most current state in case of disaster. For DR purposes we recommend to do frequent full database backups or snapshot database backups rather than one full database backup followed by a long series incremental backups.
- ▶ For high availability and DR protection of the TSM server at Tier Levels 4-6 you can backup the database using device class of type FILE on a remote mirrored disk space. This backup is quick and will allow fast recovery of the database in event of a disaster. You should do this even if you are remote mirroring the database and recovery logs, since there is still a small possibility of application corruption to all mirrored copies. In such cases you need to restore the database from backup.
- ▶ If you have multiple TSM servers in your configuration, you could use server-to-server virtual volumes to make additional copies of your database backups. This is particularly useful if the servers are in different physical locations.

8.3.3 TSM storage pools

A TSM storage pool is a collection of like media (using the same device class), that is the destination of backed-up, archived, or migrated data. Storage pools can be defined on direct access storage (disk, disk array), sequential access storage (tape drive, tape library), optical media (optical drive, optical library) and virtual volumes (server-to-server communications). TSM data storage is defined as the collection of storage pools. Backed up, archived or migrated files may initially be placed on different storage pools according to the desired policy and usually are migrated automatically to other storage pools to satisfy demands for utilization, performance, and recovery. TSM provides this ability to organize storage pools in a storage hierarchy to take advantage of different device characteristics. For example, small client files will backup fastest to a random-access (disk) storage pool, and be migrated to tape. Large database files usually provide best performance and cost benefit if streamed directly to tape.

Disk storage pools

Disk storage pools are usually used for primary storage pools. Disks as devices with direct access allow TSM to backup multiple clients simultaneously to the same disk storage pool. Client data is then usually migrated to a tape storage pool. You can ensure that files remain in a storage pool for a minimum amount of time before the server migrates them to another pool, by setting a migration delay period for a storage pool (MIGDELAY parameter) as shown in Figure 8-5. Files can temporarily remain in disk storage pools even after migration, space permitting. This is done by enabling caching for disk storage pools.

Using the parameters Migration Delay and Cache Migrated Files can be useful for some DR cases, particularly if the disk storage pool is located on a storage system with a remote copy. Important files can therefore be kept in a disk storage pool for quick access.

Define a new disk storage pool
Help

Storage Pool Name	<input type="text" value="Disk_STG_critical"/>
Description	<input type="text" value="Disk STG for critical data"/>
Media Access Status	<input type="text" value="READWRITE"/>
Maximum Size Threshold	<input type="text" value="NOLIMIT"/>
Next Storage Pool	<input type="text" value="LTO_STGP_01"/>
High Migration Threshold	<input type="text" value="80"/>
Low Migration Threshold	<input type="text" value="20"/>
Cache Migrated Files?	<input checked="" type="radio"/> YES <input type="radio"/> NO
Migration Processes	<input type="text" value="1"/>
Migration Delay	<input type="text" value="7"/>
Migration Continue	<input checked="" type="radio"/> YES <input type="radio"/> NO
Copy Storage Pool(s)	<input type="text" value="LTOCopySTG_01"/>
Continue Copy on Error?	<input checked="" type="radio"/> YES <input type="radio"/> NO
CRC Data	<input type="radio"/> YES <input checked="" type="radio"/> NO

Step 1 of 1
Finish Cancel

Figure 8-5 Define disk storage pool with migration delay

Tape storage pools

Tape storage pools in most TSM installations store the majority of the data volume. Tape storage pools can hold more data than disk storage pools. Unlike disk, the tape medium provides sequential access. TSM maintains and optimizes the utilization of tape media by the space reclamation process. Space reclamation does not directly influence the DR process; however if tape volumes are sparsely utilized due to expiring and deleted files, data recoveries will take much longer.

The TSM server has the ability to collocate client data on tape volumes. When files are moved to a collocated storage pool, TSM ensures that the files for a specific client are written to the same set of tapes. This can limit the number of tapes that must be mounted when restoring that client's system. Collocation can be done at the client level or by individual filespace. When you are deciding whether or not to enable collocation, keep in the mind:

- ▶ Non-collocation increases performance on backups because TSM does not have to select specific tapes.
- ▶ Collocation increases performance on restores because client data is confined to their own dedicated set of tapes. Therefore there is less necessity to “skip” data not needed for the restore.

Collocation is a parameter of any given sequential access storage pool. Each client whose data is placed in that storage pool will have its files collocated. From a DR perspective, collocation is recommended for storage pools containing data which has the shortest RTO. Backup data of common workstations can be held in non-collocated tape storage pools. Consider carefully whether to use collocation on copy storage pools. It will dramatically increase the number of tapes used each day to create the offsite copies.

Copy storage pools

Copy storage pools contain duplicates data of data residing in primary storage pools. The tape volumes in a copy storage pool can be sent offsite in the event of disaster. Data is copied to storage pools incrementally. Only data that has not previously been copied to the copy storage pool is copied during a backup cycle. The backing up of a primary storage pool to the copy storage pool should be done after all the daily client data has been backed up or archived.

For lower Tier disaster solutions, we recommend that you maintain two copy storage pools — one kept onsite, and one taken offsite. The onsite copy storage pool will provide availability in the case of media failure in primary storage pool. The offsite copy storage pool protects data in the case of a disaster. Alternatively, a SAN-connected tape library in a remote location (using electronic tape vaulting) reduces the need for two copy pools, because the copy pool on this library combines the benefits of the onsite and offsite storage pools. Naturally, the performance of restoring data from an offsite SAN-attached tape library should be carefully benchmarked to ensure it will meet the RTO.

Given that one to many primary storage pools can be configured in TSM, we recommend that you backup all storage pools in each single hierarchy to the same copy storage pool. This can help recovery in the case of disaster. More than one storage pool hierarchy can be backed up to the same copy storage pool.

8.4 Best practices for environment management

The criticality of data and available resources will determine the best method of disaster preparedness for the TSM server. In the event of a disaster which destroys the TSM server at the primary site, it is necessary to re-establish the TSM server at an alternative site. This alternative site can be setup as a cold, warm or hot standby site.

8.4.1 Cold standby

In this scenario the TSM server is rebuilt on a new machine with the same operating system and TSM server software, TSM database backups and configuration files. It is highly recommended that Disaster Recovery Manager be used to speed up the process of recovering the TSM server.

8.4.2 Warm standby

In this scenario the TSM server code is duplicated on another physical machine or is installed as another instance on another TSM server. The TSM server application itself is installed but not running. The standby TSM server should be at the same operating system level (including patches) and TSM version (including patches) as the original.

The speed of Switchover/Failover will depend on configuration and access. Two scenarios will be discussed here.

No access to primary TSM server volumes

In this case, there is no access to either the primary TSM server database volumes, or primary storage pool volumes. This would be because the primary TSM server has failed and there is no physical path to the database, log or stgpool volumes or to the physical tape library. In this scenario, any disk storage pool data that was not previously migrated to tape will be unavailable and the currency of information is only as good as the last TSM database (offsite) backup. The offsite database backup and copy storage pools are available.

Here is what is needed to accomplish the restore:

- ▶ TSM server database backup
- ▶ TSM copy storage pool volumes
- ▶ Configuration files

Although it is possible to recover the TSM server without configuration files (such as volume history and device configuration), the availability of these will speed up backup service resumption. DRM plays a key role in making sure these files are available and in automating the recreation of the space needed for the database and logs.

- ▶ Current size of the primary TSM server's database and logs

It is required to have the database volumes formatted and available before restoring the TSM database. Not doing so incurs the extra and lengthy overhead of having to get the correct allocation size and format them. The total size allocated must be equal to or greater than that allocated in the original TSM server. DRM automates the collection of this information and re-creating the required volumes.

- ▶ Tape drive or library compatible with the primary server

The TSM standby system must have access to a tape drive or drives that are compatible with the tape format of the copy storage pools and TSM database backups. A device configuration file is needed for access to the tape library — if it is not available, one must be manually created, adding to the service resumption time.
- ▶ The ability to assume the primary TSM server's TCP/IP address

The standby TSM server needs a method of assuming the primary TSM server's TCP/IP address; otherwise the TCPSEVERADDRESS parameters for all clients must first be changed, adding to the overhead of resumption of service.
- ▶ New storage pool volumes (optional)

If the standby TSM server is to be in service for a length of time, storage pool volumes must be available in order to receive future backups. If the storage pool volumes were not already available, formatting them will add to the time required to resume full backup service.

Access to primary TSM server volumes

In this case the primary TSM server has failed but there is still a physical path or one could be provided quickly through switching or re-cabling to the TSM database volumes and disk storage pool volumes, TSM configuration files and the physical tape library.

Here is what needs to be done in this instance:

- ▶ The primary TSM server must be completely down

The standby TSM server cannot be started unless the primary TSM server is completely down, otherwise a lock file error will occur. The assumption of the primary server TCP/IP address cannot occur until the machine itself has been turned off or the TCP/IP address associated with the TSM server disabled.
- ▶ Verify the device configuration file to speed up the tape library and driver discovery.

As a result of re-cabling of alternate SCSI or SAN paths, there is a good chance that the device configuration file used for the primary TSM server will not work because of SCSI ID reassignments. To speed up recovery a validated device configuration file should be pre-defined.

For SAN environments, the HBA card must be supported and the correct firmware loaded. Using the same HBA cards and firmware level used in the primary TSM server will facilitate a successful switchover.
- ▶ A method of assuming the primary TSM server's TCP/IP address

The standby TSM server needs a method of assuming the primary TSM server's TCP/IP address; otherwise the TCPSEVERADDRESS parameters for all clients must first be changed, adding to the overhead of resumption of service.

8.4.3 Hot standby

In this scenario a TSM server is installed and running either on a backup machine (hot standby, as shown in Figure 8-6), or as another instance (hot instance) on another TSM server.

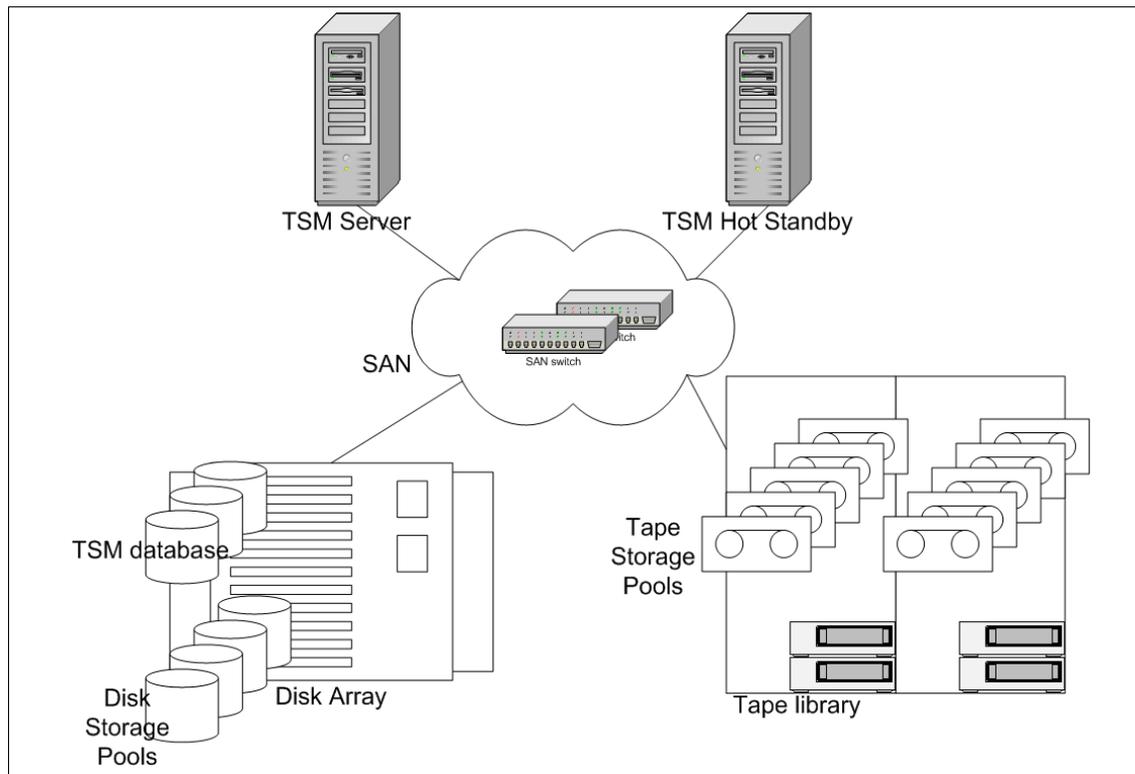


Figure 8-6 Hot standby TSM server and TSM server

In the event of a disaster, the TSM database must first be updated on the standby TSM server — basically this operation is a `DSMSERV RESTORE DB` operation. In order not to interfere with the database backup series of the primary TSM server a point in time restore should be done with a database snapshot backup. Unless a current volume history file with this database snapshot backup entry is available, a restore without volume history file operation will have to be done. The database snapshot backup should be done immediately after the primary TSM

server regular database backup is done to capture the committed changes from the log.

Requirements

- ▶ Same operating system version and PTF level.
- ▶ Same TSM version and PTF level.
- ▶ Access to tape hardware compatible with the primary server tape format
- ▶ Equal or greater database and storage pool volume capacity to the primary TSM server.
- ▶ TCP/IP address on the standby server must not be the same as the primary while primary is active, but must have the ability to assume the primary server's TCP/IP address in a disaster or change client TCPSERVERADDRESS parameters in the client options file on all TSM clients.

The next two examples show a hot standby TSM server where there is automatic synchronization of the server database using PPRC or other remote mirroring technique. First, Figure 8-7, which shows a remote TSM server running as a hot standby. The TSM database volumes, recovery log, disk storage pools, and configuration files reside on a mirrored disk array. This scenario allows the standby TSM server to assume the primary TSM server responsibilities in a minimal amount of time.

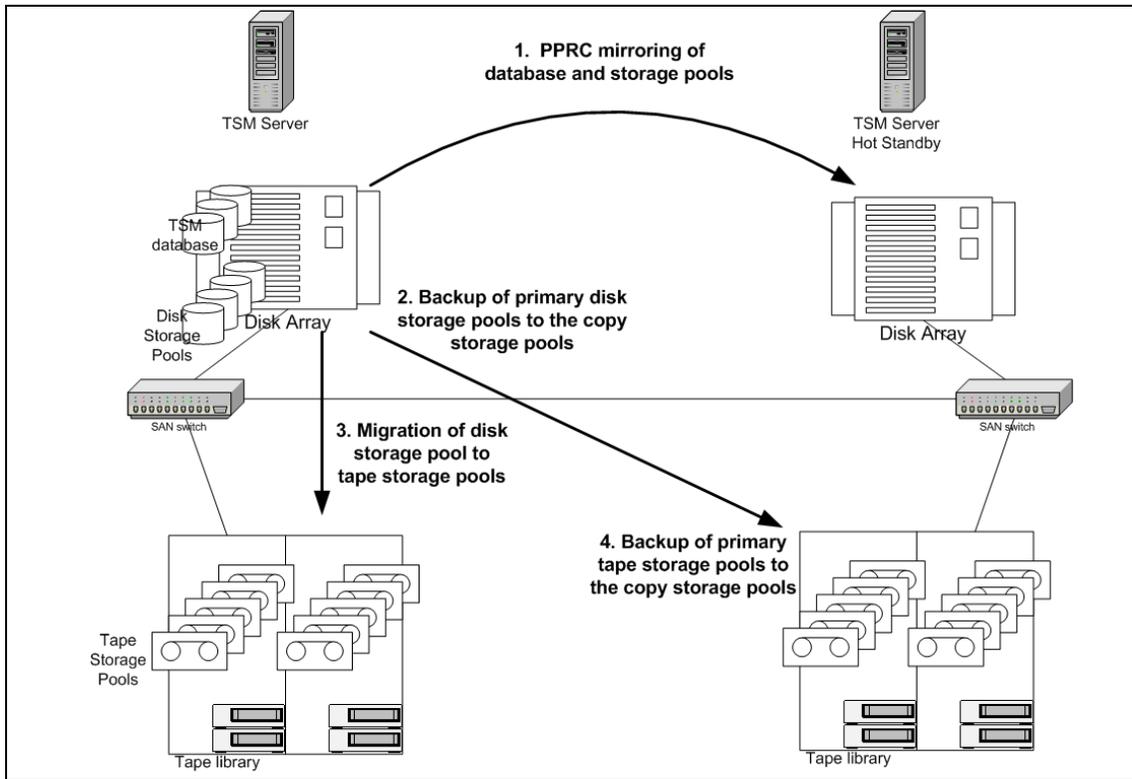


Figure 8-7 Hot standby TSM server

Figure 8-8 shows a more sophisticated setup, where two TSM servers each contain a hot instance of each other which can be used for DR in the event of either one failing. In this scenario, in normal operations, the two TSM servers run separate, independent workloads at separate sites. Each server's database and storage pools is mirrored remotely, with a standby TSM instance installed also. If a failure took out one site, the surviving site would then be able to run both TSM servers on the same system.

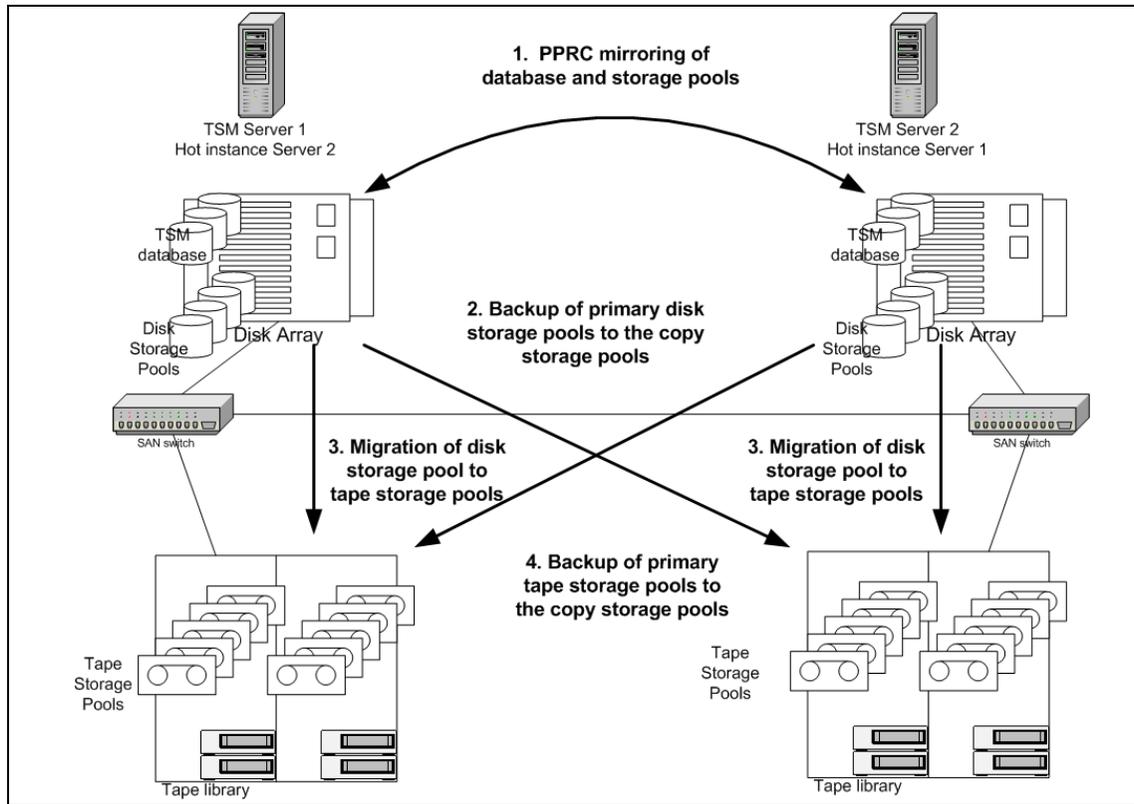


Figure 8-8 Hot instance on TSM server

Requirements

- ▶ Same operating system version and PTF level.
- ▶ Same TSM version and PTF level.
- ▶ TCP/IP address on the standby server must not be the same as the primary while primary is active, but must have the ability to assume the primary server's TCP/IP address in a disaster or change client TCPSERVERADDRESS parameters in the client options file on all TSM clients.

8.4.4 Configuration considerations

In order to implement this solution, we need to take care of the following items.

Machine type

We recommended that you use a similar class of machine with equal or greater amount of CPU power and memory for the recovery system. A smaller class of machine may be used if your standby configuration is only a short-term bridge until the original production TSM server machine is restored. This may mean an elongated backup window and that the restore of clients will need to be staged based on need. The same is true if you elect to host the standby TSM server as a second instance on a system. Normally the application will be in some type of standby mode consuming little, memory, network, CPU and I/O director resources. In the case of failure when both instances will be running actively, depending on machine size, both TSM servers may suffer degradation due to contention for resources. In this case decisions have to be made on what clients get the best service for restore and backup.

TCP/IP addressing

You have to make sure that either the recovery machine has the same TCP/IP address as the original TSM server, or that all TSM client option files have to be updated.

TSM database and log size

You need to use DRM to capture this, or record manually the sizes. They will be needed during restores.

Storage tape device

A compatible tape device will be needed to do the database and storage pool restores from. Note that if you use a smaller capacity library (or even a manual device), this may obviate doing additional tape swapping and manual mounts in the recovery site, which could potentially slow recovery times considerably.

Configuration files

Configuration files need to be made available on the standby servers – they include devconfig, volhistory, and dsmserv.opt.

Summary

IBM Tivoli Storage Manager and Disaster Recovery Manager can provide a method to restore the production TSM server for backups/restores and Space management. What needs to be considered is what scenario is best for you, cold, warm or hot standby. This will be driven the needs of your business.

8.5 Proactive hints/tips for management availability

High availability and resilience to the effects of hardware, software and environmental failures is an important requirement for a data management

solution managing business critical information. An outage of a system used to support a mission critical interactive application (for example, a customer call center application), can result in lower customer satisfaction as well as lost revenue. Consequently, the availability and speed of recovery needed in a storage management solution depends on the business requirements, the service levels that are in place and the budget available to implement the solution.

IBM Tivoli Storage Manager incorporates features that can be exploited to give a range of protection for the storage management server. This extends from hardware disk failure, through to immediate fail-over to alternative hardware in another location; without loss of stored data or interruption to the service.

8.5.1 Failover

The TSM server fully supports the process of failover to alternate hardware in the event of failure of the TSM server hardware, where the servers are clustered using the applications IBM HACMP or MSCS from Microsoft. To achieve this, the TSM server database and storage pools are allocated on shared disk between the primary and fail-over servers. The fail-over product monitors the TSM server process and hardware. In the event of failure of the primary server the fail-over process restarts the TSM server on the fail-over server.

Where a mirrored copy of the TSM server database available, as provided by the DBMIRRORWRITE SEQUENTIAL or PARALLEL and DBPAGESHADOW options, the TSM server can be restarted immediately; without the loss of service. Any currently scheduled and executing client or server tasks and operations will restart and continue.

For the Microsoft environment, TSM is a fully cluster-aware application that fully integrates with and takes advantage of MSCS's clustering and administrative capabilities. TSM uses MSCS fail-over to offer two configurations:

Active/passive configuration

In this configuration one instance of a TSM server is created that can run on either node. The active node performs the normal TSM functions and the passive node serves as an online (hot) backup. In this configuration, the server has one database, recovery log, and one set of storage pool volumes.

Active/active configuration

This configuration enables the cluster to support two independent instances of a TSM server. Although the instances typically run on separate nodes, one node can run both instances. Each TSM server has a separate database, recovery log,

and a separate set of storage pool volumes. This is because virtual servers cannot share data, nor can they share the database or recovery logs.

Tape device failover can be achieved between two servers in a cluster by dual attaching the tape devices to both server platforms. A number of different configurations are possible:

- ▶ IBM 3590 tape drives have dual SCSI ports, enabling each drive to be separately cabled to each server.
- ▶ Drives with only a single SCSI port can be connected using dual ended SCSI cables with the device connected between both primary and fail-over Tivoli Storage Manager servers.
- ▶ In a Fibre Channel environment the drives can be zoned to both servers.

8.5.2 Geographic failover

It is a small step from providing high availability through server failover that protects against individual hardware failures to configuring for geographic failover to provide full disaster tolerance. Hardware disk replication technologies, such as IBM's Peer-to-Peer Remote Copy (PPRC) or EMC's SRDF can be used to provide real-time mirroring of the TSM database to a remote site. These functions provide an online copy of the TSM database at the remote site that would be used to resume the service in the event of a disaster.

These advanced hardware replication technologies allow only complete writes to be mirrored to the remote site. This protects the TSM database from incomplete partial writes at the I/O level. However TSM database updates may be broken into multiple I/Os at the operating system level. Consequently, it is necessary to replicate both TSM managed mirrors of the TSM database to the remote site to protect against a broken chain of I/Os.

Using this approach, the TSM server can be restarted immediately following a disaster. Consequently, when access has been provided to the offsite storage pool copy tapes, recovery of client systems can commence. It provides a recovery point that is determined by the last time offsite tapes were recreated and sent to the recovery site.

Further resilience is achieved by replicating the TSM disk and tape storage pools in real time to the remote site as well as the database and logs. This provides failover without any loss of data at any level in the storage hierarchy.

To enable this in real-time, network access to the tape libraries at the remote site is required to store the offsite copy of the data. There are two approaches to this: use of TSM's virtual tape vaulting over a TCP/IP network, or by having a Fibre

Channel connection to the remote tape library. Access to remote tape allows creation of offsite tape media data copies without any manual tape movement.

8.6 DRM — what it is and what it is not

Most organizations agree that backing up data is a critical part of business continuance. Without the proper safeguards, an equipment malfunction or physical catastrophe can cause mission-critical data to be lost. To complicate matters, data protection has become a more difficult task as information is more widely distributed geographically and organizationally. In the flurry of activity following a disaster, when operational chaos can make key storage management decisions more complex, you cannot afford to worry about these issues. An airtight, organized roadmap to fast recovery must be in place. Tivoli DRM can ensure recovery and help get an enterprise functioning and operating quickly. Tivoli DRM helps you maintain business continuance by:

- ▶ Managing server database and storage pool backup volumes
- ▶ Establishing a thorough DRP for the TSM server
- ▶ Tracking and reporting client systems destroyed, in the event of a disaster
- ▶ Automating vital server recovery steps to bring your business back to normal operation
- ▶ Prioritizing client system restores.

Manage the details needed to automate recovery

The automated and integrated Tivoli DRM server helps keep essential data secure. Tivoli Storage Manager automates the planning, preparation, and execution of a DRP for data managed in a TSM environment. When you generate a DRM plan file, up-to-date information from the server is gathered to prepare for restoration. Tivoli DRM intelligently manages and tracks the movement of backup media that helps determine which media to keep onsite and which to move offsite. In the event of a disaster, DRM can locate backup media quickly — whether it is onsite, in transit, or offsite, in a vault. In addition, the client management tracking capabilities of TSM enable you to identify systems destroyed in a disaster and determine hardware and software requirements of damaged equipment. Accounting for such equipment makes it possible to properly and quickly replace what was damaged. Tivoli DRM stores vital information: the priority-level order in which to restore machines, the people to contact in case of a disaster, and the location of the operating system installation media.

Simplify the planning and audit processes

Tivoli DRM minimizes the time you spend in the DR Planning process by automatically collecting all the information needed for a recovery. By prompting users for vital information and details important to the recovery process, Tivoli DRM helps you establish a thorough DR plan. Tivoli DRM maintains the DRP in a central location — the TSM server itself, simplifying the audit process. Likewise, the clear instructions provided by Tivoli DRM make it easier to test all or part of the plan.

Consistent integrated storage management

Tivoli DRM is integrated with Tivoli Storage Manager and gathers information directly from the TSM database. It helps ensure that multiple versions of the recovery plan are in the TSM database. You can issue the commands to create and generate the Disaster Recovery Plan from the TSM server. In addition, you can use TSM scheduling features to make sure DR preparation commands take place on a regular basis. Because of the integration of DRM with TSM, you can use server-to-server communication to store remote copies of your recovery plan on virtual volumes.

Table 8-4 summarizes the main functions of DRM.

Table 8-4 DRM principal features

Management Service	What It Does	What It Means to You
Automated Disaster Recovery Plan	Provides comprehensive solutions for disaster recovery.	Safeguards all your equipment and data.
Database Management	Manages and tracks TSM server database and storage pool backup volumes.	Lets you know exactly what has been done and what may still need to be done.
Ready-to-use Guides	Guides operations staff step-by-step through the recovery process.	Reduces the chance of manual errors. Trains your staff to be knowledgeable in the event of disaster.
Extended Storage	Stores multiple versions of the disaster recovery plans inside TSM database.	Generates a flat file copy of the plan.

8.7 Example of DRM execution

This section gives a practical example of using Disaster Recovery Manager. In our lab we have a Windows 2000 TSM server, RADON, connected to an IBM 3583 LTO Library via IBM SAN Data Gateway to the SAN. The TSM server is called RADON_Server1. We will set up DRM on RADON, backup all the necessary data and restore the TSM server on to another Windows 2000 system called LEAD. The result will be that the TSM server named RADON_Server1 will run on the machine called LEAD as shown in Figure 8-9. Our situation most closely matches a disaster recovery with a cold standby machine.

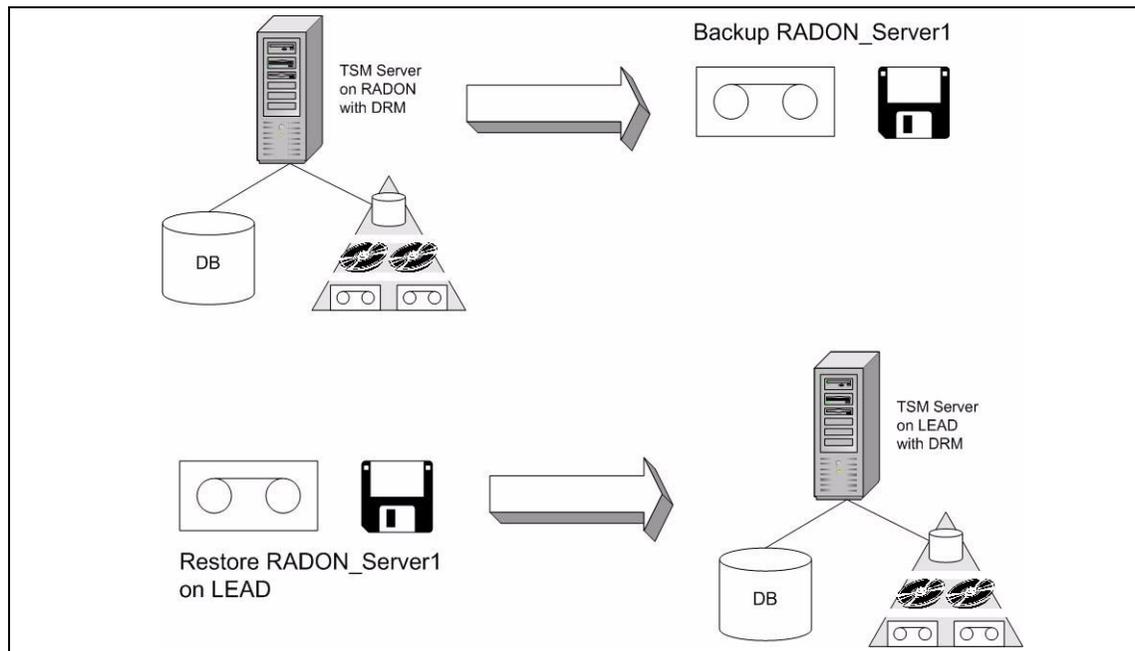


Figure 8-9 DRM lab setup

Here is a summary of the steps we will perform:

- ▶ TSM server settings
 - DRM setup
 - Register DRM license
 - Define backup copy pool
 - DRM settings
 - Daily operations
 - Backup the primary storage pools to the copy storage pool
 - Backup TSM database

- Move DRmedia
 - Remove the volumes from library
 - Run PREPARE
 - Transport DR media off-site
 - Get free media from off-site back
- ▶ Restore operations on backup server
 - Install TSM server
 - Create instruction files
 - Restore server using DRM scripts

8.7.1 DRM setup

First we need to configure Disaster Recovery Manager using the administrator interface.

Register DRM license

To use DRM, it must first be licensed in TSM. Here we register the DRM license using the TSM command-line interface.

Example 8-1 Register DRM license

```

tsm: RADON_SERVER1>register license file(drm.lic)
ANR2852I Current license information:
ANR2828I Server is licensed to support Tape Library Sharing.
ANR2827I Server is licensed to support Managed System for LAN for a quantity of
20.
ANR2827I Server is licensed to support Managed System for SAN for a quantity of
5.
ANR2853I New license information:
ANR2860I Server is licensed to support Tivoli Disaster Recovery Manager.
ANR2828I Server is licensed to support Tape Library Sharing.
ANR2827I Server is licensed to support Managed System for LAN for a quantity of
20.
ANR2827I Server is licensed to support Managed System for SAN for a quantity of
5.

```

You can check if license was registered correctly using the command **QUERY LICENSE**.

Create backup copies of primary storage pools

To take tapes to an off-site location we will need a backup of the primary storage pools. We will create one copy storage pool for all primary storage pools, and we make a backup. Example 8-2 shows using the **DEFINE STGPOOL** command to create a copy storage pool, called LTOCopySTG_01 using the device class

defined for the LTO tape library. Note the value of REUSEDELAY=14. This is important for ensuring correct recovery integrity. See Step 8 on page 197 for more information on how this relates to the database backup series expiration value.

Example 8-2 Define copy storage pool

```
tsm: RADON_SERVER1>DEFINE STGPOOL LTOcopySTG_01 CLASS1 POOLTYPE=COPY
DESCRIPTION="LTO copy storage pool 01" ACCESS=READWRITE COLLOCATE=NO
RECLAIM=100 MAXSCRATCH=0 REUSEDELAY=14 CRCDATA=NO DATAFORMAT=NATIVE
ANR2200I Storage pool LTOCOPYSTG_01 defined (device class CLASS1).
```

If we look at this new storage pool, using the **QUERY STGPOOL** command, we can see that it has been created as a copy storage pool, shown in Figure 8-3.

Example 8-3 Query storage pool

```
tsm: RADON_SERVER1>q stgpool ltocopystg_01 f=d

Storage Pool Name: LTOCOPYSTG_01
Storage Pool Type: Copy
Device Class Name: CLASS1
Estimated Capacity (MB): 1,907,340.0
Pct Util: 1.4
Pct Migr:
Pct Logical: 99.9
High Mig Pct:
Low Mig Pct:
Migration Delay:
Migration Continue:
Migration Processes:
Next Storage Pool:
Reclaim Storage Pool:
Maximum Size Threshold:
Access: Read/Write
Description: LTO copy storage pool 01
Overflow Location:
Cache Migrated Files?:
Collocate?: No
Reclamation Threshold: 100
Maximum Scratch Volumes Allowed: 10
Delay Period for Volume Reuse: 14 Day(s)
Migration in Progress?:
Amount Migrated (MB):
Elapsed Migration Time (seconds):
Reclamation in Progress?: No
Volume Being Migrated/Reclaimed:
Last Update by (administrator): IGOR
Last Update Date/Time: 07/26/2002 11:10:33
```

Storage Pool Data Format: Native
Copy Storage Pool(s):
Continue Copy on Error?:
CRC Data: No

DRM settings

Now we will configure DRM with some basic settings.

1. First we want to specify a directory where the DR recovery plans will be stored, using the **SET DRMPLANPREFIX** command. You can specify the plan prefix with a full directory path. Using the form shown in Example 8-4, our DR plans as generated by DRM will be stored in the directory C:\DRM\PLAN\ and each file will be prefixed by the string RADON. If we did not use this command, a default path would be used, which is the directory where the instance of the TSM server is running.

Example 8-4 Define DRM plan prefix

```
tsm: RADON_SERVER1>set drmplprefix c:\drm\plan\radon  
ANR6700I SET DRMPLANPREFIX command completed successfully.
```

2. Next we set the prefix for where the recovery instructions will be stored. The **DRM PREPARE** command will look for these files in this directory. This is done using the **SET DRMINSTRPREFIX** command. You can specify the instructions prefix with a full directory path. Using the form shown in Example 8-5, the recovery instruction files should be located in the directory C:\DRM\INSTRUCTION\ with each file prefixed by the string RADON. The prefix does not need to be specified — if we did not use this command, a default path would be used, which is the directory where the instance of the TSM server is running. The recovery instruction files are user generated and can contain any instructions related to the DR process. You can create those files using any plain text editor, and include the information which is relevant to your installation. Instruction files will be automatically included in the DRM plan. The standard names for the instruction plans (without prefix) are
 - RECOVERY.INSTRUCTIONS.GENERAL - for general information such as the system administrator and backup names and contact details, and passwords.
 - RECOVERY.INSTRUCTIONS.OFFSITE - for information about the offsite vault location and courier, including name, phone number, e-mail, fax, after-hours pager, and so on.
 - RECOVERY.INSTRUCTIONS.INSTALL - for TSM server installation instructions such as passwords, hardware/software requirements, fix levels, and so on.

- RECOVERY.INSTRUCTIONS.DATABASE - for TSM server database recovery instructions
- RECOVERY.INSTRUCTIONS.STGPOOL - for TSM server primary storage pool recovery instructions, for example, any priorities required.

Example 8-5 Set DRM instruction prefix

```
tsm: RADON_SERVER1>set drminstrprefix c:\drm\instruction\radon
ANR6700I SET DRMINSTRPREFIX command completed successfully.
```

3. Next, we will specify a character which will be appended to the end of the replacement volumes names in the recovery plan file. This is done using the **SET DRMPLANVPOSTFIX** command. Use of this special character means you can easily search and replace these names for the replacement primary storage pool volumes to your desired names before the recovery plan executes. In Example 8-6, we are using the default character of @.

Example 8-6 Set DRM replacement volume names postfix

```
tsm: RADON_SERVER1>set drmplnvpostfix @
ANR6700I SET DRMPLANVPOSTFIX command completed successfully.
```

4. Now, use the **SET DRMCHECKLABEL** command to specify whether TSM will read the labels of tape media when they are checked out using the **MOVE DRMEDIA** command (Example 8-7). The default value is YES.

Example 8-7 Command SET DRMCHECKLABEL

```
tsm: RADON_SERVER1>set drmchecklabel yes
ANR6700I SET DRMCHECKLABEL command completed successfully.
```

5. Now we need to indicate which primary storage pool(s) will be managed by DRM using the **SET DRMPRIMSTGPOOL** command. Those primary storage pools will be able to recovered by DRM after a disaster. In Example 8-8, we specify the storage pools DISK_STG_01 and LTO_STGP_01 to be managed by DRM.

Example 8-8 Specify primary storage pools to be managed by DRM

```
tsm: RADON_SERVER1>set drmprimstgpool DISK_STG_01,LTO_STGP_01
ANR6700I SET DRMPRIMSTGPOOL command completed successfully.
```

6. Use the **SET DRMCOPYSTGPOOL** command to indicate one or more copy storage pools to be managed by DRM. Those copy storage pools will be used to recover the primary storage pools after a disaster. The **MOVE MEDIA** and **QUERY DRMEDIA** commands will process volumes in the copy storage pools listed here by default (unless explicitly over-ridden with the COPYSTGPOOL parameter).

In Example 8-9, we specify that our copy storage pool, LTOCOPYSTG_01 (as defined in Example 8-2 on page 194) is to be managed by DRM.

Example 8-9 Specify primary storage pools to be managed by DRM

```
tsm: RADON_SERVER1>set drmcopystgpool LTOCOPYSTG_01
ANR6700I SET DRMCOPYSTGPOOL command completed successfully.
```

7. Next, use the **SET DRMCOURIERNAME** command to define the name of your courier company. If not set, this will use the default value of COURIER. Any string can be inserted here (see Example 8-10).

Example 8-10 Set the DRM couriername

```
tsm: RADON_SERVER1>set drmcouriername "Fast Leg Courier Service"
ANR6700I SET DRMCOURIERNAME command completed successfully.
```

8. Now, specify the number of days before expiration is used to expire a database backup series, using the **SET DRMDBBACKUPEXPIREDAYS** command. defines criteria for DB backup series expiration. The value applies to both a snapshot and a full plus incremental database backup series. The age of the last volume in the series must exceed the expiration value defined here to be eligible for expiration. The most recent backup series of either type is never deleted. In Example 8-11, we specify an expiration value of 14 days. To ensure that the database can be restored to an earlier level and database references to files in the storage pool are still valid, the number of days specified by this command and the number of days specified by the REUSEDELAY parameter in the copy storage pool definitions should be the same for the copy storage pools managed by DRM.

Example 8-11 Set the database backup series expiration value

```
tsm: RADON_SERVER1>set drmdbbackupexpiredays 14
ANR6700I SET DRMDBBACKUPEXPIREDAYS command completed successfully.
```

9. The **SET DRMFILEPROCESS** command (shown in Example 8-12 indicates if the **MOVE DRMEDIA** and **QUERY DRMEDIA** commands should process database backup volumes and copy storage pool volumes that have been defined with the FILE device class. The default value is NO.

Example 8-12 Set whether to process FILE device class media

```
tsm: RADON_SERVER1>set drmfileprocess no
ANR6700I SET DRMFILEPROCESS command completed successfully.
```

10. Next, specify the location where media will be stored while it is waiting to be sent to the offsite location, using the **SET DRMNOTMOUNTABLENAME** command. This location name is used by the **MOVE DRMEDIA** command to set the location

of volumes that are transitioning to the NOTMOUNTABLE state. The default value is NOTMOUNTABLE.

Example 8-13 Set NOTMOUNTABLE location name

```
tsm: RADON_SERVER1>set drmnotmountable "Waiting for Courier"
ANR6700I SET DRMNOTMOUNTABLENAME command completed successfully.
```

11. The **SET DRMRPFEXPIREDDAYS** command sets the number of days after creation that a recovery plan file which has been stored on a target server (using server-to-server communication) will be retained. This command and expiration processing only applies to recovery plan files that are created with the DEVCLASS parameter on the **PREPARE** command (that is, virtual volumes of type RPFIL and RPSNAPSHOT). The most recent files are never deleted. Example 8-14 shows changing this value from the default of 60 days back to 30 days.

Example 8-14 Set expiration period for recovery plan files

```
tsm: RADON_SERVER1>set drmrpfexpiredays 30
ANR6700I SET DRMRPFEXPIREDDAYS command completed successfully.
```

12. You can identify the vault name with the **SET DRMVAULTNAME** command, as shown in Example 8-15. You can specify any string or use the default value of VAULT.

Example 8-15 Set DRM vault name

```
tsm: RADON_SERVER1>set drmvaultname "IronVault, Fort Knox, Kentucky"
ANR6700I SET DRMVAULTNAME command completed successfully.
```

Verifying the settings

You can display and check settings of all DRM parameters with the command **QUERY DRMSTATUS**, as shown in Example 8-16.

Example 8-16 Command QUERY DRMSTATUS

```
tsm: RADON_SERVER1>q drmstatus

Recovery Plan Prefix: C:\DRM\PLAN\RADON
Plan Instructions Prefix: C:\DRM\INSTRUCTION\RADON
Replacement Volume Postfix: @
Primary Storage Pools: DISK_STG_01 LTO_STGP_01
Copy Storage Pools: LTOCOPYSTG_01
Not Mountable Location Name: Waiting for Courier
Courier Name: Fast Leg Courier Service
Vault Site Name: IronVault, Fort Knox, Kentucky
DB Backup Series Expiration Days: 14 Day(s)
```

Recovery Plan File Expiration Days: 30 Day(s)
Check Label?: Yes
Process FILE Device Type?: No
Command File Name: C:\DRM\RADON_EXEC.CMD

8.7.2 Daily operations

After setup, DRM will now require special operations which will occur after the daily TSM client backups. These include:

- ▶ Backup of the primary storage pools to the copy storage pool
- ▶ Backup the TSM database
- ▶ Generate the recovery plan
- ▶ Move DR media and DRP to the offsite location

We will now describe these steps in detail. Throughout this section, we refer to the possible states which a piece of DR media could be in. The DR media consists of the volumes used for storage pool and database backup. The states, and their normal life cycle is shown in Figure 8-10. Media changes state by use of the `MOVE DRMEDIA` command.

