Tivoli software

IBM

IBM Tivoli Storage Manager in a Clustered Environment

- Learn how to build highly available Tivoli Storage Manager environments
- Covering Linux, IBM AIX, and Microsoft Windows solutions
- Understand all aspects of clustering

Roland Tretau
Dan Edwards
Werner Fischer
Marco Mencarelli
Maria Jose Rodriguez Canales
Rosane Goldstein Golubcic Langnor

Redbooks

ibm.com/redbooks

Archived

Archived

International Technical Support Organization

IBM Tivoli Storage Manager in a Clustered Environment

June 2005

Archived

SG24-6679-00

Note: Before using this information and the product it supports, read the information in "Notices" on page xlvii.

Archived

First Edition (June 2005)

This edition applies to IBM Tivoli Storage Manager Version 5.3.

© Copyright International Business Machines Corporation 2005. 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 xv

Tables xxxiii

Examples xxxv

Notices xlvii
Trademarks xlviii

Preface xlix
The team that wrote this redbook xlix
Become a published author lii
Comments welcome lii

Part 1. Highly available clusters with IBM Tivoli Storage Manager 1

Chapter 1. What does high availability imply? 3
1.1 High availability 4
1.1.1 Downtime 4
1.1.2 High availability concepts 5
1.1.3 High availability versus fault tolerance 6
1.1.4 High availability solutions 7
1.2 Cluster concepts 8
1.3 Cluster terminology 8

Chapter 2. Building a highly available Tivoli Storage Manager cluster environment 11
2.1 Overview of the cluster application 12
2.1.1 IBM Tivoli Storage Manager Version 5.3 12
2.1.2 IBM Tivoli Storage Manager for Storage Area Networks V5.3 14
2.2 Design to remove single points of failure 16
2.2.1 Storage Area Network considerations 16
2.2.2 LAN and network interface considerations 17
2.2.3 Private or heartbeat network considerations 17
2.3 Lab configuration 17
2.3.1 Cluster configuration matrix 19
2.3.2 Tivoli Storage Manager configuration matrix 20

Chapter 3. Testing a highly available Tivoli Storage Manager cluster environment 21

© Copyright IBM Corp. 2005. All rights reserved.

3.1 Objectives 22
3.2 Testing the clusters 22
3.2.1 Cluster infrastructure tests 23
3.2.2 Application tests 23

Part 2. Clustered Microsoft Windows environments and IBM Tivoli Storage Manager Version 5.3 25

Chapter 4. Microsoft Cluster Server setup 27
4.1 Overview 28
4.2 Planning and design 28
4.3 Windows 2000 MSCS installation and configuration 29
4.3.1 Windows 2000 lab setup 29
4.3.2 Windows 2000 MSCS setup 32
4.4 Windows 2003 MSCS installation and configuration 44
4.4.1 Windows 2003 lab setup 45
4.4.2 Windows 2003 MSCS setup 48
4.5 Troubleshooting 76

Chapter 5. Microsoft Cluster Server and the IBM Tivoli Storage Manager Server 77
5.1 Overview 78
5.2 Planning and design 78
5.3 Installing Tivoli Storage Manager Server on a MSCS 79
5.3.1 Installation of Tivoli Storage Manager server 80
5.3.2 Installation of Tivoli Storage Manager licenses 86
5.3.3 Installation of Tivoli Storage Manager device driver 89
5.3.4 Installation of the Administration Center 92
5.4 Tivoli Storage Manager server and Windows 2000 118
5.4.1 Windows 2000 lab setup 118
5.4.2 Windows 2000 Tivoli Storage Manager Server configuration 123
5.4.3 Testing the Server on Windows 2000 146
5.5 Configuring ISC for clustering on Windows 2000 167
5.5.1 Starting the Administration Center console 173
5.6 Tivoli Storage Manager Server and Windows 2003 179
5.6.1 Windows 2003 lab setup 180
5.6.2 Windows 2003 Tivoli Storage Manager Server configuration 184
5.6.3 Testing the server on Windows 2003 208
5.7 Configuring ISC for clustering on Windows 2003 231
5.7.1 Starting the Administration Center console 236

Chapter 6. Microsoft Cluster Server and the IBM Tivoli Storage Manager Client 241
6.1 Overview 242

6.2 Planning and design 242
6.3 Installing Tivoli Storage Manager client on MSCS 242
6.3.1 Installation of Tivoli Storage Manager client components 243
6.4 Tivoli Storage Manager client on Windows 2000 248
6.4.1 Windows 2000 lab setup 249
6.4.2 Windows 2000 Tivoli Storage Manager Client configuration 252
6.4.3 Testing Tivoli Storage Manager client on Windows 2000 MSCS . . 275
6.5 Tivoli Storage Manager Client on Windows 2003 289
6.5.1 Windows 2003 lab setup 289
6.5.2 Windows 2003 Tivoli Storage Manager Client configurations 292
6.5.3 Testing Tivoli Storage Manager client on Windows 2003 315
6.6 Protecting the quorum database 327

Chapter 7. Microsoft Cluster Server and the IBM Tivoli Storage Manager Storage Agent 329
7.1 Overview 330
7.2 Planning and design 330
7.2.1 System requirements 330
7.2.2 System information 331
7.3 Installing the Storage Agent on Windows MSCS 331
7.3.1 Installation of the Storage Agent 332
7.4 Storage Agent on Windows 2000 MSCS 333
7.4.1 Windows 2000 lab setup 333
7.4.2 Configuration of the Storage Agent on Windows 2000 MSCS 339
7.4.3 Testing Storage Agent high availability on Windows 2000 MSCS . 367
7.5 Storage Agent on Windows 2003 MSCS 378
7.5.1 Windows 2003 lab setup 378
7.5.2 Configuration of the Storage Agent on Windows 2003 MSCS 383
7.5.3 Testing the Storage Agent high availability 398

Part 3. AIX V5.3 with HACMP V5.2 environments and IBM Tivoli Storage Manager Version 5.3 415

Chapter 8. Establishing an HACMP infrastructure on AIX 417
8.1 Overview 418
8.1.1 AIX overview 418
8.2 HACMP overview 419
8.2.1 What is HACMP? 420
8.3 HACMP concepts 421
8.3.1 HACMP terminology 421
8.4 Planning and design 422
8.4.1 Supported hardware and software 422
8.4.2 Planning for networking 423
8.4.3 Plan for cascading versus rotating 426

8.5 Lab setup . . . 427
8.5.1 Pre-installation tasks . . . 430
8.5.2 Serial network setup . . . 433
8.5.3 External storage setup . . . 436
8.6 Installation . . . 441
8.6.1 Install the cluster code . . . 441
8.7 HACMP configuration . . . 442
8.7.1 Initial configuration of nodes . . . 443
8.7.2 Resource discovery . . . 445
8.7.3 Defining HACMP interfaces and devices . . . 445
8.7.4 Persistent addresses . . . 447
8.7.5 Further cluster customization tasks . . . 448

Chapter 9. AIX and HACMP with IBM Tivoli Storage Manager Server . . 451
9.1 Overview . . . 452
9.1.1 Tivoli Storage Manager Version 5.3 new features overview . . . 452
9.1.2 Planning for storage and database protection . . . 454
9.2 Lab setup . . . 455
9.3 Installation . . . 455
9.3.1 Tivoli Storage Manager Server AIX filesets . . . 455
9.3.2 Tivoli Storage Manager Client AIX filesets . . . 456
9.3.3 Tivoli Storage Manager Client Installation . . . 456
9.3.4 Installing the Tivoli Storage Manager Server software . . . 460
9.3.5 Installing the ISC and the Administration Center . . . 464
9.3.6 Installing Integrated Solutions Console Runtime . . . 465
9.3.7 Installing the Tivoli Storage Manager Administration Center . . . 472
9.3.8 Configure resources and resource groups . . . 478
9.3.9 Synchronize cluster configuration and make resource available . . 481
9.4 Tivoli Storage Manager Server configuration . . . 486
9.5 Testing . . . 495
9.5.1 Core HACMP cluster testing . . . 496
9.5.2 Failure during Tivoli Storage Manager client backup . . . 506
9.5.3 Tivoli Storage Manager server failure during LAN-free restore . . . 510
9.5.4 Failure during disk to tape migration operation . . . 515
9.5.5 Failure during backup storage pool operation . . . 517
9.5.6 Failure during database backup operation . . . 520
9.5.7 Failure during expire inventory process . . . 523

Chapter 10. AIX and HACMP with IBM Tivoli Storage Manager Client . . 527
10.1 Overview . . . 528
10.2 Clustering Tivoli Data Protection . . . 528
10.3 Planning and design . . . 529
10.4 Lab setup . . . 531

Archived

10.5 Installation 531
10.5.1 HACMP V5.2 installation 531
10.5.2 Tivoli Storage Manager Client Version 5.3 installation 531
10.5.3 Tivoli Storage Manager Server Version 5.3 installation 531
10.5.4 Integrated Solution Console and Administration Center 531
10.6 Configuration 532
10.7 Testing server and client system failure scenarios 536
10.7.1 Client system failover while the client is backing up to the disk storage pool 536
10.7.2 Client system failover while the client is backing up to tape 540
10.7.3 Client system failover while the client is backing up to tape with higher CommTimeOut 543
10.7.4 Client system failure while the client is restoring 550

Chapter 11. AIX and HACMP with the IBM Tivoli Storage Manager Storage Agent 555
11.1 Overview 556
11.2 Planning and design 557
11.2.1 Lab setup 560
11.3 Installation 560
11.4 Configuration 561
11.4.1 Configure tape storage subsystems 561
11.4.2 Configure resources and resource groups 562
11.4.3 Tivoli Storage Manager Storage Agent configuration 562
11.5 Testing the cluster 578
11.5.1 LAN-free client system failover while the client is backing up 578
11.5.2 LAN-free client system failover while the client is restoring 584

Part 4. Clustered IBM System Automation for Multiplatforms Version 1.2 environments and IBM Tivoli Storage Manager Version 5.3 591

Chapter 12. IBM Tivoli System Automation for Multiplatforms setup . . 593
12.1 Linux and Tivoli System Automation overview 594
12.1.1 Linux overview 594
12.1.2 IBM Tivoli System Automation for Multiplatform overview 595
12.1.3 Tivoli System Automation terminology 596
12.2 Planning and design 598
12.3 Lab setup 599
12.4 Preparing the operating system and drivers 600
12.4.1 Installation of host bus adapter drivers 600
12.4.2 Installation of disk multipath driver (RDAC) 602
12.4.3 Installation of the IBMtape driver 604
12.5 Persistent binding of disk and tape devices 605
12.5.1 SCSI addresses 605

12.5.2 Persistent binding of disk devices 606
12.6 Persistent binding of tape devices. 611
12.7 Installation of Tivoli System Automation 611
12.8 Creating a two-node cluster 612
12.9 Troubleshooting and tips 614

Chapter 13. Linux and Tivoli System Automation with IBM Tivoli Storage Manager Server 617
13.1 Overview 618
13.2 Planning storage 618
13.3 Lab setup. 619
13.4 Installation 619
13.4.1 Installation of Tivoli Storage Manager Server 620
13.4.2 Installation of Tivoli Storage Manager Client. 620
13.4.3 Installation of Integrated Solutions Console 621
13.4.4 Installation of Administration Center 623
13.5 Configuration 624
13.5.1 Preparing shared storage 624
13.5.2 Tivoli Storage Manager Server configuration 625
13.5.3 Cluster resources for Tivoli Storage Manager Server 629
13.5.4 Cluster resources for Administration Center 633
13.5.5 AntiAffinity relationship 635
13.6 Bringing the resource groups online 635
13.6.1 Verify configuration 635
13.6.2 Bringing Tivoli Storage Manager Server resource group online . . 637
13.6.3 Bringing Administration Center resource group online 639
13.7 Testing the cluster 639
13.7.1 Testing client incremental backup using the GUI 639
13.7.2 Testing a scheduled client backup 642
13.7.3 Testing migration from disk storage pool to tape storage pool . . . 645
13.7.4 Testing backup from tape storage pool to copy storage pool 647
13.7.5 Testing server database backup 649
13.7.6 Testing inventory expiration 651

Chapter 14. Linux and Tivoli System Automation with IBM Tivoli Storage Manager Client 653
14.1 Overview 654
14.2 Planning and design 655
14.3 Lab setup. 656
14.4 Installation 657
14.4.1 Tivoli System Automation V1.2 installation 657
14.4.2 Tivoli Storage Manager Client Version 5.3 installation 657
14.5 Configuration 657

14.5.1 Tivoli Storage Manager Client configuration 657
14.5.2 Tivoli Storage Manager client resource configuration 660
14.6 Testing the cluster 663
14.6.1 Testing client incremental backup 664
14.6.2 Testing client restore 668

Chapter 15. Linux and Tivoli System Automation with IBM Tivoli Storage Manager Storage Agent 673
15.1 Overview 674
15.2 Planning and design 674
15.3 Installation 674
15.4 Configuration 675
15.4.1 Storage agents 675
15.4.2 Client 681
15.4.3 Resource configuration for the Storage Agent 683
15.5 Testing the cluster 687
15.5.1 Backup 687
15.5.2 Restore 695

Part 5. Establishing a VERITAS Cluster Server Version 4.0 infrastructure on AIX with IBM Tivoli Storage Manager Version 5.3 701

Chapter 16. The VERITAS Cluster Server for AIX 703
16.1 Executive overview 704
16.2 Components of a VERITAS cluster 704
16.3 Cluster resources 705
16.4 Cluster configurations 708
16.5 Cluster communication 708
16.6 Cluster installation and setup 709
16.7 Cluster administration facilities 710
16.8 HACMP and VERITAS Cluster Server compared 710
16.8.1 Components of an HACMP cluster 711
16.8.2 Cluster resources 711
16.8.3 Cluster configurations 713
16.8.4 Cluster communications 713
16.8.5 Cluster installation and setup 714
16.8.6 Cluster administration facilities 715
16.8.7 HACMP and VERITAS Cluster Server high level feature comparison summary 716

Chapter 17. Preparing VERITAS Cluster Server environment 719
17.1 Overview 720
17.2 AIX overview 720
17.3 VERITAS Cluster Server 720

Archived

17.4 Lab environment 721
17.5 VCS pre-installation. 723
17.5.1 Preparing network connectivity 723
17.5.2 Installing the Atape drivers 724
17.5.3 Preparing the storage 725
17.5.4 Installing the VCS cluster software 731

Chapter 18. VERITAS Cluster Server on AIX and IBM Tivoli Storage Manager Server. 743
18.1 Overview 744
18.2 Installation of Tivoli Storage Manager Server 744
18.2.1 Tivoli Storage Manager Server AIX filesets 744
18.2.2 Tivoli Storage Manager Client AIX filesets 745
18.2.3 Tivoli Storage Manager Client Installation. 745
18.2.4 Installing the Tivoli Storage Manager server software 749
18.3 Configuration for clustering 753
18.3.1 Tivoli Storage Manager server configuration 754
18.4 Veritas Cluster Manager configuration 757
18.4.1 Preparing and placing application startup scripts 757
18.4.2 Service Group and Application configuration 763
18.5 Testing the cluster. 770
18.5.1 Core VCS cluster testing. 770
18.5.2 Node Power Failure 770
18.5.3 Start Service Group (bring online). 772
18.5.4 Stop Service Group (bring offline). 773
18.5.5 Manual Service Group switch 775
18.5.6 Manual fallback (switch back) 777
18.5.7 Public NIC failure 778
18.5.8 Failure of the server during a client backup 781
18.5.9 Failure of the server during a client scheduled backup. 785
18.5.10 Failure during disk to tape migration operation. 785
18.5.11 Failure during backup storage pool operation. 787
18.5.12 Failure during database backup operation 791

Chapter 19. VERITAS Cluster Server on AIX with the IBM Tivoli Storage Manager StorageAgent. 793
19.1 Overview 794
19.2 Planning and design 795
19.3 Lab setup. 797
19.4 Tivoli Storage Manager Storage Agent installation 797
19.5 Storage agent configuration 798
19.6 Configuring a cluster application 804
19.7 Testing. 810

19.7.1 Veritas Cluster Server testing 810
19.7.2 Node power failure 811
19.7.3 Start Service Group (bring online). 812
19.7.4 Stop Service Group (bring offline). 814
19.7.5 Manual Service Group switch 817
19.7.6 Manual fallback (switch back) 820
19.7.7 Public NIC failure 822
19.7.8 LAN-free client system failover while the client is backing up. . . . 824
19.7.9 LAN-free client failover while the client is restoring. 831

Chapter 20. VERITAS Cluster Server on AIX with IBM Tivoli Storage Manager Client and ISC applications 839
20.1 Overview 840
20.2 Planning 840
20.3 Tivoli Storage Manager client installation 841
20.3.1 Preparing the client for high availability. 841
20.4 Installing the ISC and the Administration Center. 842
20.5 Veritas Cluster Manager configuration 857
20.5.1 Preparing and placing application startup scripts 857
20.5.2 Configuring Service Groups and applications 865
20.6 Testing the highly available client and ISC 870
20.6.1 Cluster failure during a client back up 870
20.6.2 Cluster failure during a client restore 873

Part 6. Establishing a VERITAS Cluster Server Version 4.0 infrastructure on Windows with IBM Tivoli Storage Manager Version 5.3 877

Chapter 21. Installing the VERITAS Storage Foundation HA for Windows environment. 879
21.1 Overview 880
21.2 Planning and design 880
21.3 Lab environment 880
21.4 Before VSFW installation 882
21.4.1 Installing Windows 2003 882
21.4.2 Preparing network connectivity 883
21.4.3 Domain membership 883
21.4.4 Setting up external shared disks 884
21.5 Installing the VSFW software 887
21.6 Configuring VERITAS Cluster Server 896
21.7 Troubleshooting. 902

Chapter 22. VERITAS Cluster Server and the IBM Tivoli Storage Manager Server. 903
22.1 Overview 904

Archived

22.2 Planning and design 904
22.3 Lab setup. 904
22.3.1 Installation of IBM tape device drivers 908
22.4 Tivoli Storage Manager installation 909
22.5 Configuration of Tivoli Storage Manager for VCS 909
22.5.1 Configuring Tivoli Storage Manager on the first node. 909
22.5.2 Configuring Tivoli Storage Manager on the second node. 919
22.6 Creating service group in VCS 920
22.7 Testing the Cluster 932
22.8 IBM Tivoli Storage Manager Administrative Center 933
22.8.1 Installing the Administrative Center in a clustered environment . . 933
22.8.2 Creating the service group for the Administrative Center 933
22.9 Configuring Tivoli Storage Manager devices. 945
22.10 Testing the Tivoli Storage Manager on VCS 945
22.10.1 Testing incremental backup using the GUI client 945
22.10.2 Testing a scheduled incremental backup 948
22.10.3 Testing migration from disk storage pool to tape storage pool . . 952
22.10.4 Testing backup from tape storage pool to copy storage pool . . . 955
22.10.5 Testing server database backup 960

Chapter 23. VERITAS Cluster Server and the IBM Tivoli Storage Manager Client 965
23.1 Overview 966
23.2 Planning and design 966
23.3 Lab setup. 967
23.4 Installation of the backup/archive client. 968
23.5 Configuration 969
23.5.1 Configuring Tivoli Storage Manager client on local disks 969
23.5.2 Configuring Tivoli Storage Manager client on shared disks 969
23.6 Testing Tivoli Storage Manager client on the VCS 988
23.6.1 Testing client incremental backup. 989
23.6.2 Testing client restore. 993
23.7 Backing up VCS configuration files 997

Chapter 24. VERITAS Cluster Server and the IBM Tivoli Storage Manager Storage Agent 999
24.1 Overview 1000
24.2 Planning and design 1000
24.2.1 System requirements 1000
24.2.2 System information 1001
24.3 Lab setup. 1001
24.3.1 Tivoli Storage Manager LAN-free configuration details. 1002
24.4 Installation 1004

24.5 Configuration 1004
24.5.1 Configuration of Tivoli Storage Manager server for LAN-free . . . 1005
24.5.2 Configuration of the Storage Agent for local nodes 1006
24.5.3 Configuration of the Storage Agent for virtual nodes 1010
24.6 Testing Storage Agent high availability 1015
24.6.1 Testing LAN-free client incremental backup 1015
24.6.2 Testing client restore 1021

Part 7. Appendixes 1027

Appendix A. Additional material 1029
Locating the Web material 1029
Using the Web material 1029
Requirements for downloading the Web material 1030
How to use the Web material 1030

Glossary 1031

Abbreviations and acronyms 1039

Related publications 1047
IBM Redbooks 1047
Other publications 1047
Online resources 1049
How to get IBM Redbooks 1050
Help from IBM 1051

Index 1053

Archived

Archived

Figures

2-1 Tivoli Storage Manager LAN (Metadata) and SAN data flow diagram . . 15
2-2 Multiple clients connecting through a single Storage Agent 16
2-3 Cluster Lab SAN and heartbeat networks 18
2-4 Cluster Lab LAN and hearbeat configuration 19
4-1 Windows 200 MSCS configuration 29
4-2 Network connections windows with renamed icons 32
4-3 Recommended bindings order 33
4-4 LUN configuration for Windows 2000 MSCS 35
4-5 Device manager with disks and SCSI adapters 36
4-6 New partition wizard 37
4-7 Select all drives for signature writing 37
4-8 Do not upgrade any of the disks 38
4-9 Select primary partition 38
4-10 Select the size of the partition 39
4-11 Drive mapping 39
4-12 Format partition 40
4-13 Disk configuration 40
4-14 Cluster Administrator after end of installation 43
4-15 Cluster Administrator with TSM Group 43
4-16 Windows 2003 MSCS configuration 45
4-17 Network connections windows with renamed icons 48
4-18 Recommended bindings order 49
4-19 LUN configuration for our Windows 2003 MSCS 51
4-20 Device manager with disks and SCSI adapters 52
4-21 Disk initialization and conversion wizard 53
4-22 Select all drives for signature writing 53
4-23 Do not upgrade any of the disks 54
4-24 Successfull completion of the wizard 54
4-25 Disk manager after disk initialization 55
4-26 Create new partition 55
4-27 New partition wizard 56
4-28 Select primary partition 56
4-29 Select the size of the partition 57
4-30 Drive mapping 57
4-31 Format partition 58
4-32 Completing the New Partition wizard 58
4-33 Disk configuration 59
4-34 Open connection to cluster 60

© Copyright IBM Corp. 2005. All rights reserved.

4-35 New Server Cluster wizard (prerequisites listed) 60
4-36 Clustername and domain 61
4-37 Warning message 61
4-38 Select computer 62
4-39 Review the messages 62
4-40 Warning message 63
4-41 Cluster IP address 63
4-42 Specify username and password of the cluster service account 64
4-43 Summary menu 64
4-44 Selecting the quorum disk 65
4-45 Cluster creation 65
4-46 Wizard completed 66
4-47 Cluster administrator 66
4-48 Add cluster nodes 67
4-49 Node analysis 68
4-50 Specify the password 68
4-51 Summary information 69
4-52 Node analysis 69
4-53 Setup complete 70
4-54 Private network properties 71
4-55 Configuring the heartbeat 71
4-56 Public network properties 72
4-57 Configuring the public network 72
4-58 Cluster properties 73
4-59 Network priority 73
4-60 Cluster Administrator after end of installation 74
4-61 Moving resources 75
4-62 Final configuration 75
5-1 IBM Tivoli Storage Manager InstallShield wizard 80
5-2 Language select 81
5-3 Main menu 81
5-4 Install Products menu 82
5-5 Installation wizard 83
5-6 Licence agreement 83
5-7 Customer information 84
5-8 Setup type 84
5-9 Beginning of installation 85
5-10 Progress bar 85
5-11 Successful installation 86
5-12 Reboot message 86
5-13 Install Products menu 87
5-14 License installation 87
5-15 Ready to install the licenses 88

Archived

5-16 Installation completed 88
5-17 Install Products menu 89
5-18 Welcome to installation wizard. 90
5-19 Ready to install 90
5-20 Restart the server 91
5-21 InstallShield wizard for IBM Integrated Solutions Console 93
5-22 Welcome menu 93
5-23 ISC License Agreement 94
5-24 Location of the installation CD 95
5-25 Installation path for ISC 96
5-26 Selecting user id and password for the ISC. 97
5-27 Selecting Web administration ports 98
5-28 Review the installation options for the ISC 99
5-29 Welcome 100
5-30 Installation progress bar. 101
5-31 ISC Installation ends 102
5-32 ISC services started for the first node of the MSCS 103
5-33 Administration Center Welcome menu 104
5-34 Administration Center Welcome 105
5-35 Administration Center license agreement 106
5-36 Modifying the default options 107
5-37 Updating the ISC installation path 108
5-38 Web administration port 109
5-39 Selecting the administrator user id. 110
5-40 Specifying the password for the iscadmin user id 111
5-41 Location of the administration center code 112
5-42 Reviewing the installation options 113
5-43 Installation progress bar for the Administration Center 114
5-44 Administration Center installation ends 115
5-45 Main Administration Center menu 116
5-46 ISC Services started as automatic in the second node 117
5-47 Windows 2000 Tivoli Storage Manager clustering server configuration119
5-48 Cluster Administrator with TSM Group 122
5-49 Successful installation of IBM 3582 and IBM 3580 device drivers. . . . 123
5-50 Cluster resources 124
5-51 Starting the Tivoli Storage Manager management console 124
5-52 Initial Configuration Task List. 125
5-53 Welcome Configuration wizard 126
5-54 Initial configuration preferences 126
5-55 Site environment information 127
5-56 Initial configuration 127
5-57 Welcome Performance Environment wizard 128
5-58 Performance options 128

Archived

5-59 Drive analysis 129
5-60 Performance wizard 129
5-61 Server instance initialization wizard 130
5-62 Cluster environment detection 130
5-63 Cluster group selection 131
5-64 Server initialization wizard 131
5-65 Server volume location 132
5-66 Server service logon parameters 133
5-67 Server name and password 133
5-68 Completing the Server Initialization wizard 134
5-69 Completing the server installation wizard 134
5-70 Tivoli Storage Manager Server has been initialized 135
5-71 Cluster configuration wizard 135
5-72 Select the cluster group 136
5-73 Tape failover configuration 137
5-74 IP address 137
5-75 Network name 138
5-76 Completing the Cluster configuration wizard 138
5-77 End of Tivoli Storage Manager cluster configuration 139
5-78 Tivoli Storage Manager console 139
5-79 Cluster resources 140
5-80 Cluster configuration wizard 141
5-81 Cluster group selection 141
5-82 Completing the cluster configuration wizard (I) 142
5-83 Completing the cluster configuration wizard (II) 142
5-84 Successful installation 143
5-85 Tivoli Storage Manager Group resources 143
5-86 Bringing resources online 144
5-87 Tivoli Storage Manager Group resources online 145
5-88 Services overview 146
5-89 Cluster Administrator shows resources on RADON 147
5-90 Selecting a client backup using the GUI 148
5-91 Transferring files to the server 148
5-92 Reopening the session 149
5-93 Transfer of data goes on when the server is restarted 149
5-94 Defining a new resource for IBM WebSphere application server 168
5-95 Specifying a resource name for IBM WebSphere application server . . 169
5-96 Possible owners for the IBM WebSphere application server resource 169
5-97 Dependencies for the IBM WebSphere application server resource . . 170
5-98 Specifying the same name for the service related to IBM WebSphere 171
5-99 Registry replication values 171
5-100 Successful creation of the generic resource 172
5-101 Selecting the resource name for ISC Help Service 172

5-102 Login menu for the Administration Center 173
5-103 Administration Center 174
5-104 Options for Tivoli Storage Manager 175
5-105 Selecting to create a new server connection 176
5-106 Specifying Tivoli Storage Manager server parameters 177
5-107 Filling in a form to unlock ADMIN_CENTER 178
5-108 TSMSRV01 Tivoli Storage Manager server created 179
5-109 Lab setup for a 2-node cluster 180
5-110 Cluster Administrator with TSM Group 183
5-111 3582 and 3580 drivers installed 184
5-112 Cluster resources 185
5-113 Starting the Tivoli Storage Manager management console 186
5-114 Initial Configuration Task List 187
5-115 Welcome Configuration wizard 187
5-116 Initial configuration preferences 188
5-117 Site environment information 188
5-118 Initial configuration 189
5-119 Welcome Performance Environment wizard 189
5-120 Performance options 190
5-121 Drive analysis 190
5-122 Performance wizard 191
5-123 Server instance initialization wizard 191
5-124 Cluster environment detection 192
5-125 Cluster group selection 192
5-126 Server initialization wizard 193
5-127 Server volume location 194
5-128 Server service logon parameters 194
5-129 Server name and password 195
5-130 Completing the Server Initialization wizard 196
5-131 Completing the server installation wizard 196
5-132 Tivoli Storage Manager Server has been initialized 197
5-133 Cluster configuration wizard 197
5-134 Select the cluster group 198
5-135 Tape failover configuration 199
5-136 IP address 199
5-137 Network Name 200
5-138 Completing the Cluster configuration wizard 200
5-139 End of Tivoli Storage Manager Cluster configuration 201
5-140 Tivoli Storage Manager console 201
5-141 Cluster resources 202
5-142 Cluster configuration wizard 203
5-143 Selecting the cluster group 203
5-144 Completing the Cluster Configuration wizard 204

Archived

5-145 The wizard starts the cluster configuration 204
5-146 Successful installation 205
5-147 TSM Group resources 205
5-148 Bringing resources online 206
5-149 TSM Group resources online 206
5-150 Services 207
5-151 Cluster Administrator shows resources on SENEGAL 208
5-152 Selecting a client backup using the GUI 209
5-153 Transferring files to the server 209
5-154 Reopening the session 210
5-155 Transfer of data goes on when the server is restarted 210
5-156 Schedule result 215
5-157 Defining a new resource for IBM WebSphere Application Server 232
5-158 Specifying a resource name for IBM WebSphere application server. . 232
5-159 Possible owners for the IBM WebSphere application server resource 233
5-160 Dependencies for the IBM WebSphere application server resource . . 233
5-161 Specifying the same name for the service related to IBM WebSphere 234
5-162 Registry replication values 235
5-163 Successful creation of the generic resource 235
5-164 Selecting the resource name for ISC Help Service 236
5-165 Login menu for the Administration Center 237
5-166 Administration Center 237
5-167 Options for Tivoli Storage Manager 238
5-168 Selecting to create a new server connection 238
5-169 Specifying Tivoli Storage Manager server parameters 239
5-170 Filling a form to unlock ADMIN_CENTER 240
5-171 TSMSRV03 Tivoli Storage Manager server created 240
6-1 Setup language menu 243
6-2 InstallShield Wizard for Tivoli Storage Manager Client 244
6-3 Installation path for Tivoli Storage Manager client 245
6-4 Custom installation. 245
6-5 Custom setup. 246
6-6 Start of installation of Tivoli Storage Manager client 246
6-7 Status of the installation. 247
6-8 Installation completed 247
6-9 Installation prompts to restart the server 248
6-10 Tivoli Storage Manager backup/archive clustering client (Win.2000) . 249
6-11 Tivoli Storage Manager client services 253
6-12 Generating the password in the registry 257
6-13 Result of Tivoli Storage Manager scheduler service installation 258
6-14 Creating new resource for Tivoli Storage Manager scheduler service. 260
6-15 Definition of TSM Scheduler generic service resource 260
6-16 Possible owners of the resource 261

6-17 Dependencies 261
6-18 Generic service parameters 262
6-19 Registry key replication 263
6-20 Successful cluster resource installation 263
6-21 Bringing online the Tivoli Storage Manager scheduler service 264
6-22 Cluster group resources online 264
6-23 Windows service menu 265
6-24 Installing the Client Acceptor service in the Cluster Group 267
6-25 Successful installation, Tivoli Storage Manager Remote Client Agent 268
6-26 New resource for Tivoli Storage Manager Client Acceptor service . . . 270
6-27 Definition of TSM Client Acceptor generic service resource 270
6-28 Possible owners of the TSM Client Acceptor generic service 271
6-29 Dependencies for TSM Client Acceptor generic service 271
6-30 TSM Client Acceptor generic service parameters 272
6-31 Bringing online the TSM Client Acceptor generic service 272
6-32 TSM Client Acceptor generic service online 273
6-33 Windows service menu 273
6-34 Windows 2000 filespace names for local and virtual nodes 275
6-35 Resources hosted by RADON in the Cluster Administrator 276
6-36 Event log shows the schedule as restarted 280
6-37 Schedule completed on the event log 281
6-38 Windows explorer 282
6-39 Checking backed up files using the TSM GUI 283
6-40 Scheduled restore started for CL_MSCS01_SA 284
6-41 Schedule restarted on the event log for CL_MSCS01_SA 288
6-42 Event completed for schedule name RESTORE 289
6-43 Tivoli Storage Manager backup/archive clustering client (Win.2003) . 290
6-44 Tivoli Storage Manager client services 294
6-45 Generating the password in the registry 298
6-46 Result of Tivoli Storage Manager scheduler service installation 299
6-47 Creating new resource for Tivoli Storage Manager scheduler service. 300
6-48 Definition of TSM Scheduler generic service resource 301
6-49 Possible owners of the resource 301
6-50 Dependencies 302
6-51 Generic service parameters 302
6-52 Registry key replication 303
6-53 Successful cluster resource installation 303
6-54 Bringing online the Tivoli Storage Manager scheduler service 304
6-55 Cluster group resources online 304
6-56 Windows service menu 305
6-57 Installing the Client Acceptor service in the Cluster Group 307
6-58 Successful installation, Tivoli Storage Manager Remote Client Agent 308
6-59 New resource for Tivoli Storage Manager Client Acceptor service . . . 310

6-60 Definition of TSM Client Acceptor generic service resource 310
6-61 Possible owners of the TSM Client Acceptor generic service 311
6-62 Dependencies for TSM Client Acceptor generic service 311
6-63 TSM Client Acceptor generic service parameters 312
6-64 Bringing online the TSM Client Acceptor generic service 313
6-65 TSM Client Acceptor generic service online 313
6-66 Windows service menu 314
6-67 Windows 2003 filespace names for local and virtual nodes 315
6-68 Resources hosted by SENEGAL in the Cluster Administrator 316
6-69 Scheduled incremental backup started for CL_MSCS02_TSM 317
6-70 Schedule log file: incremental backup starting for CL_MSCS02_TSM 317
6-71 CL_MSCS02_TSM loss its connection with the server 318
6-72 The schedule log file shows an interruption of the session 318
6-73 Schedule log shows how the incremental backup restarts 319
6-74 Attributes changed for node CL_MSCS02_TSM 319
6-75 Event log shows the incremental backup schedule as restarted 320
6-76 Schedule INCR_BCK completed successfully 320
6-77 Schedule completed on the event log 320
6-78 Windows explorer 321
6-79 Checking backed up files using the TSM GUI 322
6-80 Scheduled restore started for CL_MSCS02_TSM 323
6-81 Restore starts in the schedule log file for CL_MSCS02_TSM 323
6-82 Restore session is lost for CL_MSCS02_TSM 324
6-83 Schedule log file shows an interruption for the restore operation 324
6-84 Attributes changed from node CL_MSCS02_TSM to SENEGAL 324
6-85 Restore session starts from the beginning in the schedule log file . . . 325
6-86 Schedule restarted on the event log for CL_MSCS02_TSM 325
6-87 Statistics for the restore session 326
6-88 Schedule name RESTORE completed for CL_MSCS02_TSM 326
7-1 Install TSM Storage Agent 332
7-2 Windows 2000 TSM Storage Agent clustering configuration 334
7-3 Updating the driver 338
7-4 Device Manager menu after updating the drivers 339
7-5 Choosing RADON for LAN-free backup 342
7-6 Enable LAN-free Data Movement wizard for RADON 343
7-7 Allowing LAN and LAN-free operations for RADON 344
7-8 Creating a new Storage Agent 345
7-9 Storage agent parameters for RADON 346
7-10 Storage pool selection for LAN-free backup 347
7-11 Modify drive paths for Storage Agent RADON_STA 348
7-12 Specifying the device name from the operating system view 349
7-13 Device names for 3580 tape drives attached to RADON 350
7-14 LAN-free configuration summary 351

7-15 Initialization of a local Storage Agent 352
7-16 Specifying parameters for Storage Agent 352
7-17 Specifying parameters for the Tivoli Storage Manager server 353
7-18 Specifying the account information 354
7-19 Completing the initialization wizard 354
7-20 Granted access for the account 355
7-21 Storage agent is successfully initialized 355
7-22 TSM StorageAgent1 is started on RADON 356
7-23 Installing Storage Agent for LAN-free backup of shared disk drives . . 358
7-24 Installing the service related to StorageAgent2 359
7-25 Management console displays two Storage Agents 359
7-26 Starting the TSM StorageAgent2 service in POLONIUM 360
7-27 TSM StorageAgent2 installed in RADON 361
7-28 Use cluster administrator to create resource for TSM StorageAgent2 362
7-29 Defining a generic service resource for TSM StorageAgent2 362
7-30 Possible owners for TSM StorageAgent2 363
7-31 Dependencies for TSM StorageAgent2 363
7-32 Service name for TSM StorageAgent2 364
7-33 Registry key for TSM StorageAgent2 364
7-34 Generic service resource created successfully:TSM StorageAgent2 . 365
7-35 Bringing the TSM StorageAgent2 resource online 365
7-36 Adding Storage Agent resource as dependency for TSM Scheduler . 366
7-37 Storage agent CL_MSCS01_STA session for tape library sharing . . . 368
7-38 A tape volume is mounted and the Storage Agent starts sending data 368
7-39 Client starts sending files to the TSM server in the schedule log file . . 369
7-40 Sessions for TSM client and Storage Agent are lost in the activity log 369
7-41 Both Storage Agent and TSM client restart sessions in second node . 370
7-42 Tape volume is dismounted by the Storage Agent 371
7-43 The scheduled is restarted and the tape volume mounted again 371
7-44 Final statistics for LAN-free backup 372
7-45 Starting restore session for LAN-free 374
7-46 Restore starts on the schedule log file 374
7-47 Both sessions for the Storage Agent and the client lost in the server . 375
7-48 Resources are started again in the second node 375
7-49 Tape volume is dismounted by the Storage Agent 376
7-50 The tape volume is mounted again by the Storage Agent 376
7-51 Final statistics for the restore on the schedule log file 377
7-52 Windows 2003 Storage Agent configuration 378
7-53 Tape devices in device manager page 382
7-54 Device Manager page after updating the drivers 382
7-55 Modifying the devconfig option to point to devconfig file in dsmsta.opt 384
7-56 Specifying parameters for the Storage Agent 385
7-57 Specifying parameters for the Tivoli Storage Manager server 386

7-58 Specifying the account information 387
7-59 Storage agent initialized. 387
7-60 TSM StorageAgent1 is started. 388
7-61 Installing Storage Agent for LAN-free backup of shared disk drives . . 390
7-62 Installing the service attached to StorageAgent2. 390
7-63 Management console displays two Storage Agents 391
7-64 Starting the TSM StorageAgent2 service in SENEGAL. 391
7-65 TSM StorageAgent2 installed in TONGA. 392
7-66 Use cluster administrator to create a resource: TSM StorageAgent2 . 393
7-67 Defining a generic service resource for TSM StorageAgent2 393
7-68 Possible owners for TSM StorageAgent2 394
7-69 Dependencies for TSM StorageAgent2 394
7-70 Service name for TSM StorageAgent2 395
7-71 Registry key for TSM StorageAgent2 395
7-72 Generic service resource created successfully: TSM StorageAgent2 . 396
7-73 Bringing the TSM StorageAgent2 resource online. 396
7-74 Adding Storage Agent resource as dependency for TSM Scheduler . 397
7-75 Storage agent CL_MSCS02_STA mounts tape for LAN-free backup . 399
7-76 Client starts sending files to the TSM server in the schedule log file . . 399
7-77 Sessions for TSM client and Storage Agent are lost in the activity log 400
7-78 Connection is lost in the client while the backup is running 400
7-79 Both Storage Agent and TSM client restart sessions in second node . 401
7-80 Tape volume is dismounted and mounted again by the server 401
7-81 The scheduled is restarted and the tape volume mounted again 402
7-82 Final statistics for LAN-free backup 403
7-83 Activity log shows tape volume is dismounted when backup ends . . . 404
7-84 Starting restore session for LAN-free. 406
7-85 Restore starts on the schedule log file. 407
7-86 Storage agent shows sessions for the server and the client 407
7-87 Both sessions for the Storage Agent and the client lost in the server . 408
7-88 Resources are started again in the second node 409
7-89 Storage agent commands the server to dismount the tape volume. . . 409
7-90 Storage agent writes to the volume again 410
7-91 The client restarts the restore 410
7-92 Final statistics for the restore on the schedule log file 411
7-93 Restore completed and volume dismounted by the server in actlog . . 412
8-1 HACMP cluster 420
8-2 AIX Clusters - SAN (Two fabrics) and network 427
8-3 Logical layout for AIX and TSM filesystems, devices, and network. . . 428
8-4 9-pin D shell cross cable example 434
8-5 tty configuration 435
8-6 DS4500 configuration layout. 437
8-7 boot address configuration. 443

8-8 Define cluster example. 444
8-9 An add cluster node example 445
8-10 Configure HACMP Communication Interfaces/Devices panel 446
8-11 Selecting communication interfaces. 447
8-12 The Add a Persistent Node IP Label/Address panel 448
9-1 The smit install and update panel 457
9-2 Launching SMIT from the source directory, only dot (.) is required . . . 457
9-3 AIX installp filesets chosen: Tivoli Storage Manager client installation 458
9-4 Changing the defaults to preview with detail first prior to installing . . . 459
9-5 The smit panel demonstrating a detailed and committed installation . 459
9-6 AIX lslpp command to review the installed filesets 460
9-7 The smit software installation panel. 460
9-8 The smit input device panel 461
9-9 The smit selection screen for Tivoli Storage Manager filesets. 462
9-10 The smit screen showing non-default values for a detailed preview . . 463
9-11 The final smit install screen with selections and a commit installation. 463
9-12 AIX lslpp command listing of the server installp images 464
9-13 ISC installation screen 467
9-14 ISC installation screen, license agreement 467
9-15 ISC installation screen, source path 468
9-16 ISC installation screen, target path - our shared disk for this node . . . 469
9-17 ISC installation screen, establishing a login and password 470
9-18 ISC installation screen establishing the ports which will be used 470
9-19 ISC installation screen, reviewing selections and disk space required 471
9-20 ISC installation screen showing completion. 471
9-21 ISC installation screen, final summary providing URL for connection . 472
9-22 Service address configuration 479
9-23 Add a resource group 480
9-24 Add resources to the resource group. 481
9-25 Cluster resources synchronization. 482
9-26 Starting cluster services. 483
9-27 X11 clstat example. 484
9-28 clstat output 484
9-29 WebSMIT version of clstat example 485
9-30 Check for available resources 485
9-31 The Add a Custom Application Monitor panel 495
9-32 Clstop with takeover. 499
10-1 HACMP application server configuration for the client's start and stop 535
11-1 Start Server to Server Communication wizard. 563
11-2 Setting Tivoli Storage Manager server password and address 563
11-3 Select targeted server and View Enterprise Properties 564
11-4 Define Server chose under Servers section 564
11-5 Entering Storage Agent name, password, and description 565

11-6 Insert communication data 565
11-7 Click Next on Virtual Volumes panel 566
11-8 Summary panel 566
11-9 Share the library and set resetdrives to yes. 568
11-10 Define drive path panel 568
13-1 Logical drive mapping for cluster volumes 625
13-2 Selecting client backup using the GUI 640
13-3 Transfer of files starts 640
13-4 Reopening Session 641
13-5 Transferring of files continues to the second node 642
15-1 Selecting the server in the Enterprise Management panel 676
15-2 Servers and Server Groups defined to TSMSRV03 676
15-3 Define a Server - step one 677
15-4 Define a Server - step two 677
15-5 Define a Server - step three 678
15-6 Define a Server - step four 678
15-7 Define a Server - step five 679
17-1 cl_veritas01 cluster physical resource layout. 722
17-2 Network, SAN (dual fabric), and Heartbeat logical layout 723
17-3 Atlantic zoning 725
17-4 Banda zoning 726
17-5 DS4500 LUN configuration for cl_veritas01 726
17-6 Veritas Cluster Server 4.0 Installation Program 731
17-7 VCS system check results 732
17-8 Summary of the VCS Infrastructure fileset installation. 732
17-9 License key entry screen 733
17-10 Choice of which filesets to install 733
17-11 Summary of filesets chosen to install. 734
17-12 VCS configuration prompt screen 736
17-13 VCS installation screen instructions. 736
17-14 VCS cluster configuration screen. 737
17-15 VCS screen reviewing the cluster information to be set 737
17-16 VCS setup screen to set a non-default password for the admin user . 737
17-17 VCS adding additional users screen 738
17-18 VCS summary for the privileged user and password configuration . . . 738
17-19 VCS prompt screen to configure the Cluster Manager Web console . 738
17-20 VCS screen summarizing Cluster Manager Web Console settings. . . 739
17-21 VCS screen prompt to configure SNTP notification 739
17-22 VCS screen prompt to configure SNMP notification 739
17-23 VCS prompt for a simultaneous installation of both nodes 740
17-24 VCS completes the server configuration successfully 741
17-25 Results screen for starting the cluster server processes 742
17-26 Final VCS installation screen 742

18-1 The smit install and update panel 746
18-2 Launching SMIT from the source directory, only dot (.) is required . . . 746
18-3 AIX installp filesets chosen for client installation 747
18-4 Changing the defaults to preview with detail first prior to installing . . . 748
18-5 The smit panel demonstrating a detailed and committed installation . 748
18-6 AIX lslpp command to review the installed filesets 749
18-7 The smit software installation panel 749
18-8 The smit input device panel 750
18-9 The smit selection screen for filesets 751
18-10 The smit screen showing non-default values for a detailed preview . . 752
18-11 The final smit install screen with selections and a commit installation. 752
18-12 AIX lslpp command listing of the server installp images 753
18-13 Child-parent relationships within the sg_tsmsrv Service Group. 767
18-14 VCS Cluster Manager GUI switching Service Group to another node. 776
18-15 Prompt to confirm the switch 776
19-1 Administration Center screen to select drive paths 800
19-2 Administration Center screen to add a drive path 801
19-3 Administration Center screen to define DRLTO_1. 801
19-4 Administration Center screen to review completed adding drive path . 802
19-5 Administration Center screen to define a second drive path 803
19-6 Administration Center screen to define a second drive path mapping. 803
19-7 Veritas Cluster Manager GUI, sg_isc_sta_tsmcli resource relationship808
19-8 VCS Cluster Manager GUI switching Service Group to another node. 818
19-9 Prompt to confirm the switch 819
20-1 ISC installation screen 844
20-2 ISC installation screen, license agreement 844
20-3 ISC installation screen, source path 845
20-4 ISC installation screen, target path - our shared disk for this node . . . 846
20-5 ISC installation screen, establishing a login and password 847
20-6 ISC installation screen establishing the ports which will be used 847
20-7 ISC installation screen, reviewing selections and disk space required 848
20-8 ISC installation screen showing completion. 849
20-9 ISC installation screen, final summary providing URL for connection . 849
20-10 Welcome wizard screen. 851
20-11 Review of AC purpose and requirements 851
20-12 AC Licensing panel 852
20-13 Validation of the ISC installation environment 852
20-14 Prompting for the ISC userid and password 853
20-15 AC installation source directory 854
20-16 AC target source directory 854
20-17 AC progress screen 855
20-18 AC successful completion 855
20-19 Summary and review of the port and URL to access the AC. 856

20-20 Final AC screen 856
20-21 GUI diagram, child-parent relation, sg_isc_sta_tsmcli Service Group . 869
21-1 Windows 2003 VSFW configuration 881
21-2 Network connections 883
21-3 LUN configuration 885
21-4 Device manager with disks and SCSI adapters 886
21-5 Choosing the product to install 888
21-6 Choose complete installation 888
21-7 Pre-requisites - attention to the driver signing option 889
21-8 License agreement 889
21-9 License key 890
21-10 Common program options 890
21-11 Global cluster option and agents 891
21-12 Install the client components 891
21-13 Choosing the servers and path 892
21-14 Testing the installation 892
21-15 Summary of the installation 893
21-16 Installation progress on both nodes 893
21-17 Install report 894
21-18 Reboot remote server 894
21-19 Remote server online 895
21-20 Installation complete 895
21-21 Start cluster configuration 896
21-22 Domain and user selection 897
21-23 Create new cluster 897
21-24 Cluster information 898
21-25 Node validation 898
21-26 NIC selection for private communication 899
21-27 Selection of user account 899
21-28 Password information 900
21-29 Setting up secure or non secure cluster 900
21-30 Summary prior to actual configuration 901
21-31 End of configuration 901
21-32 The Havol utility - Disk signatures 902
22-1 Tivoli Storage Manager clustering server configuration 905
22-2 IBM 3582 and IBM 3580 device drivers on Windows Device Manager 908
22-3 Initial Configuration Task List 910
22-4 Welcome Configuration wizard 910
22-5 Initial configuration preferences 911
22-6 Site environment information 911
22-7 Initial configuration 912
22-8 Welcome Performance Environment wizard 912
22-9 Performance options 913

22-10 Drive analysis . . . 913
22-11 Performance wizard . . . 914
22-12 Server instance initialization wizard . . . 914
22-13 Server initialization wizard . . . 915
22-14 Server volume location . . . 916
22-15 Server service logon parameters . . . 916
22-16 Server name and password . . . 917
22-17 Completing the Server Initialization Wizard . . . 917
22-18 Completing the server installation wizard . . . 918
22-19 TSM server has been initialized . . . 918
22-20 Tivoli Storage Manager console . . . 919
22-21 Starting the Application Configuration Wizard . . . 921
22-22 Create service group option . . . 921
22-23 Service group configuration . . . 922
22-24 Change configuration to read-write . . . 922
22-25 Discovering process . . . 923
22-26 Choosing the kind of application . . . 923
22-27 Choosing TSM Server1 service . . . 924
22-28 Confirming the service . . . 924
22-29 Choosing the service account . . . 925
22-30 Selecting the drives to be used . . . 925
22-31 Summary with name and account for the service . . . 926
22-32 Choosing additional components . . . 926
22-33 Choosing other components for IP address and Name . . . 927
22-34 Specifying name and IP address . . . 927
22-35 Completing the application options . . . 928
22-36 Service Group Summary . . . 928
22-37 Changing resource names . . . 929
22-38 Confirming the creation of the service group . . . 929
22-39 Creating the service group . . . 930
22-40 Completing the wizard . . . 930
22-41 Cluster Monitor . . . 931
22-42 Resources online . . . 931
22-43 Link dependencies . . . 932
22-44 Starting the Application Configuration Wizard . . . 934
22-45 Create service group option . . . 934
22-46 Service group configuration . . . 935
22-47 Discovering process . . . 935
22-48 Choosing the kind of application . . . 936
22-49 Choosing TSM Server1 service . . . 936
22-50 Confirming the service . . . 937
22-51 Choosing the service account . . . 937
22-52 Selecting the drives to be used . . . 938

Archived

22-53 Summary with name and account for the service 938
22-54 Choosing additional components 939
22-55 Choosing other components for IP address and Name 940
22-56 Informing name and ip address 940
22-57 Completing the application options 941
22-58 Service Group Summary 941
22-59 Changing the names of the resources 942
22-60 Confirming the creation of the service group 942
22-61 Creating the service group 943
22-62 Completing the wizard 943
22-63 Correct link for the ISC Service Group 944
22-64 Accessing the administration center 944
22-65 Veritas Cluster Manager console shows TSM resource in SALVADOR946
22-66 Starting a manual backup using the GUI from RADON 946
22-67 RADON starts transferring files to the TSMSRV06 server 947
22-68 RADON loses its session, tries to reopen new connection to server . . 947
22-69 RADON continues transferring the files again to the server 948
22-70 Scheduled backup started for RADON in the TSMSRV06 server 949
22-71 Schedule log file in RADON shows the start of the scheduled backup 950
22-72 RADON loses its connection with the TSMSRV06 server 950
22-73 In the event log the scheduled backup is restarted 951
22-74 Schedule log file in RADON shows the end of the scheduled backup. 951
22-75 Every volume was successfully backed up by RADON 952
22-76 Migration task started as process 2 in the TSMSRV06 server 953
22-77 Migration has already transferred 4124 files to the tape storage pool . 953
22-78 Migration starts again in OTTAWA 954
22-79 Migration process ends successfully 954
22-80 Process 1 is started for the backup storage pool task 956
22-81 Process 1 has copied 6990 files in copy storage pool tape volume . . 956
22-82 Backup storage pool task is not restarted when TSMSRV06 is online 957
22-83 Volume 023AKKL2 defined as valid volume in the copy storage pool . 958
22-84 Occupancy for the copy storage pool after the failover 958
22-85 Occupancy is the same for primary and copy storage pools 959
22-86 Process 1 started for a database backup task 961
22-87 While the database backup process is started OTTAWA fails 961
22-88 Volume history does not report any information about 027AKKL2 . . . 962
22-89 The library volume inventory displays the tape volume as private 962
23-1 Tivoli Storage Manager backup/archive clustering client configuration 967
23-2 Starting the Application Configuration Wizard 975
23-3 Modifying service group option 976
23-4 No existing resource can be changed, but new ones can be added . . 976
23-5 Service group configuration 977
23-6 Discovering process 977

23-7 Choosing the kind of application 978
23-8 Choosing TSM Scheduler CL_VCS02_ISC service. 978
23-9 Confirming the service 979
23-10 Choosing the service account 979
23-11 Selecting the drives to be used 980
23-12 Summary with name and account for the service 980
23-13 Choosing additional components. 981
23-14 Choosing other components for Registry Replication 981
23-15 Specifying the registry key 982
23-16 Name and IP addresses 982
23-17 Completing the application options 983
23-18 Service Group Summary 983
23-19 Confirming the creation of the service group 984
23-20 Completing the wizard 984
23-21 Link after creating the new resource 985
23-22 Client Acceptor Generic service parameters 987
23-23 Final link with dependencies 988
23-24 A session starts for CL_VCS02_ISC in the activity log 989
23-25 CL_VCS02_ISC starts sending files to Tivoli Storage Manager server 990
23-26 Session lost for client and the tape volume is dismounted by server . 990
23-27 The event log shows the schedule as restarted. 991
23-28 The tape volume is mounted again for schedule to restart backup . . . 991
23-29 Schedule log shows the backup as completed 992
23-30 Schedule completed on the event log 992
23-31 Scheduled restore started for CL_MSCS01_SA 993
23-32 A session is started for restore and the tape volume is mounted 994
23-33 Restore starts in the schedule log file 994
23-34 Session is lost and the tape volume is dismounted 995
23-35 The restore process is interrupted in the client 995
23-36 Restore schedule restarts in client restoring files from the beginning . 996
23-37 Schedule restarted on the event log for CL_MSCS01_ISC 996
23-38 Restore completes successfully in the schedule log file 997
24-1 Clustered Windows 2003 configuration with Storage Agent 1002
24-2 Modifying devconfig option to point to devconfig file in dsmsta.opt . . 1006
24-3 Specifying parameters for the Storage Agent 1007
24-4 Specifying parameters for the Tivoli Storage Manager server. 1007
24-5 Specifying the account information 1008
24-6 Storage agent initialized. 1008
24-7 StorageAgent1 is started 1009
24-8 Installing Storage Agent for LAN-free backup of shared disk drives . 1011
24-9 Installing the service attached to StorageAgent2. 1011
24-10 Management console displays two Storage Agents 1012
24-11 Starting the TSM StorageAgent2 service in SALVADOR 1012

24-12 Creating StorageAgent2 resource 1013
24-13 StorageAgent2 must come online before the Scheduler 1014
24-14 Storage Agent CL_VCS02_STA session for Tape Library Sharing . . 1016
24-15 A tape volume is mounted and Storage Agent starts sending data . . 1016
24-16 Client starts sending files to the server in the schedule log file 1017
24-17 Sessions for Client and Storage Agent are lost in the activity log . . . 1017
24-18 Backup is interrupted in the client 1018
24-19 Tivoli Storage Manager server mounts tape volume in second drive 1018
24-20 The scheduled is restarted and the tape volume mounted again . . . 1019
24-21 Backup ends successfully 1019
24-22 Starting restore session for LAN-free. 1021
24-23 Restore starts on the schedule log file. 1022
24-24 Both sessions for Storage Agent and client are lost in the server . . . 1022
24-25 The tape volume is dismounted by the server 1023
24-26 The Storage Agent waiting for tape volume to be mounted by server 1023
24-27 Event log shows the restore as restarted. 1024
24-28 The client restores the files from the beginning 1024
24-29 Final statistics for the restore on the schedule log file 1025

Archived

Tables

1-1 Single points of failure 5
1-2 Types of HA solutions 7
2-1 Cluster matrix. 19
2-2 Tivoli Storage Manager configuration matrix 20
4-1 Windows 2000 cluster server configuration 30
4-2 Cluster groups for our Windows 2000 MSCS 31
4-3 Windows 2000 DNS configuration 31
4-4 Windows 2003 cluster server configuration 46
4-5 Cluster groups for our Windows 2003 MSCS 47
4-6 Windows 2003 DNS configuration 47
5-1 Windows 2000 lab ISC cluster resources 120
5-2 Windows 2000 lab Tivoli Storage Manager server cluster resources . 120
5-3 Windows 2000 Tivoli Storage Manager virtual server in our lab 121
5-4 Lab Windows 2003 ISC cluster resources 181
5-5 Lab Windows 2003 Tivoli Storage Manager cluster resources 181
5-6 Tivoli Storage Manager virtual server for our Windows 2003 lab 182
6-1 Tivoli Storage Manager backup/archive client for local nodes. 250
6-2 Tivoli Storage Manager backup/archive client for virtual nodes. 251
6-3 Windows 2003 TSM backup/archive configuration for local nodes . . . 290
6-4 Windows 2003 TSM backup/archive client for virtual nodes 291
7-1 LAN-free configuration details 335
7-2 TSM server details 337
7-3 SAN devices details 337
7-4 Windows 2003 LAN-free configuration of our lab 379
7-5 Server information 381
7-6 Storage devices used in the SAN 381
8-1 HACMP cluster topology 429
8-2 HACMP resources groups 430
10-1 Tivoli Storage Manager client distinguished configuration. 529
10-2 .Client nodes configuration of our lab 530
11-1 Storage Agents distinguished configuration. 558
11-2 .LAN-free configuration of our lab 559
11-3 Server information 560
11-4 Storage Area Network devices 560
13-1 Lab Tivoli Storage Manager server cluster resources 619
14-1 Tivoli Storage Manager client distinguished configuration. 655
14-2 Client nodes configuration of our lab 656
15-1 Storage Agents configuration. 674

© Copyright IBM Corp. 2005. All rights reserved.