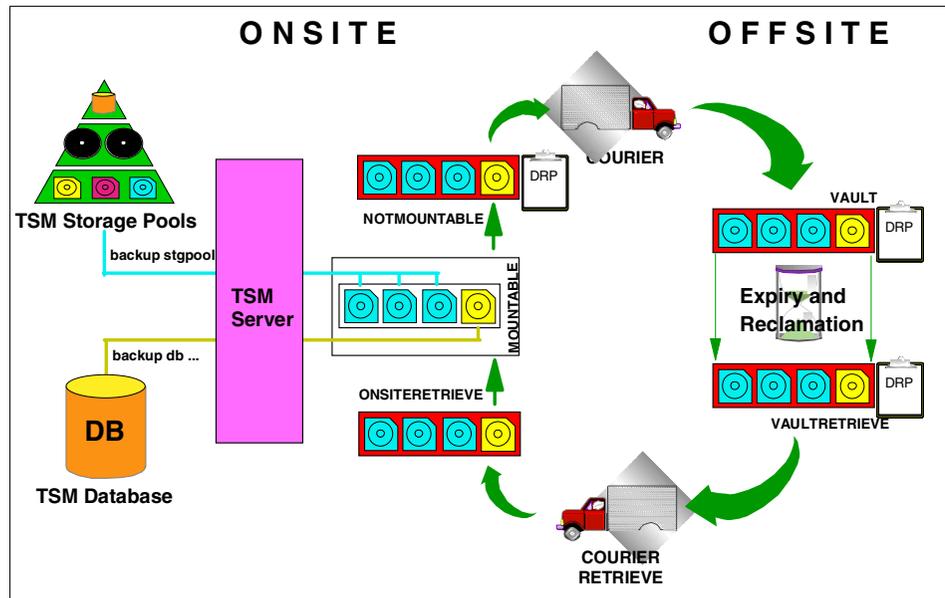


Figure 8-10 DRM media states and lifecycle

When the copy tapes are initially made, they are in state `MOUNTABLE`, meaning they are available to the TSM server. After they are ejected from the library and are ready to be picked up for transport by the Courier, they transition to the

NONMOUNTABLE state. Once picked up, they enter the COURIER state. After being received and formally acknowledged at the vault site, they move to the VAULT state. They remain in this state until the data on them expires or they are eligible for reclamation. At this stage, the now-empty tapes are in VAULTRETRIEVE state while waiting pickup for returning to the production site. While in transit, they are in COURIERRETRIEVE state. On arrival back at the site, they change to ONSITERETRIEVE until they are re-loaded to the tape library. At this time, they return to MOUNTABLE for re-use as scratch volumes. Note that according to local policy, not all states are required. You can skip states, so long as it is clearly understood which states correspond to media which is safe in the vault and can be used for a disaster recovery.

Backup primary storage pools to copy storage pool

In our setup, the TSM clients backed up to the primary storage pool, DISK_STG_01, which migrates to the pool LTO_STGP_01. We need to backup both of storage pools to our already defined copy storage pool, LTOCOPYSTG_01. We use the **BACKUP STGPPOOL** command for this, as shown in Example 8-17 and Example 8-18.

Example 8-17 Backup of primary disk storage pool

```
tsm: RADON_SERVER1>backup stgpool DISK_STG_01 LTOCOPYSTG_01 MAXPROCESS=1
PREVIEW=NO WAIT=NO
ANS8003I Process number 22 started.
```

Activity log record

```
07/26/2002 17:07:08 ANR1210I Backup of primary storage pool DISK_STG_01 to
                    copy storage pool LTOCOPYSTG_01 started as process 22.
07/26/2002 17:07:32 ANR0407I Session 585 started for administrator IGOR
                    (WinNT) (Tcp/Ip 9.1.39.7(1907)).
007/26/2002 17:08:34 ANR1212I Backup process 22 ended for storage pool
                    DISK_STG_01.
07/26/2002 17:08:34 ANR0986I Process 22 for BACKUP STORAGE POOL running in the
                    BACKGROUND processed 45924 items for a total of
                    790,208,512 bytes with a completion state of SUCCESS at
                    17:08:34.
07/26/2002 17:08:34 ANR1214I Backup of primary storage pool DISK_STG_01 to
                    copy storage pool LTOCOPYSTG_01 has ended. Files Backed
                    Up: 45924, Bytes Backed Up: 790208512, Unreadable Files:
                    0, Unreadable Bytes: 0.
```

Example 8-18 Backup of primary tape storage pool

```
tsm: RADON_SERVER1>backup stgpool LTO_STGP_01 LTOCOPYSTG_01 MAXPROCESS=1
PREVIEW=NO WAIT=NO
ANS8003I Process number 23 started.
```

Activity record log

```
07/26/2002 16:27:31 ANR2017I Administrator IGOR issued command: BACKUP STGPPOOL
LTO_STGP_01 LTOCOPYSTG_01 MAXPROCESS=1 PREVIEW=NO WAIT=NO
07/26/2002 16:27:31 ANR0984I Process 21 for BACKUP STORAGE POOL started in the
BACKGROUND at 16:27:31.
07/26/2002 16:27:31 ANR2110I BACKUP STGPPOOL started as process 21.
07/26/2002 16:27:31 ANR1210I Backup of primary storage pool LTO_STGP_01 to
copy storage pool LTOCOPYSTG_01 started as process 21.
07/26/2002 16:27:31 ANR0609I BACKUP STGPPOOL started as process 21.
07/26/2002 16:36:55 ANR1212I Backup process 21 ended for storage pool
LTO_STGP_01.
07/26/2002 16:36:55 ANR0986I Process 21 for BACKUP STORAGE POOL running in the
BACKGROUND processed 5346 items for a total of
1,119,408,390 bytes with a completion state of SUCCESS at
16:36:55.
07/26/2002 16:36:55 ANR1214I Backup of primary storage pool LTO_STGP_01 to
copy storage pool LTOCOPYSTG_01 has ended. Files Backed
Up: 5346, Bytes Backed Up: 1119408390, Unreadable Files:
0, Unreadable Bytes: 0.
```

Backup of TSM database

After successful backup of all primary storage pools, we are ready to backup the TSM database, using the **BACKUP DB** command as in Example 8-19. The database backup is written to the tape volume ABA927L1.

Example 8-19 TSM database backup

```
tsm: RADON_SERVER1>backup db DEVCLASS=CLASS1 TYPE=FULL SCRATCH=YES
ANR2280I Full database backup started as process 24.
ANS8003I Process number 24 started.
```

Activity log record

```
07/26/2002 17:23:38 ANR2017I Administrator IGOR issued command: BACKUP DB
DEVCLASS=CLASS1 TYPE=FULL SCRATCH=YES
07/26/2002 17:23:38 ANR0984I Process 24 for DATABASE BACKUP started in the
BACKGROUND at 17:23:38.
07/26/2002 17:23:38 ANR2280I Full database backup started as process 24.
07/26/2002 17:26:42 ANR8337I LTO volume ABA927L1 mounted in drive MT01
(mt1.2.0.4).
07/26/2002 17:27:04 ANR1360I Output volume ABA927L1 opened (sequence number
1).
07/26/2002 17:27:20 ANR4554I Backed up 15936 of 31905 database pages.
07/26/2002 17:27:23 ANR4554I Backed up 31872 of 31905 database pages.
07/26/2002 17:27:25 ANR1361I Output volume ABA927L1 closed.
07/26/2002 17:27:25 ANR4550I Full database backup (process 24) complete, 31905
pages copied.
07/26/2002 17:27:25 ANR0985I Process 24 for DATABASE BACKUP running in the
BACKGROUND completed with completion state SUCCESS at
```

QUERY DRMEDIA and MOVE DRMEDIA

When the TSM database backup finished successfully, we can check the status of our DR media, using the **QUERY DRMEDIA** command. This command will list the volumes which should be moved offsite, and show the volumes just used for the database and storage pool backup. Example 8-20 shows the output of the command. Since we know that volume ABA927L1 was used for the database backup (as shown in Example 8-19), we can assume that the other volume shown in the output, ABA926L1, has been used for the copy storage pool. Specifying the state as MOUNTABLE ensures that only new DR media will be displayed (as opposed to DR media which has already been take offsite, or is in transit, for example).

Example 8-20 Query the DR media

```
tsm: RADON_SERVER1>query drmedia * wherestate=mountable
```

Volume Name	State	Last Update Date/Time	Automated LibName
-----	-----	-----	-----
ABA927L1	Mountable	07/24/2002 19:27:29	LB6.0.0.3
ABA926L1	Mountable	07/26/2002 17:23:38	LB6.0.0.3

Now we want to move the DR media to the VAULT location. First, we should check, if any of the required volumes are still mounted. We can do this using the **QUERY MOUNT** command. If yes, we can dismount them with the **DISMOUNT VOLUME** command as shown in Example 8-21.

Example 8-21 Check and eject DR media

```
tsm: RADON_SERVER1>query mount
ANR8329I LTO volume ABA927L1 is mounted R/W in drive MT01 (mt1.2.0.4), status:
IDLE.
```

```
tsm: RADON_SERVER1>dismount volume ABA927L1
ANR8499I Command accepted.
```

```
Activity log record
07/26/2002 17:50:01 ANR8336I Verifying label of LTO volume ABA927L1 in drive
MT01 (mt1.2.0.4).
07/26/2002 17:51:32 ANR8468I LTO volume ABA927L1 dismounted from drive MT01
(mt1.2.0.4) in library LB6.0.0.3.
```

Send DR media off-site

Now we can mark initiate sending the DR media we have just identified to the offsite location, using the command **MOVE DRMEDIA**. By default, it will move the database backup volume plus volumes from the copy storage pool volumes specified by the **SET DRMCOPYSTGPOOL** command (as described in 6 on page 196). The command output is shown in Example 8-22 and shows how using the **REMOVE=Y** parameter automatically ejects media from the library.

Example 8-22 Eject DR media from the library

```
tsm: RADON_SERVER1>move drmedia * WHERESTATE=MOUNTABLE REMOVE=Y
TOSTATE=VAULT
Session established with server RADON_SERVER1: Windows
  Server Version 5, Release 1, Level 1.0
  Server date/time: 07/26/2002 18:09:59  Last access: 07/26/2002 17:58:41

Activity log record
07/26/2002 18:09:59 ANR2017I Administrator IGOR issued command: MOVE DRMEDIA *
                    WHERESTATE=MOUNTABLE SOURCE=DBBACKUP REMOVE=YES
                    TOSTATE=VAULT
07/26/2002 18:09:59 ANR0984I Process 25 for MOVE DRMEDIA started in the
                    BACKGROUND at 18:09:59.
07/26/2002 18:09:59 ANR0609I MOVE DRMEDIA started as process 25.
07/26/2002 18:09:59 ANR0610I MOVE DRMEDIA started by IGOR as process 25.
07/26/2002 18:09:59 ANR6696I MOVE DRMEDIA: CHECKOUT LIBVOLUME for volume
                    ABA926L1 in library LB6.0.0.3 starting.
07/26/2002 18:10:41 ANR8336I Verifying label of LTO volume ABA926L1 in drive
                    MT01 (mt1.2.0.4).
07/26/2002 18:11:15 ANR8322I 001: Remove LTO volume ABA926L1 from entry/exit
                    port of library LB6.0.0.3; issue 'REPLY' along with the
                    request ID when ready.
07/26/2002 18:12:19 ANR2017I Administrator IGOR issued command: REPLY 001
07/26/2002 18:12:19 ANR8499I Command accepted.
07/26/2002 18:12:19 ANR6697I MOVE DRMEDIA: CHECKOUT LIBVOLUME for volume
                    ABA926L1 in library LB6.0.0.3 completed successfully.
07/26/2002 18:12:19 ANR6683I MOVE DRMEDIA: Volume ABA926L1 was moved from
                    MOUNTABLE state to NOTMOUNTABLE.
07/26/2002 18:12:19 ANR6696I MOVE DRMEDIA: CHECKOUT LIBVOLUME for volume
                    ABA927L1 in library LB6.0.0.3 starting.
07/26/2002 18:13:05 ANR8336I Verifying label of LTO volume ABA927L1 in drive
                    MT02 (mt1.4.0.4).
07/26/2002 18:13:42 ANR8322I 002: Remove LTO volume ABA927L1 from entry/exit
                    port of library LB6.0.0.3; issue 'REPLY' along with the
                    request ID when ready.
07/26/2002 18:14:33 ANR2017I Administrator IGOR issued command: REPLY 002
07/26/2002 18:14:33 ANR8499I Command accepted.
07/26/2002 18:14:33 ANR6697I MOVE DRMEDIA: CHECKOUT LIBVOLUME for volume
                    ABA927L1 in library LB6.0.0.3 completed successfully.
```

```
07/26/2002 18:14:33 ANR6683I MOVE DRMEDIA: Volume ABA927L1 was moved from
MOUNTABLE state to NOTMOUNTABLE.
07/26/2002 18:14:33 ANR6682I MOVE DRMEDIA command ended: 2 volumes processed.
07/26/2002 18:14:33 ANR0611I MOVE DRMEDIA started by IGOR as process 25 has
ended.
07/26/2002 18:14:33 ANR0987I Process 25 for MOVE DRMEDIA running in the
BACKGROUND processed 2 items with a completion state of
SUCCESS at 18:14:33.
```

This has processed both volumes and marked them as state NOTMOUNTABLE since they are no longer available in the tape library. Their location has been changed to the value specified in **SET DRMNOTMOUNTABLENAME** (as described in step 10 on page 197). At this stage, the courier would arrive to collect the volumes. Once the volumes have been signed over to the courier, we need to indicate this state change to DRM, using again the **MOVE DRMEDIA** command, shown in Example 8-23. This command will change the state of the DR media volumes from NOTMOUNTABLE to COURIER, and their location to the value specified in **SET DRMCOURIERNAME** (as described in Step 7 on page 197), indicating they are in transit.

Example 8-23 Sign over the DR media to the courier

```
tsm: RADON_SERVER1>move drmedia * WHERESTATE=NOTMOUNTABLE
```

Finally, when the vault confirms that the DR media have safely arrived, use the **MOVE DRMEDIA** command one more time, as shown in Example 8-24. This sets the state of the DR media volumes to VAULT, and the location to the value specified in **SET DRMVAULTNAME** (as described in Step 12 on page 198).

Example 8-24 Confirm DR media arrives at vault

```
tsm: RADON_SERVER1>move drmedia * WHERESTATE=COURIER
```

We successfully finished backup of primary storage pools, and database backup. The media is now safely offsite. We can check the media status once again with the command **QUERY DRMEDIA**, as shown in Example 8-25.

Example 8-25 Display offsite media status

```
tsm: RADON_SERVER1>query drmedia f=d
      Volume Name: ABA926L1
      State: Vault
      Last Update Date/Time: 07/26/2002 18:12:19
      Location: IronVault, Fort Knox, Kentucky
      Volume Type: CopyStgPool
      Copy Storage Pool Name: LTOCOPYSTG_01
      Automated LibName:
```

Volume Name: ABA927L1
State: Vault
Last Update Date/Time: 07/26/2002 18:14:33
Location: IronVault, Fort Knox, Kentucky
Volume Type: DBBackup
Copy Storage Pool Name:
Automated LibName:

Figure 8-11 shows the process we have used to backup the storage pools and database, and send the media offsite.

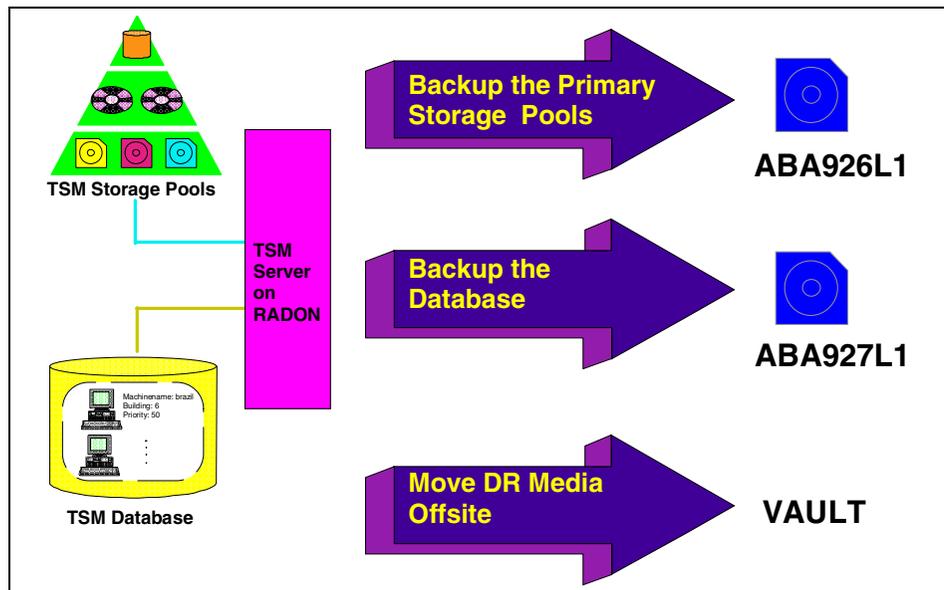


Figure 8-11 Daily operations - primary pools backup and TSM database backup

Generate the recovery plan

Now we are ready to generate the Disaster Recovery Plan as shown in Figure 8-12.

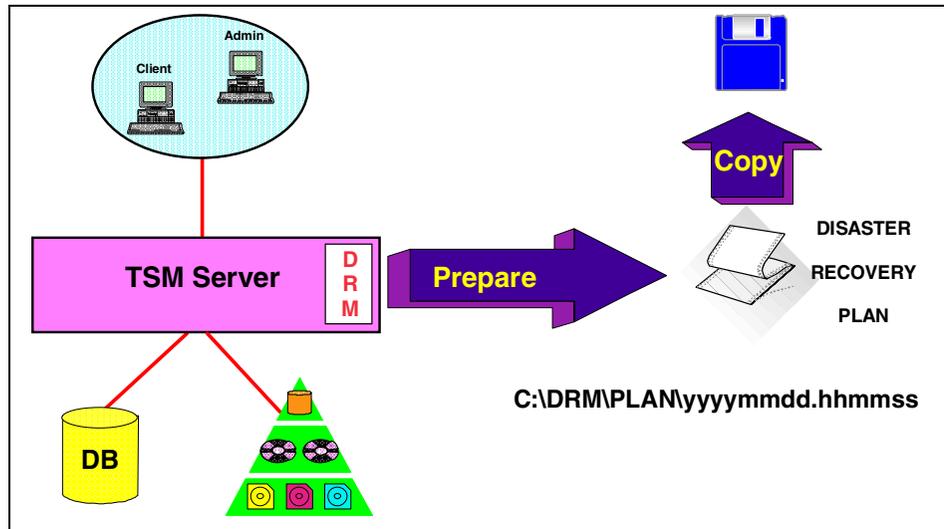


Figure 8-12 Disaster Recovery Plan generation

We generate the recovery plan using the **PREPARE** command, as in Example 8-26.

Example 8-26 Prepare the recovery plan

```
tsm: RADON_SERVER1>prepare
ANS8003I Process number 26 started.
```

Activity log record

```
07/26/2002 18:30:58 ANR2017I Administrator IGOR issued command: PREPARE
07/26/2002 18:30:58 ANR0984I Process 26 for PREPARE started in the BACKGROUND
                        at 18:30:58.
07/26/2002 18:30:58 ANR9678W C:\Program Files\tivoli\tsm\Server\dsmserve used
                        for server executable. A server is currently running as a
                        service.
07/26/2002 18:30:59 ANR6900I PREPARE: The recovery plan file
                        C:\DRM\PLAN\RADON.20020726.183058 was created.
07/26/2002 18:30:59 ANR0985I Process 26 for PREPARE running in the BACKGROUND
                        completed with completion state SUCCESS at 18:30:59.
```

The plan is now stored in a time-stamped file in the local directory with prefix as defined in **SET DRMPPLANPREFIX**, as shown in Step 1 on page 195. The file created here is called C:\DRM\PLAN\RADON.20020726.183058. A complete listing of the recovery plan output is provided in Appendix C, “DRM plan output” on page 353. We recommend for safety, that you create multiple copies of the recovery plan, stored in different locations. You can create a remote copy of the DRP by specifying a DEVCLASS on the **PREPARE** command. This DEVCLASS can only be of type SERVER, and is used to store the DRP on a target server

using server-to-server communication. Another way to create a remote copy is to over-ride the default directory location (set by **SET DRMPLANPREFIX**) by specifying an alternative (typically on a network drive) location with the **PLANPREFIX** parameter. The DR plan(s) created can also be ftp'd to another site, and even printed for emergency retrieval. Use whatever redundancy is necessary to maintain confidence that a copy of the plan will be accessible in the event of a disaster. A copy of the recovery plan should also be stored on removable media and moved to the vault location.

Returning expired volumes

Part of daily operations is also moving tapes back on-site which have expired and can therefore be returned to scratch. Volumes can be expired if they are **EMPTY** copy storage pool volumes or **EXPIRED** volumes from a database backup series. DRM automatically changes the state of such volumes to **VAULTRETRIEVE**. You can generate a list of DR media to be returned using the command:

```
QUERY DRMedia * WHERESTATE=vaultretrieve
```

Send this list to the vault for return to the primary location. When you are notified that the volumes have been given to the courier, this state change can be reflected in DRM by using:

```
MOVE DRMedia * WHERESTATE=vaultretrieve
```

This will update the media state from **VAULTRETRIEVE** to **COURIERRETRIEVE**. When the courier has delivered the volumes back to the primary location, change their statue to **ONSITERETRIEVE** with the command:

```
MOVE DRMedia * WHERESTATE=courierretrieve
```

Returned media can then be inserted back into the library for reuse with the **CHECKIN LIBVOL** command. Note that you can also use options on the **MOVE DRMEDIA** command for DR media in the **COURIERRETRIEVE** state to automatically generate a macro of **CHECKIN LIBVOL** commands. Refer to the TSM Administrator's Reference for more details.

8.7.3 Server restore setup

If a disaster happens, you can easily prepare to use the DRM generated plans and offsite stored media to restore the TSM server. Here is a summary of the process:

- ▶ Obtain the latest Disaster Recovery Plan.
- ▶ Break out the file to view, update, print, or run as macros or scripts.

- ▶ Obtain the backup volumes from the vault.
- ▶ Locate a suitable replacement machine.
- ▶ Install — restore operating system and TSM to replacement machine.
- ▶ Review the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE and RECOVERY.SCRIPT.NORMAL.MODE files, because they are important for restoring the server to a point where clients can be recovered.

Obtain the latest DR Plan

This should be retrieved from the courier or from any other location, according to your company's policy, where it has been stored. Double-check that it is the latest version of the plan by referring to the line near the top of the file.

Created by DRM PREPARE on 07/26/2002 18:30:58

The plan should be copied back to the same directory as it was stored on the original TSM server when it was created.

Break out the DR Plan

The DRP is created as a single file, which is organized in stanzas. For execution after a disaster, it needs to be broken up into its components, as shown in Figure 8-13.

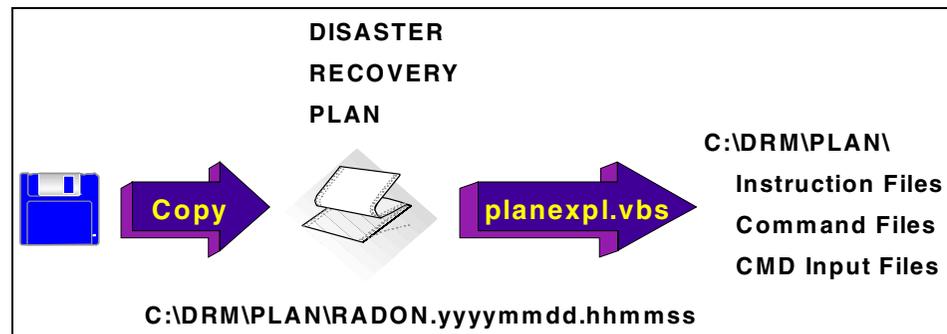


Figure 8-13 Generating instruction, command, and cmd input files from DR plan

You can use a text editor to manually divide the recovery plan into its components, or use sample scripts shipped with TSM. For Windows, a sample VBScript is in `planexpl.vbs`, shipped with DRM. For UNIX, a sample awk script is in `planexpl.awk.smp`. You should keep a copy of these scripts offsite along with the recovery plan. We recommend for you to be familiar with executing the scripts, as the plan will be large and doing a manual breakout will be time-consuming and prone to errors.

To use the sample script to break out a recovery plan on Windows, use the syntax shown in Example 8-27.

Example 8-27 Break out DRP on Windows platforms

```
cscript planexpl.vbs recoveryplanfilename
```

For AIX, the syntax is shown in Example 8-28.

Example 8-28 Break out DRP on AIX platform

```
awk -f planexpl.awk recoveryplanfilename
```

For Sun Solaris, the syntax is shown in Example 8-29.

Example 8-29 Break out DRP on Sun Solaris platform

```
nawk -f planexpl.awk recoveryplanfilename
```

The TSM server restore process is overviewed in Figure 8-14.

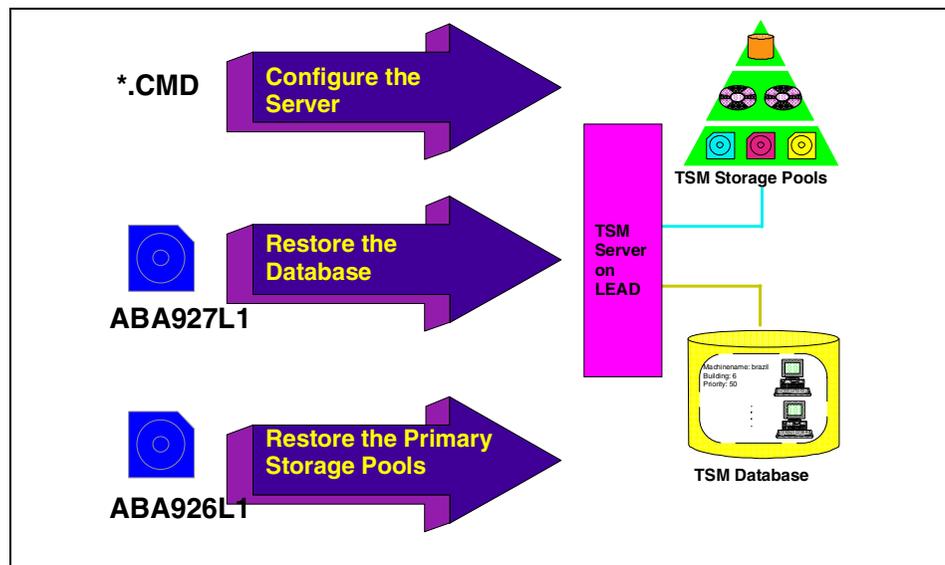


Figure 8-14 TSM server restore

Example 8-30 shows the command output from breaking out the DRP on the replacement server. You can see that many smaller files are created.

Example 8-30 Generating instruction, command and cmd files

```
C:\DRM\plan>cscript planexpl.vbs C:\DRM\plan\RADON.20020726.183058
Microsoft (R) Windows Script Host Version 5.6
Copyright (C) Microsoft Corporation 1996-2001. All rights reserved.
```

```
Planfile: C:\DRM\plan\RADON.20020726.183058
set planprefix to C:\DRM\PLAN\RADON.
Creating file C:\DRM\PLAN\RADON.SERVER.REQUIREMENTS
Creating file C:\DRM\PLAN\RADON.RECOVERY.INSTRUCTIONS.GENERAL
Creating file C:\DRM\PLAN\RADON.RECOVERY.INSTRUCTIONS.OFFSITE
Creating file C:\DRM\PLAN\RADON.RECOVERY.INSTRUCTIONS.INSTALL
Creating file C:\DRM\PLAN\RADON.RECOVERY.INSTRUCTIONS.DATABASE
Creating file C:\DRM\PLAN\RADON.RECOVERY.INSTRUCTIONS.STGPOOL
Creating file C:\DRM\PLAN\RADON.RECOVERY.VOLUMES.REQUIRED
Creating file C:\DRM\PLAN\RADON.RECOVERY.DEVICES.REQUIRED
Creating file C:\DRM\PLAN\RADON.RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE.CMD
Creating file C:\DRM\PLAN\RADON.RECOVERY.SCRIPT.NORMAL.MODE.CMD
Creating file C:\DRM\PLAN\RADON.LOG.VOLUMES
Creating file C:\DRM\PLAN\RADON.DB.VOLUMES
Creating file C:\DRM\PLAN\RADON.LOGANDDB.VOLUMES.INSTALL.CMD
Creating file C:\DRM\PLAN\RADON.LICENSE.REGISTRATION.MAC
Creating file C:\DRM\PLAN\RADON.COPYSTGPOOL.VOLUMES.AVAILABLE.MAC
Creating file C:\DRM\PLAN\RADON.COPYSTGPOOL.VOLUMES.DESTROYED.MAC
Creating file C:\DRM\PLAN\RADON.PRIMARY.VOLUMES.DESTROYED.MAC
Creating file C:\DRM\PLAN\RADON.PRIMARY.VOLUMES.REPLACEMENT.CREATE.CMD
Creating file C:\DRM\PLAN\RADON.PRIMARY.VOLUMES.REPLACEMENT.MAC
Creating file C:\DRM\PLAN\RADON.STGPOOLS.RESTORE.MAC
Creating file C:\DRM\PLAN\RADON.VOLUME.HISTORY.FILE
Creating file C:\DRM\PLAN\RADON.DEVICE.CONFIGURATION.FILE
Creating file C:\DRM\PLAN\RADON.DSMSERV.OPT.FILE
Creating file C:\DRM\PLAN\RADON.LICENSE.INFORMATION
Creating file C:\DRM\PLAN\RADON.MACHINE.GENERAL.INFORMATION
Creating file C:\DRM\PLAN\RADON.MACHINE.RECOVERY.INSTRUCTIONS
Creating file C:\DRM\PLAN\RADON.MACHINE.CHARACTERISTICS
Creating file C:\DRM\PLAN\RADON.MACHINE.RECOVERY.MEDIA.REQUIRED
```

If executed, the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE batch program will delete the following log and db volumes if they exist and then reallocate them. During a normal disaster recovery scenario this is not a problem since you are going to restore data to them from the db backup.

```
C:\TSMDB\TSMLOG01.DB
C:\TSMDB\TSMDB01.DB
C:\TSMDB\TSMDB02.DB
```

Find a replacement server and storage

Local information on hardware required is specified in the RECOVERY.INSTRUCTIONS.INSTALL stanza. The RECOVERY.DEVICES.REQUIRED stanza specifies the device type that is needed to read the backups. The SERVER.REQUIREMENTS stanza specifies the disk space required.

Obtain the recovery volumes

A list of the offsite volumes required for recovery is in the stanza RECOVERY.VOLUMES.REQUIRED. This is shown in Example 8-31. Send this list to the vault for retrieval as fast as possible.

Example 8-31 Volumes required for TSM server recovery

```
begin RECOVERY.VOLUMES.REQUIRED

Volumes required for data base restore

Location = IronVault, Fort Knox, Kentucky
Device Class = CLASS1
Volume Name =
ABA927L1

Volumes required for storage pool restore

Location = IronVault, Fort Knox, Kentucky
Copy Storage Pool = LTOCOPYSTG_01
Device Class = CLASS1
Volume Name =
ABA926L1

end RECOVERY.VOLUMES.REQUIRED
```

We can see that volume ABA926L1 has a copy storage pool entry, whereas volume ABA927L1 does not. This is because ABA927L1 was used for a database backup. We can confirm this by looking at the volume history section of the recovery plan (VOLUME.HISTORY.FILE stanza). We will need the volume ABA927L1 initially for restoring the TSM database.

Restore the operating system and install TSM

Restore the operating system and install TSM server software on the replacement server. If your tape library uses a vendor-supplied driver (as opposed to the inbuilt ADMSCSI driver), you have to install it before starting TSM server recovery. The media names required for recovery and their locations

are specified in the RECOVERY.INSTRUCTIONS.INSTALL stanza and the MACHINE.RECOVERY.MEDIA.REQUIRED stanza. Ensure that the environment is the same as when the Disaster Recovery Plan file was created. The environment includes:

- ▶ The directory structure and location of the TSM server executables, enrollment certificates, administrative command line client, and disk formatting utility.
- ▶ The directory structure for TSM server instance-specific files and the database, log, and storage pool volumes.
- ▶ The directory structure and the files created when the Disaster Recovery Plan file was split into multiple files.

Review the TSM macros

Review the TSM macros contained in the recovery plan. If at the time of the disaster, the courier had not picked up the previous night's database and storage pool incremental backup volumes but they were not destroyed, and hence will be available for recovery, remove the entry for the storage pool backup volumes from the COPYSTGPOOL.VOLUMES.DESTROYED file.

If some required storage pool backup volumes could not be retrieved from the vault, remove the volume entries from the COPYSTGPOOL.VOLUMES.AVAILABLE file.

If all primary volumes were destroyed, no changes are required to the PRIMARY.VOLUMES.DESTROYED script and TSM macro files.

In the PRIMARY.VOLUMES.REPLACEMENT.CREATE.CMD we have to specify replacement volumes for the primary disk storage volumes. In, Example 8-32, the primary storage pool volume was originally located on H:\TSMDATA\STG_POOL_01.DSM. The new location for the primary disk storage pool volume is C:\TSMDATA\STG_POOL_01.DSM. Create **dsmfmt** entries for all storage pool volumes.

Example 8-32 File PRIMARY.VOLUMES.REPLACEMENT.CREATE.CMD

```
@echo off

rem Purpose: Create replacement volumes for primary storage pools that
rem use device class DISK.
rem Recovery administrator: Edit this section for your replacement
rem volume names. New name must be unique, i.e. different from any
rem original or other new name.

rem Set the TSM management console directory.
pushd "C:\Program Files\tivoli\tsm\console\"
```

```
echo Replace H:\TSMDATA\STG_POOL_01.DSM DISK 2,048.0M in DISK_STG_01
dsmfmt -data "C:\TSMDATA\STG_POOL_01.DSM" 2048
```

```
rem Restore the previous working directory.
popd
```

In the file PRIMARY.VOLUMES.REPLACEMENT.MAC we need to define replacement volumes for destroyed primary storage pool volumes.

Example 8-33 File PRIMARY.VOLUMES.REPLACEMENT.MAC

```
/* Purpose: Define replacement primary storage pool volumes for either: */
/* 1. Original volume in a storage pool whose device class was DISK. */
/* 2. Original volume in a storage pool with MAXSCRATCH=0. */
/* 3. Original volume in a storage pool and volume scratch=no. */
/* Recovery administrator: Edit this section for your replacement */
/* volume names. New name must be unique, i.e. different from any */
/* original or other new name. */

/* Replace H:\TSMDATA\STG_POOL_01.DSM DISK 2,048.0M in DISK_STG_01 */
def vol DISK_STG_01 "C:\TSMDATA\STG_POOL_01.DSM" acc=READW

/* Replace ABA920L1 CLASS1 190,734.0M in LTO_STGP_01 */
def vol LTO_STGP_01 "ABA990L1" acc=READW
```

Review the device configuration

Review the device configuration file to ensure that the hardware configuration at the recovery site is the same or equivalent to the original site. Any differences (for example, device special file names) must be updated in the device configuration file. Table 8-5 shows our before and after picture for device configuration. We can see that different SCSI addresses were assigned to the medium changer and tape drives on the replacement server, compared to the original server.

Table 8-5 Review of the TSM device configuration

TSM server on RADON	TSM server on LEAD
IBM 3583 Scalable LTO tape library	IBM 3583 Scalable LTO tape library
Medium changer address lb1.6.0.4	Medium changer address lb0.6.0.4
Tape drive mt1.2.0.4	Tape drive mt0.2.0.4

TSM server on RADON	TSM server on LEAD
Tape drive mt1.4.0.4	Tape drive mt0.4.0.4

The IBM 3583 is an automated library and we will have to manually place the database backup volumes into the library (since there is no TSM server to check them in) and update the configuration information to identify the element within the library where the volumes are placed. This allows the server to locate the required database backup volumes. In Example 8-34 we added a line to the device configuration file (DEVICE.CONFIGURATION.FILE stanza in the DR Plan) with the location of tape volume ABA927L1, and the actual element address 0016. This element address corresponds to I/O station slot 1 in the library. For information on the element addresses for your particular devices, consult your tape library vendor documentation and the Tivoli Web site on device support:

http://www.tivoli.com/support/storage_mgr/requirements.html

We changed the DEFINE statements for the library and paths to reflect the actual new device special files for the library and tape drives.

Example 8-34 DEVICE.CONFIGURATION.FILE

```

/* Device Configuration */
DEFINE DEVCLASS CLASS1 DEVTYPE=LTO FORMAT=DRIVE MOUNTLIMIT=DRIVES MOUNTWAIT=60
MOUNTRETENTION=60 PREFIX=ADSM LIBRARY=LB6.0.0.3
DEFINE DEVCLASS FILEONFAST DEVTYPE=FILE FORMAT=DRIVE MAXCAPACITY=512000K
MOUNTLIMIT=1 DIRECTORY="H:\TSMDBBACKUP\" SHARED=NO
SET SERVERNAME RADON_SERVER1
SET SERVERPASSWORD 18c0be89
DEFINE LIBRARY LB6.0.0.3 LIBTYPE=SCSI SHARED=NO
DEFINE DRIVE LB6.0.0.3 MT01 ELEMENT=256 ONLINE=Yes
DEFINE DRIVE LB6.0.0.3 MT02 ELEMENT=257 ONLINE=Yes
/* LIBRARYINVENTORY SCSI LB6.0.0.3 ABA920L1 4096 101*/
/* LIBRARYINVENTORY SCSI LB6.0.0.3 ABA922L1 4101 101*/
/* LIBRARYINVENTORY SCSI LB6.0.0.3 ABA928L1 4097 101*/
/* LIBRARYINVENTORY SCSI LB6.0.0.3 ABA929L1 4100 101*/
/* LIBRARYINVENTORY SCSI LB6.0.0.3 ABA990L1 4102 101*/
/* LIBRARYINVENTORY SCSI LB6.0.0.3 ABA927L1 0016 101*
DEFINE PATH RADON_SERVER1 LB6.0.0.3 SRCTYPE=SERVER DESTTYPE=LIBRARY
DEVICE=1b0.6.0.4 ONLINE=YES
DEFINE PATH RADON_SERVER1 MT01 SRCTYPE=SERVER DESTTYPE=DRIVE LIBRARY=LB6.0.0.3
DEVICE=mt0.2.0.4 ONLINE=YES
DEFINE PATH RADON_SERVER1 MT02 SRCTYPE=SERVER DESTTYPE=DRIVE LIBRARY=LB6.0.0.3
DEVICE=mt0.4.0.4 ONLINE=YES

end DEVICE.CONFIGURATION.FILE

```

Start of restore TSM server scripts

To restore the TSM server to a point where clients can be recovered from copy storage pools, invoke the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script file. Enter the script file name at the command prompt and follow with administrator name and password as parameters. The output is shown in Example 8-35. The script formats volumes for the database and recovery log, then restores the database using the volume in the tape library.

Example 8-35 Invoke RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE

```
C:\DRM\plan>RADON.RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE.CMD admin admin
    1 file(s) copied.
    1 file(s) copied.
    1 file(s) copied.
ANR0900I Processing options file C:\PROGRA~1\tivoli\tsm\server1\dsmserve.opt.
ANR7800I DSMSEV generated at 09:45:38 on May 30 2002.

Tivoli Storage Manager for Windows
Version 5, Release 1, Level 1.0

Licensed Materials - Property of IBM

5698-ISE (C) Copyright IBM Corporation 1999,2002. 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 C:\TSMDB\TSMLOG01.DB: 135266304 bytes.
Allocated space for C:\TSMDB\TSMDB01.DB: 1074790400 bytes.
Allocated space for C:\TSMDB\TSMDB02.DB: 1074790400 bytes.
ANR0300I Recovery log format started; assigned capacity 128 megabytes.
ANR0301I Recovery log format in progress; 4 megabytes of 128.
ANR0301I Recovery log format in progress; 8 megabytes of 128.
...
...
ANR0301I Recovery log format in progress; 124 megabytes of 128.
ANR0301I Recovery log format in progress; 128 megabytes of 128.
ANR0302I Recovery log formatting took 5000 milliseconds.
ANR0303I Format rate: 6553.6 pages/second.
ANR0304I Page service time: 0.2 ms.
ANR0305I Recovery log format complete.
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.
ANR0992I Server installation complete.

ANR0900I Processing options file C:\PROGRA~1\tivoli\tsm\server1\dsmserve.opt.
```

ANR7800I DSMSEV generated at 09:45:38 on May 30 2002.

Tivoli Storage Manager for Windows
Version 5, Release 1, Level 1.0

Licensed Materials - Property of IBM

5698-ISE (C) Copyright IBM Corporation 1999,2002. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

ANR8200I TCP/IP driver ready for connection with clients on port 1500.
ANR0200I Recovery log assigned capacity is 128 megabytes.
ANR0201I Database assigned capacity is 2048 megabytes.
ANR4600I Processing volume history file C:\PROGRA~1\TIVOLI\TSM\SERVER1\VOLHIST.
OUT.
ANR4620I Database backup series 7 operation 0 device class CLASS1.
ANR4622I Volume 1: ABA927L1.
ANR4634I Starting point-in-time database restore to date 07/26/2002 17:23:38.
ANR8337I LTO volume ABA927L1 mounted in drive MT01 (mt0.2.0.4).
ANR1363I Input volume ABA927L1 opened (sequence number 1).
ANR4646I Database capacity required for restore is 1024 megabytes.
ANR4649I Reducing database assigned capacity to 1024 megabytes.
ANR4638I Restore of backup series 7 operation 0 in progress.
ANR4639I Restored 15872 of 31905 database pages.
ANR4639I Restored 31808 of 31905 database pages.
ANR4640I Restored 31905 pages from backup series 7 operation 0.
ANR0306I Recovery log volume mount in progress.
ANR4641I Sequential media log redo pass in progress.
ANR4642I Sequential media log undo pass in progress.
ANR1364I Input volume ABA927L1 closed.
ANR4644I A full backup will be required for the next database backup operation.
ANR4635I Point-in-time database restore complete, restore date 07/26/2002
17:23:38.
ANR8468I LTO volume ABA927L1 dismounted from drive MT01 (mt0.2.0.4) in library
LB6.0.0.3.

Wait for the server to start. Ensure that the Administrative command
line client option file is set up to communicate with this server, then
press enter to continue recovery script execution.
Press any key to continue . . .

At this stage, a second command window is opened which starts the TSM server
as shown in Example 8-36.

Example 8-36 The start of TSM server after database recovery

ANR0900I Processing options file C:\PROGRA~1\tivoli\tsm\server1\dsmsevr.opt.
ANR7800I DSMSEV generated at 09:45:38 on May 30 2002.

Tivoli Storage Manager for Windows
Version 5, Release 1, Level 1.0

Licensed Materials - Property of IBM

5698-ISE (C) Copyright IBM Corporation 1999,2002. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

ANR0990I Server restart-recovery in progress.
ANR0200I Recovery log assigned capacity is 128 megabytes.
ANR0201I Database assigned capacity is 1024 megabytes.
ANR0306I Recovery log volume mount in progress.
ANR0287W Contents of the page shadow file dbpgshdw.bdt are not valid.
ANR0285I Database page shadowing started using file dbpgshdw.bdt.
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.
ANR2100I Activity log process has started.
ANR4726I The NAS-NDMP support module has been loaded.
ANR4726I The ServerFree support module has been loaded.
ANR2102I Activity log pruning started: removing entries prior to 07/27/2002
00:00:00.
ANR9969E Unable to open volume H:\TSMDATA\STG_POOL_01.DSM. The most likely
reason is that another TSM server is running and has the volume allocated.
ANR1311E Vary-on failed for disk volume H:\TSMDATA\STG_POOL_01.DSM - unable to
access disk device.
ANR0984I Process 1 for EXPIRATION started in the BACKGROUND at 10:07:37.
ANR0811I Inventory client file expiration started as process 1.
ANR2803I License manager started.
ANR2718W Schedule manager disabled.
ANR8260I Named Pipes driver ready for connection with clients.
ANR8200I TCP/IP driver ready for connection with clients on port 1500.
ANR8280I HTTP driver ready for connection with clients on port 1580.
ANR2860I Server is licensed to support Tivoli Disaster Recovery Manager.
ANR2827I Server is licensed to support Managed System for LAN for a quantity of
20.
ANR2827I Server is licensed to support Managed System for SAN for a quantity of
5.
ANR2841W Server is NOT IN COMPLIANCE with license terms.
ANR0984I Process 2 for AUDIT LICENSE started in the BACKGROUND at 10:07:37.
ANR2820I Automatic license audit started as process 2.
ANR0993I Server initialization complete.
ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli is now ready for use.
ANR2825I License audit process 2 completed successfully - 6 nodes audited.
ANR0987I Process 2 for AUDIT LICENSE running in the BACKGROUND processed 6
items with a completion state of SUCCESS at 10:07:37.

```
ANR2841W Server is NOT IN COMPLIANCE with license terms.
TSM:RADON_SERVER1>
ANR2103I Activity log pruning completed: 4847 records removed.
ANR8840E Unable to open device 1b1.6.0.4 with error -1.
ANR8440E Initialization failed for SCSI library LB6.0.0.3; will retry in 2
minute(s).
ANR0812I Inventory file expiration process 1 completed: examined 1933 objects,
deleting 1896 backup objects, 0 archive objects, 0 DB backup volumes, and 0
recovery plan files. 0 errors were encountered.
ANR0987I Process 1 for EXPIRATION running in the BACKGROUND processed 1896
items with a completion state of SUCCESS at 10:07:59.
ANR0407I Session 1 started for administrator ADMIN (WinNT) (Tcp/Ip
9.1.38.186(2-125)).
```

The scripts have finished successfully, the TSM database was restored, and the TSM server starts. In our case, the library and drive paths has been altered in the `DEVICE.CONFIGURATION.FILE`, which causes the `Unable to open device` error. To correct this, we update the library and drives paths directly in that command window as shown in Example 8-37.

Example 8-37 Commands UPDATE PATH

```
UPDATE PATH RADON_SERVER1 LB6.0.0.3 SRCTYPE=SERVER DESTTYPE=LIBRARY
DEVICE=1b0.6.0.4 ONLINE=YES
ANR2017I Administrator SERVER_CONSOLE issued command: UPDATE PATH RADON_SERVER1
LB6.0.0.3 SRCTYPE=SERVER DESTTYPE=LIBRARY DEVICE=1b0.6.0.4 ONLINE=YES
ANR1722I A path from RADON_SERVER1 to LB6.0.0.3 has been updated.
TSM:RADON_SERVER1>
UPDATE PATH RADON_SERVER1 MT01 SRCTYPE=SERVER DESTTYPE=DRIVE LIBRARY=LB6.0.0.3
DEVICE=mt0.2.0.4 ONLINE=YES
ANR2017I Administrator SERVER_CONSOLE issued command: UPDATE PATH RADON_SERVER1
MT01 SRCTYPE=SERVER DESTTYPE=DRIVE LIBRARY=LB6.0.0.3 DEVICE=mt0.2.0.4
ONLINE=YES
ANR1722I A path from RADON_SERVER1 to LB6.0.0.3 MT01 has been updated.
TSM:RADON_SERVER1>
UPDATE PATH RADON_SERVER1 MT02 SRCTYPE=SERVER DESTTYPE=DRIVE LIBRARY=LB6.0.0.3
DEVICE=mt0.4.0.4 ONLINE=YES
ANR2017I Administrator SERVER_CONSOLE issued command: UPDATE PATH RADON_SERVER1
MT02 SRCTYPE=SERVER DESTTYPE=DRIVE LIBRARY=LB6.0.0.3 DEVICE=mt0.4.0.4
ONLINE=YES
ANR1722I A path from RADON_SERVER1 to LB6.0.0.3 MT02 has been updated.
```

Once the device configuration is set correctly you can mount copy storage pool volumes upon request, check in the volumes in advance, or manually place the volumes in the library and ensure consistency by issuing the **AUDIT LIBRARY** command.

After the audit of the library, the TSM server is ready to provide restore of TSM clients using the copy storage pool volumes directly. If you plan to stay for some time on the backup TSM machine, you will need to create primary storage pools and primary storage pool volumes to backup data. You can restore storage pools and volumes using the script RECOVERY.SCRIPT.NORMAL.MODE.COMD.

Restore primary storage pools

To restore primary storage pools from copy storage pools, execute RECOVERY.SCRIPT.NORMAL.MODE script file. If client machines are damaged, you may want to delay this action until after all clients are recovered.

Note: This action is optional because TSM can access the copy storage pool volumes directly to restore client data. Using this feature, you can minimize client recovery time, because server primary storage pools do not have to be restored first.

Enter the script file name at the command prompt and follow with the administrator name and password as parameters. This script creates replacement primary storage pool volumes, defines them to TSM and restores them from the copy storage pool volumes.

Example 8-38 Invoke RECOVERY.SCRIPT.NORMAL.MODE

```
C:\DRM\plan>C:\DRM\plan\RADON.RECOVERY.SCRIPT.NORMAL.MODE.COMD admin admin  
Replace H:\TSMDATA\STG_POOL_01.DSM DISK 2,048.0M in DISK_STG_01 by
```

Windows NT Server Volume Extent/Volume Formatting Program

Licensed Materials - Property of IBM

5698-ISE (C) Copyright IBM Corporation 1990, 2002. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

Activity log record

```
07/28/2002 13:21:24 ANR2017I Administrator ADMIN issued command: DEFINE  
VOLUME DISK_STG_01 C:\TSMDATA\STG_POOL_01.DSM acc=READW  
07/28/2002 13:21:24 ANR2206I Volume C:\TSMDATA\STG_POOL_01.DSM defined in  
storage pool DISK_STG_01 (device class DISK).  
07/28/2002 13:21:24 ANR2017I Administrator ADMIN issued command: DEFINE  
VOLUME LTO_STGP_01 ABA990L1 acc=READW  
07/28/2002 13:21:24 ANR1305I Disk volume C:\TSMDATA\STG_POOL_01.DSM varied  
online.  
07/28/2002 13:21:24 ANR2206I Volume ABA990L1 defined in storage pool  
LTO_STGP_01 (device class CLASS1).  
07/28/2002 13:21:24 ANR2017I Administrator ADMIN issued command: RESTORE
```

```

                                STGPOOL DISK_STG_01
07/28/2002 13:21:24 ANR0984I Process 7 for RESTORE STORAGE POOL started in
the BACKGROUND at 13:21:24.
07/28/2002 13:21:24 ANR1230I Restore of primary storage pool DISK_STG_01
started as process 7.
07/28/2002 13:21:24 ANR2110I RESTORE STGPOOL started as process 7.
07/28/2002 13:21:24 ANR2017I Administrator ADMIN issued command: RESTORE
STGPOOL LTO_STGP_01
07/28/2002 13:21:24 ANR0984I Process 8 for RESTORE STORAGE POOL started in
the BACKGROUND at 13:21:24.
07/28/2002 13:21:24 ANR1254I Removable volume ABA926L1 is required for
restore processing.
07/28/2002 13:21:24 ANR1230I Restore of primary storage pool LTO_STGP_01
started as process 8.
07/28/2002 13:21:24 ANR2110I RESTORE STGPOOL started as process 8.
07/28/2002 13:21:24 ANR0405I Session 187 ended for administrator ADMIN
(WinNT).
07/28/2002 13:21:24 ANR1254I Removable volume ABA924L1 is required for
restore processing.
07/28/2002 13:22:01 ANR8337I LTO volume ABA924L1 mounted in drive MT02
(mt0.4.0.4).
07/28/2002 13:22:27 ANR8337I LTO volume ABA990L1 mounted in drive MT01
07/28/2002 14:21:04 ANR1234I Restore process 8 ended for storage pool
LTO_STGP_01.
07/28/2002 14:21:04 ANR0986I Process 8 for RESTORE STORAGE POOL running in
the BACKGROUND processed 183051 items for a total of 26,484,716,732 bytes with
a completion state of SUCCESS at 14:21:04.
07/28/2002 14:21:04 ANR1238I Restore of primary storage pool LTO_STGP_01 has
ended. Files Restored: 183051, Bytes Restored:
26484716732, Unreadable Files: 0, Unreadable Bytes: 0.
07/28/2002 14:21:04 ANR2208I Volume ABA920L1 deleted from storage pool
LTO_STGP_01.
07/28/2002 14:21:05 ANR1341I Scratch volume ABA920L1 has been deleted from
storage pool LTO_STGP_01.
07/28/2002 14:26:40 ANR0986I Process 7 for RESTORE STORAGE POOL running in
the BACKGROUND processed 46317 items for a total of 864,725,234 bytes with a
completion state of SUCCESS at 14:26:40.
07/28/2002 14:26:40 ANR1238I Restore of primary storage pool DISK_STG_01 has
ended. Files Restored: 46317, Bytes Restored: 864725234, Unreadable Files: 0,
Unreadable Bytes: 0.
07/28/2002 14:26:40 ANR2208I Volume H:\TSMDATA\STG_POOL_01.DSM deleted from
storage pool DISK_STG_01. (mt0.2.0.4).

```

As an alternative, you can use the recovery script as a guide and manually run each step.

The ongoing results of the TSM server restore process using the DRM-generated scripts is logged in the following files.

- ▶ LOGANDDDB.VOLUMES.INSTALL.LOG
- ▶ LICENSE.REGISTRATION.LOG
- ▶ COPYSTGPOOL.VOLUMES.AVAILABLE.LOG
- ▶ PRIMARY.VOLUMES.DESTROYED.LOG
- ▶ COPYSTGPOOL.VOLUMES.DESTROYED.LOG
- ▶ PRIMARY.VOLUMES.REPLACEMENT.CREATE.LOG
- ▶ PRIMARY.VOLUMES.REPLACEMENT.LOG
- ▶ STGPOOLS.RESTORE.LOG

Summary of example DR plan

Using TSM Disaster Recovery Manager, we were able quickly to restore TSM server named RADON_Server1 on another set of hardware using the offsite backup of the TSM database and the copy storage pool volumes. Once the server is up and running, you should use DRM procedures to return the retrieved vault recovery volumes so they are available in the event of another disaster. You will of course resume normal backup and disaster protection operations.



Overview of bare metal recovery techniques

The next several chapters will document different techniques used to perform bare metal recoveries (recovery of a machine's entire set of data, including operating system, applications and data files) on a variety of platforms. These platforms are Solaris, AIX, Windows 2000 and Linux.

For some of these platforms, there may be more than one option for bare metal recovery. Therefore, this short chapter will provide a framework of concepts to help guide when one or more of these solutions may be used. This chapter will also discuss some common alternatives to a full-blown bare metal recovery (BMR) solution.

Keep in mind that there may be some hosts that are of such criticality to your organization that using backup data to recover from is too slow. In those instances solutions such as replicated data or applications will provide faster recovery, but incur more cost. This concept has been discussed several times within the DR chapters. In all cases, the value of a host to the organization will guide the type of DR support.

Important: In the following chapters we suggest some approaches for DR on various operating system platforms. These have not been thoroughly tested, and are not intended to be definitive or exhaustive examples, but simply to give some starting points and ideas, which have worked in our ITSO simple lab environment. You should consult your operating system documentation and perhaps engage a service provider to develop and test detailed procedures tailored to your own environment.

9.1 Bare metal recovery steps

Any type of BMR will have essentially the same steps. These are:

1. Boot the hardware
2. Partition the hard disk space
3. Access the backup data
4. Restore backup data
5. Post-restore configuration

What will differentiate the types of BMR solutions will be *how* each of these steps is accomplished.

9.1.1 Manual bootstrap BMR

For example, the most inexpensive BMR solution is to manually reinstall a minimal copy of the base operating system (to bootstrap the BMR process), install the backup tool and then to restore the backed up data (see Figure 9-1).

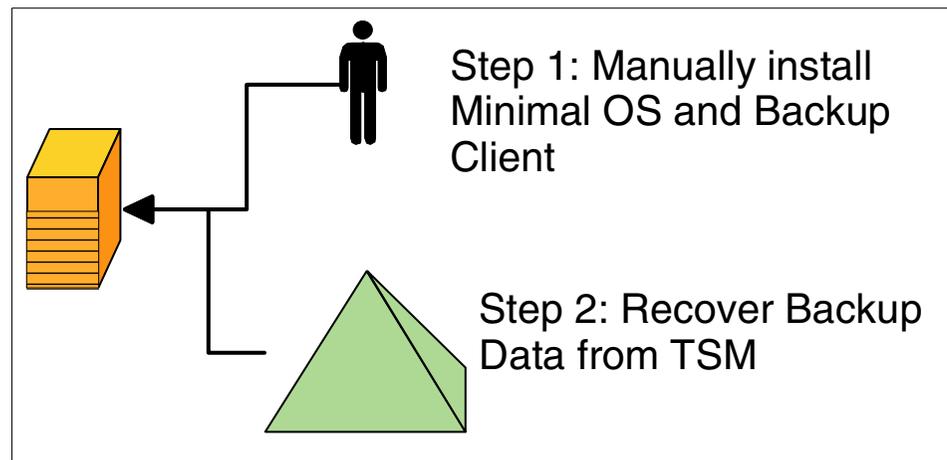


Figure 9-1 Manual bootstrap BMR

While this has the advantage of little to no up front cost, it requires more manual effort and an estimated half hour to one and a half hours per host to install the operating system, before it is possible to begin restoring the actual data (depending on the operating system). In the case of more complex UNIX hosts, this time may be significantly longer. This type of solution may be suitable for a small number of critical hosts or larger numbers of non-critical hosts.

9.1.2 Operating system image bootstrap BMR

Next up the ladder from a manual process is one that uses third party software or native utilities to expedite the recovery of the minimal operating system installation (see Figure 9-2). Although this seems like a minor change from a manual process, it has some interesting variations that can make for an effective BMR solution without an unreasonable expenditure.

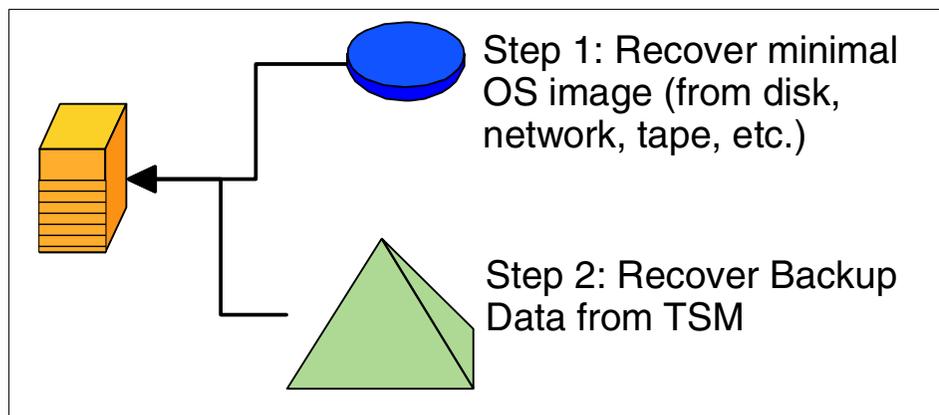


Figure 9-2 Operating system image bootstrap BMR

After a fast restore of the operating system from the additional media, the normal backup engine is accessed to restore the entire machine to its last state. The third party software in this case can be imaging software or other specialized clients. Enterprise-class UNIX operating systems often have this software included as a utility in the base configuration (for example, the `mksysb` command on AIX). By using a generic operating system image simply as a bootstrap into the BMR process, this image can serve many hosts therefore, avoiding the administrative overhead of keeping current and complete images for all hosts.

A common variation on this solution is to utilize an existing infrastructure that supports other necessary IT tasks. For example, if a standardized operating system build is installed using disk-imaging software, then that same solution can obviously be exploited for BMR. All that is necessary is to include the backup/restore software on the standard load. Likewise, if a network distribution infrastructure is in place for software distribution, it can be exploited to distribute a minimal operating system installation in a recovery scenario. Of course, if a network distribution system was utilized, it would have to be at least partially available when needed to meet the defined disaster recovery requirements for your organization.

Another less common variation on an image bootstrap BMR is to have a secondary hard disk (or even a partition on the same volume that hosts the live

operating system) installed that contains a minimal emergency operating system with sufficient capabilities to run local recovery utilities or the enterprise backup/restore client. This is often a very good way to exploit older hardware (hard disks too small for a current operating system load, for example) to increase uptime. However, once again, this solution does not address all Disaster Recovery scenarios.

The use of a minimal operating system installed recovered to its own partition (or secondary operating system that was already installed) can be used in combination with hot image backups from the enterprise backup/restore product. Hot image backups can be performed on a running server (including the operating system partition) but these hot backups cannot be restored onto the same partition currently running the operating system. By using another partition to run a minimal operating system and the backup/restore software, that image can be restored to another inactive partition. Restoring those files that could not be backed up during the hot image backup will complete the BMR. This data would include system configuration databases, registries, directories, and so on.

9.1.3 Live media bootstrap BMR

With many Operating Systems, it is possible to boot and run a fully operating instance of the operating system from removable media, such as CD-ROM or DVD. This live operating system instance would have been previously configured with the necessary drivers and other software to adequately access the normal backup data (file system or image) and restore it down to the host. See Figure 9-3.

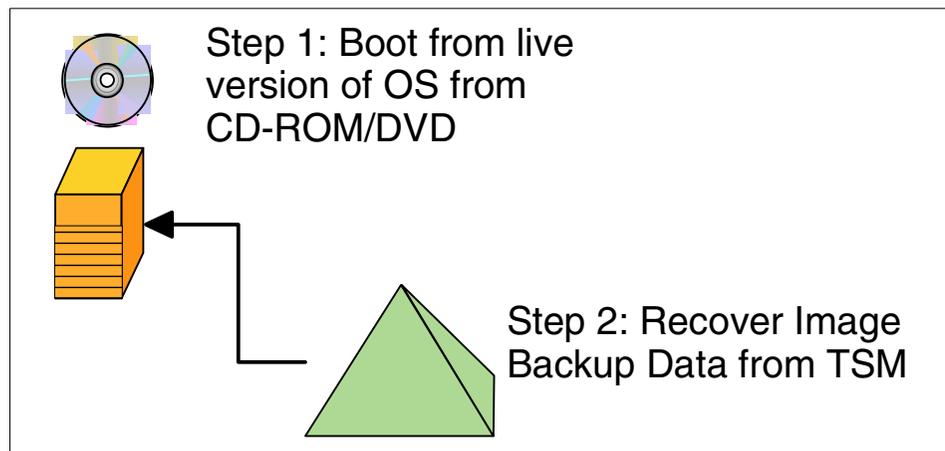


Figure 9-3 Live media bootstrap BMR

Currently, this technology is exploited best by UNIX operating systems, such as Linux. In the past, Windows has not been able to be adequately accessed using this method, but Microsoft is enhancing product line with both an embedded version of Windows XP and an enhanced OEM distribution toolkit. Potentially, either of these could be exploited for a live CD recovery utility.

9.1.4 Automated BMR suite

It is possible to tie in many of the processes described above into a more seamless solution. For example, VERITAS Bare Metal Restore combines normal network boot technologies with an intelligent recovery operating system that automates the BMR process. For example, a host would boot off the network (if supported for that platform, otherwise local boot media is used), download a minimal operating system and then interface with TSM to restore the machines original data. Bare Metal Restore will provide the capability to recover critical hosts simultaneously. However, it is important to plan the configuration of such a complex tool appropriately to plan for all Disaster Recovery contingencies.



Solaris client bare metal recovery

In this chapter we present how to use the **ufsdump** and **ufsrestore** commands to recover a Sun Solaris 2.7 7 client environment that uses TCP/IP to communicate with TSM server:

We discuss the kind of information that should be backed up in preparation for a disaster, how to prepare for the **ufsrestore** restore, and how to use this technique in conjunction with the previously saved information to recover a Solaris system after a disaster.

10.1 Sun Solaris Version 7 recovery

In this chapter we focus on Solaris Version 7, which was the only version tested, however the techniques will probably work on other Solaris versions.

10.1.1 Product Overview

Solaris Version 7 is Sun Microsystems's implementation of the UNIX operating system. Solaris is typically used in medium-scale to large-scale environment, and is a full multitasking, multiuser operating system. To understand the issues involved in Solaris client recovery, you must first understand some of the specifics of the Solaris Operating System.

The Solaris operating environment runs on two types of hardware, or platforms: SPARC and x86 (Intel). The Solaris operating environment runs on both 64-bit and 32-bit address spaces. IBM TSM client version 5.1 and 4.2 for Sun Solaris is supported only on SPARCstations based on sun4u architecture.

10.1.2 Backup

The most frequently used utilities for system backup and restore on the Solaris platform are **ufsdump** and **ufsrestore**. Solaris Version 7 also provides other utilities for backup and restore of files, and file systems. Commands for backing up and restoring files and file systems are summarized in Table 10-1.

Table 10-1 *Commands for copying files and files systems*

Task to do	Command	Notes
Back up complete or individual file systems to local or remote tape	ufsdump command	
Restore complete file systems or individual files from removable media to a working directory	ufsrestore command	
Copy, list, and retrieve files on tape	tar , cpio , or pax command	
Copy, list, and retrieve files on diskette	tar , cpio , or pax command	
Copy master disk to a clone disk	dd command	

Task to do	Command	Notes
Backup and restore a NIS+ master server	nisbackup and nisrestore commands	
Back up complete file system for all systems on a network from server	Solstice Backup Software	

The **ufsdump** command provides a full, cumulative incremental, or discrete incremental backups. TSM provides more granular backups of data files, but cannot be used to restore the root system after a disaster. Therefore, we will use **ufsrestore** to restore the root file systems, then restore user data with the TSM client.

10.2 Pre disaster preparation

The disaster we are preparing for is the complete loss of the TSM client, where provision has to made to restore everything from the bare metal up. In this section we discuss the information to collect and save before a disaster to enable a bare metal restore.

10.2.1 TSM backups

For this test we divided all data on the system into user data and system data. We used **ufsdump** and **ufsrestore** to back up and restore the system data and managed the user data with TSM. We defined system data as the file systems necessary to the base operating system, this is usually the `/`, and `/usr` file systems. We defined user data as anything that is not included in system data.

For the purpose of this test we used the default include/exclude list for a TSM UNIX client. We also set the copy group to use shared static option to avoid fuzzy backups. For more information about using and configuring the TSM client on Sun Solaris, see *Tivoli Storage Manager for UNIX Backup-Archive Clients Installation and User's Guide*, GC32-0789.

10.2.2 Operating environment of the Solaris client

In preparation for disaster recovery we recommend that you collect and save offsite the information listed below in addition to the TSM backup data. Depending on your installation, you could save this information by using the DRM feature, or you could save it on hardcopy stored offsite. The way in which you save the information is not as important as ensuring that it is saved somewhere where it can be easily retrieved along with the client recovery media

when a disaster occurs. Further information on parameters for the commands in this section can be obtained by looking at the man pages or Solaris online documentation.

Hardware configuration

Use the command `/usr/platform/`uname -i`/sbin/prtdiag -v` command to collect information about the machine. Other useful commands are:

- `prtconf -v`
- `psrinfo -v`
- `isainfo -v`
- `dmesg`
- `iostat -En`
- `netstat -in`

Logical storage configuration

Information regarding the storage setup of the client machine is vital to system recovery. Devices are referenced in three ways in the Solaris environment.

- ▶ **Physical device name** – Represents the full device pathname in the device information hierarchy. Physical device names are displayed using the following commands:

- `dmesg`
- `format`
- `sysdef`
- `prtconf`

Physical device files are found in the `/devices` directory.

- ▶ **Instance name** – Represents the kernel's abbreviation name for every possible device on the system. For example, `sd0` and `sd1` represent the instance names of two disk devices. Instance names are mapped in the `/etc/path_to_inst` file and are displayed using the following commands:

- `dmesg`
- `sysdef`
- `prtconf`

- ▶ **Logical device name** – Used by system administrators with most file system commands to refer to devices. Logical device files in the `/dev` directory are symbolically linked to physical device files in the `/devices` directory.

All disk, and partitioning information should be recorded and stored before the disaster in either hardcopy or within DRM. Information should be gained by the commands `df` or `prvtoc`.

Optional: statistics for post disaster recovery verification

You may want to collect and save information that you can use after a disaster once the restore process is complete to validate that all information has indeed been restored correctly. Each installation will have its own ideas about the best way to verify that the recovery is complete.

10.2.3 Creating ufsdump backup

The following steps provide the general steps for backing up file systems using the **ufsdump** command. The examples show specific uses of options and arguments.

1. Become superuser.
2. Boot the system to run level S (single-user mode). Switch machine on and use the **boot** command shown in Example 10-1.

Example 10-1 Start system in single user mode

```
OK boot -s
Boot device: /pci@1f,4000/scsi@3/disk@6,0:f File and args: -s
SunOS Release 5.7 Version Generic_106541-08 [UNIX(R) System V Release 4.0]
Copyright (c) 1983-1999, Sun Microsystems, Inc.
TSI: gfxp0 is GFX*P @ 1280x1024
configuring network interfaces: hme0
Hostname: sol-e

INIT: SINGLE USER MODE

Type control-d to proceed with normal startup,
(or give root password for system maintenance):
single-user privilege assigned to /dev/console.
Entering System Maintenance Mode

Jul 29 11:40:28 su: `su -root` succeeded for root on /dev/console
Sun Microsystems Inc. SunOS 5.7 Generic October 1998
You have mail.
root@solemio:/:

INIT: SINGLE USER MODE
```

3. Check the file system for consistency with the **fsck** command, shown in Example 10-2. Running the **fsck** command checks for consistency of file systems in case something like a power failure has left the files in an inconsistent state.

Example 10-2 File-system check command fsck

```
# fsck /dev/rdisk/c0t0d0s0
```

```

** /dev/rdisk/c0t0d0s0
** Last Mounted on /
** Phase 1 - Check Blocks and Sizes
** Phase 2 - Check Pathnames
** Phase 3 - Check Connectivity
** Phase 4 - Check Reference Counts
** Phase 5 - Check Cyl groups
3929 files, 430921 used, 62767 free (567 frags, 7775 blocks, 0.1%
fragmentation)
#

```

4. Identify the device name of an attached tape drive. We are using /dev/rmt/0.
5. Insert a tape that is not write protected into the tape drive.
6. Back up file systems using the **ufsdump** command. We need a full backup of root file system. This is specified by option 0, as in Example 10-3.

Example 10-3 Back up of root file system

```

root@solemio:/: ufsdump 0ucf /dev/rmt/0 /
DUMP: Writing 63 Kilobyte records
DUMP: Date of this level 0 dump: Mon Jul 29 17:56:00 2002
DUMP: Date of last level 0 dump: the epoch
DUMP: Dumping /dev/rdisk/c0t0d0s0 (sol-e:/) to /dev/rmt/0
DUMP: Mapping (Pass I) [regular files]
DUMP: Mapping (Pass II) [directories]
DUMP: Estimated 867660 blocks (423.66MB)
DUMP: Dumping (Pass III) [directories]
DUMP: Dumping (Pass IV) [regular files]
DUMP: 867508 blocks (423.66MB) on 1 volume at 5790 KB/s
DUMP: DUMP IS DONE
DUMP: Level 0 dump on Mon Jul 29 17:59:49 2002
root@solemio:/:

```

7. If prompted, remove the tape and replace with another volume.
8. Label each tape with the volume number, level, date, system name, disk slice, and file system.
9. Bring the system back to run level 3 by pressing Ctrl-d.
10. Verify the backup was successful by using the **ufsrestore -t** command to display the tape contents.

Use the same procedure to backup the /usr file system.

10.2.4 Restore root file system using ufsrestore command

The **ufsrestore** command copies files to disk, relative to the current working directory, from backups created using the **ufsdump** command. You can use **ufsrestore** to reload an entire file system hierarchy from a level 0 dump and incremental dumps that follow it or to restore one or more single files from any dump tape. If **ufsrestore** is run as superuser, files are restored with their original owner, last modification time, and mode (permissions).

Before you start to restore files or file systems, you need to know or have:

- ▶ Which tapes you need
- ▶ The raw device name on which you want to restore the file system
- ▶ The type of tape drive you will use
- ▶ The device name for the tape drive
- ▶ Bootable Solaris environment CD

When a disaster strikes, first of all we need to restore the / and /usr file systems. Then we can restore user data using TSM client. Here is the procedure.

1. Boot from the Solaris 7 CD to single user mode as shown in Example 10-4.

Example 10-4 Boot from CD to single user mode

```
ok boot cdrom -s
Boot device: /pci@1f,4000/scsi@3/disk@6,0:f File and args: -s
SunOs Release 5.7 Version Generic_106541-08 [UNIX(R) System V Release 4.0]
Copyright (c) 1983-1999, Sun Microsystems, Inc.
TSI: gfxp0 is GFX*P @ 1280x1024
Configuring /dev and /devices
-
INIT: SINGLE USER MODE
#
```

2. Format a partition as ufs file system where root directory will be restored.

Example 10-5 Format partition using newfs command

```
# newfs /dev/rdisk/c0t0d0s0
newfs: construct a new file system /dev/rdisk/c0t0d0s0: (y/n)? y
/dev/rdisk/c0t0d0s0: 6144200 sectors in 1711 cylinders of 27 tracks, 133
sectors 3000.1MB in 107 cyl groups (16 c/g, 28.05MB/g, 4544 i/g)
super-block backups (for fsck -F ufs -o b=#) at:
32, 57632, 115232, 172832, 230432, 288032, 345632, 403232, 460832, 518432,
576032, 633632, 691232, 748832, 806432, 864032, 921632, 979232, 1036832,
1094432, 1152032, 1209632, 1267232, 1324832, 1382432, 1440032, 1497632,
1555232, 1612832, 1670432, 1728032, 1785632, 1838624, 1896224, 1953824,
2011424, 2069024, 2126624, 2184224, 2241824, 2299424, 2357024, 2414624,
2472224, 2529824, 2587424, 2645024, 2702624, 2760224, 2817824, 2875424,
2933024, 2990624, 3048224, 3105824, 3163424, 3221024, 3278624, 3336224,
```

```
3393824, 3451424, 3509024, 3566624, 3624224, 3677216, 3734816, 3792416,
3850016, 3907616, 3965216, 4022816, 4080416, 4138016, 4195616, 4253216,
4310816, 4368416, 4426016, 4483616, 4541216, 4598816, 4656416, 4714016,
4771616, 4829216, 4886816, 4944416, 5002016, 5059616, 5117216, 5174816,
5232416, 5290016, 5347616, 5405216, 5462816, 5515808, 5573408, 5631008,
5688608, 5746208, 5803808, 5861408, 5919008, 5976608, 6034208, 6091808,
#
```

3. Check the partition where the root partition will be restored as shown in Example 10-6.

Example 10-6 Checking partition using fsck command

```
# fsck /dev/rdisk/c0t0d0s0
** /dev/rdisk/c0t0d0s0
** Last Mounted on /
** Phase 1 - Check Blocks and Sizes
** Phase 2 - Check Pathnames
** Phase 3 - Check Connectivity
** Phase 4 - Check Reference Counts
** Phase 5 - Check Cyl groups
2 files, 9 used, 3009585 free (17 frags, 376196 blocks, 0.0% fragmentation)
#
```

4. Now that we have a partition ready to restore the operating system, mount it and change directory to it, as in Example 10-7.

Example 10-7 Mount partition for restore

```
# mount /dev/dsk/c0t0d0s0 /a
# cd /a
# ls
lost+found
#
```

5. Now we are ready to restore partition from the tape using the **ufsrestore** command as in Example 10-8.

Example 10-8 Restore root partition using ufsrestore

```
# ufsrestore rvf /dev/rmt/0
Verify volume and initialize maps
Media block size is 126
Dump date: Mon Jul 29 17:59:49 2002
Dumped from: the epoch
Level 0 dump of / on sol-e:/dev/dsk/c0t0d0s0
Label: none
Begin level 0 restore
Initialize symbol table.
Extract directories from tape
```

```
Calculate extraction list.
Warning: ./lost+found: File exists
extract file ./usr
extract file ./var
extract file ./opt
extract file ./scratch
extract file ./scratch/lotti
extract file ./scratch/oracle
extract file ./tmp
extract file ./platform
extract file ./platform/SUNW,Ultra-1
...
...
Add links
Set directory mode, owner, and times.
Check the symbol table.
Check pointing the restore
#
```

6. When partition is restored, we have to remove auxiliary file created by **ufsrestore**, shown in Example 10-9.

Example 10-9 Remove restoresymtable

```
# rm restoresymtable
#
```

7. Now we can umount file system, as in Example 10-10.

Example 10-10 Umount file system

```
# cd /
# umount /a
#
```

8. Then reinstall the boot sector, Example 10-11.

Example 10-11 Reinstallation of boot sector

```
# cd /usr/platform/$(uname -m)/lib/fs/ufs
# installboot bootblk /dev/rdisk/c0t0d0s0
#
```

9. We can one more time check disk partition using **fsck** command; see Example 10-6 on page 236.
10. Repeat Steps 2 through 7, to restore the /usr partition. You should use the same partition device name as previously so this will match the backup information on the root partition. Use the tape made with **ufsdump** of the /usr file system to restore it.

11. The last step of the restore process is to re-boot to normal mode.

You have now restored your root file system. This should include the TSM client executable and configuration files, presuming they were installed in the root file system as well. If so, you can then use the TSM client to restore user data (that is, outside the root file systems). Otherwise, first re-install the TSM client, then restore the data.



AIX client bare metal recovery

This section describes the primary ways to perform a bare metal restore an AIX system. The three methods discussed are:

- ▶ **mksysb**, an AIX-provided tool which creates bootable backups of the root volume group
- ▶ Network Installation Management (NIM), an installable AIX component which provides network-based install, boot and upgrade services
- ▶ SysBack — a separately priced product from IBM which provides basic volume-group based backups.

More information on these topics can be obtained from the AIX system documentation, from the redbook *NIM: From A to Z in AIX 4.3*, SG24-5524 and from the SysBack Web site:

<http://sysback.services.ibm.com/>

11.1 AIX system backup overview

Bare metal restores of AIX systems can be achieved by using three primary methods.

- ▶ **Mksysb:** The `mksysb` command is the standard AIX command for making system backups of the rootvg volume group. Mksysb backups can be made to tape, disk, CD-R, or files.
- ▶ **Network Installation Management (NIM):** NIM is an installable feature of AIX, which provides software distribution and installation services to networked clients. In addition to backing up networked system mksysb images, NIM can be implemented to allow multiple AIX systems to boot over a network connection for bare metal restore operations. TSM can be used to protect the data in the NIM environment for additional data protection.
- ▶ **SysBack:** SysBack is a separately licensed application that provides AIX system backup and recovery using networks and tape resources.

Each solution offers a variety of implementation methods and characteristics. Table 11-1 lists general characteristics of each software tool.

Table 11-1 *Mksysb, NIM, and SysBack characteristics*

Characteristic	mksysb	NIM	SysBack
Centrally Managed Backups	No	Yes, with scripting and TSM	Yes
Unprompted network boot and bare metal restores	No	Yes	Yes, using NIM services
SMIT interface	Yes	Yes	Yes
Backup/Restore support for Tape, CD and DVD	Yes	N/A	Yes
Exclude list support for backups	Yes	Yes	Yes

11.1.1 AIX root volume group overview

The AIX root volume group contains the operating system, installable applications and many configuration files. Backup strategies for this operating system data depends on the enterprise's recovery strategy. Full system image backups provide the most efficient means to recover a complete system. The

tools we discuss provide bootable recovery tapes, bootable recovery CD's, or networked based bootable images. Enterprise disaster recovery requirements must map to the AIX bare metal restore solution of choice.

To understand how system backup methods function, it is important to have an understanding of AIX file system structure and data placement. Figure 11-1 shows the basic root volume group file system structure for AIX.

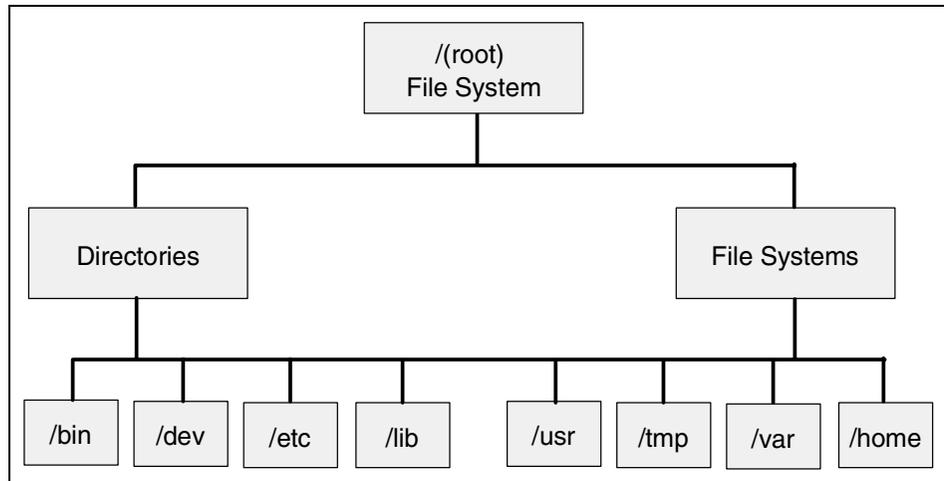


Figure 11-1 Typical AIX file system for system data

System data makes up the operating system and its extensions. The system data is always to be kept in the system file systems, namely / (root), /bin, /dev, /etc, /lib, /usr, /tmp, and /var.

It is good practice to keep user (that is non-system) data out of the root volume group file systems. User data might be files belonging to individual users, or it might be application-specific data. User data tends to change much more often than operating system data.

At a minimum, we suggest weekly full system backups to bootable media for disaster recovery purposes. This media should be kept on and offsite for a range of disaster scenarios. By using TSM, more frequent incremental backups of user data can be automated.

11.2 Using the mksysb for bare metal restore

The `mksysb` command is used to clone the rootvg volume group, which contains the operating system. The root volume group contains the following:

- ▶ Startup executables
- ▶ Base operating system commands and files
- ▶ System configuration information
- ▶ Optional software products

This command creates a bootable backup image which is written to a tape device, a CD-ROM, or to a file for use by NIM. All mounted JFS (Journaled File Systems) and JFS2 (Enhanced Journaled File Systems) in the root volume group are backed up. The cloned system image will have the same configuration information as the original system. There are many options for the **mksysb** command. For details, please refer to the man page for this command.

Hint: In AIX 5.1.1, if you modify the bos.inst.data file before creating the **mksysb** output, using `RECOVER_DEVICE=no`, then device specific customization will not be performed upon restoration. This option would be useful in a coldsite disaster recovery scenario, where recovery equipment configurations are often not identical to the primary site.s

Figure 11-2 shows the general data layout of a **mksysb** tape.

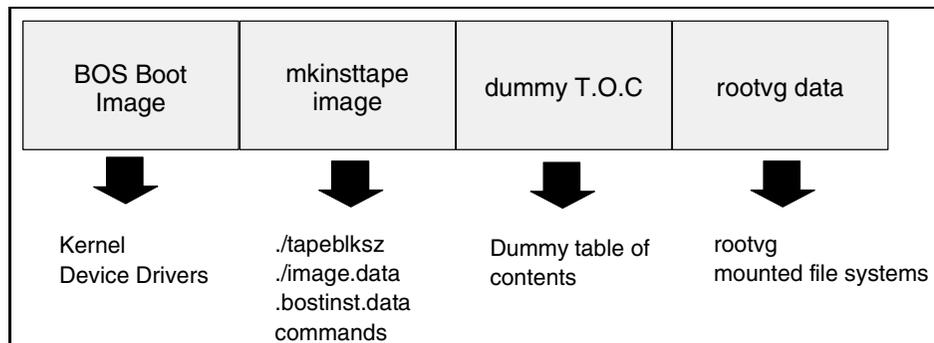


Figure 11-2 Layout of a **mksysb** tape

As the **mksysb** tape is read during a system restore, the kernel and device drivers are loaded, the system is prepared to receive the restored data, and the rootvg data is restored.

11.2.1 Exclude files for **mksysb** backup

On each AIX machine, the `/etc/exclude.rootvg` file can be created to exclude files or directories from the rootvg backup. This can limit the amount of extraneous data in the **mksysb** backup image and speed the overall bare metal recovery process. Example 11-1 shows an example of an `/etc/exclude.rootvg` file.

Example 11-1 Sample /etc/exclude.rootvg file

```
/examples/  
^./tmp/  
^./logs/  
*.log
```

The exclude files option can be chosen from the SMIT interface, which by default refers to the /etc/exclude.rootvg file. From the command line, use the **-e** option for the **mksysb** command.

11.2.2 Saving additional volume group definitions

Most systems have one or more additional AIX volume groups configured for application data, which also need to be restored during a bare metal restore. The **mksysb** command only backs up the root volume group, therefore additional consideration is required for this extra data.

The AIX **savevg** command finds and backs up all files belonging to a specified volume group. It creates a volume group definition file by calling the **mkvgdata** command. **Savevg** alone does not provide enough granularity to just save volume group configuration data. The actual files in the volume group can be backed up with TSM. In addition to this, the specific volume group configuration needs to be saved so that it can be re-created at a later date to receive the restored data in the event of a bare metal restore. We present here a script, **vg_recovery**, which saves volume group configuration data in an easily restorable format.

The **vg_recovery** script saves the volume group configuration data, without actually backing up any data within the volume group. Example 11-2 shows the content of the **vg_recovery** script. Details of how to download a softcopy of this script are given in Appendix D, “Additional material” on page 387.

Example 11-2 The vg_recovery script

```
# Desc:   Program to build savevg filesystem data for future system recoveries  
# Build exclude list  
#  
#  
for i in `lsvg -o`  
do  
  
    if [ $i != "rootvg" ]  
    then  
        # Build filesystem exclude list  
        lsvg -l $i | egrep -v "LV|N\|A" | awk '{print $7}' > /etc/exclude.$i  
        if [[ $? != 0 ]]; then  
            exit 1  
        fi  
    fi  
done
```

```

fi

# execute savevg to build vg recovery data
savevg -i -m -e -v -f /usr/local/vgdata/$i $i
if [[ $? != 0 ]]; then
    exit 1
fi

fi

done

```

Tip: If the `vg_recovery` script does not execute, run `#chmod +x vg_recovery` in the directory where you have saved the script.

The `vg_recovery` script saves volume group configuration data in the `/usr/local/vgdata` directory under each volume group name. This data is included in the `mksysb` creation and will be available after completion of the `mksysb` restore. To re-create the volume group information saved by the `vg_recovery` script, you can use SMIT (as shown in Example 11-3) or the `restvg` command. You have the option to specify alternate disk volume names if these are different in the recovery system.

Example 11-3 Restoring a volume group using the `restvg` command

Remake a Volume Group

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

```

                                                    [Entry Fields]
* Restore DEVICE or FILE                               [ /usr/local/vgdata/vg01 ]
+ /
  SHRINK the filesystems?                               no                +
  PHYSICAL VOLUME names                                [ ]                  +
    (Leave blank to use the PHYSICAL VOLUMES listed
    in the vname.data file in the backup image)
  Use existing MAP files?                               yes                +
  Physical partition SIZE in megabytes                  [ ]
+ #
    (Leave blank to have the SIZE determined
    based on disk size)
  Number of BLOCKS to read in a single input           [ ]
#
    (Leave blank to use a system default)

```

This command can also be executed on the command line via:

```
#restvg -q -f /usr/local/vgdata/vg01
```

The volume group configuration to be rebuilt is provided as the only argument to the `restvg` command. The `/usr/local/vgdata` directory contains volume group information about the previous disk volumes and logical volume layout. The `restvg` command will use this information to rebuild volume group on the same disk set.