16-1 HACMP/VERITAS Cluster Server feature comparison 716
16-2 HACMP/VERITAS Cluster Server environment support 718
19-1 Storage Agent configuration for our design 795
19-2 .LAN-free configuration of our lab 796
19-3 Server information 797
19-4 Storage Area Network devices 797
20-1 Tivoli Storage Manager client configuration 840
21-1 Cluster server configuration 881
21-2 Service Groups in VSFW 882
21-3 DNS configuration 882
22-1 Lab Tivoli Storage Manager server service group 906
22-2 ISC service group 906
22-3 Tivoli Storage Manager virtual server configuration in our lab 907
23-1 Tivoli Storage Manager backup/archive client for local nodes 968
23-2 Tivoli Storage Manager backup/archive client for virtual node 968
24-1 LAN-free configuration details 1003
24-2 TSM server details 1004
24-3 SAN devices details 1004
A-1 Additional material 1030

Archived

Examples

5-1 Activity log when the client starts a scheduled backup 150
5-2 Schedule log file shows the start of the backup on the client 150
5-3 Error log when the client lost the session 151
5-4 Schedule log file when backup is restarted on the client 151
5-5 Activity log after the server is restarted 152
5-6 Schedule log file shows backup statistics on the client 153
5-7 Disk storage pool migration started on server 155
5-8 Disk storage pool migration started again on the server 155
5-9 Disk storage pool migration ends successfully 156
5-10 Starting a backup storage pool process. 157
5-11 After restarting the server the storage pool backup does not restart . . 158
5-12 Starting a database backup on the server 161
5-13 After the server is restarted database backup does not restart 162
5-14 Volume history for database backup volumes 163
5-15 Library volumes 163
5-16 Starting inventory expiration 165
5-17 No inventory expiration process after the failover 165
5-18 Starting inventory expiration again. 166
5-19 Activity log when the client starts a scheduled backup 211
5-20 Schedule log file shows the start of the backup on the client 211
5-21 Error log when the client lost the session 213
5-22 Schedule log file when backup is restarted on the client 213
5-23 Activity log after the server is restarted 213
5-24 Schedule log file shows backup statistics on the client 214
5-25 Restore starts in the event log 216
5-26 Restore starts in the schedule log file of the client. 216
5-27 The session is lost in the client 217
5-28 The client reopens a session with the server. 217
5-29 The schedule is restarted in the activity log 218
5-30 Restore final statistics 218
5-31 The activity log shows the event failed 218
5-32 Disk storage pool migration started on server 220
5-33 Disk storage pool migration started again on the server 220
5-34 Disk storage pool migration ends successfully 221
5-35 Starting a backup storage pool process. 222
5-36 Starting a database backup on the server 225
5-37 After the server is restarted database backup does not restart 226
5-38 Starting inventory expiration 227

© Copyright IBM Corp. 2005. All rights reserved.

5-39 No inventory expiration process after the failover 229
5-40 Starting inventory expiration again. 230
6-1 Session started for CL_MSCS01_SA 277
6-2 Schedule log file shows the client sending files to the server 277
6-3 The client loses its connection with the server. 278
6-4 Schedule log file shows backup is restarted on the client 278
6-5 A new session is started for the client on the activity log. 280
6-6 Schedule log file shows the backup as completed 281
6-7 Schedule log file shows the client restoring files 284
6-8 Connection is lost on the server. 285
6-9 Schedule log for the client starting the restore again. 286
6-10 New session started on the activity log for CL_MSCS01_SA 287
6-11 Schedule log file on client shows statistics for the restore operation. . 288
8-1 /etc/hosts file after the changes 431
8-2 The edited /usr/es/sbin/etc/cluster/rhosts file. 431
8-3 The AIX bos filesets that must be installed prior to installing HACMP. 431
8-4 The lslpp -L command 432
8-5 The RSCT filesets required prior to HACMP installation 432
8-6 The AIX fileset that must be installed for the SAN discovery function . 432
8-7 SNMPD script to switch from v3 to v2 support. 433
8-8 HACMP serial cable features. 433
8-9 lsdev command for tape subsystems. 437
8-10 The lspv command output 438
8-11 The lscfg command 438
8-12 mkvg command to create the volume group 438
8-13 mklv commands to create logical volumes 439
8-14 mklv commands used to create the logical volumes 439
8-15 The logform command. 439
8-16 The crfs commands used to create the filesystems. 439
8-17 The varyoffvg command 439
8-18 The importvg command. 440
8-19 The chvg command 440
8-20 The varyoffvg command 440
8-21 The mkvg command 440
8-22 The chvg command 440
8-23 The varyoffvg command 441
8-24 The importvg command. 441
8-25 APAR installation check with instfix command. 442
9-1 The tar command extraction 465
9-2 setupISC usage 465
9-3 The tar command extraction 472
9-4 startInstall.sh usage. 472
9-5 Command line installation for the Administration Center. 473

9-6 lssrc -g cluster 483
9-7 Stop the initial server installation instance 486
9-8 Files to remove after the initial server installation 486
9-9 The server stanza for the client dsm.sys file 487
9-10 The variables which must be exported in our environment 487
9-11 dsmfmt command to create database, recovery log, storage pool files 488
9-12 The dsmserv format prepares db & log files and the dsmserv.dsk file 488
9-13 Starting the server in the foreground 488
9-14 Our server naming and mirroring. 488
9-15 The define commands for the diskpool 489
9-16 An example of define library, define drive and define path commands 489
9-17 Library parameter RESETDRIVES set to YES 489
9-18 The register admin and grant authority commands 489
9-19 The register admin and grant authority commands 490
9-20 Copy the example scripts on the first node 490
9-21 Setting running environment in the start script. 490
9-22 Stop script setup instructions 491
9-23 Modifying the lock file path. 492
9-24 dsmadmc command setup 492
9-25 ISC startup command 492
9-26 ISC stop sample script 492
9-27 Monitor script example. 494
9-28 Verify available cluster resources 496
9-29 Takeover progress monitor 499
9-30 Post takeover resource checking. 500
9-31 Monitor resource group moving 501
9-32 Resource group state check 502
9-33 Monitor resource group moving 502
9-34 Resource group state check 503
9-35 Monitor resource group moving 504
9-36 Resource group state check 505
9-37 Client sessions starting 506
9-38 Query sessions for data transfer 506
9-39 client stops sending data 507
9-40 The restarted Tivoli Storage Manager accept client rejoin. 507
9-41 The client reconnect and continue operations 508
9-42 Scheduled backup case. 509
9-43 Query event result 510
9-44 Register node command 511
9-45 Define server using the command line. 511
9-46 Define path commands 511
9-47 Client sessions starting 511
9-48 Tape mount for LAN-free messages 512

9-49 Query session for data transfer 512
9-50 Storage unmount the tapes for the dropped server connection. 512
9-51 client stops receiving data 513
9-52 The restarted Tivoli Storage Manager rejoin the Storage Agent. 514
9-53 Library recovery for Storage Agent 514
9-54 New restore operation 514
9-55 Volume mounted for restore after the recovery 515
9-56 Migration restarts after a takeover 516
9-57 Migration process ending. 517
9-58 Tivoli Storage Manager restarts after a takeover. 518
9-59 Tivoli Storage Manager restarts after a takeover. 520
9-60 Search for database backup volumes 522
9-61 Expire inventory process starting. 524
9-62 Tivoli Storage Manager restarts. 524
9-63 Database and log volumes state 525
9-64 New expire inventory execution. 525
10-1 dsm.opt file contents located in the application shared disk 532
10-2 dsm.sys file contents located in the default directory. 533
10-3 Current contents of the shared disk directory for the client 534
10-4 The HACMP directory which holds the client start and stop scripts. . . 534
10-5 Selective backup schedule 536
10-6 Client sessions starting 537
10-7 Client session cancelled due to the communication timeout. 537
10-8 The restarted client scheduler queries for schedules (client log). 537
10-9 The restarted client scheduler queries for schedules (server log) 538
10-10 The restarted backup operation. 538
10-11 Client sessions starting 540
10-12 Monitoring data transfer through query session command 540
10-13 Query sessions showing hanged client sessions. 541
10-14 The client reconnect and restarts incremental backup operations. . . . 541
10-15 The Tivoli Storage Manager accept the client new sessions 542
10-16 Query event showing successful result.. 543
10-17 Client sessions starting 544
10-18 The client and restarts and hits MAXNUMMP 545
10-19 Hanged client session with an output volume 546
10-20 Old sessions cancelling work in startup script 546
10-21 Hanged tape holding sessions cancelling job 548
10-22 Event result 549
10-23 Restore schedule. 550
10-24 Client sessions starting 551
10-25 The server log during restore restart 552
10-26 The Tivoli Storage Manager client log 553
10-27 Query server for restartable restores. 554

11-1 lsdev command for tape subsystems 561
11-2 Set server settings from command line 563
11-3 Define server using the command line 567
11-4 Define paths using the command line 569
11-5 Local instance dsmsta.opt 569
11-6 The dsmsta setstorageserver command 569
11-7 The dsmsta setstorageserver command for clustered Storage Agent . 569
11-8 The devconfig.txt file 570
11-9 Clustered Storage Agent devconfig.txt 570
11-10 The /usr/tivoli/tsm/client/ba/bin/dsm.sys file 570
11-11 Example scripts copied to /usr/es/sbin/cluster/local/tsmsrv, first node 571
11-12 Our Storage Agent with AIX server startup script 572
11-13 Application server start script 572
11-14 Copy from /usr/tivoli/tsm/server/bin to /usr/es/sbin/cluster/local/tsmsrv 573
11-15 Our Storage Agent with non-AIX server startup script 574
11-16 Application server start script 577
11-17 Storage agent stanza in dsm.sys 577
11-18 Application server stop script 578
11-19 Client sessions starting 579
11-20 Output volumes open messages 579
11-21 Client sessions transferring data to Storage Agent 579
11-22 The ISC being restarted 580
11-23 The Tivoli Storage Manager Storage Agent is restarted 580
11-24 CL_HACMP03_STA reconnecting 580
11-25 Trace showing pvr at work with reset 581
11-26 Tape dismounted after SCSI reset 582
11-27 Extract of console log showing session cancelling work 582
11-28 The client schedule restarts 583
11-29 Server log view of restarted restore operation 583
11-30 Client sessions starting 584
11-31 Tape mount and open messages 585
11-32 Checking for data being received by the Storage Agent 585
11-33 ISC restarting 586
11-34 Storage agent restarting 586
11-35 Tivoli Storage Manager server accepts new sessions, unloads tapes 586
11-36 Extract of console log showing session cancelling work 587
11-37 The client restore re issued 588
11-38 Server log of new restore operation 588
11-39 Client restore terminating successfully 589
12-1 Verifying the kernel version information in the Makefile 601
12-2 Copying kernel config file 601
12-3 The grub configuration file /boot/grub/menu.lst 603
12-4 Verification of RDAC setup 604

12-5 Installation of the IBMtape driver 604
12-6 Device information in /proc/scsi/IBMtape and /proc/scsi/IBMchanger . 605
12-7 Contents of /proc/scsi/scsi 608
12-8 SCSI devices created by scsidev. 608
12-9 UUID changes after file system is created. 609
12-10 Devlabel configuration file /etc/sysconfig/devlabel. 610
12-11 Installation of Tivoli System Automation for Multiplatforms 611
12-12 Configuration of the disk tie breaker 613
12-13 Displaying the status of the RecoveryRM with the lssrc command . . . 615
13-1 Installation of Tivoli Storage Manager Server 620
13-2 Stop Integrated Solutions Console and Administration Center 624
13-3 Necessary entries in /etc/fstab for the Tivoli Storage Manager server 625
13-4 Cleaning up the default server installation. 626
13-5 Contents of /tsm/files/dsmserv.opt. 626
13-6 Server stanza in dsm.sys to enable the use of dsmadmc 626
13-7 Setting up necessary environment variables 627
13-8 Formatting database, log, and disk storage pools with dsmfmt 627
13-9 Starting the server in the foreground 627
13-10 Set up servername, mirror db and log, and set logmode to rollforward 628
13-11 Definition of the disk storage pool 628
13-12 Definition of library devices 628
13-13 Registration of TSM administrator 629
13-14 Extract of the configuration file sa-tsmserver.conf 630
13-15 Verification of tape and medium changer serial numbers with sginfo . 631
13-16 Execution of cfgtsmserver to create definition files 632
13-17 Executing the SA-tsmserver-make script. 632
13-18 Extract of the configuration file sa-tsmadmin.conf 633
13-19 Execution of cfgtsmadminc to create definition files 634
13-20 Configuration of AntiAffinity relationship 635
13-21 Validation of resource group members 635
13-22 Persistent and dynamic attributes of all resource groups 636
13-23 Output of the lsrel command 637
13-24 Changing the nominal state of the SA-tsmserver-rg to online 638
13-25 Output of the getstatus script. 638
13-26 Changing the nominal state of the SA-tsmadminc-rg to online 639
13-27 Log file /var/log/messages after a failover 641
13-28 Activity log when the client starts a scheduled backup 643
13-29 Schedule log file showing the start of the backup on the client 643
13-30 Error log file when the client looses the session 643
13-31 Schedule log file when backup restarts on the client. 644
13-32 Activity log after the server is restarted 644
13-33 Schedule log file showing backup statistics on the client. 644
13-34 Disk storage pool migration starting on the first node 646

13-35 Disk storage pool migration starting on the second node 646
13-36 Disk storage pool migration ends successfully 647
13-37 Starting a backup storage pool process. 647
13-38 After restarting the server the storage pool backup doesn't restart . . . 648
13-39 Starting a database backup on the server 650
13-40 After the server is restarted database backup does not restart 650
13-41 Starting inventory expiration 651
14-1 dsm.opt file contents located in the application shared disk 658
14-2 Stanza for the clustered client in dsm.sys 659
14-3 Creation of the password file TSM.PWD 659
14-4 Creation of the symbolic link that point to the Client CAD script 661
14-5 Output of the lsrg -m command before configuring the client 661
14-6 Definition file SA-nfsserver-tsmclient.def 662
14-7 Output of the lsrel command 663
14-8 Output of the lsrg -m command while resource group is online. 663
14-9 Session for CL_ITSAMP02_CLIENT starts 664
14-10 Schedule log file during starting of the scheduled backup. 664
14-11 Activity log entries while diomede fails. 665
14-12 Schedule log file dsmsched.log after restarting the backup. 665
14-13 Activity log entries while the new session for the backup starts 667
14-14 Schedule log file reports the successfully completed event. 667
14-15 Activity log entries during start of the client restore 668
14-16 Schedule log entries during start of the client restore 668
14-17 Activity log entries during the failover 669
14-18 Schedule log entries during restart of the client restore. 669
14-19 Activity log entries during restart of the client restore 671
14-20 Schedule log entries after client restore finished 671
15-1 Installation of the TIVsm-stagent rpm on both nodes 675
15-2 Clustered instance /mnt/nfsfiles/tsm/StorageAgent/bin/dsmsta.opt. . . 679
15-3 The dsmsta setstorageserver command 680
15-4 The dsmsta setstorageserver command for clustered STA. 680
15-5 The devconfig.txt file 681
15-6 Clustered Storage Agent dsmsta.opt 681
15-7 dsm.opt file contents located in the application shared disk 681
15-8 Server stanza in dsm.sys for the clustered client. 682
15-9 Creation of the password file TSM.PWD 683
15-10 Creation of the symbolic link that points to the Storage Agent script. . 684
15-11 Output of the lsrg -m command before configuring the Storage Agent 684
15-12 Definition file SA-nfsserver-tsmsta.def. 684
15-13 Definition file SA-nfsserver-tsmclient.def 685
15-14 Output of the lsrel command 686
15-15 Output of the lsrg -m command while resource group is online. 687
15-16 Scheduled backup starts 687

15-17 Activity log when scheduled backup starts 689
15-18 Activity log when tape is mounted 690
15-19 Activity log when failover takes place 690
15-20 Activity log when tsmclientctrl-cad script searches for old sessions . . 691
15-21 dsmwebcl.log when the CAD starts 691
15-22 Actlog when CAD connects to the server 691
15-23 Actlog when Storage Agent connects to the server 692
15-24 Schedule log when schedule is restarted 692
15-25 Activity log when the tape volume is mounted again 693
15-26 Schedule log shows that the schedule completed successfully. 694
15-27 Scheduled restore starts 695
15-28 Actlog when the schedule restore starts 696
15-29 Actlog when resources are stopped at diomede 697
15-30 Schedule restarts at lochness 698
15-31 Restore finishes successfully 699
17-1 Atlantic .rhosts 724
17-2 Banda .rhosts 724
17-3 atlantic /etc/hosts file 724
17-4 banda /etc/hosts file 724
17-5 The AIX command lscfg to view FC disk details 725
17-6 The lspv command output 726
17-7 The lscfg command 727
17-8 The mkvg command to create the volume group. 727
17-9 The mklv commands to create the logical volumes 728
17-10 The mklv commands used to create the logical volumes 728
17-11 The logform command 728
17-12 The crfs commands used to create the file systems 728
17-13 The varyoffvg command 728
17-14 The importvg command 729
17-15 The chvg command 729
17-16 The varyoffvg command 729
17-17 The mkvg command to create the volume group. 729
17-18 The mklv commands to create the logical volumes 730
17-19 The logform command 730
17-20 The crfs commands used to create the file systems 730
17-21 The chvg command 730
17-22 The varyoffvg command 730
17-23 .rhosts file. 731
17-24 VCS installation script 731
17-25 The VCS checking of installation requirements 734
17-26 The VCS install method prompt and install summary 740
18-1 The AIX rmitab command 754
18-2 Stop the initial server installation instance 754

18-3 The variables which must be exported in our environment 754
18-4 Files to remove after the initial server installation 755
18-5 The server stanza for the client dsm.sys file 755
18-6 dsmfmt command to create database, recovery log, storage pool files 756
18-7 The dsmserv format command to prepare the recovery log 756
18-8 An example of starting the server in the foreground 756
18-9 The server setup for use with our shared disk files 756
18-10 The define commands for the diskpool 756
18-11 An example of define library, define drive and define path commands 757
18-12 The register admin and grant authority commands 757
18-13 /opt/local/tsmsrv/startTSMsrv.sh 758
18-14 /opt/local/tsmsrv/stopTSMsrv.sh 759
18-15 /opt/local/tsmsrv/cleanTSMsrv.sh 760
18-16 /opt/local/tsmsrv/monTSMsrv.sh 762
18-17 Adding a Service Group sg_tsmsrv 763
18-18 Adding a NIC Resource 763
18-19 Configuring an IP Resource in the sg_tsmsrv Service Group 763
18-20 Adding the LVMVG Resource to the sg_tsmsrv Service Group 764
18-21 Configuring the Mount Resource in the sg_tsmsrv Service Group . . . 764
18-22 Adding and configuring the app_tsmsrv Application 766
18-23 The sg_tsmsrv Service Group: /etc/VRTSvcs/conf/config/main.cf file . 767
18-24 The results return from hastatus 770
18-25 hastatus log from the surviving node, Atlantic 771
18-26 tail -f /var/VRTSvcs/log/engine_A.log from surviving node, Atlantic . . 771
18-27 The recovered cluster using hastatus 771
18-28 Current cluster status from the hastatus output 772
18-29 hagrp -online command 772
18-30 hastatus of the online transition for the sg_tsmsrv. 772
18-31 tail -f /var/VRTSvcs/log/engine_A.log 773
18-32 Verify available cluster resources using the hastatus command 773
18-33 hagrp -offline command 775
18-34 hastatus output for the Service Group OFFLINE 775
18-35 tail -f /var/VRTSvcs/log/engine_A.log 775
18-36 hastatus output prior to the Service Groups switching nodes 775
18-37 hastatus output of the Service Group switch 777
18-38 tail -f /var/VRTSvcs/log/engine_A.log from surviving node, Atlantic . . 777
18-39 hastatus output of the current cluster state 778
18-40 hargrp -switch command to switch the Service Group back to Banda. 778
18-41 /var/VRTSvcs/log/engine_A.log segment for the switch back to Banda778
18-42 /var/VRTSvcs/log/engine_A.log output for the failure activity 779
18-43 hastatus of the ONLINE resources 780
18-44 /var/VRTSvcs/log/engine_A.log output for the recovery activity 780
18-45 hastatus of the online resources fully recovered from the failure test . 781

18-46 hastatus | grep ONLINE output 781
18-47 Client sessions starting 782
18-48 client stops sending data 782
18-49 Cluster log demonstrating the change of cluster membership status . 783
18-50 engine_A.log online process and completion summary. 783
18-51 The restarted Tivoli Storage Manager accept client rejoin. 784
18-52 The client reconnect and continue operations 784
18-53 Command query mount and process. 786
18-54 Actlog output showing the mount of volume ABA990 786
18-55 Actlog output demonstrating the completion of the migration 787
18-56 q mount output. 788
18-57 q process output 788
18-58 VCS hastatus command output after the failover 789
18-59 q process after the backup storage pool command has restarted. . . . 790
18-60 q mount after the takeover and restart of Tivoli Storage Manager. . . . 790
19-1 The dsmsta setstorageserver command 798
19-2 The devconfig.txt file 798
19-3 dsmsta.opt file change results 799
19-4 dsm.sys stanzas for Storage Agent configured as highly available. . . 799
19-5 /opt/local/tsmsta/startSTA.sh 804
19-6 /opt/local/tsmsta/stopSTA.sh 805
19-7 /opt/local/tsmsta/cleanSTA.sh 806
19-8 monSTA.sh script. 806
19-9 VCS commands to add app_sta application into sg_isc_sta_tsmcli . . 807
19-10 The completed /etc/VRTSvcs/conf/config/main.cf file 808
19-11 The results return from hastatus 811
19-12 hastatus log from the surviving node, Atlantic 811
19-13 tail -f /var/VRTSvcs/log/engine_A.log from surviving node, Atlantic . . 812
19-14 The recovered cluster using hastatus 812
19-15 Current cluster status from the hastatus output. 813
19-16 hagrp -online command 813
19-17 hastatus of online transition for sg_isc_sta_tsmcli Service Group. . . . 813
19-18 tail -f /var/VRTSvcs/log/engine_A.log 814
19-19 Verify available cluster resources using the hastatus command 814
19-20 hagrp -offline command 817
19-21 hastatus output for the Service Group OFFLINE. 817
19-22 tail -f /var/VRTSvcs/log/engine_A.log 817
19-23 hastatus output prior to the Service Groups switching nodes 817
19-24 hastatus output of the Service Group switch 819
19-25 tail -f /var/VRTSvcs/log/engine_A.log from surviving node, Atlantic . . 820
19-26 hastatus output of the current cluster state 820
19-27 hargrp -switch command to switch the Service Group back to Banda. 821
19-28 /var/VRTSvcs/log/engine_A.log segment for the switch back to Banda821

19-29 /var/VRTSvcs/log/engine_A.log output for the failure activity 822
19-30 hastatus of the ONLINE resources 823
19-31 /var/VRTSvcs/log/engine_A.log output for the recovery activity 824
19-32 hastatus of the online resources fully recovered from the failure test . 824
19-33 Client selective backup schedule configured on TSMSRV03 825
19-34 Client sessions starting 825
19-35 Tivoli Storage Manager server volume mounts 825
19-36 The sessions being cancelled at the time of failure 826
19-37 TSMSRV03 actlog of the cl_veritas01_sta recovery process 826
19-38 Server process view during LAN-free backup recovery 828
19-39 Extract of console log showing session cancelling work 829
19-40 dsmsched.log output showing failover transition, schedule restarting . 829
19-41 Backup during a failover shows a completed successful summary . . . 830
19-42 Restore schedule 831
19-43 Client restore sessions starting 832
19-44 Query the mounts looking for the restore data flow starting 832
19-45 Query session command during the transition after failover of banda . 833
19-46 The server log during restore restart 833
19-47 Addition restore session begins, completes restore after the failover . 835
19-48 dsmsched.log output demonstrating the failure and restart transition . 836
19-49 Server sessions after the restart of the restore operation. 836
19-50 dsmsched.log output of completed summary of failover restore test . . 837
20-1 /opt/IBM/ISC/tsm/client/ba/bin/dsm.opt file content 841
20-2 /usr/tivoli/tsm/client/ba/bin/dsm.sys stanza, links clustered dsm.opt file841
20-3 The path and file difference for the passworddir option 842
20-4 The tar command extraction 843
20-5 Integrated Solutions Console installation script 843
20-6 Administration Center install directory 850
20-7 /opt/local/tsmcli/startTSMcli.sh. 857
20-8 /opt/local/tsmcli/stopTSMcli.sh 859
20-9 /opt/local/tsmcli/cleanTSMcli.sh 863
20-10 /opt/local/isc/startISC.sh 863
20-11 /opt/local/isc/stopISC.sh. 864
20-12 /opt/local/isc/cleanISC.sh 864
20-13 /opt/local/isc/monISC.sh 864
20-14 Changing the OnlineTimeout for the ISC 865
20-15 Adding a Service Group 865
20-16 Adding an LVMVG Resource 865
20-17 Adding the Mount Resource to the Service Group sg_isc_sta_tsmcli . 866
20-18 Adding a NIC Resource 866
20-19 Adding an IP Resource 866
20-20 VCS commands to add tsmcad application to the sg_isc_sta_tsmcli . 867
20-21 Adding app_isc Application to the sg_isc_sta_tsmcli Service Group. . 867

20-22 Example of the main.cf entries for the sg_isc_sta_tsmcli 867
20-23 Client sessions starting 870
20-24 Volume opened messages on server console 870
20-25 Server console log output for the failover reconnection 871
20-26 The client schedule restarts. 871
20-27 q session shows the backup and dataflow continuing 872
20-28 Unmounting the tape once the session is complete 872
20-29 Server actlog output of the session completing successfully 872
20-30 Schedule a restore with client node CL_VERITAS01_CLIENT 873
20-31 Client sessions starting 874
20-32 Mount of the restore tape as seen from the server actlog 874
20-33 The server log during restore restart 875
20-34 The Tivoli Storage Manager client log 875
23-1 Registering the node password 971
23-2 Creating the schedule on each node 973

Archived

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.

© Copyright IBM Corp. 2005. All rights reserved.

Trademarks

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

AFS®
AIX®
AIX 5L™
DB2®
DFS™
Enterprise Storage Server®
ESCON®
@server®
@server®
HACMP™
IBM®
ibm.com®
iSeries™
PAL®
PowerPC®
pSeries®
RACF®
Redbooks™
Redbooks (logo) ™
SANergy®
ServeRAID™
Tivoli®
TotalStorage®
WebSphere®
xSeries®
z/OS®
zSeries®

The following terms are trademarks of other companies:

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.

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

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

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

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

Archived

Preface

This IBM® Redbook is an easy-to-follow guide which describes how to implement IBM Tivoli® Storage Manager Version 5.3 products in highly available clustered environments.

The book is intended for those who want to plan, install, test, and manage the IBM Tivoli Storage Manager Version 5.3 in various environments by providing best practises and showing how to develop scripts for clustered environments.

The book covers the following environments: IBM AIX® HACMP™, IBM Tivoli System Automation for Multiplatforms on Linux® and AIX, Makeshift Cluster Server on Windows® 2000 and Windows 2003, and VERITAS Storage Foundation HA on AIX, and Windows Server 2003 Enterprise Edition.

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.

Archived

© Copyright IBM Corp. 2005. All rights reserved.

The team, from left to right: Werner, Marco, Roland, Dan, Rosane, and Maria.

Roland Tretau is a Project Leader with the IBM International Technical Support Organization, San Jose Center. Before joining the ITSO in April 2001, Roland worked in Germany as an IT Architect with a major focus on open systems solutions and Microsoft® technologies. He holds a Master's degree in Electrical Engineering with an emphasis in telecommunications. He is a Red Hat Certified Engineer (RHCE) and a Microsoft Certified Systems Engineer (MCSE), and he holds a Masters Certificate in Project Management from The George Washington University School of Business and Public Management.

Dan Edwards is a Consulting I/T Specialist with IBM Global Services, Integrated Technology Services, and is based in Ottawa, Canada. He has over 27 years experience in the computing industry, with the last 15 years spent working on Storage and UNIX solutions. He holds multiple product certifications, including Tivoli, AIX, and Oracle. He is also an IBM Certified Professional, and a member of the I/T Specialist Certification Board. Dan spends most of his client contracting time working with Tivoli Storage Manager, High Availability, and Disaster Recovery solutions.

Werner Fischer is an IT Specialist in IBM Global Services, Integrated Technology Services in Austria. He has 3 years of experience in the high availability field. He has worked at IBM for 2 years, including 1 year at the EMEA Storage ATS (Advanced Technical Support) in Mainz, Germany. His areas of expertise include planning and implementation of Linux high availability clusters, SAN disk and tape solutions, and hierarchical storage management environments. Werner holds a graduate degree in computer and media security from the University of Applied Sciences of Upper Austria in Hagenberg where he now also teaches as assistant lecturer.

Marco Mencarelli is an IT Specialist in IBM Global Services, Integrated Technology Services, Italy. He has 6 years of experience in planning and implementing Tivoli Storage Manager and HACMP. His areas of expertise include AIX, Disaster Recovery solutions, several Tivoli Data Protection products, and implementation of storage solutions.

Rosane Goldstein Golubcic Langnor is an IT Specialist in Brazil working for IBM Global Services. She has been working since 2000 with Tivoli Storage Manager, and her areas of expertise include planning and implementing Windows servers, backup solutions, and storage management. She is a Microsoft Certified System Engineer (MCSE).

Maria Jose Rodriguez Canales is an IT Specialist in IBM Global Services, Integrated Technology Services, Spain. She has 12 years of experience in IBM Storage Subsystem implementations for mainframe and open environments. Since 1997, she has specialized in Tivoli Storage Manager, working in areas as diverse as AIX, Linux, Windows, and z/OS®, participating in many projects to back up databases and mail or file servers over LAN and SAN networks. She holds a degree in Physical Science from the Complutense University, in Madrid.

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

Yvonne Lyon, Deanna Polm, Sangam Racherla, Leslie Parham, Emma Jacobs
International Technical Support Organization, San Jose Center

Tricia Jiang, Freddy Saldana, Kathy Mitton, Jo Lay, David Bohm, Jim Smith
IBM US

Thomas Lumpp, Enrico Jödecke, Wilhelm Blank
IBM Germany

Christoph Mitasch
IBM Austria

Michelle Corry, Nicole Zakhari, Victoria Krischke
VERITAS Software

Archived

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 e-mail 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

Archived

Part 1

Highly available clusters with IBM Tivoli Storage Manager

In this part of the book, we discuss our basic setup and explain how we approached the different high availability clusters solutions with IBM Tivoli Storage Manager.

© Copyright IBM Corp. 2005. All rights reserved.

Archived

1

What does high availability imply?

In this chapter, we discuss high availability concepts and terminology.

Archived

© Copyright IBM Corp. 2005. All rights reserved.

1.1 High availability

In today's complex environments, providing continuous service for applications is a key component of a successful IT implementation. High availability is one of the components that contributes to providing continuous service for the application clients, by masking or eliminating both planned and unplanned systems and application downtime. This is achieved through the elimination of hardware and software single points of failure (SPOFs).

A high availability solution will ensure that the failure of any component of the solution, either hardware, software, or system management, will not cause the application and its data to be unavailable to the user.

High availability solutions should eliminate single points of failure (SPOFs) through appropriate design, planning, selection of hardware, configuration of software, and carefully controlled change management discipline.

1.1.1 Downtime

The downtime is the time frame when an application is not available to serve its clients. We can classify the downtime as:

- Planned:
  - Hardware upgrades
  - Repairs
  - Software updates/upgrades
  - Backups (offline backups)
  - Testing (periodic testing is required for cluster validation)
  - Development
- Unplanned:
  - Administrator errors
  - Application failures
  - Hardware failures
  - Environmental disasters

A high availability solution is based on well-proven clustering technology, and consists of two components:

- **High availability:** The process of ensuring an application is available for use through the use of duplicated and/or shared resources.
- **Cluster multi-processing:** Multiple applications are running on the same nodes with shared or concurrent access to the data.

Archived

1.1.2 High availability concepts

What needs to be protected? Ultimately, the goal of any IT solution in a critical environment is to provide continuous service and data protection.

High availability is just one building block in achieving the continuous operation goal. The high availability is based on the availability of the hardware, software (operating system and its components), application, and network components.

For a high availability solution, you need:

- Redundant servers
- Redundant networks
- Redundant network adapters
- Monitoring
- Failure detection
- Failure diagnosis
- Automated failover
- Automated reintegration

The main objective of a highly available cluster is eliminate Single Points of Failure (SPOFs) (see Table 1-1).

Table 1-1 Single points of failure

Cluster object	Eliminated as a single point of failure by:
Node (servers)	Multiple nodes
Power supply	Multiple circuits and/or power supplies
Network adapter	Redundant network adapters
Network	Multiple networks to connect nodes
TCP/IP subsystem	A non- IP network to back up TCP/IP
Disk adapter	Redundant disk adapters
Disk	Redundant hardware and disk mirroring or RAID technology
Application	Configuring application monitoring and backup node(s) to acquire the application engine and data

Each of the items listed in Table 1-1 in the Cluster Object column is a physical or logical component that, if it fails, will result in the application being unavailable for serving clients.

Archived

1.1.3 High availability versus fault tolerance

The systems for the detection and handling of the hardware and software failures can be defined in two groups:

- Fault-tolerant systems
- High availability systems

Fault-tolerant systems

The systems provided with fault tolerance are designed to operate virtually without interruption, regardless of the failure that may occur (except perhaps for a complete site being down due to a natural disaster). In such systems, *all* components are at least duplicated for either software or hardware.

Thus, CPU, memory, and disks have a special design and provide continuous service, even if one sub-component fails.

Such systems are very expensive and extremely specialized. Implementing a fault tolerant solution requires a lot of effort and a high degree of customizing for all system components.

In places where *no* downtime is acceptable (life support and so on), fault-tolerant equipment and solutions are required.

High availability systems

The systems configured for high availability are a combination of hardware and software components configured in such a way as to ensure automated recovery in case of failure with a minimal acceptable downtime.

In such systems, the software involved detects problems in the environment, and then provides the transfer of the application on another machine, taking over the identity of the original machine (node).

Thus, it is very important to eliminate all single points of failure (SPOFs) in the environment. For example, if the machine has only one network connection, a second network interface should be provided in the same node to take over in case the primary adapter providing the service fails.

Another important issue is to protect the data by mirroring and placing it on shared disk areas accessible from any machine in the cluster.

The high availability cluster software provides the framework and a set of tools for integrating applications in a highly available system.

Archived

Applications to be integrated in a cluster will require customizing, not at the application level, but rather at the cluster software and operating system platform levels. In addition to the customizing, significant testing is also needed prior to declaring the cluster as production ready.

The cluster software products we will be using in this book are flexible platforms that allow integration of generic applications running on AIX, Linux, Microsoft Windows platforms, and providing for high available systems at a reasonable cost.

1.1.4 High availability solutions

The high availability (HA) solutions can provide many advantages compared to other solutions. In Table 1-2, we describe some HA solutions and their characteristics.

Table 1-2 Types of HA solutions

Solutions	Standalone	Enhanced Standalone	High Availability Clusters	Fault-Tolerant Computers
Downtime	Couple of days	Couple of hours	Couple of minutes	Never stop
Data Availability	Last full Backup	Last transaction	Last transaction	No loss of data

High availability solutions offer the following benefits:

- Standard components
- Can be used with the existing hardware
- Work with just about any application
- Work with a wide range of disk and network types
- Excellent availability at reasonable cost
- Proven solutions, most are mature technologies (HACMP, VCS, MSCS)
- Flexibility (most applications can be protected using HA clusters)
- Using “of the shelf” hardware components

Considerations for providing high availability solutions include:

- Thorough design and detailed planning
- Elimination of single points of failure
- Selection of appropriate hardware
- Correct implementation (no “shortcuts”)
- Disciplined system administration practices
- Documented operational procedures
- Comprehensive testing

Archived

1.2 Cluster concepts

The basic concepts can be classified as follows:

- Cluster topology:

  Contains basic cluster components nodes, networks, communication interfaces, communication devices, and communication adapters.

- Cluster resources:

  Entities that are being made highly available (for example, file systems, raw devices, service IP labels, and applications). Resources are grouped together in resource groups/service groups, which the cluster software keeps highly available as a single entity.

  Resource groups can be available from a single node (active-passive) or, in the case of concurrent applications, available simultaneously from multiple nodes (active-active).

- Failover:

  Represents the movement of a resource group from one active node to another node (backup node) in response to a failure on that active node.

- Fallback:

  Represents the movement of a resource group back from the backup node to the previous node, when it becomes available. This movement is typically in response to the reintegration of the previously failed node.

1.3 Cluster terminology

To understand the correct functionality and utilization of cluster solutions, it is necessary to know some important terms:

- Cluster:

  Loosely-coupled collection of independent systems (nodes) organized into a network for the purpose of sharing resources and communicating with each other.

  These individual nodes are together responsible for maintaining the functionality of one or more applications in case of a failure of any cluster component.

- Node:

  A machine running any operational systems and a cluster software defined as part of a cluster. Each node has a collection of resources (disks, file systems, IP address(es), and applications) that can be transferred to another node in the cluster in case the node fails.

Archived

- Resource:

  Resources are logical components of the cluster configuration that can be moved from one node to another. All the logical resources necessary to provide a highly available application or service are grouped together in a resource group.

  The components in a resource group move together from one node to another in the event of a node failure. A cluster may have more than one resource group, thus allowing for efficient use of the cluster nodes

- Takeover:

  This is the operation of transferring resources between nodes inside the cluster. If one node fails due to a hardware problem or operational system crash, its resources and applications will be moved to another node.

- Clients:

  A client is a system that can access the application running on the cluster nodes over a local area network. Clients run a client application that connects to the server (node) where the application runs.

- Heartbeating:

  In order for a cluster to recognize and respond to failures, it must continually check the health of the cluster. Some of these checks are provided by the heartbeat function. Each cluster node sends heartbeat messages at specific intervals to other cluster nodes, and expects to receive heartbeat messages from the nodes at specific intervals. If messages stop being received, the cluster software recognizes that a failure has occurred.

- Heartbeats can be sent over:
  - TCP/IP networks
  - Point-to-point networks
  - Shared disks.

Archived

Archived

2

Building a highly available Tivoli Storage Manager cluster environment

In this chapter we discuss and demonstrate the building of a highly available Tivoli Storage Manager cluster.

© Copyright IBM Corp. 2005. All rights reserved.

2.1 Overview of the cluster application

Here we introduce the technology that we will work with throughout this book — the IBM Tivoli Storage Manager products, which we have used for our clustered applications. Any concept or configuration that is specific to a particular platform or test scenario will be discussed in the pertinent chapters.

2.1.1 IBM Tivoli Storage Manager Version 5.3

In this section we provide a brief overview of Tivoli Storage Manager Version 5.3 features. If you would like more details on this new version, please refer to the following IBM Redbook: *IBM Tivoli Storage Manager Version 5.3 Technical Guide,* SG24-6638-00.

Tivoli Storage Manager V5.3 new features overview

IBM Tivoli Storage Manager V5.3 is designed to provide significant improvements to the ease of use and ease of administration and serviceability characteristics. These enhancements help you improve the productivity of personnel administering and using IBM Tivoli Storage Manager. Additionally, the product is easier to use for new administrators and users.

- Improved application availability:
  - IBM Tivoli Storage Manager for Space Management: HSM for AIX JFS2,enhancements to HSM for AIX and Linux GPFS
  - IBM Tivoli Storage Manager for application products update
- Optimized storage resource utilization:
  - Improved device management, SAN attached device dynamic mapping, native STK ACSLS drive sharing and LAN-free operations, improved tape checkin and checkout, and label operations, and new device support
  - Disk storage pool enhancements, collocation groups, proxy node support, improved defaults, reduced LAN-free CPU utilization, parallel reclamation and migration
- Enhanced storage personnel productivity:
  - New Administrator Web GUI
  - Task-oriented interface with wizards to simplify tasks such as scheduling, managing server maintenance operations (storage pool backup, migration, reclamation), and configuring devices

Archived

- Health monitor which shows status of scheduled events, the database and recovery log, storage devices, and activity log messages
- Calendar-based scheduling for increased flexibility of client and administrative schedules
- Operational customizing for increased ability to control and schedule server operations

Server enhancements, additions, and changes

This section lists all the functional enhancements, additions, and changes for the IBM Tivoli Storage Manager Server introduced in Version 5.3, as follows:

- ACSLS Library Support Enhancements
- Accurate SAN Device mapping for UNIX® Servers
- ACSLS Library Support Enhancements
- Activity Log Management
- Check-In and Check-Out Enhancements
- Collocation by Group
- Communications Options
- Database Reorganization
- Disk-only Backup
- Enhancements for Server Migration and Reclamation Processes
- IBM 3592 WORM Support
- Improved Defaults
- Increased Block Size for Writing to Tape
- LAN-free Environment Configuration
- NDMP Operations
- Net Appliance SnapLock Support
- New Interface to Manage Servers: Administration Center
- Server Processing Control in Scripts
- Simultaneous Write Inheritance Improvements
- Space Triggers for Mirrored Volumes
- Storage Agent and Library Sharing Fallover
- Support for Multiple IBM Tivoli Storage Manager Client Nodes
- IBM Tivoli Storage Manager Scheduling Flexibility

Archived

Client enhancements, additions, and changes

This section lists all the functional enhancements, additions, and changes for the IBM Tivoli Storage Manager Backup Archive Client introduced in Version 5.3, as follows:

- Include-exclude enhancements
- Enhancements to query schedule command
- IBM Tivoli Storage Manager Administration Center
- Support for deleting individual backups from a server file space
- Optimized option default values
- New links from the backup-archive client Java™ GUI to the IBM Tivoli Storage Manager and Tivoli Home Pages
- New options, Errorlogmax and Schedlogmax, and DSM_LOG environment variable changes
- Enhanced encryption
- Dynamic client tracing
- Web client enhancements
- Client node proxy support [asnodename]
- Java GUI and Web client enhancements
- IBM Tivoli Storage Manager backup-archive client for HP-UX Itanium 2
- Linux for zSeries® offline image backup
- Journal based backup enhancements
- Single drive support for Open File Support (OFS) or online image backups.

2.1.2 IBM Tivoli Storage Manager for Storage Area Networks V5.3

IBM Tivoli Storage Manager for Storage Area Networks is a feature of Tivoli Storage Manager that enables LAN-free client data movement. This feature allows the client system to directly write data to, or read data from, storage devices attached to a storage area network (SAN), instead of passing or receiving the information over the local area network (LAN).

Data movement is thereby off-loaded from the LAN and from the Tivoli Storage Manager server, making network bandwidth available for other uses.

Archived

The new version of Storage Agent supports communication with Tivoli Storage Manager clients installed on other machines. You can install the Storage Agent on a client machine that shares storage resources with a Tivoli Storage Manager server as shown in Figure 2-1, or on a client machine that does not share storage resources but is connected to a client machine that does share storage resources with the Tivoli Storage Manager server.

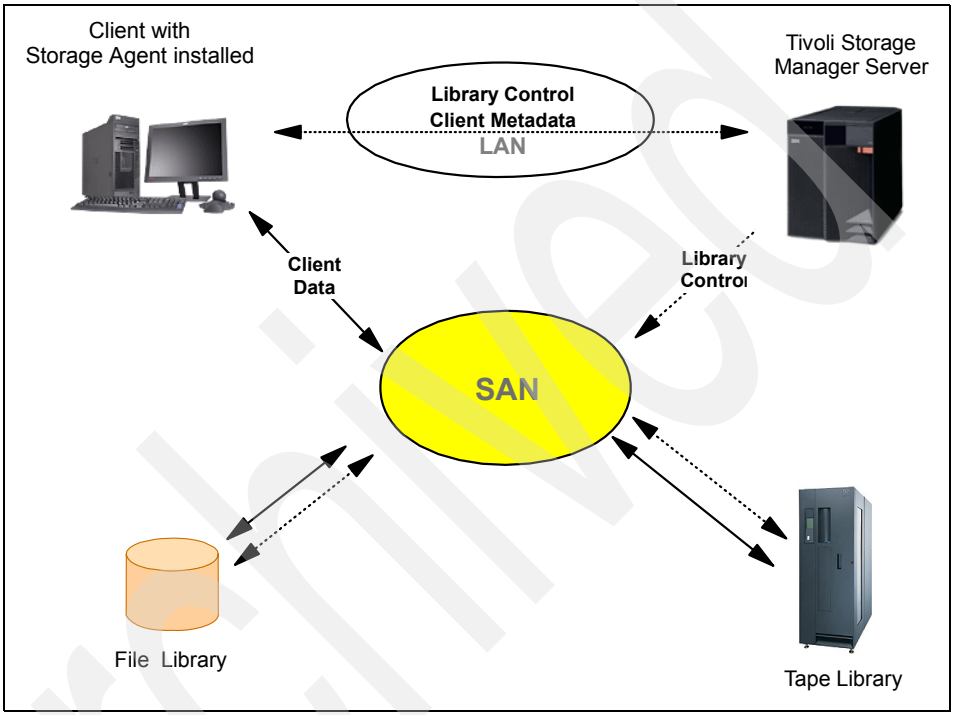

Figure 2-1 Tivoli Storage Manager LAN (Metadata) and SAN data flow diagram

Archived

Figure 2-2 shows multiple clients connected to a client machine that contains the Storage Agent.

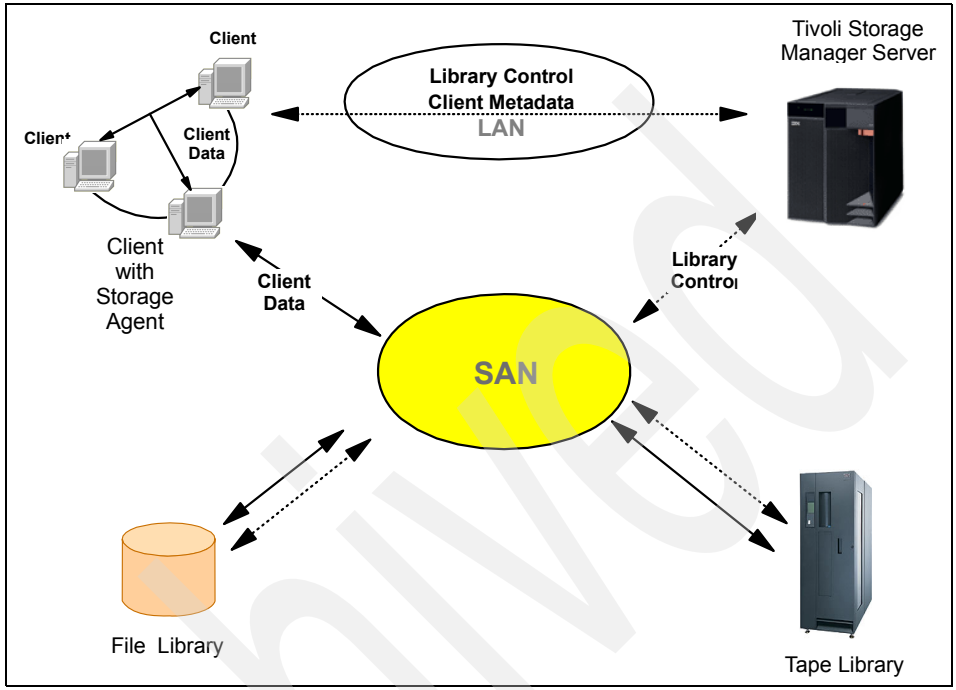

Figure 2-2 Multiple clients connecting through a single Storage Agent

2.2 Design to remove single points of failure

When designing our lab environment for this book, we focused on eliminating as many single points of failure possible, within the cost and physical constraints which existed.

2.2.1 Storage Area Network considerations

Today, many of the physical device issues — which challenged highly available configurations in the past — have been removed with the implementation of SAN devices. Although these devices still utilize the SCSI command set, most of these challenges were physical connection limitations, and some challenges still exist in the architecture, primarily SCSI reserves.

Archived

The Tivoli Storage Manager V5.3 addresses most of the device reserve challenges; however, this is currently limited to the AIX server platform only. In the cases of other platforms, such as Linux, we have provided SCSI device resets within the starting scripts.

When planning the SAN, we will build redundancy into the fabrics, allowing for dual HBAs connecting to each fabric. We will keep our disk and tape on separate fabrics, and will also create separate aliases and zones each device separately. Our intent with this design is to isolate bus or device reset activity, as well as limiting access to the resources, to only those host systems which require that access.

2.2.2 LAN and network interface considerations

In most cases, multiple Network Interface Cards (NICs) are required for these configurations. Depending on the cluster software, at least two NICs that will be used for public network traffic will be required.

There are many options for configuring redundancy at the NIC layer, which will vary depending on the operating system platform. It is important to keep in mind that building redundancy into the design is critical, and is what brings value to the highly available cluster solution.

2.2.3 Private or heartbeat network considerations

Most clustering software will require two NICs for the private network which carry the heartbeat traffic (keep-alive packets). Some products will allow the use of RS232 or disk heartbeat solutions.

2.3 Lab configuration

First, we will diagram our layout, then review the connections, adapters, and ports required to ensure that we have the appropriate hardware to connect our environment, removing any single points of failure. Our final result for the complete lab SAN environment is shown in Figure 2-3.

Archived

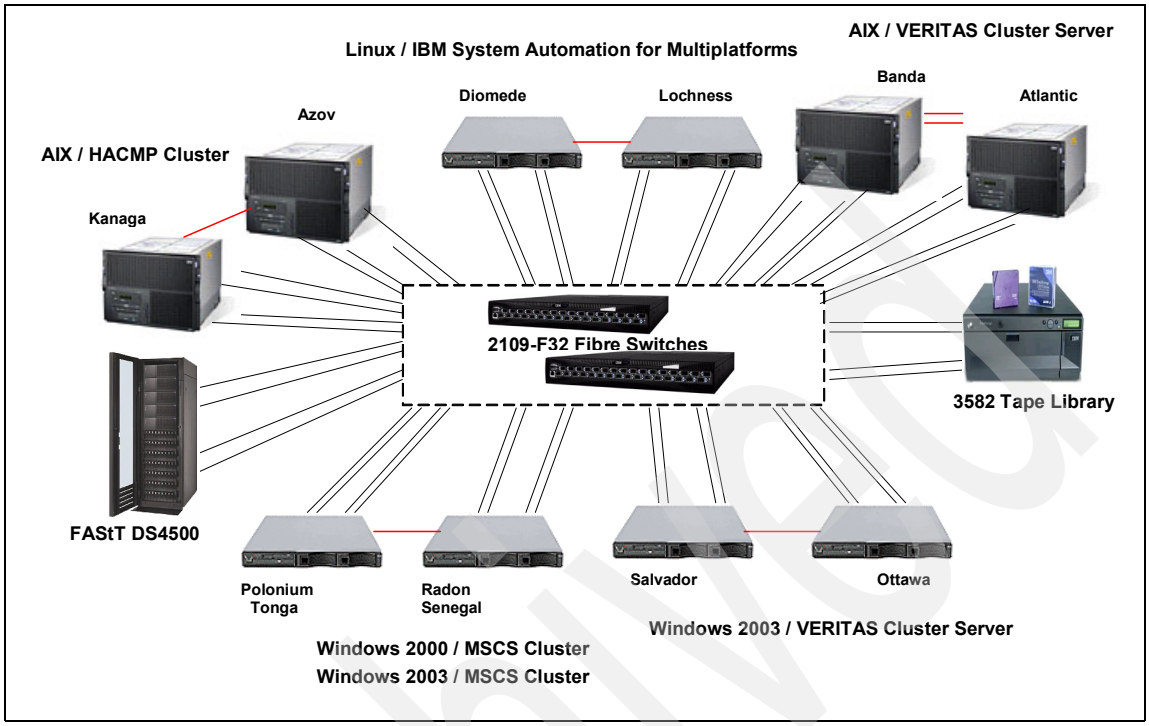

Figure 2-3 Cluster Lab SAN and heartbeat networks

Our connections for the LAN environment for our complete lab are shown in Figure 2-4.

Archived

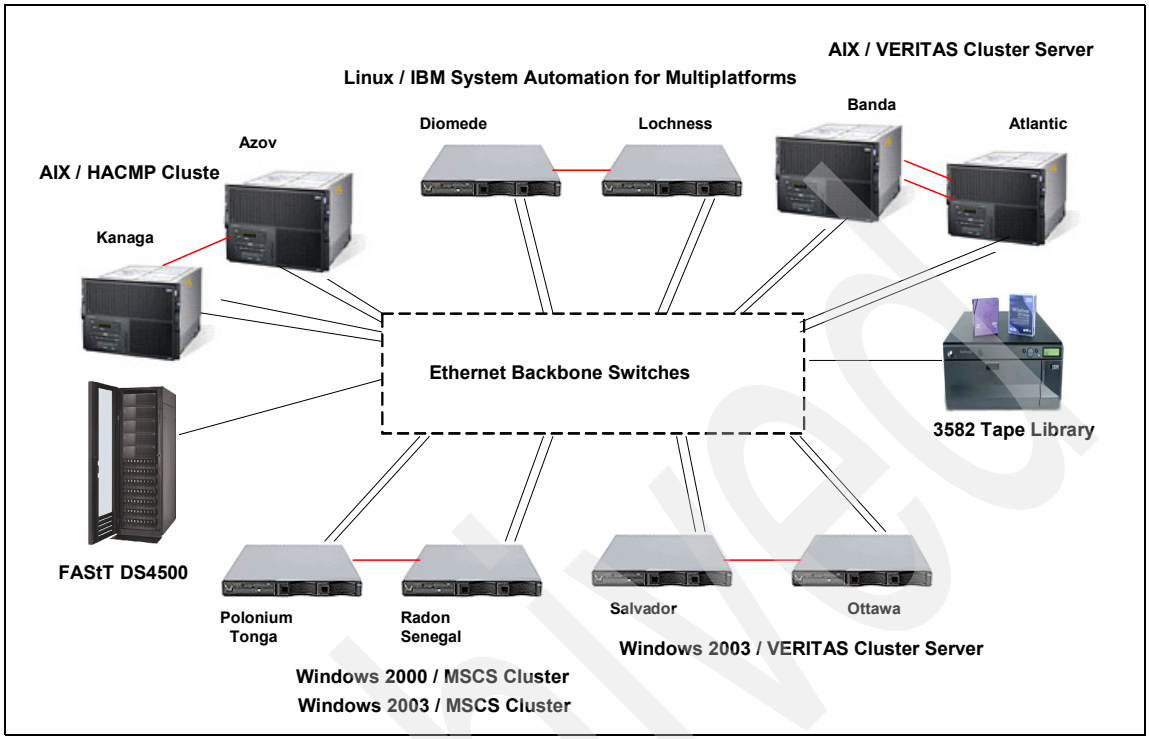

Figure 2-4 Cluster Lab LAN and hearbeat configuration

2.3.1 Cluster configuration matrix

In the following chapters we reference many different configurations, on multiple platforms. We illustrate the various configurations in Table 2-1.