11.2.3 Classic mksysb to tape

To create a `mksysb` image to a tape device, select:

SMIT -> System Storage Management (Physical and Logical Storage) -> System Backup Manager -> Back Up the System -> Back Up this system to tape / file

Select the device and appropriate options and run the command. Example 11-4 shows a sample execution.

Example 11-4 Mksysb interface on AIX

Back Up the System

Type or select values in entry fields.

Press Enter AFTER making all desired changes.

	[Entry Fields]
WARNING: Execution of the <code>mksysb</code> command will result in the loss of all material previously stored on the selected output medium. This command backs up only <code>rootvg</code> volume group.	
* Backup DEVICE or FILE	[/dev/rmt1]
+/ Create MAP files?	no +
EXCLUDE files?	no +
List files as they are backed up?	no +
Generate new <code>/image.data</code> file?	yes +
EXPAND <code>/tmp</code> if needed?	no +

This command can also be executed from the command line as

```
#mksysb -i /dev/rmt0.
```

The procedure for creating `mksysb` tapes is simple and can be performed using locally attached tape devices. Writing `mksysb` images to tape has certain

advantages and limitations. Bootable tape backups are easily moved from site to site and full system restore performance is acceptable. Tape media can be re-used for additional mksysb images. This backup procedure can be performed using a portable tape device, an internal tape drive, or a SAN attached tape device which supports manual mounting. As a rule, the simpler the tape configuration, the better.

11.2.4 Mksysb to CD-ROM or DVD

You can also use DVD (DVD-R or DVD-RAM) or CD (CD-R or CD-RW) devices to create bootable **mksysb** media on AIX 5L version 5.1. AIX 4.3.3 supports this procedure for CD-R devices only. This media type is relatively inexpensive and is easy to store or transport for disaster recovery.

The **mkcd** command allows users to write bootable mksysb images to writable CD media. CD-R drives and media have become very inexpensive. Each CD-R disk can contain up to 700 MB of data and DVD media holds up to 4.7 GB per surface, which is usually sufficient for operating system data in a mksysb format. The highest quality media is recommended.

Many CD-R devices are available on the market. Currently, four have been tested with AIX for use in this capacity.

- ▶ Yamaha CRW4416SX - CD-RW
- ▶ RICOH MP 6201S- CD-R
- ▶ Panasonic 7502-B - CD-R
- ▶ Young Minds CD Studio - CD-R

To create bootable mksysb DVD or CD media, the following requirements should be met on your system:

- ▶ Suitable DVD or CD devices are configured (either via **cfgmgr** or through **smit**)
- ▶ GNU **mkisofs**/cdrecord tools installed
- ▶ Appropriate symbolic links are created
- ▶ Sufficient disk space is provided (for mksysb file generation)

A mksysb CD will only allow a new and complete overwrite install, because it recreates the rootvg and restores the mksysb image. By default, the **mkcd** command creates three file systems in the root volume group for the mksysb image, the CD/DVD file system, and the CD images. This data is then written to media using the **mkcd** command. Specific procedures for using the **mkcd** command with DVD or CDR devices is provided in the redbook *Managing AIX Server Farms*, SG24-6606 and also in the AIX man pages.

11.2.5 The use of TSM in addition to mksysb procedures

The use of TSM for system data backups should not replace the creation of bootable mksysb tapes or disks, but adds a higher level of redundancy for specific elements included in system data backups. For machines in production environments, the device and system configuration data will not change too frequently. Other files, such as /etc/passwd or /etc/hosts may change more frequently.

Note: The backup and restore method described below should **not** include /etc/objrepos, /usr/lpp/ /unix, or ../core files.

As shown in Example 11-5, specific system files can be excluded in the client options file through the use of exclude statements.

Example 11-5 Sample dsm.opt file exclude statements

```
*****
Servername tsm1

DOMAIN "/"
DOMAIN "/usr"
DOMAIN "/var"
DOMAIN "/home"
DOMAIN "/hacmp"
DOMAIN ALL-LOCAL
EXCLUDE /usr/lpp/.../*
EXCLUDE ../objrepos/.../*
EXCLUDE /unix
EXCLUDE ../core
```

The use of TSM backup archive client can compliment standard **mksysb** procedures to lower the frequency of system backups needed for relatively static operating system configurations. The list of files to be included should be carefully compiled and tested. With **mksysb** and TSM, the system restore procedure would be:

- ▶ Boot and install system from latest **mksysb** media
- ▶ Make necessary changes to device configurations (if needed)
- ▶ Restore system data and user data with the TSM client (using -ifnewer option)

Using TSM to backup mksysb images

The **mksysb** command can also re-direct its output to a file. For example, you could designate a particular AIX system with a large NFS-mounted file system to be the destination for multiple mksysb images from many AIX systems. You could

then use the TSM client to backup or archive these files to TSM storage pools, and from there to offsite copy pools. The archive function might be particularly useful here as you could enter an appropriate description including the name of the client and date stamp. If you do this, remember you will have to first recover the correct **mksysb** image to the TSM client, and then use it to network boot and install the required AIX client system for bare metal restore. This will add extra time over having traditional **mksysb** media available (for example, on tape or CD). However it is certainly a valid option for providing redundant backups of the **mksysb** images.

11.2.6 Bare metal restore using mksysb media

To initiate a bare metal recovery from a **mksysb** CD, DVD or tape, the system must be set to boot directly from that device, using appropriate function keys or SMS menus during the initial boot sequence. Refer to the documentation for your specific server model for details of how to do this. Make sure the media is in the correct device, and follow the system prompts to re-load the system from the **mksysb** image.

Tip: You can boot from the AIX installation CD if your **mksysb** media fails to boot. The initial Welcome screen includes an option to enter a maintenance mode in which you can continue the installation from your backup **mksysb** media.

Once the **mksysb** install procedure is initiated, a few prompts for terminal settings and language definitions appear. After these, the bare metal restore of AIX continues without prompts until the system is restored.

11.3 Using NIM for bare metal restore

NIM (Network Installation Management) provides central administration for AIX operating system data backups and restores, and can also be used to distribute AIX system software updates to multiple servers. NIM is an extremely flexible and powerful application which can be implemented to support large numbers of AIX clients. For DR purposes, NIM can be used to recover multiple AIX systems via network connections to the NIM master server.

The NIM application defines one NIM master and many NIM clients. The NIM master is responsible for storing NIM client information (**mksysb** backups and other resources), network information, and specific client machine definitions. NIM allows one machine to act as the master for a given set of clients, but many instances of the NIM master can coexist within a large AIX environment.

TSM compliments a NIM installation by backing up NIM data and providing on and offsite copies of NIM master server data so that NIM clients can still be restored in a full site DR scenario. Figure 11-3 shows the relationship between TSM, NIM, and AIX clients.

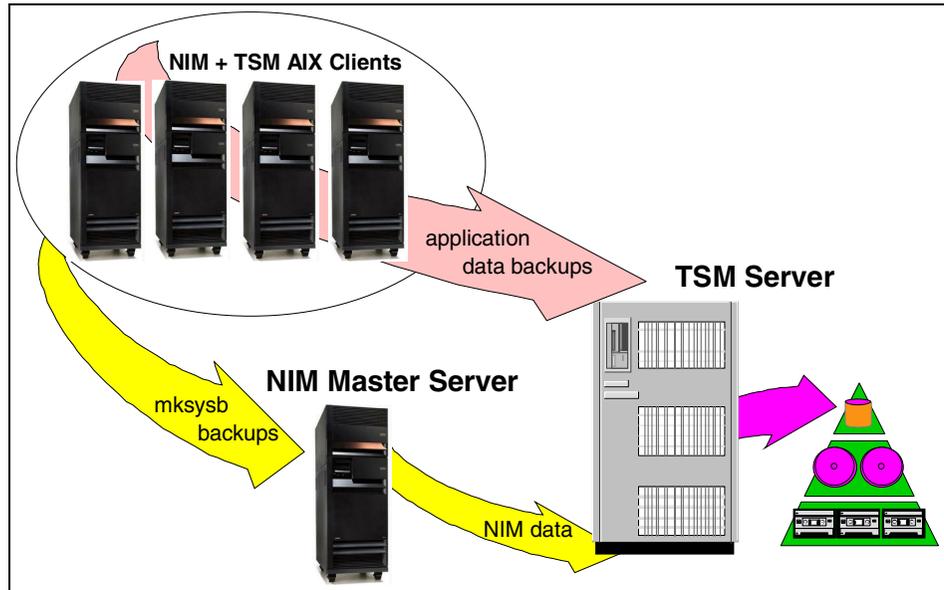


Figure 11-3 AIX client dataflow with NIM and TSM

NIM controls the backup and restore of AIX system data for each NIM client. A TSM client is configured for the NIM master server to backup additional copies and versions of the NIM client resource data. Specific elements within NIM are also backed up to provide DR capabilities to the NIM server itself.

11.3.1 Disaster Recovery using NIM and TSM

In the event a production site is unusable, NIM and TSM can be used together to restore AIX systems and application data. Figure 11-4 shows the recovery process for AIX clients using NIM and TSM.

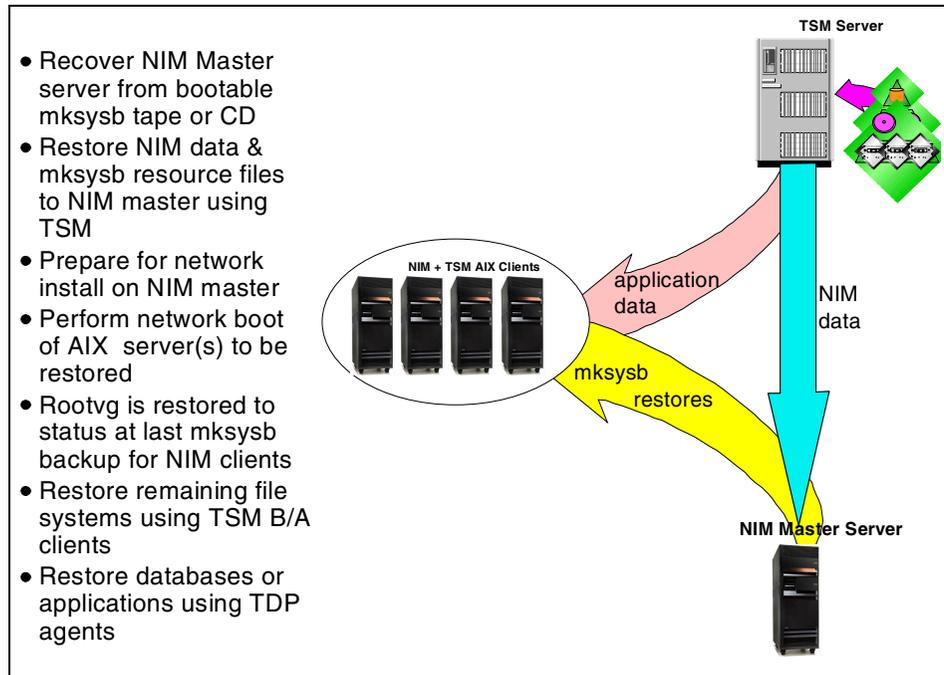


Figure 11-4 Bare metal restore of AIX clients using NIM and TSM

In a large DR scenario, NIM can be used to restore multiple AIX system images at once. Network bandwidth determines the speed by which multiple machines can be restored and must be accounted for in the Disaster Recovery Plan of events. With proper systems design and planning, NIM bare metal restore can outperform other system recovery methods.

11.3.2 Basic NIM Setup

Using NIM is relatively simple for an experienced AIX administrator. The NIM master server is installed and configured and NIM clients are configured prior to performing NIM backup procedures. The following procedure provides the basic process for installing NIM for DR of AIX systems. For more detailed information, consult the AIX system documentation. In our setup here, we define a NIM Master, *crete*, and a NIM client, *sicily*.

NIM Master server planning

On the NIM master server, the rootvg must have at least 4.5 GB of free space, because the NIM installation will create two new file systems, /tftpboot and /export/nim in the root volume group. Space should also be allocated for an additional file system /export/mksysbs/ which is used to store the current NIM

client backups. Manually create the /export/mksysbs journaled file system in the root volume group. Make certain this file system supports the backup image sizes for all current NIM clients, and also make sure there is sufficient space within the root volume group.

TFTP, bootp, and NFS must be enabled on the NIM master server to run NIM. For security and performance reasons, we recommend that you build a dedicated server for NIM in a medium to large AIX environment.

Network hostname resolution is vital for the NIM master to be able to manage the NIM client machines. If these machines are not on the local DNS server, an entry for each NIM client hostname must be made in the /etc/hosts file. Name resolution is also vital when mounting file systems from the resource servers.

NIM Master server setup

Once the current version of AIX is configured and operating system fixes are applied, the future NIM master server is ready for installation. Using CD1 of the AIX operating system disks, use smitty to install the bos.sysmgmt.nim.master filesets. This procedure will take some time and prompt the user for multiple AIX CD volumes.

Once the NIM master filesets are installed, configure the NIM master by using:

```
#smitty nim_config_env
```

The output from the basic NIM master setup screen is shown in Example 11-6.

Example 11-6 Basic NIM Master setup

Configure a Basic NIM Environment (Easy Startup)

Type or select values in entry fields.

Press Enter AFTER making all desired changes.

[TOP]	[Entry Fields]	
Initialize the NIM Master:		
* Primary Network Interface for the NIM Master	[en0]	+
Basic Installation Resources:		
* Input device for installation images	[cd0]	+
* LPP_SOURCE Name	[lpp_source1]	
* LPP_SOURCE Directory	[/export/lpp_source]	+
Create new filesystem for LPP_SOURCE?	[yes]	+
Filesystem SIZE (MB)	[650]	
#		
VOLUME GROUP for new filesystem	[rootvg]	+
* SPOT Name	[spot1]	
* SPOT Directory	[/export/spot]	+

```

    Create new filesystem for SPOT?                [yes]                +
Filesystem SIZE (MB)                             [350]                #
    VOLUME GROUP for new filesystem              [rootvg]             +

Create Diskless/Dataless Machine Resources?      [no]                  +
Specify Resource Name to Define:
  ROOT (required for diskless and dataless)     [root1]
  DUMP (required for diskless and dataless)     [dump1]
  PAGING (required for diskless)                [paging1]
  HOME (optional)                               [home1]
  SHARED_HOME (optional)                       [shared_home1]
  TMP (optional)                                [tmp1]
Diskless/Dataless resource directory            [/export/dd_resource]
  Create new filesystem for resources?          [yes]                +
Filesystem SIZE (MB)                             [150]                #
    VOLUME GROUP for new filesystem              [rootvg]             +
Define NIM System Bundles?                      [yes]                +

Add Machines from a Definition File?            [no]                  +
Specify Filename                                 []

* Remove all newly added NIM definitions        [yes]                +
  and filesystems if any part of this
  operation fails?

```

This command will take some time to build the NIM master file systems and resources. The LPPSOURCE and SPOT resources are built during the NIM master creation. As long as the required space is available in the root volume group, the installation will complete. Next, configure the NIM clients on each host machine.

Note: The LPPSOURCE and SPOT resources are AIX version specific, so for heterogeneous AIX environments, multiple LPPSOURCE and SPOT resources must be configured to support system restores. When implementing a NIM master to support a heterogeneous AIX environment, choose clear names and locations for the additional NIM LPPSOURCE and SPOT files.

NIM Client machine setup

On the NIM client machine, the NIM client software (bos.sysmgt.nim.client fileset) is normally already installed with AIX. If not, install it. As shown in Example 11-7, configure the NIM client definitions by using the following command:

```
#smitty nimit
```

Example 11-7 NIM Client configuration

Configure Network Installation Management Client Fileset

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

```

                                     [Entry Fields]
* Machine Name                       [sicily]
* Primary Network Install Interface   [en0]          +
* Host Name of Network Install Master [crete]

Hardware Platform Type                [chrp]
Kernel to use for Network Boot        [up]          +
IPL ROM Emulation Device              []

/
Comments                              []

Alternate Port Numbers for Network Communications
(reserved values will be used if left blank)
Client Registration                    []
# Client Communications                 []
```

Be certain to select the appropriate NIM master machine name. By default NIM master settings, the basic client definition data will be sent to the NIM master server database automatically. NIM uses NFS and TFTP services to move NIM and systems data between the server and client. Make certain NFS mounts are allowed on the NIM client machine. On the NIM master, check to make sure the new NIM client is defined as a standalone machine by running the `lsnim` command shown in Example 11-8.

Example 11-8 List standalone NIM clients on NIM master

```
root@crete > lsnim -t standalone
sicily    machines    standalone
brazil    machines    standalone
```

Define a NIM client specific mksysb resource

The next step defines a NIM client `mksysb` resource to the NIM master server. In this case, the NIM resource is a `mksysb` image produced on the NIM client, which is sent to the NIM master directory specified below. In Example 11-9, this directory is located in the `/export/mksysbs` file system we created earlier.

Note: Each time a `mksysb` is created for a NIM client, new NIM resource names must be created. For each new `mksysb` resource, the NIM resource name and NIM location of resource name can be identical. These names can not be used to rename previous `mksysb` resource versions because duplicate resource names cannot be stored in the NIM database.

We suggest that NIM administrators use the following naming system to avoid NIM database conflicts:

NIM Resource Name:

<hostname>_<MMDDYYYY>_mksysb

NIM Location of Resource Name:

<hostname>_<MMDDYYYY>_mksysb

To conserve disk space on the NIM master server, we suggest older versions of `mksysb` files be deleted. TSM can be used to backup and archive older `mksysb` versions from the NIM master server as needed.

To define a resource to the NIM master, use the command:

```
#smitty nim_mkres
```

Select `mksysb` from the first list of options and then define the resource to the NIM master as shown in Example 11-9.

Example 11-9 Client mksysb backup definition to the NIM master

Define a Resource

Type or select values in entry fields.

Press Enter AFTER making all desired changes.

```
[TOP]                                     [Entry Fields]
* Resource Name                           [Sicily_07262002_mksysb]
]
* Resource Type                             mksysb
* Server of Resource                       [master] +
* Location of Resource                     [/export/mksysbs/Sicily>
/
Comments                                   []

Source for Replication                     [] +
-OR-
System Backup Image Creation Options:
  CREATE system backup image?             yes +
  NIM CLIENT to backup                     [sicily] +
```

PREVIEW only?	no	+
IGNORE space requirements?	no	+
EXPAND /tmp if needed?	yes	+
Create MAP files?	yes	+
Number of BLOCKS to write in a single output	[]	
#		
(leave blank to use system default)		
Use local EXCLUDE file?	no	+
(specify no to include all files in backup)		
-OR-		
EXCLUDE_FILES resource	[]	+

The EXCLUDE_FILES resource option shown in Example 11-9 can be used to limit the size of the `mksysb` image for each NIM client. This resource file will be located on the NIM master.

First create the `/export/exclude/Exclude_files` file on the NIM master. Then, on the NIM master, define this NIM resource as shown in Example 11-10.

Example 11-10 Defining the Exclude_files resource on the NIM Master

Define a Resource

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

	[Entry Fields]
* Resource Name	[Basic_Exclude]
* Resource Type	exclude_files
* Server of Resource	[master]
+	
* Location of Resource	
[/export/exclude/Exclude_files]	/
Comments	[Standard Exclude List]
Source for Replication	[]
+	

Define a SPOT resource on the NIM Master

The SPOT (Shared Product Object Tree) is a fundamental resource in the NIM environment. It provides the network boot support for all clients. The SPOT varies in size from 100 MB up to, and sometimes in excess of 300 MB depending on the software installed. All device support is installed in the SPOT. SPOTs are used to support all NIM network boot operations including the `bos_inst` operation for AIX bare metal restore.

To define a SPOT resource, use smitty: **NIM -> Perform NIM Administration Tasks -> Manage Resources -> Define a Resource**. Select SPOT.

The SPOT resource is associated along with the mksysb and lppsource resources of the NIM client. The SPOT resource will be installed from images in the image source (AIX media, LPPSOURCE, etc.) and must match the operating system version of the AIX client. The SPOT resource resides on the NIM Master server and is stored in the /export/spot/spot1 directory. The SPOT is used to boot the NIM client during network boot procedures.

When a SPOT is created, network boot images are constructed in the /tftpboot directory of the NIM master server. During the network boot process for bare metal restore, the NIM client uses TFTP to obtain the boot image from the server. After the boot image is loaded into memory at the client, the SPOT is mounted in the clients RAM file system to provide all additional software support required to complete the installation operation. One SPOT resource can be used redundantly for a client and does not need to be recreated with every mksysb backup.

Review NIM machine characteristics on the NIM master

Once the NIM clients and NIM master are configured, access the NIM master database to view the network resources. In a client boot scenario, it is important to ensure that the NIM master recognizes the NIM client (and references associated NIM resource files like mksysb) by the network MAC address on the client. This information is stored in the NIM database and can be displayed by: **smitty NIM -> Perform NIM Administration Tasks -> Manage Machines -> Change/Show Characteristics of a Machine**. Select the appropriate machine name.

In Example 11-11, the network adapter hardware MAC address (000629B91F31) is shown for a NIM client.

Example 11-11 Sample view of NIM machine characteristics

Change/Show Characteristics of a Machine

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[TOP]	[Entry Fields]
Machine Name	[sicily]
* Hardware Platform Type	[chrp] +
* Kernel to use for Network Boot	[up] +
Machine Type	standalone
Network Install Machine State	currently running
Network Install Control State	ready for a NIM operat>
Primary Network Install Interface	

Network Name	network1	
Host Name	[Sicily.almaden.ibm.com]	
Network Adapter Hardware Address	[000629B91F31]	
Network Adapter Logical Device Name	[ent0]	
Cable Type	N/A	+
IPL ROM Emulation Device	[]	+
CPU Id	[000161CF4C00]	
Comments	[]	
Force	no	+

In a Disaster Recovery situation, it is probable that some, if not all, machines or devices are going to be replaced. A well built Disaster Recovery Plan will prepare recovery LAN configurations to support existing hostnames, subnets, and address schemes. For each different server and/or network adapter in a NIM Disaster Recovery environment, the machine hardware will have new MAC addresses, which will need to be manually updated in the NIM machine characteristics screen shown above. Network card MAC addresses are sometimes included in product documentation or are labeled on the actual device. New hardware MAC addresses must also be updated in the /etc/bootptab client entries shown in Example 11-12.

Run NIM Client update procedure on NIM Master

For the standalone NIM client to boot successfully over the network, the bootp request from the NIM client must find its entry in the /etc/bootptab file on the NIM master. This entry usually includes client information about the hostname, bootfile information, and network address configuration as shown in Example 11-12.

Example 11-12 Sample etc/bootptab NIM client entry on NIM master

```
Sicily.almaden.ibm.com:bf=/tftpboot/Sicily.almaden.ibm.com:ip=9.1.38.154:ht=ethernet:ha=000629B91F31:sa=9.1.38.191:sm=255.255.254.0:
```

This file references the client.info file, which specifies exactly which NIM resources (SPOT, LPPSOURCE, MKSYSB) the NIM Master will use to restore the NIM client. The client.info file listing for our NIM client is shown in Example 11-13.

Example 11-13 Sample client.info file output

```
export NIM_NAME=sicily
export NIM_HOSTNAME=Sicily.almaden.ibm.com
export NIM_CONFIGURATION=standalone
export NIM_MASTER_HOSTNAME=Crete.almaden.ibm.com
export NIM_MASTER_PORT=1058
```

```

export NIM_REGISTRATION_PORT=1059
export RC_CONFIG=rc.bos_inst
export NIM_BOSINST_ENV="/../SPOT/usr/lpp/bos.sysmgt/nim/methods/c_bosinst_env"
export
NIM_BOSINST_RECOVER="/../SPOT/usr/lpp/bos.sysmgt/nim/methods/c_bosinst_env -a
hostname=Sicily.almaden.ibm.com"
export SPOT=Crete.almaden.ibm.com:/usr
export NIM_CUSTOM="/../SPOT/usr/lpp/bos.sysmgt/nim/methods/c_script -a
location=Crete.almaden.ibm.com:/export/nim/scripts/sicily.script"
export NIM_BOS_IMAGE=/NIM_BOS_IMAGE
export NIM_BOS_FORMAT=mksysb
export NIM_HOSTS=" 9.1.38.154:Sicily.almaden.ibm.com
9.1.38.191:Crete.almaden.ibm.com "
export NIM_MOUNTS="
Crete.almaden.ibm.com:/export/lpp_source/lpp_source1:/SPOT/usr/sys/inst.images:
dir Crete.almaden.ibm.com:/export/mksysbs/test2:/NIM_BOS_IMAGE:file "

```

The client.info file, which is named after the client hostname, describes the NIM resources to be mounted to the NIM client during a bare metal restore.

To make certain each NIM client is registered with a client.info and bootptab listings on the NIM master, we follow this procedure for each NIM client or group of NIM clients: **smitty nim -> Perform Software Installation and Maintenance Tasks -> Install and Update Software -> Install the Base Operating System on Standalone Clients.**

Select the target to be the standalone client you are configuring (sicily in our case), choose mksysb for Installation TYPE, and choose the mksysb and spot resources most recently created for the NIM client (under MKSYSB and LSPOT). Also select the LPPSOURCE file which matches the NIM client operating system version. Example 11-14 shows the general settings for performing this procedure.

Example 11-14 Sample screen for installing the BOS for Standalone clients

```

* Installation Target                sicily
* Installation TYPE                  mksysb
* SPOT                               sicily_spot
* LPP_SOURCE                         lpp_source1
  MKSYSB                             sicily_mksysb2

  BOSINST_DATA to use during installation  []          +
  IMAGE_DATA to use during installation    []          +
RESOLV_CONF to use for network configuration []          +
  Customization SCRIPT to run after installation []        +
    ACCEPT new license agreements?        [no]         +
  Remain NIM client after install?       [yes]         +
  PRESERVE NIM definitions for resources on [yes]         +

```

```

    this target?

FORCE PUSH the installation?                [no]                +
Initiate reboot and installation now?       [no]                +
-OR-
Set bootlist for installation at the        [yes]               +
next reboot?
Additional BUNDLES to install                []                  +
-OR-
Additional FILESETS to install              []                  +
(bundle will be ignored)

installp Flags
  COMMIT software updates?                  [no]                +
  SAVE replaced files?                      [no]                +
  AUTOMATICALLY install requisite software? [no]                +
  EXTEND filesystems if space needed?       [yes]               +
  OVERWRITE same or newer versions?        [no]                +
  VERIFY install and check file sizes?     [no]                +
ACCEPT new license agreements?              [no]                +
  Preview new LICENSE agreements?          [no]                +

Group controls (only valid for group targets):
  Number of concurrent operations           []
#
  Time limit (hours)                       []
#

Schedule a Job                              [no]                +
YEAR                                        []
#
MONTH                                       []
+#
DAY (1-31)                                  []
+#
HOUR (0-23)                                 []
+#
MINUTES (0-59)                             []                +

```

Since this procedure simply updates client settings on the NIM master server, it is important to make certain the installp flags, initiate reboot, and set bootlist settings are set **exactly** as they are shown above. This procedure ensures that the NIM master will use the appropriate resources during a bare metal restore client. Each time a new **mksysb** resource is created for a NIM client, this procedure must be run on the NIM master.

Check status of the NIM client machine

On the NIM client, verify the NIM client initialization and setup has completed successfully, enter the command shown in Example 11-15.

Example 11-15 Sample output for NIM client verification

```
root@sicily:/: nimclient -l -l sicily
sicily:
  class           = machines
  type            = standalone
  platform        = chrp
  netboot_kernel = up
  if1             = network1 Sicily.almaden.ibm.com 000629B91F31 ent0
  cable_type1    = N/A
  Cstate         = BOS installation has been enabled
  prev_state     = ready for a NIM operation
  Mstate        = currently running
  boot           = boot
  lpp_source     = lpp_source1
  mksysb        = sicily_mksysb2
  nim_script     = nim_script
  spot          = sicily_spot
  iplrom_emu    = /dev/fd0
  cpuid         = 000161CF4C00
  control       = master
```

This command provides a useful way to make sure all appropriate client definitions are made and there are no errors in the NIM client setup. We see here that the client is ready to install its base operating system from the mksysb image called sicily_mksysb2.

Determine IPL-ROM network boot capability for the NIM client

In a bare metal restore situation, the NIM client will have to boot from the network using a network adapter. Some older AIX machines do not support this function without the use of an IPL-ROM diskette. Check your system documentation for details.

All versions of the IPL-ROM can search local devices for an AIX boot image. Only BOOTP-enabled IPL-ROM can use a network interface to search for a remote boot image. Machines manufactured before 1993 do not have BOOTP-enabled IPL-ROM. To determine if a NIM client machine requires IPL-ROM emulation, run the command:

```
#bootinfo -q <network adapter name>
```

If the response is 1, the adapter is network boot enabled. If the response is 0, then you must create the IPL-ROM emulation media (usually floppy disk) for each NIM client, by following the given procedure on the NIM Master.

Insert a formatted diskette into the floppy drive on the NIM master and use **smitty IPL-ROM -> Create IPL-ROM emulation media**. Select the target device to create the IPL-ROM emulation media (diskette or tape), the pathname of the emulation image, and the default value for the boot mode.

IPL-ROM media requirements should be determined during the Disaster Recovery Planning process. During a bare metal restore using NIM, IPL-ROM media must be available to non-BOOTP enabled systems. If Disaster Recovery hardware is provided by a third-party provider, check to find out the hardware type of the recovery hardware. If the hardware support is unknown, create offsite copies of IPL-ROM floppy disks just in case.

11.3.3 AIX bare metal restore using NIM

In the event a NIM client needs to be recovered, the client can boot and install from its most recent **mksysb** image on the network attached NIM server. Based on the client machine's MAC address, the server will TFTP the appropriate boot image to the client and NFS mount the installation resources and **mksysb** data to the client.

Important Reminder: If a new machine or adapter is in use, the network card MAC addresses must also be updated in the NIM master server `/etc/bootptab` client entries and also in the NIM master configuration for client machine characteristics.

Boot procedures for current AIX systems

In a disaster recovery scenario, the AIX machine is booted over the network using standard bootp procedures for the machine type. The client obtains the boot image via TFTP (a thin version of FTP) and begins running a mini-kernel in a file system in RAM.

The hardware platform and kernel type client determine the procedure required to boot the machine over the network. There may be `rspc`, `rs6k`, or `chrp`-based hardware. To determine the platform of a machine, use the `# bootinfo - p` command. To determine the kernel type of a running machine, use the `bootinfo -z` command. The following section describes the boot procedures for an `rspc` type machine. Other system type boot sequences are described in the redbook *NIM: From A to Z in AIX 4.3*, SG24-5524.

Network booting an rspc machine

1. Begin with machine powered off.
2. Bring the machine up to System Management Services using the SMS diskette, or, once the graphic images start appearing on the screen, press the **F1 or 1** key.

Note: For ASCII terminals, press the F4 key as words representing the AIX icons appear. The relevant function key will depend on the type and model of rspc machine, so refer to your User Guide.

For later models of rspc, the functionality of the SMS diskette is incorporated into the firmware, which is accessed by pressing the F1 key.

3. The SMS menu is displayed.
4. Select the **Remote Initial Program Load Setup**.
5. From the Network Parameters screen, select the IP parameters option.
6. Specify the **IP address** of:
 - NIM client
 - NIM Master server
 - Gateway
 - Subnet
7. Select the **network adapter** to be used as the client's boot device.
8. **Ping Test** the network connection (Note: heavy network traffic may cause this command to fail).
9. Select **Multiboot ->Select Boot Device -> Configure 1st Boot Device** from the SMS main menu.
10. Select the network adapter to be used for the network boot. Be sure to select the correct one.

Note: If the recovery system relies on IPL-ROM support, select the floppy drive (fd0) as the first device and the appropriate network adapter as the second device.

11. Exit from SMS menu and commence NIM network boot and installation.

Recover AIX system via network boot

Once the network boot procedure is initiated from the AIX client machine, a few prompts for terminal settings and language definitions appear. After these, the bare metal restore of AIX over the network continues without prompts until the system is restored.

11.3.4 NIM administration

NIM provides many capabilities for operating system data backup, recovery and maintenance. The procedure defined above describes the basic setup for backing up and recovering standalone AIX systems with NIM. All of the procedures provided here can be performed via command line. NIM automation is achievable with use of NIM commands, scripts, and **cron**.

The NIM database should be regularly backed up using the **smit nim_backup_db** fastpath. By default, the NIM database is backed up to `/etc/objrepos/nimdb.backup`. In a recovery scenario, the NIM server is restored using the **smit nim_restore_db** fastpath, assuming the `nimdb.backup` file is already available in its default directory. A NIM database should only be restored to the most recent version used in a NIM environment.

For anyone new to NIM, we highly recommend the following IBM Redbooks and manuals for detailed information about NIM setup, administration, troubleshooting, and customization:

- ▶ *NIM: From A to Z in AIX 4.3*, SG24-5524
- ▶ *AIX 5L Version 5.1 Network Installation Management Guide and Reference*, SC23-4385
- ▶ *AIX Version 4.3 Network Installation Management Guide and Reference*, SC23-4113

Using TSM with NIM

TSM complements NIM with the ability to provide versioning, additional copies, and offsite vaulting of NIM master and client data. TSM copygroup and archive policies will depend on the Disaster Recovery requirements for the particular environment. Retention periods for deleted versions of **mksysb** files should be small if multiple versions already exist on the NIM master server.

To configure a TSM client to support a NIM master environment, we suggest the following procedure:

Install the TSM backup client onto the NIM master server and configure TSM to provide regular backups of the following directories and files:

- ▶ `/export/`
- ▶ `/ftpboot/`
- ▶ `/etc/objrepos/nimdb.backup`
- ▶ `/etc/objrepos/nim_attr`
- ▶ `/etc/objrepos/nim_attr.vc`
- ▶ `/etc/objrepos/nim_object`
- ▶ `/etc/objrepos/nim_object.vc`
- ▶ `/etc/niminfo`

- ▶ /etc/bootptab
- ▶ /etc/export

The NIM objects in the /etc/objrepos directory will be included in any case in the `mksysb` backups of the NIM master server. Depending on the size of the NIM environment, and frequency of NIM `mksysb` resource creations, the NIM master server environment could become huge. TSM can help to limit the space requirements on the NIM master and provide full recovery capabilities for the NIM master environment.

Note: The system data for the NIM master server must be routinely backed up to a bootable `mksysb` tape or CD/DVD for its own disaster recovery. We suggest that users unmount or exclude the /export directory prior to creating the NIM master `mksysb` if this is being backed up by TSM.

11.4 SysBack overview

System Backup and Recovery Version 5.1 for AIX, also known as SysBack, is a separately sold and licensed product for AIX. It provides administrators with a simple way to backup up or recover system data from a SMIT or command line interface. SysBack can be used to create full system backups for AIX systems, and scheduling and scripting capabilities are available to automate system backups.

SysBack can be used to reinstall the system to its original device configuration, including the volume group and logical volume placement on disk and attached devices. SysBack can also be used to clone system images for restoration on different systems in a disaster recovery scenario. SysBack also provides the ability to automate versioning and incremental backups of system data, and operations can be performed over non-local tape resources.

11.4.1 Network Boot installations using SysBack

SysBack can be implemented for bare metal restore capability in a networked AIX environment. For network boot installations, SysBack uses the BOOTP protocol to enable a client to communicate with the SysBack boot server. The client sends a BOOTP request across the network to the server. Then, the SysBack boot server responds with the information that the client needs to contact the server, and subsequently creates access the network boot image.

SysBack can boot through classic network boot or NIM resource boot procedures. If SysBack is using classic boot procedure, the boot server AIX level must match the clients being restored. If using the NIM resource boot procedure,

SysBack utilizes the NIM SPOT and LPPSOURCE resources in combination with SysBack boot/installation operations.

For more information about SysBack, please refer to the following Web site.

<http://sysback.services.ibm.com/>



Windows 2000 client bare metal recovery

This chapter discusses performing bare metal system recovery of a Windows 2000 machine. It describes methods for the crucial step of collecting machine information before a disaster occurs, and how to make this information available in the event of a disaster (for example, using DRM). Finally, we provide a step-by-step basic bare metal recovery procedure. The redbook, *Deploying the Tivoli Storage Manager Client in a Windows 2000 Environment*, SG24-6141, discusses bare metal recovery procedures related to Windows 2000 in detail and should be consulted for more information on this topic.

12.1 Windows 2000 client bare metal restore using TSM

TSM provides integration with documented Windows 2000 APIs for backup of system objects. This allows for a complete bare metal recovery in combination with the boot/ system partition and other data partitions.

However, to restore your system, you need to have previously gathered and saved certain machine-specific characteristics, such as network and disk partition information. Therefore, we discuss in detail methods for collecting this information using operating system utilities and storing this information within DRM. We then provide detailed instructions for recovery of Windows 2000 client, step-by-step, in conjunction with TSM.

Again, *Deploying the Tivoli Storage Manager Client in a Windows 2000 Environment*, SG24-6141, should also be reviewed for a comprehensive review of Windows 2000 and bare metal restore. It discusses additional considerations for restoring an Active Directory, Domain Controllers, and DFS.

12.1.1 Collecting client machine information for Disaster Recovery

Collecting and recording information about your client systems will greatly help your ability to restore a Windows 2000 machine after a disaster to a pre-disaster state. There are several add-on tools, utilities and features built into Windows 2000 that can assist you with information collection, including `msinfo32`, `srvinfo` (found in the Windows 2000 Resource Kit), `diskmap`, the Disk Management interface, `ipconfig`, or writing your own Windows Management Instrumentation application (see the Windows 2000 SDK for more information about WMI). The information that should be collected for the client system should include:

- ▶ Hard drive partition information, for example, number and type of partitions, disk size, drive letters, amount of data used per volume, boot partition, and system directory
- ▶ System hostname
- ▶ TCP/IP networking information, for example, IP address, subnet mask, default gateway, DNS information
- ▶ Windows 2000 Service Pack levels

Scripts or batch files can be used to automate the collection of client information for users not skilled in these kinds of system level commands. Client system information should then be stored offsite for potential use during a Disaster Recovery procedure. Client system information can be imported into DRM (via scripts discussed in 12.1.3, “Storing system information for DRM access” on page 271) or DRM administrators can be given access to system information collected into a text file and backed up by the TSM backup-archive client.

The msinfo32 command

The command **msinfo32** is a feature of Windows 2000 that enables you to collect detailed system hardware and configuration data that can be used for problem determination. This same information can be useful for Disaster Recovery purposes since it documents the original system's configuration. Just storing **msinfo32** information does not prepare you for Disaster Recovery; you still need a predetermined and rehearsed strategy. However, having this information can be one more weapon in your Disaster Recovery arsenal.

msinfo32 has both a GUI and command line interface. This section will focus on the command line interface since it can be scheduled for periodic execution using mechanisms like the Tivoli Storage Manager Backup/ Archive Client scheduler, scripts, or batch files. Generally **msinfo32** will provide most of the required information listed above and is installed by default with Windows 2000. It can be run by entering this at the command-line:

```
C:\Program Files\Common Files\Microsoft Shared\MSInfo\msinfo32.exe
```

You can use **msinfo32** to display configuration information interactively via the GUI interface, or generate a text file report via GUI or command line batch invocation. The text file is organized in *categories* and *subcategories* stanzas which are delimited with [category] and [subcategory] headings. There are many categories and subcategories of information in the report including *Hardware Resources*, *IRQs*, *Memory*, *Drivers*, and *Environment Variables*.

To run **msinfo32** and collect it in a report execute the following commands as shown in Example 12-1.

Example 12-1 How to run msinfo32

```
cd \Program Files\Common Files\Microsoft Shared\MSInfo  
<prompt>msinfo32 /report msinfo32.txt /categories +all
```

A portion of the output for our BMR client machine follows in Example 12-2 — notice the first [System Information] stanza.

Example 12-2 Example of msinfo32 report output

```
System Information report written at: 07/22/2002 10:11:11 AM  
[System Information]
```

```
[ Following are sub-categories of this main category ]
```

```
[System Summary]
```

```
ItemValue
```

```
OS NameMicrosoft Windows 2000 Server
Version5.0.2195 Service Pack 2 Build 2195
OS ManufacturerMicrosoft Corporation
System NameGALLIUM
System ManufacturerIBM
System Modelserver xSeries 330 -[867411X]-
System TypeX86-based PC
Processorx86 Family 6 Model 11 Stepping 1 GenuineIntel ~1128 Mhz
Processorx86 Family 6 Model 11 Stepping 1 GenuineIntel ~1128 Mhz
BIOS VersionIBM BIOS Ver 0.0
Windows DirectoryC:\WINNT
System DirectoryC:\WINNT\System32
Boot Device\Device\Harddisk0\Partition1
LocaleUnited States
User NameGALLIUM\Administrator
Time ZonePacific Daylight Time
Total Physical Memory3,866,068 KB
Available Physical Memory3,564,772 KB
Total Virtual Memory9,660,368 KB
Available Virtual Memory9,223,604 KB
Page File Space5,794,300 KB
Page FileC:\pagefile.sys
```

[Hardware Resources]

You probably should save the whole report, but if there are sections you are sure would not be useful you may want to delete them. If you type `msinfo32 /?` you can see various invocation options. If the `/categories` option did not seem to have the granularity you desired, a script can be used to extract selected information. A sample VBScript script called `msinfoextract.vbs` that can pull subcategories out of the report is provided in “Reducing msinfo32 output using a VBScript” on page 351.

Note: Note, running `msinfo32` and generating a report may take some time. In our case it took about a half minute to generate the report.

On our test client system, we created a batch file to automatically save system information to a text file and then start the TSM backup/archive client. We created an icon on our desktop with a link to this batch file which can be used as a replacement for our TSM backup/archive launch icon. The sample batch file is shown in Example 12-3.

Example 12-3 Batch file for saving machine information and starting TSM client

```
@echo off
echo.
echo SAVING MACHINE INFORMATION FOR DISASTER RECOVERY
```

```

echo.
pause
cd c:\Program Files\Common Files\Microsoft Shared\MSInfo
start msinfo32 /report msinfo32bat.txt /categories +all
cd c:\Program Files\Tivoli\TSM\baclient
start dsm

```

12.1.2 Collect partition and logical volume information with diskmap

The **diskmap** utility is a command-line disk mapping tool that allows users to view the partition and logical volume structures of the disks. It displays a map of the disk and produces a report about the disk's configuration. It provides information about the disk characteristics and a description of each partition and logical volume on the disk. The diskmap tool can be downloaded from Microsoft at:

<http://www.microsoft.com/windows2000/techinfo/reskit/tools/existing/diskmap-o.asp>

An example of using diskmap and the output from our main system volume is given in Example 12-4.

Example 12-4 Example diskmap output

```
<prompt> diskmap /d0 > diskmapd0.txt
```

```

Cylinders  HeadsPerCylinder  SectorsPerHead  BytesPerSector  MediaType
      2212                255                63                512             12
TrackSize = 32256, CylinderSize = 8225280, DiskSize = 18194319360 (17351MB)

```

```
Signature = 0x5a3c8bb3
```

```

      StartingOffset  PartitionLength  StartingSector  PartitionNumber
*                32256          18186061824          63                1

```

```
MBR:
```

*	Starting			Ending			System ID	Relative Sector	Total Sectors
	Cylinder	Head	Sector	Cylinder	Head	Sector			
*	0	1	1	1023	254	63	0x07	63	35519652
	0	0	0	0	0	0	0x00	0	0
	0	0	0	0	0	0	0x00	0	0
	0	0	0	0	0	0	0x00	0	0

12.1.3 Storing system information for DRM access

You can store the **msinfo32** information with regular TSM client backups of your application and systems data. Using this method the **msinfo32** report file is picked up by the normal TSM Backup/Archive Client incremental backup processing.

You can save the **msinfo32** report file in any directory that is not excluded by the TSM Backup/Archive Client. We recommend using a standard location where members of the organization will know to look.