Table 2-1 Cluster matrix

Cluster Name	TSM Name	Node A	Node B	Platform	Cluster SW
cl_mscs01	tsmsrv01	radon	polonium	win2000 sp4	MSCS
cl_mscs02	tsmsrv02	senegal	tonga	win2003 sp1	MSCS
cl_hacmp01	tsmsrv03	azov	kanaga	AIX V5.3	HACMP V5.2
cl_veritas01	tsmsrv04	atlantic	banda	AIX V5.2 ml4	VCS V4.0
cl_VCS02	tsmsrv06	salvador	ottawa	win2003 sp1	VSFW V4.2
cl_itsamp01	tsmsrv05	lochness	diomede	RH ee3	ITSAMP V1.2 fp3
cl_itsamp02	tsmsrv07	azov	kanaga	AIX V5.3	ITSAMP V1.2 fp3

Archived

2.3.2 Tivoli Storage Manager configuration matrix

All the Tivoli Storage Manager Server configurations will be using a 25 GB diskpool protected by hardware RAID-5. We illustrate some configuration differences, as shown in Table 2-2.

Table 2-2 Tivoli Storage Manager configuration matrix

TSM Name	TSM DB & LOG Mirror	Mirroring Method	DB Page Shadowing	Mirroring Mode	Logmode
tsmsrv01	NO	HW Raid-5	YES	N/A	Roll Forward
tsmsrv02	YES	TSM	YES	Parallel	Roll Forward
tsmsrv03	YES	TSM	NO	Sequential	Roll Forward
admcnt01	N/A	HW Raid-5	N/A	N/A	N/A
tsmsrv04	YES	AIX	YES	na	Roll Forward
tsmsrv06	YES	TSM	YES	Parallel	Roll Forward
tsmsrv05	YES	TSM		Parallel	Roll Forward
tsmsrv07	YES	AIX	YES	Parallel	Roll Forward

Archived

3

Testing a highly available Tivoli Storage Manager cluster environment

In this chapter we discuss the testing of our cluster configurations. We focus on two layers of testing:

- Cluster infrastructure
- Application (Tivoli Storage Manager Server, Client, StorageAgent) failure and recovery scenarios

© Copyright IBM Corp. 2005. All rights reserved.

3.1 Objectives

Testing highly available clusters is a science. Regardless of how well the solution is architected or implemented, it all comes down to how well you test the environment. If the tester does not understand the application and its limitations, or doesn't understand the cluster solution and its implementation, there will be unexpected outages.

The importance of creative, thorough testing cannot be emphasized enough. The reader should not invest in cluster technology unless they are prepared to invest in the testing time, both pre-production and post-production. Here are the major task items involved in testing a cluster:

- Build the testing scope.
- Build the test plan.
- Build a schedule for testing of the various application components.
- Document the initial test results.
- Hold review meetings with the application owners, discuss and understand the results, and build the next test plans.
- Retest as required from the review meetings.
- Build process documents, including dataflow and an understanding of failure situations with anticipated results.
- Build recovery processes for the most common user intervention situations.
- Prepare final documentation.

Important: Planning for the appropriate testing time in a project is a challenge, and is often the forgotten or abused phase. It is our team's experience that the testing phase must be at least two times the total implementation time for the cluster (including the customizing for the applications.

3.2 Testing the clusters

As we will emphasize throughout this book, testing is critical towards building a successful (and reliable) Tivoli Storage Manager cluster environment.

Archived

3.2.1 Cluster infrastructure tests

The following cluster infrastructure tests should be performed:

- Manual failover for the core cluster
- Manual failback for the core cluster
- Start each Resource Group (Service Group)
- Stop each Resource Group (Service Group)
- Test FC adapter failure
- Test FC adapter recovery
- Test public NIC failure
- Test public NIC recovery
- Test private NIC failure
- Test private NIC recovery
- Test disk heartbeat failure
- Test disk heatbeat recovery
- Test power failure of each node
- Test power failure recovery of each node

Ensuring that a reliable, predictable, highly available cluster has been designed and implemented, these would be considered a minimal set of cluster infrastructure tests.

For each of these tests, a document detailing the testing process and resulting behavior should be produced. Following this regimen will ensure that issues will surface, be resolved, and be retested, thus producing final documentation.

3.2.2 Application tests

Resource Group (or Service Group) testing includes the complete Application (Tivoli Storage Manager component) and all the associated resources supporting the application.

Tivoli Storage Manager Server tests

These tests are designed around Tivoli Storage Manager server failure situations. The Tivoli Storage Manager server is highly available:

- Server nodeA fails during a scheduled client backup to diskpool.
- Server recovers on nodeB during a scheduled client backup to diskpool.
- Server nodeA fails during a migration from disk to tape.

Archived

- Server node recovers on nodeB after the migration failure.
- Server nodeA fails during a backup storage pool tape to tape operation.
- Server recovers on nodeB after the backup storage pool failure.
- Server nodeA fails during a full DB backup to tape.
- Server recovers on nodeB after the full DB backup failure.
- Server nodeA fails during an expire inventory.
- Server recovers on nodeB after failing during an expire inventory.
- Server nodeA fails during a StorageAgent backup to tape.
- Server recovers on nodeB after failing during a StorageAgent backup to tape.
- Server nodeA fails during a session serving as a library manager for a library client.
- Server recovers on nodeB after failing as a library manager.

Tivoli Storage Manager Client tests

These are application tests for a highly available Tivoli Storage Manager client:

- Client nodeA fails during a scheduled backup.
- Client recovers on nodeB after failing during a scheduled backup.
- Client nodeA fails during a client restore.
- Client recovers on nodeB after failing during a client restore.

Tivoli Storage Manager Storage Agent tests

These are application tests for a highly available Tivoli Storage Manager Storage Agent (and the associated Tivoli Storage Manager client):

- StorageAgent nodeA fails during a scheduled backup to tape.
- StorageAgent recovers on nodeB after failing during a scheduled backup.

Archived

Part 2

Clustered Microsoft Windows environments and IBM Tivoli Storage Manager Version 5.3

In this part of the book, we discuss the implementation of Tivoli Storage Manager products with Microsoft Cluster Server (MSCS) in Windows 2000 and 2003 Server environments.

© Copyright IBM Corp. 2005. All rights reserved.

Archived

4

Microsoft Cluster Server setup

This chapter provides general information about the tasks needed to set up *Microsoft Cluster Services* (MSCS) in the following environments:

- Two servers with Windows 2000 Advanced Server
- Two servers with Windows 2003 Enterprise Server

© Copyright IBM Corp. 2005. All rights reserved.

4.1 Overview

Microsoft Cluster Service (MSCS) is one of the Microsoft solutions for high availability, where a group of two or more servers together form a single system, providing high availability, scalability, and manageability for resources and applications. For a generic approach on how to set up a Windows 2003 cluster, please refer to the following Web site:

http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/clustering/confclus.mspx

4.2 Planning and design

Our software/hardware should meet the requirements established by Microsoft:

- For Windows 2000 servers:
  - Microsoft Windows 2000 Advanced Server or Microsoft Windows 2000 Datacenter Server installed on all computers in the cluster and belonging to a same domain. We recommend to apply all latest available service packs and patches for each node.
- For Windows 2003 servers:
  - Microsoft Windows Server 2003 Enterprise Edition or Windows 2003 Datacenter Edition installed on all computers in the cluster and belonging to a same domain. We recommend to apply all latest available service packs and patches for each node.
- At least two network adapter cards on each node. Since we want a high available environment, we do not use multiport network adapters. Also, we do not use teaming for the heartbeat. If it is necessary fault tolerance, we can use two network adapter cards.
- An SCSI or Fibre Channel adapter.
- One or more external disks on either an SCSI or Fibre Channel bus.
- A *Domain Name System* (DNS) server.
- An account in the domain that belongs to the local administrators group on each node, that will be used to start MSCS service.

All nodes should belong to the same domain, have access to the domain controllers and DNS servers in the network. However, it is still possible to not have a Windows based environment with domain controllers. In this case, we will need to set up at least 2 servers as domain controllers and DNS servers.

Archived

All hardware used in the solution must be on the *Hardware Compatibility List* (HCL) that we can find at `http://www.microsoft.com/hcl`, under "*cluster*". For more information, see the following articles from *Microsoft Knowledge Base*:

- *309395 The Microsoft Support Policy for Server Clusters and the Hardware*
- *304415 Support for Multiple Clusters Attached to the Same SAN Device*

4.3 Windows 2000 MSCS installation and configuration

In this section we describe all the tasks and our lab environment to install and configure MSCS in two Windows 2000 Advanced Servers, POLONIUM and RADON.

4.3.1 Windows 2000 lab setup

Figure 4-1 shows the lab we use to set up our Windows 2000 Microsoft Cluster Services:

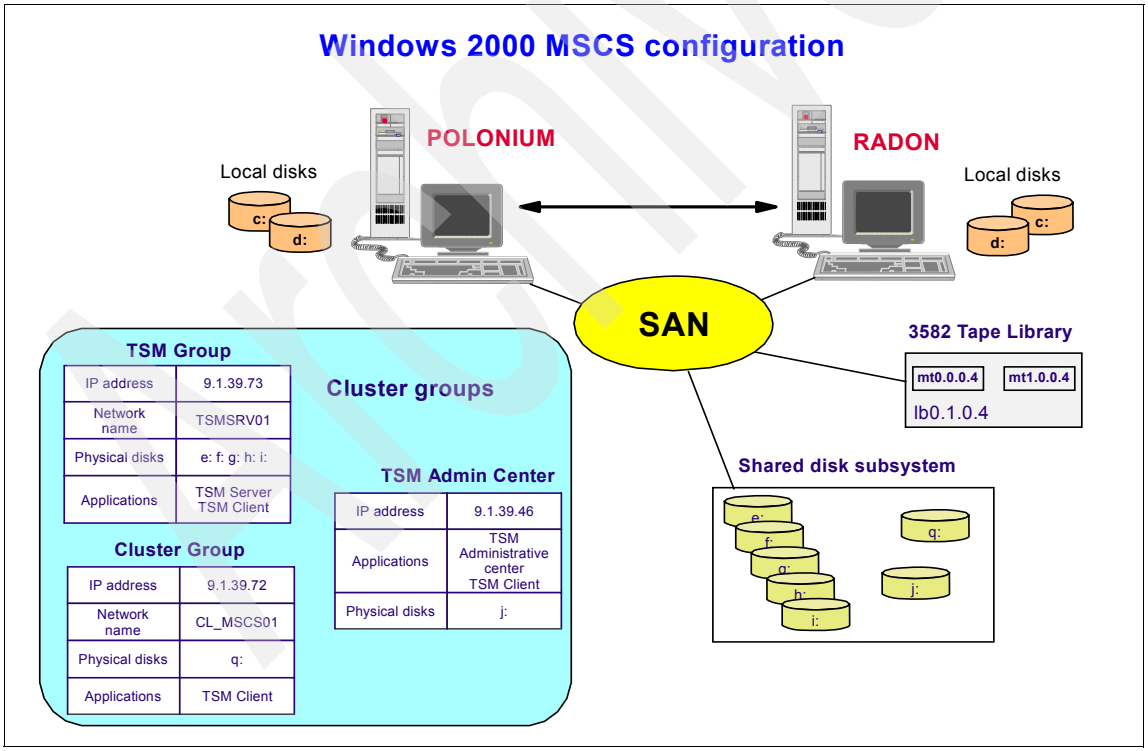

Figure 4-1 Windows 200 MSCS configuration

Table 4-1, Table 4-2, and Table 4-3 describe our lab environment in detail.

Table 4-1 Windows 2000 cluster server configuration

MSCS Cluster	
Cluster name	CL_MSCS01
Cluster IP address	9.1.39.72
Network name	CL_MSCS01
Node 1	
Name	POLONIUM
Private network IP address	10.0.0.1
Public network IP address	9.1.39.187
Node 2	
Name	RADON
Private network IP address	10.0.0.2
Public network IP address	9.1.39.188

Archived

Table 4-2 Cluster groups for our Windows 2000 MSCS

Cluster Group 1	
Name	Cluster Group
IP address	9.1.39.72
Network name	CL_MSCS01
Physical disks	q:
Applications	TSM Client
Cluster Group 2	
Name	TSM Admin Center
Physical disks	j:
IP address	9.1.39.46
Applications	IBM WebSphere® Application Server ISC Help Service TSM Client
Cluster Group 3	
Name	TSM Group
IP address	9.1.39.73
Network name	TSMSRV01
Physical disks	e: f: g: h: i:
Applications	TSM Server, TSM client

Table 4-3 Windows 2000 DNS configuration

Domain	
Name	TSMW2000
Node 1	
DNS name	polonium.tsmw2000.com
Node 2	
DNS name	radon.tsmw200.com

Archived

4.3.2 Windows 2000 MSCS setup

We install Windows 2000 Advanced or Database Server on each of the machines that form the cluster. At this point, we do not need to have the shared disks attached to the servers yet. If we have, it is better to shut them down to avoid corruption.

Network setup

After we install the OS, we turn on both servers and we set up the networks with static IP addresses.

One adapter is to be used only for internal cluster communications, also known as heartbeat. It needs to be in a different network from the public adapters. We use a cross-over cable in a two-node configuration, or a dedicated hub if we have more servers in the cluster.

The other adapters are for all other communications and should be in the public network.

For ease of use we rename the network connections icons to **Private** (for the heartbeat) and **Public** (for the public network) as shown in Figure 4-2.

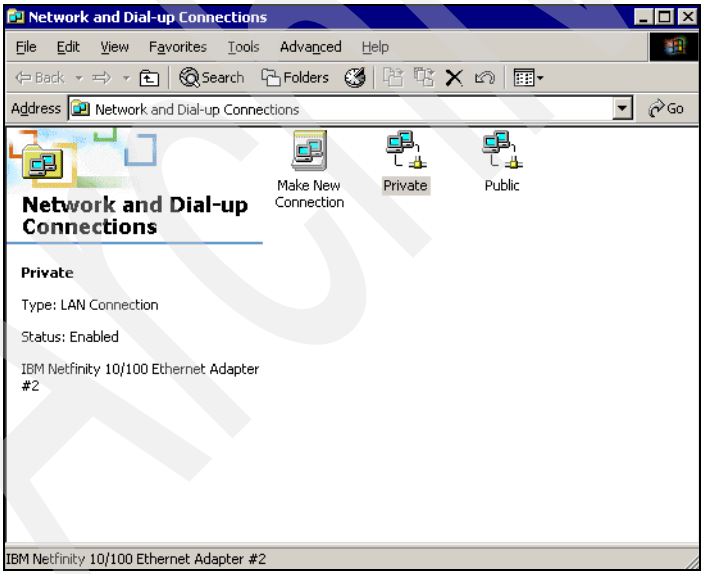

Figure 4-2 Network connections windows with renamed icons

Archived

We also recommend to set up the binding order of the adapters, leaving the public adapter in the top position. We go to the **Advanced** menu on the Network and Dial-up Connections menu and in the **Connections** box, we change to the order shown in Figure 4-3.

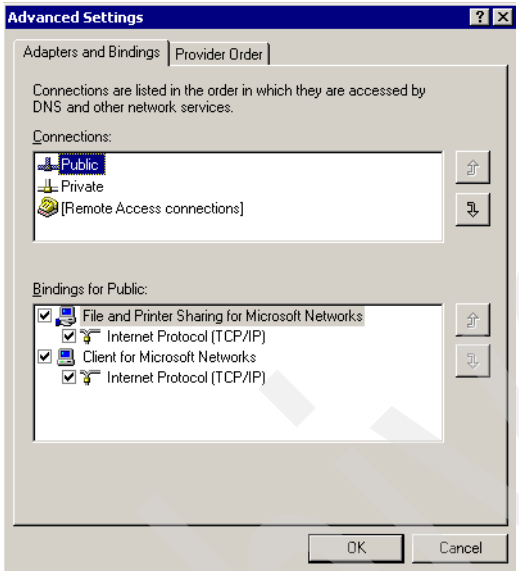

Figure 4-3 Recommended bindings order

Private network configuration

When setting up the private network adapter, we choose any static IP address that is not on the same subnet or network as the public network adapter. For the purpose of this book, we use 10.0.0.1 and 10.0.0.2 with 255.255.255.0 mask.

Also, we make sure we have the following configuration in the TCP/IP properties:

- There should be **no default gateway**.
- In the **Advanced** button, DNS tab, we uncheck the option **Register this connection's addresses in DNS**.
- In the **Advanced** button, **WINS** tab, we click **Disable NetBIOS over TCP/IP**. If we receive a message: "This connection has an empty primary WINS address. Do you want to continue?", we should click **Yes**.
- On the **Properties** tab of the network adapter, we manually set the speed to 10 Mbps/Half duplex.

We must make sure these settings are set up for all the nodes.

Archived

Public network configuration

We do not have to use DHCP so that cluster nodes will not be inaccessible if the DHCP server is unavailable.

We set up TCP/IP properties including DNS and WINS addresses.

Connectivity testing

We test all communications between the nodes on the public and private networks using the **`ping`** command locally and also on the remote nodes for each IP address.

We make sure name resolution is also working. For that we ping each node using the node's machine name. Also we use PING -a to do reverse lookup.

Domain membership

All nodes must be members of the same domain and have access to a DNS server. In this lab we set up the servers both as domain controllers as well as DNS Servers. If this is your scenario, use **`dcpromo.exe`** to promote the servers to domain controllers.

Promoting the first server

These are the steps:

1. We set up our network cards so that the servers point to each other for primary DNS resolution, and to themselves for secondary resolution.
2. We run **`dcpromo`** and create a new domain, a new tree and a new forest.
3. We take note of the password used for the administrator account.
4. We allow the setup to install DNS server.
5. We wait until the setup finishes and boot the server.
6. We configure the DNS server and create a *Reverse Lookup Zones* for all our network addresses. We make them active directory integrated zones.
7. We define new hosts for each of the nodes with the option of creating the associated pointer (PTR) record.
8. We test DNS using **`nslookup`** from a command prompt.
9. We look for any error messages in the event viewer.

Promoting the other servers

These are the steps:

1. We run **`dcpromo`** and join the domain created above, selecting **Additional domain controller for an existing domain**.

2. We use the password set up in step 3 on page 34 above.
3. When the server boots, we install **DNS server**.
4. We check if DNS is replicated correctly using **`nslookup`**.
5. We look for any error messages in the event viewer.

Setting up a cluster user account

Before going on and installing the cluster service, we create a cluster user account that will be required to bring the service up. This account should belong to the administrators group on each node. For security reasons we set the password settings to **User Cannot Change Password** and **Password Never Expires**.

Setting up external shared disks

When we install the SCSI/fibre adapter, we always use the same slot for all servers.

> **Attention:** While configuring shared disks, we have always only one server up at a time, to avoid corruption. To proceed, we shut down all servers, turn on the storage device and turn on only one of the nodes.

On the DS4500 side we prepare the LUNs that will be designated to our servers. A summary of the configuration is shown in Figure 4-4.

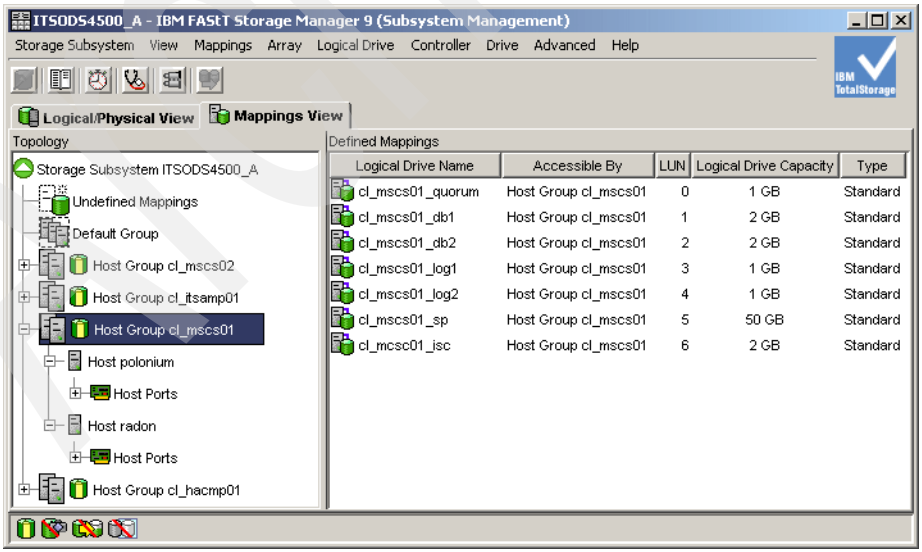

Figure 4-4 LUN configuration for Windows 2000 MSCS

We install the necessary drivers according to the manufacturer's manual, so that Windows recognizes the storage disks. The device manager should look similar to Figure 4-5 on the Disk drivers and SCSI and RAID controllers icons.

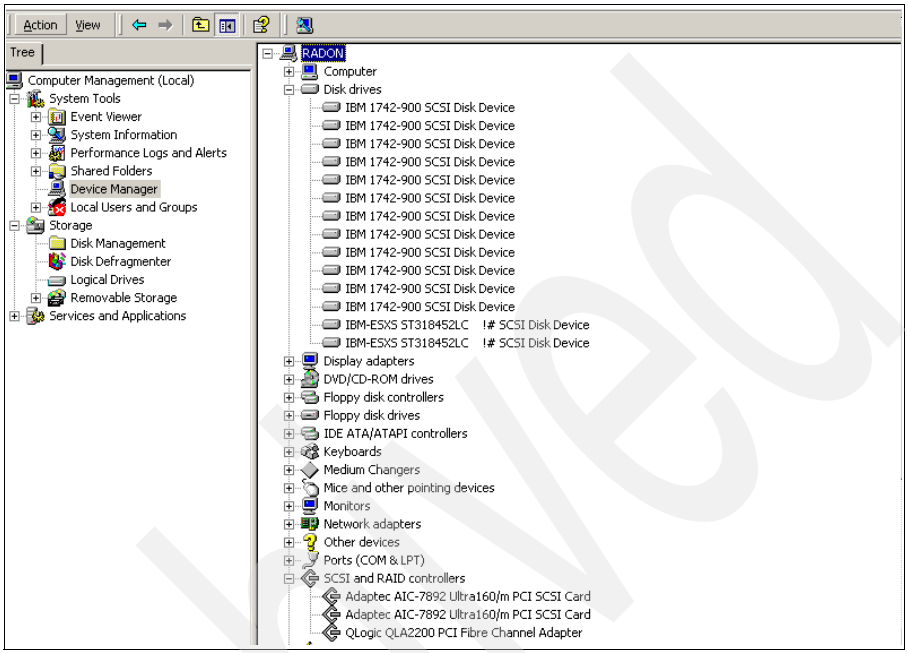

Figure 4-5 Device manager with disks and SCSI adapters

Configuring shared disks

To configure the shared disks:

1. We double-click **Disk Management** and the Write Signature and Upgrade Disk Wizard (Figure 4-6) begins:

Archived

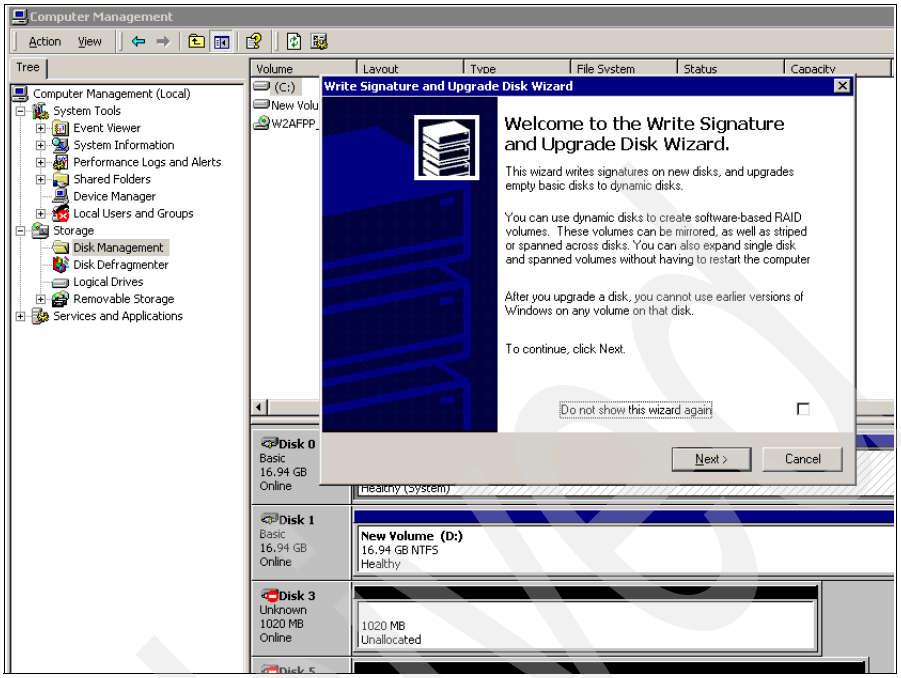

Figure 4-6 New partition wizard

2. We select all disks for the **Write Signature** part in Figure 4-7.

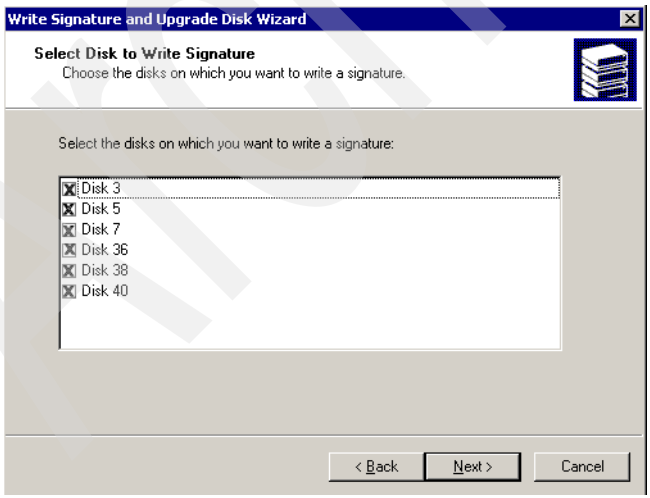

Figure 4-7 Select all drives for signature writing

3. We do not upgrade any of the disks to dynamic in Figure 4-8. In case we upgrade them, to be capable of resetting the disk to basic, we should right-click the disk we want to change, and we choose **Revert to Basic Disk**.

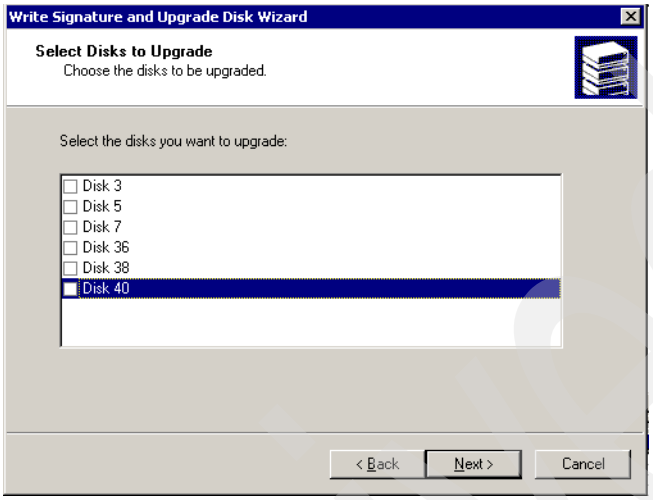

Figure 4-8 Do not upgrade any of the disks

4. We right-click each of the unallocated disks and the Create Partition Wizard begins. We select **Primary Partition** in Figure 4-9.

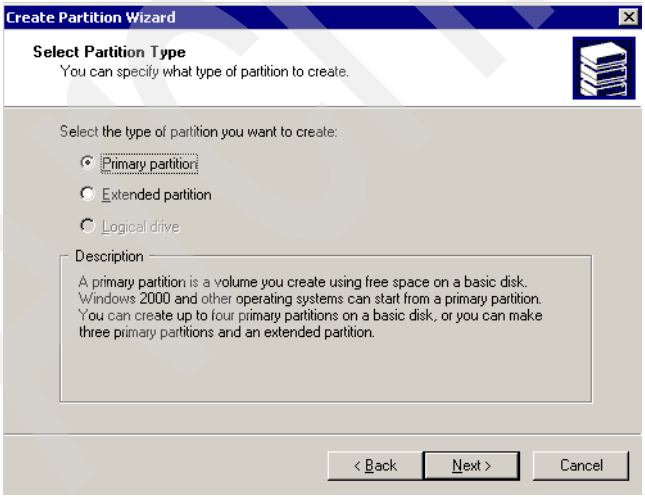

Figure 4-9 Select primary partition

5. We assign the partition size in Figure 4-10. We recommend to use only one partition per disk, assigning the maximum size.

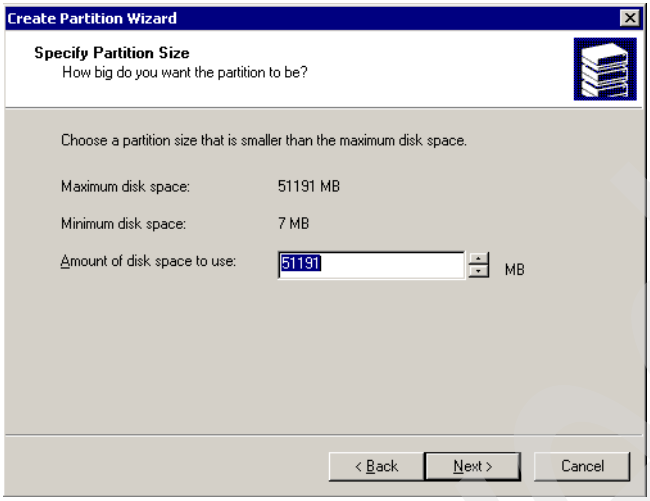

Figure 4-10 Select the size of the partition

6. We make sure to assign a drive mapping (Figure 4-11). This is *crucial* for the cluster to work. For the cluster quorum disk, we recommend to use drive q: and the name Quorum, for clarity reasons.

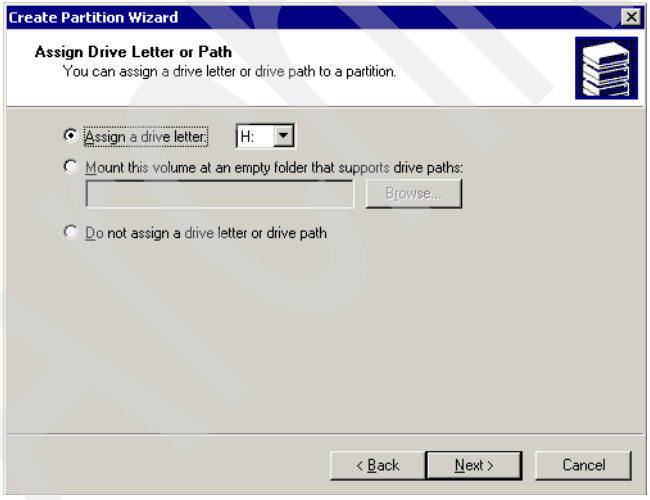

Figure 4-11 Drive mapping

7. We format the disk using NTFS (Figure 4-12) and we give it a name that reflects the application we will be setting up.

Archived

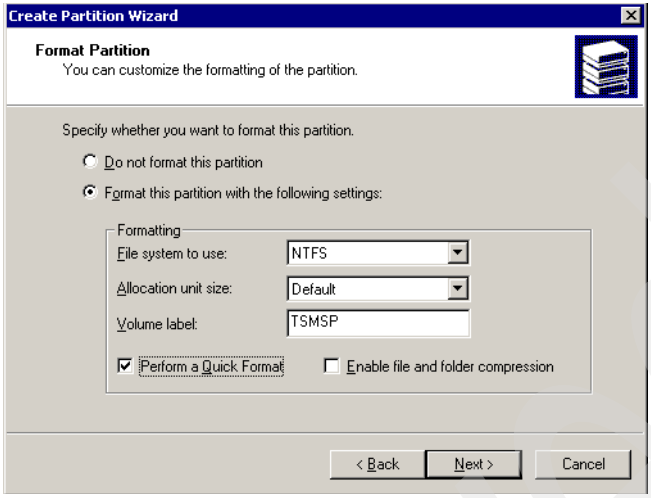

Figure 4-12 Format partition

8. We verify that all shared disks are formatted as NTFS and are healthy. We write down the letters assigned to each partition (Figure 4-13).

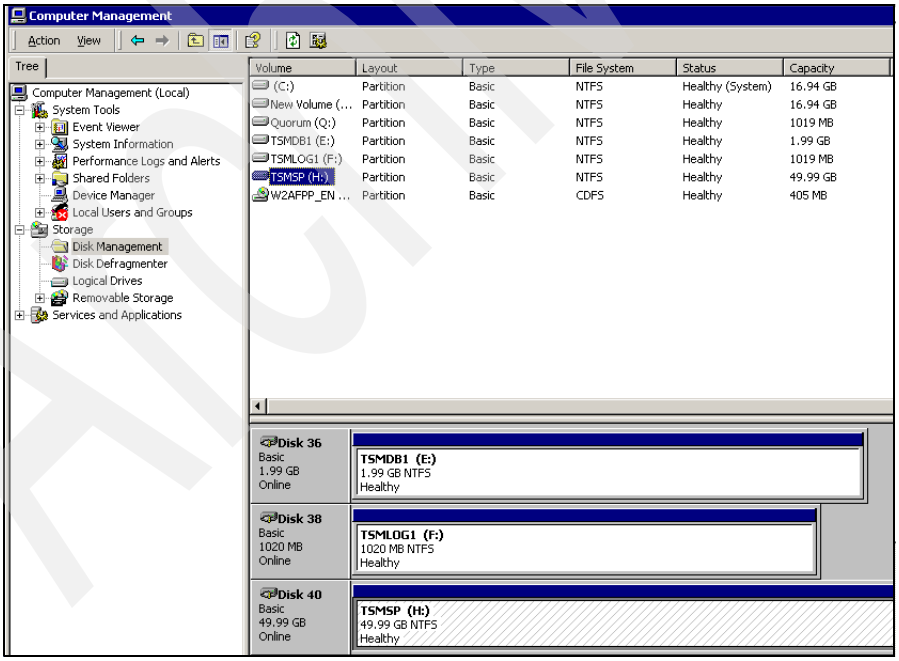

Figure 4-13 Disk configuration

9. We check disk access using the Windows Explorer menu. We create any file on the drives and we also try to delete it.
10. We repeat steps 2 to 6 for each shared disk.
11. We turn off the first node and turn on the second one. We check the partitions: if the letters are not set correctly, we change them to match the ones set up on the first node. We also test write/delete file access from the other node.

Windows 2000 cluster installation

Now that all of the environment is ready, we run the MSCS setup. The installation of the first node is different from the setup of the following nodes. Since the shared disks are still being recognized by both servers (with no sharing management yet), we turn on only the first node before starting the installation. This avoids disk corruption.

First node installation

To install MSCS in the first node:

From **Control Panel** → **Add/Remove Software** → **Add/Remove Windows Components**, we select **Cluster Service** and click **Next**.

> **Tip:** If you are using ServeRAID™ adapters, install the cluster service from the ServeRAID CD using `\programs\winnt\cluster\setup.exe`

12. We select **Next** to choose the Terminal Services Setup to accept the Remote administration mode.
13. The Cluster Service Configuration Wizard will start. We click **Next**.
14. We push the button **I understand** to accept the hardware notice and we click **Next**.
15. We select **The first node in the cluster** and click **Next**.
16. We give the cluster a name.
17. We type the username, password and domain created in “Setting up a cluster user account” on page 35. We click **Next**.
18. We choose the disks that will form the cluster and click **Next**.
19. We select the disk that will be the quorum disk (cluster management), **drive q:** and we click **Next**.
20. We click **Next** on the Configure Cluster Networks menu.

Archived

21. We configure the networks as follows:
    Private network for **internal cluster communications only**
    Public network for **all communications**
22. We set the network priority with the private network on the top.
23. We type the virtual TCP/IP address (the one that will be used by clients to access the cluster).
24. We **click Finish** and wait until the wizard completes the configuration. At completion we receive a notice saying the cluster service has started and that we have successfully completed the wizard.
25. We verify that the cluster name and IP address have been added to DNS. If they have not, we should do it manually.
26. We verify our access to the Cluster Management Console (**Start** →**Programs** → **Administrative Tools** → Cluster Administrator).
27. We keep this server up and bring the second node up to start the installation on it.

Second node installation

1. We repeat steps 1 to 4 of “First node installation” on page 41.
2. We select **The second or next node in the cluster** on the Create or Join a Cluster menu of the wizard, and we click **Next**.
3. We type our cluster name and we click **Next**.
4. We type the password for the cluster user and we click **Next**.
5. We click **Finish** and wait until the wizard completes the configuration. At completion we will receive a notice saying the cluster service has started successfully and that we have successfully completed the wizard.
6. It is necessary to repeat these steps for the remaining nodes, in case we had more than two nodes.

Windows 2000 cluster configuration

When the installation is complete the cluster looks like Figure 4-14, with one group resource for each disk. We may change this distribution, creating new groups with more than one disk resource, to best fit our environment.

Archived

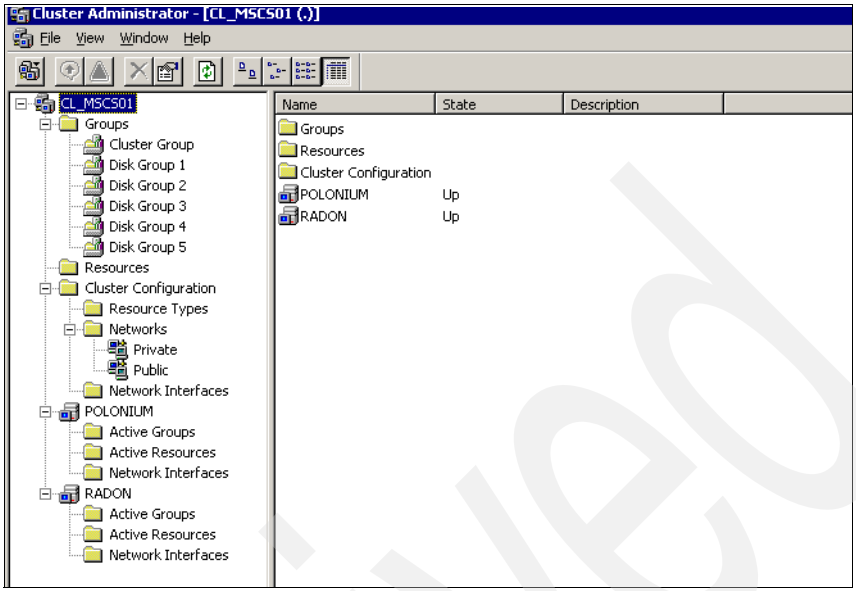

Figure 4-14 Cluster Administrator after end of installation

The next step is to group disks together so that we have only two groups: **Cluster Group** with the cluster name, ip and quorum disk, and **TSM Group** with all the other disks as shown in Figure 4-15.

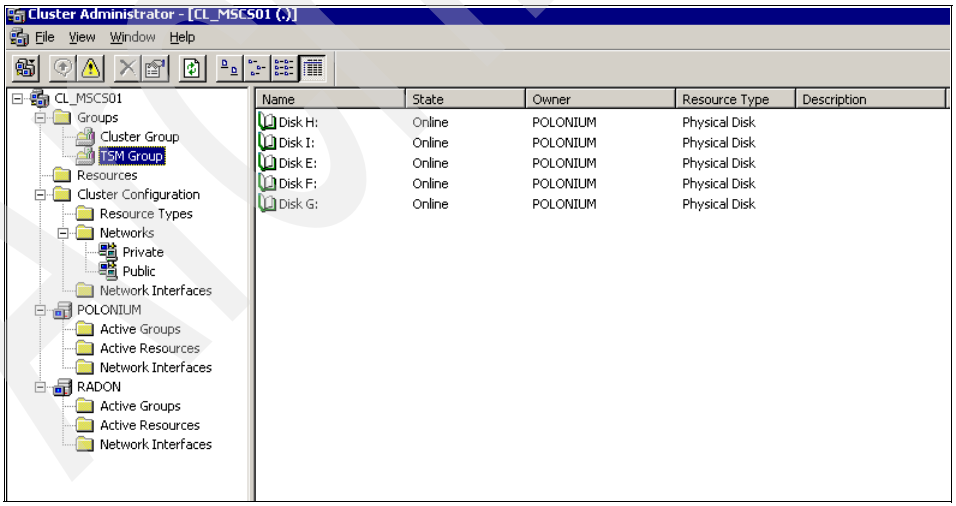

Figure 4-15 Cluster Administrator with TSM Group

Archived

In order to move disks from one group to another, we right-click the disk resource and we choose **Change Group**. Then we select the name of the group where the resource should move to.

> **Tip:** Microsoft recommends that for all Windows 2000 clustered environments, a change is made to the registry value for DHCP media sense so that if we lose connectivity on both network adapters, the network role in the server cluster for that network would not change to All Communications (Mixed Network). We set the value of DisableDHCPMediaSense to 1 in the following registry key:
>
> `HKLM\SYTEM\CurrentControlSetting\services\tcpip\parameters`
>
> For more information about this issue, read the article "254651 Cluster network role changes automatically in the Microsoft Knowledge Base".

Testing the cluster

To test the cluster functionality, we use the Cluster Administrator menu and we perform the following tasks:

- Moving groups from one server to another. We verify that resources fail over and are brought online on the other node.
- Moving all resources to one node and stopping the Cluster service. We verify that all resources fail over and come online on the other node
- Moving all resources to one node and shutting it down. We verify that all resources fail over and come online on the other node.
- Moving all resources to one node and removing the public network cable from that node. We verify that the groups will fail over and come online on the other node.

4.4 Windows 2003 MSCS installation and configuration

In this section we describe all the tasks and our lab environment to install and configure MSCS in two Windows 2003 Enterprise Servers, SENEGAL and TONGA.

Archived

4.4.1 Windows 2003 lab setup

Figure 4-16 shows the lab we use to set up our Windows 2003 Microsoft Cluster Services:

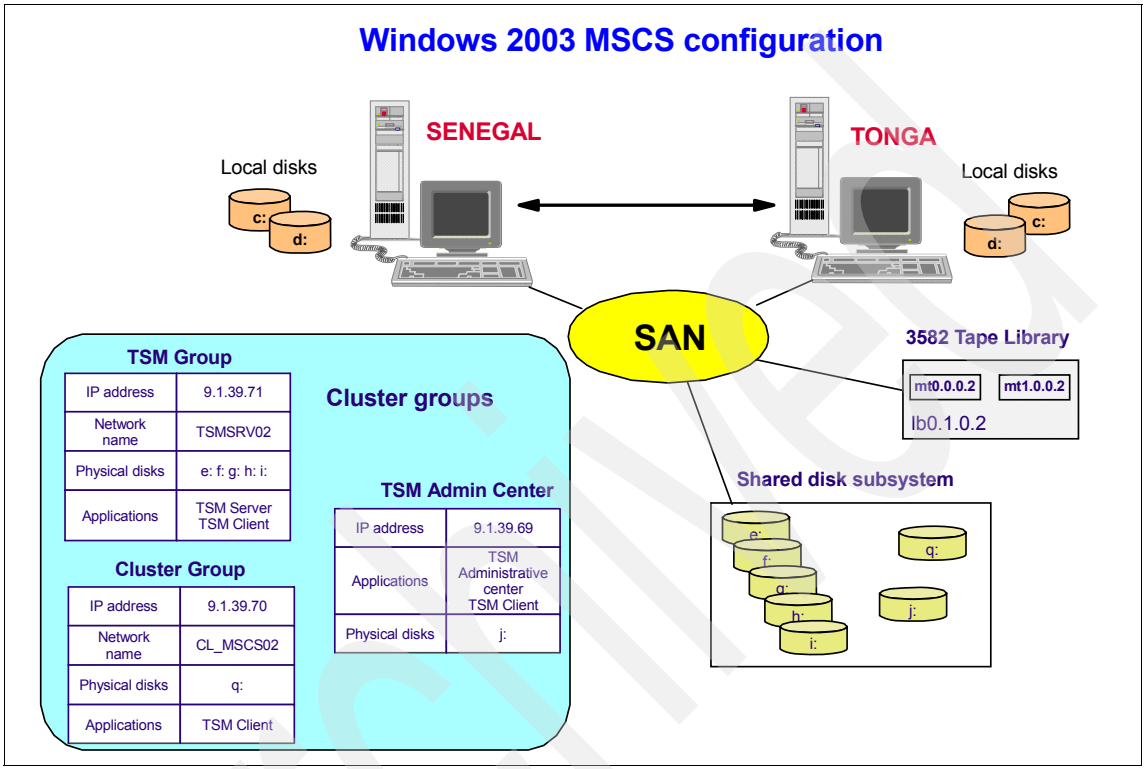

Figure 4-16 Windows 2003 MSCS configuration

Archived

Table 4-4, Table 4-5, and Table 4-6 describe our lab environment in detail.

Table 4-4 Windows 2003 cluster server configuration

MSCS Cluster	
Cluster name	CL_MSCS02
Cluster IP address	9.1.39.70
Network name	CL_MSCS02
Node 1	
Name	SENEGAL
Private network IP address	10.0.0.1
Public network IP address	9.1.39.166
Node 2	
Name	TONGA
Private network IP address	10.0.0.2
Public network IP address	9.1.39.168

Archived

Table 4-5 Cluster groups for our Windows 2003 MSCS

Cluster Group 1	
Name	Cluster Group
IP address	9.1.39.70
Network name	CL_MSCS02
Physical disks	q:
Cluster Group 2	
Name	TSM Admin Center
IP address	9.1.39.69
Physical disks	j:
Applications	IBM WebSphere Application Center ISC Help Service TSM Client
Cluster Group 3	
Name	TSM Group
IP address	9.1.39.71
Network name	TSMSRV02
Physical disks	e: f: g: h: i:
Applications	TSM Server, TSM client

Table 4-6 Windows 2003 DNS configuration

Domain	
Name	TSMW2003
Node 1	
DNS name	senegal.tsmw2000.com
Node 2	
DNS name	tonga.tsmw200.com

Archived

4.4.2 Windows 2003 MSCS setup

We install Windows 2003 Enterprise or Datacenter Edition on each of the machines that form the cluster. At this point, we do not need to have the shared disks attached to the servers yet. But if we did, it is best to shut them down to avoid corruption.

Network setup

After we install the OS, we turn on both servers and we set up the networks with static IP addresses.

One adapter is to be used only for internal cluster communications, also known as heartbeat. It needs to be in a different network from the public adapters. We use a cross-over cable in a two-node configuration, or a dedicated hub if we had more servers in the cluster.

The other adapters are for all other communications and should be in the public network.

For ease of use, we rename the network connections icons to **Private** (for the heartbeat) and **Public** (for the public network) as shown in Figure 4-17.

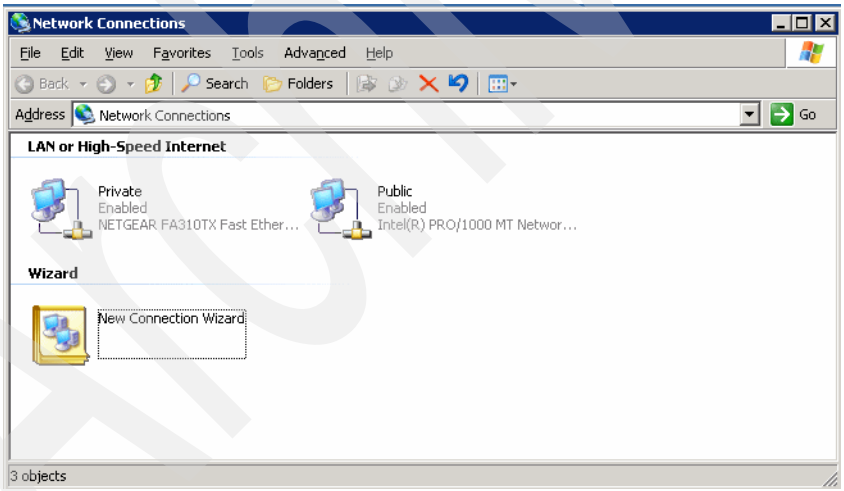

Figure 4-17 Network connections windows with renamed icons

We also recommend to set up the binding order of the adapters, leaving the public adapter in the top position. In the Network Connections menu, we select **Advanced** → **Advanced Settings**. In the Connections box, we change to the order shown below in Figure 4-18.

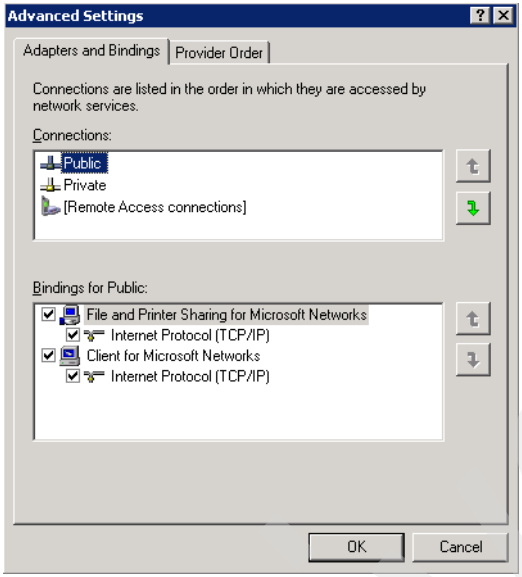

Figure 4-18 Recommended bindings order

Private network configuration

When setting up the private network adapter, we choose any static IP address that is not on the same subnet or network as the public network adapter. For the purpose of this book, we use 10.0.0.1 and 10.0.0.2 with 255.255.255.0 mask.

Also, we must make sure to have the following configuration in the TCP/IP properties:

- There should be **no default gateway**.
- In the Advanced button, DNS tab, we uncheck the option **Register this connection's addresses in DNS**.
- In the **Advanced** button, WINS tab, we click **Disable NetBIOS over TCP/IP**. If we receive a message: "This connection has an empty primary WINS address. Do you want to continue?", we should click **Yes**.
- On the **Properties** tab of the network adapter, we manually set the speed to 10 Mbps/Half duplex.

We make sure these settings are set up for all the nodes.

Public network configuration

We do not use DHCP so that cluster nodes will not be inaccessible if the DHCP server is unavailable.

Archived

We set up TCP/IP properties including DNS and WINS addresses.

Connectivity testing

We test all communications between the nodes on the public and private networks using the **ping** command locally and also on the remote nodes for each IP address.

We make sure name resolution is also working.For that, we ping each node using the node's machine name. We also use **`PING -a`** to do reverse lookup.

Domain membership

All nodes must be members of the same domain and have access to a DNS server. In this lab we set up the servers both as domain controllers and DNS Servers. If this is our scenario, we should use **dcpromo.exe** to promote the servers to domain controllers.

Promoting the first server

1. We set up our network cards so that the servers point to each other for primary DNS resolution and to themselves for secondary resolution.
2. We run **`dcpromo`** and we create a new domain, a new tree and a new forest.
3. We take note of the password used for the administrator account.
4. We allow the setup to install DNS server.
5. We wait until the setup finishes and we boot the server.
6. We configure DNS server and we create a Reverse Lookup Zones for all our network addresses. We make them active directory integrated zones.
7. We define new hosts for each of the nodes with the option of creating the associated pointer (PTR) record.
8. We test DNS using **`nslookup`** from a command prompt.
9. We look for any error messages in the event viewer.

Promoting the other servers

To promote the rest of the servers:

1. We run **`dcpromo`** and we join to the domain created above, selecting **Additional domain controller for an existing domain**.
2. We use the password established in step 3 on Promoting the first server.
3. After the server boots, we install the DNS server.
4. We check if DNS is replicated correctly and we test using **`nslookup`**.
5. We look for any error messages in the event viewer.

Setting up a cluster user account

Before we go on and installing the cluster service, we create a cluster user account that will be required to bring the service up. This account should belong to the administrators group on each node. For security reasons we set the password setting to **User Cannot Change Password** and **Password Never Expires**.

Setting up external shared disks

When we install the SCSI/fibre adapter, we always use the same slot for all servers.

> **Attention:** While configuring shared disks, we have always only one server up at a time, to avoid corruption. To proceed, we shut down all servers, turn on the storage device, and turn on only one of the nodes.

On the DS4500 side, we prepare the LUNs that will be designated to our servers. A summary of the configuration is shown in Figure 4-19.

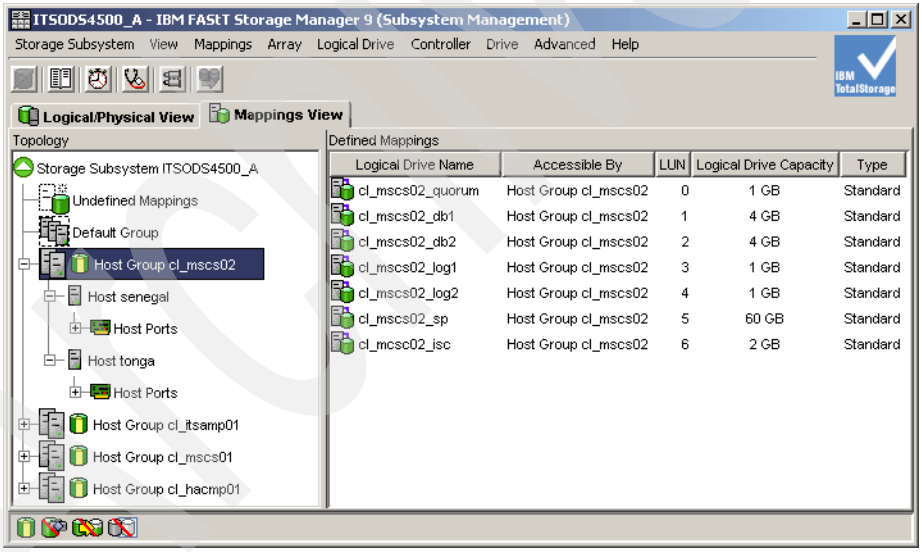

Figure 4-19 LUN configuration for our Windows 2003 MSCS

Archived

We install the necessary drivers according to the manufacturer's manual, so that Windows recognizes the storage disks. Device manager should look similar to Figure 4-20 on the items Disk drivers and SCSI and RAID controllers.

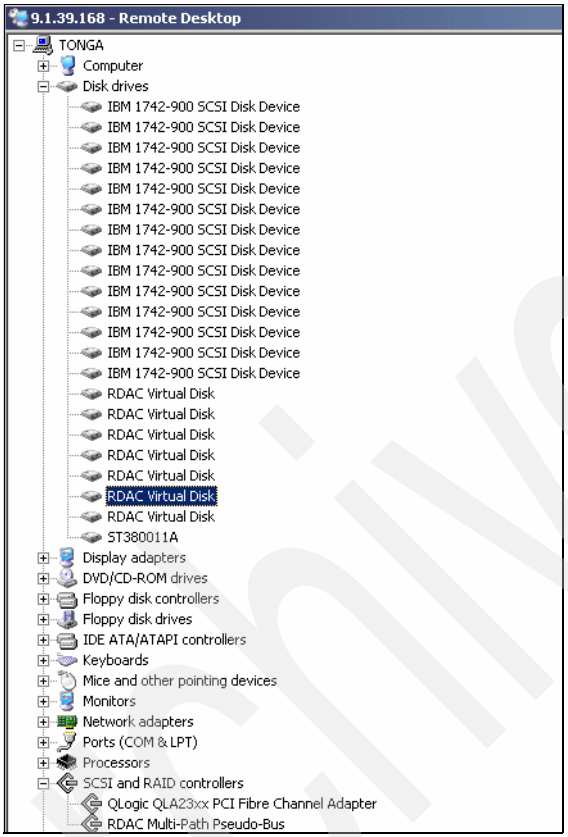

Figure 4-20 Device manager with disks and SCSI adapters

Archived

Configuring shared disks

To configure the shared disks:

1. We double click **Disk Management** and the Write Signature and Upgrade Disk Wizard (Figure 4-21) begins.

Figure 4-21 Disk initialization and conversion wizard

2. We select all disks for the **Write Signature** part in Figure 4-22.

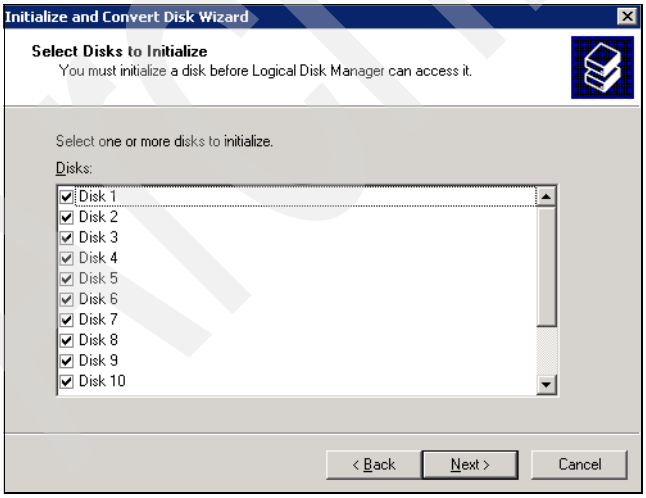

Figure 4-22 Select all drives for signature writing

Archived

3. We do not upgrade any of the disks to dynamic in Figure 4-23. In case we want to upgrade them, to reset the disk to basic, we should right-click the disk we want to change, and choose **Revert to Basic Disk**.

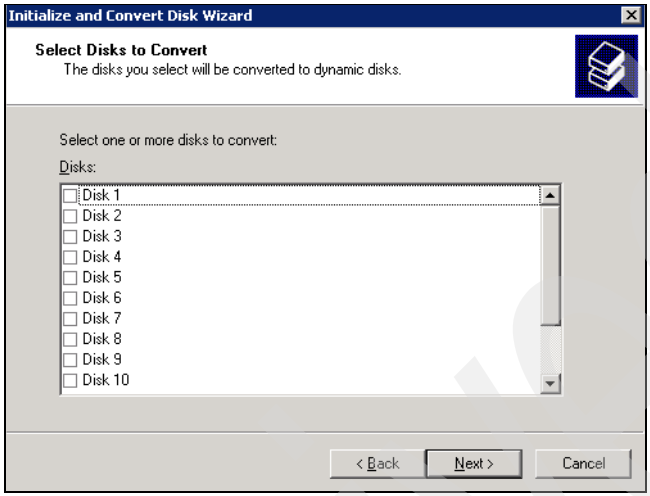

Figure 4-23 Do not upgrade any of the disks

4. We click **Finish** when the wizard completes as shown in Figure 4-24.

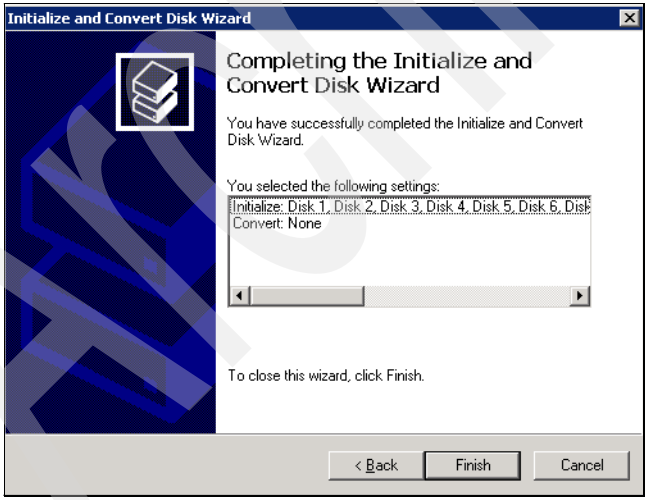

Figure 4-24 Successfull completion of the wizard

Archived

5. The disk manager will show now all disks online, but with unallocated partitions, as shown in Figure 4-25.

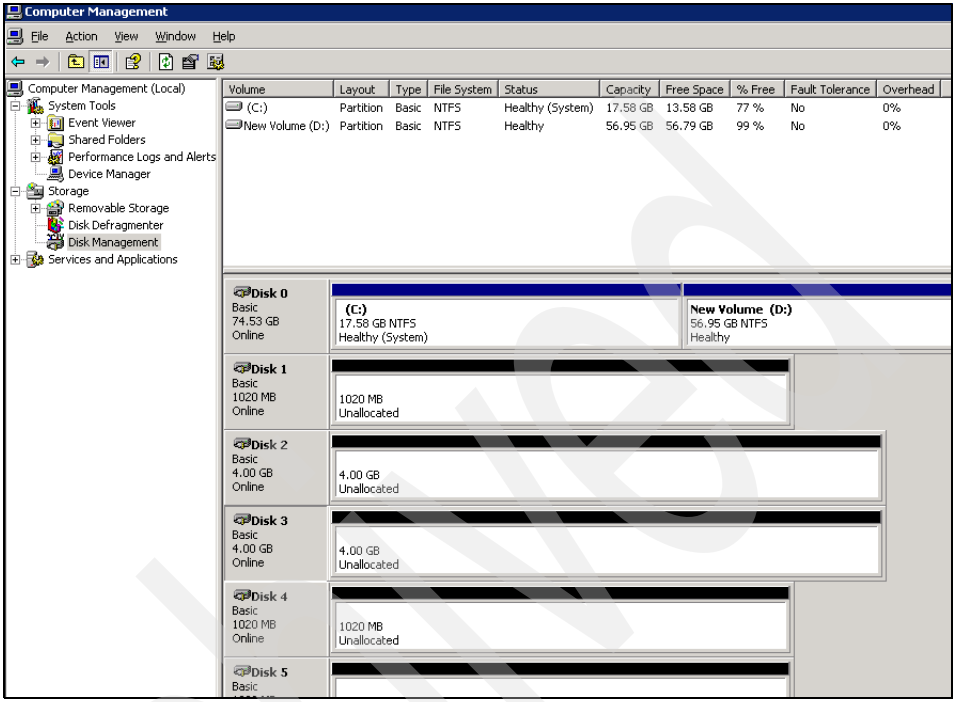

Figure 4-25 Disk manager after disk initialization

6. We right-click each of the unallocated disks and select **New Partition** in Figure 4-26.

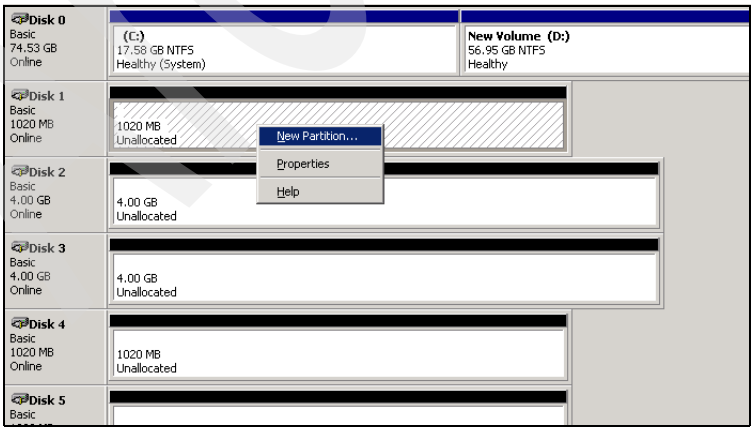

Figure 4-26 Create new partition

Archived

7. The New Partition wizard begins in Figure 4-27.

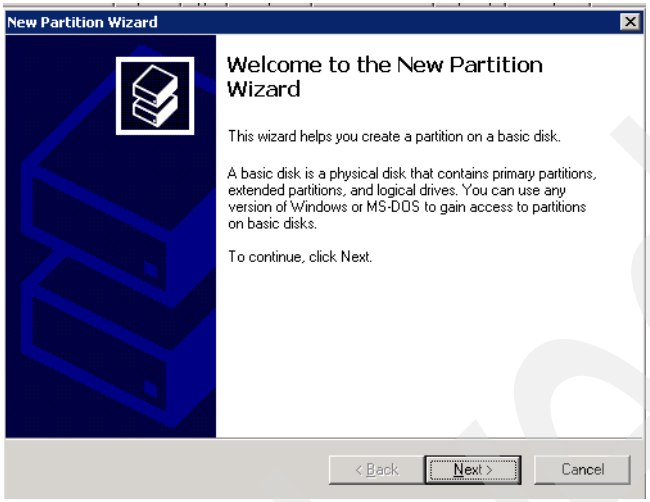

Figure 4-27 New partition wizard

8. We select **Primary Partition type** in Figure 4-28.

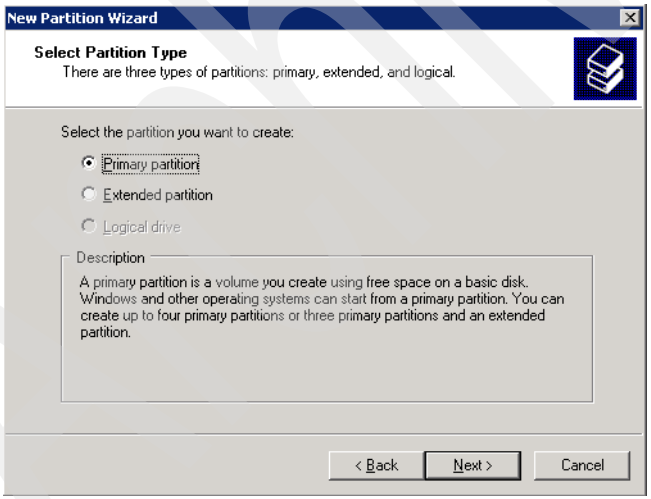

Figure 4-28 Select primary partition

Archived

9. We assign the partition size in Figure 4-29. We recommend only one partition per disk, assigning the maximum size.

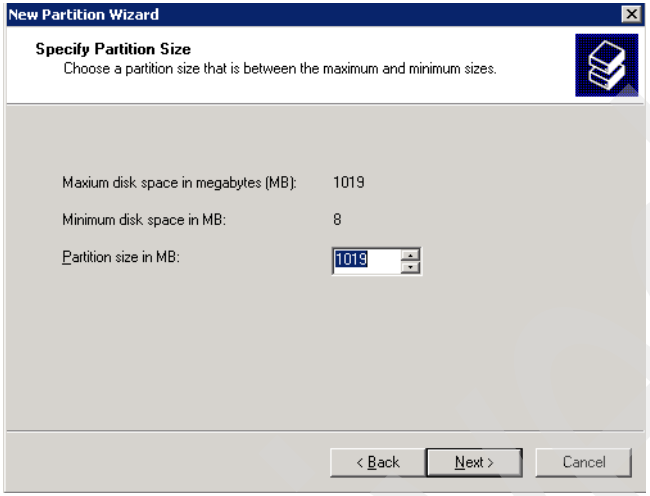

Figure 4-29 Select the size of the partition

10. We make sure to assign a drive mapping (Figure 4-30). This is crucial for the cluster to work. For the cluster quorum disk we recommend to use drive Q and the name Quorum, for clarity.

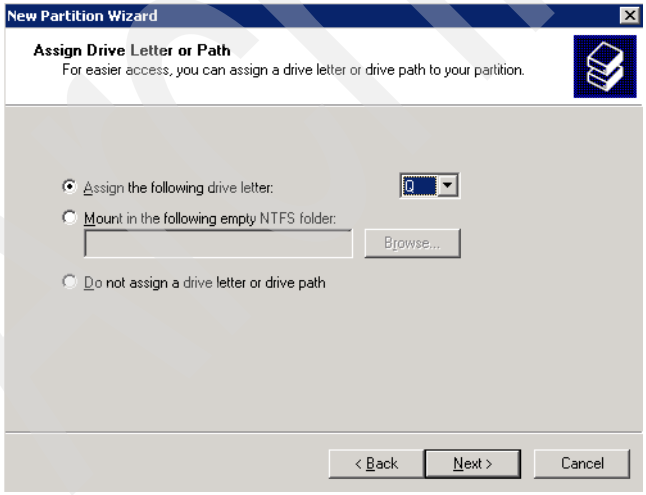

Figure 4-30 Drive mapping

11. We format the disk using NTFS in Figure 4-31, and we give a name that reflects the application we are setting up.

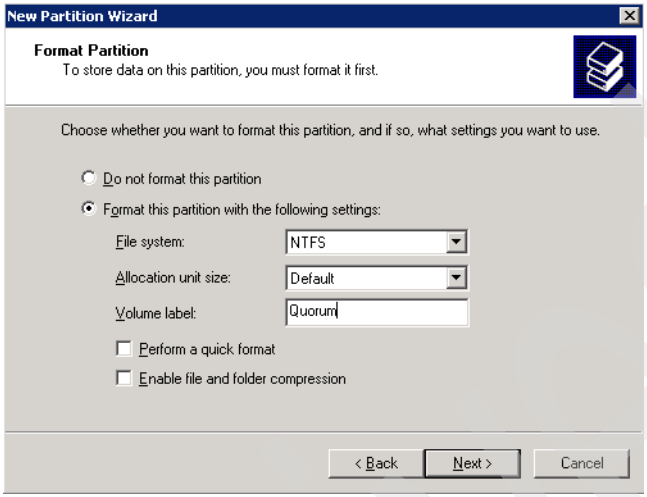

Figure 4-31 Format partition

12. The wizard shows the options we selected. To complete the wizard, we click **Finish** in Figure 4-32.

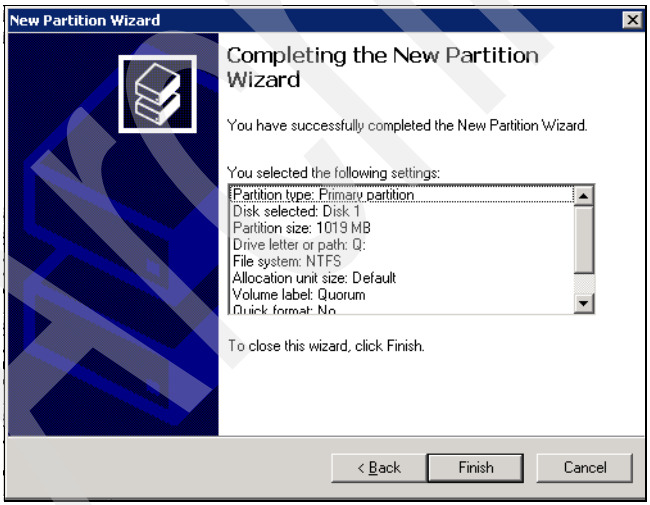

Figure 4-32 Completing the New Partition wizard

13. We verify that all shared disks are formatted as NTFS and are healthy and we write down the letters assigned to each partition in Figure 4-33.

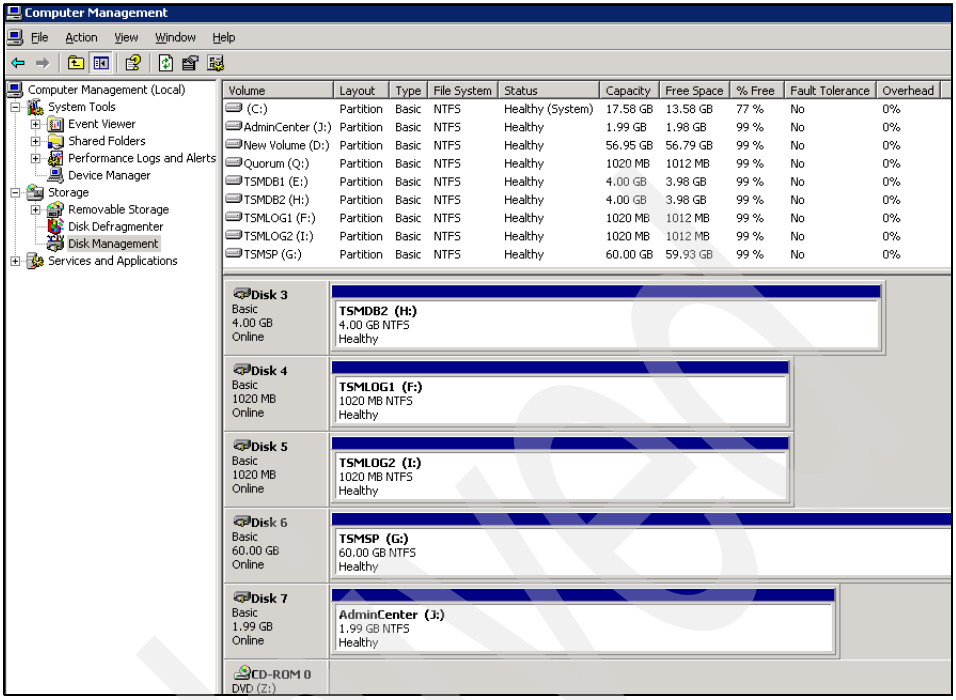

Figure 4-33 Disk configuration

14. We check disk access in Windows Explorer. We create any file on the drives and we also try to delete them.
15. We repeat steps 2 to 11 for every shared disk
16. We turn off the first node and turn on the second one. We check the partitions. If the letters are not set correctly, we change them to match the ones we set up on the first node. We also test write/delete file access from the other node.

Windows 2003 cluster setup

When we install Windows 2003 Enterprise or Datacenter editions, the Cluster Service is installed by default. So at this point no software installation is needed. We will use the Cluster Administrator to configure our environment.

Since the shared disks are still being recognized by both servers but with no sharing management, just one server should be turned on when we set up the first cluster node, to avoid disk corruption.

First node setup

To set up the first node:

1. We click **Start** → **All Programs** → **Administrative Tools** → **Cluster Administrator**. On the Open Connection to Cluster menu in Figure 4-34, we select **Create new cluster** and click **OK**

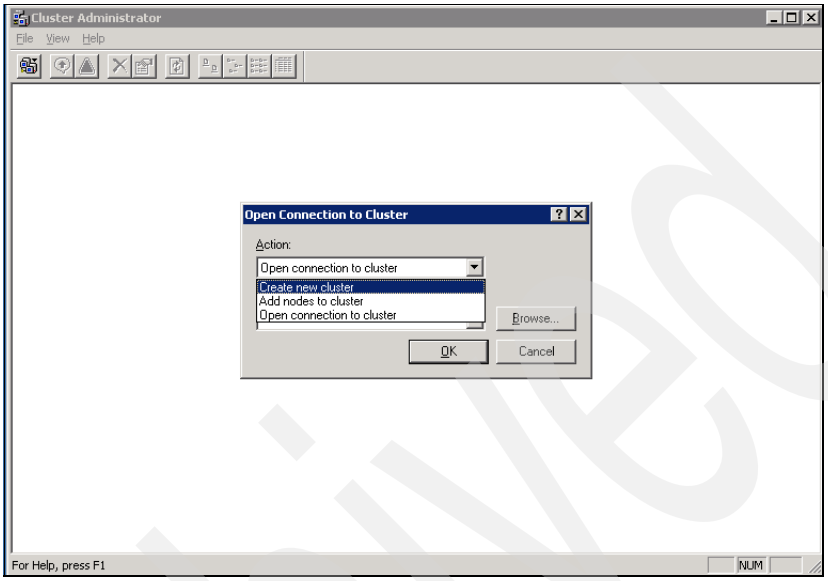

Figure 4-34 Open connection to cluster

2. The New Server Cluster Wizard starts. We check if we have all information necessary to configure the cluster (Figure 4-35). We click **Next**.

Figure 4-35 New Server Cluster wizard (prerequisites listed)

3. We type the unique NetBIOS clustername (up to 15 characters). Refer to Figure 4-36 for this information. The Domain is already typed based on the computer domain membership information when the server is set up.

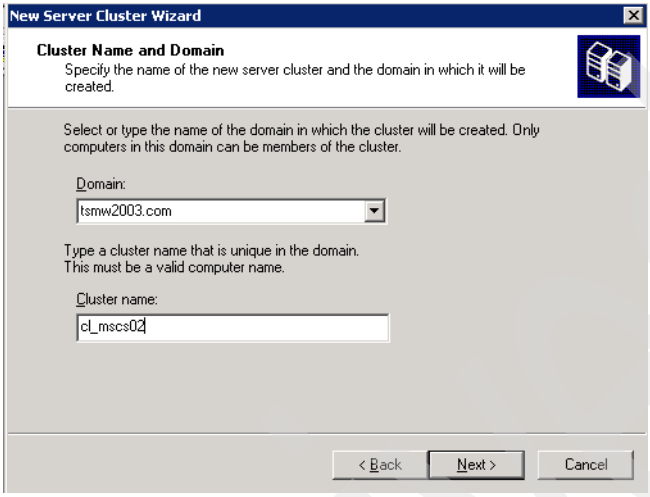

Figure 4-36 Clustername and domain

4. If we receive the message shown in Figure 4-37, we should analyze our application to see if the special characters will not affect it. In our case, Tivoli Storage Manager can handle the underscore character.

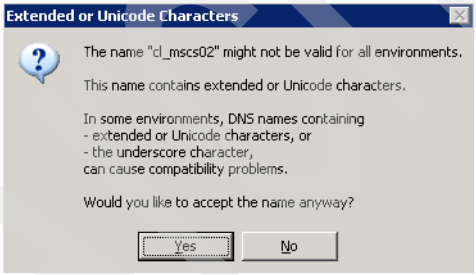

Figure 4-37 Warning message

5. Since in Windows 2003, it is possible to set up the cluster remotely, we confirm the name of the server that we are now setting the cluster up, as shown in Figure 4-38, and we click **Next**.

Archived

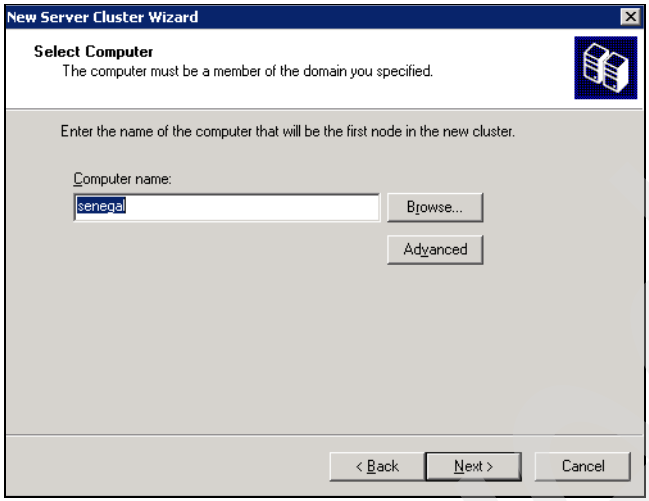

Figure 4-38 Select computer

6. The wizard starts analyzing the node looking for possible hardware or software problems. At the end, we review the warnings or error messages, clicking the **Details** button (Figure 4-39).

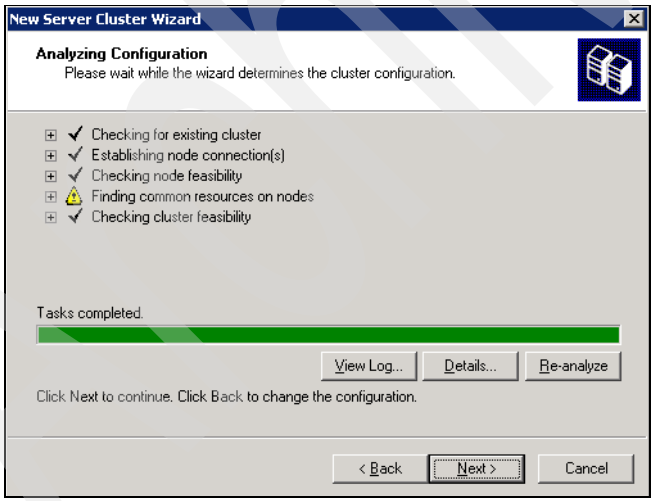

Figure 4-39 Review the messages

7. If there is anything to be corrected, we must run Re-analyze after corrections are made. As shown on the Task Details menu in Figure 4-40, this warning message is expected because the other node is down, as it should be.

We can continue our configuration. We click **Close** on the Task Details menu and **Next** on the Analyzing Configuration menu.

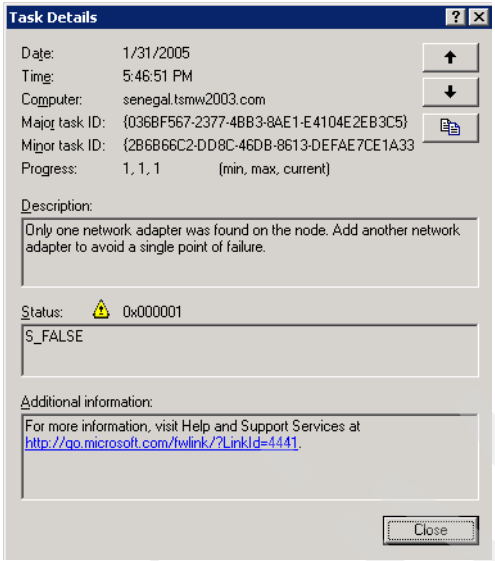

Figure 4-40 Warning message

8. We enter the cluster IP address. Refer to Figure 4-41.

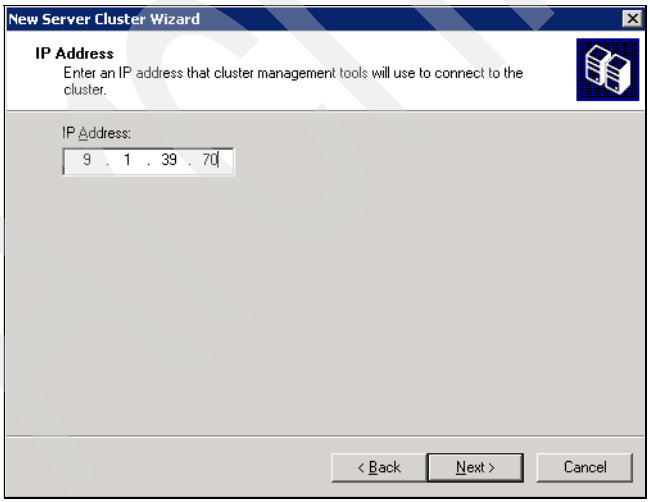

Figure 4-41 Cluster IP address

Archived

9. Next (Figure 4-42), we type the username and password of the cluster service account created in “Setting up a cluster user account” on page 51.

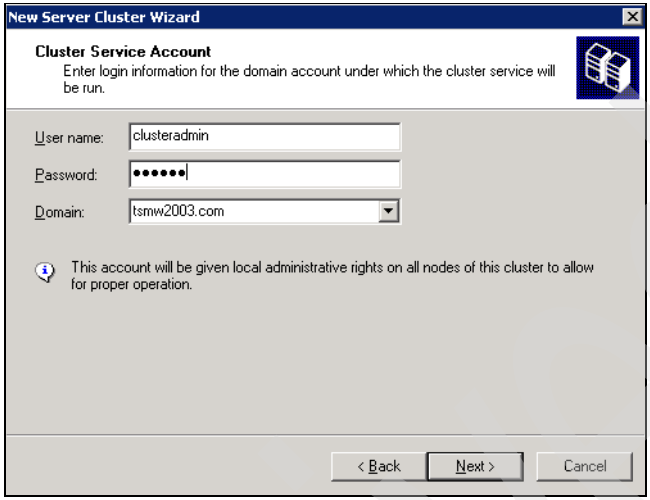

Figure 4-42 Specify username and password of the cluster service account

10. We review the information shown on the Proposed Cluster Configuration menu in Figure 4-43.

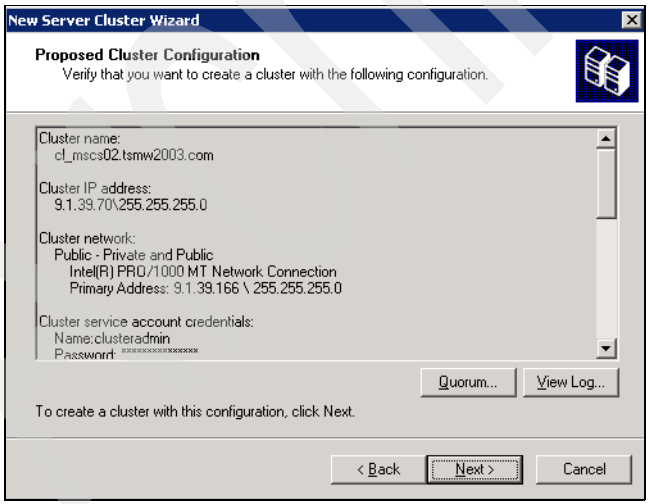

Figure 4-43 Summary menu

11. We click the **Quorum** button if it is necessary to change the disk that will be used for the Quorum (Figure 4-44). As default, the wizard automatically selects the drive that has the smallest partition larger than 50 MB. If everything is correct, we click **Next**.

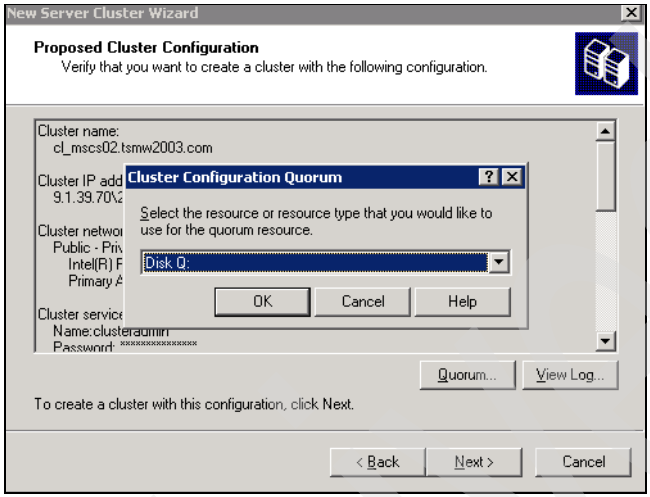

Figure 4-44 Selecting the quorum disk

12. We wait until the wizard finishes the creation of the cluster. We review any error or warning messages and we click **Next** (Figure 4-45).

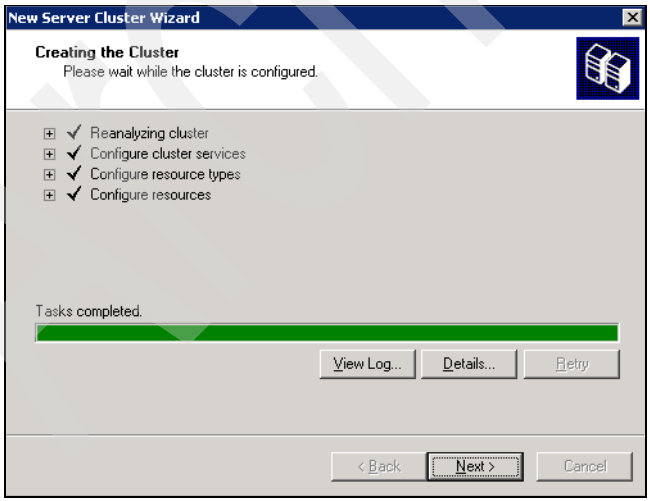

Figure 4-45 Cluster creation

Archived

13. We click **Finish** in Figure 4-46 to complete the wizard.

Figure 4-46 Wizard completed

14. We open the Cluster Administrator and we check the installation. We click **Start** → **Programs** →**Administrative Tools** →**Cluster Administrator** and expand all sessions. The result is shown in Figure 4-47. We check that the resources are all online.

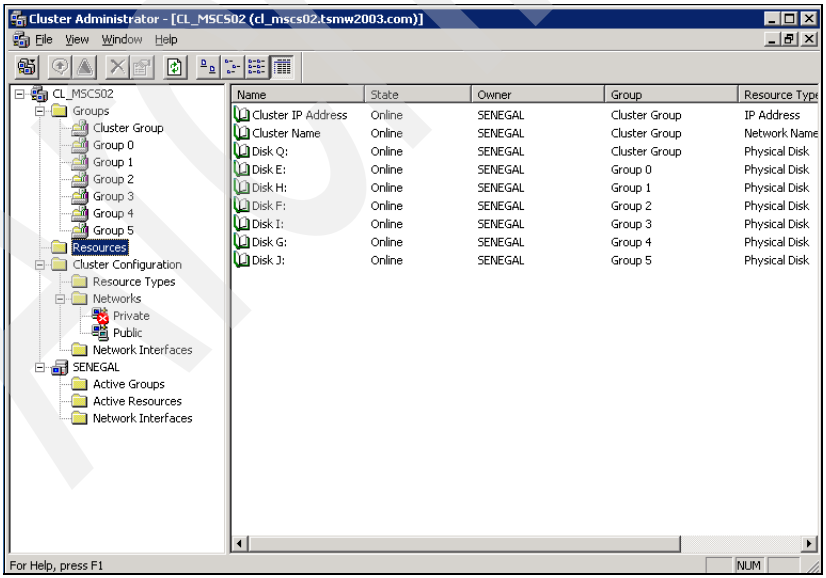

Figure 4-47 Cluster administrator

15. We leave this server turned on and bring the second node up to continue the setup.

Second node setup

The setup of the following nodes takes less time. The wizard configures network settings based on the first node configuration.

1. We open the Cluster Administrator (**Start** → **Programs** →**Administrative Tools** →**Cluster Administrator)**. We select **File** →**New** →**Node**.
2. We click Next on the Welcome to the Add Node Wizard menu
3. We type the computer name of the machine we are adding and we click **Add**. If there are more nodes, we can add them all here. We click **Next** (Figure 4-48).

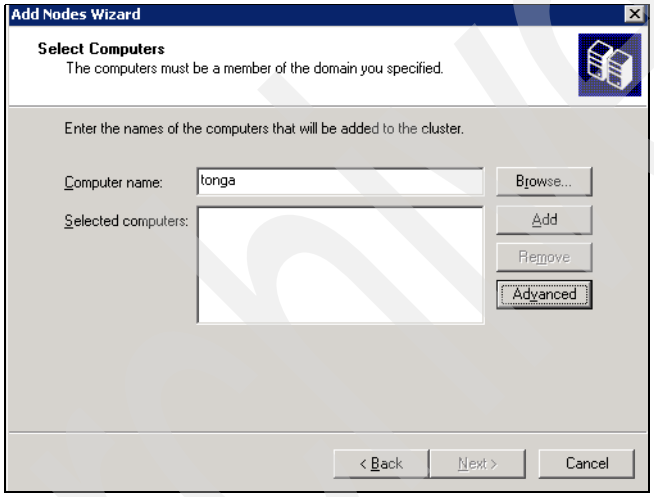

Figure 4-48 Add cluster nodes

Archived

4. The wizard starts checking the node. We check the messages and we correct the problems if needed (Figure 4-49).

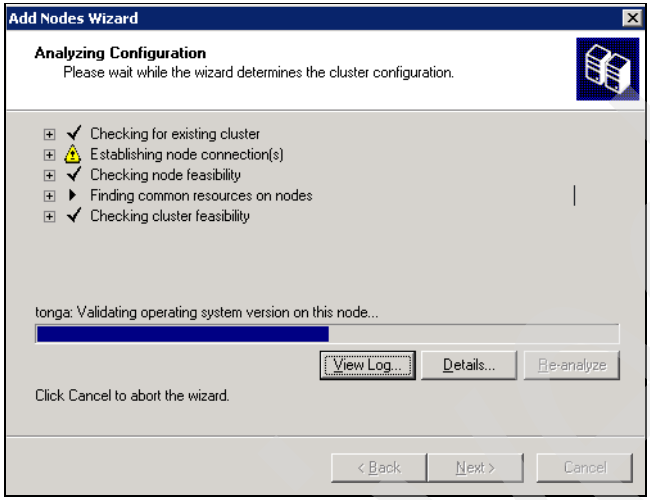

Figure 4-49 Node analysis

5. We type the password for the cluster service user account created in “Setting up a cluster user account” on page 51 (Figure 4-50).

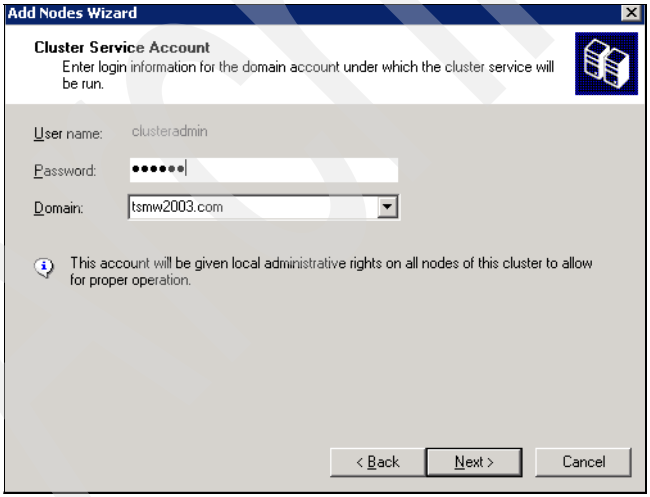

Figure 4-50 Specify the password

6. We review the summary information and we click **Next** (Figure 4-51).

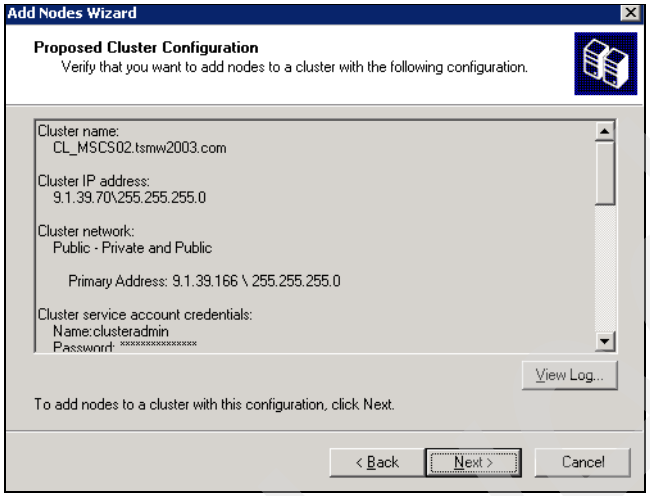

Figure 4-51 Summary information

7. We wait until the wizard finishes the analysis of the node. We review and correct any errors and we click **Next** (Figure 4-52).

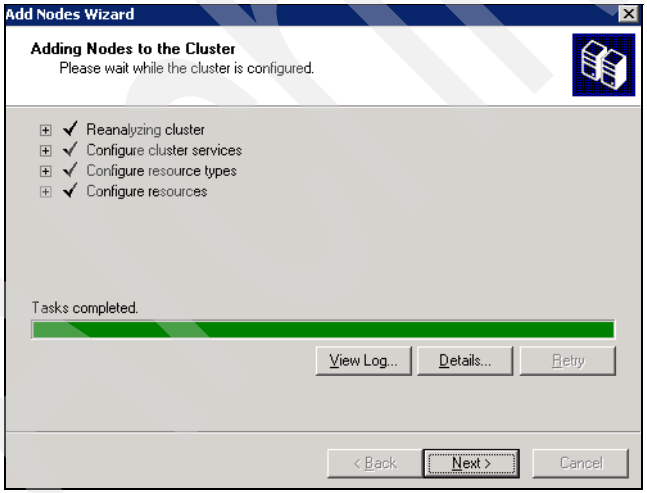

Figure 4-52 Node analysis

Archived

8. We click **Finish** to complete the setup (Figure 4-53).

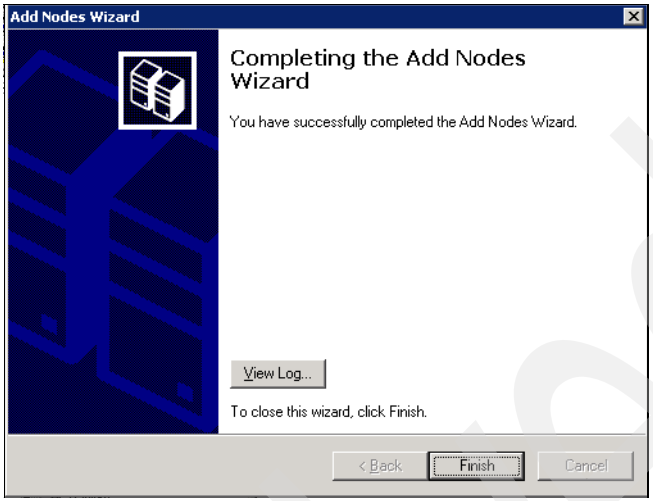

Figure 4-53 Setup complete

Configure the network roles of each adapter

The adapters can be configured for internal communications of the cluster (private network), for client access only (public network) or for all communications (mixed network). For a two-node cluster as the one we have in this lab, the private adapter is used for internal cluster communications only (heartbeat) and the public adapter is used for all communications.

To set up these roles, we follow these steps:

1. We open the Cluster Administrator. In the left panel, we click **Cluster Configuration** → **Network.** We right-click **Private** and we choose **Properties** as shown in Figure 4-54.

Archived

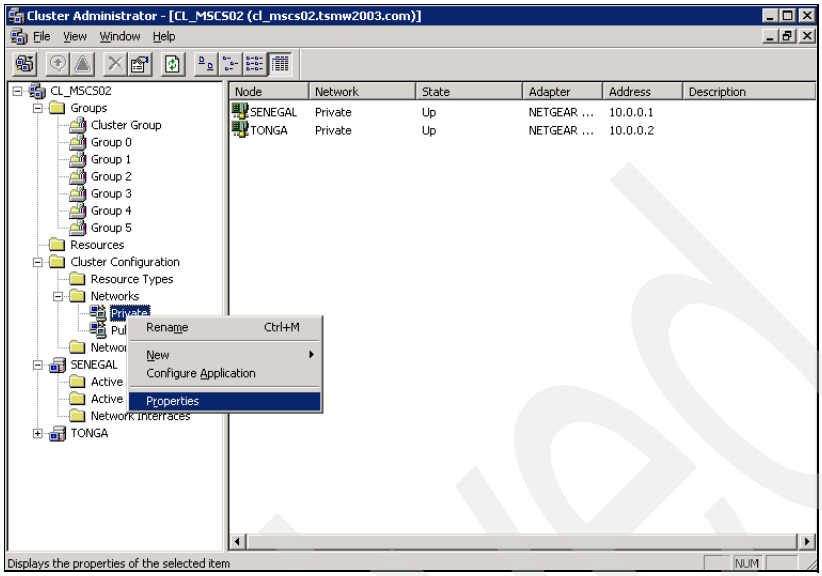

Figure 4-54 Private network properties

2. We choose **Enable this network for cluster use** and **Internal cluster communications only (private network)** and we click **OK** (Figure 4-55).

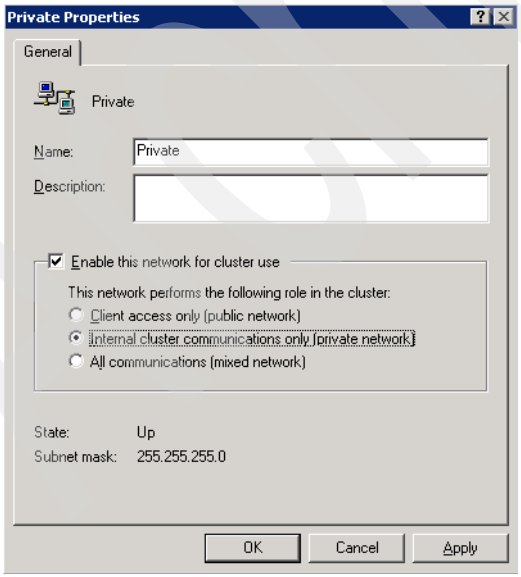

Figure 4-55 Configuring the heartbeat

Archived

3. We right-click **Public** and we choose **Properties** (Figure 4-56).

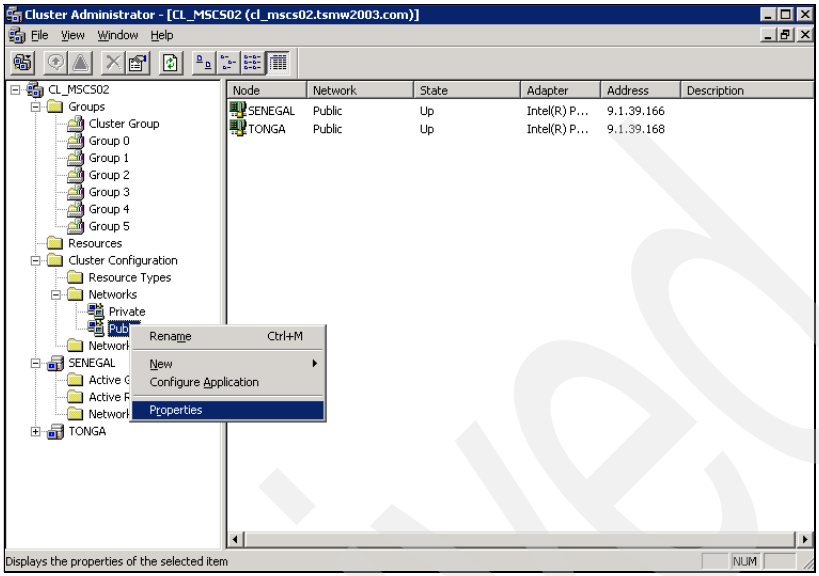

Figure 4-56 Public network properties

4. We choose **Enable this network for cluster use** and **All communications (mixed network)** and we click **OK** (Figure 4-57).

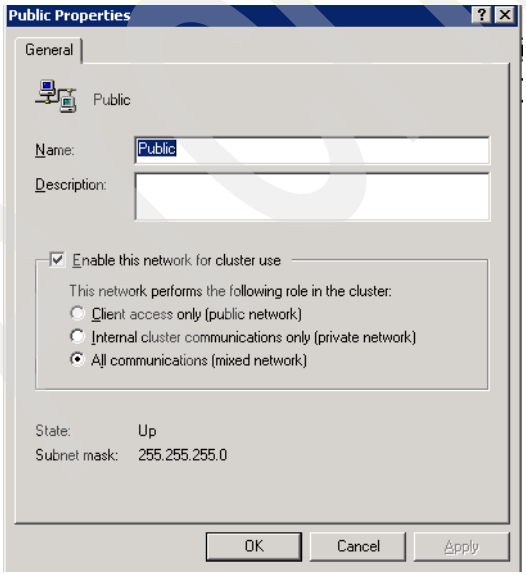

Figure 4-57 Configuring the public network

5. We set the priority of each network for the communication between the nodes. We right-click the cluster name and choose **Properties** (Figure 4-58).

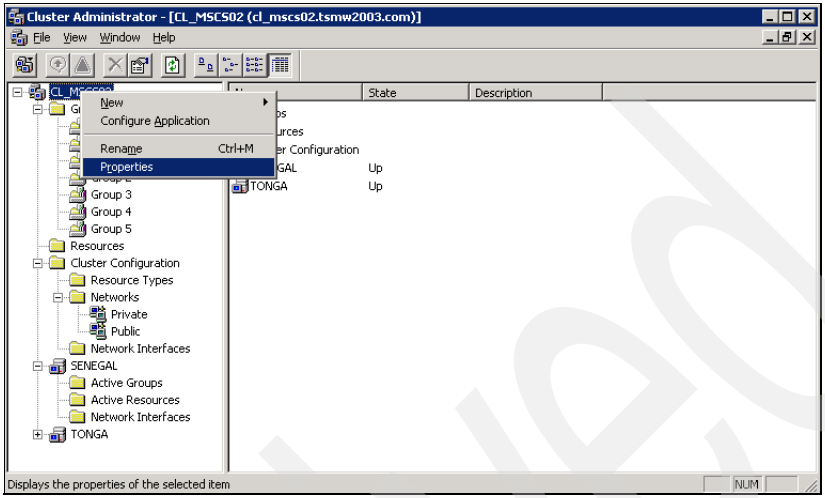

Figure 4-58 Cluster properties

6. We choose the **Network Priority** tab and we use the **Move Up** or **Move Down** buttons so that the Private network comes at the top as shown in Figure 4-59 and we click **OK**.

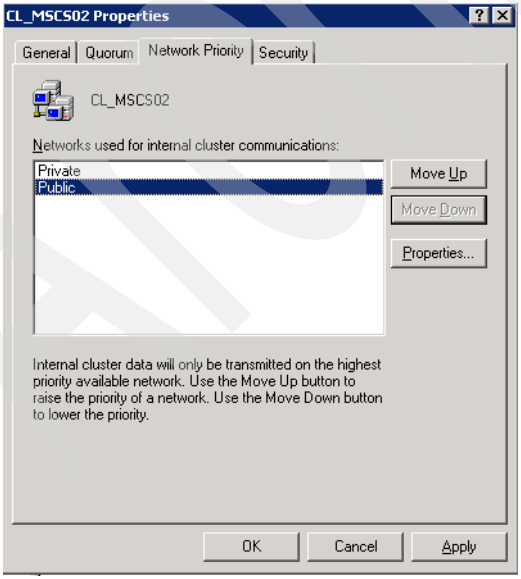

Figure 4-59 Network priority

Archived

Windows 2003 cluster configuration

When the installation is complete the cluster looks like Figure 4-60, with one group resource for each disk. We may change this distribution, creating new groups with more than one disk resource, to best fit our environment.

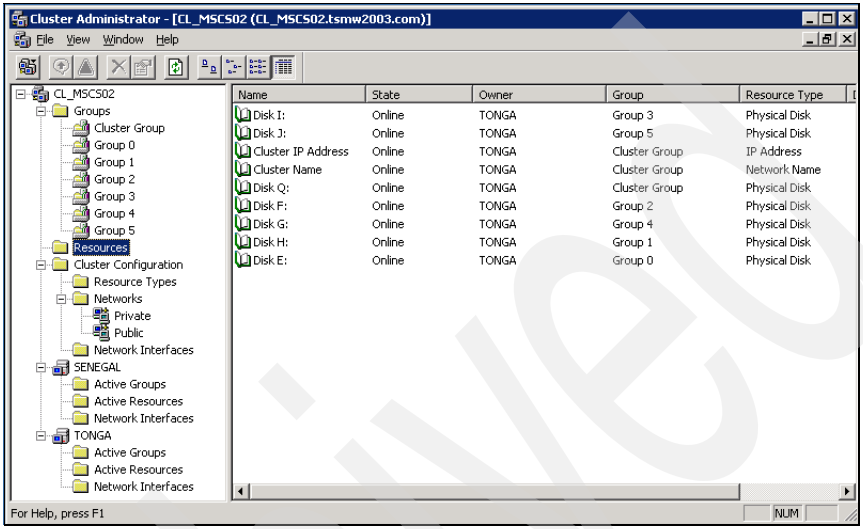

Figure 4-60 Cluster Administrator after end of installation

The next step is to group disks together for each application. Cluster Group should have the cluster name, ip and quorum disk, and we create, for the purpose of this book, two other groups: Tivoli Storage Manager Group with disks E through I and Tivoli Storage Manager Admin Center with disk J.

1. We use the Change Group option as shown in Figure 4-61.

Archived

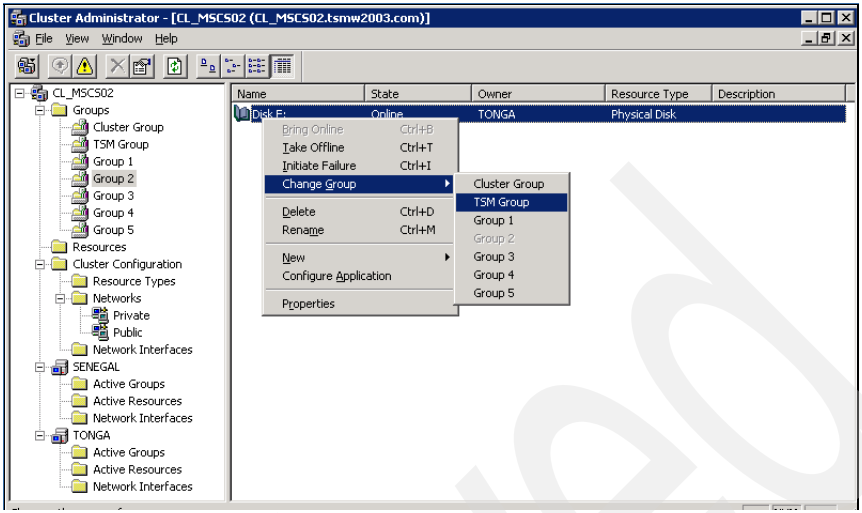

Figure 4-61 Moving resources

2. We reply **Yes** twice to confirm the change.
3. We delete the groups that become empty, with no resource. The result is shown in Figure 4-62.

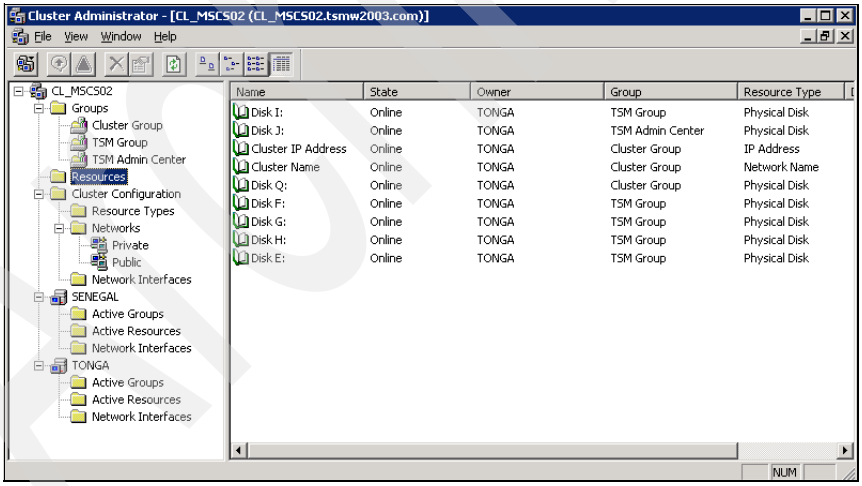

Figure 4-62 Final configuration

Archived

Tests

To test the cluster functionality, we use the Cluster Administrator and we perform the following tasks:

- Move groups from one server to another. Verify that resources failover and are brought online on the other node.
- Move all resources to one node and stop the Cluster service. Verify that all resources failover and come online on the other node.
- Move all resources to one node and shut it down. Verify that all resources failover and come online on the other node.
- Move all resources to one node and remove the public network cable from that node. Verify that the groups will failover and come online on the other node.

4.5 Troubleshooting

The cluster log is a very useful troubleshooting tool. It is enabled by default and its output is printed as a log file in %SystemRoot%Cluster.

DNS plays an important role in the cluster functionality. Many of the problems can be avoided if we make sure that DNS is well configured. Fail to create reverse lookup zones has been one of the main reasons for the cluster setup failure.

Archived

5

Microsoft Cluster Server and the IBM Tivoli Storage Manager Server

This chapter discusses how we set up Tivoli Storage Manager server to work in *Microsoft Cluster Services* (MSCS) environments for high availability.

We use our two Windows MSCS environments described in Chapter 4:

- Windows 2000 MSCS formed by two servers: POLONIUM and RADON
- Windows 2003 MSCS formed by two servers: SENEGAL and TONGA.

© Copyright IBM Corp. 2005. All rights reserved.

5.1 Overview

In an MSCS environment, independent servers are configured to work together in order to enhance the availability of applications using shared disk subsystems.

Tivoli Storage Manager server is an application with support for MSCS environments. Clients can connect to the Tivoli Storage Manager server using a virtual server name.

To run properly, Tivoli Storage Manager server needs to be installed and configured in a special way, as a shared application in the MSCS.

This chapter covers all the tasks we follow in our lab environment to achieve this goal.

5.2 Planning and design

When planning our Tivoli Storage Manager server cluster environment, we should:

- Choose the cluster configuration that best fits our high availability needs.
- Identify disk resources to be used by Tivoli Storage Manager. We should not partition a disk and use it with other applications that might reside in the same server, so that a problem in any of the applications will not affect the others. We have to remember that the quorum disk should also reside on a separate disk, with at least 500 MB. We should not use the quorum disk for anything but the cluster management.
- Have enough IP addresses. Each node in the cluster uses two IP addresses (one for the heartbeat communication between the nodes and another one on the public network). The cluster virtual server uses a different IP address and Tivoli Storage Manager server also uses one (minimum 6 for a two-server cluster).
- Create one separate cluster resource for each Tivoli Storage Manager instance, with the corresponding disk resources.
- Check disk space on each node for the installation of Tivoli Storage Manager server. We highly recommend that the same drive letter and path be used on each machine.
- Use an additional shared SCSI bus so that Tivoli Storage Manager can provide tape drive failover support.

Archived

Note: Refer to “Appendix A” of the *IBM Tivoli Storage Manager for Windows: Administrator’s Guide* for instructions on how to manage SCSI tape failover.

For additional planning and design information, refer to *Tivoli Storage Manager for Windows Installation Guide* and *Tivoli Storage Manager Administrator’s Guide.*

Notes:

- Service Pack 3 is required for backup and restore of SAN File Systems.
- Windows 2000 hot fix 843198 is required to perform open file backup together with *Windows Encrypting File System* (EFS) files.

5.3 Installing Tivoli Storage Manager Server on a MSCS

In order to implement Tivoli Storage Manager server to work correctly on a Windows 2000 MSCS or Windows 2003 MSCS environment as a virtual server in the cluster, it is necessary to perform these tasks:

1. Installation of Tivoli Storage Manager software components on each node of the MSCS, on local disk.
2. If necessary, installation of the correct tape drive and tape medium changer device drivers on each node of the MSCS.
3. Installation of the new administrative Web interface, the Administration Center console, to manage the Tivoli Storage Manager server.
4. Configuration of Tivoli Storage Manager server as a clustered application, locating its database, recovery log and disk storage pool volumes on shared resources.
5. Testing the Tivoli Storage Manager server.