Once there is a backup copy of the **msinfo32** report for this machine in the TSM server you probably want to allow other users, such as the members of your Disaster Recovery team, to access it. This assumes that they have TSM Backup/Archive Client Node IDs registered for them. In Example 12-5, the TSM Backup/Archive Client command line (the GUI can also be used) is used to permit a TSM Client Node ID called DRTEAM to access the **msinfo32.txt** file backed up from the directory `c:\program files\common files\microsoft shared\msinfo`.

Example 12-5 Setting client system information report access in TSM

```
dsmc set acc backup "c:\program files\common files\microsoft  
shared\msinfo\msinfo32.txt" DRTEAM
```

Assuming you have authorized it, a TSM Backup/Archive Client user on another node could restore your **msinfo32** report to a temporary directory on their machine so that it can be referred to an alternative location while the destroyed machine is rebuilt.

12.1.4 Inserting client machine information into DRM

Client machine information can be inserted in DRM for potential use later during disaster recovery. The TSM **INSERT MACHINE** command allows an administrator to store machine characteristics and/or recovery instructions in DRM. Using this method the data can easily be queried and updated by an administrator, without requiring access to storage pool volumes, and the information can potentially be a part of a Tivoli Disaster Recovery Manager **PREPARE** recovery plan file.

We outline below how machine information can be inserted into DRM. We assume that the client machine is already defined as a DRM client machine (using the **DEFINE MACHINE** command). You can use the administrative command line or the Web-based GUI to insert the clients machine information into DRM. Figure 12-1 shows how to use the Web-based GUI to insert client machine information.

Insert Machine : GALLIUM SEQUENCENUMBER

Sequence number

Characteristics

Recovery instructions

Step 1 of 1

Figure 12-1 Using DRM INSERT MACHINE via TSM Web-based interface

Alternatively, you can insert machine information using the command line as shown in Example 12-6.

Example 12-6 Using TSM command line to insert machine information

```
insert machine gallium 1 char="MACHINE OWNER: ITSO SAN JOSE LAB"
```

12.1.5 Using machchar.vbs to insert machine reports into DRM

Since the **INSERT MACHINE** command allows for adding one line of information at a time we can use a script or macro to easily add many lines of information at a time in DRM. Scripts are provided with TSM and can be found in the directory C:\Program Files\tivoli\tsm\server,

A script called machchar.vbs takes a text file and create a TSM macro file of **INSERT MACHINE** commands. This macro can then be run by the TSM administrator to load a DRM MACHINE table with the information. Example 12-7 uses the machchar.vbs script and the machine information report (msinfo32bat.txt) to create a macro (msinfo32bat.mac) that inserts multiple lines of client information automatically. A VBScript is run from the Windows command line as shown.

Example 12-7 Running machchar.vbs to start machine information collection

Usage:

```
cscript machchar.vbs machinename inputmachinefilename outputmacrofilename
```

Example

```
cscript "C:\Program Files\tivoli\tsm\server\machchar.vbs" gallium "C:\Program Files\Common Files\Microsoft Shared\MSInfo\msinfo32bat.txt" "C:\Program Files\Common Files\Microsoft Shared\MSInfo\msinfo32bat.mac"
```

Now, from the TSM administrative command line run this command (Example 12-8):

Example 12-8 Running the TSM macro to insert machine date into DRM

```
macro "c:\Program Files\Common Files\Microsoft Shared\MSInfo\msinfo32.bat.mac"
```

To view machine characteristics added to DRM, use the **QUERY MACHINE** command as shown in Example 12-9.

Example 12-9 Querying DRM for Client Machine Information

```
query machine gallium f=d
```

Figure 12-2 shows the machine client information now made available in DRM for the machine called GALLIUM. We only show the machine summary stanza from our **msinfo32** output information. However, typically you would also include other important stanzas from this output.

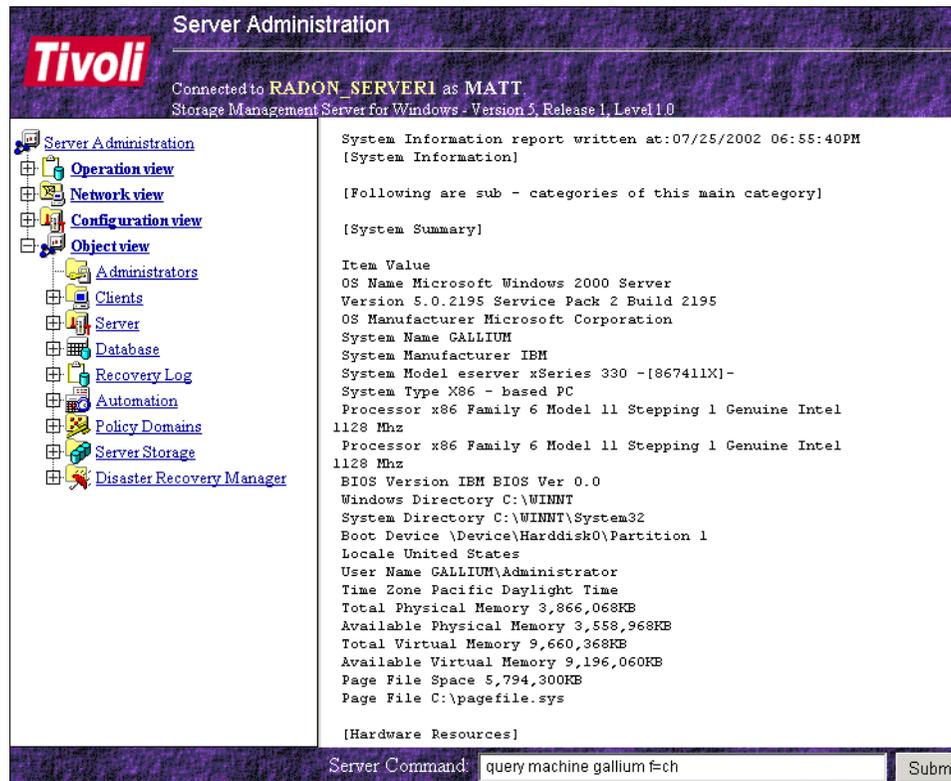


Figure 12-2 DRM machine information output

- ▶ Cluster Database (cluster node only)
- ▶ COM+ database
- ▶ Registry
- ▶ System and boot files
- ▶ System volume

To back up all partitions, including the system partition, TSM requires access to regular files, the directory or the local drive in order to back them up. System objects for Windows 2000 show up in the backup and restore screens of the TSM GUI according to what services are available on the system that TSM is attempting to back up.

Note: From the command line interface for the Backup/Archive client, the **BACKUP SYSTEMOBJECT** command can be used to back up all valid system objects on a Windows 2000 system.

A whole system object can be backed up as an entity, or individual components of the system object can be selected to be backed up. Microsoft indicates that the way to restore objects is to restore the whole System State at the same time from the same backup. This implies that the whole system must be restored to restore just one object. We generally recommend backing up and restoring of the whole system object. For example, restoring the Active directory will also restore the COM+ database, effectively back-leveling it at the same time. Take this into consideration before starting a restore of the system object.

Important: The TSM backup/archive client does not allow an image backup of system objects or an image restore of a boot / system partition (C:) to its original location. In order to perform BMR with image backups, it is necessary to restore the image from another partition. After the image is recovered, the registry would have to be manually copied from SystemDrive\ADSM.SYS to SystemRoot\System32\Config and the restored partition booted. Finally, any other post-BMR steps, such as restoring and activating the Active Directory, would have to be performed.

The directory SystemDrive\ADSM.SYS is created by TSM to store temporary copies of some, but not all, system objects during the backup and restore process. System objects that cannot be directly backed up are first exported from their original location to directories and files under SystemDrive\ADSM.SYS. Then, TSM backs up the directories and files under SystemDriveADSM.SYS.

Important: The SystemDrive\ADSM.SYS directory must not be excluded from backup by the include-exclude list.

The components that use the directory SystemDrive\ADSM.SYS and where the copies of the Windows 2000 object are stored are:

- ▶ Registry contents are in SystemDrive\ADSM.SYS\w2kreg
- ▶ Cluster database are in SystemDrive\ADSM.SYS\clusterdb
- ▶ COM+ database are in SystemDrive\ADSM.SYS\compdb
- ▶ Event logs are in SystemDrive\ADSM.SYS\eventlog

During restore processing of the Registry, a copy of the current Registry is first placed into the directory SystemDrive\ADSM.SYS\w2kreg.sav. This directory may be able to be used to restore the Registry in the event of a restore failure.

Note: You need to be logged in as a user with Backup Data, Restore Data, and Audit Security to backup any local system data (such as registry and system files). For any network-centric system objects, such as the Active Directory, Domain Administrator rights are required.

Example 12-12 provides a recommended include-exclude modification to the TSM client options file in order to backup system objects and avoid errors associated with backing up system objects directly from the system directory.

Example 12-12 Include-exclude modification to backup system objects

```
Exclude "**:\...\system32\config\...\*"
Include *:\adsm.sys\...\*
```

The method for backing up our system object files, boot / system partition (C:), and data partitions via the TSM Backup/Archive GUI is shown in Figure 12-3. Note, system objects are automatically included in the ALL-LOCAL domain for backup.

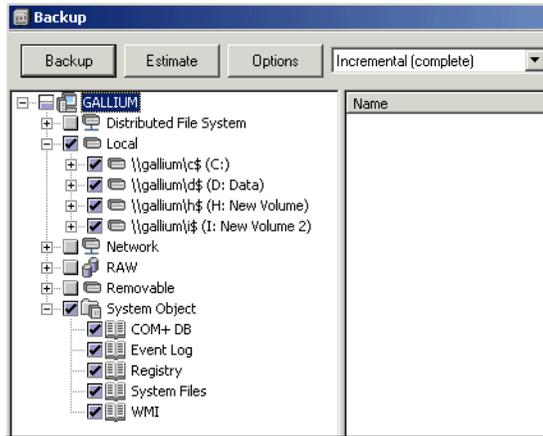


Figure 12-3 Backup Selection of Local Domain and System Objects

To confirm that all of the System Object components were backed up as a single entity, run the command **QUERY SYSTEMOBJECT**. This will display the backup timestamps for each component and it should be obvious if a *rogue* backup of an isolated System Object component exists. Although this is an unlikely event, it is still worth taking the time to check, given the problems that it could cause. Example 12-13 shows the results of the **QUERY SYSTEMOBJECT** command.

Example 12-13 System Object Verification using QUERY SYSTEMOBJECT

```
tsm> q systemobject
Session established with server RADON_SERVER1: Windows
Server Version 5, Release 1, Level 1.0
Data compression forced off by the server
Server date/time: 07/21/2002 21:03:40 Last access: 07/20/2002 21:28:21

      Size      Backup Date      Mgmt Class A/I File
      ----      -
      16,161 B  07/20/2002 19:58:39  DEFAULT    A  SYSTEM OBJECT\GALLIUM\COMPDB
      558,348 B  07/20/2002 19:58:41  DEFAULT    A  SYSTEM OBJECT\GALLIUM\EVENTLOG
      217,248,482 B  07/20/2002 19:57:19  DEFAULT    A  SYSTEM OBJECT\GALLIUM\SYSDFILES
      10,637,492 B  07/20/2002 19:58:40  DEFAULT    A  SYSTEM OBJECT\GALLIUM\REGISTRY
      740,749 B  07/20/2002 19:57:13  DEFAULT    A  SYSTEM OBJECT\GALLIUM\WMI

tsm>
```

12.1.7 Basic Windows 2000 BMR procedure with TSM

Now that we have ensured we have collected all the important information needed for a bare metal restore, we can focus on the specific recovery procedure. We recommend that this procedure be tested in your environment for you to become familiar with this process and so that any steps unique to your environment can be added. The basic Windows 2000 recovery procedure is as follows (note: some of these steps require system restarts):

Table 12-1 Basic Windows 2000 BMR Procedure with TSM

Windows 2000 BMR Procedure with TSM	
1	Boot from Windows 2000 Installation CD.
2	Create a System Partition (C:).
3	Set Computer Name and Complete a Basic Windows 2000 Install.
4	Set IP Address Information to Original System Settings.
5	Test Network Connection.
6	Install Windows 2000 Service Packs to Original Levels.
7	Install Any Drivers Related to Partition Restore Processes (for example, adapter drivers for external disk arrays).
8	Create and Configure Drive Partitions to Original State.
9	Install TSM Client and Configure Pointer to TSM Server.
10	Perform a TSM restore of the boot / system partition (c:).
11	Perform a TSM restore of the Windows Systems Objects.
12	Perform a TSM restore of all other system partitions.
13	Verify that restore is valid.

Our restore procedure assumes that a TSM server is installed with appropriate code levels and patches, that the client node is registered with the server, and that backups of all partitions and system objects have been taken. For the restore, access to a Windows 2000 Installation CD and TSM Backup/Archive Clients Installation CD is also required. The new hardware should be identical to the original hardware. In some cases, it may be possible to restore to hardware which is slightly different, for example, different disk capacities. We recommend thoroughly testing the restore process if different hardware is to be used. The system's hardware components should already be correctly installed and configured. This includes, but is not limited to:

- ▶ System has power to all components
- ▶ Keyboard, mouse and monitor are connected
- ▶ Network controllers are installed and connected to the network
- ▶ Cabling of disk controllers and array controllers is complete
- ▶ Check that the hardware firmware is at the correct level (ideally this firmware should be at the same level as when the backup was taken)

Our environment consists of Windows 2000 Advanced Server (with service pack 2) installed on an IBM Intel-based server. Our TSM server is at V5.1.1 and our TSM Backup/Archive client software is V5.1.0. We have a system partition (C:) and 3 other partitions for data and applications (D:, H:, I:). To simulate a disaster in our environment we deleted the data partitions (D:, H:, I:) and formatted the system partition (C:). Since Windows 2000 generally does not allow formatting of the C partition, we booted with the Windows 2000 CD in emergency repair mode and then formatted the C partition.

Restore procedure

1. Boot from Windows 2000 Installation CD.
 - Windows 2000 is supplied on CD-ROM. The CD is also a bootable entity. Ensure the boot order in your system BIOS are set to CD-ROM first.
 - Windows 2000 will automatically boot if your hard drive is blank, otherwise you will get a prompt on screen for a few seconds saying Press any key to boot from CD. At that point hit a key CD.
2. Create a system partition (C:).
 - Make the Windows 2000 boot partition the same size as the system being recovered. Remember this information should have already been collected using `msinfo32`, found within the [Storage] [Drives] stanzas. The partition should also be formatted using the same file system, either FAT32 or NTFS.
 - The partition should have the same drive letter. The Windows 2000 operating system folder must be named the same as the system being recovered (this will usually be WINNT).
3. Set computer name and complete a basic Windows 2000 install.
 - Configure the server with the same computer name as the system being recovered — GALLIUM in our case.
 - Place the server into a temporary workgroup (use a workgroup name that does not exist). Do not make the server a domain member.
 - Set the time and date correctly.
 - There is no need to install additional services or applications that will not be used by the recovery process. For example, Terminal Server or Macintosh services. Installing such items will only increase the amount of time required to complete the operating system installation and may in fact add unnecessary complications to the restore process. To speed up the install time, you can deselect components which are installed by default, such as Accessories, Internet Information Server, Indexing Service and Script Debugger.
4. Set TCP/IP address information to the original system settings.

- Configure IP address(s) of network cards, including default gateway, subnet mask, and DNS settings. Much of this information can be accessed from the **msinfo32** output collected in advance. Note, if you experienced a site disaster your DNS settings may be different.
 - It is only necessary to perform a minimal install of Windows 2000 with just the networking components required to get the system on the network.
5. Test the network connection.
 - Make sure that you are able to connect over the network to the TSM server. An effective way to test this connection is to **ping** the TSM Server from the client machine.
 6. Install Windows 2000 Service Packs to Original Levels.
 - Install any service packs, patches or drivers that were running on the original system that directly interact with components used by the restore process. For example, network card drivers, disk controller drivers, operating system patches.
 7. Install any drivers related to partition restore processes (for example, adapter device drivers for external disk arrays).
 - In some environments, particularly those using SAN attached disk systems, Fibre Channel HBA device drivers need to be installed to connect to disk partitions that have been externally allocated to the client system.
 - For example, in our environment, an IBM TotalStorage FastT700 Storage Server provided the LUNs where our H: and I: partitions resided. Our host was attached to the SAN with a QLogic HBA. Therefore, we had to install the correct device driver at this point in preparation for the next step.
 8. Create and configure drive partitions to original state.
 - Recreate the same number of disk partitions that were on the original system. Ensure the following partition properties match the original system:
 - Partition type (primary or logical/extended)
 - File system type (FAT32 or NTFS)
 - Disk type (basic or dynamic)
 - Drive letters (D:, E:, F: and so on)

Note: Although, it is desirable to get the partition sizes to match the original system, this is not absolutely crucial. As long as there is sufficient space to restore all the data, it should not affect the success of the recovery.

9. Install TSM Client software and configure pointer to TSM Server.

- For compatibility reasons, ensure that the version of the TSM client used for the restore matches the one used for backup.
 - Install the TSM client into the same path and folder as the original system.
 - Only configure the Backup/Archive client. It is not necessary to configure the TSM Scheduler or Web client.
 - Ensure the TSM client node name is the same as it was on the original system.
 - Configure the client to point to the TSM server.
10. Perform a TSM restore of the boot / system partition (C:).
- Start the TSM Backup/Archive client and select the restore option.
 - Expand File Level and select the drive designated as the Windows 2000 boot / system partition (this is usually the C drive).
 - Select the **Restore Options** button and set the collision options for files that already exist as shown in Figure 12-4.

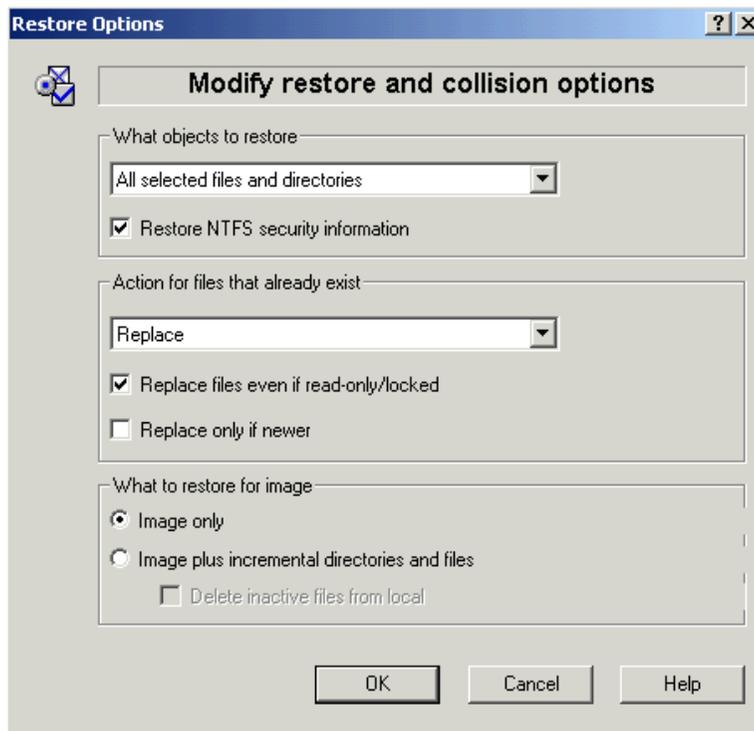


Figure 12-4 Windows 2000 System Partition Restore Options

- Continue the restore process. When prompted, select to restore files to their original location.
- The restore should run cleanly without errors.
- At the end of the restore, **do not select to reboot.**

Tip: Before starting the restore, confirm the consistency of the System Object backup by running the command **QUERY SYSTEMOBJECT** from the Tivoli Storage Manager client command line.

11. If the system is *not* a Domain Controller, perform a TSM restore of the Windows System Objects.
 - Select the System Object for restore. Do not select individual objects for restore.
 - Continue the restore process.
 - At the end of the restore, select to reboot.
12. If the system *is* a Domain Controller:
 - Restart the machine in Directory Recovery mode.
 - Restore the Active Directory. (If you wish to do an authoritative restore, use the NTDSUTIL utility to accomplish this. This is not usually desired in this scenario).
 - Restore any other appropriate system objects, such as the Certificate DB, SYSVOL, and so on. Which of these you recovery will depend on what services you were running on the client prior to its failure.

Note: The event logs are not restored back into the Operating System (that is, they do not become active). They are restored into the folder `\adsm.sys\eventlog`.

To view the logs you should point the event log viewer to the appropriate log file in this folder.

13. Perform a TSM restore of all other system partitions.
 - After the reboot, the system should generally appear to be in its original system state.
 - Restore any other data onto other drives on the system (D:, E:, and so on).
14. Verify that restore is complete and valid.

- Confirm that the system is a domain member again. From the desktop select **My Computer -> Properties -> Network Identification**. Also validate this by logging onto the system with a domain-based account.
- Check Windows 2000 event logs for errors. In particular check for service and device driver failure. If the event logs were restored, check back to see if any errors that may be occurring are new or are just a legacy of the original system.
- Check that all services show the correct status; that is, running, stopped, automatic, manual, and so forth.
- Check that locally defined user and group accounts are present. If the system is part of a domain, check that domain-based accounts are members of local groups (check to see that the domain-based global administrators group is part of the local administrators group).
- Check that print queues are present and functioning.
- Check security on objects, for example, print queues and NTFS files and folders.
- Check that the time zone and system time is correct.
- Ask all users who use the system to check that their profiles have been restored.

Tip: The approximate time for our restore was 3 hours.

12.2 Using third party products for BMR with TSM

The integration capabilities made possible by APIs from Tivoli enable ISVs to more easily adapt their bare metal recovery solutions to IBM software, so joint customers can customize, better secure and extend the functionality of their storage environment. Vendors exploiting these APIs and other IBM Tivoli Storage Manager API functionality include Cristie, St. Bernard Software, UltraBac Software and VERITAS.

An overview of applications from Cristie, Ultrabac, and VERITAS for bare metal recovery is given below. In each case we discuss how TSM can be integrated with these solutions.

12.2.1 Cristie

Cristie is a provider of data storage and backup solutions, integrated with IBM Tivoli Storage Manager to provide a Bare Metal Restore (BMR) solution for Window users with Solaris and Linux versions available shortly. The combined

functionality enables customers to recover a Windows NT or 2000 operating system to a new disk drive using only floppy disks and a disaster recovery backup stored in the IBM Tivoli Storage Manager storage hierarchy.

Cristie PC-BaX Professional Edition is a fully featured 32-bit multi threaded backup and recovery application designed to give optimum performance and portability. It fully supports Windows file systems and backs up all Local Registry, Security and long file names. The user interface, data format and command script language are common across OS/2, Windows 9X/ME, NT and Windows 2000 enabling easy migration between platforms. Cristie provides for free trial downloads from their Web site. More information can be obtained from the Cristie Web site at:

<http://www.cristie.co.uk>

12.2.2 Ultrabac

UltraBac Software integrate its backup and disaster recovery software for Windows based servers and workstations with IBM Tivoli Storage Manager using the TSM API. The combination of these products allows customers to take advantage of UltraBac's client level backup features while maintaining the enterprise level backup functionality already provided by IBM Tivoli Storage Manager. Customers will also be able to exploit backup and recovery features found in UltraBac, such as Image Disaster Recovery.

UltraBac manages the administration of the backup and restore processes, while Tivoli manages the media that these backups are written to. This allows users with existing Tivoli Storage Manager installations to utilize these processes for their backups, as well as providing an avenue for UltraBac users desiring enterprise-level media management facilities to combine the strengths of UltraBac and Tivoli offerings. The device is built on top of the Tivoli Storage Manager API for a robust interface with the Tivoli Storage Manager architecture, and has been qualified for TSM API levels 4.2 and above.

Further information about Ultrabac and trial downloads can be found at:

<http://www.ultrabac.com/>

12.2.3 VERITAS BMR

The product, VERITAS Bare Metal Restore (BMR) is integrated with Tivoli Storage Manager. This product was formally known as The Kernel Group (TKG) BMR. The two products work together to provide an automated process to recover a system in the event of a disaster. The operating system, system configuration files, and all backed-up data are restored. When BMR is used, clients files are backed up using TSM. Before the backup is performed, BMR

saves the client system configuration file. This information allows BMR to completely recover a system from just the TSM backup data. Note that VERITAS BMR supports bare metal recovery for various client operating systems, including Windows 2000, AIX and Solaris.

BMR uses the existing TSM server and adds BMR server, boot server and file server components. The BMR server manages the process to restore a client. It makes the appropriate boot image and file systems available to the client, ensures that the boot server and file server are properly configured and generates a customized client boot script. At restore time, the file server makes the necessary file systems available to the client via NFS. This includes the file system that contains the necessary operating system commands and libraries, the BMR client code and the TSM client code.

The TSM server provides the backup files required to restore the system. In order to restore the system from TSM, all client files must be backed up to the TSM server. The standard TSM client is used to restore all files from the TSM server, including the operating system, applications, configuration files and users files.

Additional information on VERITAS Bare Metal Restore is available on the Web site:

<http://www.veritas.com>



Linux client bare metal recovery

This chapter discusses considerations, preparation, and specific procedures for accomplishing a bare metal recovery of a Linux system. Once the system itself is restored we discuss the use of TSM for recovering the remainder of the data on your system.

13.1 Linux Red Hat bare metal recovery and TSM

There may be several methods to achieve a bare metal recovery of a Linux system. This section focusses on providing basic instructions on how to recover Linux Redhat, in conjunction with TSM using a simple procedure.

Before a bare metal restore procedure is even considered, careful collection of configuration information required for a restore must be performed. In addition, a backup of important boot and system files must also be taken. The following sections cover those topics.

13.1.1 Collecting client machine information for Disaster Recovery

In preparation for DR we recommend that you collect and save offsite the information listed below in addition to normal TSM backup data. Depending on your installation, you could save this information by using the DRM feature, or you could save it as a hardcopy printout stored offsite. The way in which you save the information is not as important as ensuring that it is saved somewhere where it can be easily retrieved along with the client recovery media when a disaster occurs.

Collecting and recording information about your client systems will greatly help your ability to restore your Linux machine after a disaster to a pre-disaster state. There are several commands, utilities and features built into Linux that can assist you with information collection, for example, **fdisk**, the **/proc** directory, **df**, **ifconfig**, and the **/etc/sysconfig** directory. The information that should be collected for the client system should include:

- ▶ Hard drive partition information, for example, number and type of partitions, disk size, mount points, amount of data used per volume, boot partition, and root directory
- ▶ System hostname
- ▶ TCP/IP networking information including IP address, subnet mask, default gateway, DNS information
- ▶ Operating System levels

Scripts or batch files can be used to automate the collection of client information for users not skilled to run these kinds of system level commands. Client system information should then be stored offsite for potential use during a disaster recovery. Client system information can then be imported into DRM (via scripts discussed in Chapter 12, “Windows 2000 client bare metal recovery” on page 267) or DRM administrators can be given access to system information collected into a text file and backed up by the TSM Backup/Archive client.

The following examples provide descriptions and actual output from commands and files built into Linux that can provide crucial input needed for recovery.

The /proc directory

The /proc directory contains virtual files that are windows into the current state of the running Linux kernel. This allows the user to access a large amount of information, effectively providing the kernel's point-of-view within the system. In addition, the user can use the /proc directory to communicate particular configuration changes to the kernel. The /proc/version (shown in Example 13-1) virtual file indicates the level of the Linux operating system and kernel.

Example 13-1 Using /proc/version to view operating system information

```
#cat /proc/version
```

```
Linux version 2.4.2-2 (root@porky.devel.redhat.com) (gcc version 2.96 20000731  
(Red Hat Linux 7.1 2.96-79)) #1 Sun Apr 8 20:41:30 EDT  
2001
```

The /proc/pci virtual file (shown in Example 13-2) lists the PCI devices installed on the system. The information provided by this file may help resolve driver and support issues.

Example 13-2 Using /proc/pci to view PCI device information

```
#cat /proc/pci
```

PCI devices found:

```
Bus 0, device 0, function 0:  
  Host bridge: Intel Corporation 440BX/ZX - 82443BX/ZX Host bridge (rev 3).  
  Master Capable. Latency=32.  
  Prefetchable 32 bit memory at 0xec000000 [0xefffffff].  
Bus 0, device 1, function 0:  
  PCI bridge: Intel Corporation 440BX/ZX - 82443BX/ZX AGP bridge (rev 3).  
  Master Capable. Latency=64. Min Gnt=136.  
Bus 0, device 2, function 0:  
  ISA bridge: Intel Corporation 82371AB PIIX4 ISA (rev 2).  
Bus 0, device 2, function 1:  
  IDE interface: Intel Corporation 82371AB PIIX4 IDE (rev 1).  
  Master Capable. Latency=32.  
  I/O at 0xffff0 [0xffff].  
Bus 0, device 2, function 2:  
  USB Controller: Intel Corporation 82371AB PIIX4 USB (rev 1).  
  IRQ 11.  
  Master Capable. Latency=48.  
  I/O at 0x1000 [0x101f].  
Bus 0, device 2, function 3:
```

```

Bridge: Intel Corporation 82371AB PIIX4 ACPI (rev 2).
  IRQ 9.
Bus 0, device 3, function 0:
  Ethernet controller: Intel Corporation 82557 [Ethernet Pro 100] (rev 5).
  IRQ 11.
  Master Capable. Latency=64. Min Gnt=8.Max Lat=56.
  Prefetchable 32 bit memory at 0xf3dff000 [0xf3dfffff].
  I/O at 0x7c60 [0x7c7f].
  Non-prefetchable 32 bit memory at 0xf3f00000 [0xf3ffffff].
Bus 0, device 16, function 0:
  Token ring network controller: IBM 16/4 Token ring UTP/STP controller (rev
70).
  IRQ 10.
  Master Capable. Latency=48. Min Gnt=6.Max Lat=120.
  I/O at 0x7800 [0x78ff].
  Non-prefetchable 32 bit memory at 0xf3eff700 [0xf3eff7ff].
  Non-prefetchable 32 bit memory at 0xf3eff800 [0xf3efffff].
Bus 1, device 1, function 0:
  VGA compatible controller: S3 Inc. Trio 64 3D (rev 1).
  IRQ 9.
  Master Capable. Latency=32. Min Gnt=4.Max Lat=255.
  Non-prefetchable 32 bit memory at 0xf4000000 [0xf7ffffff].

```

System partition information

Perhaps one of the most important pieces of system information we can obtain is the logical storage and partition information provided by the **fdisk** command as shown in Example 13-3. In the event of a disaster we may need to rebuild partition tables directly from this output. The bare metal restore procedures given later in this chapter require rebuilding of system partitions.

Example 13-3 Using fdisk to output system partition information

```

#fdisk -l

Disk /dev/hda: 255 heads, 63 sectors, 2482 cylinders
Units = cylinders of 16065 * 512 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/hda1    *           1           7        56196   83  Linux
/dev/hda2                8        2482    19880437+   5  Extended
/dev/hda5                8          203    1574338+   82  Linux swap
/dev/hda6           204          895    5558458+   83  Linux
/dev/hda7           896         1587    5558458+   83  Linux
/dev/hda8          1588         1620     265041    83  Linux
/dev/hda9          1621         1653     265041    83  Linux

```

The `/etc/fstab` file provides mapping information between devices and mount points at boot time. In the event of a disaster, the `/etc/fstab` file will help logically rebuild the systems device configuration and help to easily determine where our devices are mounted. Example 13-4 shows the contents of this file on our system.

Example 13-4 viewing mounted file systems in /etc/fstab

```
#cat /etc/fstab
```

LABEL=/	/	ext2	defaults	1 1
LABEL=/boot	/boot	ext2	defaults	1 2
LABEL=/home	/home	ext2	defaults	1 2
/dev/fd0	/mnt/floppy	auto	noauto,owner	0 0
LABEL=/usr	/usr	ext2	defaults	1 2
LABEL=/var	/var	ext2	defaults	1 2
none	/proc	proc	defaults	0 0
none	/dev/pts	devpts	gid=5,mode=620	0 0
/dev/hda5	swap	swap	defaults	0 0
/dev/cdrom	/mnt/cdrom	iso9660	noauto,owner,kudzu,0	0 0
/dev/sda4	/mnt/zip250.0	auto	noauto,owner,kudzu	0 0

The `df` command, shown in Example 13-4, gives us an understanding of how much data is used in our system, within the various filesystems. It maps the partitions, given by the `fdisk` command to the filesystem names and mountpoints in `/etc/fstab`. This output could also help us compare our restored system with the original system to see if the appropriate amounts of data were restored.

Example 13-5 Using df to output filesystem size, percent used, and mount points

```
#df
```

Filesystem	1k-blocks	Used	Available	Use%	Mounted on
/dev/hda8	256667	75388	168027	31%	/
/dev/hda1	54416	3485	48122	7%	/boot
/dev/hda6	5471188	62544	5130724	2%	/home
/dev/hda7	5471188	743808	4449460	15%	/usr
/dev/hda9	256667	16334	227081	7%	/var

Network configuration information

Network communication configuration can be obtained using the `ifconfig` command (Example 13-6) and the `/etc/sysconfig/network` file (Example 13-7). TCP/IP address information, hostnames, and type of network card may be crucial to a restore process.

Example 13-6 Using ifconfig to display network card settings

```
#ifconfig

eth0      Link encap:Ethernet  HWaddr 00:06:29:AE:6E:DD
          inet addr:9.1.38.168  Bcast:9.1.39.255  Mask:255.255.254.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:160 errors:0 dropped:0 overruns:0 frame:0
          TX packets:36 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:100
          Interrupt:11 Base address:0x7c60

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:6 errors:0 dropped:0 overruns:0 frame:0
          TX packets:6 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
```

Example 13-7 Viewing system network settings in the /etc/sysconfig/network file

```
#cat /etc/sysconfig/network

NETWORKING=yes
HOSTNAME=tonga
GATEWAY=9.1.38.1
```

13.1.2 Saving client machine information for DRM

Client machine information can be inserted into DRM for potential use later during disaster recovery. The **INSERT MACHINE** command used with DRM allows an administrator to insert machine characteristics and/or recovery instructions into DRM. Using this method the data can easily be queried and updated by an administrator, and the information can potentially be a part of a Tivoli Disaster Recovery Manager **PREPARE** recovery plan file. This is the detailed Disaster Recovery file that is generated by the DRM **PREPARE** command. In order to include a machine's system information or recovery instructions in the DRM plan file, the client system must be marked as **adsmserver=yes** using the **UPDATE MACHINE** command, as shown in Example 12-10 on page 275.

A script called machchar.vbs (provided on the TSM Windows NT/2000 server platform) takes a text file and create a TSM macro file of **INSERT MACHINE** commands. This macro can then be run by the TSM administrator to load a DRM machine table with the information.

You can include a file with system information with regular TSM client backups of your application and systems data. Using this method, the system information report file is picked up by normal TSM Backup/Archive Client incremental backup processing. This makes it a good approach for storing descriptions of TSM Backup/Archive Client protected machines. The information should also be stored offsite — therefore the recommendation to insert it into DRM.

Specific instruction for inserting machine information into DRM as well as a more detailed discussion of the machchar.vbs script are given in Chapter 12, “Windows 2000 client bare metal recovery” on page 267.

13.1.3 Backing up your Linux system

One of the first preparation steps for a bare metal restore on Linux is to backup operating system files using a common packaging utility, for example, **tar** or **gzip**. The backup of the system files to a removable media device, for example a ZIP drive or SCSI tape drive is done so that these system files can be restored during the bare metal recovery process.

To backup our system files we used an Iomega 250 USB ZIP Drive, which was automatically recognized by the base Red Hat Linux 7.1 install. We used **tar** to archive our boot and system files. The system files we backed up are listed in Table 13-1.

Table 13-1 Linux system directories backed up as tar files

/boot	/dev	/etc	/bin
/sbin	/lib	/usr/sbin	/usr/bin

Generally a good rule of thumb is to backup any directories required at boot time (a quick scan of the boot initialization file `/etc/rc.sysinit` will reveal them). We used **tar** to package these directories as shown in Example 13-8. Once these files were generated we copied them to our mounted zip drive (`/mnt/zip250.0/`).

Example 13-8 Using the tar command to package system directories

```
#tar cvfp bootdirectory.tar /boot
#tar cvfp devdirectory.tar /dev
#tar cvfp etcdirectory.tar /etc
#tar cvfp bindirectory.tar /bin
#tar cvfp sbindirectory.tar /sbin
#tar cvfp libdirectory.tar /lib
#tar cvfp usrsbindirectory.tar /usr/sbin
#tar cvfp usrbindirectory.tar /usr/bin
```

Attention: The amount of system file data you package depends on your system and size of Linux installation. In some cases the amount of system file data you need to backup may exceed the capacity of your removable media device. If this is the case, you may be required to use multiple removable media cartridges or larger capacity cartridges (for example, an external SCSI tape devices).

In addition to backing up system files to a directly attached removable media device, we made a full backup of our system to our TSM server, as shown in Figure 13-1.

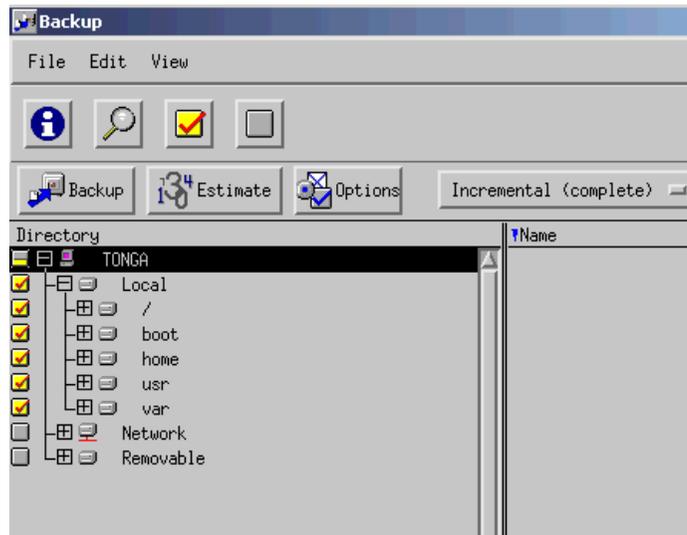


Figure 13-1 Linux full system TSM backup

13.1.4 Preparing a BMR bootable diskette

Finally, we need a rescue diskette or some form of bootable media that will allow us to have a *mini-root* that we can boot to, and execute common commands. Within the *mini-root* environment we will be able to rebuild our partition tables, and restore our system files to a root directory. Although many rescue disk images and tools for creating rescue disks are available, we chose **tomsrtbt**, which can be found at:

<http://www.toms.net/rb>

Full instructions for downloading and creating the **tomsrtbt** rescue disk can be found at that URL. Note, there is also a bootable rescue CD utility found at this URL. We downloaded **tomsrtbt-2.0.103.tar.gz** for a Linux/GNU installation and

unpacked it. To create the bootable rescue diskette, place a diskette into your floppy drive and in the tomsrtbt-2.0.103 directory, run:

Example 13-9 Creating a bootable rescue diskette using tomsrtbt

```
./install.s
```

13.2 Basic Linux BMR procedure with TSM

Now that we have ensured we collected all the important information needed for a bare metal restore, we can focus on the specific recovery procedure. Note that this procedure is given as a suggestion only. We found that it worked in our own specific environment, however it has not been tested for more general application. We strongly recommend that this procedure be tested and modified accordingly in your environment for you to become familiar with this process (so that the steps can be performed efficiently in the event of a disaster) and that any steps unique to your environment can be understood. Parts of this procedure have been adapted from: *Unix Backup and Recovery*, by W. Curtis Preston and Gigi Estabrook; and *Linux Complete Backup and Recovery HOWTO*, by Charles Curley, found at:

<http://www.tldp.org/HOWTO/Linux-Complete-Backup-and-Recovery-HOWTO/>

A basic Linux Red Hat recovery procedure is summarized in Table 13-2.

Table 13-2 Basic Linux Red Hat BMR Procedure with TSM

Linux Redhat BMR Preparation and Restore Procedure with TSM	
1	Boot from rescue media or bootable “mini-root” diskette/CD, e.g., ‘tomsrtbt’.
2	Install any removable media drivers required and mount media device.
3	Use fdisk in bootable “mini-root” partition to rebuild the partition table from pre-disaster information.
4	Mount boot partition in a temporary /root directory.
5	Restore backed up system files to temporary /root directory.
6	Replace ‘/’ (root) directory with temporary /root directory using chroot command.
7	Reboot (to a restored system environment).
8	Install TSM Client and configure pointer to TSM Server.
9	Perform a TSM restore of all other file systems/data.
10	Verify that restore is valid.

Our restore procedure assumes that the TSM server is installed with appropriate code levels and patches, that the TSM client node is registered with the TSM server, and that backups of all partitions and system objects have been taken. For restore, access to a TSM Backup/Archive for UNIX Clients installation CD is also required. The system's hardware components should already be correctly installed and configured. Ideally if you are restoring on new hardware it should be identical (or at least very similar) to the original hardware. This includes, but is not limited to:

- ▶ System has power to all components
- ▶ Keyboard, mouse and monitor are connected
- ▶ Network controllers are installed and connected to the network
- ▶ Cabling of disk controllers and array controllers is complete
- ▶ Hardware firmware is at the correct level (ideally this firmware should be at the same level as when the backup was taken)

Our environment consists of Linux Red Hat 7.1 Server installed on an Intel PC. Our TSM server has TSM version 5.1.1 installed and our TSM Backup/Archive client software was at version 5.1.0. We have a boot/system partition on hd1 and 6 other Linux partitions configured. To simulate a disaster in our environment, we deleted the partition table by running the `dd` command shown in Example 13-10. Subsequently, we were not able to boot our system.

Example 13-10 Deleting the Linux partition table

```
# dd if=/dev/zero of=/dev/hda bs=512 count=1
```

Restore procedure

1. Boot from rescue media or bootable “mini-root” diskette/CD, for example, **tomsrtbt**.

With the tomsrtbt rescue diskette in the floppy drive during startup, the system will boot directly to a root shell with several utilities available.

Tip: The following utilities are made available by the **tomsrtbt** rescue disk.

The 2.0.36 kernel, 3c589_cs, BusLogic, CVF, DEC_ELCP, EEXPRESS, EEXPRESS_PRO, EL2, EL3, EXT2, FAT, FAT32, FD, IDE, IDECD, IDEFLOPPY, IDEPCMCIA, IDETAPE, ISO9660, JOLIET, LOOP, MATH_EMULATION, MINIX, MSDOS, NE2000, NFS, PROC, RAM, SD, SERIAL, SLIP, SMC, SR, TR, ULTRA, VFAT, VORTEX, WD80x3, ah152x_cs, aha152x, aha1542, aic7xxx, ash, awk, badblocks, bdflush, bZip©2, cardbus, cardmgr, cat, ce, ce.help, chatr, chgrp, chmod, chown, chroot, clear, cmp, cp, cpio, cut, date, dd, ddate, debugfs, df, dirname, dmesg, dmsdos, ds, du, dumpe2fs, dutil, e2fsck, eata, echo, egrep, elvis, emacs, extend, false, fdflush, fdformat, fdisk, fdomain, filesize, find, fmt, fsck.ext2, fsck.msdos, fstab, grep, gzip, halt, head, hexedit, hostname, i82365, ifconfig, ifport, ile, init, inittab, insmod, kill, killall5, ksyms, length, less, libc.so.5.4.13, lilo, lilo.conf, ln, loadkeys, login, losetup, ls, lsattr, mawk, memtest, mingetty, miterm, mkdir, mkdosfs, mke2fs, mkfifo, mkfs.minix, mklost+found, mknod, mkswap, mnsed, more, mount, mt, mv, nc, ncr53c8xx, nmclan_cs, nfs, pax, pcmcia, pcmcia_core, pcnnet_cs, ping, plip, ppa, printf, ps, pwd, qllogic_cs, qllogicfas, reboot, reset, rm, rmdir, rmmmod, route, rsh, rshd, script, scsi_info, seagate, sed, serial_cs, setserial, sh, slattach, sleep, slip, snarf, sort, split, stty, swapoff, swapon, sync, tail, tar, tcic, tee, telnet, test, touch, tune2fs, umount, update, vi, vi.help, and wc.