Some of these tasks are exactly the same for Windows 2000 or Windows 2003. For this reason, and to avoid duplicating the information, in this section we describe these common tasks. The specifics of each environment are described in sections “Tivoli Storage Manager server and Windows 2000” on page 118 and “Tivoli Storage Manager Server and Windows 2003” on page 179, also in this chapter.

Archived

5.3.1 Installation of Tivoli Storage Manager server

The installation of Tivoli Storage Manager server on an MSCS environment follows the same rules as in any other single Windows server. It is necessary to install the software on local disk in each node belonging to the same cluster.

In this section we describe this installation process. The same tasks apply to both Windows 2000 environments as well as Windows 2003.

We use the same disk drive letter and installation path on each node:

```
c:\Program Files\Tivoli\tsm\server
```

To install the Tivoli Storage Manager server component, we follow these steps:

1. On the first node of each MSCS, we run **setup.exe** from the Tivoli Storage Manager CD. The following panel displays (Figure 5-1).

Figure 5-1 IBM Tivoli Storage Manager InstallShield wizard

2. We click **Next**.

Archived

3. The language menu displays. The installation wizard detects the OS language and defaults to it (Figure 5-2).

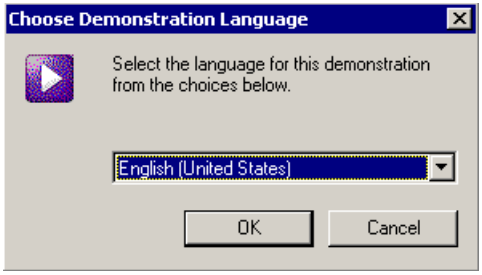

Figure 5-2 Language select

4. We select the appropriate language and click **OK**.
5. Next, the Tivoli Storage Manager Server installation menu displays (Figure 5-3).

Figure 5-3 Main menu

6. We select **Install Products**.

Archived

7. We are presented with the four Tivoli Storage Manager packages as shown in Figure 5-4.

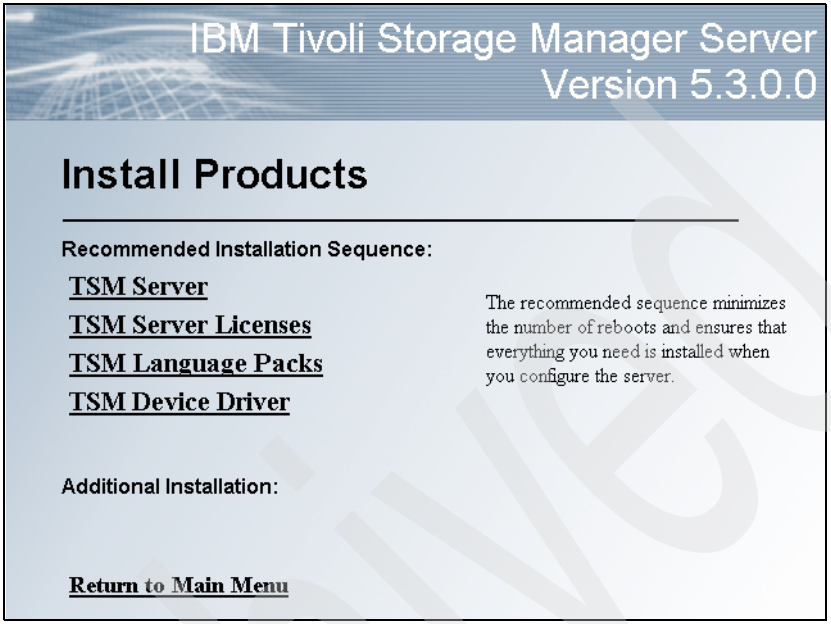

Figure 5-4 Install Products menu

We recommend to follow the installation sequence below:

a. Install Tivoli Storage Manager Server package first.

b. Install Tivoli Storage Manager Licenses package.

c. If needed, install the Tivoli Storage Manager Language Package (Optional).

d. Finally, install the Tivoli Storage Manager Device Driver if the devices need to be managed by this driver.

   We do not need Tivoli Storage Manager device driver for IBM Tape Libraries because they use their own IBM Windows drivers. However, the installation of Tivoli Storage Manager device driver is recommended because with the device information menu of the management console, we can display the device names used by Tivoli Storage Manager for the medium changer and tape drives. We only have to be sure that, after the installation, Tivoli Storage Manager device driver is not started at boot time if we do not need it to manage the tape drives.

In Figure 5-4 we first select the TSM Server package as recommended.

Archived

8. The installation wizard starts and the following menu displays (Figure 5-5).

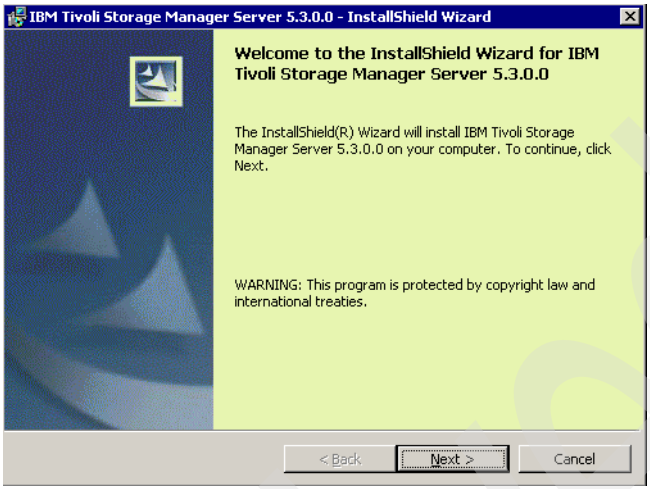

Figure 5-5 Installation wizard

9. We select **Next** to start the installation.

10. We accept the license agreement and click **Next** (Figure 5-6).

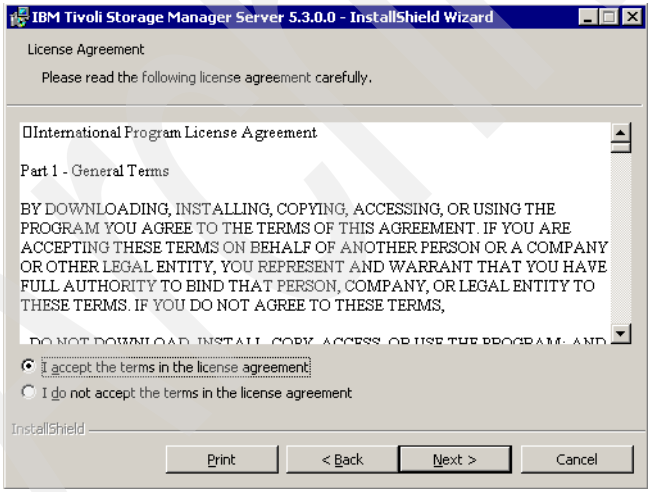

Figure 5-6 Licence agreement

Archived

11. We enter our customer information data now and click **Next** (Figure 5-7).

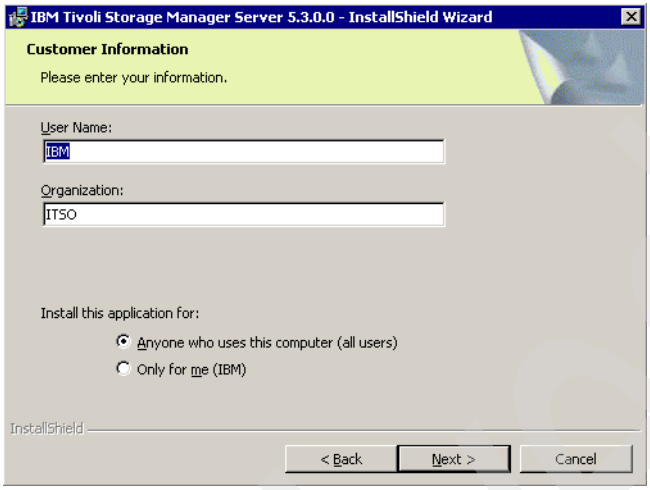

Figure 5-7 Customer information

12. We choose **Complete** installation and click **Next** (Figure 5-8).

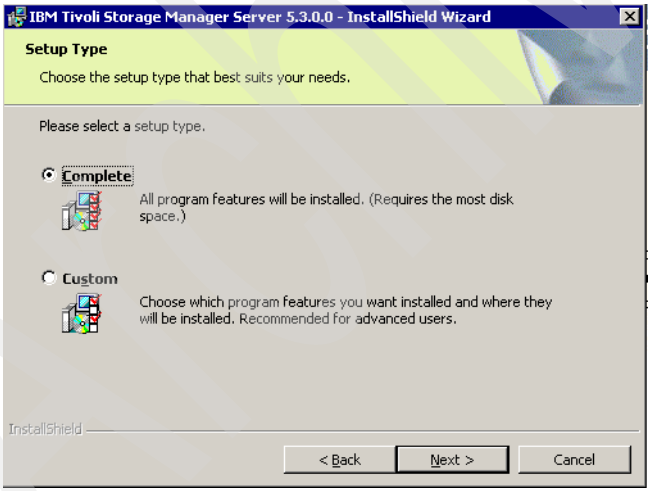

Figure 5-8 Setup type

Archived

13. The installation of the product begins (Figure 5-9).

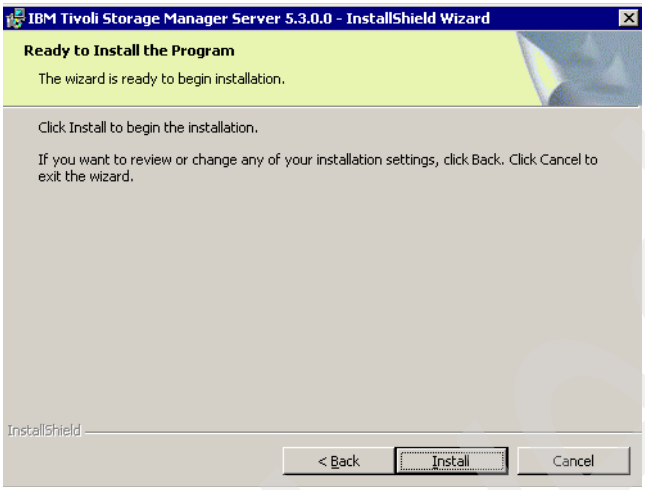

Figure 5-9 Beginning of installation

14. We click **Install** to start the installation.

15. The progress installation bar displays next (Figure 5-10).

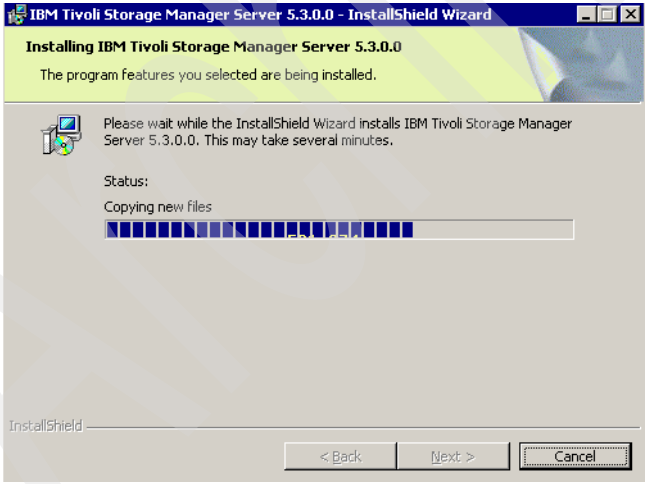

Figure 5-10 Progress bar

16. When the installation is completed, the successful message in Figure 5-11 displays. We click **Finish**.

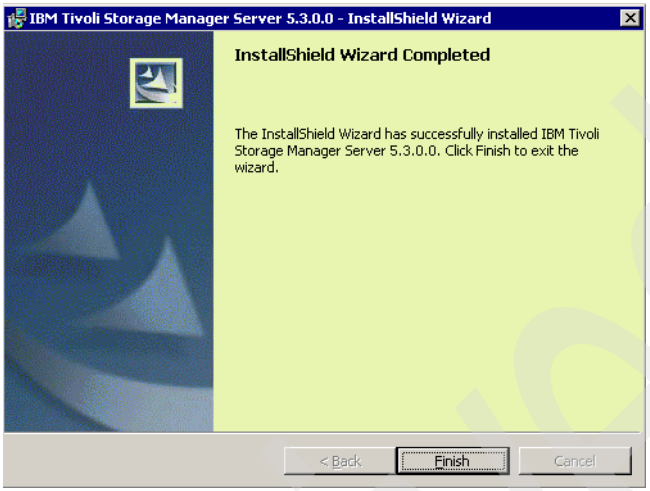

Figure 5-11 Successful installation

The Tivoli Storage Manager server is installed.

Note: A warning menu displays after the installation prompting to restart the server as shown in Figure 5-12. As we will install the remaining Tivoli Storage Manager packages, we do not need to restart the server at this point. We can do this after the installation of all the packages.

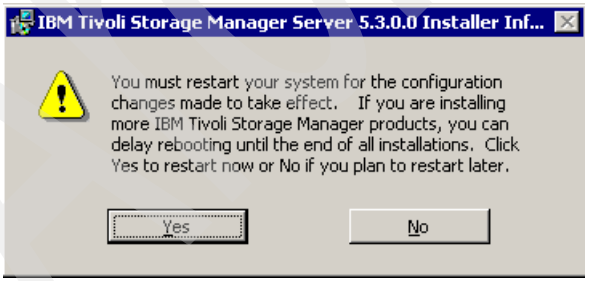

Figure 5-12 Reboot message

5.3.2 Installation of Tivoli Storage Manager licenses

In order to install the license package, in the main installation menu shown in Figure 5-13, select **TSM Server Licenses**.

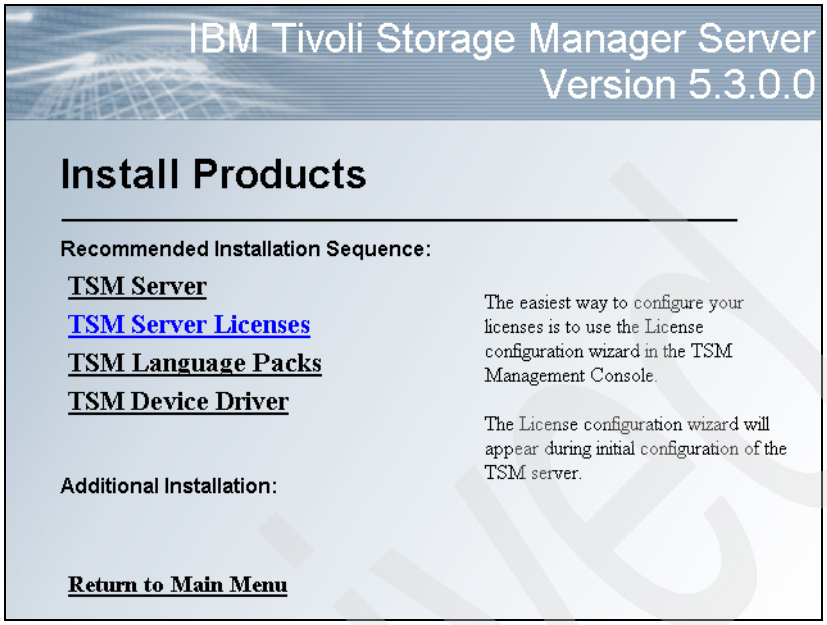

Figure 5-13 Install Products menu

The following sequence of menus displays:

1. The first panel is the Welcome Installation Wizard menu (Figure 5-14).

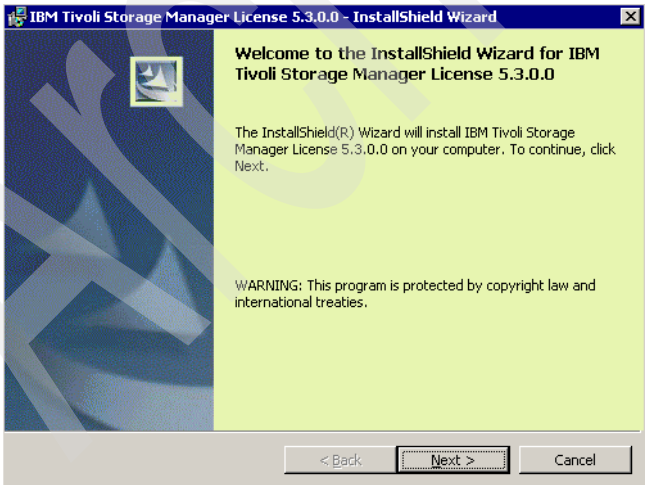

Figure 5-14 License installation

2. We click **Next**.

3. We fill in the User Name and Organization fields as shown in Figure 5-7 on page 84.
4. We select to run the **Complete** installation as shown in Figure 5-8 on page 84.
5. And finally the installation menu displays (Figure 5-15).

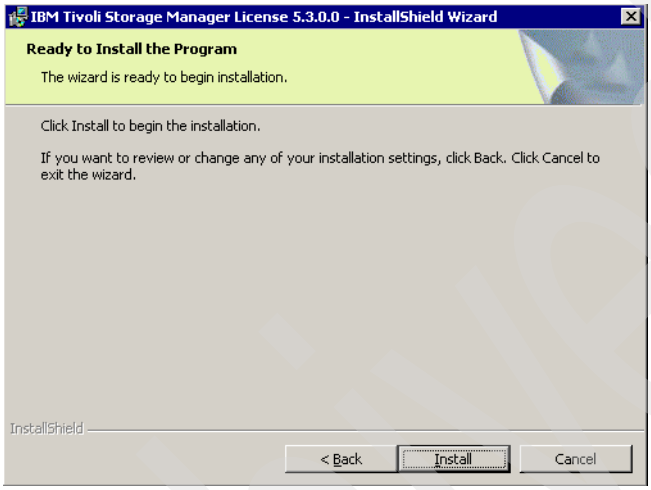

Figure 5-15 Ready to install the licenses

6. We click **Install**.
7. When the installation ends, we receive this informational menu (Figure 5-16).

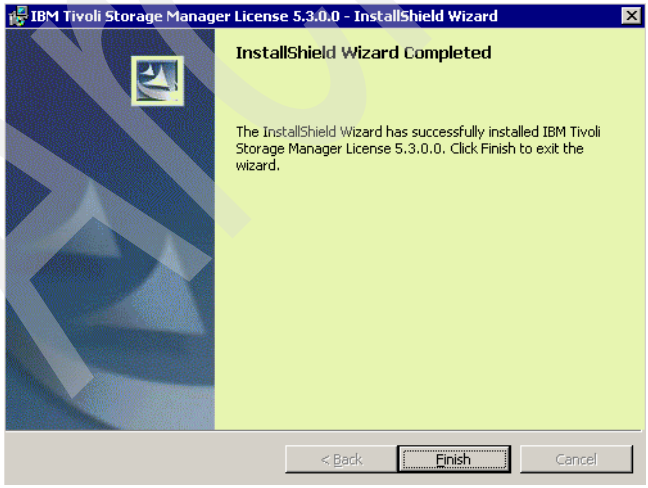

Figure 5-16 Installation completed

Archived

8. We click **Finish**. The Tivoli Storage Manager license package is installed.

5.3.3 Installation of Tivoli Storage Manager device driver

The installation of Tivoli Storage Manager device driver *is not a must*. Check Tivoli Storage Manager documentation for devices that need this driver. If the devices will be handled by OS drivers there is no need to install it.

However it is a recommended option because it helps to see the device names from the Tivoli Storage Manager and from the Windows OS perspectives when using the management console. We do not need to start the Tivoli Storage Manager device driver to get this information, just install it and disable it.

To install the driver, we follow these steps:

1. We go into the main installation menu (Figure 5-17).

Figure 5-17 Install Products menu

2. We select **TSM Device Driver**.
3. We click **Next** on the Welcome Installation Wizard menu (Figure 5-18).

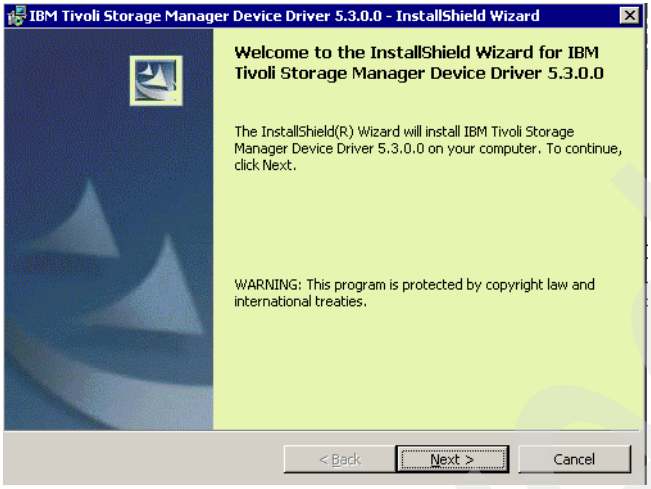

Figure 5-18 Welcome to installation wizard

4. We type the User Name and Organization fields as shown in Figure 5-7 on page 84.
5. We select to run the **Complete** installation as shown in Figure 5-8 on page 84.
6. The wizard is ready to start the installation. We click **Install** (Figure 5-19).

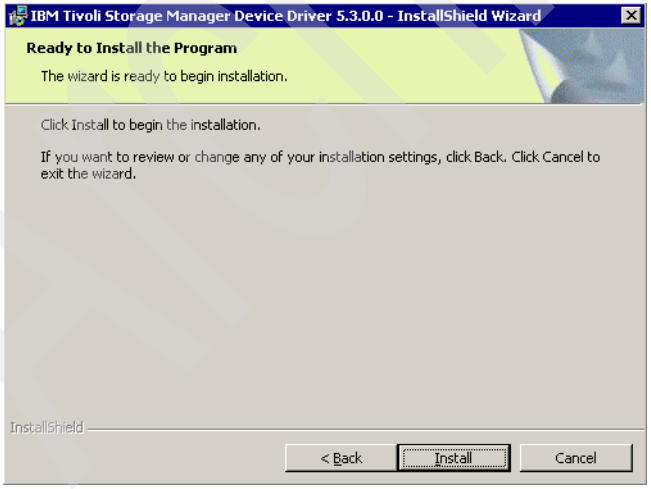

Figure 5-19 Ready to install

7. When the installation completes, we can see the same menu as shown in Figure 5-11 on page 86. We click **Finish**.
8. Finally, the installation wizard prompts to restart this server. This time, we select **Yes** (Figure 5-20).

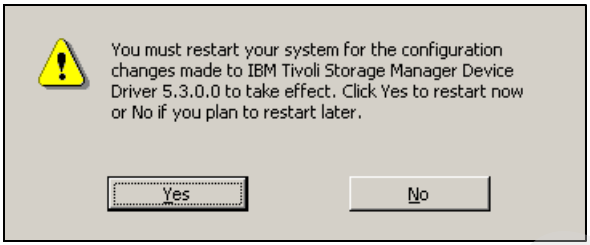

Figure 5-20 Restart the server

9. We *must* follow the same process on the second node of each MSCS, installing the same packages and using the same local disk drive path used on the first node. After the installation completes on this second node, we restart it.

Important: Remember that when we reboot a server that hosts cluster resources, they will automatically be moved to the other node. We need to be sure not to reboot both servers at the same time. We wait until the resources are all online on the other node.

We follow all these tasks in our Windows 2000 MSCS (nodes POLONIUM and RADON), and also in our Windows 2003 MSCS (nodes SENEGAL and TONGA).

Refer to “Tivoli Storage Manager server and Windows 2000” on page 118 and “Tivoli Storage Manager Server and Windows 2003” on page 179 for the configuration tasks on each of these environments.

Archived

5.3.4 Installation of the Administration Center

Since IBM Tivoli Storage Manager V5.3.0, the administrative Web interface has been replaced with the Administration Center. This is a Web-based interface to centrally configure and manage any Tivoli Storage Manager V5.3.0 server.

IBM Tivoli Storage Manager Administration Center consists of two components:

- The Integrated Solutions Console (ISC)
- The Administration Center

ISC allows you to install components provided by multiple IBM applications, and access them from a single interface. It is a requirement to install the Administration Center.

Installing the ISC and Administration Center for clustering

The Administration Center is not a clustered application and is not officially supported as a clustered application in Windows environments. However, in our lab we follow a procedure that allows us to install and configure it as a clustered application.

We first install both components in the first node of each MSCS, then we move the resources and follow a special method to install the components in the second node.

In this section we describe the common tasks for any MSCS (Windows 2000 or Windows 2003). The specifics for each environment are described in “Configuring ISC for clustering on Windows 2000” on page 167 and “Configuring ISC for clustering on Windows 2003” on page 231.

Installation of ISC in the first node

These are the tasks we follow to install the ISC in the first node of each MSCS:

1. We check the node that hosts the shared disk where we want to install the ISC.
2. We run **`setupISC.exe`** from the CD. The welcome installation menu displays (Figure 5-21).

Archived

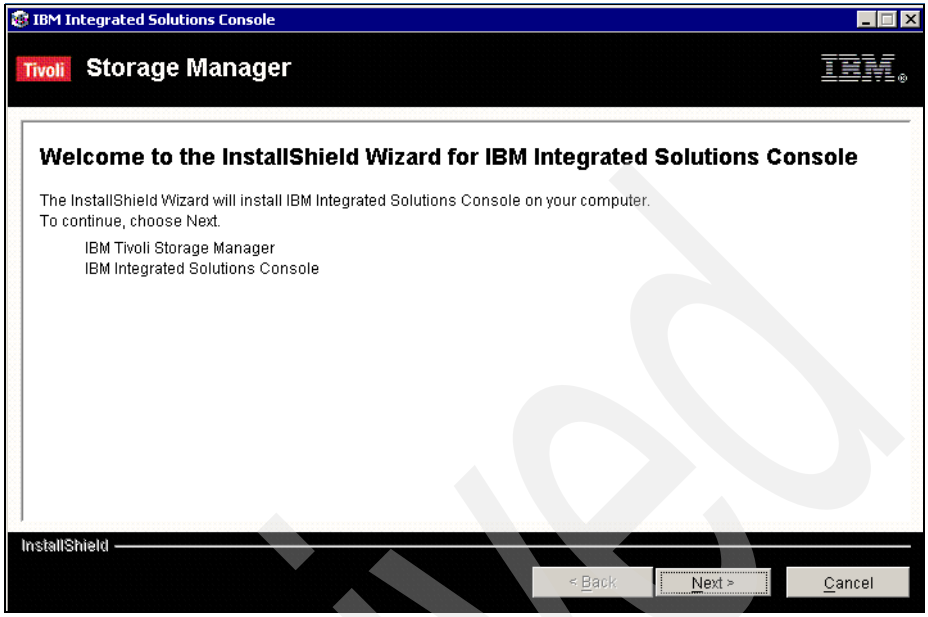

Figure 5-21 InstallShield wizard for IBM Integrated Solutions Console

3. In Figure 5-21 we click **Next** and the menu in Figure 5-22 displays.

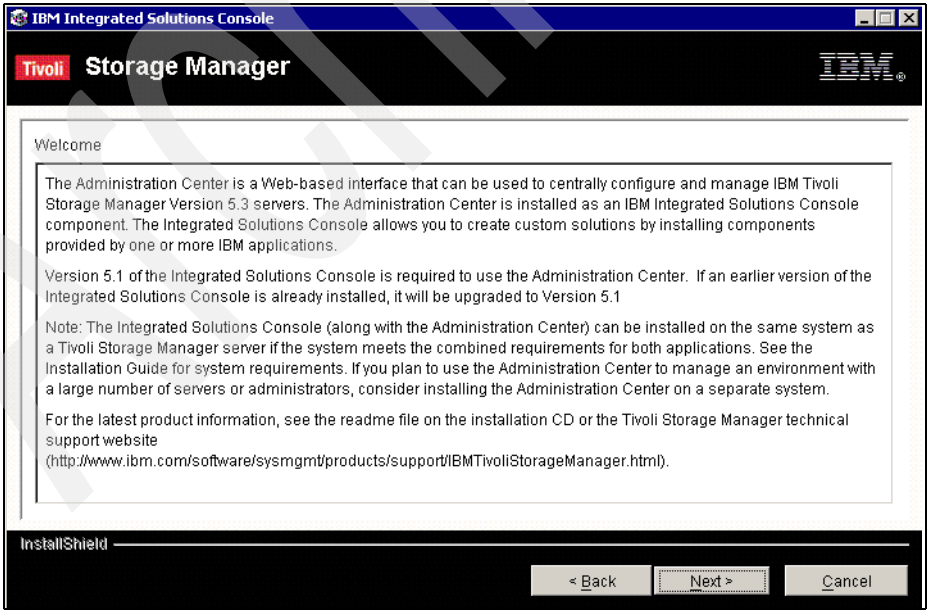

Figure 5-22 Welcome menu

Archived

4. In Figure 5-22 we click **Next** and we get the following menu (Figure 5-23).

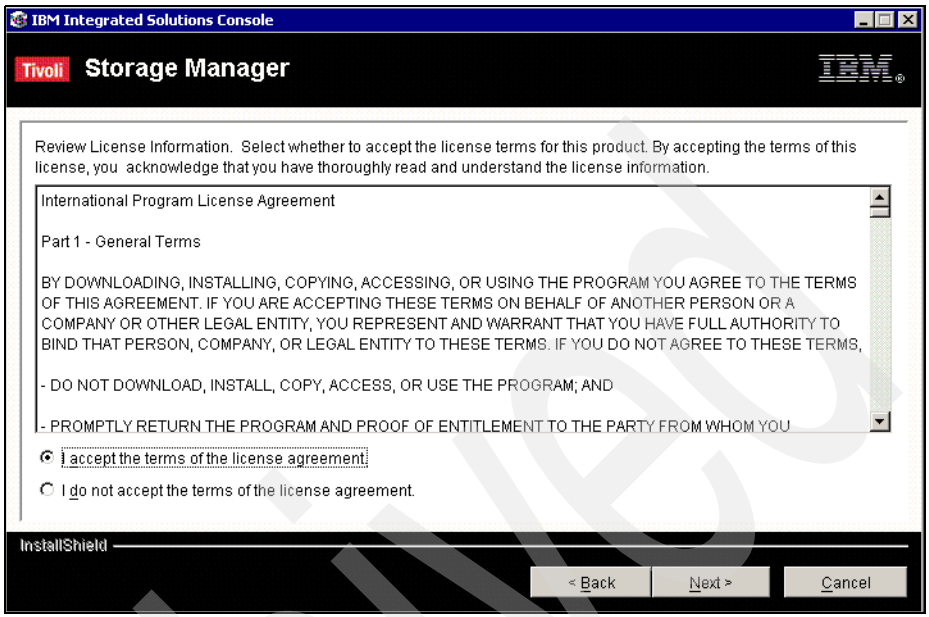

Figure 5-23 ISC License Agreement

5. In Figure 5-23 we select **I accept the terms of the license agreement** and click **Next**. Then, the following menu displays (Figure 5-24).

Archived

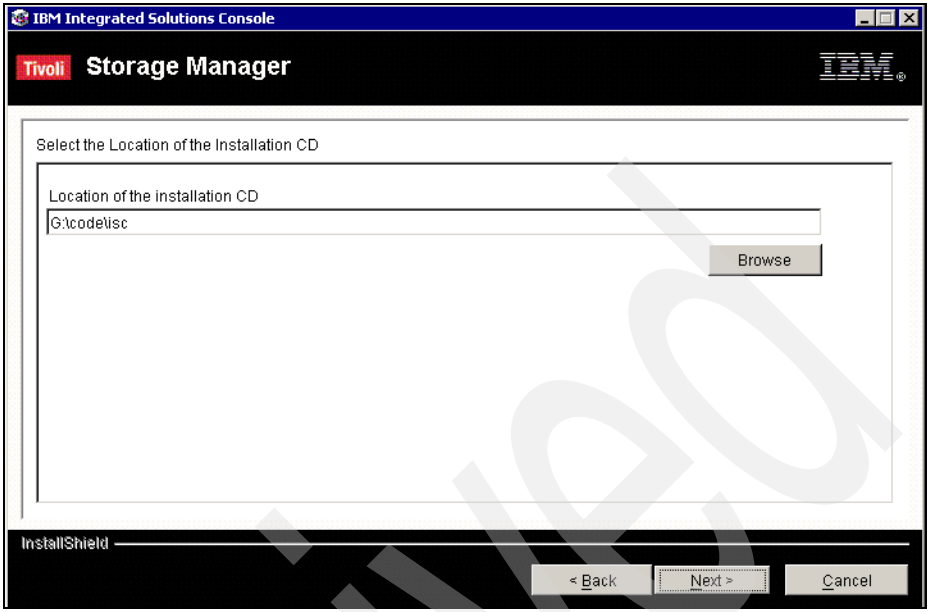

Figure 5-24 Location of the installation CD

6. In Figure 5-24 we type the path where the installation files are located and click **Next**. The following menu displays (Figure 5-25).

Archived

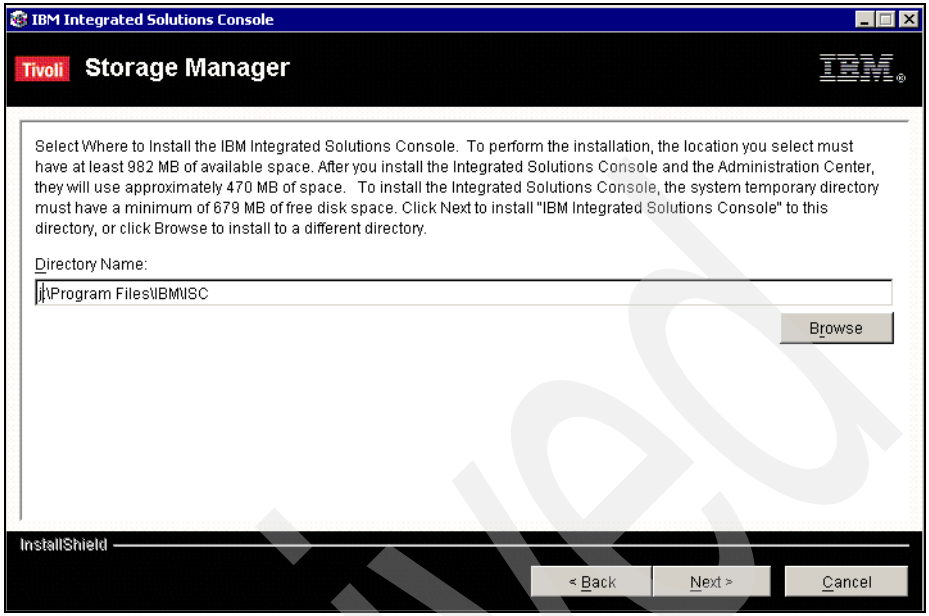

Figure 5-25 Installation path for ISC

7. In Figure 5-25 we type the installation path for the ISC. We choose a shared disk, j:, as the installation path. Then we click **Next** and we see the following panel (Figure 5-26).

Archived

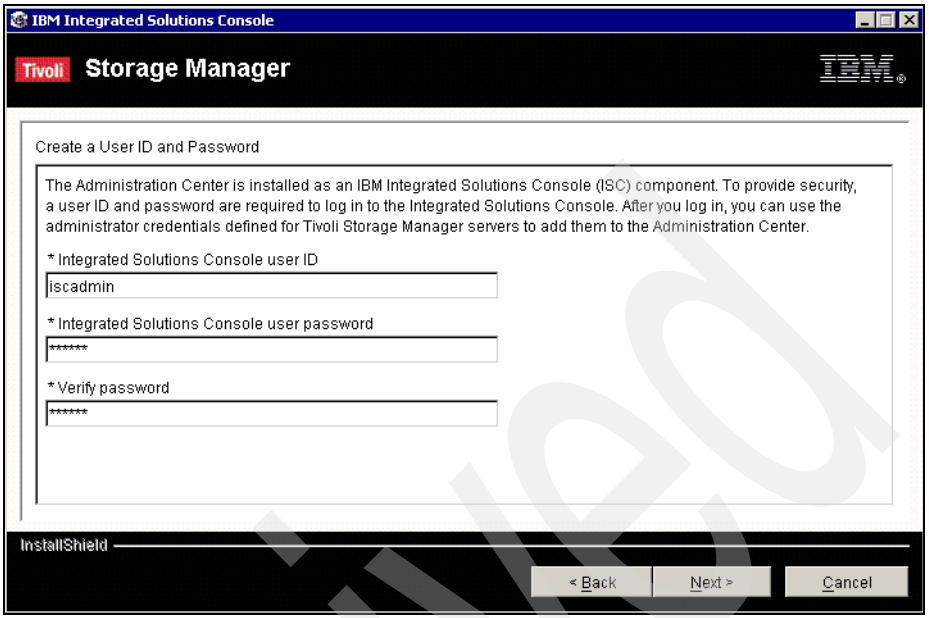

Figure 5-26 Selecting user id and password for the ISC

8. In Figure 5-26 we specify the user ID and password for connection to the ISC. Then, we click **Next** to go to the following menu (Figure 5-27).

Archived

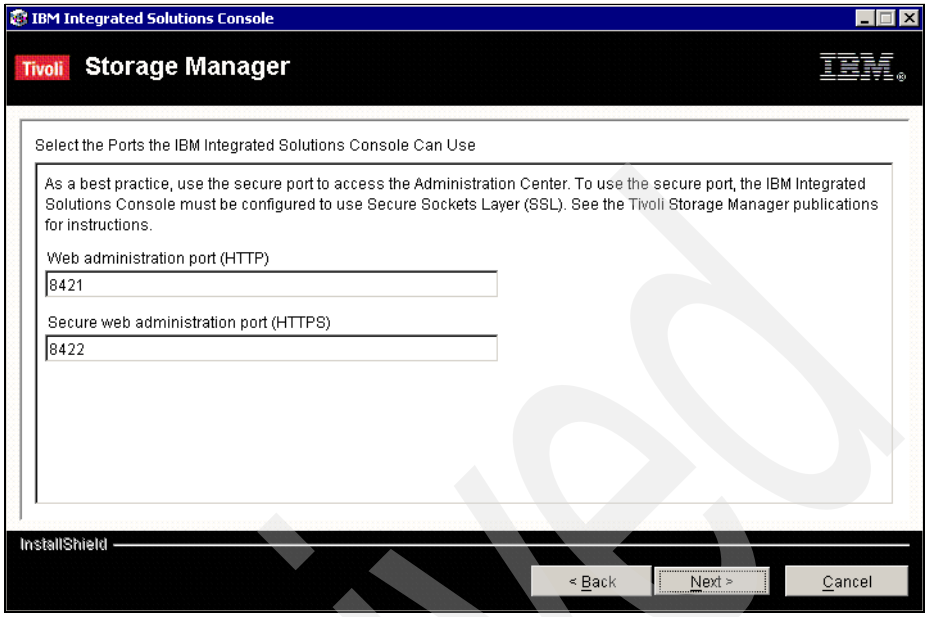

Figure 5-27 Selecting Web administration ports

9. In Figure 5-27 we leave the default Web administration and secure Web administration ports and we click **Next** to go on with the installation. The following menu displays (Figure 5-28).

Archived

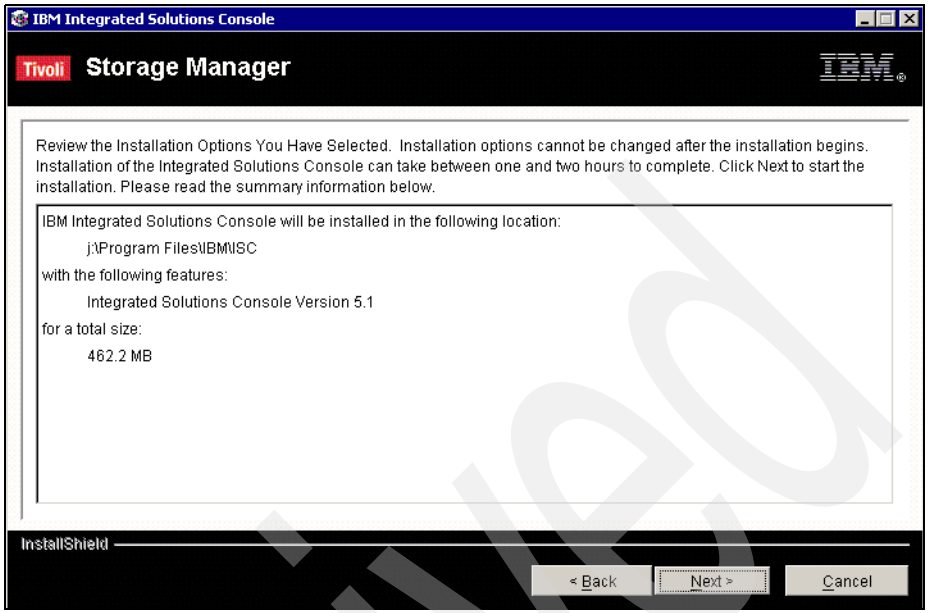

Figure 5-28 Review the installation options for the ISC

10. In Figure 5-28 we click **Next** after checking the information as valid. A welcome menu displays (Figure 5-29).

Archived

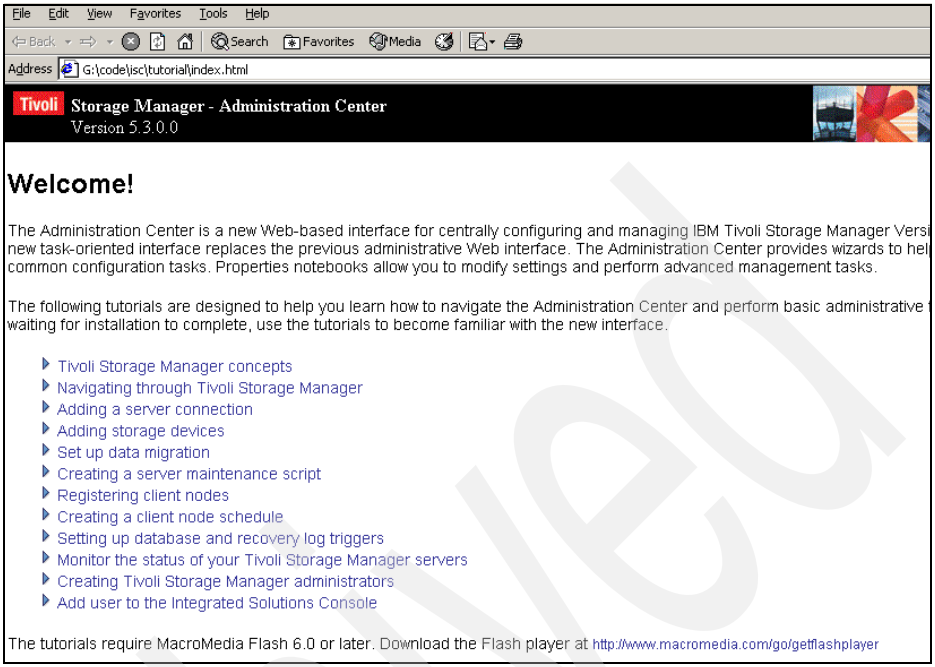

Figure 5-29 Welcome

11. We close the menu in Figure 5-29 and the installation progress bar displays (Figure 5-30).

Archived

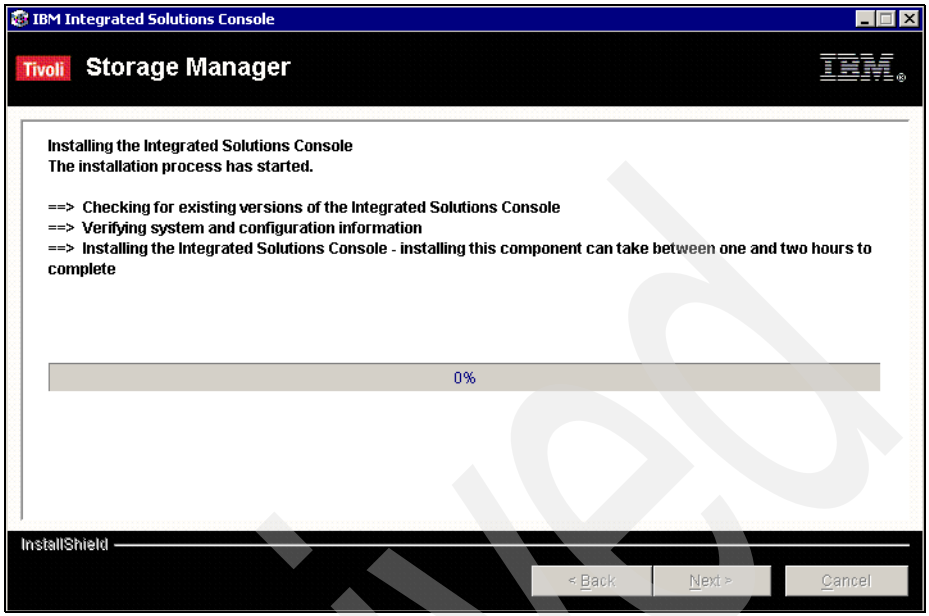

Figure 5-30 Installation progress bar

12. The installation ends and the panel in Figure 5-31 displays.

Archived

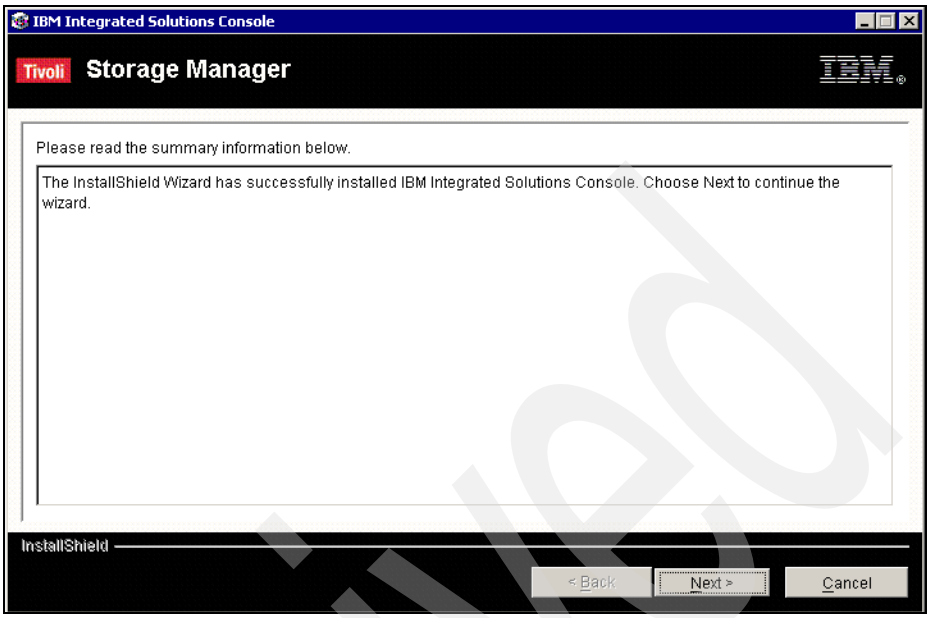

Figure 5-31 ISC Installation ends

13. We click **Next** in Figure 5-31 and an installation summary menu appears. We click **Finish** on it.

The ISC is installed in the first node of each MSCS.

Archived

The installation process creates and starts two Windows services for ISC. These services are shown in Figure 5-32.

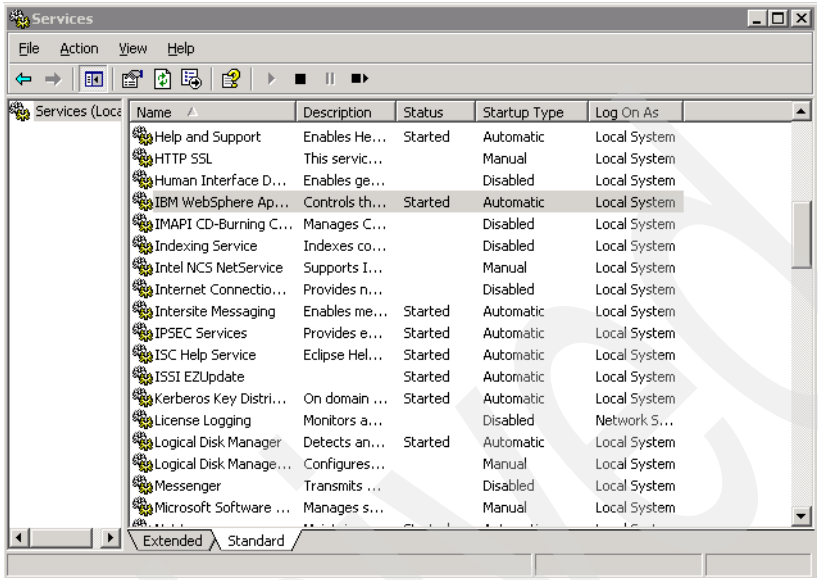

Figure 5-32 ISC services started for the first node of the MSCS

The names of the services are:

- IBM WebSphere Application Server V5 - ISC Runtime Services
- ISC Help Service

Now we proceed to install the Administration Center.

Archived

Installation of the administration center in the first node

These are the tasks we follow to achieve the Administration Center installation in the first node of each cluster.

1. We run **`setupac.exe`** from the CD. The welcome installation menu displays (Figure 5-33).

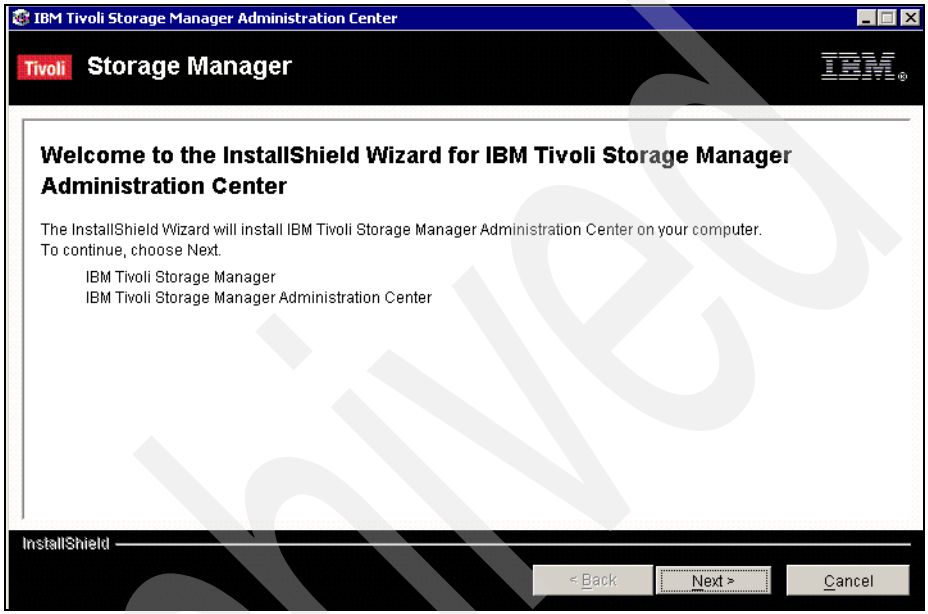

Figure 5-33 Administration Center Welcome menu

2. To start the installation we click **Next** in Figure 5-33 and the following menu displays (Figure 5-34).

Archived

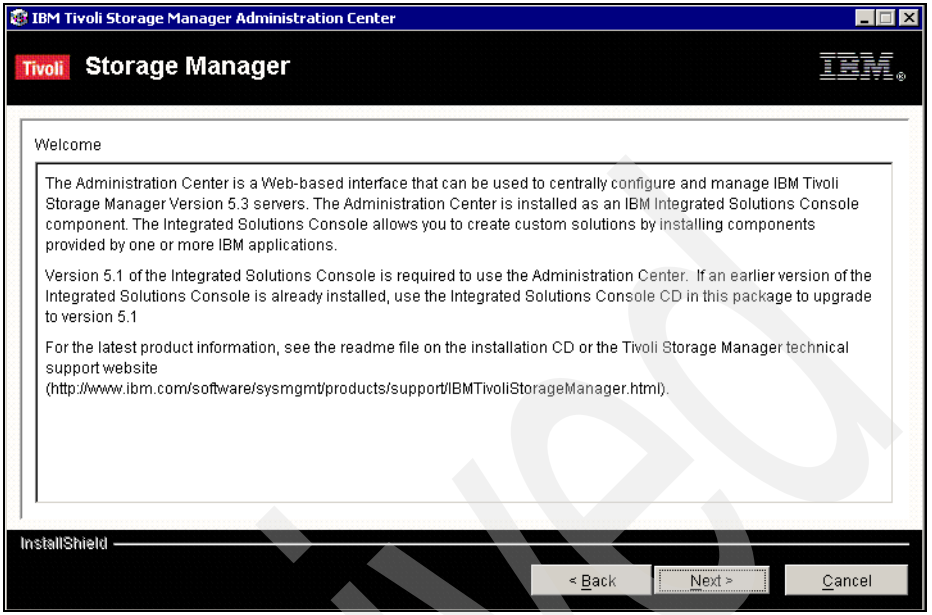

Figure 5-34 Administration Center Welcome

3. In Figure 5-34 we click **Next** to go on with the installation. The following menu displays (Figure 5-35).

Archived

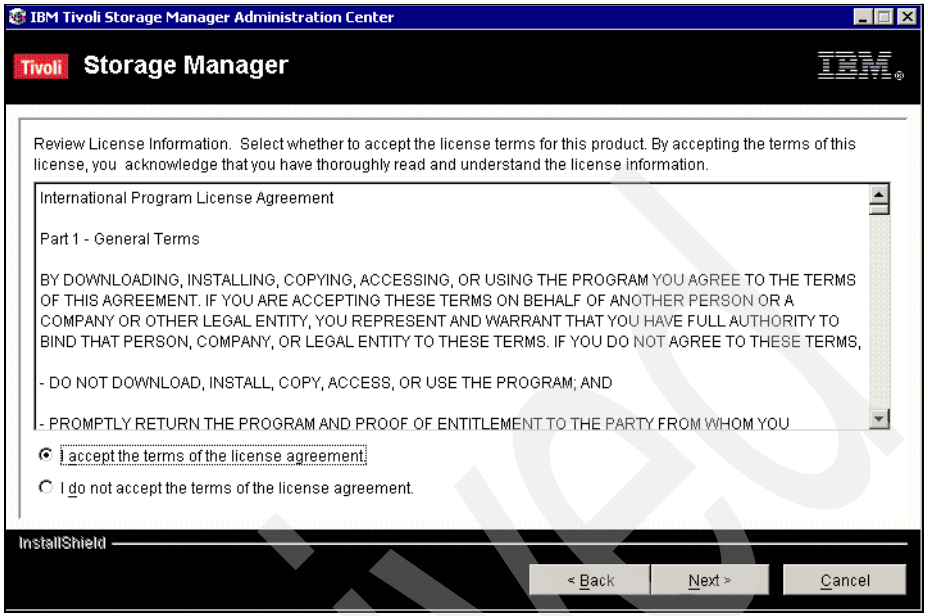

Figure 5-35 Administration Center license agreement

4. The license agreement displays as shown in Figure 5-35. We select **I accept the terms of the license agreement** and we click **Next** to follow with the installation process (Figure 5-36).

Archived

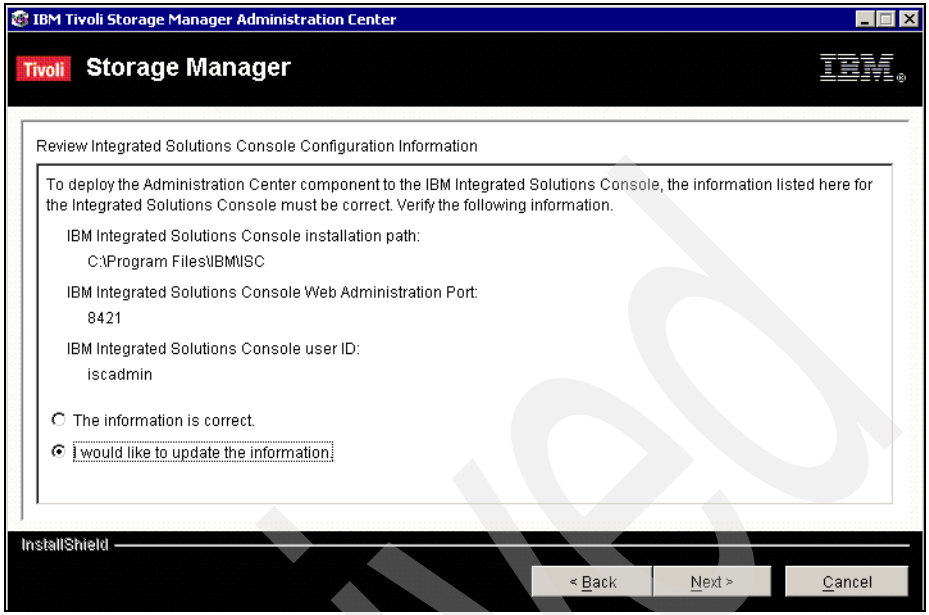

Figure 5-36 Modifying the default options

5. Since we did not install the ISC in the local disk, but in the j: disk drive, we select **I would like to update the information** in Figure 5-36 and we click **Next** (Figure 5-37).

Archived

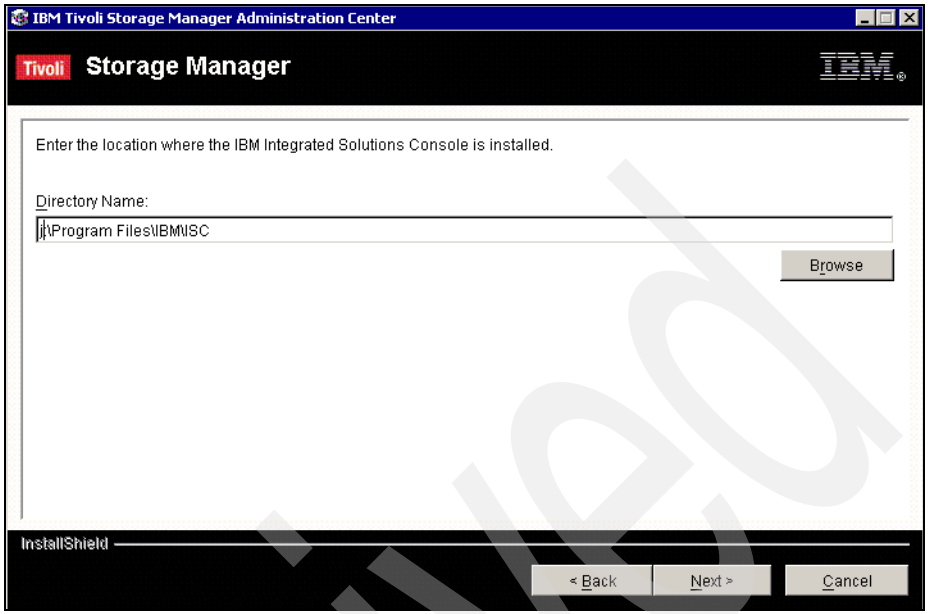

Figure 5-37 Updating the ISC installation path

6. We specify the installation path for the ISC in Figure 5-37 and then we click **Next** to follow with the process. The Web administration port menu displays (Figure 5-38).

Archived

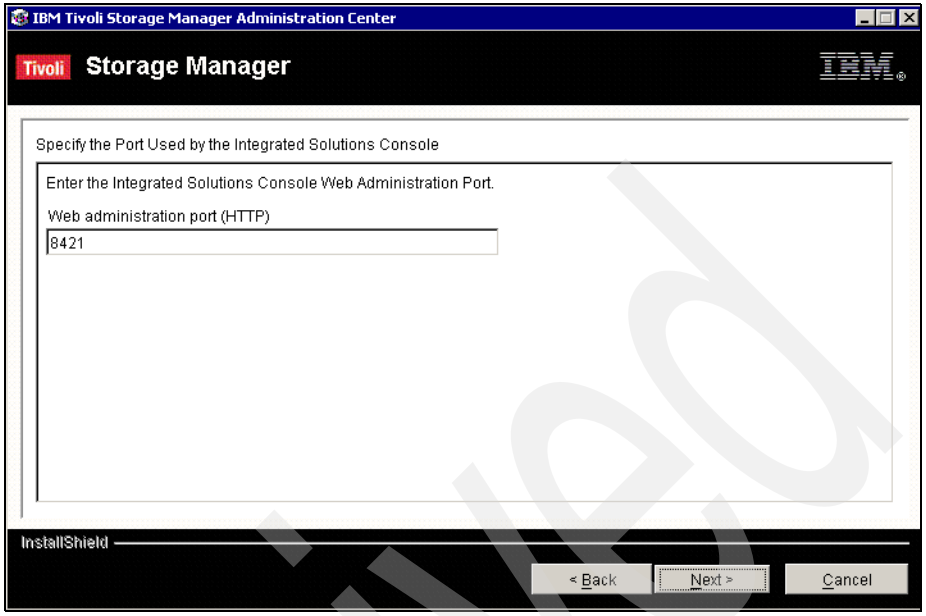

Figure 5-38 Web administration port

7. We leave the default port and we click **Next** in Figure 5-38 to get the following menu (Figure 5-39).

Archived

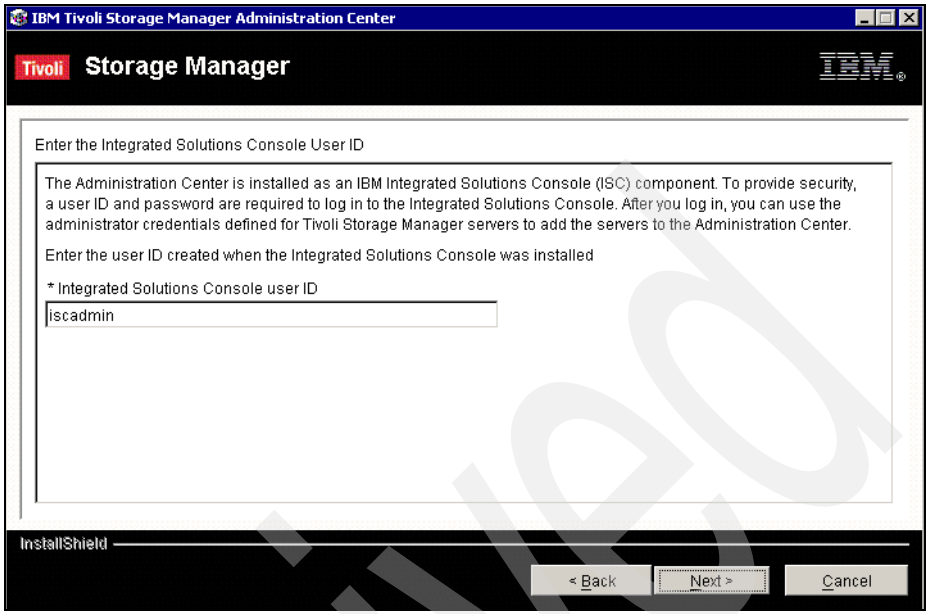

Figure 5-39 Selecting the administrator user id

8. We type the same the user ID created at ISC installation and we click **Next** in Figure 5-39. Then we must specify the password for this user ID in the following menu (Figure 5-40).

Archived

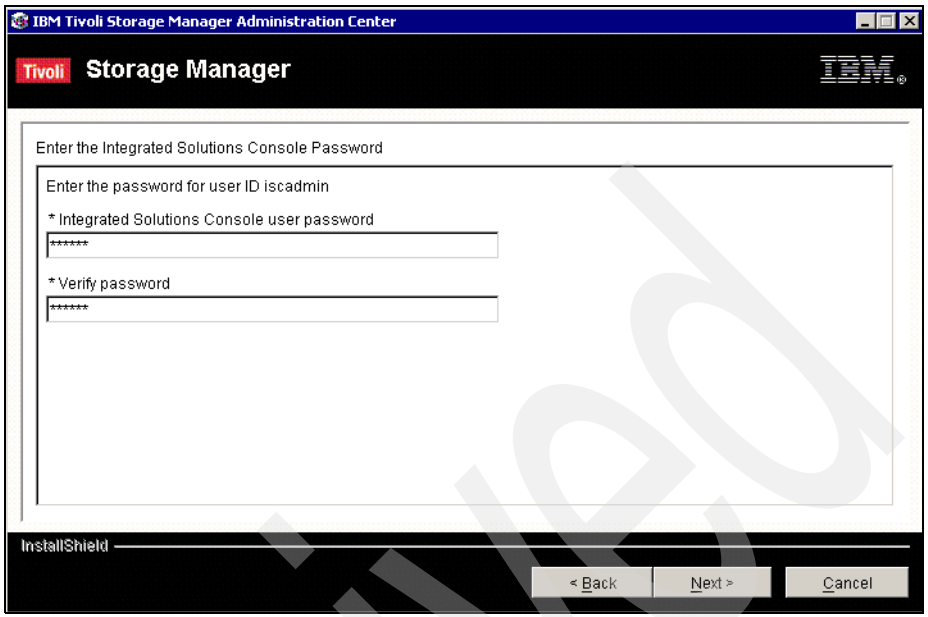

Figure 5-40 Specifying the password for the iscadmin user id

9. We type the password twice for verification in Figure 5-40 and we click **Next** (Figure 5-41).

Archived

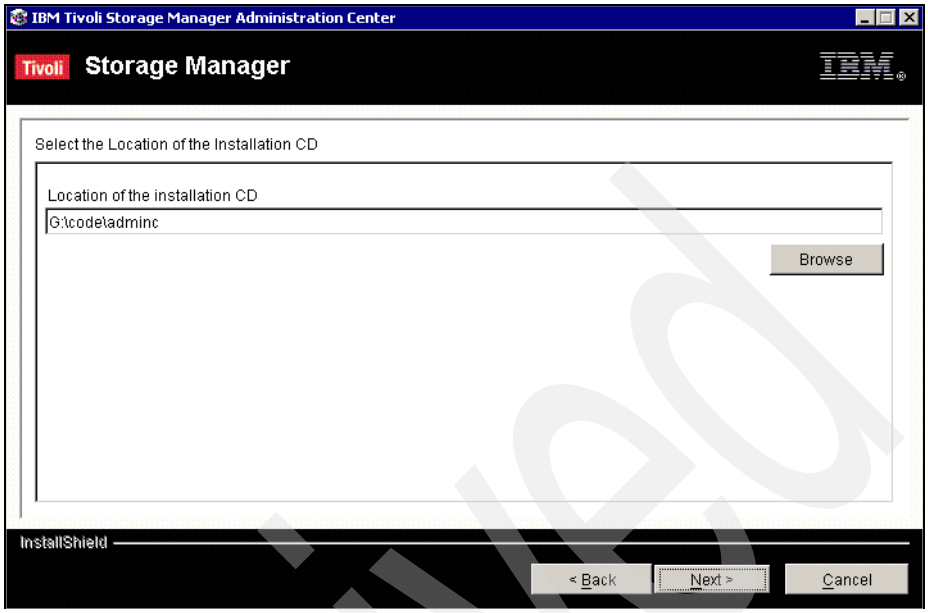

Figure 5-41 Location of the administration center code

10. Finally, in Figure 5-41 we specify the location of the installation files for the Administration Center code and we click **Next**. The following panel displays (Figure 5-42).

Archived

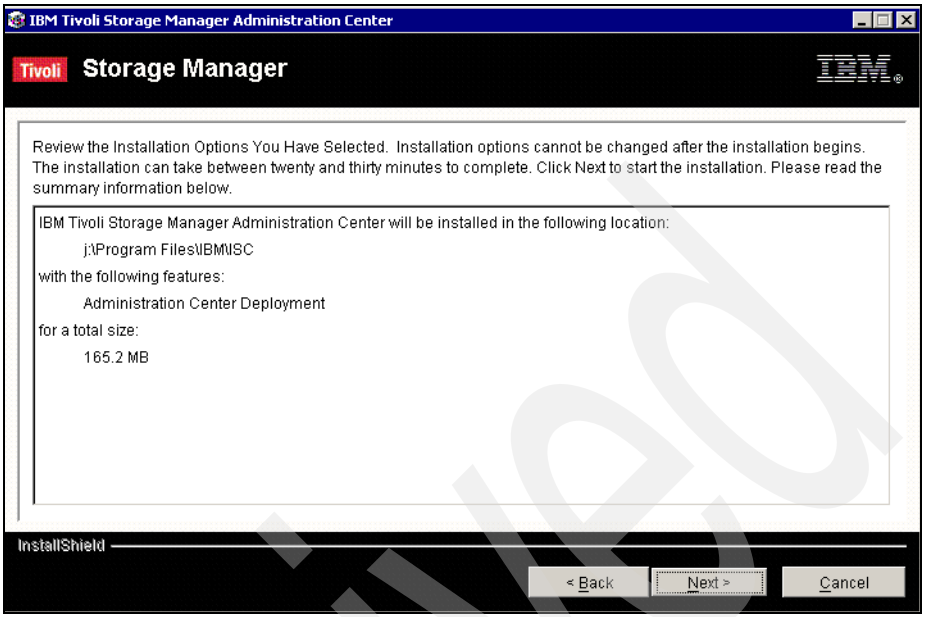

Figure 5-42 Reviewing the installation options

Archived

11. We check the installation options in Figure 5-42 and we select **Next** to start the installation. The installation progress bar displays as shown in Figure 5-43.

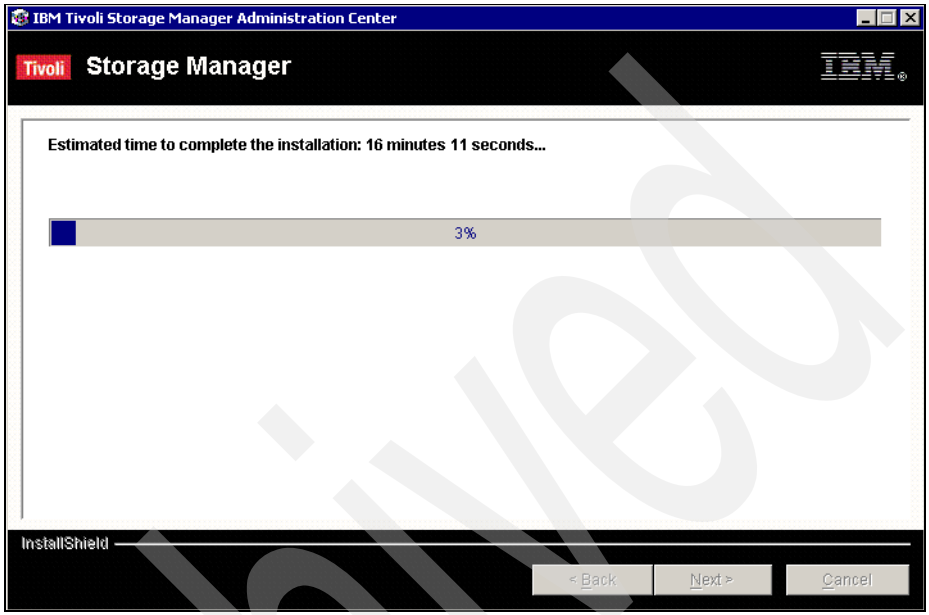

Figure 5-43 Installation progress bar for the Administration Center

Archived

12. When the installation ends, we receive the following panel, where we click **Next** (Figure 5-44).

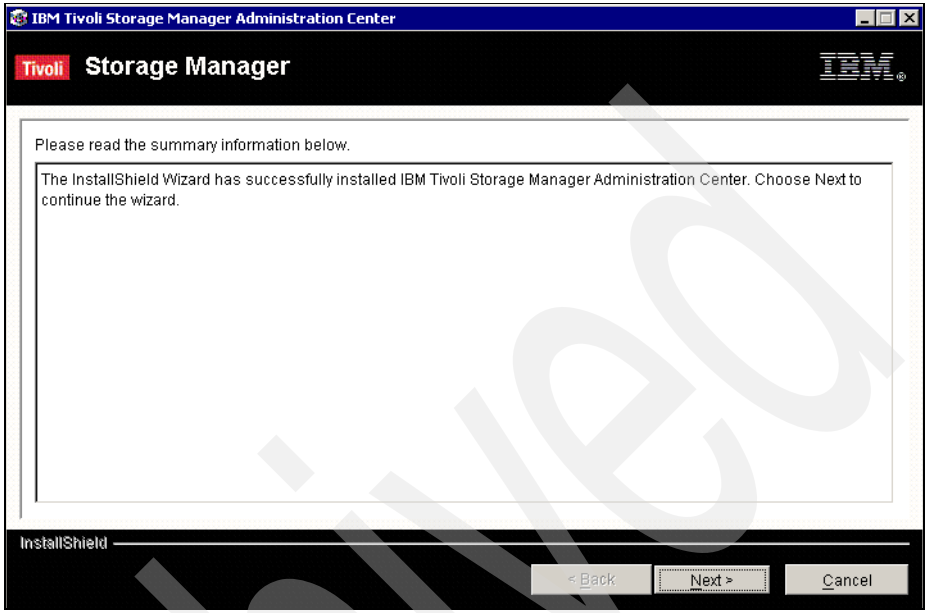

Figure 5-44 Administration Center installation ends

13. An installation summary menu displays next. We click **Next** in this menu.
14. After the installation, the administration center Web page displays, prompting for a user id and a password as shown in Figure 5-45. We close this menu.

Archived

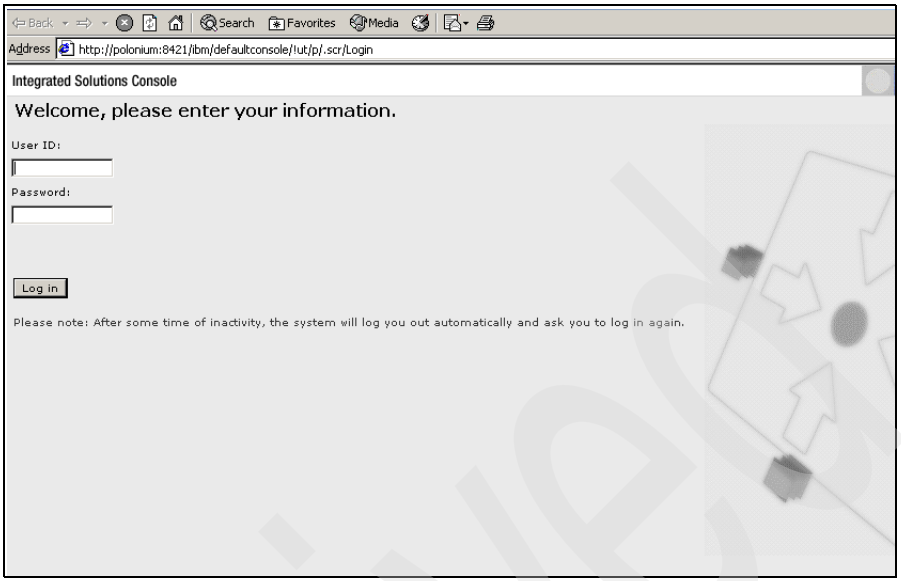

Figure 5-45 Main Administration Center menu

Installation of ISC in the second node

Before installing the ISC and Administration Center in the second node, we need to run three tasks in the first node of each MSCS:

1. Changing the ISC services to *manual* start.
2. Stopping both ISC services.
3. Shutting down the node.

The default startup type for ISC services is set to *Automatic*. Since we want to install this application as a cluster application, we must change to *Manual*.

We also need to stop both services and shut down the first node to make sure that the installation in the second node is correct and there is no shared information between them.

To install the ISC code in the second node of each MSCS, we first delete the `ISC` folder, with all its data and executable files, under `j:\program files\IBM`. We follow this method because if we do not, the installation process fails.

When the `ISC` folder is completely removed, we proceed with the installation of the ISC code, following the steps 2 to 13 of “Installation of ISC in the first node” on page 92.

Important: Do not forget to select the same shared disk and installation path for this component, such as we did in the first node.

The installation process creates and starts in this second node the same two Windows services for ISC, created in the first node, as we can see in Figure 5-46.

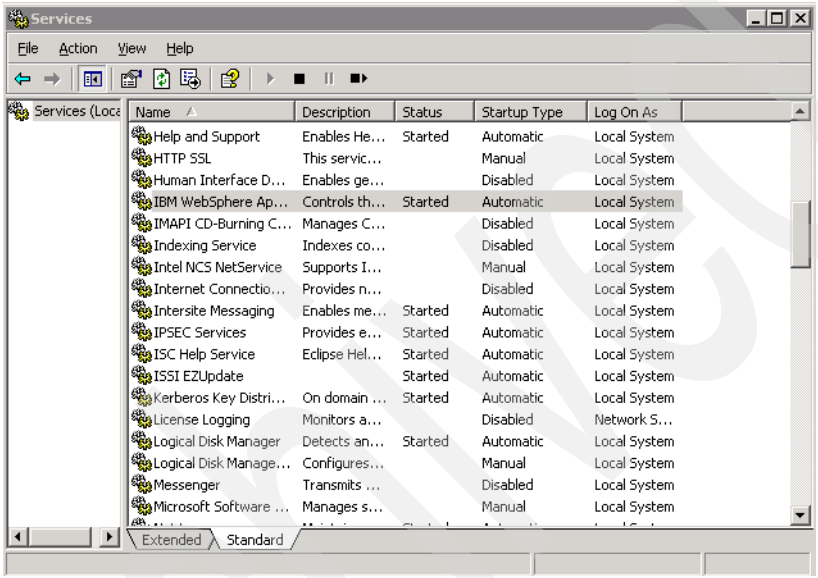

Figure 5-46 ISC Services started as automatic in the second node

Now we proceed to install the Administration Center.

Installation of the Administration Center in the second node

In order to install the Administration Center in the second node of each MSCS, we proceed with steps 1 to 14 of “Installation of the administration center in the first node” on page 104.

Important: Do not forget to select the same shared disk and installation path for this component, just like we did in the first node.

When the installation ends, we are ready to configure the ISC component as a cluster application. To achieve this goal we need to change the two ISC services to *Manual* startup type, and to stop both of them.

The final task is starting the first node, and, when it is up, we need to restart this second node for the registry updates to take place in this machine.

Refer to “Configuring ISC for clustering on Windows 2000” on page 167 and “Configuring ISC for clustering on Windows 2003” on page 231 for the specifics of the configuration on each MSCS environment.

5.4 Tivoli Storage Manager server and Windows 2000

The Tivoli Storage Manager server installation process was described on “Installing Tivoli Storage Manager Server on a MSCS” on page 79, at the beginning of this chapter.

In this section we describe how we configure our Tivoli Storage Manager server software to be capable of running in our Windows 2000 MSCS, the same cluster we installed and configured in 4.3, “Windows 2000 MSCS installation and configuration” on page 29.

5.4.1 Windows 2000 lab setup

Our clustered lab environment consists of two Windows 2000 Advanced Servers. Both servers are domain controllers as well as DNS servers.

Archived

Figure 5-47 shows our Tivoli Storage Manager clustered server configuration.

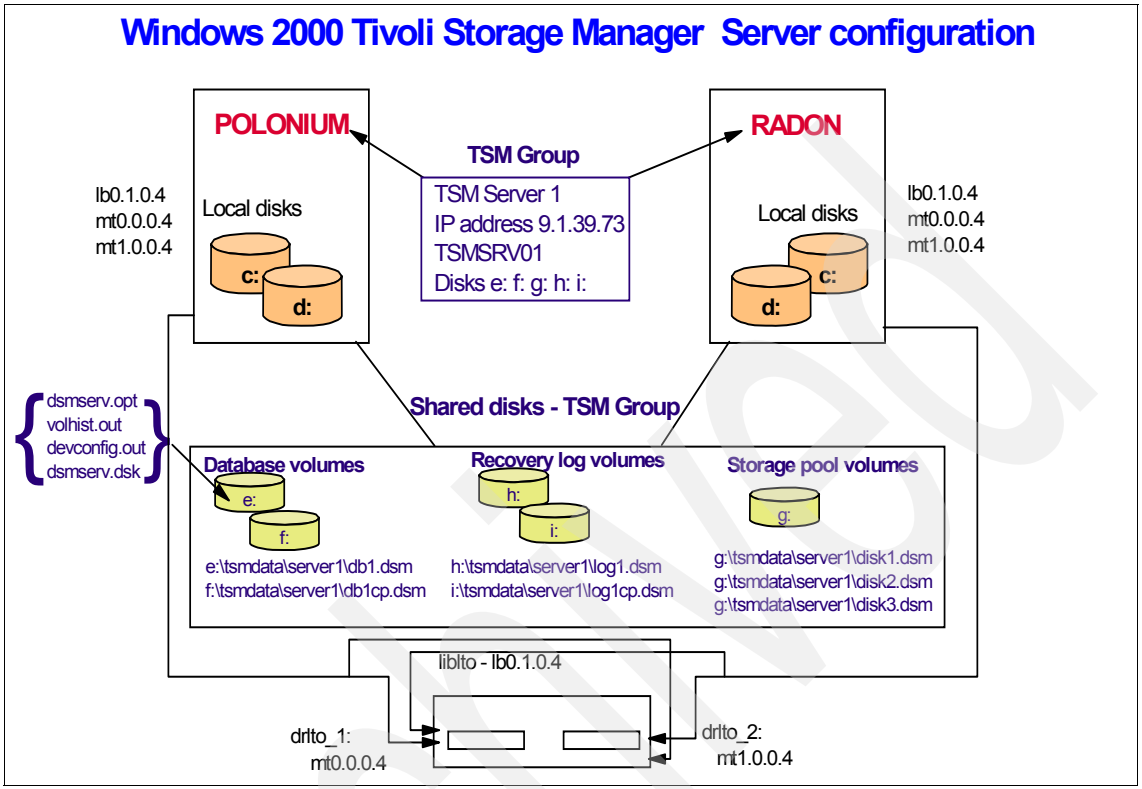

Figure 5-47 Windows 2000 Tivoli Storage Manager clustering server configuration

Archived

Refer to Table 4-1 on page 30, Table 4-2 on page 31, and Table 4-3 on page 31 for specific details of our MSCS configuration.

Table 5-1, Table 5-2, and Table 5-3, below, show the specifics of our Windows 2000 MSCS environment, Tivoli Storage Manager virtual server configuration, and ISC configuration that we use for the purpose of this section.

Table 5-1 Windows 2000 lab ISC cluster resources

Resource Group TSM Admin Center	
ISC name	ADMCNT01
ISC IP address	9.1.39.46
ISC disk	j:
ISC service names	IBM WebSphere Application Server V5 - ISC Runtime Service ISC Help Service

Table 5-2 Windows 2000 lab Tivoli Storage Manager server cluster resources

Resource Group TSM Group	
TSM server name	TSMSRV01
TSM server IP address	9.1.39.73
TSM database disks[a]	e: h:
TSM recovery log disks	f: i:
TSM storage pool disk	g:
TSM service name	TSM Server 1

a. We choose two disk drives for the database and recovery log volumes so that we can use the Tivoli Storage Manager mirroring feature.

Archived

Table 5-3 Windows 2000 Tivoli Storage Manager virtual server in our lab

Server parameters	
Server name	TSMSRV01
High level address	9.1.39.73
Low level address	1500
Server password	itsosj
Recovery log mode	roll-forward
Libraries and drives	
Library name	LIBLTO
Drive 1	DRLTO_1
Drive 2	DRLTO_2
Device names	
Library device name	lb0.1.0.4
Drive 1 device name	mt0.0.0.4
Drive 2 device name	mt1.0.0.4
Primary Storage Pools	
Disk Storage Pool	SPD_BCK (nextstg=SPT_BCK)
Tape Storage Pool	SPT_BCK
Copy Storage Pool	
Tape Storage Pool	SPCPT_BCK
Policy	
Domain name	STANDARD
Policy set name	STANDARD
Management class name	STANDARD
Backup copy group	STANDARD (default, DEST=SPD_BCK)
Archive copy group	STANDARD (default)

Archived

Before installing the Tivoli Storage Manager server on our Windows 2000 cluster, the TSM Group *must only* contains disk resources, as we can see in the Cluster Administrator menu in Figure 5-48.

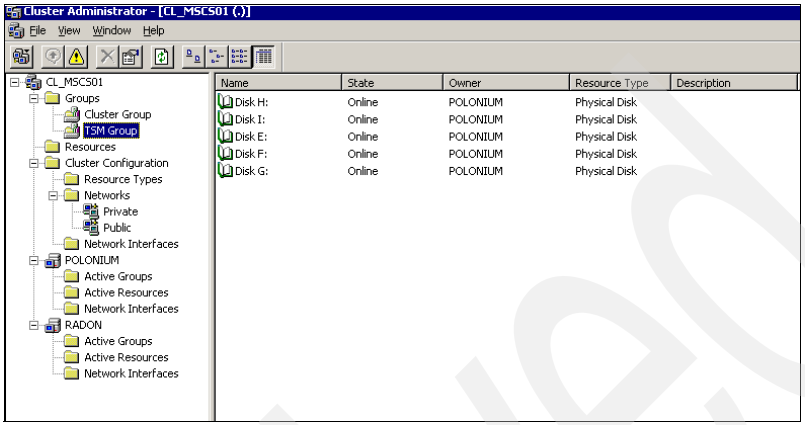

Figure 5-48 Cluster Administrator with TSM Group

Installation of IBM tape device drivers on Windows 2000

As we can see in Figure 4-1 on page 29, our two Windows 2000 servers are attached to the *Storage Area Network*, so that both can see the IBM 3582 Tape Library as well as its two IBM 3580 tape drives.

Since IBM Tape Libraries use their own device drivers to work with Tivoli Storage Manager, we have to download and install the last available version of the IBM LTO drivers for 3582 Tape Library and 3580 Ultrium 2 tape drives.

We use the folder “drivers_lto” to download the IBM drivers. Then, we use the Windows device manager menu, right-click one of the drives and select **Properties →Driver →Update driver**. We must specify the path where to look for the drivers, the “drivers_lto” folder, and follow the installation process menus.

We do not show the whole installation process in this book. Refer to the *IBM Ultrium Device Drivers Installation and User’s Guide* for a detailed description of this task.

After the successful installation of the drivers, both nodes recognize the 3582 medium changer and the 3580 tape drives as shown in Figure 5-49:

Archived

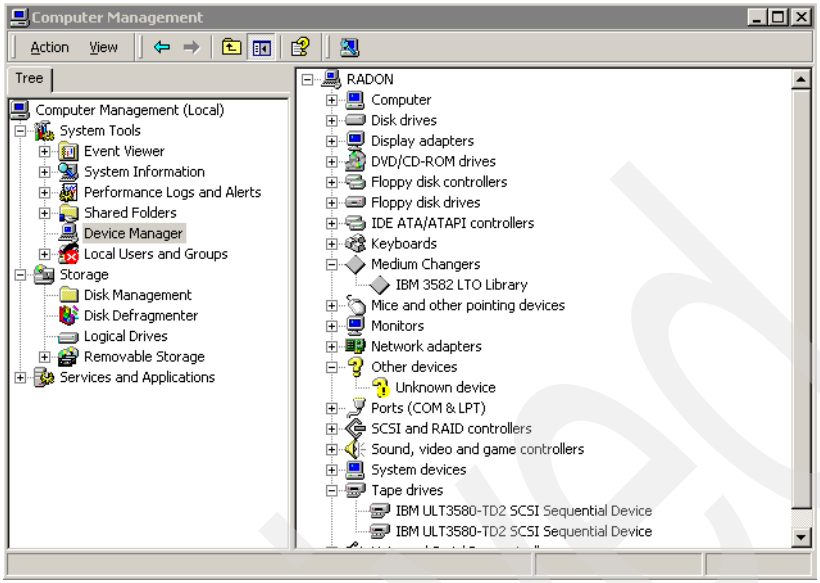

Figure 5-49 Successful installation of IBM 3582 and IBM 3580 device drivers

5.4.2 Windows 2000 Tivoli Storage Manager Server configuration

When the installation of the Tivoli Storage Manager packages on both nodes of the cluster is completed, we can proceed with the configuration.

The configuration tasks are performed on each node of the cluster. The steps vary depending upon whether it is the first node we are configuring or the second one.

When we start the configuration procedure on the first node, the Tivoli Storage Manager server instance is created and started. On the second node, the procedure will allow this server to host that instance.

> **Important:** it is necessary to install a Tivoli Storage Manager server on the first node before configuring the second node. If we do not that, the configuration will fail.

Configuring the first node

We start configuring Tivoli Storage Manager on the first node. To perform this task, resources must be hosted by this node. We can check this issue, opening the cluster administrator from **Start** → **Programs** → **Administrative Tools** → **Cluster Administrator** (Figure 5-50).

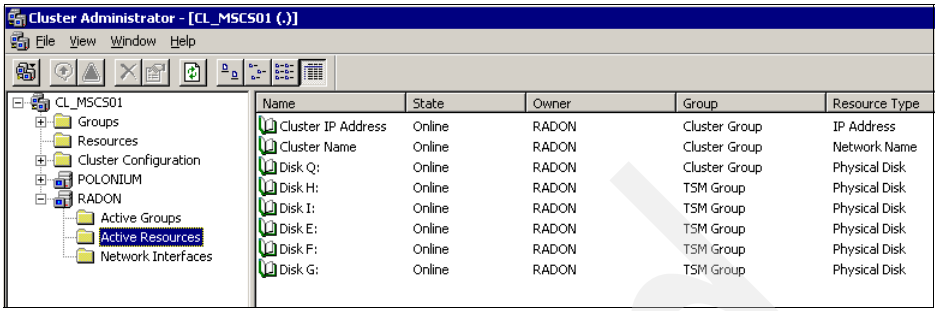

Figure 5-50 Cluster resources

As shown in Figure 5-50, RADON hosts all the resources of the TSM Group. That means we can start configuring Tivoli Storage Manager on this node.

Attention: Before starting the configuration process, we copy mfc71u.dll and mvscr71.dll files from the Tivoli Storage Manager \console directory (normally c:\Program Files\Tivoli\tsm\console) into our c:\%SystemRoot%\cluster directory on each cluster node involved. If we do not do that, the cluster configuration will fail. This is caused by a new Windows compiler (VC71) that creates dependencies between tsmsvrrsc.dll and tsmsvrrscex.dll and mfc71u.dll and mvscr71.dll. Microsoft has not included these files in its service packs.

1. To start the initialization, we open the Tivoli Storage Manager Management Console as shown in Figure 5-51.

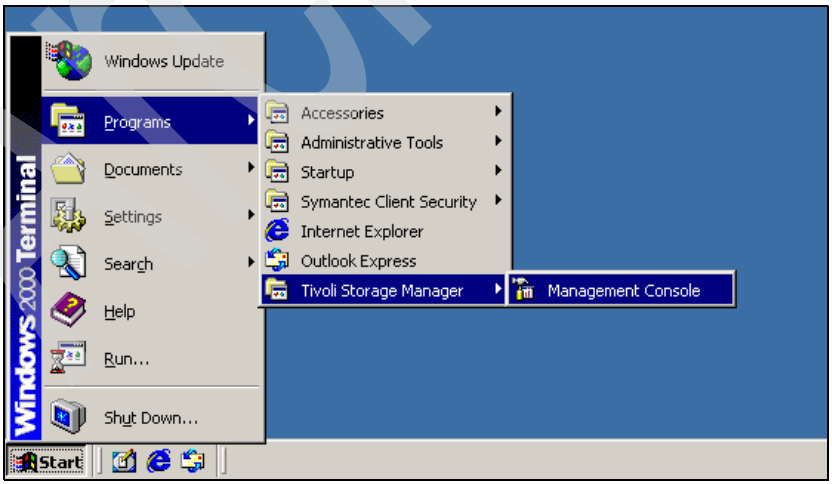

Figure 5-51 Starting the Tivoli Storage Manager management console

2. The Initial Configuration Task List for Tivoli Storage Manager menu, Figure 5-52, shows a list of the tasks needed to configure a server with all basic information. To let the wizard guide us throughout the process, we select **Standard Configuration**. This will also enable automatic detection of a clustered environment. We then click **Start**.

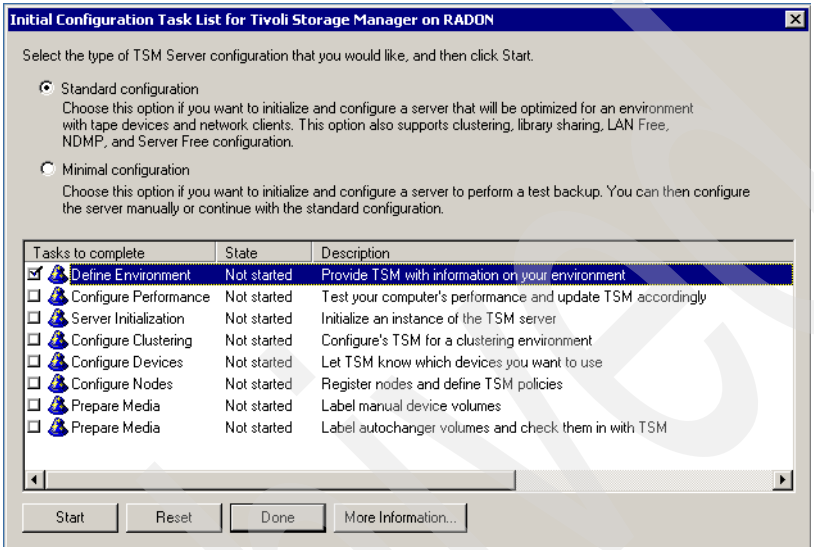

Figure 5-52 Initial Configuration Task List

Archived

3. The Welcome menu for the first task, **Define Environment**, displays (Figure 5-53). We click **Next**.

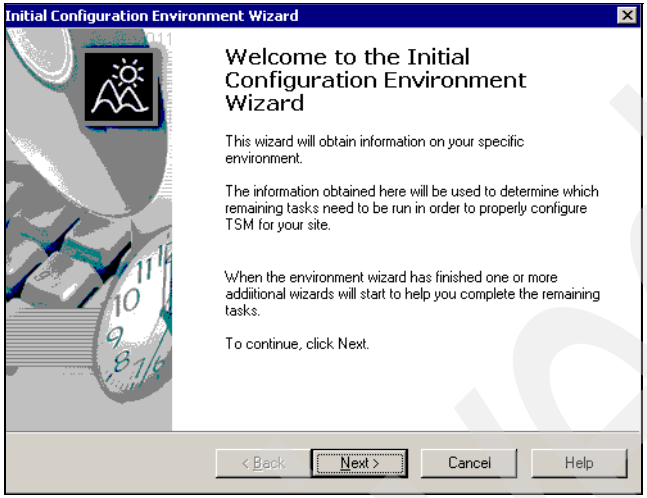

Figure 5-53 Welcome Configuration wizard

4. To have additional information displayed during the configuration, we select **Yes** and click **Next** as shown in Figure 5-54.

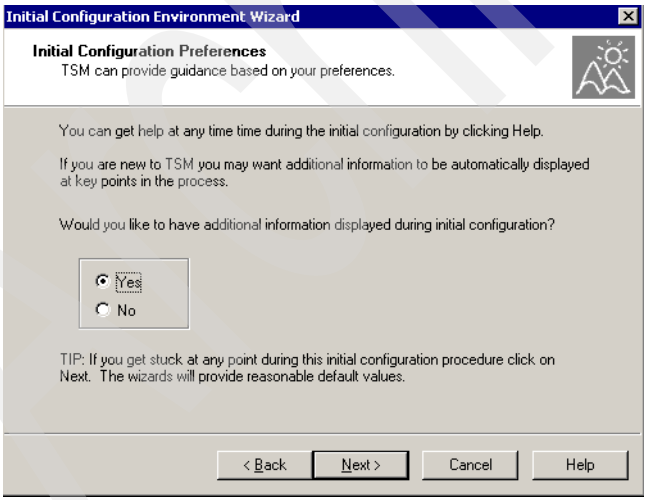

Figure 5-54 Initial configuration preferences

5. Tivoli Storage Manager can be installed Standalone (for only one client), or Network (when there are more clients). In most cases we have more than one client. We select **Network** and then click **Next** as shown in Figure 5-55.

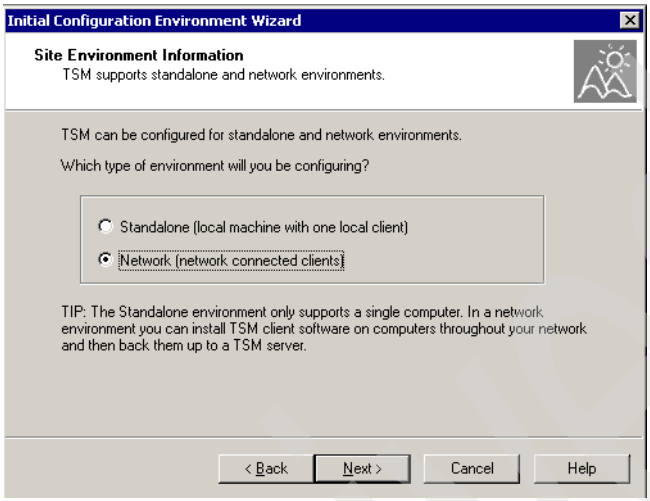

Figure 5-55 Site environment information

6. The Initial Configuration Environment is done. We click **Finish** in Figure 5-56:

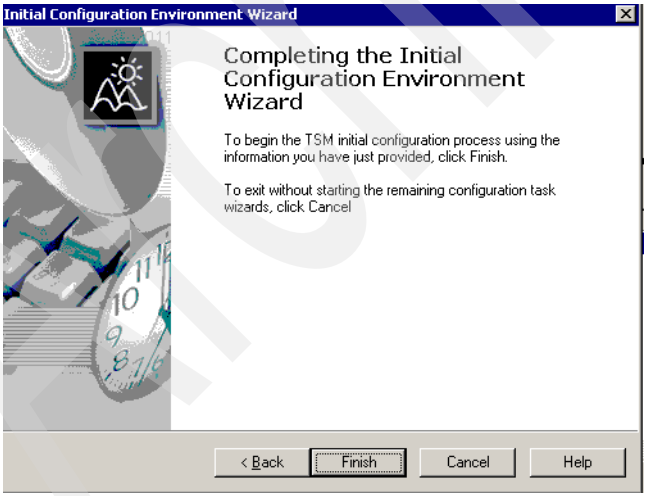

Figure 5-56 Initial configuration

Archived

7. The next task is to complete the Performance Configuration Wizard. We click **Next** (Figure 5-57).

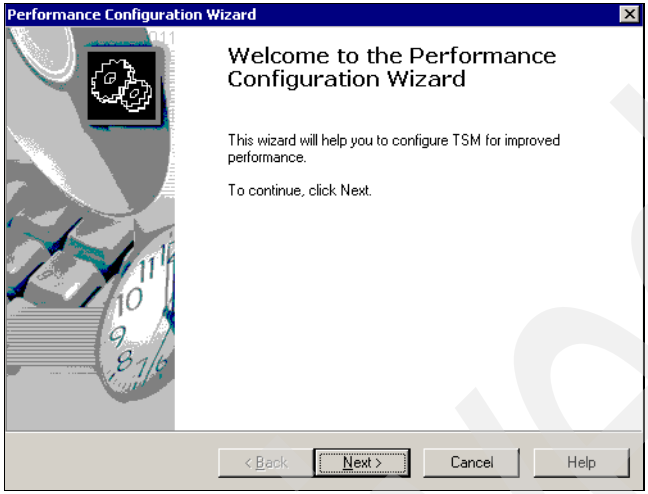

Figure 5-57 Welcome Performance Environment wizard

8. In Figure 5-58 we provide information about our own environment. Tivoli Storage Manager will use this information for tuning. For our lab we used the defaults. In a real installation, it is necessary to select the values that best fit that environment. We click **Next**.

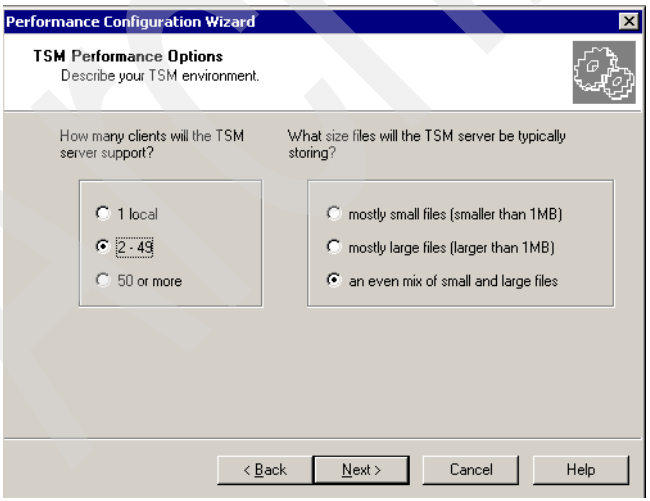

Figure 5-58 Performance options

9. The wizard starts to analyze the hard drives as shown in Figure 5-59. When the process ends, we click **Finish**.

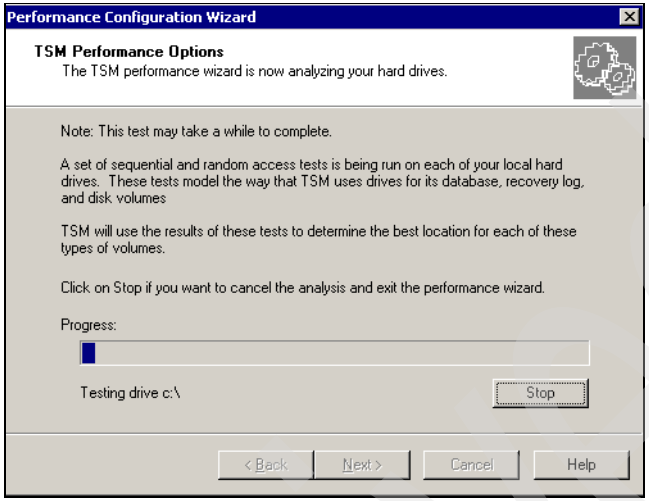

Figure 5-59 Drive analysis

10. The Performance Configuration task completes (Figure 5-60).

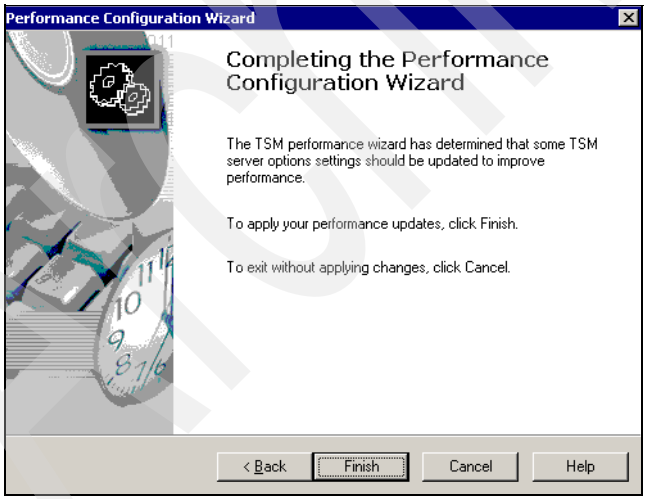

Figure 5-60 Performance wizard

11. Next step is the initialization of the Tivoli Storage Manager server instance. We click **Next** (Figure 5-61).

Archived

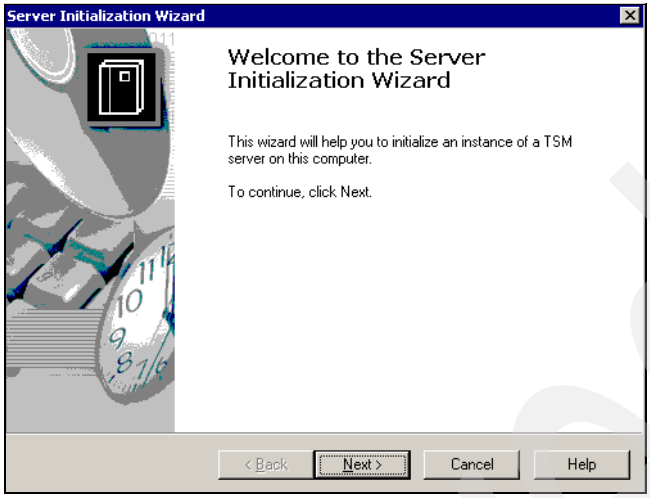

Figure 5-61 Server instance initialization wizard

12. The initialization process detects that there is a cluster installed. The option **Yes** is already selected. We leave this default in Figure 5-62 and we click **Next** so that Tivoli Storage Manager server instance is installed correctly.

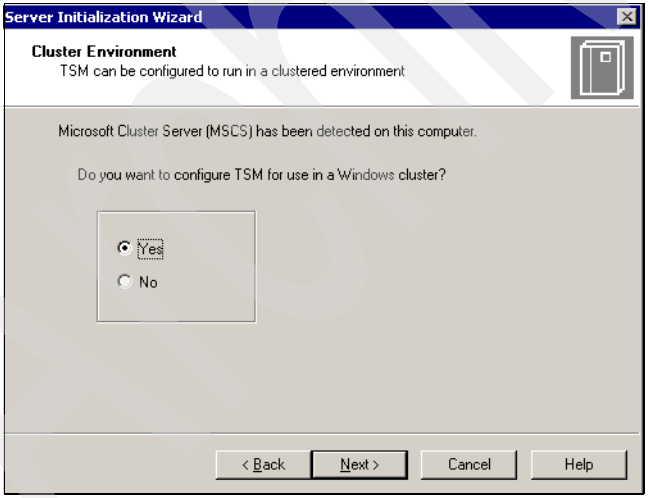

Figure 5-62 Cluster environment detection

13. We select the cluster group where Tivoli Storage Manager server instance will be created. This cluster group initially must contain only disk resources. For our environment this is TSM Group. Then we click **Next** (Figure 5-63).

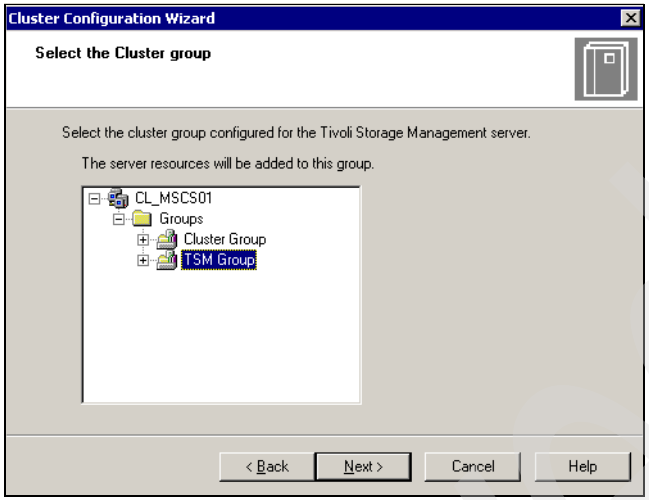

Figure 5-63 Cluster group selection

Important: The cluster group chosen here must match the cluster group used when configuring the cluster in Figure 5-72 on page 136.

14. In Figure 5-64 we select the directory where the files used by Tivoli Storage Manager server will be placed. It is possible to choose any disk on the Tivoli Storage Manager cluster group. We change the drive letter to use e: and click **Next**.

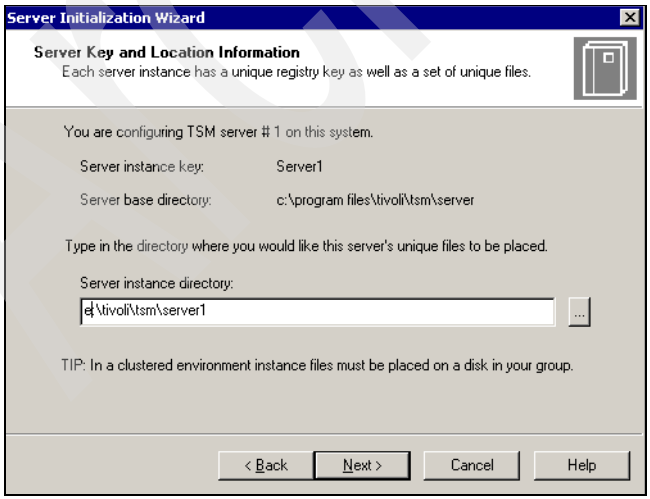

Figure 5-64 Server initialization wizard

15. In Figure 5-65 we type the complete path and sizes of the initial volumes to be used for database, recovery log and disk storage pools. Refer to Table 5-2 on page 120 where we describe our cluster configuration for Tivoli Storage Manager server.

    A specific installation should choose its own values.

    We also check the two boxes on the two bottom lines to let Tivoli Storage Manager create additional volumes as needed.

    With the selected values we will initially have a 1000 MB size database volume with name db1.dsm, a 500 MB size recovery log volume called log1.dsm, and a 5 GB size storage pool volume of name disk1.dsm. If we need, we can create additional volumes later.

    We input our values and click **Next** (Figure 5-65).

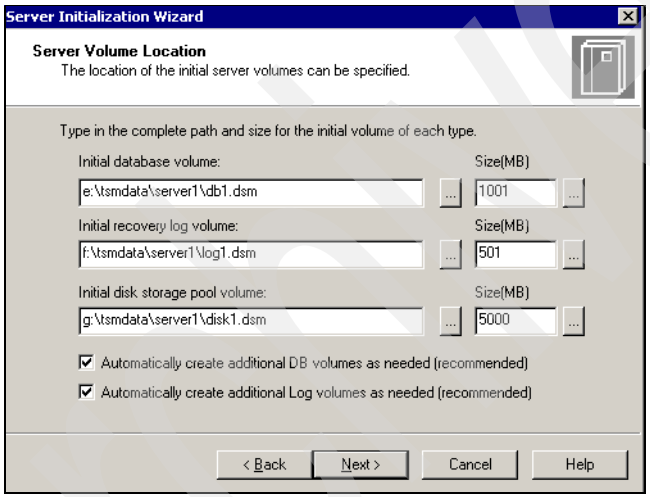

Figure 5-65 Server volume location

16. On the server service logon parameters shown in Figure 5-66 we select the Windows account and user id that Tivoli Storage Manager server instance will use when logging onto Windows. We recommend to leave the defaults and click **Next**.

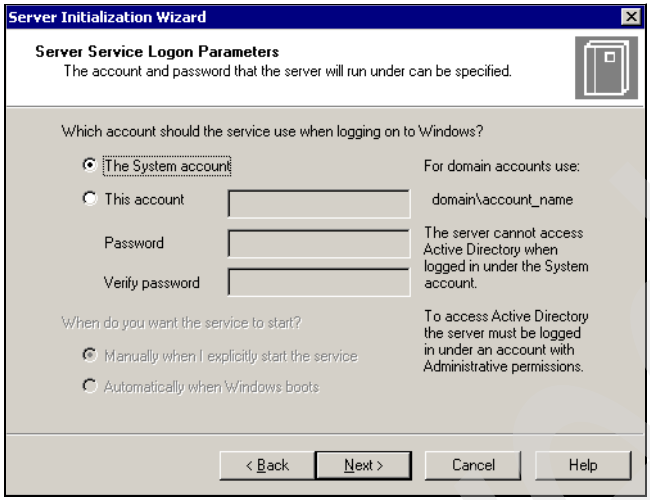

Figure 5-66 Server service logon parameters

17. In Figure 5-67, we assign the server name that Tivoli Storage Manager will use as well as its password. The server password is used for server-to-server communications. We will need it later on with Storage Agent.This password can also be set later using the administrator interface. We click **Next**.

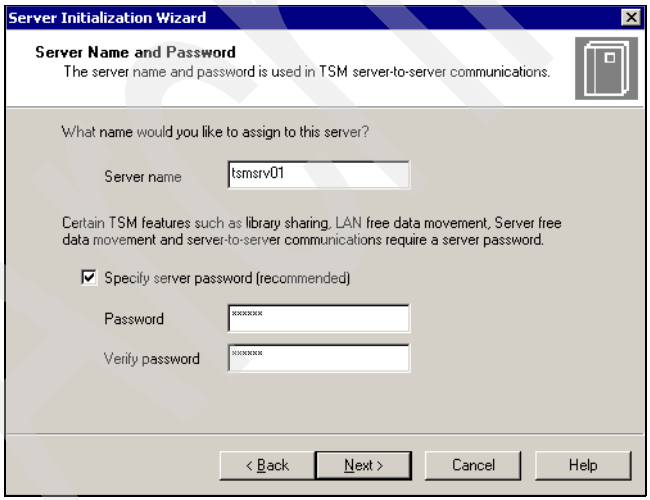

Figure 5-67 Server name and password

Important: the server name we select here must be the same name we will use when configuring Tivoli Storage Manager on the other node of the MSCS.