When the system begins initializing several prompts appear to specify display mode, keyboard type, and others. In our case all default settings were suitable. Once this initialization is complete the following prompts, shown in Example 13-11 will appear. You are now in the root directory of the “mini-root shell.

Example 13-11 tomsrtbt root shell login

...Login as root. Remove the floppy. AltF1-AltF4 for consoles.

```
tty1 tomsrtbt login: root
```

```
The default "root" password is "xxxx",
edit /etc/passwd to change it, or edit
settings.s to change it permanently
```

```
Password:
```

```
Today is Pungenday, the 19th day of Bureaucracy in the YOLD 3165
#
```

2. Install any Removable media device drivers required and mount the media device.

Since our ZIP drive has a USB host attachment, we required USB add-ons to be installed in our root shell (similarly this would be an appropriate time to install drivers or add-ons for any other device that you plan to extract your system files from). The appropriate add-ons (usbcore.o, usb-uhci.o,

usb-storage.o) were downloaded from the Internet (on another system in our lab) and copied to a floppy diskette. The diskette with our ZIP drive drivers was mounted, the device drivers were installed, and the ZIP drive was mounted as shown in Example 13-12. The ZIP drive will then be mounted and we will have access to our archived system files.

Example 13-12 Floppy drive mount, driver install, zip drive mount commands

```
#mount /dev/fd0H1440 /mnt

#insmod /mnt/usbcore.o

#insmod /mnt/usb-uhci.o

#insmod /mnt/usb-st~1.o (note file length limitation that results in “~” )

mkdir /mnt/zip

#mount -t vfat /dev/sda4 /mnt/zip
```

3. Use the **fdisk** command in the bootable *mini-root* to rebuild the partition table (based on **fdisk -l** partition information gathered before the disaster).

We need to repair (or rebuild) our root disk. The **fdisk** utility can now be used to repartition the root disk to its exact format before the disaster. In order for us to do this we must have the output from the **fdisk -l** command taken before the disaster. Luckily we collected this information in Example 13-3. While **fdisk** may seem cryptic at first, it is actually quite simple. Essentially **fdisk** is used to rebuild the partition table, character by character. You just need to be careful when following each instruction below.

Example 13-13 Using fdisk to repartition the root disk to pre-disaster state

Original **fdisk -l** output is given here again for reference.

Disk /dev/hda: 255 heads, 63 sectors, 2482 cylinders
Units = cylinders of 16065 * 512 bytes

Device	Boot	Start	End	Blocks	Id	System
/dev/hda1	*	1	7	56196	83	Linux
/dev/hda2		8	2482	19880437+	5	Extended
/dev/hda5		8	203	1574338+	82	Linux swap
/dev/hda6		204	895	5558458+	83	Linux
/dev/hda7		896	1587	5558458+	83	Linux
/dev/hda8		1588	1620	265041	83	Linux
/dev/hda9		1621	1653	265041	83	Linux

We can now begin the process of repartitioning the root disk.

```
# dd if=/dev/zero of=/dev/hda1 bs=1024 count=1
```

```
# fdisk /dev/hda
```

```
Command (m for help): p
```

```
Disk /dev/hda: 255 heads, 63 sectors, 2482 cylinders  
Units = cylinders of 16065 * 512 bytes
```

```
Device Boot Start End Blocks Id System
```

Note, this printout indicates no partitions are currently defined.

```
Command (m for help): n
```

```
Command action
```

```
e extended
```

```
p primary partition (1-4)
```

```
p
```

```
Partition number (1-4): 1
```

```
First cylinder (1-2482): 1
```

```
Last cylinder or +size or +sizeM or +sizeK ([1]-2482): 7
```

```
Command (m for help): n
```

```
Command action
```

```
e extended
```

```
p primary partition (1-4)
```

```
e
```

```
Partition number (1-4): 2
```

```
First cylinder (8-2482): 8
```

```
Last cylinder or +size or +sizeM or +sizeK (8-2482): 2482
```

```
Command (m for help): n
```

```
Command action
```

```
l logical (5 or over)
```

```
p primary partition (1-4)
```

```
l
```

```
First cylinder (8-2482): 8
```

```
Last cylinder or +size or +sizeM or +sizeK (8-2482): 203
```

```
Command (m for help): n
```

```
Command action
```

```
l logical (5 or over)
```

```
p primary partition (1-4)
```

```
l
```

```
First cylinder (204-2482): 204
```

```
Last cylinder or +size or +sizeM or +sizeK (8-2482): 895
```

```
Command (m for help): n
```

```
Command action
```

```

    l logical (5 or over)
    p primary partition (1-4)
]
First cylinder (896-2482): 896
Last cylinder or +size or +sizeM or +sizeK (896-2482): 1587

Command (m for help): n
Command action
    l logical (5 or over)
    p primary partition (1-4)
]
First cylinder (1588-2482): 1588
Last cylinder or +size or +sizeM or +sizeK (1588-2482): 1620

Command (m for help): n
Command action
    l logical (5 or over)
    p primary partition (1-4)
]
First cylinder (1621-2482): 1621
Last cylinder or +size or +sizeM or +sizeK (1621-2482): 1653

Command (m for help): p

```

Disk /dev/hda: 255 heads, 63 sectors, 2482 cylinders
Units = cylinders of 16065 * 512 bytes

Device	Boot	Start	End	Blocks	Id	System
/dev/hda1		1	7	56196	83	Linux
/dev/hda2		8	2482	19880437+	5	Extended
/dev/hda5		8	203	1574338+	83	Linux
/dev/hda6		204	895	5558458+	83	Linux
/dev/hda7		896	1587	5558458+	83	Linux
/dev/hda8		1588	1620	265041	83	Linux
/dev/hda9		1621	1653	265041	83	Linux

We must also change the partition type of /dev/hda5 to "Linux Swap" using the 't' option and set the bootable flag on /dev/hda1 with the 'a' option.

```

Command (m for help): t
Partition number (1-5): 5
Hex code (type L to list codes): 82
Changed system type of partition 5 to 82 (Linux swap)

Command (m for help): a
Partition number (1-5): 1

Command (m for help): p

```

Disk /dev/hda: 255 heads, 63 sectors, 2482 cylinders
Units = cylinders of 16065 * 512 bytes

Device	Boot	Start	End	Blocks	Id	System
/dev/hda1	*	1	7	56196	83	Linux
/dev/hda2		8	2482	19880437+	5	Extended
/dev/hda5		8	203	1574338+	82	Linux swap
/dev/hda6		204	895	5558458+	83	Linux
/dev/hda7		896	1587	5558458+	83	Linux
/dev/hda8		1588	1620	265041	83	Linux
/dev/hda9		1621	1653	265041	83	Linux

Finally we need to write the disk label to the disk:

```
Command (m for help): w
The partition table has been altered!
```

```
Calling ioctl() to re-read partition table.
hda: hda1 hda2 < hda5 hda6 hda7 hda8 hda9 >
hda: hda1 hda2 < hda5 hda6 hda7 hda8 hda9 >
Syncing disks.
#
```

4. Mount boot partition in a temporary /root directory.

Now that we have a properly partitioned and labelled system disk, we can create a file system on which we can mount a temporary /root directory. This process is shown in Example 13-14.

Example 13-14 Creating a file system and /root directory

```
#mke2fs /dev/hda1

#mkdir /root

#mount /dev/hda1 /root
```

5. Restore backed up system files to temporary /root directory.

We are now ready to restore our system files that we had previously backed up using **tar** to our zip drive. For now, we will restore all those files to the /root directory as shown in Example 13-15.

Example 13-15 Restoring our system files to the /root directory

```
#cd /root

#tar xvf /mnt/zip/bootdirectory.tar
#tar xvf /mnt/zip/devdirectory.tar
#tar xvf /mnt/zip/etcdirectory.tar
```

```
#tar xvf /mnt/zip/bindirectory.tar
#tar xvf /mnt/zip/sbindirectory.tar
#tar xvf /mnt/zip/libdirectory.tar
#tar xvf /mnt/zip/usrbindirectory.tar
#tar xvf /mnt/zip/usrbindirectory.tar
```

6. Replace '/' (root) directory with temporary /root directory using **chroot** command.

We now need to restore the lilo boot block. LILO (Linux LOader), places itself in a boot sector of your hard disk. LILO installs a boot loader that will be activated the next time you boot. The **chroot** command here takes two arguments: the first is a directory name, which will cause that directory to become the new root directory, that is, the starting point for path searches of pathnames beginning with '/'. The second argument is a command which will then be run. To run the **lilo** command without causing errors, we use **chroot** followed by the **lilo** command. This is shown in Example 13-16.

Example 13-16 Replacing the '/' directory and restoring the lilo boot block

```
# chroot /root /sbin/lilo
```

7. Reboot (to a restored system environment).

The system should now be rebooted and any rescue diskette or CD removed. Depending on your original system install and the amount of system files you were able to restore, you may see some errors upon booting. These errors may not be critical at this point.

8. Install TSM Client software and configure it to point to the TSM Server.

From the command line, perform an install of the TSM Backup/Archive Client for UNIX systems. In Example 13-17 we are installing the TSM client from CD.

Example 13-17 Performing a command line install of a TSM Backup/Archive Client

```
#mount /dev/cdrom /mnt/cdrom
```

```
#rpm -i TIVsm-API.i386.rpm
```

```
#rpm -i TIVsm-BA.i386.rpm
```

Note, at this point, configuration of your `dsm.opt` file and `dsm.sys` file should be configured to point to the TSM Server your client is registered with.

9. Perform a TSM restore of all other file systems and data.

You can proceed with restoring the remainder of your system from the TSM Backup/Archive client. What you restore depends on what applications or

data you had on your original system. We restored the following directories as shown in Example 13-18.

Example 13-18 Performing a command line restore of files systems and data

```
#cd /opt/tivoli/tsm/client/ba/bin

#dsmc restore /home/-subdir=yes -replace=all -tapeprompt=no
#dsmc restore /misc/-subdir=yes -replace=all -tapeprompt=no
#dsmc restore /usr/-subdir=yes -replace=all -tapeprompt=no
#dsmc restore /var/-subdir=yes -replace=all -tapeprompt=no
```

Note, if you want to restore files from your /opt directory be sure to exclude the /opt/tivoli directory to avoid termination of your TSM client session.

10. Verify that the restore is valid.

- Check event/error logs for issues. In particular check for process and device driver failure.
- Check that locally defined user and group accounts are present.
- Check that print queues are present and functioning.
- Check security and permissions on files and directories.
- Check that the time zone and system time is correct.
- Ask all users who use the system to check that their profiles have been restored.
- Test applications for operation.

Tip: The approximate time for our restore was 3 hours.



Putting it all together

In this final chapter, we will bring together many of the ideas we have covered so far to give you some guidance on designing real solutions. This chapter will include:

- ▶ Some TSM implementation scenarios, describing how TSM solutions can provide functionality and availability at the different DR tier levels
- ▶ Two case studies which show how real customers mapped to a TSM design
- ▶ Summary of all the processes and some pointers on application backup

14.1 TSM implementation scenarios

Several factors contribute to the way TSM is implemented for Disaster Recovery. Scenarios range from high availability mirrored site models to remote TSM storage management implementations. The scenarios listed in this section represent a range of environments and technologies. The flexibility of TSM allows a wide variety of implementation methods and architectural models. These scenarios are intended to inspire ideas for your individual Disaster Recovery environment. Actual implementation of multi-site TSM environments requires the careful planning and technical expertise provided by consulting organizations such as IBM Global Services.

14.1.1 Storage networking considerations

The common thread between all electronic vaulting scenarios is the use of advanced networking technology to move data from one site to another. The following TSM implementation scenarios offer flexibility in the use of these technologies for site to site data transfer. Figure 14-1 shows the general characteristics of IP, DWDM, and SAN storage networking implementations across long distances.

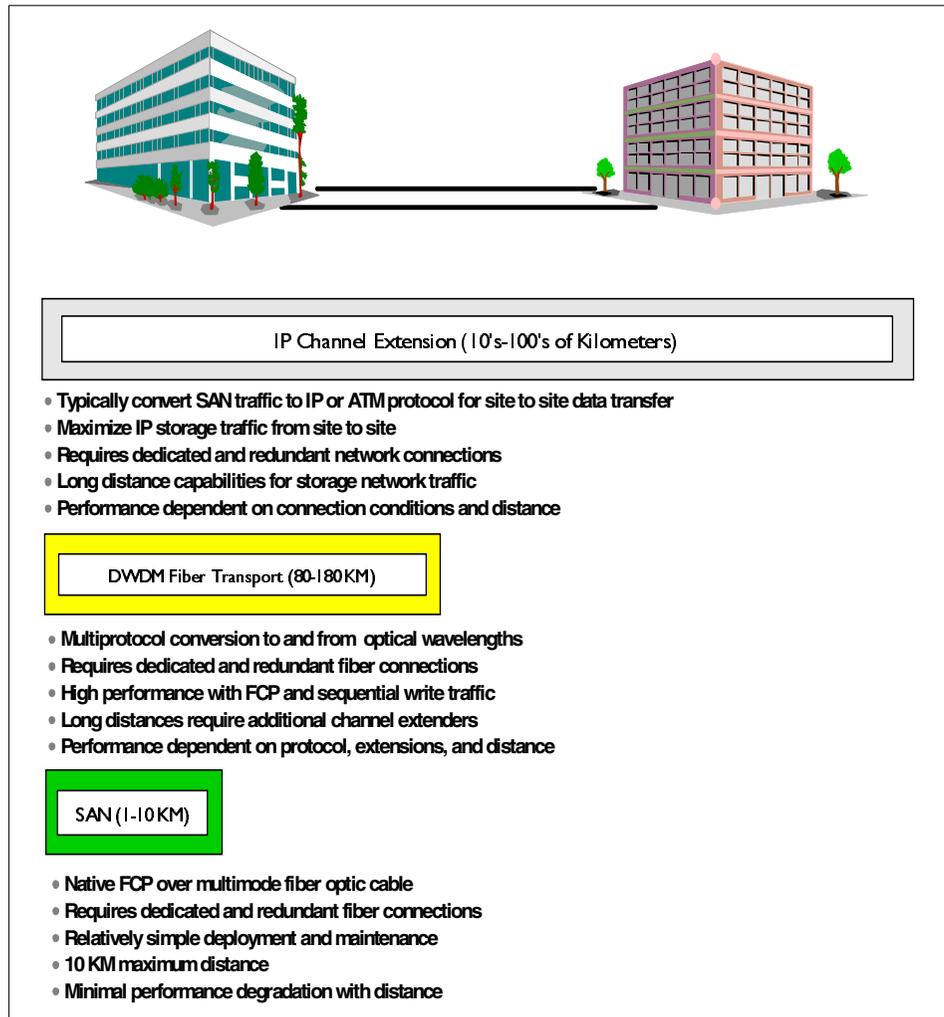


Figure 14-1 Storage network extension technologies

IP and ATM wide area networks have been in use for a long time, but an emerging trend is the use of channel extenders to transfer storage traffic to these WAN environments. Companies such as CNT offer the ability to translate FCP and IP storage network traffic across long distances using IP or ATM protocols. The protocol conversion is transparent to a SAN network fabric. For very long site to site distance requirements, this technology clearly provides a unique capability. Telecommunications network connections ranging from T-1 to OC-48 can be used with channel extenders. While latency over long distances can be an issue, availability of network connections is usually not.

The financial industry has used DWDM (dense wavelength division multiplexing) technology over fibre optic connections for years to transfer large amounts of data from site to site. We see this technology moving quickly into the mainstream due to the high speed and bandwidth characteristics of DWDM, the abundance of fibre in certain metropolitan areas, and rapid advancement of DWDM devices from CNT, Cisco, and IBM. Fibre optic connections can be owned and maintained by companies or leased from managed service providers. Managed fibre price varies significantly from one geography to another. DWDM provides very high bandwidth capabilities by merging dozens of protocol streams into light wavelengths, which are multiplexed into a fibre optic signal and de-multiplexed at the receiving end, back to the original protocol streams. The protocol conversion is transparent to a SAN network fabric. Signal degradation stems from protocol characteristics, distance, and physical disruptions to line placement. With increased demand, availability of fibre lines will continue to improve.

For campus environments, native SAN technology provides a robust and easily managed solution for site to site data transfer. Using long wave GBICs and 9 micron single mode fibre cable, FCP traffic can be routed up to 10 km across optical fibre SAN connected devices. By using repeaters (switches with long wave ports) and dedicated fibres, we can achieve distances of about 40 km, however access points to the fibre must be available approximately every 10 km, which is a limiting factor in many cases. Nonetheless, the native SAN solution provides a scalable and cost effective solution, where cable can be physically run site to site without easement or access issues.

Advanced storage networking technologies are discussed in detail in the IBM Redbook *Introduction to SAN Distance Solutions*, SG24-6408. Decision criteria for these types of networking technologies include cost, management complexity, additional skills development, bandwidth requirements, business requirements for availability, and distance requirements from site to site for disaster recovery.

14.1.2 Remote disk copy and mirroring considerations

Consideration must also be given to advanced copy functions for disk subsystems. For PPRC implementations, ESCON connectivity must be made between each disk subsystem. Direct ESCON connectivity ranges to 3 km. ESCON traffic can be routed through IP channel extenders or DWDM+ fibre infrastructure to achieve longer remote disk copy distances between sites. Other disk subsystems support remote mirroring over native Fibre Channel, such as the IBM TotalStorage FAS*T*product line.

Note: Synchronous mirroring across long distances using PPRC can introduce performance degradation for the primary TSM database transactions, depending on I/O and network conditions. Instead, asynchronous remote mirroring achieved through FlashCopy + PPRC may reduce the write penalty and latency issues for remote disk operations. IBM currently supports PPRC synchronous copy implementations up to 103 km.

The PPRC Extended Distance (PPRC-XD) operation can also be implemented for non-synchronous writes over very long distances through channel extenders. TSM database performance will not significantly degrade using this kind of implementation.

14.1.3 Tier 6 mirrored site implementation with TSM

Enterprise scale high availability is ultimately achieved through the implementation of mirrored production sites. Each transaction is processed in a single commit scope between each site, so logically there is one live version of all production data on each site at any given point in time. In the event of a single site disaster, recovery time and data loss are minimal. TSM complements this kind of environment with the ability to backup and retain versions of the operating environment. Data corruption is still a possibility, even for a high availability environment and maintaining availability at each site relates intricately to storage management. Figure 14-2 illustrates the use of TSM in a mirrored dual site environment.

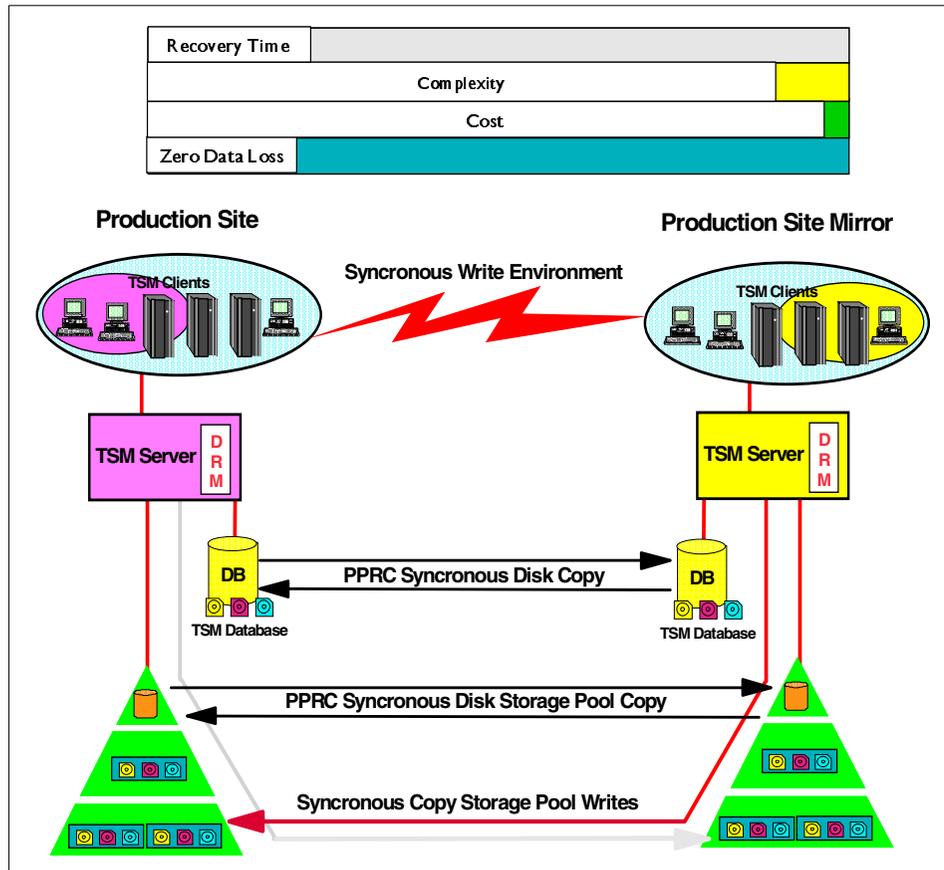


Figure 14-2 Tier 6 TSM implementation, mirrored production sites

TSM is implemented at each production facility, however each TSM server handles a unique workload for backup and recovery operations. Together, the full site is backed up and protected. Additionally, TSM database volumes and disk storage pools are mirrored synchronously between sites to respectively provide additional redundancy for each TSM environment and critical data. In the event of a site failure, DRM data and the TSM database mirror can be used to restore the lost TSM production instance in the remaining site. On each site, the TSM infrastructure must be sized to support capacity and workload requirements for the production and recovery environments.

A mirrored site environment typically staffs both production sites 24x7. The TSM environment can be automated to a high degree, however storage administration staff must be available for both sites at any given time to monitor operations, troubleshoot problems, and restore the lost TSM environment in the case of a single site disaster.

14.1.4 Tier 5 hotsite with TSM electronic vaulting

A recovery hotsite is typically staffed 24x7 with critical staff to maintain operations. In this implementation scenario, TSM simultaneously sends backup data to primary storage pools on the production site and copy storage pools on the recovery site. TSM database volumes and disk storage pool data is mirrored synchronously to the hotsite to provide maximum availability for the TSM environment and critical data. Figure 14-3 illustrates this kind of environment.

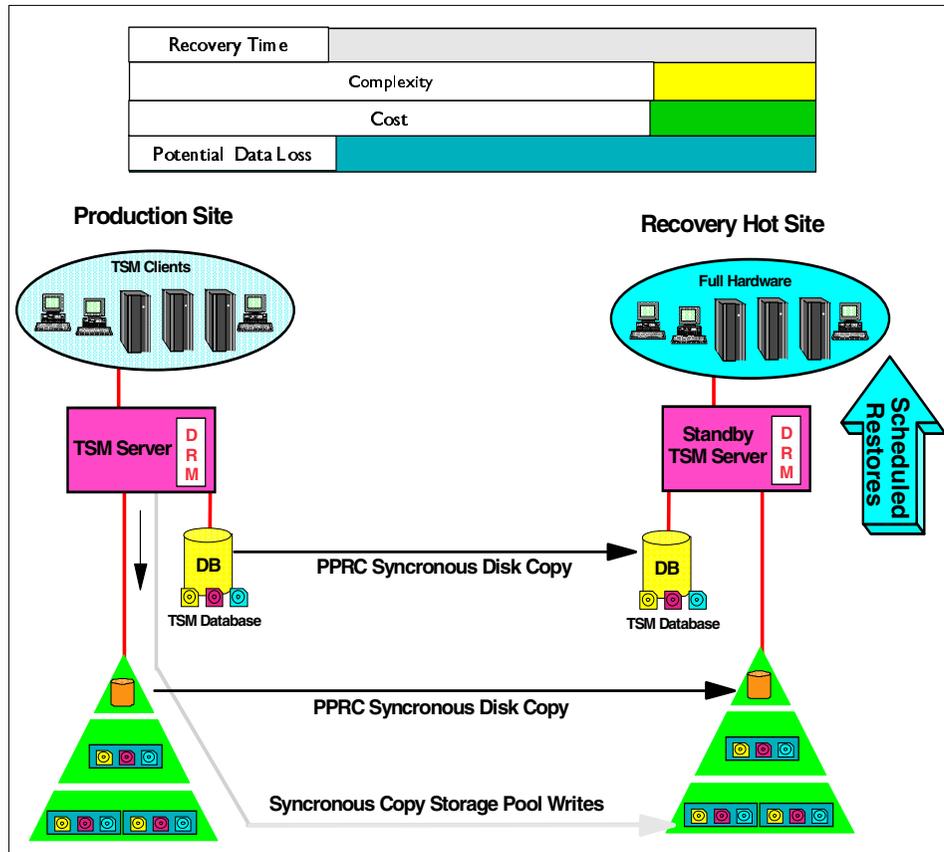


Figure 14-3 Tier 5 hotsite with TSM electronic vaulting

In the event of a production site disaster, TSM can be quickly restored using DRM and the mirrored copy of the TSM database. Primary disk pool volumes and copy storage pool volumes can then be accessed at the recovery site and hotsite systems can be restored according to the Disaster Recovery Plan.

In many hotsite environments, client hostname, LAN, and SAN address schema mirror the production site. As a result, the hotsite pre-configured instance of TSM can use a split mirror production TSM database and perform periodic restores of critical data to critical hotsite systems. This type of procedure should be performed while the primary site TSM environment is relatively inactive, such as during peak day time production hours when production backups are infrequent. Scheduled restores of critical systems in a hotsite keeps the hotsite staff busy, minimizes critical system restore time, and continually tests the functionality of both TSM and the Disaster Recovery Plan.

14.1.5 Tier 5 hotsite with clustered TSM and electronic vaulting

In addition to the hotsite scenario explained above, TSM servers can be clustered between sites to maintain high availability of the TSM environment. Figure 14-4 shows the general layout of this environment.

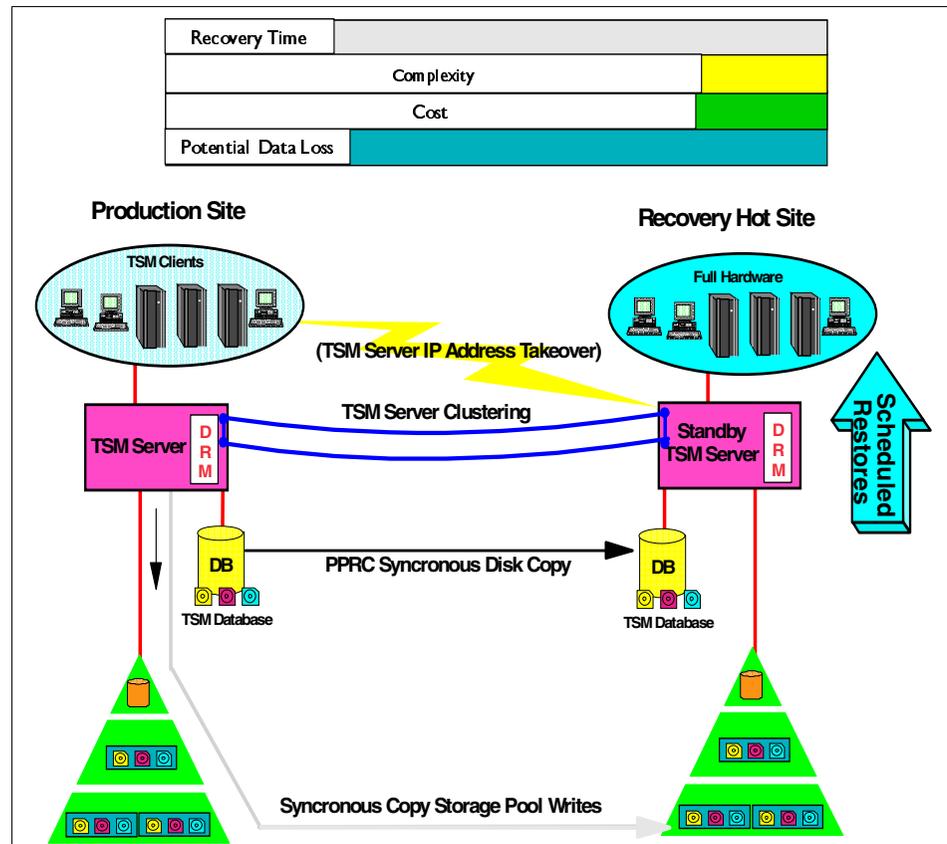


Figure 14-4 Clustered TSM Servers, with site to site electronic vaulting

In the event of a production site failure or the loss of the primary TSM server, the following sequence of events occurs:

1. Primary TSM Server fails or production site disaster.
2. Clustering software performs network address takeover.
3. DNS updates for WAN environment (can be scripted and automated).
4. TSM database mirror volumes are varied on for standby TSM server.
5. DRM scripts are invoked by cluster software to automate the restore of TSM.
6. TSM hotsite takes over backup operations over WAN and recovers data from storage pools as needed.
7. If primary site is lost, TSM can then restore data to recovery site clients with transition to the recovery site LAN environment.

While clustering software adds tremendous value to a TSM environment, we advise careful testing and change management procedures for the TSM environment when using clustering. Simple changes to system environments or network congestion can accidentally trigger system failover. WAN infrastructures must also be designed to handle peak production workloads.

14.1.6 Tier 4 warmsite TSM implementation with electronic vaulting

Warmsite environments provide equipment for critical systems and ample space and infrastructure for to reconstruct a full production environment. Compared to hotsites, recovery times are longer for critical systems, however the cost and complexity of the implementation are significantly lower. Figure 14-5 illustrates a warmsite environment using TSM and electronic vaulting methods for disaster recovery.

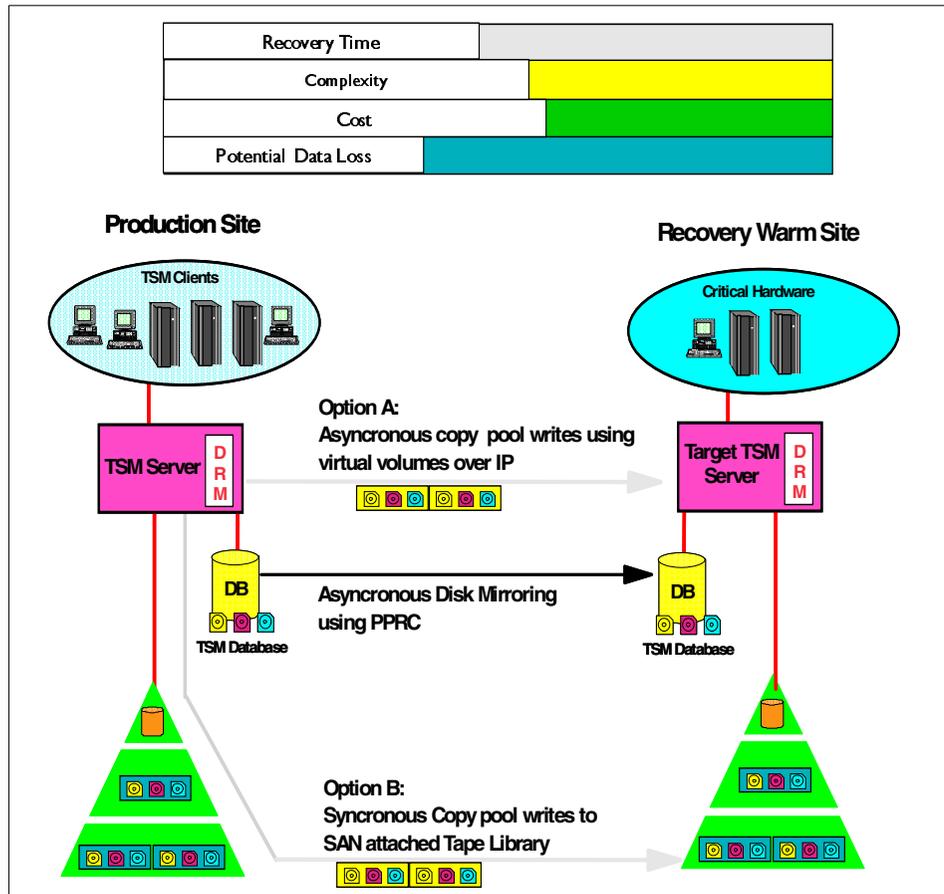


Figure 14-5 Asynchronous TSM DB mirroring and copy storage pool mirroring

Using PPRC disk copy functions, TSM databases can be asynchronously mirrored or copied to remote disk storage infrastructure.

Using server-to-server communications and virtual volumes, production data can be archived to the recovery site TSM environment. The target TSM server is seen as a device class for the production site TSM server. Virtual volumes are transmitted using IP communications only. To recover the client data, the TSM server would need to be recovered as an additional instance on the recovery site TSM server, so the virtual volumes could then be reclaimed and restored to clients at the recovery site.

The second option is to write copy storage pool data synchronously to the offsite tape library using a SAN infrastructure. In the event of a disaster, the TSM DRM

output can be used to automate TSM server recovery at the warm site and copy storage pool data can be recovered and restored to critical systems.

14.1.7 Tier 3 warmsite TSM implementation

In an environment where distance or cost limit electronic vaulting capabilities, TSM can still support warmsite recovery operations with basic IP services and a combination of electronic and manual vaulting. Figure 14-6 illustrates this scenario.

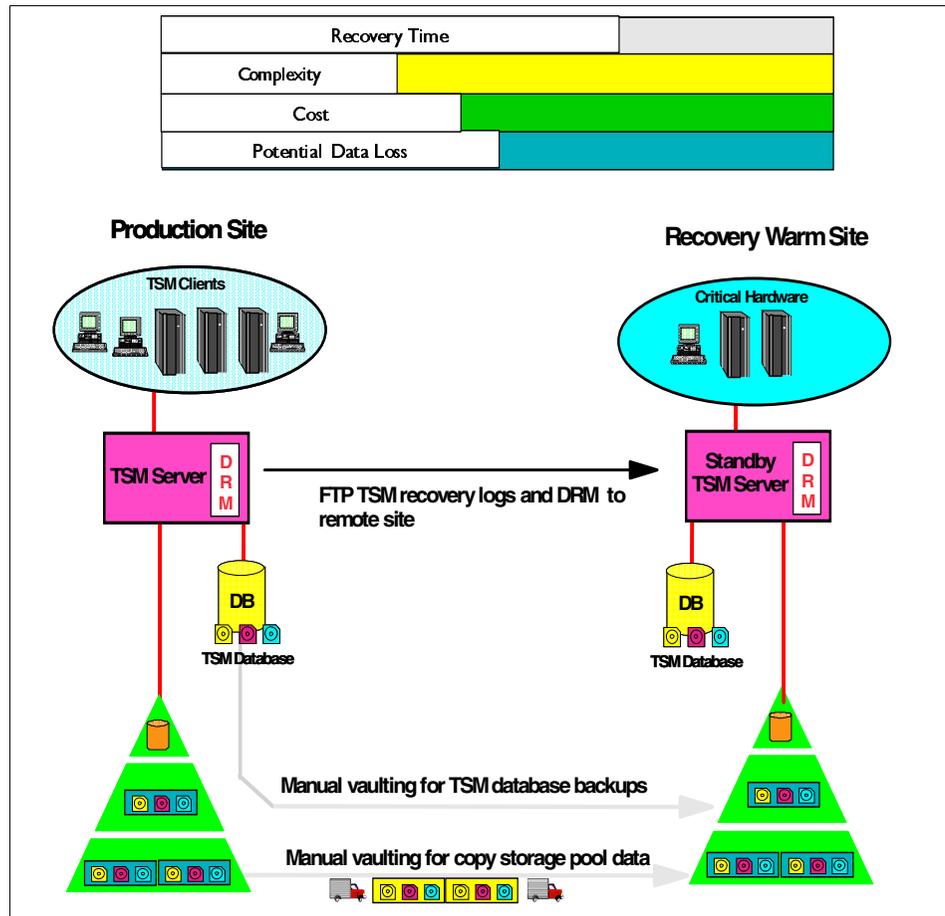


Figure 14-6 Electronic vaulting of TSM DRM, manual DB, and copy pool vaulting

Using a secure WAN network, the DRM output can be sent to the remote site using FTP services. The TSM database backups can be manually vaulted along with the copy storage pool data to the warm site environment for disaster

recovery. In the event of a disaster, current DRM plans help automate the recovery of TSM on the standby TSM server. Once the TSM database is restored, the copy storage pool data can be used to restore critical systems. This type of TSM implementation can be easily modified to support growth into electronic vaulting and hotsite implementations should the business requirements for availability increase.

14.1.8 Dual production sites using TSM

Many large organizations use two or more data center environments. Multiple production sites can leverage TSM for Disaster Recovery, as shown in Figure 14-7.

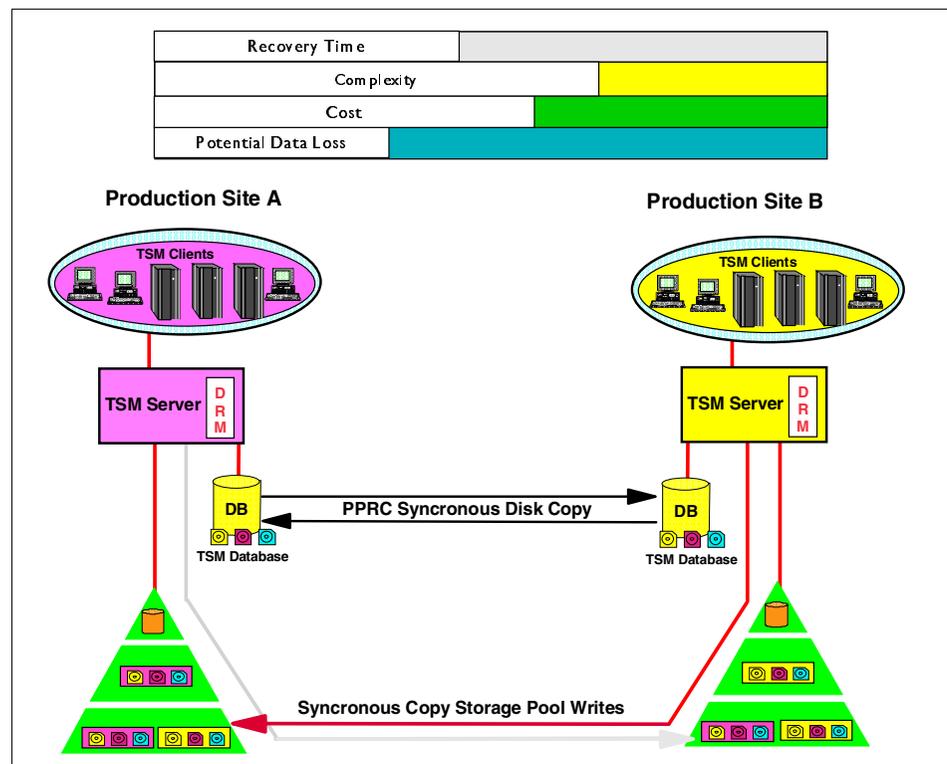


Figure 14-7 Dual production sites, electronic TSM DB and copy pool vaulting

This implementation is similar to the mirrored site implemented described in 14.1.3, “Tier 6 mirrored site implementation with TSM” on page 311. Basically, each TSM environment functions independently, and vaults its data to the alternate TSM site. In the case of a disaster at an individual site, the existing TSM environment is used to recover the lost TSM environment and associated

production systems. On each site, the TSM infrastructure must be sized to support capacity and workload requirements for the production and recovery environments.

14.1.9 Remote TSM server at warmsite, with manual offsite vaulting

Another implementation of TSM, which is popular with Storage Service Providers (SSPs), is the remote implementation of a TSM environment. This method relies solely on storage networking connections for data backup, as illustrated in Figure 14-8.

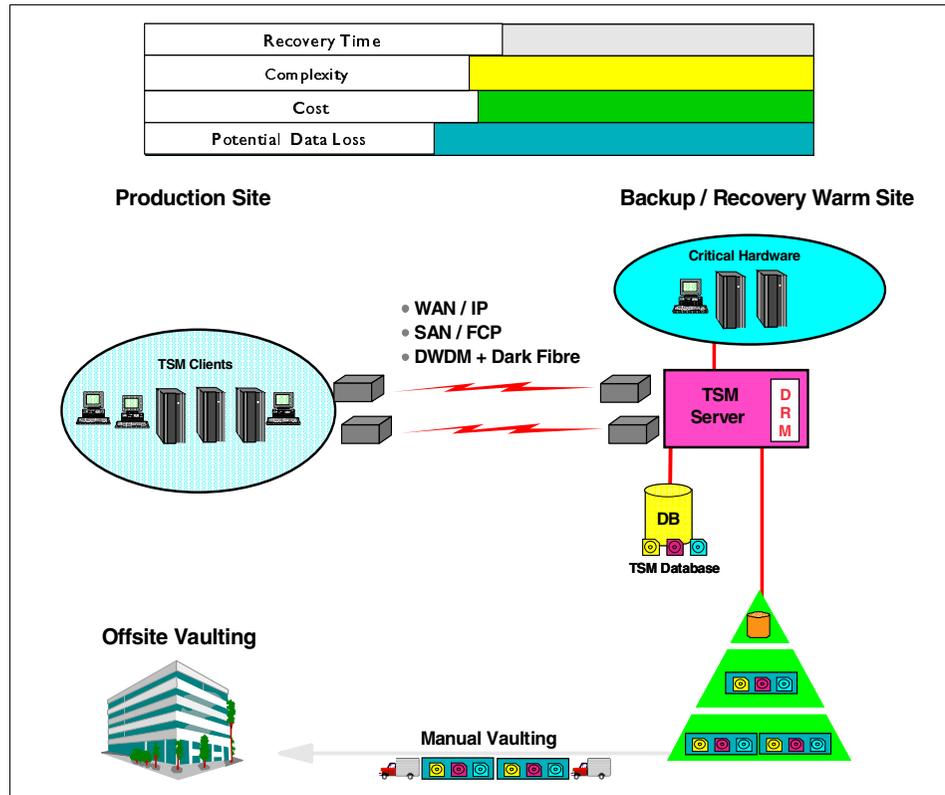


Figure 14-8 Remote TSM Server at warm site, plus manual offsite vaulting

Since traditional TSM client implementations rely on IP protocol, backups can be implemented across IP networks. Using channel extenders or DWDM technologies, multiple types of TSM client data stream (FCP and IP) can be transmitted site to site across considerable distances. In this scenario, the onus is on the network architecture and its ability to handle full and incremental backup traffic. TSM database backups and copy storage pool data can then be manually

vaulted from the recovery site to an offsite location. In the event of a production site disaster, the TSM environment will be ready to restore data to critical systems in the warm site environment.

14.2 Case studies

Here are two case studies, based on real data, which show how TSM was used to meet customer DR requirements.

14.2.1 Case study 1 - background

XYZ Bank (the “Bank”) is a financial institution holding a universal banking licence and carrying out business activities through a network of branches and branch-offices in the country. The Bank provides a wide range of banking and financial services to entrepreneurial entities, to individuals and institutional clients.

The Bank has a central IT environment based on an IBM mainframe, IBM TotalStorage Enterprise Storage Server (ESS), and an IBM TotalStorage Virtual Tape Server (VTS) with 3494 tape library. The Bank has also established up a backup center in a remote location, and in cooperation with IBM, has developed Disaster Recovery procedures for central IT processing. In the current situation, the Bank can switch central IT processing to the backup center and continue business within 4 hours. This is the RTO as defined by Bank management.

The Bank also has a decentralized IT environment, consisting of different server platforms including Windows NT, Windows 2000, Sun Solaris, and AIX. Some of the application servers are clustered. Applications running on those servers including Microsoft SQL, Microsoft Exchange, and Lotus Notes. Data backup is provided by TSM Server Version 4.2, installed on a Windows 2000 platform. The main backup storage for these servers is an IBM TotalStorage Enterprise Tape Library located in the primary location, along with all the decentralized servers. Figure 14-9 shows the local setup, where the TSM server migrates the disk storage pools to storage pools on the tape library.