18. We click **Finish** in Figure 5-68 to start the process of creating the server instance.

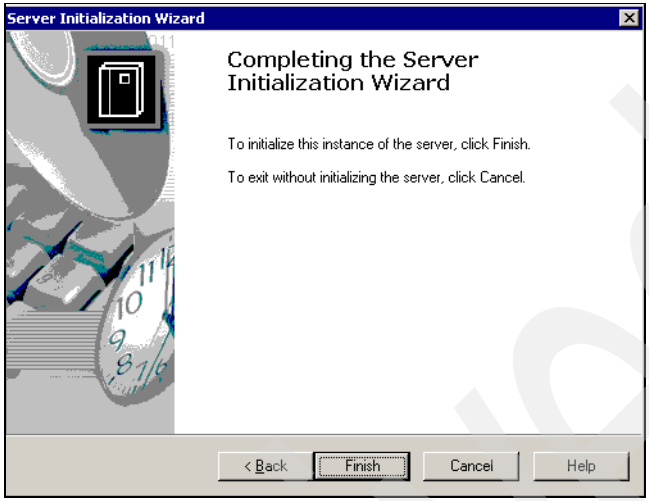

Figure 5-68 Completing the Server Initialization wizard

19. The wizard starts the process of the server initialization and shows a progress bar (Figure 5-69).

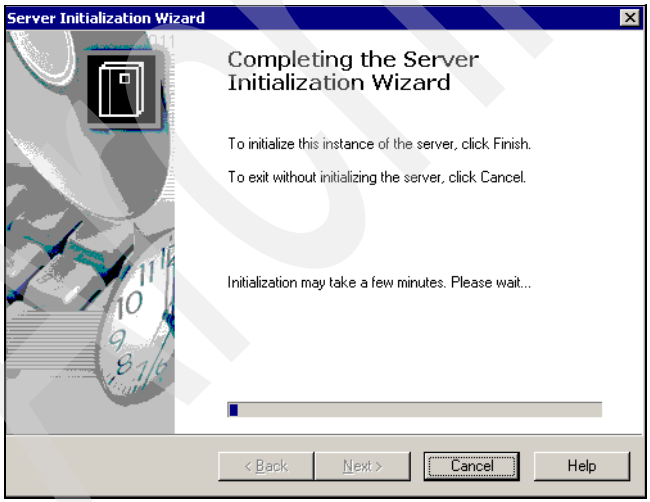

Figure 5-69 Completing the server installation wizard

20. If the initialization ends without any errors, we receive the following informational message. We click **OK** (Figure 5-70).

Archived

Figure 5-70 Tivoli Storage Manager Server has been initialized

21. The next task the wizard performs is the Cluster Configuration. We click **Next** on the welcome page (Figure 5-71).

Figure 5-71 Cluster configuration wizard

22. We select the cluster group where Tivoli Storage Manager server will be configured and click **Next** (Figure 5-72).

Important: *Do not forget* that the cluster group we select here, must match the cluster group used during the server initialization wizard process in Figure 5-63 on page 131.

Archived

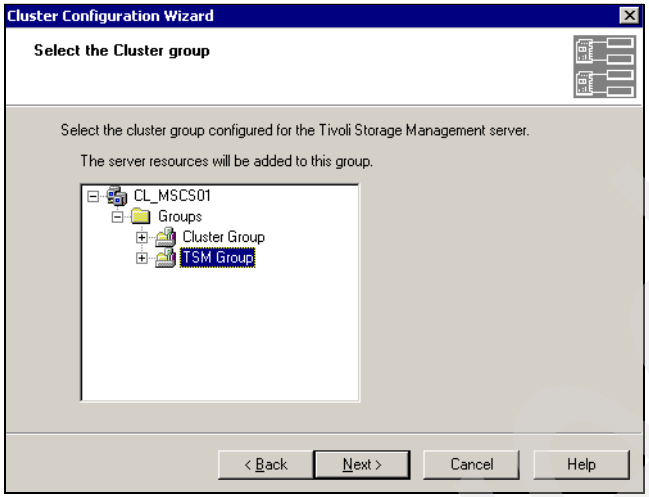

Figure 5-72 Select the cluster group

23. In Figure 5-73 we can configure Tivoli Storage Manager to manage tape failover in the cluster.

> **Note:** MSCS does not support the failover of tape devices. However, Tivoli Storage Manager can manage this type of failover using a shared SCSI bus for the tape devices. Each node in the cluster must contain an additional SCSI adapter card. The hardware and software requirements for tape failover to work and the configuration tasks are described in “Appendix A” of the *Tivoli Storage Manager for Windows Administrator’s Guide*.

Our lab environment does not meet the requirements for tape failover support, so we select **Do not configure TSM to manage tape failover** and then click **Next**.

Archived

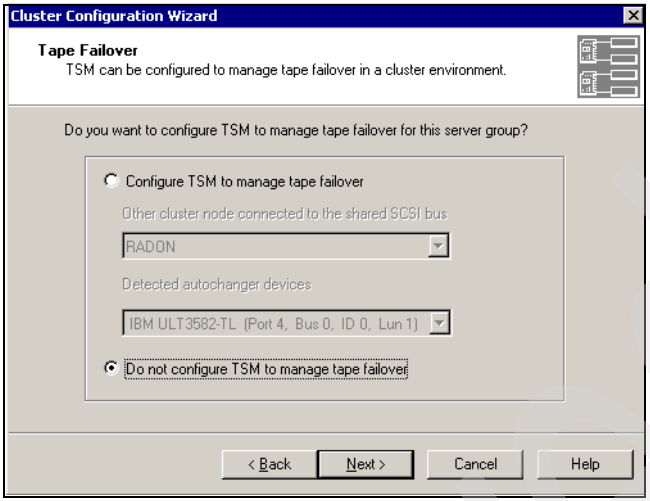

Figure 5-73 Tape failover configuration

24. In Figure 5-74 we enter the IP Address and Subnest Mask that Tivoli Storage Manager virtual server will use in the cluster. This IP address must match the IP address selected in our planning and design worksheets (see Table 5-2 on page 120).

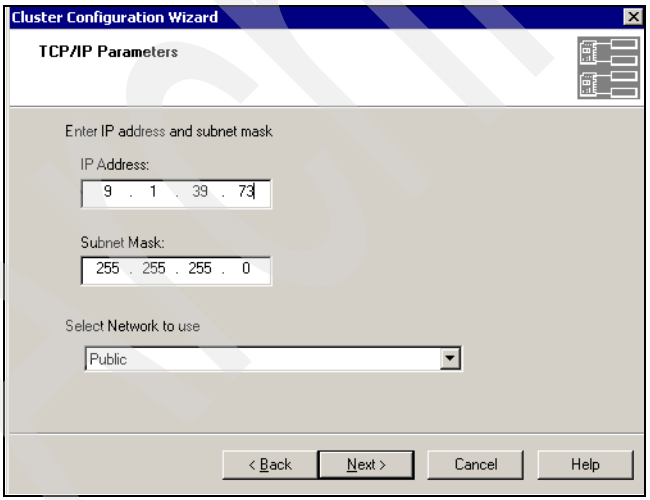

Figure 5-74 IP address

25. In Figure 5-75 we enter the Network name. This must match the network name we selected in our planning and design worksheets (see Table 5-2 on page 120). We enter TSMSRV01 and click **Next**.

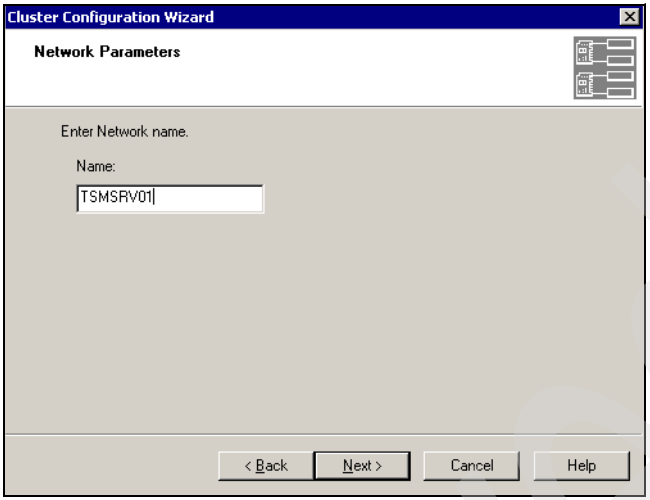

Figure 5-75 Network name

26. On the next menu we check that everything is correct and we click **Finish**. This completes the cluster configuration on RADON (Figure 5-76).

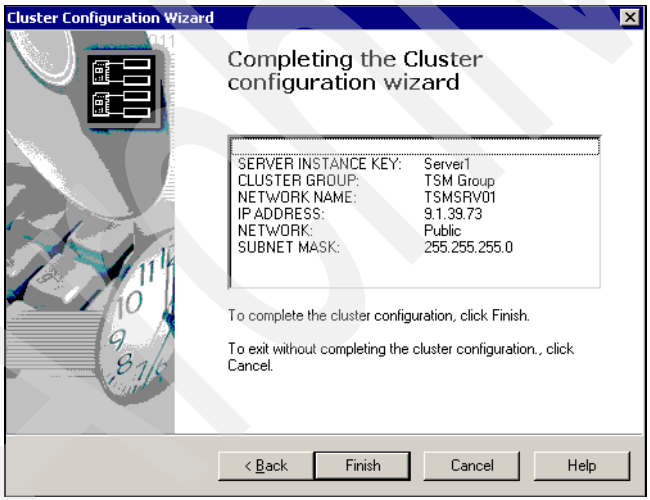

Figure 5-76 Completing the Cluster configuration wizard

27. We receive the following informational message and click **OK** (Figure 5-77).

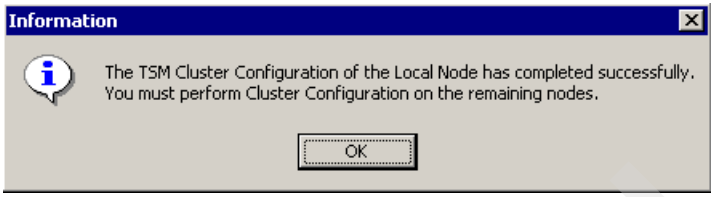

Figure 5-77 End of Tivoli Storage Manager cluster configuration

At this time, we can continue with the initial configuration wizard, to set up devices, nodes, and media. However, for the purpose of this book we will stop here. These tasks are the same ones we would follow in a regular Tivoli Storage Manager server. So, we click **Cancel** when the Device Configuration welcome menu displays.

So far Tivoli Storage Manager server instance is installed and started on RADON. If we open the Tivoli Storage Manager console, we can check that the service is running as shown in Figure 5-78.

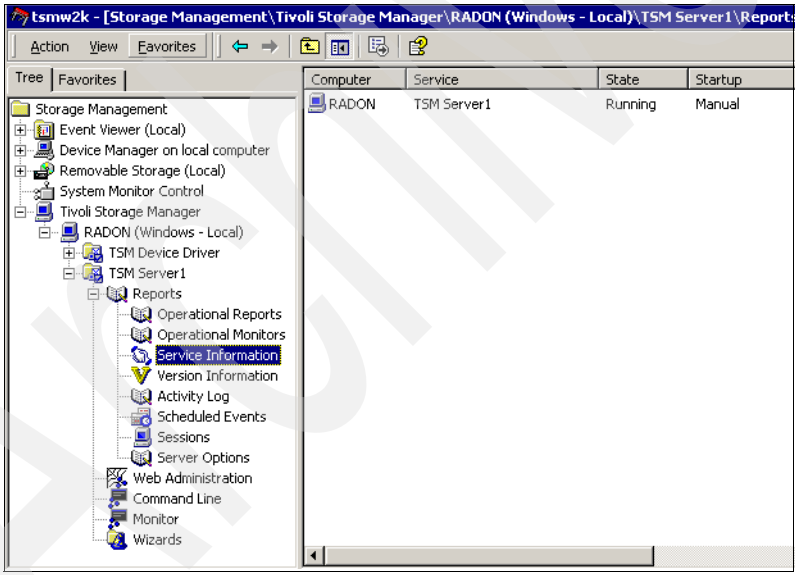

Figure 5-78 Tivoli Storage Manager console

Important: Before starting the initial configuration for Tivoli Storage Manager on the second node, we must stop the instance on the first node.

28. We stop the Tivoli Storage Manager server instance on RADON before going on with the configuration on POLONIUM.

Configuring the second node

In this section we describe how we configure Tivoli Storage Manager on the second node of the MSCS. We follow the same process as for the first node. The only difference is that the Tivoli Storage Manager server instance was already created on the first node. Now the installation will allow the second node to host that server instance.

1. First of all we move the Tivoli Storage Manager cluster group to the second node using the Cluster Administrator. Once moved, the resources should be hosted by POLONIUM, as shown in Figure 5-79:

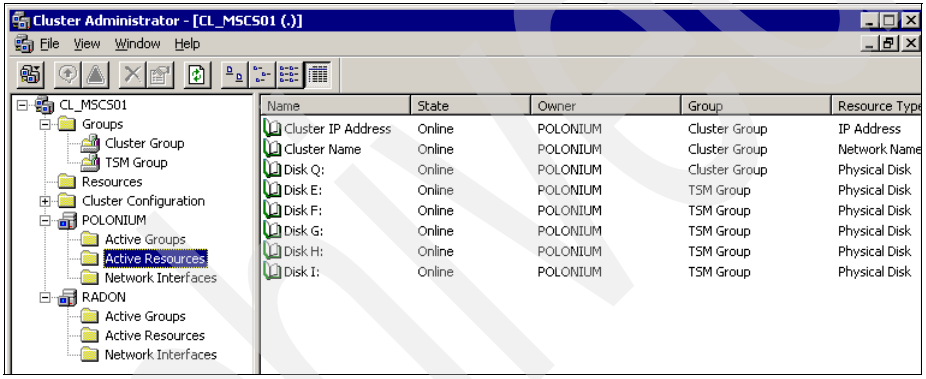

Figure 5-79 Cluster resources

> **Note:** As we can see in Figure 5-79, the IP address and network name resources for the TSM group are not created yet. We still have only disk resources in the TSM resource group. When the configuration ends in POLONIUM, the process will create those resources for us.

2. We open the Tivoli Storage Manager console to start the initial configuration on the second node and follow the same steps (1 to 18) from section “Configuring the first node” on page 123, until we get into the **Cluster Configuration** Wizard in Figure 5-80. We click **Next**.

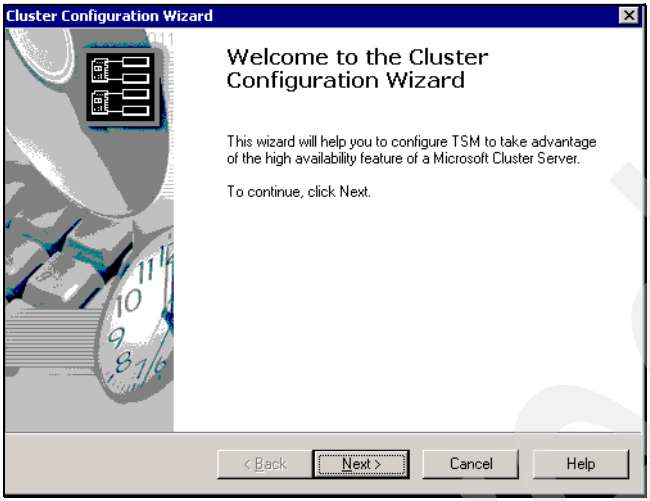

Figure 5-80 Cluster configuration wizard

3. On the Select **Cluster Group** menu in Figure 5-81, we select the same group, the TSM Group, and then we click **Next**.

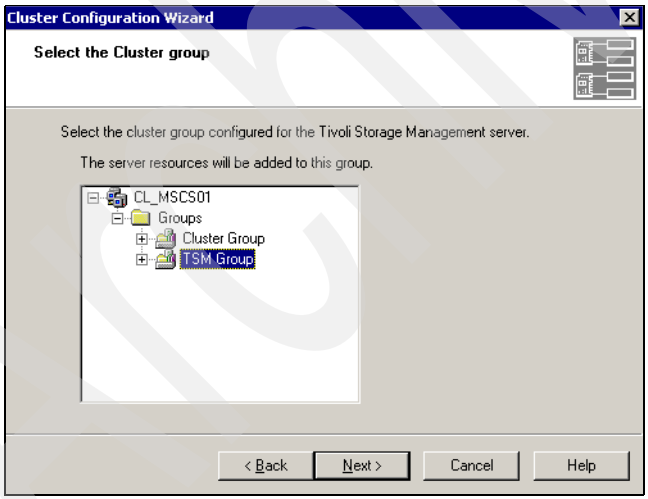

Figure 5-81 Cluster group selection

4. In Figure 5-82 we check that the information reported is correct and then we click **Finish**.

Archived

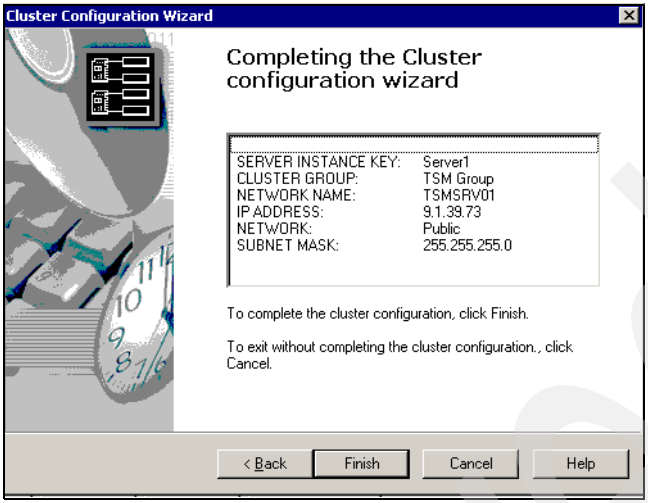

Figure 5-82 Completing the cluster configuration wizard (I)

5. The wizard starts the configuration for the server as shown in Figure 5-83.

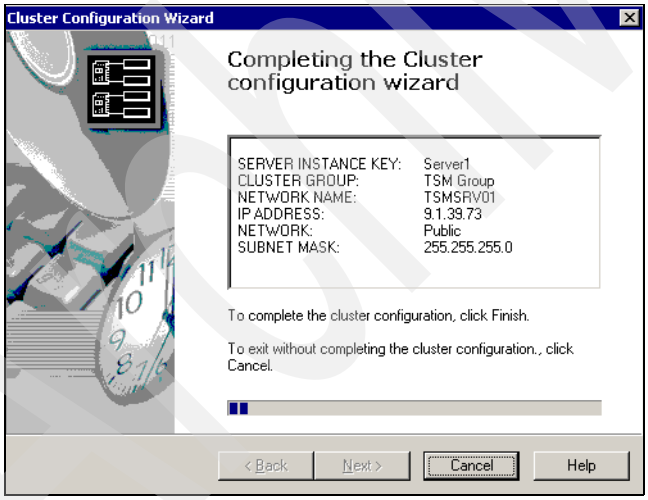

Figure 5-83 Completing the cluster configuration wizard (II)

6. When the configuration is successfully completed, the following message displays. We click **OK** (Figure 5-84).

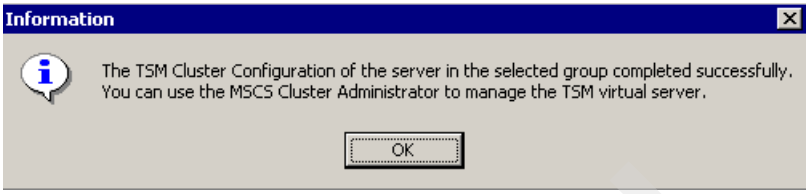

Figure 5-84 Successful installation

Validating the installation

After the wizard completes, we manage the Tivoli Storage Manager virtual server using the MSCS Cluster Administrator.

When we open the MSCS Cluster Administrator to check the results of the process followed on this node, we can see that there are three new resources, as shown in Figure 5-85, created by the wizard:

- TSM Group IP Address: The one we specified in Figure 5-74 on page 137.
- TSM Group Network name: The one specified in Figure 5-75 on page 138.
- TSM Group Server: The Tivoli Storage Manager server instance.

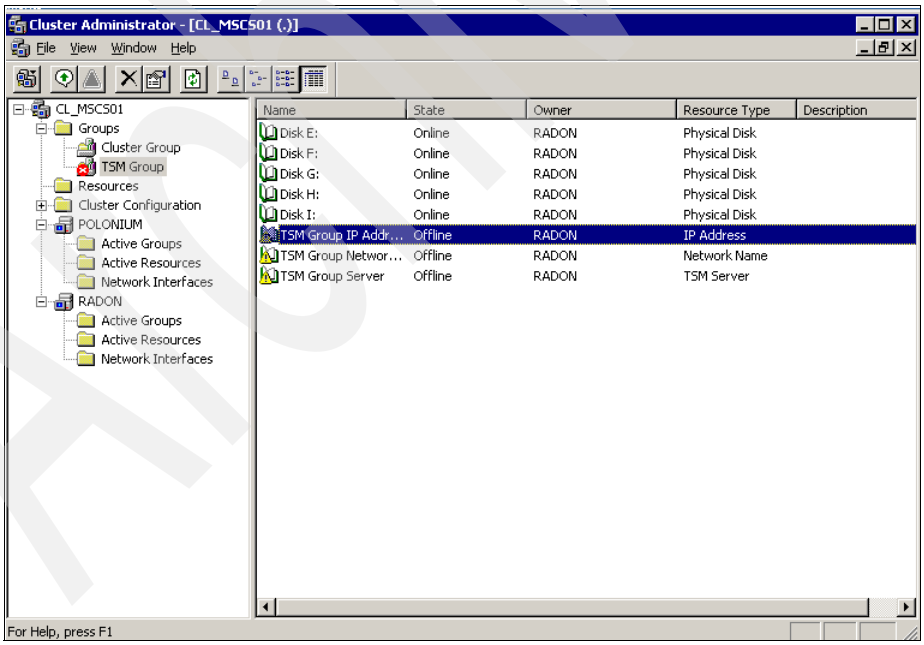

Figure 5-85 Tivoli Storage Manager Group resources

The TSM Group cluster group is offline because the new resources are offline. Now we must bring online every resource on this group, as shown in Figure 5-86.

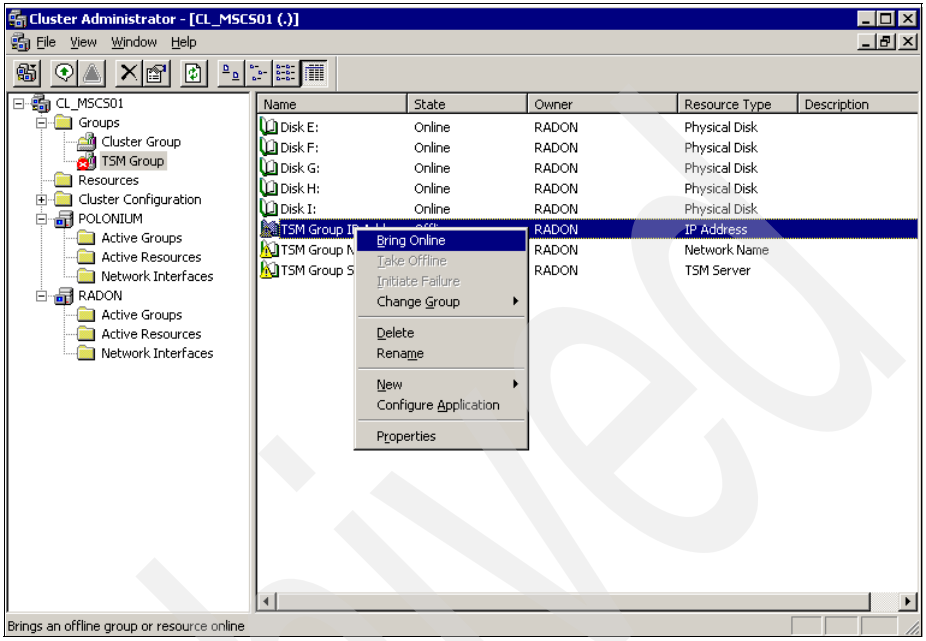

Figure 5-86 Bringing resources online

Archived

In Figure 5-87 we show how to bring online the TSM Group IP Address. The same process should be done for the remaining resources.

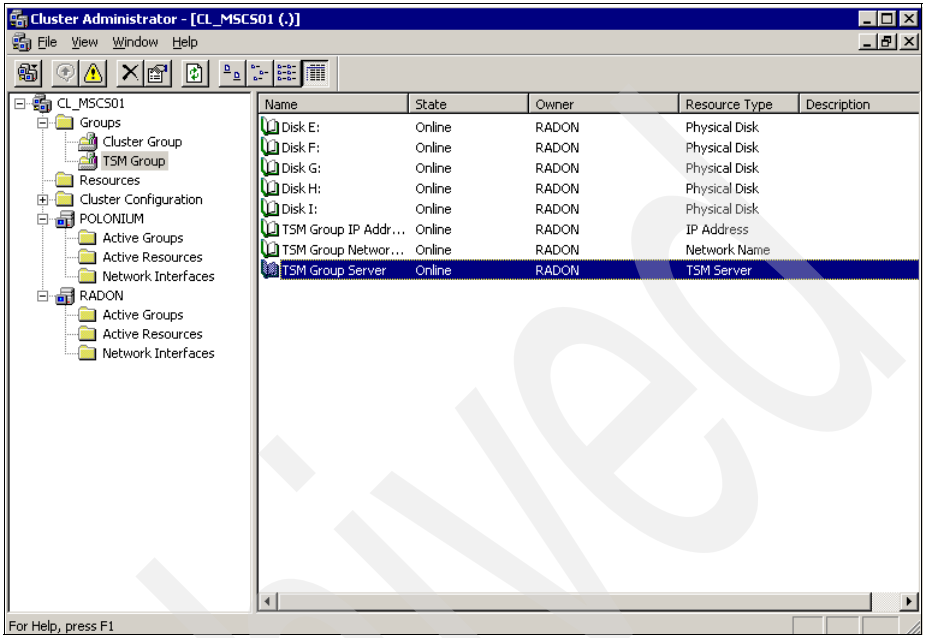

Figure 5-87 Tivoli Storage Manager Group resources online

Now the Tivoli Storage Manager server instance is running on RADON, which is the node which hosts the resources. If we go into the Windows services menu, Tivoli Storage Manager server instance is started as shown in Figure 5-88.

Archived

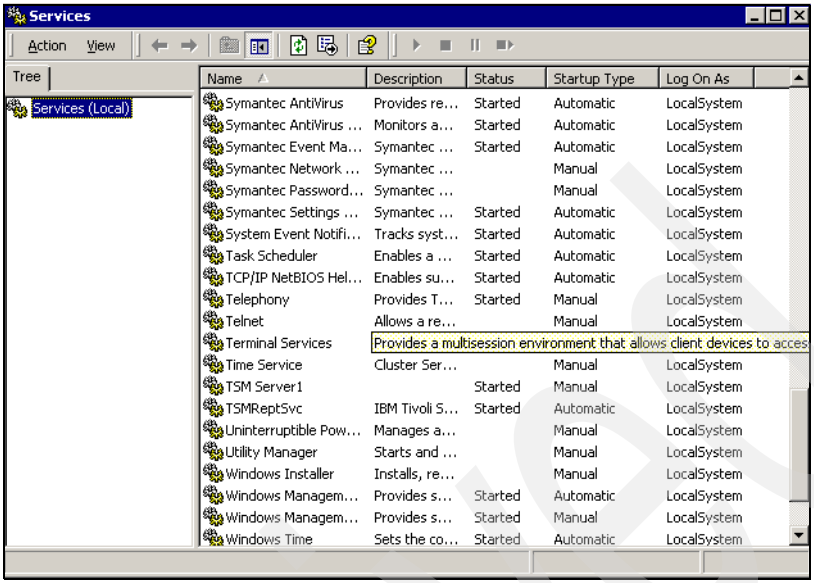

Figure 5-88 Services overview

We move the resources between groups to certify that the configuration is working properly.

> **Important:** Do not forget always to manage the Tivoli Storage Manager server instance using the Cluster Administrator menu, to bring it online or offline.

5.4.3 Testing the Server on Windows 2000

In order to check the high availability of Tivoli Storage Manager server in our lab environment, we must do some testing.

Our objective with these tests is to show how Tivoli Storage Manager on a clustered environment manage its own resources to achieve high availability and how it can respond after certain kinds of failures that affect these shared resources.

Testing client incremental backup using the GUI

Our first test uses the Tivoli Storage Manager GUI to start an incremental backup.

Objective

The objective of this test is to show what happens when a client incremental backup starts using the Tivoli Storage Manager GUI, and suddenly the node which hosts the Tivoli Storage Manager server fails.

Activities

To do this test, we perform these tasks:

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager server. RADON does, as we see in Figure 5-89:

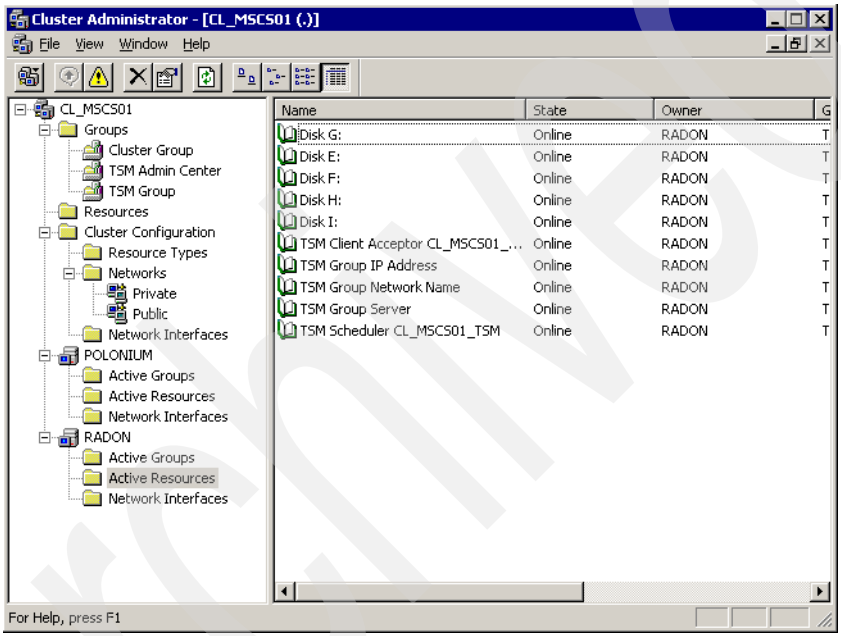

Figure 5-89 Cluster Administrator shows resources on RADON

2. We start an incremental backup from a Windows 2003 Tivoli Storage Manager client with nodename SENEGAL using the GUI. We select the local drives, the System State and the System Services as shown in Figure 5-90.

Archived

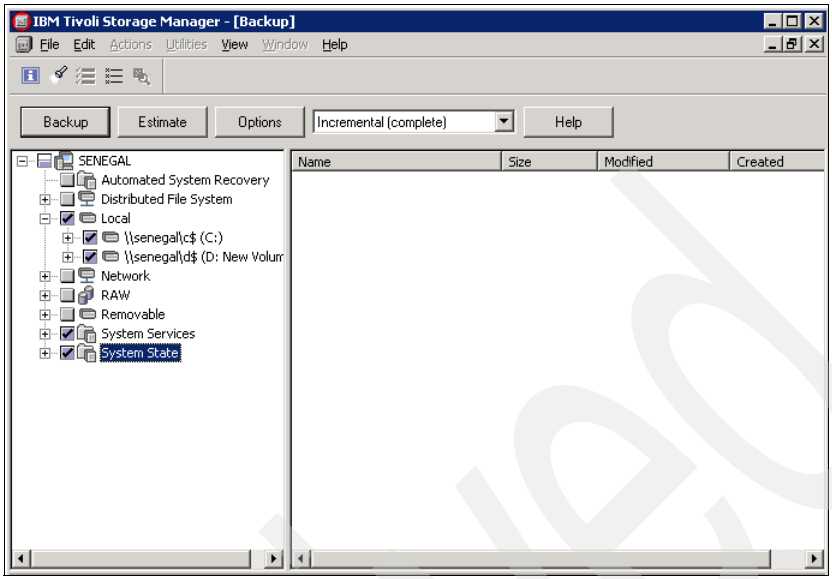

Figure 5-90 Selecting a client backup using the GUI

3. The transfer of files starts as we can see in Figure 5-91.

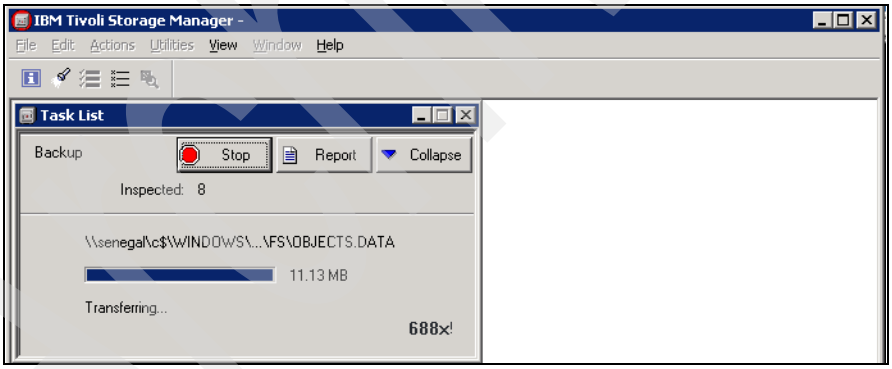

Figure 5-91 Transferring files to the server

Archived

4. While the client is transferring files to the server, we force a failure on RADON, the node that hosts the Tivoli Storage Manager server. In the client, backup is held and we receive a “reopening session” message on the GUI as we can see in Figure 5-92.

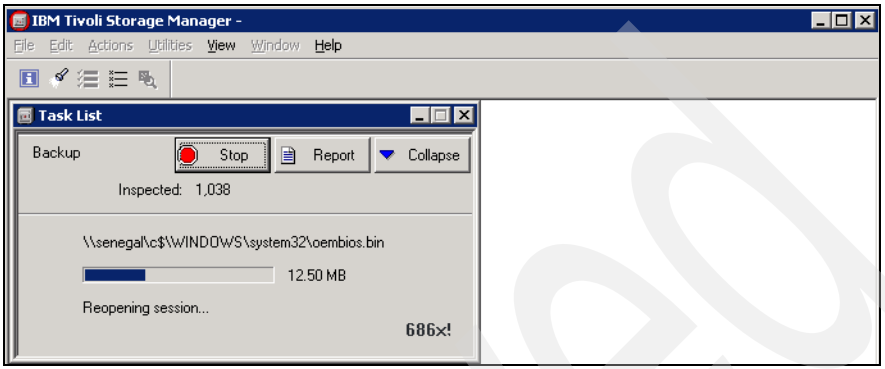

Figure 5-92 Reopening the session

5. When the Tivoli Storage Manager server restarts on POLONIUM, the client continues transferring data to the server (Figure 5-93).

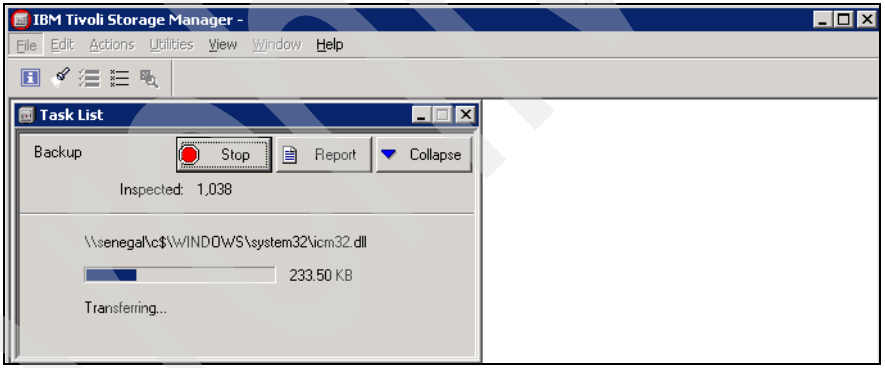

Figure 5-93 Transfer of data goes on when the server is restarted

6. The incremental backup ends successfully.

Results summary

The result of the test shows that when we start a backup from a client and there is an interruption that forces Tivoli Storage Manager server to fail, the backup is held and when the server is up again, the client reopens a session with the server and continues transferring data.

Note: In the test we have just described, we used a disk storage pool as the destination storage pool. We also tested using a tape storage pool as the destination and we got the same results. The only difference is that when the Tivoli Storage Manager server is again up, the tape volume it was using on the first node is unloaded from the drive and loaded again into the second drive, and the client receives a “media wait” message while this process takes place. After the tape volume is mounted, the backup continues, ending successfully.

Testing a scheduled client backup

The second test consists of a scheduled backup.

Objective

The objective of this test is to show what happens when a scheduled client backup is running and suddenly the node which hosts the Tivoli Storage Manager server fails.

Activities

We perform these tasks:

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager cluster group: POLONIUM.
2. We schedule a client incremental backup operation using the Tivoli Storage Manager server scheduler and this time we associate the schedule to a virtual client in our Windows 2000 cluster with nodename CL_MSCS01_SA.
3. A session starts for CL_MSCS01_SA as shown in Example 5-1.

Example 5-1 Activity log when the client starts a scheduled backup

```
01/31/2005 11:28:26  ANR0406I Session 7 started for node CL_MSCS01_SA (WinNT)
(Tcp/Ip radon.tsmw2000.com(1641)). (SESSION: 7)
01/31/2005 11:28:27  ANR2017I Administrator ADMIN issued command: QUERY SESSION
(SESSION: 3)
01/31/2005 11:28:27  ANR0406I Session 8 started for node CL_MSCS01_SA (WinNT)
(Tcp/Ip radon.tsmw2000.com(1644)). (SESSION: 8)
```

4. The client starts sending files to the server as shown in Example 5-2.

Example 5-2 Schedule log file shows the start of the backup on the client

```
Executing scheduled command now.
01/31/2005 11:28:26 Node Name: CL_MSCS01_SA
01/31/2005 11:28:26 Session established with server TSMSRV01: Windows
01/31/2005 11:28:26   Server Version 5, Release 3, Level 0.0
01/31/2005 11:28:26   Server date/time: 01/31/2005 11:28:26  Last access:
01/31/2005 11:25:26
```

```
01/31/2005 11:28:26 --- SCHEDULEREC OBJECT BEGIN INCR_BACKUP 01/31/2005
11:24:11
01/31/2005 11:28:26 Incremental backup of volume '\\cl_mscs01\j$'
01/31/2005 11:28:37 Directory-->  0 \\cl_mscs01\j$\ [Sent]
01/31/2005 11:28:37 Directory-->  0 \\cl_mscs01\j$\Program Files [Sent]
01/31/2005 11:28:37 Directory-->  0 \\cl_mscs01\j$\RECYCLER [Sent]
01/31/2005 11:28:37 Directory-->  0 \\cl_mscs01\j$\System Volume Information
[Sent]
01/31/2005 11:28:37 Directory-->  0 \\cl_mscs01\j$\TSM [Sent]
01/31/2005 11:28:37 Directory-->  0 \\cl_mscs01\j$\TSM_Images [Sent]
01/31/2005 11:28:37 Directory-->  0 \\cl_mscs01\j$\Program Files\IBM [Sent]
```

5. While the client continues sending files to the server, we force POLONIUM to fail. The following sequence occurs:
   a. In the client, the backup is interrupted and errors are received as shown in Example 5-3.

Example 5-3 Error log when the client lost the session

```
01/31/2005 11:29:27 ANS1809W Session is lost; initializing session reopen
procedure.
01/31/2005 11:29:28 ANS1809W Session is lost; initializing session reopen
procedure.
01/31/2005 11:29:47 ANS5216E Could not establish a TCP/IP connection with
address '9.1.39.73:1500'. The TCP/IP error is 'Unknown error' (errno = 10061).
01/31/2005 11:29:47 ANS4039E Could not establish a session with a TSM server or
client agent.  The TSM return code is -50.
01/31/2005 11:30:07 ANS5216E Could not establish a TCP/IP connection with
address '9.1.39.73:1500'. The TCP/IP error is 'Unknown error' (errno = 10061).
01/31/2005 11:30:07 ANS4039E Could not establish a session with a TSM server or
client agent.  The TSM return code is -50.
```

   b. In the Cluster Administrator menu, POLONIUM is not in the cluster and RADON begins to bring the resources online.
   c. After a while the resources are online on RADON.
   d. When the Tivoli Storage Manager server instance resource is online (hosted by RADON), client backup restarts against the disk storage pool as shown on the schedule log file in Example 5-4.

Example 5-4 Schedule log file when backup is restarted on the client

```
01/31/2005 11:29:28 Normal File-->          80,090 \\cl_mscs01\j$\Program
Files\IBM\ISC\AppServer\java\include\jni.h  ** Unsuccessful **
```

01/31/2005 11:29:28 ANS1809W Session is lost; initializing session reopen procedure.

01/31/2005 11:31:23 ... successful

01/31/2005 11:31:23 Retry # 1 Directory--> 0 \\cl_mscs01\j$\Program Files\IBM\ISC\AppServer\installedApps\DefaultNode\wps_facade.ear\wps_facade.war\WEB-INF [Sent]

01/31/2005 11:31:23 Retry # 1 Normal File--> 53 \\cl_mscs01\j$\Program Files\IBM\ISC\AppServer\installedApps\DefaultNode\wps_facade.ear\wps_facade.war\META-INF\MANIFEST.MF [Sent]

---

e. Example 5-5 shows messages that are received on the Tivoli Storage Manager server activity log after restarting.

Example 5-5 Activity log after the server is restarted

```
01/31/2005 11:31:15      ANR2100I Activity log process has started.
01/31/2005 11:31:15      ANR4726I The NAS-NDMP support module has been loaded.
01/31/2005 11:31:15      ANR4726I The Centera support module has been loaded.
01/31/2005 11:31:15      ANR4726I The ServerFree support module has been
loaded.
01/31/2005 11:31:15      ANR2803I License manager started.
01/31/2005 11:31:15      ANR0993I Server initialization complete.
01/31/2005 11:31:15      ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli
is now ready for use.
01/31/2005 11:31:15      ANR2828I Server is licensed to support Tivoli Storage
Manager Basic Edition.
01/31/2005 11:31:15      ANR2560I Schedule manager started.
01/31/2005 11:31:15      ANR8260I Named Pipes driver ready for connection with
clients.
01/31/2005 11:31:15      ANR8200I TCP/IP driver ready for connection with
clients on port 1500.
01/31/2005 11:31:15      ANR8280I HTTP driver ready for connection with clients
on port 1580.
01/31/2005 11:31:15      ANR4747W The web administrative interface is no longer
supported. Begin using the Integrated Solutions Console instead.
01/31/2005 11:31:15      ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK3.DSM
varied online.
01/31/2005 11:31:15      ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK1.DSM
varied online.
01/31/2005 11:31:15      ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK2.DSM
varied online.
01/31/2005 11:31:22      ANR0406I Session 3 started for node CL_MSCS01_SA
(WinNT) (Tcp/Ip tsmsrv01.tsmw2000.com(1784)). (SESSION: 3)
01/31/2005 11:31:22      ANR1639I Attributes changed for node CL_MSCS01_SA: TCP
Address from 9.1.39.188 to 9.1.39.73. (SESSION: 3)
```

```
01/31/2005 11:31:28      ANR8439I SCSI library LIBLTO is ready for operations.
```

6. When the backup ends, the client sends the final statistics messages we show on the schedule log file in Example 5-6.

Example 5-6 Schedule log file shows backup statistics on the client

```
01/31/2005 11:35:50 Successful incremental backup of '\\cl_mscs01\j$'

01/31/2005 11:35:50 --- SCHEDULEREC STATUS BEGIN

01/31/2005 11:35:50 Total number of objects inspected:   17,875

01/31/2005 11:35:50 Total number of objects backed up:   17,875

01/31/2005 11:35:50 Total number of objects updated:         0

01/31/2005 11:35:50 Total number of objects rebound:         0

01/31/2005 11:35:50 Total number of objects deleted:         0

01/31/2005 11:35:50 Total number of objects expired:         0

01/31/2005 11:35:50 Total number of objects failed:          0

01/31/2005 11:35:50 Total number of bytes transferred:     1.14 GB

01/31/2005 11:35:50 Data transfer time:                  24.88 sec

01/31/2005 11:35:50 Network data transfer rate:       48,119.43 KB/sec

01/31/2005 11:35:50 Aggregate data transfer rate:      2,696.75 KB/sec

01/31/2005 11:35:50 Objects compressed by:                   0%

01/31/2005 11:35:50 Elapsed processing time:          00:07:24

01/31/2005 11:35:50 --- SCHEDULEREC STATUS END

01/31/2005 11:35:50 --- SCHEDULEREC OBJECT END INCR_BACKUP
01/31/2005 11:24:11

01/31/2005 11:35:50 ANS1512E Scheduled event 'INCR_BACKUP' failed.
Return code = 12.

01/31/2005 11:35:50 Sending results for scheduled event 'INCR_BACKUP'.

01/31/2005 11:35:50 Results sent to server for scheduled event
'INCR_BACKUP'.
```

Archived

Attention: the scheduled event can end as *failed* with return code = 12 or as *completed* with return code = 8. It depends on the elapsed time until the second node of the cluster brings the resource online. In both cases, however, the backup completes successfully for each drive, as we can see in the first line of the schedule log file in Example 5-6.

Results summary

The test results show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a scheduled backup started from one client is restarted after the failover on the other node of the MSCS.

In the event log, the schedule can display *failed* instead of *completed*, with a return code = 12, if the elapsed time since the first node lost the connection, is too long. In any case, the incremental backup for each drive ends successfully.

Note: In the test we have just described, we used a disk storage pool as the destination storage pool. We also tested using a tape storage pool as destination and we got the same results. The only difference is that when the Tivoli Storage Manager server is again up, the tape volume it was using on the first node is unloaded from the drive and loaded again into the second drive, and the client receives a “media wait” message while this process takes place. After the tape volume is mounted, backup continues and ends successfully.

Testing migration from disk storage pool to tape storage pool

Our third test is a server process: migration from disk storage pool to tape storage pool.

Objective

The objective of this test is to show what happens when a disk storage pool migration process starts on the Tivoli Storage Manager server and the node that hosts the server instance fails.

Activities

For this test, we perform these tasks:

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager cluster group: RADON.

2. We update the disk storage pool (SPD_BCK) high threshold migration to 0. This forces migration of backup versions to its next storage pool, a tape storage pool (SPT_BCK).
3. A process starts for the migration task, and Tivoli Storage Manager prompts the tape library to mount a tape volume as shown in Example 5-7.

Example 5-7 Disk storage pool migration started on server

```
01/31/2005 10:37:36      ANR0984I Process 8 for MIGRATION started in the BACKGROUND at 10:37:36.
(PROCESS: 8)

01/31/2005 10:37:36      ANR1000I Migration process 8 started for storage pool
SPD_BCK automatically, highMig=0, lowMig=0, duration=No. (PROCESS: 8)

01/31/2005 10:37:36      ANR0513I Process 8 opened output volume 020AKKL2.
(PROCESS: 8)

01/31/2005 10:37:45      ANR8330I LTO volume 020AKKL2 is mounted R/W in
drive DRLTO_2 (mt1.0.0.4), status: IN USE. (SESSION: 6)

01/31/2005 10:37:45      ANR8334I        1 matches found. (SESSION: 6)
```

4. While migration is running, we force a failure on RADON. The following sequence occurs:
   a. In the Cluster Administrator menu, RADON is not in the cluster and POLONIUM begins to bring the resources online.
   b. After a few minutes, the resources are online on POLONIUM.
   c. When the Tivoli Storage Manager Server instance resource is online (hosted by POLONIUM), the tape volume is unloaded from the drive. Since the high threshold is still 0, a new migration process is started and the server prompts to mount the same tape volume as shown in Example 5-8.

Example 5-8 Disk storage pool migration started again on the server

```
01/31/2005 10:40:15      ANR0984I Process 2 for MIGRATION started in the
BACKGROUND at 10:40:15. (PROCESS: 2)
01/31/2005 10:40:15      ANR1000I Migration process 2 started for storage pool
SPD_BCK automatically, highMig=0, lowMig=0, duration=No. (PROCESS: 2)
01/31/2005 10:42:05      ANR8439I SCSI library LIBLTO is ready for operations.
01/31/2005 10:42:34      ANR8337I LTO volume 020AKKL2 mounted in drive DRLTO_1
(mt0.0.0.4). (PROCESS: 2)
01/31/2005 10:42:34      ANR0513I Process 2 opened output volume
020AKKL2.(PROCESS: 2)
01/31/2005 10:43:01      ANR8330I LTO volume 020AKKL2 is mounted R/W in drive
DRLTO_1 (mt0.0.0.4), status: IN USE. (SESSION: 2)
```

```
01/31/2005 10:43:01      ANR8334I         1 matches found. (SESSION: 2)
```

Attention: The migration process is not really restarted when the server failover occurs, as we can see by comparing the process numbers for migration between Example 5-7 and Example 5-8. However, the tape volume is unloaded correctly after the failover and loaded again when the new migration process starts on the server.

5. The migration ends successfully, as we show on the activity log taken from the server in Example 5-9.

Example 5-9 Disk storage pool migration ends successfully

```
01/31/2005 10:46:06      ANR1001I Migration process 2 ended for storage pool
SPD_BCK. (PROCESS: 2)
01/31/2005 10:46:06      ANR0986I Process 2 for MIGRATION running in the
BACKGROUND processed 39897 items for a total of 5,455,876,096 bytes with a
completion state of SUCCESS at 10:46:06. (PROCESS: 2)
```

Results summary

The results of our test show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a migration process which is started on the server before the failure, starts again using a new process number when the second node on the MSCS brings the Tivoli Storage Manager server instance online. This is true if the high threshold is still set to the value that caused the migration process to start.

Testing backup from tape storage pool to copy storage pool

In this section we test another internal server process, backup from a tape storage pool to a copy storage pool.

Objective

The objective of this test is to show what happens when a backup storage pool process (from tape to tape) is started on the Tivoli Storage Manager server and the node that hosts the resource fails.

Activities

For this test, we perform these tasks:

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager cluster group: POLONIUM.
2. We run the following command to start a storage pool backup from our primary tape storage pool SPT_BCK to our copy storage pool SPCPT_BCK:

```
ba stg spt_bck spcpt_bck
```

3. A process starts for the storage pool backup task and Tivoli Storage Manager prompts to mount two tape volumes as shown in Example 5-10.

Example 5-10 Starting a backup storage pool process

```
01/31/2005 14:35:09      ANR0984I Process 4 for BACKUP STORAGE POOL started in the BACKGROUND at
14:35:09. (SESSION: 16, PROCESS: 4)

01/31/2005 14:35:09      ANR2110I BACKUP STGPOOL started as process 4.
(SESSION: 16, PROCESS: 4)

01/31/2005 14:35:09      ANR1210I Backup of primary storage pool SPT_BCK to
copy storage pool SPCPT_BCK started as process 4. (SESSION: 16,
PROCESS: 4)

01/31/2005 14:35:09      ANR1228I Removable volume 020AKKL2 is required for
storage pool backup. (SESSION: 16, PROCESS: 4)

01/31/2005 14:35:43      ANR8337I LTO volume 020AKKL2 mounted in drive
DRLTO_1 (mt0.0.0.4). (SESSION: 16, PROCESS: 4)

01/31/2005 14:35:43      ANR0512I Process 4 opened input volume 020AKKL2.
(SESSION: 16, PROCESS: 4)

01/31/2005 14:36:12      ANR8337I LTO volume 021AKKL2 mounted in drive
DRLTO_2 (mt1.0.0.4). (SESSION: 16, PROCESS: 4)

01/31/2005 14:36:12      ANR1340I Scratch volume 021AKKL2 is now defined in
storage pool SPCPT_BCK. (SESSION: 16, PROCESS: 4)

01/31/2005 14:36:12      ANR0513I Process 4 opened output volume
021AKKL2.(SESSION: 16, PROCESS: 4)
```

4. While the process is started and the two tape volumes are mounted on both drives, we force a failure on POLONIUM and the following sequence occurs:
   a. In the Cluster Administrator menu, POLONIUM is not in the cluster and RADON begins to bring the resources online.
   b. After a few minutes the resources are online on RADON.
   c. When the Tivoli Storage Manager Server instance resource is online (hosted by RADON), the tape library dismounts the tape volumes from the drives. However, in the activity log there is no process started and there is no track of the process that was started before the failure in the server, as we can see in Example 5-11.

Archived

Example 5-11 After restarting the server the storage pool backup does not restart

```
01/31/2005 14:37:54      ANR4726I The NAS-NDMP support module has been loaded.
01/31/2005 14:37:54      ANR4726I The Centera support module has been loaded.
01/31/2005 14:37:54      ANR4726I The ServerFree support module has been
loaded.
01/31/2005 14:37:54      ANR2803I License manager started.
01/31/2005 14:37:54      ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK1.DSM
varied online.
01/31/2005 14:37:54      ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK3.DSM
varied online.
01/31/2005 14:37:54      ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK2.DSM
varied online.
01/31/2005 14:37:54      ANR2828I Server is licensed to support Tivoli Storage
Manager Basic Edition.
01/31/2005 14:37:54      ANR8260I Named Pipes driver ready for connection with
clients.
01/31/2005 14:37:54      ANR8200I TCP/IP driver ready for connection with
clients on port 1500.
01/31/2005 14:37:54      ANR8280I HTTP driver ready for connection with clients
on port 1580.
01/31/2005 14:37:54      ANR4747W The web administrative interface is no longer
supported. Begin using the Integrated Solutions Console instead.
01/31/2005 14:37:54      ANR0993I Server initialization complete.
01/31/2005 14:37:54      ANR2560I Schedule manager started.
01/31/2005 14:37:54      ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli
is now ready for use.
01/31/2005 14:38:04      ANR8779E Unable to open drive mt0.0.0.4, error
number=170.
01/31/2005 14:38:24      ANR2017I Administrator ADMIN issued command: QUERY
PROCESS(SESSION: 3)
01/31/2005 14:38:24      ANR0944E QUERY PROCESS: No active processes found.
(SESSION: 3)
```

Attention: When the server restarts on the other node, an error message is received on the activity log where Tivoli Storage Manager tells it is unable to open one drive as we can see in Example 5-11. However, both tapes are unloaded correctly from the two drives.

5. The backup storage pool process does not restart again unless we start it manually.
6. If the backup storage pool process sent enough data before the failure so that the server was able to commit the transaction into the database, when the Tivoli Storage Manager server starts again in the second node, those files already backed up into the copy storage pool tape volume and committed in the server database, are valid copied versions.

However, there are still files not copied from the primary tape storage pool. If we want to be sure that the server copies all the files from this primary storage pool, we need to repeat the command. Those files committed as copied in the database will not be copied again.

This happens both using roll-forward recovery log mode as well as normal recovery log mode. In our particular test, there was no tape volume in the copy storage pool before starting the backup storage pool process in the first node, because it was the first time we used this command.

If we look at Example 5-10 on page 157, there is an informational message in the activity log telling us that the scratch volume 021AKKL2 is now defined in the copy storage pool.

- When the server is again online in the second node, we run the command:

  ```
  q content 021AKKL2
  ```

  The command reports information. This means some information was committed before the failure.

- To be sure that the server copies the rest of the files, we start a new backup from the same primary storage pool, SPT_BCK to the copy storage pool, SPCPT_BCK.

- When the backup ends, we use the following commands:

  ```
  q occu stg=spt_bck
  q occu stg=spcpt_bck
  ```

  Both commands should report the same information it there are no more primary storage pools.

7. If the backup storage pool task did not process enough data to commit the transaction into the database, when the Tivoli Storage Manager server starts again in the second node, those files copied in the copy storage pool tape volume before the failure are not recorded in the Tivoli Storage Manager server database. So, if we start a new backup storage pool task, they will be copied again.

   If the tape volume used for the copy storage pool before the failure was taken from the scratch pool in the tape library (as in our case), it is given back to scratch status in the tape library.

   If the tape volume used for the copy storage pool before the failure had already data belonging to back up storage pool tasks from other days, the tape volume is kept in the copy storage pool but the new information written on it is not valid.

   If we want to be sure that the server copies all the files from this primary storage pool, we need to repeat the command.

This happens both using roll-forward recovery log mode as well as normal recovery log mode.

In a test we made where the transaction was not committed into the database, also with no tape volumes in the copy storage pool, the server also mounted a scratch volume that was defined in the copy storage pool. However, when the server started on the second node after the failure, the tape volume was deleted from the copy storage pool.

Results summary

The results of our test show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a backup storage pool process (from tape to tape) started on the server before the failure, does not restart when the second node on the MSCS brings the Tivoli Storage Manager server instance online.

Both tapes are correctly unloaded from the tape drives when the Tivoli Storage Manager server is again online, but the process is not restarted unless we run the command again.

Depending on the amount of data already sent when the task failed (if it was committed to the database or not), the files backed up into the copy storage pool tape volume before the failure, will either be reflected on the database, or not.

If enough information was copied to the copy storage pool tape volume so that the transaction was committed before the failure, when the server restarts in the second node, the information is recorded in the database and the files figure as valid copies.

If the transaction was not committed, there is no information in the database about the process and the files backed up into the copy storage pool before the failure, will need to be copied again.

This situation happens either if the recovery log is set to roll-forward mode or it is set to normal mode.

In any of the cases, to be sure that all information is copied from the primary storage pool to the copy storage pool, we should repeat the command.

There is no difference between a scheduled backup storage pool process or a manual process using the administrative interface. In our lab we tested both methods and the results were the same.