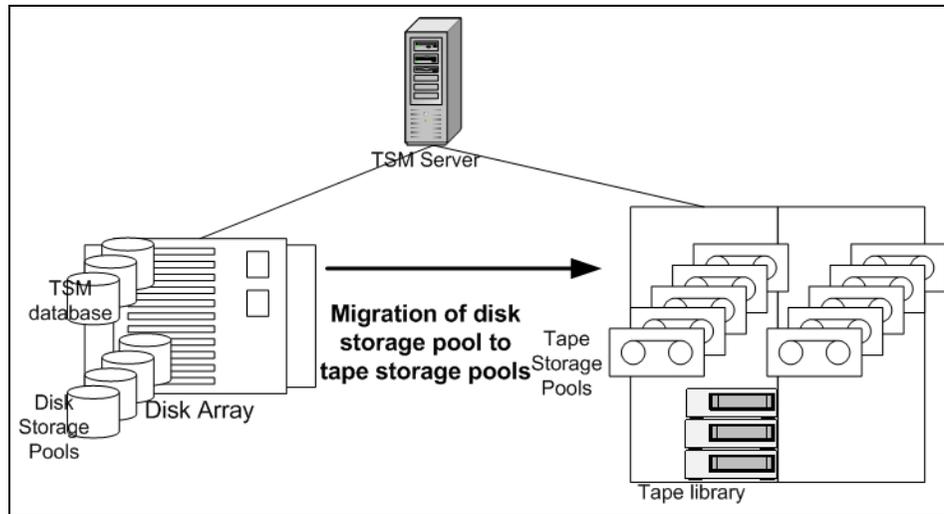


Figure 14-9 Starting situation at XYZ Bank

In this set up the Bank can restore any application server in the event of some local disaster like a hard drive crash, or individual server loss. In the event of loss of the TSM server, backup of the database and configuration files has been prepared for Disaster Recovery, but the estimated recovery time for the TSM server was too long.

Problem

The Bank asked IBM to help develop a Disaster Recovery solution for the decentralized servers. An RTO of 4 hours is established for the most important application servers. The Bank also decided to start develop a new branch office solution based on a Linux platform, using DB2 database and WebSphere. The backup, archive, and DR solution to has to cover all new servers and applications.

Solution

IBM analyzed the current situation in the Bank and suggested implementing a solution with the following steps:

- ▶ Equip the IBM TotalStorage Enterprise Tape Library in the backup location with additional tape drives for the TSM server and connect the TSM server and the drives using existing SAN connection between the two locations.
- ▶ Define this second library on the TSM server and create copy storage pools on the second library on remote location.
- ▶ Backup the primary storage pools from the tape library in the primary location to copy storage pools on the tape library in the backup location.

- ▶ Define ESS disk space in the primary location for the TSM server, place the TSM database, recovery log, and disk storage pools on it, and mirror these disk volumes using PPRC to the backup location.

These steps will increase the availability of the TSM server and data, and implement prerequisites for the Disaster Recovery solution for existing servers. See Figure 14-10 for an illustration of how protection for the TSM server has been improved.

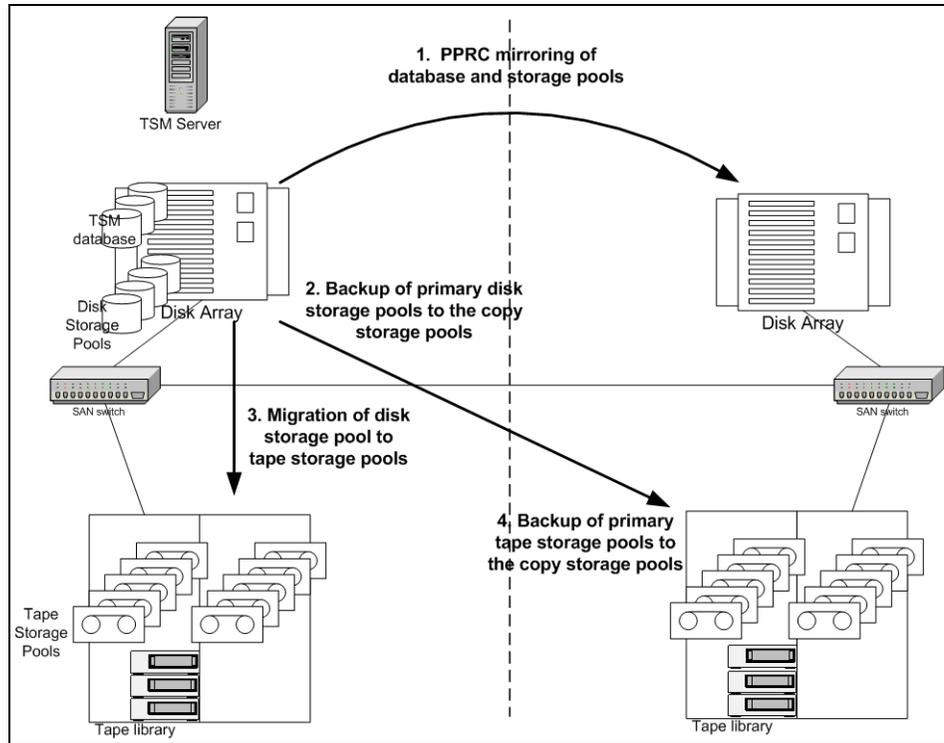


Figure 14-10 TSM data mirrored on backup location

In this situation, the Bank can now perform a Disaster Recovery for the TSM server, providing another server is available to replace it. The next series of recommended steps should solve all the remaining requirements.

- ▶ Buy a second TSM server for the backup location, and add new tape drives for the second TSM server to both libraries.
- ▶ Define disk space on the ESS for second TSM server, and mirror this disk space using PPRC to the other (primary) location.
- ▶ Geographically disperse clustered servers.

- ▶ Set up both TSM servers, so that will have defined hot instances for each other.
- ▶ The new application servers can be backed up by the second TSM server, which will have primary tape storage pools in the backup location, and copy tape storage pools in the primary location. The final situation will be as shown in Figure 14-11.

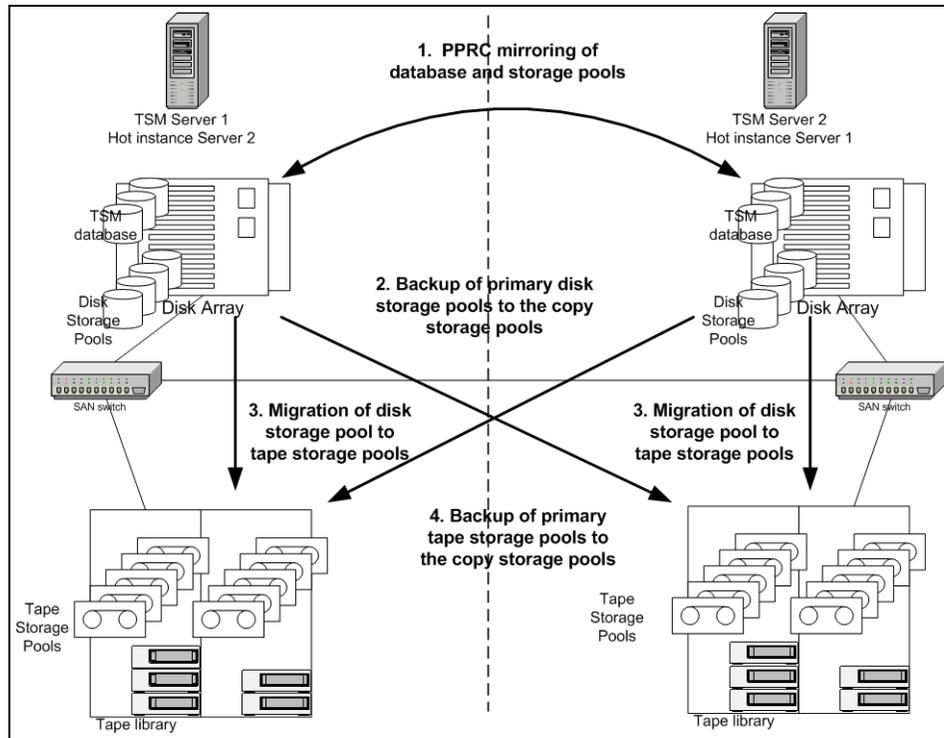


Figure 14-11 Recommended solution for Bank - final situation

14.2.2 Case study 2 - background

ABC Bank (the “Bank”) is the leading locally headquartered commercial bank on island in the Pacific Ocean islands. The company and its subsidiaries provide varied financial services to businesses, consumers, and governments in the Pacific rim.

The Bank’s main IT operations center includes more than 150 open system servers, based on Windows, Solaris, AIX, and HP platforms. A team of six administrators was using a variety of local tape drives and media for data and operating system backups. Network based backups were done using three different storage management applications, which tended to consume large amounts of tape and network resources. These backup products also experienced problems with restoring data in a reasonable time scale. Nobody was able to guarantee an RTO in the case of a major disaster caused by volcanic eruption or a tsunami. The administrator’s skills were sufficient to perform a bare metal restore in the case of a single server crash, but very little enterprise policy existed for disaster recovery. The data backup policy was decentralized and no enterprise standards for backup, retention, or archive existed.

Problem

In this situation the Bank asked an IBM Business Partner to develop a solution to provide data consolidation, centralized backup, and increased disaster recover ability. Disaster recovery RTO requirements for critical servers was a 2 hour time frame, but the existing infrastructure barely provided restore capabilities within a 12 hour window. Furthermore, enterprise storage growth estimates were 30% per year, and no current strategy or product provided a manageable solution. Current software license agreements for backup software were soon to expire, and a critical decision for enterprise storage management strategy had to be made.

On a business level, the Bank required the development of a hot site data center within 10 KM of the main production facility to provide DR capabilities for the enterprise. Rigorous regulatory requirements mandated that critical data must be continuously available and recoverable.

Solution

In the first phase, the IBM Business Partner analyzed the banking IT environment and daily processing to estimate the total amount of data stored, daily change rate of data, and the overall volume of data storage projected for a 3 year time frame. The bank had approximately 150 TB of data, of which around 5 TB changed daily change. Approximately 50% of the data resided in databases. Since the majority of the Bank’s clients were located in the Pacific rim, the Bank

decided to build a backup location on the opposite site of the same island as the primary location.

The IBM Business Partner designed a solution where data was consolidated on IBM TotalStorage ESS disk arrays, with a backup and disaster recovery solution based on TSM servers and two IBM TotalStorage Enterprise Tape Libraries in each of the primary and backup locations. The IBM Business Partner also recommended to the Bank to move some critical servers to the backup location, and by this action, decrease the impact and required recoveries in the case of local disaster.

The solution was designed to provide site to site storage management capabilities using TSM, and disaster recovery capabilities for the hot site TSM environment. In the event of a primary site loss, TSM can restore the primary site data to standby recovery servers in the backup site. On each site, the TSM infrastructure should be sized to support capacity and workload requirements for the production and recovery environments.

14.3 TSM and DR from beginning to end - summary

To help you “put it all together” we provide here a summary of concepts presented in this book. In our experience, it’s the big picture, an understanding of your environment and requirements, and an understanding what the best solution is, that is the most the difficult task — not necessarily the technical aspects.

We understand that we have packed many DR and TSM concepts into this guide. As such we do not want to reiterate all those details here. We refer you to the rest of this redbook for that. Instead we provide you with a description that lists the components that should be considered for DR and TSM. Figure 14-12 presents this summary.

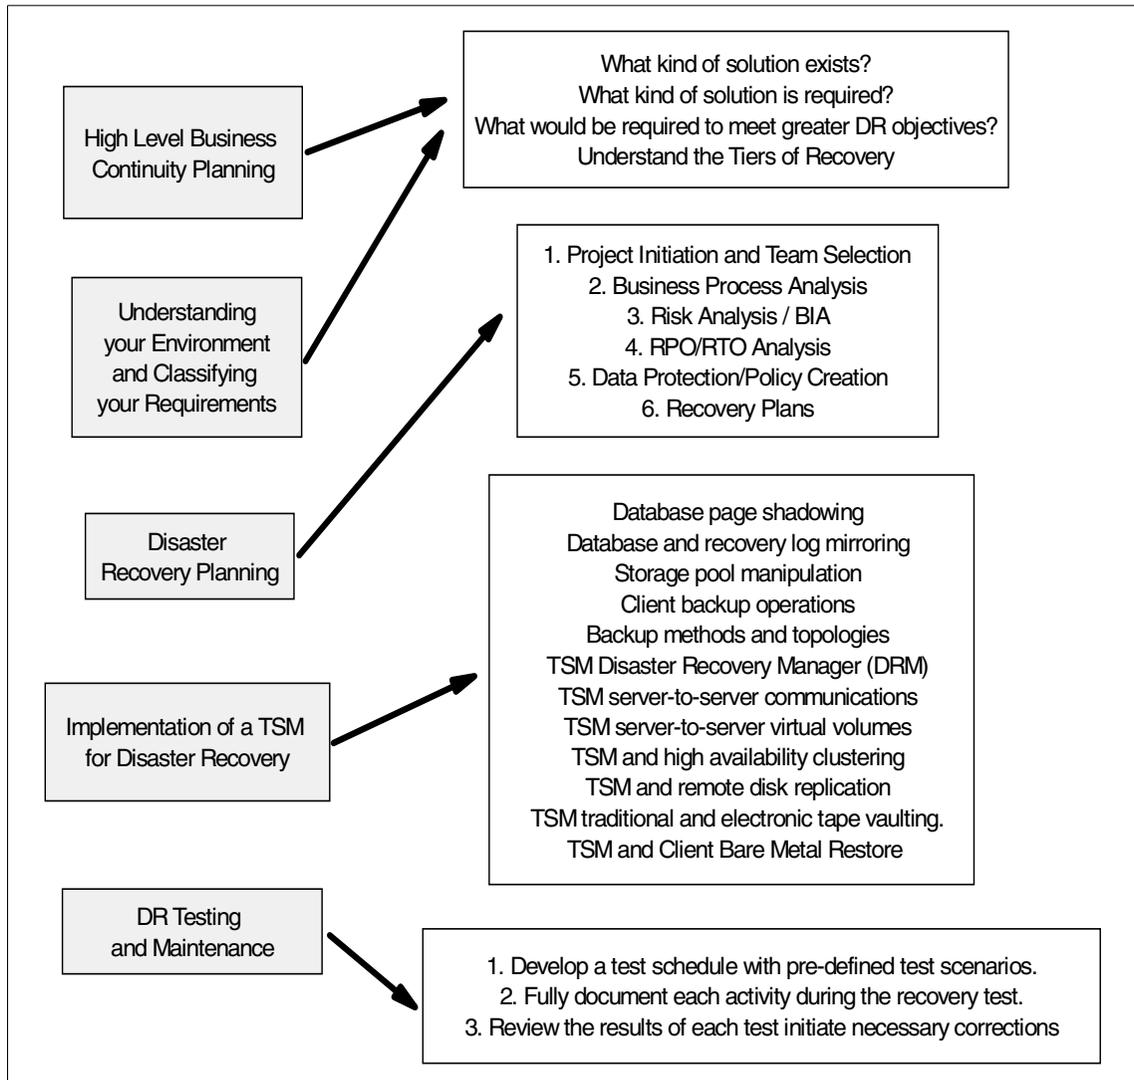


Figure 14-12 Summary of TSM and Disaster Recovery considerations

14.4 Application and database backup Redbooks

We acknowledge that DR strategies require special considerations for particular applications and databases that you may be using in your environment. Tivoli provides a number of added products which interface TSM with particular application environments. These are the products formerly known as Tivoli Data Protection (TDP). The current list includes:

- ▶ Tivoli Storage Manager for Application Servers
- ▶ Tivoli Storage Manager for Databases
- ▶ Tivoli Storage Manager for Enterprise Resource Management
- ▶ Tivoli Storage Manager for Hardware
- ▶ Tivoli Storage Manager for Mail

Detailed discussion of these products is beyond the scope of this book, however we recommend the following Redbooks for more details on TSM's capabilities for application and database backup.

- ▶ *Backing Up Oracle Using Tivoli Storage Management*, SG24-6249
- ▶ *Backing Up DB2 Using Tivoli Storage Manager*, SG24-6247
- ▶ *Backing Up Lotus Domino R5 Using Tivoli Storage Management*, SG24-5247
- ▶ *Using Tivoli Data Protection for Microsoft Exchange Server*, SG24-6147
- ▶ *R/3 Data Management Techniques Using Tivoli Storage Manager*, SG24-5743
- ▶ *Using Tivoli Data Protection for Microsoft SQL Server*, SG24-6148
- ▶ *Using Tivoli Storage Manager to Back Up Lotus Notes*, SG24-4534
- ▶ *Backing up WebSphere Application Server Using Tivoli Storage Management*, REDP0149



Part 3

Appendixes



DR and Business Impact Analysis Planning Templates

This appendix provides a template for Disaster Recovery Planning and the business impact analysis. These tools are intended to educate and help teams organize the process of plan development. Every organization's unique requirements must be reflected in a highly customized final plan. We provide a softcopy of this plan so you can modify it for your enterprise. The details for accessing this are in Appendix D, "Additional material" on page 387.

Disaster Recovery Plan Template

This sample format provides a template for preparing an information technology Disaster Recovery Plan. The template is intended to be used as a guide, and the DR planner should modify the format as necessary to meet the system's contingency requirements and comply with internal policies. Where practical, the guide provides instructions for completing specific sections. Text is added in certain sections; however, this information is intended only to suggest the type of information that may be found in that section. The text is not comprehensive and should be modified to meet specific organization and system considerations. The IT Disaster Recovery Plan should be marked with the appropriate security label, such as "Official Use Only".

A.1 Introduction

The introduction of the plan establishes the purpose, applicability, scope, and assumptions for the Disaster Recovery Plan. This plan template can scale from one system to the enterprise.

A.1.1 Purpose

This {system name} Disaster Recovery Plan establishes procedures to recover the {system name} following a disruption. The following objectives have been established for this plan:

- ▶ Maximize the effectiveness of contingency operations through an established plan that consists of the following phases:
 - Notification/Activation phase to detect and assess damage and to activate the plan
 - Recovery phase to restore temporary IT operations and recover damage done to the original system
 - Reconstitution phase to restore IT system processing capabilities to normal operations.
- ▶ Identify the activities, resources, and procedures needed to carry out {system name} processing requirements during prolonged interruptions to normal operations.
- ▶ Assign responsibilities to designated {Organization name} personnel and provide guidance for recovering {system name} during prolonged periods of interruption to normal operations.
- ▶ Ensure coordination with other {Organization name} staff who will participate in the Disaster Recovery Planning strategies. Ensure coordination with external points of contact and vendors who will participate in the Disaster Recovery Planning strategies.

A.1.2 Applicability

The {system name} Disaster Recovery Plan applies to the functions, operations, and resources necessary to restore and resume {Organization name}'s {system name} operations as it is installed at {primary location name, City, State}.

The {system name} Disaster Recovery Plan applies to {Organization name} and all other persons associated with {system name} as identified under A.2.3, "Responsibilities" on page 337.

The {system name} DRP is supported by {plan name}, which provides the {purpose of plan}. Procedures outlined in this plan are coordinated with and support the {plan name}, which provides {purpose of plan}.

A.1.3 Scope

The plan scope outlines the planning principles, assumptions, policy references, and a record of changes.

A.1.3.1 Planning Principles

Various scenarios were considered to form a basis for the plan, and multiple assumptions were made. The applicability of the plan is predicated on two key principles:

- ▶ The {Organization name}'s facility in {City, State}, is inaccessible;
- ▶ Therefore, {Organization name} is unable to perform {system name} processing for the Department.
- ▶ A valid contract exists with the alternate site that designates that site in {City, State}, as the {Organization name}'s alternate operating facility.
 - {Organization name} will use the alternate site building and IT resources to recover {system name} functionality during an emergency situation that prevents access to the original facility.
 - The designated computer system at the alternate site has been configured to begin processing {system name} information.
 - The alternate site will be used to continue {system name} recovery and processing throughout the period of disruption, until the return to normal operations.

A.1.3.2 Assumptions

Based on these principles, the following assumptions were used when developing the IT Disaster Recovery Plan:

- ▶ The {system name} is inoperable at the {Organization name} computer center and cannot be recovered within 48 hours.
- ▶ Key {system name} personnel have been identified and trained in their emergency response and recovery roles; they are available to activate the {system name} Disaster Recovery Plan.
- ▶ Preventive controls (for example, generators, environmental controls, waterproof tarps, sprinkler systems, fire extinguishers, and fire department assistance) are fully operational at the time of the disaster.

- ▶ Computer center equipment, including components supporting {system name}, are connected to an uninterruptible power supply (UPS) that provides 45 minutes to 1 hour of electricity during a power failure.
- ▶ {system name} hardware and software at the {Organization name} original site are unavailable for at least 48 hours.
- ▶ Current backups of the application software and data are intact and available at the offsite storage or disaster recovery facility.
- ▶ The equipment, connections, and capabilities required to operate {system name} are available at the alternate site in City, State.
- ▶ Service agreements are maintained with {system name} hardware, software, and communications providers to support the emergency system recovery.

The {system name} Disaster Recovery Plan does not apply to the following situations:

- ▶ Overall recovery and continuity of business operations. The Business Resumption Plan (BRP) and Continuity of Operations Plan (COOP) are appended to the plan.
- ▶ Emergency evacuation of personnel. The Occupant Evacuation Plan (OEP) is appended to the plan.
- ▶ Any additional constraints should be added to this list.

A.1.4 References/requirements

This {system name} Disaster Recovery Plan complies with the {Organization name}'s IT DR Planning policy as follows:

The organization shall develop a DR Planning capability to meet the needs of critical supporting operations in the event of a disruption extending beyond 72 hours. The procedures for execution of such a capability shall be documented in a formal Disaster Recovery Plan and shall be reviewed at least annually and updated as necessary. Personnel responsible for target systems shall be trained to execute contingency procedures. The plan, recovery capabilities, and personnel shall be tested to identify weaknesses of the capability at least annually.

A.2.2 Line of succession

The {Organization name} sets forth an order of succession, in coordination with the order set forth by the department to ensure that decision-making authority for the {system name} Disaster Recovery Plan is uninterrupted. The Chief Information Officer (CIO), {Organization name} is responsible for ensuring the safety of personnel and the execution of procedures documented within this {system name} Disaster Recovery Plan. If the CIO is unable to function as the overall authority or chooses to delegate this responsibility to a successor, the Deputy CIO shall function as that authority. Continue description of succession as applicable.

A.2.3 Responsibilities

The following teams have been developed and trained to respond to a contingency event affecting the IT system.

The Disaster Recovery Plan establishes several teams assigned to participate in recovering {system name} operations. The {team name} is responsible for recovery of the {system name} computer environment and all applications. Members of the team name include personnel who are also responsible for the daily operations and maintenance of {system name}. The team leader title directs the {team name}.

Continue to describe each team, their responsibilities, leadership, and coordination with other applicable teams during a recovery operation.

The relationships of the team leaders involved in system recovery and their member teams are illustrated in the figure.

Insert hierarchical diagram of recovery teams. Show team names and leaders; do not include actual names of personnel.

Describe each team separately, highlighting overall recovery goals and specific responsibilities. Do not detail the procedures that will be used to execute these responsibilities. These procedures will be itemized in the appropriate phase sections.

A.3 Notification and Activation phase

This phase addresses the initial actions taken to detect and assess damage inflicted by a disruption to {system name}. Based on the assessment of the event, the plan may be activated by the Disaster Recovery Planning coordinator.

In an emergency, the {Organization name}'s top priority is to preserve the health and safety of its staff before proceeding to the Notification and Activation procedures.

Contact information for key personnel is located in Appendix A. The notification sequence is listed below:

- ▶ The first responder is to notify the Disaster Recovery Planning Coordinator. All known information must be relayed to the Disaster Recovery Planning coordinator.
- ▶ The systems manager is to contact the Damage Assessment Team Leader and inform them of the event. The Disaster Recovery Planning coordinator is to instruct the Team Leader to begin assessment procedures.
- ▶ The Damage Assessment Team Leader is to notify team members and direct them to complete the assessment procedures outlined below to determine the extent of damage and estimated recovery time. If damage assessment cannot be performed locally because of unsafe conditions, the Damage Assessment Team is to follow the outline below.

A.3.1 Damage assessment procedures

Detailed procedures should be outlined to include activities to determine the cause of the disruption; potential for additional disruption or damage; affected physical area and status of physical infrastructure; status of IT equipment functionality and inventory, including items that will need to be replaced; and estimated time to repair services to normal operations.

- ▶ Upon notification from the disaster recovery planning coordinator, the Damage Assessment Team Leader is to ...
- ▶ The Damage Assessment Team is to

A.3.2 Alternate assessment procedures

- ▶ Upon notification from the Disaster Recovery Planning coordinator, the Damage Assessment Team Leader is to ...
- ▶ The Damage Assessment Team is to
 - When damage assessment has been completed, the Damage Assessment Team Leader is to notify the Disaster Recovery Planning coordinator of the results.
 - The Disaster Recovery Planning Coordinator is to evaluate the results and determine whether the Disaster Recovery Plan is to be activated and if relocation is required.

- Based on assessment results, the Disaster Recovery Planning Coordinator is to notify assessment results to civil emergency personnel (for example, police, fire) as appropriate.

A.3.3 Criteria for activating the DR plan

The DRP is to be activated if one or more of the following criteria are met:

1. {System name} will be unavailable for more than X hours
2. Facility is damaged and will be unavailable for more than Y hours
3. Other criteria, as appropriate.

If the plan is to be activated, the DR Planning coordinator is to notify all Team Leaders and inform them of the details of the event and if relocation is required.

Upon notification from the DR Planning coordinator, Team Leaders are to notify their respective teams. Team members are to be informed of all applicable information and prepared to respond and relocate if necessary.

The DR Planning coordinator is to notify the off-site storage facility that a contingency event has been declared and to ship the necessary materials (as determined by damage assessment) to the alternate site.

The DR Planning coordinator is to notify the Alternate site that a contingency event has been declared and to prepare the facility for the Organization's arrival.

The DR Planning coordinator is to notify remaining personnel (via notification procedures) on the general status of the incident.

A.4 Recovery operations

This section provides procedures for recovering the application at the alternate site, whereas other efforts are directed to repair damage to the original system and capabilities.

The following procedures are for recovering the {system name} at the alternate site. Procedures are outlined per team required. Each procedure should be executed in the sequence it is presented to maintain efficient operations.

Primary Recovery Objective

State the first recovery objective as determined by the Business Impact Analysis (BIA). For each team responsible for executing a function to meet this objective, state the team names and list their respective procedures.

- ▶ {team name}
 - Team Recovery Procedures
- ▶ {team name}
 - Team Recovery Procedures
- ▶ {team name}
 - Team Recovery Procedures

Secondary Recovery Objective

State the second recovery objective as determined by the BIA. For each team responsible for executing a function to meet this objective, state the team names and list their respective procedures.

- ▶ {team name}
 - Team Recovery Procedures
- ▶ {team name}
 - Team Recovery Procedures
- ▶ {team name}
 - Team Recovery Procedures

Tertiary Recovery Objectives

State the remaining recovery objectives (as determined by the BIA). For each team responsible for executing a function to meet this objective, state the team names and list their respective procedures.

A.5 Return to normal operations

This section discusses activities necessary for restoring {system name} operations at the {Organization name}'s original or new site. When the computer center at the original or new site has been restored, {system name} operations at the alternate site must be transitioned back. The goal is to provide a seamless transition of operations from the alternate site to the computer center.

Original or New Site Restoration

Procedures should be outlined, per necessary team, to restore or replace the original site so that normal operations may be transferred. IT equipment and telecommunications connections should be tested.

- ▶ {team name}
 - Team Resumption Procedures
- ▶ {team name}

- Team Resumption Procedures

A.5.1 Concurrent processing

Procedures should be outlined, per necessary team, to operate the system in coordination with the system at the original or new site. These procedures should include testing the original or new system until it is functioning properly and the contingency system is shut down gracefully.

- ▶ {team name}
 - Team Resumption Procedures
- ▶ {team name}
 - Team Resumption Procedures

A.5.2 Plan deactivation

Procedures should be outlined, per necessary team, to clean the alternate site of any equipment or other materials belonging to the organization, with a focus on handling sensitive information. Materials, equipment, and backup media should be properly packaged, labeled, and shipped to the appropriate location(s). Team members should be instructed to return to the original or new site.

- ▶ {team name}
 - Team Testing Procedures
- ▶ {team name}
 - Team Testing Procedures

A.6 Plan appendices

The appendices included should be based on system and plan requirements.

- ▶ Personnel Contact List
- ▶ Vendor Contact List
- ▶ Equipment and Specifications
- ▶ Service Level Agreements and Memorandums of Understanding
- ▶ IT Standard Operating Procedures
- ▶ Business Impact Analysis
- ▶ Related Disaster Recovery Plans
- ▶ Emergency Management Plan

- ▶ Occupant Evacuation Plan
- ▶ Continuity of Operations Plan.
- ▶ Storage Management policy
- ▶ Equipment inventory and configuration data
- ▶ Bare metal restore procedures

Business Impact Analysis Template

This sample template is designed to assist the user in performing a BIA on an IT system. The BIA is an essential step in developing the IT Disaster Recovery Plan. The template is meant only as a basic guide and may not apply to all systems. The user may modify this template or the general BIA approach as required to best accommodate the specific system.

Table A-2 Preliminary system information

Organization:	Date BIA Completed:
System Name:	BIA Contact:
System Manager Contact:	
System Descriptions: {Discussion of the system purpose and architecture, including system diagrams}	
Internal System Contacts {Identify the individuals, positions, or offices within your organization that depend on or support the system; also specify their relationship to the system}	Responsibilities
▶ ▶ ▶	▶ ▶ ▶
External System Contacts {Identify the individuals, positions, or offices outside your organization that depend on or support the system; also specify their relationship to the system}	Responsibilities
▶ ▶ ▶	▶ ▶ ▶

Table A-3 Identify system resources

Identify the specific hardware, software, and other resources that comprise the system, including quantity and type.
<p>Hardware Resources</p> <ul style="list-style-type: none"> ▶ ▶ ▶
<p>Software Resources</p> <ul style="list-style-type: none"> ▶ ▶ ▶
<p>Other Resources</p> <ul style="list-style-type: none"> ▶ ▶ ▶

Table A-4 Identify critical roles

List the critical roles identified in Table A-2.
<ul style="list-style-type: none"> ▶ ▶ ▶ ▶ ▶

Table A-5 Link critical roles to critical resources

Identify the IT resources needed to accomplish the roles listed in Table A-4.	
Critical Role	Critical Resource

Table A-6 Identify outage impacts and allowable outage times

Characterize the impact on critical roles if a critical resource is unavailable; also, identify the maximum acceptable period that the resource could be unavailable before unacceptable impacts resulted.		
Resource	Outage Impact (financial)	Allowable Outage Time

Table A-7 Prioritize resource recovery

List the priority associated with recovering a specific resource, based on the outage impacts and allowable outage times provided in Table A-6. Use quantitative or qualitative scale (for example, high/medium/low, 1-5, A/B/C).	
Resource	Recovery Priority



B

Windows BMR configuration scripts

This appendix contains two scripts:

machchar.vbs - Generates a macro file containing insert machine commands. This script and resulting macro allow you to insert large files containing machine characteristics or machine instructions.

msinfoextract.vbs - Extracts subcategories of information from a report generated by MSINFO32.

Insert machine characteristics into DRM

The script in Example B-1 is provided with TSM on the Windows platform and is called machchar.vbs. It generates a macro file containing **INSERT MACHINE** commands. This script and resulting macro allow you to insert large files containing machine characteristics or machine instructions (with a modification to the script) a single time.

Example: B-1 Example of VBScript command to insert machine characteristics

```
*****
'Tivoli Disaster Recovery Manager for Windows NT/2000 Sample Script
'
'Read machine characteristics from an input file and build an output file
'that is a TSM macro.The TSM macro contains statements which are
'TSM commands to insert client machine information into the ADSM server
'database.The TSM macro is used with the TSM administrative client.
'
'Invoke with:
'cscript machchar.vbs machinename inputmachinefilename outputmacrofilename
'where:
'machinename is the name of a machine that has previously
'been defined to the TSM server with the
'DEFINE MACHINE command
'inputmachinefilename is the name of the input file which contains
'the client machine characteristics.This file
'would typically be built on the client machine
'then the file would be transferred to the
'Windows machine where the TSM Administrative
'client is installed.
'outputmacrofilename is the name of the output file in an existing
'directory which will be the TSM macro.The
'TSM macro will consist of a series of commands
'to insert machine characteristics into the TSM
'server database.For example:
'
'INSERT MACHINE mch1 n characteristics='xxx...'
'
'where:
'n represents the sequence number
'this line will have in the
'TSM server database
'xxx...'represents a single line from
'the input file
'
'NOTE:The maximum length of a line of machine
'characteristics is 1024
'
'Example usage:
```

```

'cscript machchar.vbs mch1 c:\client1 \clientinfo.txt c:\client1
\clientinfo.mac
'*****
Dim args
Dim MACHINENAME,INFILE,OUTFILE
dim fso
dim fo,fi
dim SEQUENCE
Dim CRLF
CRLF =Chr(13)&Chr(10)
Const ForReading =1,ForWriting =2
'*****
'Get input arguments:MACHINENAME =machinename
'INFILE =inputmachinefilename
'OUTFILE =outputmacrofilename
'*****
set args =Wscript.Arguments
If args.Count <3 Then
Wscript.Echo _
"usage:cscript machchar.vbs machinename inputmachinefilename
outputmacrofilename"&CRLF &_
"example:cscript machchar.vbs mch1 c:\client1 \clientinfo.txt c:\client1
\clientinfo.mac"
Wscript.Quit(1)
Else
MACHINENAME =args.Item(0)
INFILE =args.Item(1)
OUTFILE =args.Item(2)
End if
Set fso =CreateObject("Scripting.FileSystemObject")
'*****
'Create the TSM macro file.
'*****
Set fo =fso.OpenTextFile(OUTFILE,ForWriting,True)
Wscript.Echo "Creating TSM macro file:"&OUTFILE
'*****
'Place a TSM command in the TSM macro to delete any existing machine
'characteristics for this machine from the TSM server database.
'*****
fo.WriteLine "delete machine "&MACHINENAME &"type=characteristics"
'*****
'Read a line from the input machine characteristics file,add the TSM
'command to insert the line of machine characteristics into the TSM server
'database,and write the result to the output TSM macro.
'*****
SEQUENCE =1
Set fi =fso.OpenTextFile(INFILE,ForReading,False)
Do While fi.AtEndOfStream <>True
INLINE =fi.ReadLine

```

```
fo.WriteLine "insert machine "&MACHINENAME &"&SEQUENCE &"char='"&INLINE &"'"
SEQUENCE =SEQUENCE +1
Loop
'*****
'Close the files.
'*****
fo.Close
fi.Close
```

Reducing msinfo32 output using a VBScript

The VBScript script in Example B-2, (called msinfoextract.vbs) extracts subcategories of information from a report generated by `msinfo32`.

Using this script, you could build your own report of just the subcategories you want, for example:

- ▶ `Cscript msinfoextract.vbs msinfo32.rpt system summary > msinfo32x.rpt`
- ▶ `Cscript msinfoextract.vbs msinfo32.rpt drives >> msinfo32x.rpt`
- ▶ `Cscript msinfoextract.vbs msinfo32.rpt adapter >> msinfo32x.rpt`

An alternative approach might be to write a script that eliminates the stanzas you specify.

Example: B-2 VBScript to Reduce Output from MSINFO32 Report

```
*****
' Extract subcategories from a Windows 2000 MSINFO32 report.
' Assumes the subcategories are delimited between '[subcategoryname]'
' and the next '['. Results are written to stdout.
*****
' How to generate an msinfo32 report from command line on Win2k.
' cd \program files\common files\microsoft shared\msinfo
' msinfo32 /report msinfo32.rpt /categories +all
' or an abbreviated version:
' msinfo32 /report msinfo32.rpt /categories +systemsummary+components
*****

Dim args, MSINFOFILE, SUBCATNAME
Dim fso, fi, fo
Dim WORDS, FIRSTLINE, WRITING
Dim CRLF
CRLF = Chr(13) & Chr(10)
FIRSTLINE = TRUE
WRITING = FALSE

set args = Wscript.Arguments
If args.Count < 2 Then
    Wscript.Echo _
        "usage: cscript msinfoextract.vbs msinfo32file subcategoryname" & CRLF & _
        "example: cscript msinfoextract.vbs c:\data\msinfo32.rpt Drives"
    Wscript.Quit(1)
Else
    MSINFOFILE = args.Item(0)
    SUBCATNAME = "[" & args.Item(1) & "]"
End If
```

```
Set fso = CreateObject("Scripting.FileSystemObject")
Set fi = fso.OpenTextFile(MSINFOFILE, 1, False, -2)
Do While fi.AtEndOfStream <> True
  ALINE = fi.ReadLine
  if FIRSTLINE eqv TRUE then
    Wscript.Echo ALINE
    FIRSTLINE = FALSE
  end if
  if Left(ALINE,1) = "[" then
    WRITING = FALSE
    if ALINE = SUBCATNAME then WRITING = TRUE
  end if
  if WRITING = TRUE then Wscript.Echo ALINE
loop
fi.close
```

```

*
* Syntax
* +-----+-----+
* | HTTPort      | port_addr      |
* +-----+-----+
*
COMMmethod HTTP
HTTPPort 1580
*
* =====
*
* TCP/IP
*
* *****
* TCPPort
*
* Specifies the TCP/IP port address of a TSM server.
*
* Syntax
* +-----+-----+
* | TCPPort      | port_addr      |
* +-----+-----+
*
COMMmethod TCP/IP
TCPPort 1500
*
* *****
* TCPWindow size
*
* Specifies the amount of data to send or receive
* before TCP/IP exchanges acknowledgements with the client node.
* This actual window size that is used in a session will be the
* minimum size of the server and client window sizes.
* Larger window sizes may improve performance
* at the expense of memory usage.
*
* Syntax
* +-----+-----+
* | TCPWindow size | window_size    |
* +-----+-----+
*
TCPWindow size 63
*
* *****
* TCPNODELAY
*
* Specifies whether the server should send small amounts
* of data or allow TCP/IP to buffer the data.
* Disallowing buffering may improve throughput at the expense

```

```

*   of more packets being sent over the network.
*
*   Syntax
*   +-----+-----+
*   | TCPNODELAY      | YES | NO |
*   +-----+-----+
*
TCPNODELAY Yes
*
* =====
*
* NAMEDPIPE
*
* *****
* NAMEdpipeName
*
* Specifies the name of the TSM server's named pipe.
*
* Syntax
* +-----+-----+
* | NAMEdpipeName   | name |
* +-----+-----+
*
COMMmethod NAMEDPIPE
NAMEdpipeName \\.\pipe\Server1
*
* =====
*
* *****
* NPBUFFERSIZE
*
* Specifies the size of the named pipes communication buffer size
* in KB.
*
* Syntax
* +-----+-----+
* | NPBUFFERSIZE    | value |
* +-----+-----+
*
NPBUFFERSIZE 8
*
* =====
*
* *****
* SECUREPIPES
*
* Specifies whether or not to use secure named pipes (NT Unified Logon.)
*
* Specify a value of Yes or No

```

```

*
* Syntax
* +-----+-----+
* | SECUREPIPES      | value                |
* +-----+-----+
*
SECUREPipes No
*
* =====
* *****
* ADMSGROUP
*
* Specifies the Windows NT Group name to use for authentication.
*
* Syntax
* +-----+-----+
* | ADMSGROUP        | groupname            |
* +-----+-----+
*
ADMSGROUPname admsserver
*
* *****
* NPAUDITSuccess
*
* Specifies whether or not to audit successful use of secure named pipes
*
* Specify a value of Yes or No
*
* Syntax
* +-----+-----+
* | NPAUDITSuccess   | value                |
* +-----+-----+
*
NPAUDITSuccess No
*
* *****
* NPAUDITFailure
*
* Specifies whether or not to audit a failed attempt to use of secure
* named pipes
*
* Specify a value of Yes or No
*
* Syntax
* +-----+-----+
* | NPAUDITFailure   | value                |
* +-----+-----+
*

```

NPAUDITFailure No

```
*
* =====
*
* MSGINTERVAL
*
* *****
* MSGINTERval
*
* Specifies the number of minutes to wait between issuing a mount-tape
* tape message on the TSM server console.
*
* Syntax
* +-----+-----+
* | MSGINTERval      | value |
* +-----+-----+
```

MSGINTERval 1

```
*
* =====
*
* MAXSESSIONS
*
* *****
* MAXSessions
*
* Specifies the number of simultaneous client sessions.
*
* Syntax
* +-----+-----+
* | MAXSessions      | value |
* +-----+-----+
```

MAXSessions 25

```
*
* =====
*
* BUFPOOLSIZE
*
* *****
* BUFPoolsize
*
* Specifies the size of the database buffer pool in Kbytes.
*
* Syntax
* +-----+-----+
* | BUFPoolsize      | value |
* +-----+-----+
```

```

BUFPoolsize 262144
*
* =====
*
* LOGPOOLSIZE
*
* *****
* LOGPoolsize
*
* Specifies the size of the log buffer pool in Kbytes.
*
* Syntax
* +-----+-----+
* | LOGPoolsize      | value |
* +-----+-----+
*
LOGPoolsize 512
*
* =====
*
* COMMTIMEOUT
*
* *****
* COMMTimeout
*
* Specifies the communication timeout value in seconds.
*
* Syntax
* +-----+-----+
* | COMMTimeout      | value |
* +-----+-----+
*
COMMTimeout 60
*
* =====
*
* IDLETIMEOUT
*
* *****
* IDLETimeout
*
* Specifies the number of minutes that a client session can be idle
* before its session will be canceled.
*
* Syntax
* +-----+-----+
* | IDLETimeout      | value |
* +-----+-----+
*

```

```

IDLETimeout 15
*
* =====
*
* TXNGroupmax
*
* *****
* TXNGroupmax
*
* Specifies the number of files tranferred as a group between commit
* points.
*
* Syntax
* +-----+-----+
* | TXNGroupmax | numfiles |
* +-----+-----+
*
TXNGroupmax 256
*
* =====
*
* DATEFORMAT
*
* *****
* DATEformat
*
* Specifies the format in which date references will be displayed.
*
* Syntax
* +-----+-----+
* | DATEformat | value |
* +-----+-----+
*
DATEformat 1
*
* =====
*
* TIMEFORMAT
*
* *****
* TIMEformat
*
* Specifies the format in which time references will be displayed.
*
* Syntax
* +-----+-----+
* | TIMEformat | value |
* +-----+-----+
*

```

```

TIMEformat 1
*
* =====
*
* NUMBERFORMAT
*
* *****
* Numberformat
*
* Specifies the format in which number references will be displayed.
*
* Syntax
* +-----+-----+
* | Numberformat | value |
* +-----+-----+
*
Numberformat 1
*
* =====
*
* MESSAGEFORMAT
*
* *****
* MESSageformat
*
* Specifies the format in which messages will be displayed.
*
* Syntax
* +-----+-----+
* | MESSageformat | value |
* +-----+-----+
*
MESSageformat 1
*
* =====
*
* LANGUAGE
*
* *****
* LANGUage
*
* Specifies the language to use for help and error messages.
*
* Syntax
* +-----+-----+
* | LANGUage | name |
* +-----+-----+
*
LANGUage AMENG

```

```

*
* =====
*
* EXPINTERVAL
*
* *****
* EXPInterval
*
* Specifies the number of hours between automatic inventory expiration
* runs.
*
* Syntax
* +-----+-----+
* | EXPInterval | value |
* +-----+-----+
*
EXPInterval 24
*
* =====
*
* EXPQUIET
*
* *****
* EXPQuiet
*
* Reduces the number of policy change messages generated during
* expiration processing.
*
* Specify a value of Yes or No
*
* Syntax
* +-----+-----+
* | EXPQuiet | value |
* +-----+-----+
*
EXPQuiet Yes
*
* =====
*
* MIRRORREAD
*
* *****
* MIRRORRead
*
* Specifies the mode used for reading recovery log pages or data base
* log pages
*
* Syntax
* +-----+-----+-----+-----+
* | MIRRORRead | LOG | DB | Normal | Verify |

```

```

* +-----+-----+-----+
*
MIRRORRead DB Normal
*
* =====
*
* MIRRORWRITE
*
* *****
* MIRRORWrite
*
* Specifies how mirrored volumes are accessed when the server writes
* pages to the recovery log or database during normal processing.
*
* Syntax
* +-----+-----+-----+
* | MIRRORWrite | LOG | DB | Sequential | Parallel |
* +-----+-----+-----+
*
MIRRORWrite DB Sequential
*
* =====
*
* *****
* SELFTUNEBUFPOOLSIZE
*
* Specifies whether TSM can automatically tune the database buffer pool
* size. If you specify YES, TSM resets the buffer cache statistics at
* the start of expiration processing. After expiration completes, if
* cache hit statistics are less than 98, TSM increases the database
* buffer pool size to try to increase the cache hit percentage.
* The default is NO.
*
* Syntax
* +-----+-----+-----+
* | SELFTUNEBUFpoolsize | YES | NO |
* +-----+-----+-----+
*
SELFTUNEBUFpoolsize No
*
* =====
*
* *****
* DBPAGESHADOW
*
* Specifies whether database page shadowing is enabled. If database
* page shadowing is enabled TSM mirrors every write to a database page.
* You can enable shadowing only if database volume mirrors are written
* to in parallel (that is, the MIRRORWRITE DB option is set to PARALLEL.