Testing server database backup

The following test is a server database backup.

Objective

The objective of this test is to show what happens when a Tivoli Storage Manager server database backup process starts on the Tivoli Storage Manager server and the node that hosts the resource fails.

Activities

For this test, we perform these tasks:

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager cluster group: RADON.
2. We run the following command to start a full database backup:

   ```
   ba db t=full devc=cllto_1
   ```

3. A process starts for database backup and Tivoli Storage Manager prompts to mount a scratch tape volume as shown in Example 5-12.

Example 5-12 Starting a database backup on the server

```
01/31/2005 14:51:50      ANR0984I Process 4 for DATABASE BACKUP started in the BACKGROUND at 14:51:50.
(SESSION: 11, PROCESS: 4)

01/31/2005 14:51:50      ANR2280I Full database backup started as process 4.
(SESSION: 11, PROCESS: 4)

01/31/2005 14:51:59      ANR2017I Administrator ADMIN issued command:
QUERY PROCESS (SESSION: 11)

01/31/2005 14:52:11      ANR2017I Administrator ADMIN issued command:
QUERY PROCESS (SESSION: 11)

01/31/2005 14:52:18      ANR8337I LTO volume 022AKKL2 mounted in drive
DRLTO_1 (mt0.0.0.4). (SESSION: 11, PROCESS: 4)

01/31/2005 14:52:18      ANR0513I Process 4 opened output volume 022AKKL2.
(SESSION: 11, PROCESS: 4)

01/31/2005 14:52:18      ANR2017I Administrator ADMIN issued command:
QUERY PROCESS (SESSION: 11)

01/31/2005 14:52:21      ANR1360I Output volume 022AKKL2 opened (sequence
number 1). (SESSION: 11, PROCESS: 4)

01/31/2005 14:52:23      ANR4554I Backed up 7424 of 14945 database pages.
(SESSION: 11, PROCESS: 4)
```

Archived

4. While the backup is running we force a failure on RADON. The following sequence occurs:
   a. In the Cluster Administrator menu, RADON is not in the cluster and POLONIUM begins to bring the resources online.
   b. After a few minutes the resources are online on POLONIUM.
   c. When the Tivoli Storage Manager Server instance resource is online (hosted by POLONIUM), the tape volume is unloaded from the drive by the tape library automatic system. There is an error message, ANR8779E, where the server reports it is unable to open the drive where the tape volume was mounted before the failure, but there is no process started on the server for any database backup, as we can see in Example 5-13.

Example 5-13 After the server is restarted database backup does not restart

```
01/31/2005 14:53:58      ANR4726I The NAS-NDMP support module has been loaded.
01/31/2005 14:53:58      ANR4726I The Centera support module has been loaded.
01/31/2005 14:53:58      ANR4726I The ServerFree support module has been
loaded.
01/31/2005 14:53:58      ANR0984I Process 1 for EXPIRATION started in the
BACKGROUND at 14:53:58. (PROCESS: 1)
01/31/2005 14:53:58      ANR2803I License manager started.
01/31/2005 14:53:58      ANR0811I Inventory client file expiration started as
process 1. (PROCESS: 1)
01/31/2005 14:53:58      ANR8260I Named Pipes driver ready for connection with
clients.
01/31/2005 14:53:58      ANR2560I Schedule manager started.
01/31/2005 14:53:58      ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli
is now ready for use.
01/31/2005 14:53:58      ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK1.DSM
varied online.
01/31/2005 14:53:58      ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK3.DSM
varied online.
01/31/2005 14:53:58      ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK2.DSM
varied online.
01/31/2005 14:53:59      ANR2828I Server is licensed to support Tivoli Storage
Manager Basic Edition.
01/31/2005 14:53:59      ANR8280I HTTP driver ready for connection with clients
on port 1580.
01/31/2005 14:53:59      ANR8200I TCP/IP driver ready for connection with
clients on port 1500.
01/31/2005 14:54:09      ANR8779E Unable to open drive mt0.0.0.4, error
number=170.
01/31/2005 14:54:46      ANR8439I SCSI library LIBLTO is ready for operations.
01/31/2005 14:56:36      ANR2017I Administrator ADMIN issued command: QUERY
PROCESS (SESSION: 3)
```

```
01/31/2005 14:56:36      ANR0944E QUERY PROCESS: No active processes
found.(SESSION: 3)
```

5. We query the volume history looking for information about the database backup volumes, using the command:

   ```
   q volh t=dbb
   ```

   However, there is no record for the tape volume 022AKKL2, as we can see in Example 5-14.

Example 5-14 Volume history for database backup volumes

```
tsm: TSMSRV01>q volh t=dbb

       Date/Time: 01/30/2005 13:10:05
     Volume Type: BACKUPFULL
   Backup Series: 3
Backup Operation: 0
      Volume Seq: 1
    Device Class: CLLT0_1
     Volume Name: 020AKKL2
 Volume Location:
         Command:

tsm: TSMSRV01>
```

6. However, if we query the library inventory, using the command:

   ```
   q libvol
   ```

   The tape volume is reported as private and last used as dbbackup, as we see in Example 5-15.

Example 5-15 Library volumes

```
tsm: TSMSRV01>q libvol

Library Name  Volume Name  Status      Owner       Last Use   Home     Device
                                                              Element  Type
------------  -----------  ----------  ----------  ---------  -------  ------
LIBLT0        020AKKL2     Private     TSMSRV01    DbBackup   4,096    LT0
LIBLT0        021AKKL2     Private     TSMSRV01    Data       4,097    LT0
LIBLT0        022AKKL2     Private     TSMSRV01    DbBackup   4,098    LT0
LIBLT0        023AKKL2     Private     TSMSRV01    Data       4,099    LT0
LIBLT0        026AKKL2     Private     TSMSRV01               4,102    LT0
LIBLT0        027AKKL2     Private     TSMSRV01               4,116    LT0
LIBLT0        028AKKL2     Private     TSMSRV01               4,104    LT0
LIBLT0        029AKKL2     Private     TSMSRV01               4,105    LT0
LIBLT0        030AKKL2     Private     TSMSRV01               4,106    LT0
```

Archived

```
LIBLTO        031AKKL2     Private     TSMSRV01                4,107    LTO
LIBLTO        032AKKL2     Private     TSMSRV01                4,108    LTO
LIBLTO        033AKKL2     Private     TSMSRV01                4,109    LTO
LIBLTO        034AKKL2     Private     TSMSRV01                4,110    LTO
LIBLTO        036AKKL2     Private     TSMSRV01                4,112    LTO
LIBLTO        037AKKL2     Private     TSMSRV01                4,113    LTO
LIBLTO        038AKKL2     Private     TSMSRV01                4,114    LTO
LIBLTO        039AKKL2     Private     TSMSRV01                4,115    LTO

tsm: TSMSRV01>
```

7. We update the library inventory for 022AKKL2 to change its status to scratch, using the command:

   ```
   upd libvol liblto 022akkl2 status=scratch
   ```

8. We repeat the database backup command, checking that it ends successfully.

Results summary

The results of our test show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a database backup process started on the server before the failure, does not restart when the second node on the MSCS brings the Tivoli Storage Manager server instance online.

The tape volume is correctly unloaded from the tape drive where it was mounted when the Tivoli Storage Manager server is again online, but the process does not end successfully. It is not restarted unless we run the command.

There is no difference between a scheduled process or a manual process using the administrative interface.

> **Important:** The tape volume used for the database backup before the failure is not useful. It is reported as a private volume in the library inventory but it is not recorded as valid backup in the volume history file. It is necessary to update the tape volume in the library inventory to scratch and start again a new database backup process.

Testing inventory expiration

In this section we test another server task: an inventory expiration process.

Objective

The objective of this test is to show what happens when Tivoli Storage Manager server is running the inventory expiration process and the node that hosts the server instance fails.

Activities

For this test, we perform these tasks:

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager cluster group: POLONIUM.
2. We run the following command to start an inventory expiration process:

   ```
   expire inventory
   ```

3. A process starts for inventory expiration as shown in Example 5-16.

Example 5-16 Starting inventory expiration

```
02/01/2005 12:35:26      ANR0984I Process 3 for EXPIRE INVENTORY started in the
BACKGROUND at 12:35:26. (SESSION: 13, PROCESS: 3)
02/01/2005 12:35:26      ANR0811I Inventory client file expiration started as
process 3. (SESSION: 13, PROCESS: 3)
02/01/2005 12:35:26      ANR4391I Expiration processing node RADON, filespace
\\radon\c$, fsId 2, domain STANDARD, and management class DEFAULT - for BACKUP
type files. (SESSION: 13, PROCESS: 3)
02/01/2005 12:35:26      ANR4391I Expiration processing node RADON, filespace
SYSTEM OBJECT, fsId 3, domain STANDARD, and management class DEFAULT - for
BACKUP type files. (SESSION: 13, PROCESS: 3)
02/01/2005 12:35:27      ANR2017I Administrator ADMIN issued command: QUERY
PROCESS (SESSION: 13)
02/01/2005 12:35:27      ANR4391I Expiration processing node POLONIUM,
filespace SYSTEM OBJECT, fsId 1, domain STANDARD, and management class DEFAULT
- for BACKUP type files. (SESSION: 13, PROCESS: 3)
02/01/2005 12:35:30      ANR4391I Expiration processing node POLONIUM,
filespace \\polonium\c$, fsId 3, domain STANDARD, and management class DEFAULT
- for BACKUP type files. (SESSION: 13, PROCESS: 3)
```

4. While Tivoli Storage Manager server is expiring objects, we force a failure on the node that hosts the server instance. The following sequence occurs:
   a. In the Cluster Administrator menu POLONIUM is not in the cluster and RADON begins to bring the resources online.
   b. After a few minutes the resources are online on RADON.
   c. When the Tivoli Storage Manager Server instance resource is online (hosted by RADON), the inventory expiration process is not started any more. There are no errors in the activity log, just the process is not running. The last message received from the Tivoli Storage Manager server before the failure, as shown in Example 5-17, tells us it was expiring objects for POLONIUM node. After that, the server starts on the other node and there is no process started.

Example 5-17 No inventory expiration process after the failover

```
02/01/2005 12:35:30      ANR4391I Expiration processing node POLONIUM,
filespace \\polonium\c$, fsId 3, domain STANDARD, and management class DEFAULT
- for BACKUP type files. (SESSION: 13, PROCESS: 3)
02/01/2005 12:36:10      ANR2100I Activity log process has started.
02/01/2005 12:36:10      ANR4726I The NAS-NDMP support module has been loaded.
02/01/2005 12:36:10      ANR4726I The Centera support module has been loaded.
02/01/2005 12:36:10      ANR4726I The ServerFree support module has been
loaded.
02/01/2005 12:36:11      ANR2803I License manager started.
02/01/2005 12:36:11      ANR0993I Server initialization complete.
02/01/2005 12:36:11      ANR8260I Named Pipes driver ready for connection with
clients.
02/01/2005 12:36:11      ANR2560I Schedule manager started.
02/01/2005 12:36:11      ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli
is now ready for use.
02/01/2005 12:36:11      ANR8200I TCP/IP driver ready for connection with
clients on port 1500.
02/01/2005 12:36:11      ANR8280I HTTP driver ready for connection with clients
on port 1580.
02/01/2005 12:36:11      ANR2828I Server is licensed to support Tivoli Storage
Manager Basic Edition.
02/01/2005 12:36:11      ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK3.DSM
varied online.
02/01/2005 12:36:11      ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK2.DSM
varied online.
02/01/2005 12:36:23      ANR8439I SCSI library LIBLTO is ready for operations.
02/01/2005 12:36:58      ANR0407I Session 3 started for administrator ADMIN
(WinNT) (Tcp/Ip radon.tsmw2000.com(1415)). (SESSION: 3)
02/01/2005 12:37:37      ANR2017I Administrator ADMIN issued command: QUERY
PROCESS (SESSION: 3)
02/01/2005 12:37:37      ANR0944E QUERY PROCESS: No active processes found.
(SESSION: 3)
```

5. If we want to start the process again, we just have to run the same command. Tivoli Storage Manager server run the process and it ends successfully, as shown in Example 5-18.

Example 5-18 Starting inventory expiration again

```
02/01/2005 12:37:43      ANR2017I Administrator ADMIN issued command: EXPIRE INVENTORY  (SESSION: 3)
```

02/01/2005 12:37:43 ANR0984I Process 1 for EXPIRE INVENTORY started in the BACKGROUND at 12:37:43. (SESSION: 3, PROCESS: 1)

02/01/2005 12:37:43 ANR0811I Inventory client file expiration started as process 1. (SESSION: 3, PROCESS: 1)

02/01/2005 12:37:43 ANR4391I Expiration processing node POLONIUM, filespace \\polonium\c$, fsId 3, domain STANDARD, and management class DEFAULT - for BACKUP type files. (SESSION: 3, PROCESS: 1)

02/01/2005 12:37:43 ANR4391I Expiration processing node POLONIUM, filespace \\polonium\c$, fsId 3, domain STANDARD, and management class DEFAULT - for BACKUP type files. (SESSION: 3, PROCESS: 1)

02/01/2005 12:37:44 ANR0812I Inventory file expiration process 1 completed: examined 117 objects, deleting 115 backup objects, 0 archive objects, 0 DB backup volumes, and 0 recovery plan files. 0 errors were encountered. (SESSION: 3, PROCESS: 1)

02/01/2005 12:37:44 ANR0987I Process 1 for EXPIRE INVENTORY running in the BACKGROUND processed 115 items with a completion state of SUCCESS at 12:37:44. (SESSION: 3, PROCESS: 1)

Results summary

The results of our test show that after a failure on the node that hosts the Tivoli Storage Manager server instance, an inventory expiration process started on the server before the failure, does not restart when the second node on the MSCS brings the Tivoli Storage Manager server instance online.

There is no error inside the Tivoli Storage Manager server database and we can restart the process again when the server is online.

5.5 Configuring ISC for clustering on Windows 2000

In 5.3.4, “Installation of the Administration Center” on page 92 we already described how we installed the Administration Center components on each node of the MSCS.

In this section we describe the method we use to configure the Integrated Solution Console (ISC) as a clustered application on our MSCS Windows 2000. We need to create two new resources for the ISC services, in the cluster group where the shared disk used to install the code is located.

1. First we check that both nodes are again up and the two ISC services are stopped on them.
2. We open the Cluster Administrator menu and select the TSM Admin Center cluster group, the group that the shared disk j: belongs to. Then we select

New →**Resource**, to create a new generic service resource as shown in Figure 5-94.

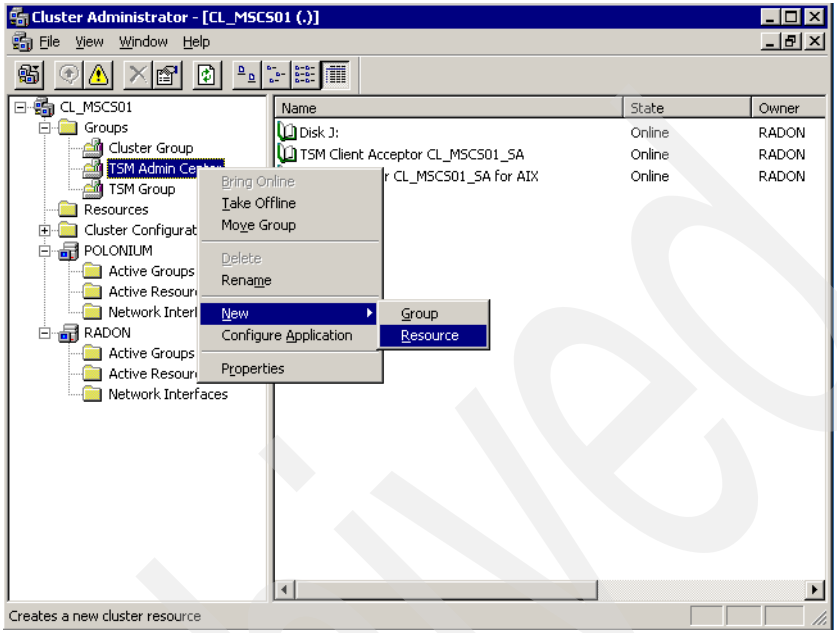

Figure 5-94 Defining a new resource for IBM WebSphere application server

3. We want to create a Generic Service resource related to the IBM WebSphere Application Server. We select a name for the resource and we choose **Generic Service** as resource type in Figure 5-95 and we click **Next**.

Archived

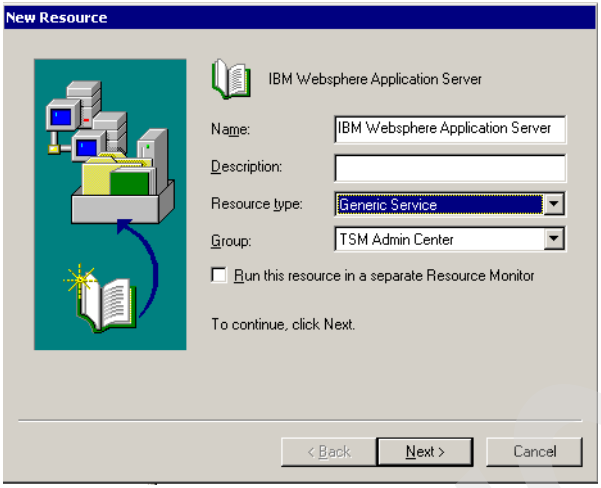

Figure 5-95 Specifying a resource name for IBM WebSphere application server

4. We leave both nodes as possible owners for the resource as shown in Figure 5-96 and we click **Next**.

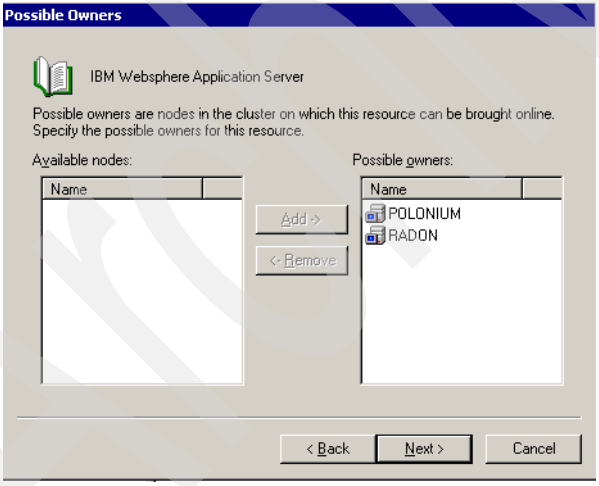

Figure 5-96 Possible owners for the IBM WebSphere application server resource

Archived

5. We select **Disk J** and **IP address** as dependencies for this resource and we click **Next** as shown in Figure 5-97.

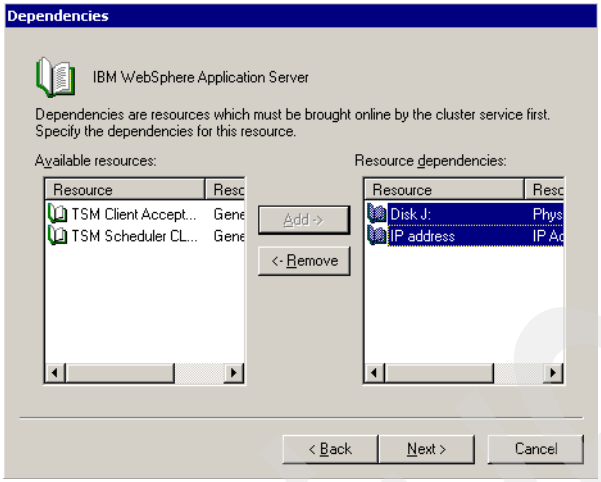

Figure 5-97 Dependencies for the IBM WebSphere application server resource

Important: The cluster group where the ISC services are defined must have an IP address resource. When the generic service is created using the Cluster Administrator menu, we use this IP address as dependency for the resource to be brought online. In this way, when we start a Web browser to connect to the WebSphere Application server, we use the IP for the cluster resource, instead of the local IP address for each node.

6. We type the real name of the IBM WebSphere Application Server service in Figure 5-98.

Archived

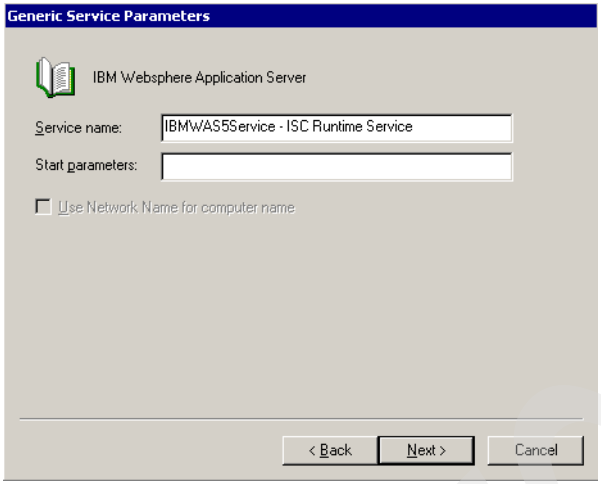

Figure 5-98 Specifying the same name for the service related to IBM WebSphere

Attention: Make sure to specify the correct name in Figure 5-98. In the Windows services menu, the name displayed for the service is not the real service name for it. Therefore, right-click the service and select **Properties** to check the service name for Windows.

7. We do not use any Registry key values to be replicated between nodes. We click **Next** in Figure 5-99.

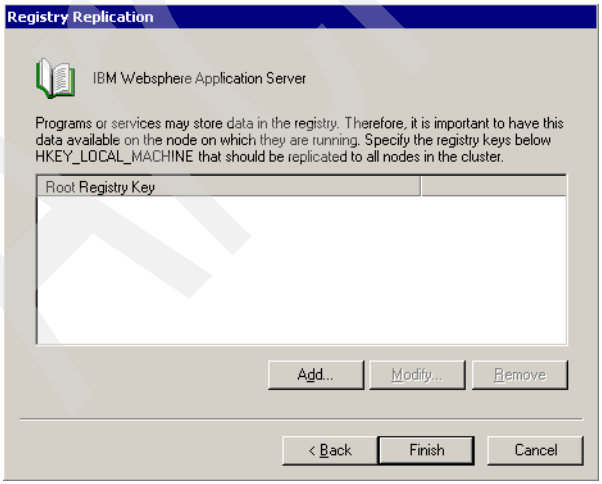

Figure 5-99 Registry replication values

Archived

8. The creation of the resource is successful as we can see in Figure 5-100. We click **OK** to finish.

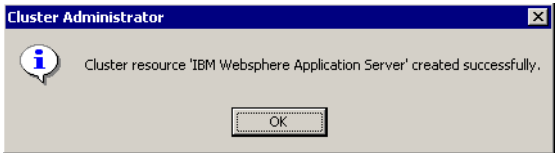

Figure 5-100 Successful creation of the generic resource

9. Now we bring this resource online.
10. The next task is the definition of a new Generic Service resource related to the ISC Help Service. We proceed using the same process as for the IBM WebSphere Application server.
11. We use ISC Help services as the name of the resource as shown in Figure 5-101.

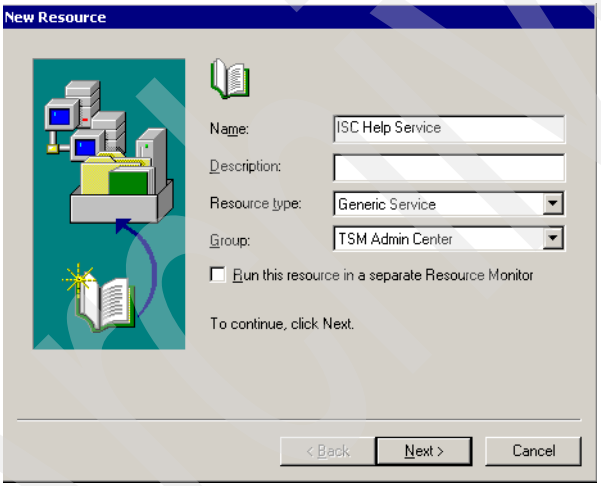

Figure 5-101 Selecting the resource name for ISC Help Service

12. As possible owners we select both nodes, in the dependencies menu we select the IBM WebSphere Application Server resource, and we do not use any Registry keys replication.
13. After the successful installation of the service, we bring it online using the Cluster Administrator menu.

Archived

14. At this moment both services are online in POLONIUM, the node that hosts the resources. To check that the configuration works correctly we proceed to move the resources to RADON. Both services are now started in this node and stopped in POLONIUM.

5.5.1 Starting the Administration Center console

After the installation and configuration of ISC and Administration Center components in both nodes we are ready to start the Administration Center console to manage any Tivoli Storage Manager server.

We use the IP address related to the TSM Admin Center cluster group, which is the group where the ISC shared installation path is located.

1. In order to start an administrator Web session using the administrative client, we open a Web browser and type:

   `http://9.1.39.46:8421/ibm/console`

   The login menu appears as shown in Figure 5-102.

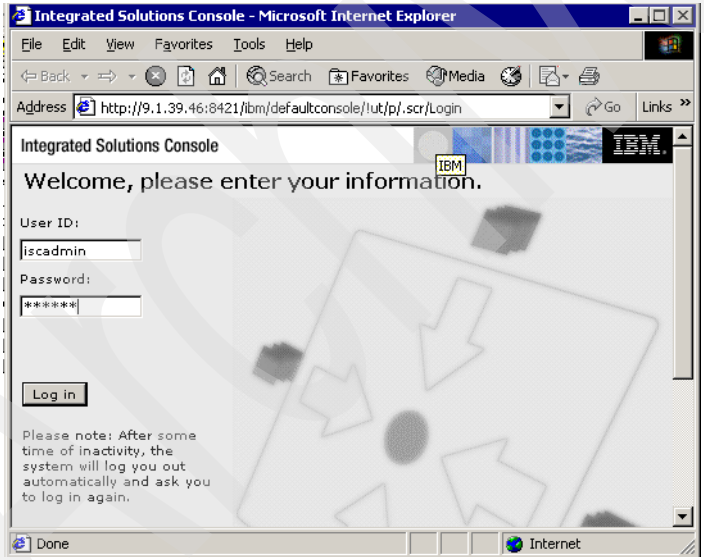

Figure 5-102 Login menu for the Administration Center

Archived

2. We type the user id and password that we chose at ISC installation in Figure 5-26 on page 97 and the panel in Figure 5-103 displays.

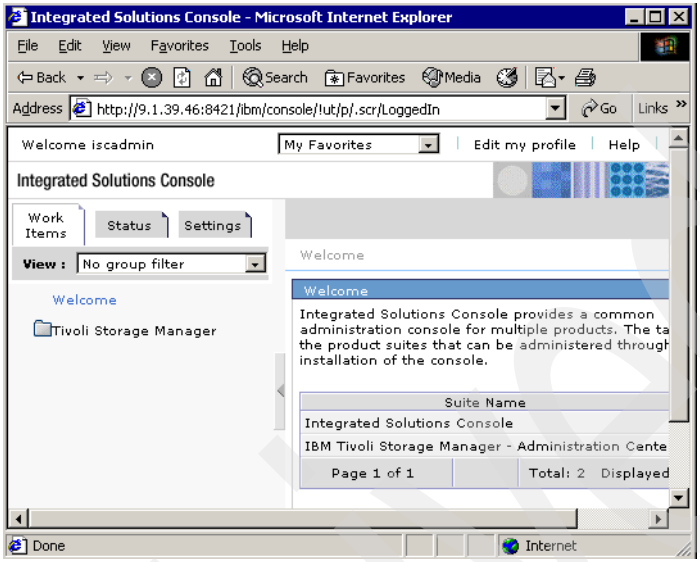

Figure 5-103 Administration Center

3. In Figure 5-103 we open the Tivoli Storage Manager folder on the right and the panel in Figure 5-104 is displayed.

Archived

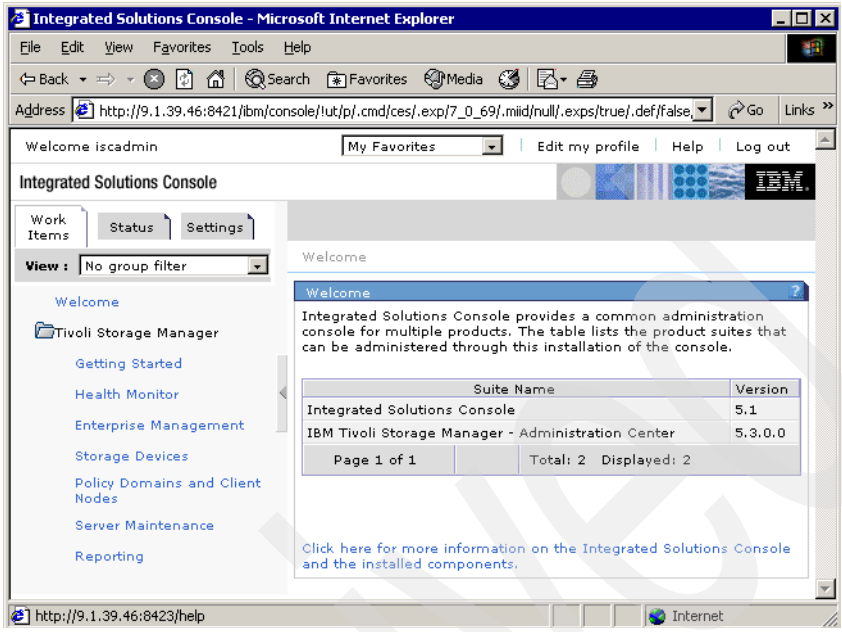

Figure 5-104 Options for Tivoli Storage Manager

4. We first need to create a new Tivoli Storage Manager server connection. To do this, we use Figure 5-104. We select Enterprise Management on this menu, and this takes us to the following menu (Figure 5-105).

Archived

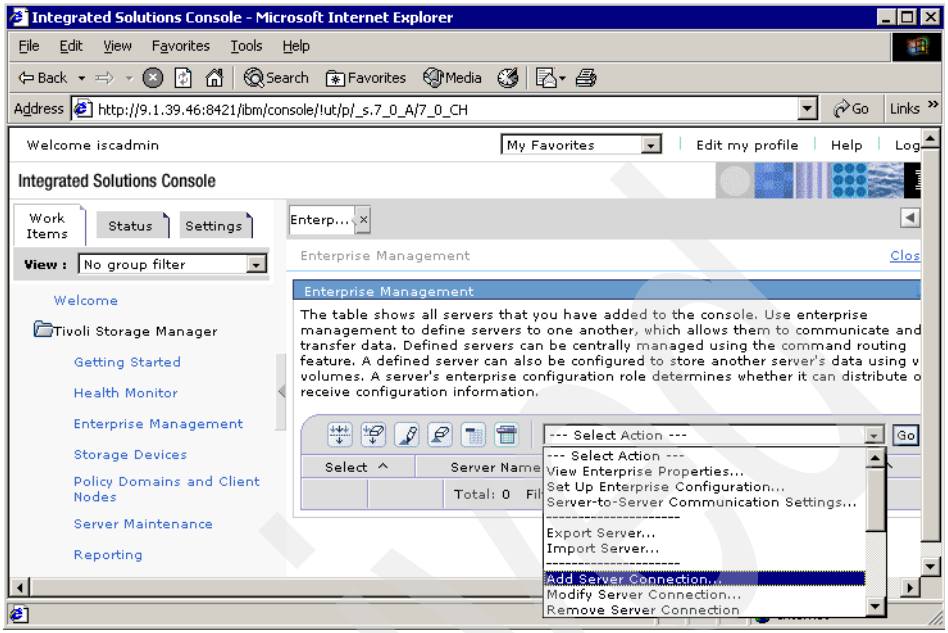

Figure 5-105 Selecting to create a new server connection

5. In Figure 5-105, if we open the pop-up menu such as we show, we have several options. To create a new server connection we select **Add Server Connection** and then we click Go. The following menu displays (Figure 5-106).

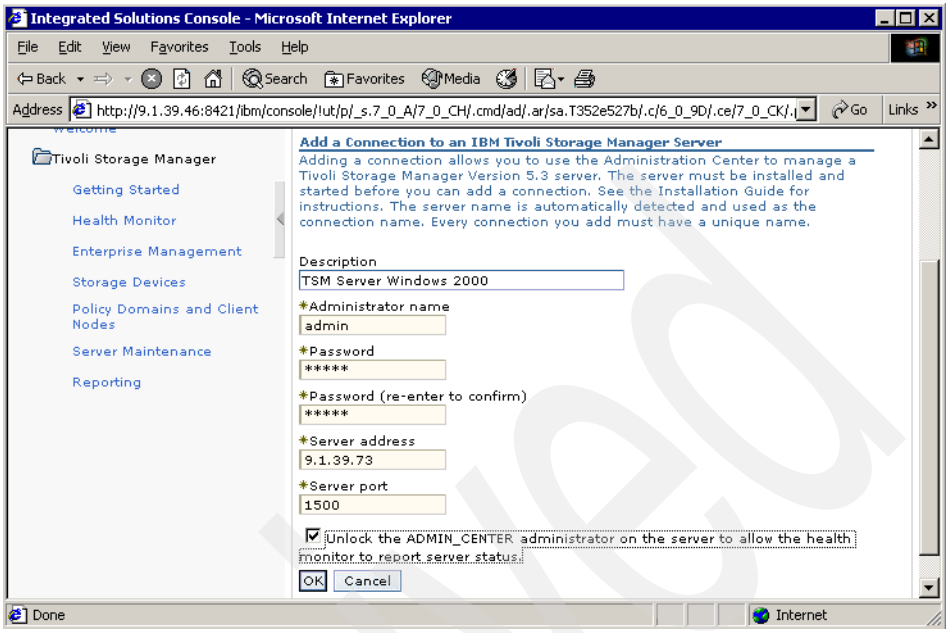

Figure 5-106 Specifying Tivoli Storage Manager server parameters

6. In Figure 5-106 we specify a Description (optional) as well as the Administrator name and Password to log into this server. We also specify the TCP/IP server address of our Windows 2000 Tivoli Storage Manager server and its TCP port. Since we want to unlock the ADMIN_CENTER administrator to allow the health monitor to report server status, we check the box and then we click **OK**.

Archived

7. An information menu displays, prompting us to fill in the form below to configure the health monitor. We type the information and then we click **OK**, as shown in Figure 5-107.

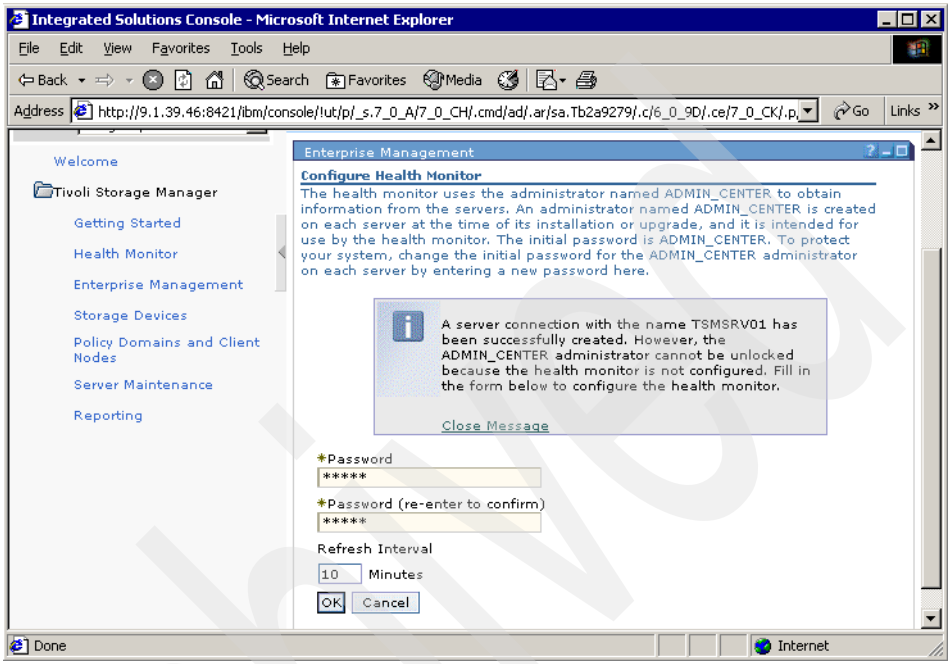

Figure 5-107 Filling in a form to unlock ADMIN_CENTER

Archived

8. And finally, Figure 5-108 shows us where we can see the connection to TSMSRV01 server. We are ready to manage this server using the different options and commands that the Administration Center provides us.

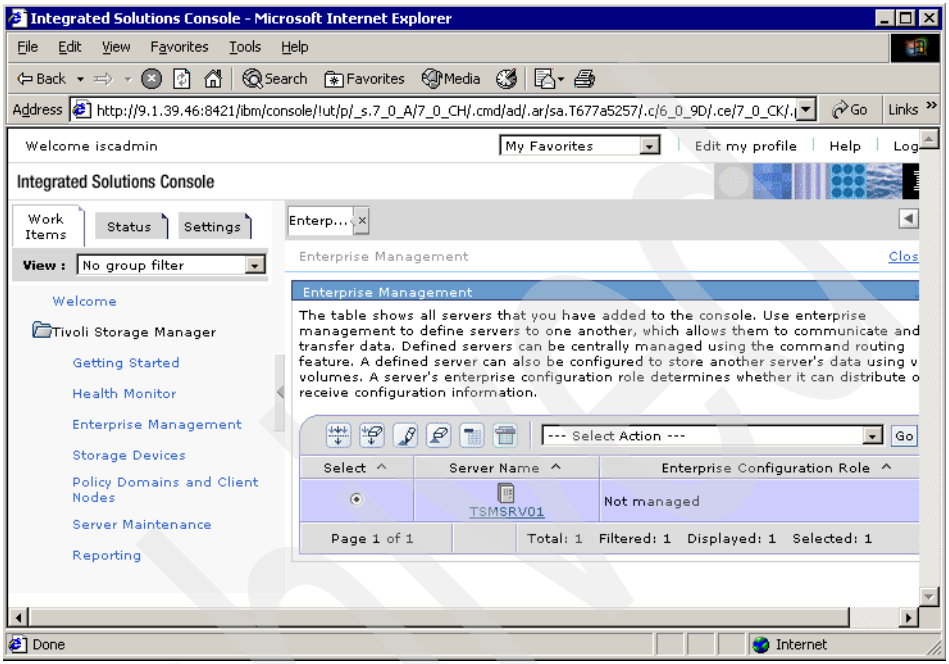

Figure 5-108 TSMSRV01 Tivoli Storage Manager server created

5.6 Tivoli Storage Manager Server and Windows 2003

The Tivoli Storage Manager server installation process was described on “Installing Tivoli Storage Manager Server on a MSCS” on page 79, at the beginning of this chapter.

In this section we describe how we configure the Tivoli Storage Manager server software to be capable of running in our MSCS Windows 2003, the same cluster we installed and configured in 4.4, “Windows 2003 MSCS installation and configuration” on page 44.

5.6.1 Windows 2003 lab setup

Our clustered lab environment consists of two Windows 2003 Enterprise Servers. Both servers are domain controllers as well as DNS servers.

Figure 5-109 shows the Tivoli Storage Manager server configuration for our Windows 2003 cluster.

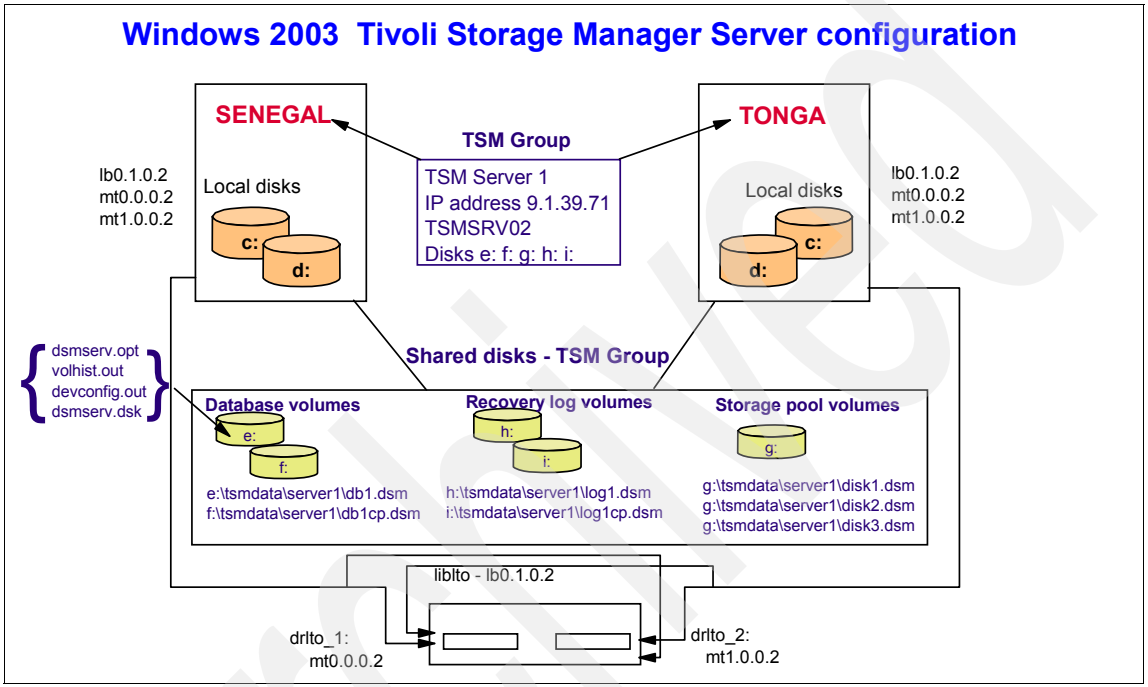

Figure 5-109 Lab setup for a 2-node cluster

Archived

Refer to Table 4-4 on page 46, Table 4-5 on page 47, and Table 4-6 on page 47 for specific details of the Windows 2003 cluster configuration.

For this section, we use the configuration shown below in Table 5-4, Table 5-5, and Table 5-6.

Table 5-4 Lab Windows 2003 ISC cluster resources

Resource Group TSM Admin Center	
ISC name	ADMCNT02
ISC IP address	9.1.39.69
ISC disk	j:
ISC services name	IBM WebSphere Application Server V5 - ISC Runtime Service ISC Help Service

Table 5-5 Lab Windows 2003 Tivoli Storage Manager cluster resources

Resource Group TSM Group	
TSM Cluster Server Name	TSMSRV02
TSM Cluster IP	9.1.39.71
TSM database disks *	e: h:
TSM recovery log disks *	f: i:
TSM storage pool disk	g:
TSM service name	TSM Server 1

* We choose two disk drives for the database and recovery log volumes so that we can use the Tivoli Storage Manager mirroring feature

Archived

Table 5-6 Tivoli Storage Manager virtual server for our Windows 2003 lab

Server parameters	
Server name	TSMSRV02
High level address	9.1.39.71
Low level address	1500
Server password	itsosj
Recovery log mode	Roll-forward
Libraries and drives	
Library name	LIBLTO
Drive 1	DRLTO_1
Drive 2	DRLTO_2
Device names	
Library device name	lb0.1.0.2
Drive 1 device name	mt0.0.0.2
Drive 2 device name	mt1.0.0.2
Primary Storage Pools	
Disk Storage Pool	SPD_BCK (nextstg=SPT_BCK)
Tape Storage Pool	SPT_BCK
Copy Storage Pool	
Tape Storage Pool	SPCPT_BCK
Policy	
Domain name	STANDARD
Policy set name	STANDARD
Management class name	STANDARD
Backup copy group	STANDARD (default, DEST=SPD_BCK)
Archive copy group	STANDARD (default)

Archived

Before installing the Tivoli Storage Manager server on our Windows 2003 cluster, the TSM Group *must only* contains disk resources, such as we can see in the Cluster Administrator menu in Figure 5-110.

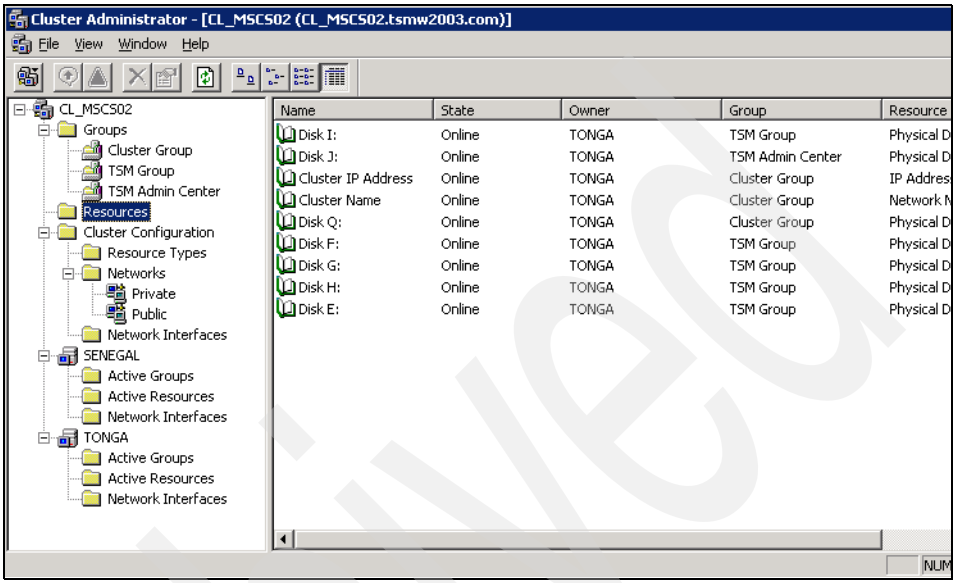

Figure 5-110 Cluster Administrator with TSM Group

Installation of IBM tape device drivers on Windows 2003

As we can see in Figure 4-16 on page 45, our two Windows 2003 servers are attached to the SAN, so that both can see the IBM 3582 Tape Library as well as its two IBM 3580 tape drives.

Since IBM Tape Libraries use their own device drivers to work with Tivoli Storage Manager, we have to download and install the last available version of the IBM LTO drivers for 3582 Tape Library and 3580 Ultrium 2 tape drives.

We use the folder “drivers_lto” to download the IBM drivers. Then, we use the Windows device manager menu, right-click one of the drives and select **Update driver**. We must specify the path where to look for the drivers, the “drivers_lto” folder, and follow the installation process menus.

We do not show the whole installation process in this book. Refer to the *IBM Ultrium Device Drivers Installation and User’s Guide* for a detailed description of this task.

Archived

After the successful installation of the drivers, both nodes recognize the 3582 medium changer and the 3580 tape drives as shown in Figure 5-111.

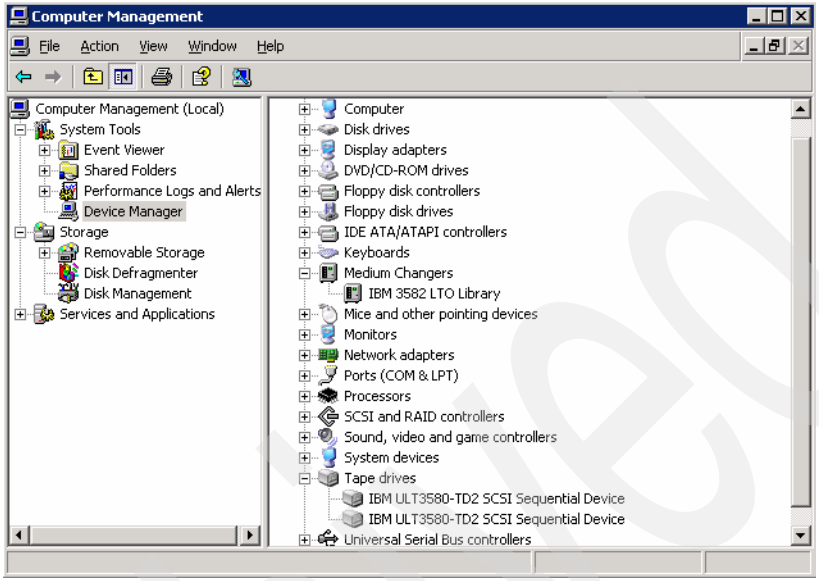

Figure 5-111 3582 and 3580 drivers installed

5.6.2 Windows 2003 Tivoli Storage Manager Server configuration

When installation of Tivoli Storage Manager packages on both nodes of the cluster is completed, we can proceed with the configuration.

The configuration tasks are performed on each node of the cluster. The steps vary depending upon whether it is the first node we are configuring or the second one.

When we start the configuration procedure on the first node, the Tivoli Storage Manager server instance is created and started. On the second node, the procedure will allow this server to host that instance.

Important: It is necessary to install a Tivoli Storage Manager server on the first node before configuring the second node. If we do not do that, the configuration will fail.

Archived

Configuring the first node

We start configuring Tivoli Storage Manager on the first node. To perform this task, resources must be hosted by this node. We can check this issue opening the cluster administrator from **Start** → **Programs** → **Administrative Tools** → **Cluster Administrator** (Figure 5-112).

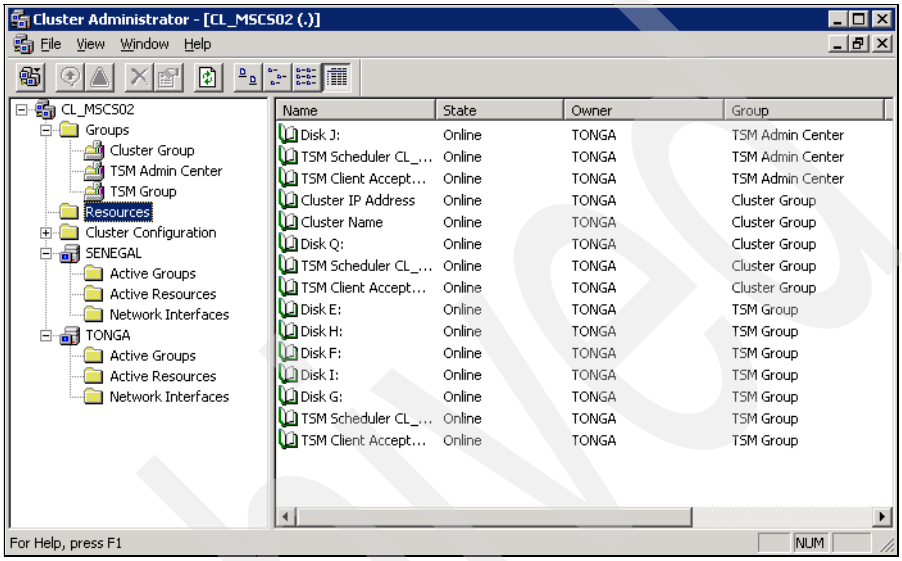

Figure 5-112 Cluster resources

As shown in Figure 5-112, TONGA hosts all the resources of the TSM Group. That means we can start configuring Tivoli Storage Manager on this node.

Attention: Before starting the configuration process, we copy mfc71u.dll and mvscr71.dll from the Tivoli Storage Manager \console directory (normally c:\Program Files\Tivoli\tsm\console) into c:\%SystemRoot%\cluster directory on each cluster node involved. If we do not do that, the cluster configuration will fail. This is caused by a new Windows compiler (VC71) that creates dependencies between tsmsvrrsc.dll and tsmsvrrscex.dll and mfc71u.dll and mvscr71.dll. Microsoft has not included these files in its service packs.

1. To start the initialization, we open the Tivoli Storage Manager Management Console as shown in Figure 5-113.

Figure 5-113 Starting the Tivoli Storage Manager management console

2. The Initial Configuration Task List for Tivoli Storage Manager menu, Figure 5-114, shows a list of the tasks needed to configure a server with all basic information. To let the wizard guide us throughout the process, we select **Standard Configuration**. This will also enable automatic detection of a clustered environment. We then click **Start**.

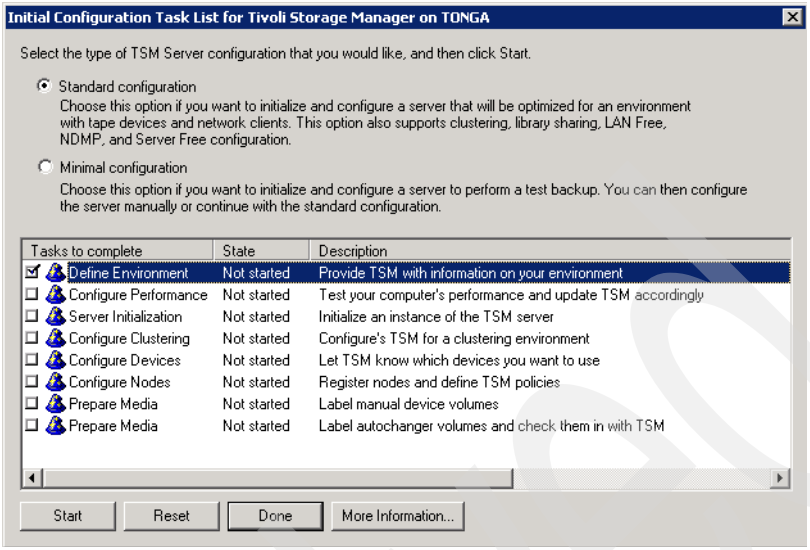

Figure 5-114 Initial Configuration Task List

3. The Welcome menu for the first task, **Define Environment**, displays as shown in Figure 5-115. We click **Next**.

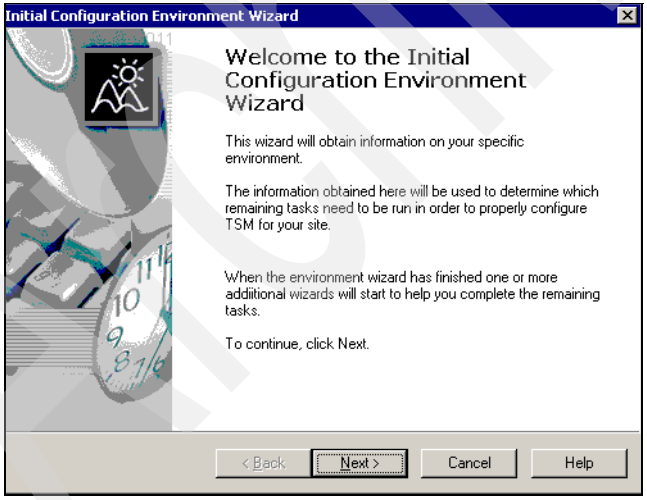

Figure 5-115 Welcome Configuration wizard

Archived

4. To have additional information displayed during the configuration, we select **Yes** and click **Next** in Figure 5-116.

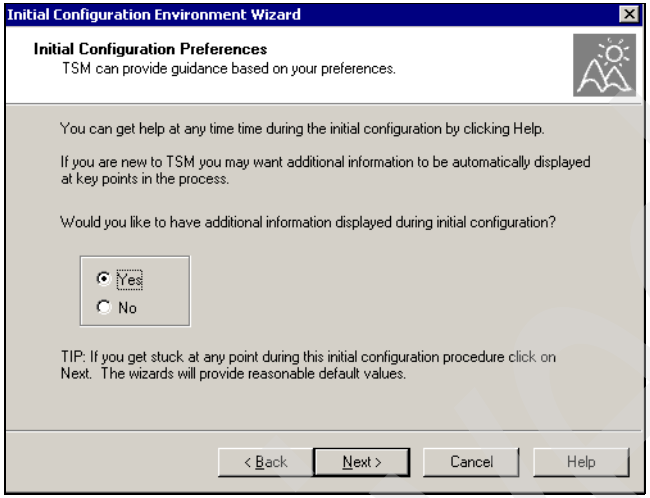

Figure 5-116 Initial configuration preferences

5. Tivoli Storage Manager can be installed **Standalone** (for only one client), or **Network** (when there are more clients). In most cases we have more than one client. We select **Network** and then click **Next** as shown in Figure 5-117.

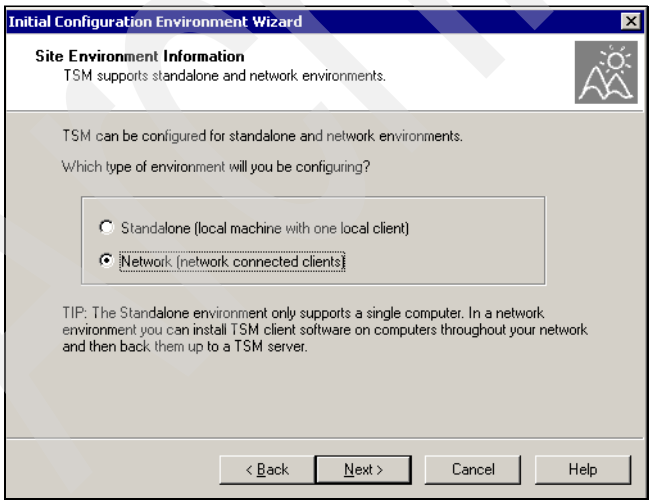

Figure 5-117 Site environment information

Archived

6. The **Initial Configuration Environment** is done. We click **Finish** in Figure 5-118.