```

```

* The default is YES.
*
* Syntax
* +-----+-----+-----+
* | DBPAGEShadow | YES | NO |
* +-----+-----+-----+
*
DBPAGEShadow Yes
*
* =====
*
* *****
* DBPAGESHADOWFILE
*
* Specifies the name of the file to use for database page shadowing.
* If database page shadowing is enabled, the page shadow either goes
* to the default file or is created in the directory that the server
* is running from. The default file, DBPGSHDW.BDT, resides in the
* directory where the server is installed.
*
* Syntax
* +-----+-----+-----+
* | DBPAGESHADOWFile | value |
* +-----+-----+-----+
*
DBPAGESHADOWFile "dbpgshdw.bdt"
*
* =====
* MIRRORREAD
*
* *****
* MIRRORRead
*
* Specifies the mode used for reading recovery log pages or data base
* log pages
*
* Syntax
* +-----+-----+-----+
* | MIRRORRead | LOG | DB | Normal | Verify |
* +-----+-----+-----+
*
MIRRORRead LOG Normal
*
* =====
* MIRRORWRITE
*
* *****
* MIRRORWrite

```

*
 * Specifies how mirrored volumes are accessed when the server writes
 * pages to the recovery log or database during normal processing.
 *

* Syntax

```
* +-----+-----+-----+
* | MIRRORWrite | LOG | DB | Sequential | Parallel |
* +-----+-----+-----+
```

*
 MIRRORWrite LOG Parallel

```
*
* =====
```

* MOVEBATCHSIZE

```
*
* *****
* MOVEBatchsize
```

* Use this entry field to specify the number of files that are to be
 * moved and grouped together in a batch within the same transaction.
 *

* Specify a number between 1 and 256.

*
 * The default value is 32.
 *

* Syntax

```
* +-----+-----+
* | MOVEBatchsize | value |
* +-----+-----+
```

* MOVEBatchsize 40

```
*
* =====
```

* MOVESIZETHRESHOLD

```
*
* *****
* MOVESizethreshold
```

* Use this entry field to specify a threshold, in megabytes, for the amount
 * of data moved as a batch within the same server transaction. When this
 * threshold is reached, no more files are added to the current batch. A
 * new transaction is then started after the current batch is moved.
 *

* Specify a number between 1 and 500 (megabytes).

*
 * The default value is 1 (megabyte).
 *

* Syntax

```

* +-----+-----+
* | MOVESizethreshold | value |
* +-----+-----+
*
MOVESizethresh 500
*
* =====
*
* *****
* SELFTUNETXNSIZE
*
* Specifies whether TSM can automatically change the values of the
* TXNGROUPMAX, MOVEBATCHSIZE, and MOVESIZETHRESH server options.
* TSM sets the TXNGROUPMAX option to optimize client-server
* throughput and sets the MOVEBATCHSIZE and MOVESIZETHRESH options
* to their maximum to optimize server throughput. The default is NO.
*
* Syntax
* +-----+-----+
* | SELFTUNETXNsize | YES | NO |
* +-----+-----+
*
SELFTUNETXNsize No
*
* =====
*
* STATUSMSGCNT
*
* *****
* STATusmsgcnt
*
* Use this entry field to specify the number of records (times 1000)
* that will be processed between status messages during DSMSERV DUMPDB
* and DSMSERV LOADDDB commands.
*
* Specify a number between 1 and 10000 (this number is multiplied by 1000).
*
* The default value is 10.
*
* Syntax
* +-----+-----+
* | STATusmsgcnt | value |
* +-----+-----+
*
STATusmsgcnt 1
*
* =====
*
* VOLUMEHISTORY

```

```

* *****
* VOLUMEHistory <filename>
*
* Specifies the name of a file that should contain sequential
* volume history information when it is changed by the server.
* Sequential volume history information is used by the administrator
* and server processes during server database recovery.
*
* More than one of these parameters may be specified to record
* sequential volume history information to multiple files
*
* Syntax
* +-----+-----+
* | VOLUMEHistory | filename |
* +-----+-----+
*
* VOLUMEHistory "volhist.out"
* The previous line was replaced by PREPARE to provide a fully qualified
* file name.
VOLHIST "C:\PROGRA~1\TIVOLI\TSM\SERVER1\VOLHIST.OUT"

*
* =====
*
* DEVCONFIG
* *****
* DEVCONFig <filename>
*
* Specifies the name of a file that should contain device
* configuration information when it is changed by the server.
* Device configuration information is used by the
* server processes during server database recovery or load and
* DSMSErv DUMPDB processing.
*
* More than one of these parameters may be specified to record
* device configuration information to multiple files.
*
* Syntax
* +-----+-----+
* | DEVCONFig | filename |
* +-----+-----+
*
* DEVCONFig "devcnfg.out"
* The previous line was replaced by PREPARE to provide a fully qualified
* file name.
DEVCONF "C:\PROGRA~1\TIVOLI\TSM\SERVER1\DEVCONF.OUT"

*
* =====

```

```

*
* *****
* RESTOREINTERVAL
*
* Specifies the restore interval.
*
* Syntax
* +-----+-----+
* | RESTOREINTERVAL      | value |
* +-----+-----+
*
RESTOREINTERval 1440
*
* =====
* *****
* USELARGEBuffers
*
* Specifies whether or not to use large buffers.
*
* Specify a value of Yes or No
*
* Syntax
* +-----+-----+
* | USELARGEBUFFERS      | values |
* +-----+-----+
*
USELARGEbuffers Yes
*
* =====
*
* MISC
*
* *****
* DISABLESCeds
*
* Specifies whether or not administrative and client schedules are
* disabled during a TSM server recovery scenario
*
* Specify a value of Yes or No
*
* Syntax
* +-----+-----+
* | DISABLESC             | value |
* +-----+-----+
*
DISABLESCHEDS No
*
* =====

```

```

*
* MISC
*
* *****
* EVENTSERVER
*
* Specifies whether at startup this server should contact the TSM
* event server.
*
* Specify a value of Yes or No
*
* Syntax
* +-----+-----+
* | EVENTSERVER | value |
* +-----+-----+
*
EVENTSERVER Yes
*
* =====
*
* REQSYSAUTHFILE
*
* Specifies whether, system authority is required for administrative
* commands that cause the server to write to an external file.
*
* Syntax
* +-----+-----+
* | REQSYSauthoutfile | YES | NO |
* +-----+-----+
*
REQSYS Yes
*
* =====
*
* ENABLE3590LIBRARY
*
* *****
* ENABLE3590LIBrary
*
* When 3590 support is enabled, the TSM server automatically begins to
* use the category with a number that is one greater than the existing
* scratch category code that was specified on the TSM server
* DEFINE LIBRARY command.
*
* Specify a value of Yes or No
*
* Syntax
* +-----+-----+
* | ENABLE3590LIBrary | value |

```

```

* +-----+-----+
*
ENABLE3590 Yes
*
* =====
*
* 3494SHARED
*
* *****
* 3494SHARED
*
* Specifies whether to use extra polling when drives are being shared.
* Default is No.
*
* Specify a value of Yes or No
*
* Syntax
* +-----+-----+
* | 3494SHARED | value |
* +-----+-----+
*
3494SHARED Yes
*
* =====
*
* ASSISTVCRRECOVERY
*
* Specifies whether TSM assists an IBM 3570 or 3590 drive in
* recovering from a lost or corrupted Vital Cartridge Records (VCR)
* condition. If YES, the default, TSM server will locate directly
* to the end-of-data to allow the drives to restore the VCR.
*
* Syntax
* +-----+-----+
* | ASSISTVCRREcovery | YES | NO |
* +-----+-----+
*
ASSISTVCRREcovery Yes
*
* =====
*
* QUERYAuth
*
* *****
*
* Specifies the administrative authority level that should be required
* to issue QUERY or SQL SELECT commands. By default any administrator
* can issue a QUERY or SELECT command. If you would like to restrict
* the use of these commands to administrators with a specific
* authority level, this option can be specified with the level of

```

```

* authority desired.
*
* Syntax
* +-----+-----+-----+-----+
* | QUERYAuth | NOne | SYstem | POLicy | STorage | OPERator | ANalyst |
* +-----+-----+-----+-----+
*
* Parameters
*
*     NONE     Specifies that any administrator can issue
*              QUERY or SELECT commands without requiring
*              any administrative authority.
*
*     SYSTEM   Specifies that administrators must have SYSTEM
*              authority to issue QUERY or SELECT commands
*
*     POLICY   Specifies that administrators must have POLICY
*              authority over one or more policy domains
*              (or SYSTEM authority) to issue QUERY or SELECT commands
*
*     STORAGE  Specifies that administrators must have STORAGE
*              authority over one or more storage pools
*              (or SYSTEM authority) to issue QUERY or SELECT commands
*
*     OPERATOR Specifies that administrators must have OPERATOR
*              authority (or SYSTEM authority) to issue QUERY or SELECT
*              commands
*
*     The default value is NONE
*
* Examples
*     QUERYAUTH     SYSTEM
*     QUERYAUTH     OPERATOR
*
* =====
* QUERYAuth NONE
*
* =====
*
* ADREGISTER
*
* Specifies whether the TSM registers itself with Active Directory
* when the server starts up. The default is No.
*
* Syntax
* +-----+-----+-----+
* | ADREGISTER | YES | NO |
* +-----+-----+-----+
*

```

ADREGISTER No

```
*
* =====
*
* ADUNREGISTER
*
* Specifies whether the TSM unregisters itself with Active Directory
* when the server halts. The default is No.
*
* Syntax
* +-----+-----+
* | ADUNREGISTER | YES | NO |
* +-----+-----+
```

ADUNREGISTER No

```
*
* =====
*
* ADSETDC
*
* Specifies the name or address of the domain controller that
* Active Directory is installed on. If this parameter is not provided
* the default action is to attempt to automatically detect the
* domain controller that the machine is registered to.
*
* Syntax
* +-----+
* | ADSETDC | DomainController name or TCP/IP address |
* +-----+
```

ADCOMMENT

```
*
* Specifies the comment used when registering the server with
* Active Directory. If this parameter is not provided a default
* comment will be generated during server registration.
```

```
*
* Syntax
* +-----+
* | ADCOMMENT | comment |
* +-----+
```

* The following option was added by PREPARE.
DISABLESCHEDS YES

* The following option was added by PREPARE.
DISABLESCHEDS YES



D

Additional material

This redbook refers to additional material that can be downloaded from the Internet as described below.

Locating the Web material

The Web material associated with this redbook is available in softcopy on the Internet from the IBM Redbooks Web server. Point your Web browser to:

<ftp://www.redbooks.ibm.com/redbooks/SG246844>

Alternatively, you can go to the IBM Redbooks Web site at:

ibm.com/redbooks

Select the **Additional materials** and open the directory that corresponds with the redbook form number, SG246844.

Using the Web material

The additional Web material that accompanies this redbook includes the following file:

<i>File name</i>	<i>Description</i>
SG246844files.zip	Zipped additional materials

Once unzipped, the following files will be available:

- ▶ **DRP_BIA Templates.doc** The BIA template (in MS-Word format) shown in Appendix A, “DR and Business Impact Analysis Planning Templates” on page 331.
- ▶ **vg_recovery** AIX script for saving VG information, presented in 11.2.2, “Saving additional volume group definitions” on page 243.

How to use the Web material

Create a subdirectory (folder) on your workstation, and unzip the contents of the Web material zip file into this folder. The **vg_recovery** script can only be executed on a system with the AIX Operating System installed. The **DRP_BIA Templates.doc** file requires a system with Microsoft Word installed to view and modify.

Abbreviations and acronyms

ACL	Access Control List	HIPAA	Health Insurance Portability and Accountability Act
API	Application Programming Interface	HSM	Hierarchical Storage Management
ATM	Asynchronous Transfer Mode	HVAC	Heating, Ventilation and Air-Conditioning
BCP	Business Continuity Planning/Plan	I/O	Input Output
BIA	Business Impact Analysis	IBM	International Business Machines Corporation
BIOS	Basic Input Output System	IRQ	Interrupt Request
BMR	Bare Metal Recovery	ISV	Independent Software Vendor
BRP	Business Resumption Plan	IT	Information Technology
COOP	Continuity of Operations Plan	ITSO	International Technical Support Organization
CPU	Central Processing Unit	JFS	Journalled File System
DBA	Database administrator	LAN	Local Area Network
DFS	Distributed File System	LDPRS	Living Disaster Recovery Planning System
DLT	Digital Linear Tape	LILO	Linux Loader
DNS	Domain Name Server	LPAR	Logical Partition
DR	Disaster Recovery	LPAR	Logical Partitioning
DRM	Disaster Recovery Manager	LTO	Linear Tape Open
DRP	Disaster Recovery Plan	LUN	Logical Unit Number
DR Planning	Disaster Recovery Planning	LVSA	Logical Volume Storage Agent
DVD	Digital Video Disc	MAN	Metropolitan Area Network
DWDM	Dense Wavelength Division Multiplexing	MB	Megabyte
EDI	Electronic Data Interchange	MTBF	Mean Time Between Failure
EMP	Electromagnetic Pulse	NAS	Network Attached Storage
ESCON	Enterprise Systems Connection	NDMP	Network Data Management Protocol
FCP	Fibre-Channel Protocol	NIM	Network Installation Management
Gbps	Gigabit per second	NRO	Network Recovery Objective
HACMP	High Availability Clustered Multi-Processing	NSM	Network Storage Manager
HBA	Host Bus Adapter		

NTFS	NT Filesystem	WAN	Wide Area Network
OEP	Occupant Evaluation Plan	WMI	Windows Management Instrumentation
OS	Operating System	WORM	Write Once, Read Many
PASE	Portable Application Solutions Environment		
PCI	Peripheral Component Interconnect		
PPRC	Peer to Peer Remote Copy		
PPRC-XD	PPRC Extended Distance		
PSM	Persistent Storage Manager		
PTAM	Pickup Truck Access Method		
QIC	Quarter Inch Cartridge		
RAID	Redundant Array Of Independent Disk		
RMAN	Recovery Manager		
RPO	Recovery Point Objective		
RTO	Recovery Time Objective		
SAN	Storage Area Network		
SCSI	Small Computer Systems Interface		
SDG	SAN Data Gateway		
SDK	Software Developers Kit		
SLA	Service Level Agreement		
SPOF	Single Point of Failure		
SPOT	Shared Product Object Tree		
SRDF	Symmetrix Remote Data Facility		
SSA	Serial Storage Architecture		
SSP	Storage Service Provider		
TCP/IP	Transmission Control Protocol/Internet Protocol		
TDP	Tivoli Data Protection		
TOC	Table of Contents		
TSM	IBM Tivoli Storage Manager		
UPS	Uninterruptible Power Supply		
USB	Universal Serial Bus		
VTS	Virtual Tape Server		

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 393.

- ▶ *Deploying the Tivoli Storage Manager Client in a Windows 2000 Environment*, SG24-6141
- ▶ *Getting Started with Tivoli Storage Manager: Implementation Guide*, SG24-5416
- ▶ *Tivoli Storage Manager Version 5.1 Technical Guide*, SG24-6554
- ▶ *Tivoli Storage Management Concepts*, SG24-4877
- ▶ *Introduction to SAN Distance Solutions*, SG24-6408
- ▶ *NIM: From A to Z in AIX 4.3*, SG24-5524
- ▶ *Backing Up Oracle Using Tivoli Storage Management*, SG24-6249
- ▶ *Backing Up DB2 Using Tivoli Storage Manager*, SG24-6247
- ▶ *Backing Up Lotus Domino R5 Using Tivoli Storage Management*, SG24-5247
- ▶ *Using Tivoli Data Protection for Microsoft Exchange Server*, SG24-6147
- ▶ *R/3 Data Management Techniques Using Tivoli Storage Manager*, SG24-5743
- ▶ *Using Tivoli Data Protection for Microsoft SQL Server*, SG24-6148
- ▶ *Using Tivoli Storage Manager to Back Up Lotus Notes*, SG24-4534
- ▶ *Backing up WebSphere Application Server with Tivoli Storage Management*, REDP0149
- ▶ *Managing AIX Server Farms*, SG24-6606

Other resources

These publications are also relevant as further information sources:

- ▶ *IBM Tivoli Storage Manager for AIX Administrator's Reference V5.1*, GC32-0769
- ▶ *IBM Tivoli Storage Manager for Windows Administrator's Reference V5.1*, GC32-0783
- ▶ *IBM Tivoli Storage Manager for Windows Administrator's Guide V5.1*, GC32-0782
- ▶ *IBM Tivoli Storage Manager for UNIX Backup-Archive Clients Installation and User's Guide V5.1*, GC32-0789
- ▶ *AIX 5L Version 5.1 Network Installation Management Guide and Reference*, SC23-4385
- ▶ *AIX Version 4.3 Network Installation Management Guide and Reference*, SC23-4113
- ▶ *Unix Backup and Recovery*, by W. Curtis Preston and Gigi Estabrook. O'Reilly & Associates, 1st edition, December 1999. ISBN: 1565926420.
- ▶ *Linux Complete Backup and Recovery HOWTO*, by Charles Curley, found at: <http://www.tldp.org/HOWTO/Linux-Complete-Backup-and-Recovery-HOWTO/>

Referenced Web sites

These Web sites are also relevant as further information sources:

- ▶ Public Domain Linux bootable floppy disk
<http://www.toms.net/rb>
- ▶ VERITAS
<http://www.veritas.com/>
- ▶ Ultrabac
<http://www.ultrabac.com/>
- ▶ Cristie
<http://www.cristie.co.uk>
- ▶ IBM Tivoli Storage Management homepage
<http://www.tivoli.com/products/solutions/storage/news.html>
- ▶ IBM SysBack
<http://sysback.services.ibm.com>
- ▶ Microsoft Windows 2000 Resource Kit tools downloads
<http://www.microsoft.com/windows2000/techinfo/reskit/tools/>

- ▶ Tivoli Technical Product Documents
http://www.tivoli.com/support/public/Prodman/public_manuals/td/TD_PROD_LIST.html
- ▶ Tivoli Storage Manager Products Technical Support
http://www.tivoli.com/support/storage_mgr/requirements.html
- ▶ Tivoli Storage Manager Versions 5.1 and 4.2 Device List For AIX, HP-UX, SUN, and Windows Servers
http://www.tivoli.com/support/storage_mgr/devices/all.html
- ▶ Windows 2000 Diskmap.exe
<http://www.microsoft.com/windows2000/techinfo/reskit/tools/existing/diskmap-o.asp>
- ▶ FTP root at index.storsys.ibm.com
<http://index.storsys.ibm.com>

How to get IBM Redbooks

You can order hardcopy Redbooks, as well as view, download, or search for Redbooks at the following Web site:

ibm.com/redbooks

You can also download additional materials (code samples or diskette/CD-ROM images) from that site.

IBM Redbooks collections

Redbooks are also available on CD-ROMs. Click the CD-ROMs button on the Redbooks Web site for information about all the CD-ROMs offered, as well as updates and formats.

Index

A

- ACL 105
- adaptive subfile backup 16
- advanced coupled systems 28
- AIX 144, 168
 - /etc/bootptab 257
 - /etc/exclude.rootvg 242
 - bare metal recovery 248, 261
 - boot from mkysyb media 248
 - bootp 251, 261
 - IPL-ROM emulation 261
 - JFS 242
 - network boot 240, 256
 - NIM 240, 248
 - rebuild volume group 245
 - root volume group 241
 - save volume group configuration 243
 - SMIT 240, 243
 - SysBack 240
 - system image backups 240
- AIX commands
 - bootinfo 260
 - cron 263
 - lsnim 253
 - mkcd 246
 - mkysyb 67, 166, 226, 240–241, 245, 253
 - mkvgdata 243
 - nimclient 260
 - restvg 244
 - savevg 243
- alternate site architecture 91
- alternate site costs 93
- alternate site partnering 93
- API 18–19, 285
- archive 12, 98
- asynchronous mirroring 152
- ATM 156, 309
- audit 51, 59, 75
- automated switchover 28
- availability 44, 90

B

- backup and restore of data 12

- backup-centric planning 96
- bare metal recovery 10, 19, 138, 165, 223, 285
 - AIX 240
 - alternate image restore 227
 - automated BMR suite 228
 - Linux 290
 - live media bootstrap 227
 - manual bootstrap 225
 - OS image bootstrap 226
 - Solaris 231
 - Windows 2000 267
- bare metal restore products
 - Cristie 285
 - UltraBac 286
 - VERITAS 228, 286
- BCP 5
- BIA 7
- BIA. See Business Impact Analysis
- BMR 285
- bootp 257, 264
- bottleneck 96
- BRP 335
- business continuance 22
- business continuity 9, 40, 43, 54, 91, 96
- Business Continuity Planning 5, 9, 43–44, 48, 61, 95
- Business Impact Analysis 7, 9, 30, 49, 52, 82, 88, 90, 97–98, 157, 167, 339, 343
- Business Resumption Plan 335

C

- cache 86
- capacity planning 43, 85, 104, 167
- change management 47, 85
- channel extenders 31, 34
- client options file 151
- client/server architecture 98
- clients
 - bare metal recovery 285
- client-server 13
- cluster resources 149
- clustering 29, 84, 144, 150, 165, 168
- CNT 309

coldsite 10, 79, 90–92, 180–181, 242
contingency plan 23, 38
Continuity of Operations Plan 335
COOP 335
copy storage pool 141, 153
cost of downtime 6, 90, 155
Cristie 19, 285

D

damage assessment team 57

data

- archive 98
- availability 39, 88–89, 97
- backup and restore 12, 92
- classification 50–51, 96, 98, 101
- corruption 40
- critical 51, 104
- important 51
- management 93
- management policy 80
- mirrored 92
- non-critical 51
- parity 86
- protection 38, 47, 87
- recovery sequence 114
- replication 29
- retention 98
- security 98, 155
- semi-important 51
- sources 49
- transfer rates 42
- vaulting 92

data center 83

- availability 79
- environmentals 80
- infrastructure 80
- mirrored 90, 92
- planning 81
- preventative controls 82
- security 80, 83
- space planning 81

database integrity 66

DB2 18, 130

DES 18

direct losses 7

Disaster Recovery

- active secondary site 26, 34
- Approval Board 73

costs 167

data synchronization 67

database applications 65

dual site 27, 90

electronic vaulting 25–26

hotsite 24

legislative requirements 38

mirrored site 91–92

outsourcing 60

regulatory requirements 98

roll-back 65–66, 68

service provider 67

solution costs 29

synchronized datacenters 11, 27

test review 70

test scenarios 64, 68

testing 167

tiers 21, 64–65, 169, 177

TSM 95

zero data loss 28

Disaster Recovery Manager. See DRM under TSM

Disaster Recovery Plan 5, 9, 23, 37, 43–44, 52, 54, 109, 190, 332

- activation 57, 92

- audience 54

- audit 75, 191

- background information section 55

- change management 71

- covering letter 78

- damage assessment 57

- design appendices 60

- distribution list 77

- Document Owner 73

- establishing currency 72

- format 55

- generation with DRM 205

- maintenance 63, 71

- notification/activation 56, 94

- obsolescence 75

- obsolete version 72

- reconstitution phase 59

- recovery phase 58

- recovery scenarios 56

- secure destruction 75

- Task Owners 73

- testing 63, 65

- updating 74

- version control 74

Disaster Recovery Plan file 135

- Disaster Recovery Planning 48, 95
- disaster tolerance 10
- disk
 - advanced copy functions 106
 - hot swap 168
 - hot-spare 87, 169
 - I/O 112
 - instant copy 130
 - locally attached 40
 - mirroring 40, 86, 105, 119, 152, 169
 - performance 87
 - RAID-1 86
 - RAID-3 87
 - RAID-5 87
 - redundancy 86, 169
 - remote replication 87, 151, 155, 189
 - SAN attached 169
 - striping 86
 - synchronous mirroring 311
- distribution list 72
- DLT 171
- DNS 11, 47, 89
- downtime 7, 30
 - by industry 8
 - cost to a business 6
- DR Planning 5
- DRM 31–32, 85
 - COPYSTGPOOL.VOLUMES.AVAILABLE 212
 - COPYSTGPOOL.VOLUMES.DESTROYED 212
 - DEVICE.CONFIGURATION.FILE 214, 218
 - MACHINE.RECOVERY.MEDIA.REQUIRED 212
 - PREPARE 137, 195, 198, 206, 275
 - PRIMARY.VOLUMES.DESTROYED 212
 - PRIMARY.VOLUMES.REPLACEMENT.CREATE.CMD 212
 - PRIMARY.VOLUMES.REPLACEMENT.MAC 213
 - RECOVERY.DEVICES.REQUIRED 211
 - RECOVERY.INSTRUCTIONS.DATABASE 196
 - RECOVERY.INSTRUCTIONS.GENERAL 195
 - RECOVERY.INSTRUCTIONS.INSTALL 195, 211–212
 - RECOVERY.INSTRUCTIONS.OFFSITE 195
 - RECOVERY.INSTRUCTIONS.STGPOOL 196
 - RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE 208, 215
 - RECOVERY.SCRIPT.NORMAL.MODE 208
 - RECOVERY.SCRIPT.NORMAL.MODE.CMD 219
 - RECOVERY.VOLUMES.REQUIRED 211
 - SERVER.REQUIREMENTS 211
 - VOLUME.HISTORY.FILE 211
- dual attach tape drive 189
- DVD 227
- DWDM 31, 34, 41, 43, 89, 112, 155–156, 308, 319

E

- electronic tape vaulting 31, 152–154, 180
- electronic vaulting 25–26, 34, 92, 121, 142, 315
- ESCON 155, 310
- ESS 87, 106, 126, 130
- EXTENDED COPY 130

F

- facilities costs 93
- facilities planning 90
- failover 145, 148, 188
- Fibre Channel protocol (FCP) 41, 112
- fibre extenders 155
- FICON 156
- financial industry 38
- firewall 90
- FlashCopy 18, 87, 106, 126, 130, 169, 311
- fsck 233
- ftp 207, 317
- full database backup 122

G

- Gantt 114
- geographic failover 189
- Gigabit Ethernet 41, 168
- GNU mkisofs 246
- gzip 295

H

- HACMP 35, 84, 144–145, 188
- HAGEO 35
- hardware 10
- hardware redundancy 10
- HBA 168, 182
- heartbeat 145
- Hierarchical Space Management (HSM) 13, 39
- high availability 10, 35, 84, 89, 92, 144, 187

high bandwidth connections 29, 31, 35
HIPAA 38
hotsite 10, 24–26, 32, 34, 64, 79, 90–92, 154, 180, 183, 313
hot-swappable devices 84, 86
HVAC 82

I

IBM Business Continuity and Recovery Services 93
IBM Network Storage Manager 164
IBM SAN Data Gateway 129, 192
IBM Tivoli Storage Manager. See TSM
incremental backups 17
indirect losses 7
Informix 18
infrastructure planning 85, 95
infrastructure redundancy 81
insourcing 60
instant copy 130
intelligent disk subsystem 18
IP extenders 89
iSCSI 156
ISV 19
IT outages 5

K

Kerberos 18

L

LAN 41
LAN performance 111
Linux
 /etc/fstab 293
 /etc/rc.sysinit 295
 /etc/sysconfig 290
 /etc/sysconfig/network 293
 /proc 290–291
 bare metal recovery 289–290, 295
 LILO 304
 mini-root 296, 299
 recovery diskette 296
 Red Hat 295
 system configuration 290
 system partition information 292
Linux commands
 chroot 297, 304
 dd 298

df 290, 293
fdisk 290, 292–293, 300
ifconfig 290, 293
tomsrvtb 296
logical partition (LPAR) 39
logical volume 106, 133
Logical Volume Storage Agent 134
Lotus Domino 18
Lotus Notes 18
LTO 171

M

MAN 155
Metropolitan Area Network (MAN) 43
Microsoft Exchange 18
Microsoft Project 114
Microsoft SQL Server 18
MIGDELAY 178
mission-critical 9, 51, 84, 113, 170
mkisofs 246
MSCS 35, 84, 144, 148, 150, 188
MTBF 83

N

NDMP 132
NetView 89
network appliances 132
network architecture 87
network attached storage (NAS) 41, 131, 156, 169
network bandwidth 143
network boot 228
network failover planning 11, 47, 88
Network Installation Management (NIM) 240, 248
network monitoring 89
Network Recovery Objective (NRO) 11, 52, 58
network redundancy 88, 112
network security 90
network topologies 111
network transfer rates 42
NFS 251
NTFS 129, 175

O

Occupant Evacuation Plan 335
OEP 335
offsite data storage 154, 170
offsite storage 23–24, 26, 32, 62, 98, 120, 135, 153,

156
onsite hardware inventory 83
onsite parts inventory 83
OpenView 89
optical fiber 43
optical media 13, 38, 120, 169
Oracle 18, 130, 170
OS/400 168
outages 5, 30, 38, 43, 188
outsourcing 60, 81

P

Persistent Storage Manager 131
Pickup Truck Access Method (PTAM) 24–25, 32
power infrastructure 81
PPRC 31, 34–35, 87, 106, 152, 155, 175, 184, 189, 310, 316, 322
PPRC Extended Distance 311
PPRC-XD 311
preventative controls 82
primary storage pool 141, 153
progressive backup 16, 50
project initiation 44

Q

QIC 171

R

R/3 130
RAID 86–87, 169
RAID-1 86
RAID-3 87
RAID-5 87, 113
RDBMS 39
recovery performance 96
recovery plans 47
Recovery Point Objective (RPO) 10, 30, 46, 52, 64, 96, 98, 113, 175
recovery time 43, 50, 64
Recovery Time Objective (RTO) 10, 24, 30, 46, 48–49, 52, 64, 96, 98, 100, 103, 107–108, 111, 113, 151, 172, 180, 320
Redbooks Web site 393
 Contact us xxvi
redundant networks 112
remote access 90
remote disk replication 154

removable media 121
resource contention 114
risk analysis 9, 30, 45, 48, 81
RMAN 170
roll-back 65
routers 87

S

SAN 88, 112, 155, 171
 and backup 171
 transmission distance 155
 zone 189
SAN Data Gateway 129
SANergy 172
SAP R/3 18, 170
SCSI 41, 168, 189
SCSI reserve 145
SCSI-3 130
sequential access storage 121, 169
server architecture 84
service level agreement 90, 188
Serviceguard 84
services 167
SHARE User Group 21–22
SMIT 240
Snapshot 169
software distribution 226, 240, 248
Solaris 229
 device representation 232
 logical storage configuration 232
 restore root file systems 235
 save system configuration 232
Solaris commands
 boot 233
 dmesg 232
 fsck 236
 installboot 237
 iostat 232
 isainfo 232
 netstat 232
 newfs 235
 nisbackup 231
 nisrestore 231
 prtconf 232
 prvtoc 232
 psrinfo 232
 sysdef 232
 ufsdump 229–230, 233–234

- ufsrestore 229–230, 234, 236
- SPARC 230
- split-mirror backup 126
- SRDF 31, 34–35, 87, 106, 152, 189
- SSA 169
- standards 85
- Storage Area Network (SAN) 41
- storage growth 39, 164, 167
- storage management 164
- storage management challenges 39, 164, 171
- storage management planning 48, 95
- Storage Service Provider 26, 92, 319
- strategic planning 80, 104, 109
- Sunguard 93
- supporting infrastructure 49
- switches 87
- synchronous mirroring 152
- SysBack 240, 264
- system dependencies 49
- system software 85

T

- tape automation 170
- tape cartridge volume 109
- tape device driver 211
- tape device failover 189
- tape drive dual attach 189
- tape drive performance 107–108, 169
- tape header information 110
- tape vaulting 62, 153
- tar 295
- target server 140
- TCP/IP 41, 141–142
- TCPSERVERADDRESS 182
- TDP 17, 126
- TDP for NDMP 132
- technology density 81
- technology selection 80
- telecommunications 93
- TFTP 251, 256
- The Kernel Group 286
- TimeFinder 18, 87, 126
- Tivoli Data Protection. See TDP
- Tivoli Disaster Recovery Manager. See DRM under TSM
- Tivoli Storage Management 4
- Tivoli Storage Manager. See TSM
- TKG 286

- training and testing 47
- transactions 65
- TSM 68, 80, 84, 92, 164, 228, 230, 241, 248, 268
 - adaptive subfile backup 16, 101, 124–125
 - administration 14, 166, 193, 272
 - ADSMSCSI device driver 211
 - and HACMP 144–145
 - and MSCS 148, 150
 - API 18–19, 170, 285
 - architecture 98
 - archive 12, 17, 33, 100, 127, 142
 - archive copy group 141
 - backup 17, 100–101, 117
 - backup and restore 12
 - backup methods 124
 - backup set 101, 127
 - backup window 170, 172, 187
 - Backup/Archive client 13, 269, 295
 - backup-centric planning 96
 - caching 178
 - client backup methods 124
 - client configuration information 135, 137–138
 - client data 113, 141
 - client interface 15
 - client option set 139
 - client options file 104, 147, 184, 186–187, 247, 278, 304
 - client policy 100
 - client Web browser interface 16
 - clients 100
 - clustering applications 144, 188
 - CLUSTERNODE 147, 151
 - code base 164
 - collocation 98, 102, 109, 156, 159, 179
 - command routing 138, 140
 - command-line interface 14, 140, 193
 - COMMRESTARTDURATION 146
 - COMMRESTARTINTERVAL 146
 - configuration files 177, 181, 187
 - configuration manager 139–140
 - configure DRM 195
 - copy groups 100–101, 103, 231
 - copy storage pool 13, 31, 98, 102, 112, 120–121, 142, 156–158, 170, 172, 180–181, 189, 193, 248
 - CPU power 187
 - CPU utilization 124
 - data flow 98
 - data movement 13

database 14, 31, 104–105, 113, 117–118, 140, 156, 173, 177, 191
 database backup 113, 121, 141–143, 154, 175–176, 181, 191, 197, 201
 database backup expiration 194, 197
 database mirroring 119, 123, 152–154, 173, 184, 188–189
 database page shadowing 118, 188
 database restore 181
 database size 187
 database snapshot backup 183
 DBMIRRORWRITE 188
 DBPAGESHADOW 118
 dedicated networks 112
 dedicated system 168
 DEVCONFIG 123
 device class 103, 141, 178, 316
 device configuration 31, 122–123, 177, 181–182, 187, 213
 device support 167
 device types 11, 121
 directory management class 104
 DIRMC 104
 Disaster Recovery 12, 18, 95
 Disaster Recovery plan file 31, 123, 141, 143, 190, 272, 275, 294
 Disaster Recovery tiers 31
 disk device class 105
 disk storage pool 105, 166, 178
 display DRM settings 198
 distributed configuration 138
 DRM 31, 113, 123, 134, 137, 143, 153, 166, 170, 176, 181, 187, 190, 231, 268, 272, 294
 electronic tape vaulting 153–154, 189
 electronic vaulting 142
 enterprise administration 15, 31, 138
 enterprise console 140
 event receiver 140
 expiration 101, 153, 157, 207
 Extended Edition 19, 132, 134, 142, 172
 failover 188
 FILE device class 177, 197
 geographic failover 189
 high availability 188
 HSM 13, 117, 148
 I/O requirements 113
 image backup 101, 124, 126, 133, 277
 incremental database backup 121
 incremental forever 17
 instant archive 127
 LAN-based clients 111, 127
 LAN-free backup 102, 111, 124, 128, 167, 172
 log sizing 187
 logical volume backup 133
 Logical Volume Storage Agent 134
 logmode 174, 176
 managed server 140
 management class 100–101, 103, 105, 110, 133, 141
 media states 199
 metadata 14, 117, 173
 MIGDELAY 106
 migration 98, 102, 110, 112, 157, 178
 monitoring 138
 multiple server instances 185, 187–188
 multiple tape drives 170
 NAS backup 131
 NDMP backup 132
 normal recovery log 174
 offsite data copy 93, 120, 122, 135, 153, 156, 170, 180, 189, 205
 Original Block File 134
 parallel backup 17
 partial write 118, 120, 189
 performance measurements 109
 planning 104
 platforms 11, 164
 point-in-time recovery 121, 133, 189
 policy 17, 96, 99, 110, 113, 120, 139, 178
 policy domain 100, 104, 141
 policy drivers 98
 policy sets 100
 PREPARE 206
 primary storage pool 13, 98, 102, 120, 142, 170, 178
 primary storage pools 156
 profile subscription 140
 progressive backup 16, 50, 101, 124
 rebuild server 181
 reclamation 98, 102, 109, 157–158, 170, 172, 179
 recovery log 31, 104, 117–118, 166, 173
 recovery log mirroring 119, 123, 184
 recovery log setup 122
 recovery log sizing 175
 recovery performance 99, 102, 109
 recovery scheduling 114
 regular DRM operations 199

- remote disk replication 151
- remote server 141
- resource contention 187
- resource utilization 114
- restore priorities 190
- restore time 172, 179
- retention policies 111, 170
- REUSEDELAY 194, 197
- ROLLFORWARD 174, 176
- rollforward recovery log 123, 174
- SAN exploitation 172
- scalability 164, 167
- scheduled operations 146, 150, 188
- Scheduler Service 147, 151
- scheduling 113, 191
- SCSI reserve 150
- security 18
- selective backup 125
- server 13, 117
- server consolidation 142
- server group 140
- server options 121, 123, 177, 187
- server policy 101
- server rebuild 166, 190
- server recovery 121, 135, 191, 207, 215
- server restart 189
- server sizing 113
- server-free backup 102, 124, 129–130, 167, 172
- server-to-server communication 31–32, 138–139, 177–178, 191, 198, 207, 316
- SHRSTATIC 231
- snapshot database backup 122, 175–177, 197
- snapshot image backup 134
- source server 140–141, 143
- split-mirror backup 126, 130
- SQL interface 173
- Storage Agent 128
- storage hierarchy 178, 189
- storage pool backup 121, 141, 153, 191, 193, 200
- storage pool hierarchy 13, 120, 180
- storage pool sizing 110
- storage pools 120, 156, 178, 189, 248
- synchronous storage pool 102
- tape device class 107
- tape device failover 189
- tape library 187
- tape library sharing 167
- tape library sizing 109
- tape storage pool 107, 110, 178–179
- target server 141, 143
- TCPSERVERADDRESS 184, 186
- throughput 167
- transaction 117, 173
- UNIX client 231
- vaulting 92, 153
- versioning 100, 110
- versions 110, 125
- virtual volumes 31, 33, 121, 138, 140, 142, 177–178, 191
- volume history 31, 122–123, 153, 177, 181, 183, 187, 211
- volume management 93, 169
- volume tracking 153, 190, 199
- Web browser interface 14, 140
- Windows 2000 APIs 268
- TSM commands
 - AUDIT LIBRARY 218
 - BACKUP DB 201
 - BACKUP DEVCONFIG 123
 - BACKUP STGPOOL 102, 200
 - BACKUP SYSTEMOBJECT 277
 - BACKUP VOLHISTORY 123
 - CHECKIN LIBVOL 207
 - COPY STGPOOL 121
 - DEFINE DEVCLASS 141
 - DEFINE MACHINE 272
 - DEFINE STGPOOL 193
 - DISMOUNT VOLUME 202
 - dsmfmt 212
 - EXPORT 141, 167
 - IMPORT 141, 167
 - INSERT MACHINE 272, 294, 348
 - MOVE DATA 110
 - MOVE DRMEDIA 196–197, 199, 203–204, 207
 - MOVE MEDIA 196
 - MOVE NODEDATA 110
 - PREPARE 195, 198, 206, 275, 294
 - QUERY DRMEDIA 196–197, 202, 204, 207
 - QUERY DRMSTATUS 198
 - QUERY LICENSE 193
 - QUERY MACHINE 274
 - QUERY MOUNT 202
 - QUERY RPFIL 143
 - QUERY STGPOOL 194
 - QUERY SYSTEMOBJECT 279, 284
 - REGISTER LICENSE 193

RESTORE DB 183
SET DRMCHECKLABEL 196
SET DRMCOPYSTGPOOL 196, 203
SET DRMCOURIERNAME 197, 204
SET DRMDBBACKUPEXPIREDAYS 197
SET DRMFILEPROCESS 197
SET DRMINSTRPREFIX 195
SET DRMNOTMOUNTABLENAME 197, 204
SET DRMPLANPREFIX 195, 206–207
SET DRMPLANVPOSTFIX 196
SET DRMPRIMSTGPOOL 196
SET DRMRPFEXPIREDAYS 198
SET DRMVaultNAME 198, 204
UPDATE MACHINE 275, 294
two-phase commit 27

U

UltraBac 19, 286
Unicenter 89
UNIX commands 233
 awk 208
 cpio 230
 dd 230
 df 232
 gzip 295
 pax 230
 tar 230, 295
UPS 81
USB 299

V

vaulting 24, 32, 61
VERITAS 19, 228, 286
virtualization 169
vital record retention, archive and retrieval 12

W

WAN 90
warmsite 10, 79, 90–92, 180–181, 315, 317
WebSphere Application Server 18
Windows 2000
 Active Directory 268, 284
 APIs 268, 276
 bare metal recovery 267
 configuration information 268, 272
 DFS 268
 Disk Management 268

domain controller 268, 284
event logs 284
installing 67
Logical Volume Storage Agent 134
NTDSUTIL 284
partition information 268, 271, 276
registry 276
Resource Kit 268
snapshot image backup 134
System Objects 268, 276, 284
SYSVOL 284
VBScript 208, 270, 273, 351
Windows 2000 commands
 diskmap 268, 271
 ipconfig 268
 msinfo32 268–270, 281, 351
 srvinfo 268
Windows Management Instrumentation 268
Windows XP 228
Winhdows 2000
 Active Directory 276
WORM 38

Z

zero data loss 28, 35
ZIP drive 121, 295, 299
zoning 189



Disaster Recovery Strategies with Tivoli Storage Management

(0.5" spine)
0.475" <-> 0.875"
250 <-> 459 pages



Disaster Recovery Strategies with Tivoli Storage Management



Redbooks

Keeping your TSM server and clients safe from disaster

How and why to build a disaster recovery plan

Testing your disaster recovery plan

Disasters, by their very nature, cannot be predicted, in either their intensity, timing, or effects. However, all enterprises can and should prepare for whatever might happen in order to protect themselves against loss of data or, worse, their entire business. It is too late to start preparing after a disaster occurs. This IBM Redbook will help you protect against a disaster — taking you step by step through the planning stages, with templates for sample documents. It explores the role that IBM Tivoli Storage Manager plays in disaster protection and recovery, from both the client and server side. Plus, it describes basic sample procedures for bare metal recovery of some popular operating systems, such as Windows 2000, AIX, Solaris, and Linux.

This book is written for any computing professional who is concerned about protecting their data and enterprise from disaster. It assumes you have basic knowledge of storage technologies and products, in particular, IBM Tivoli Storage Manager.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks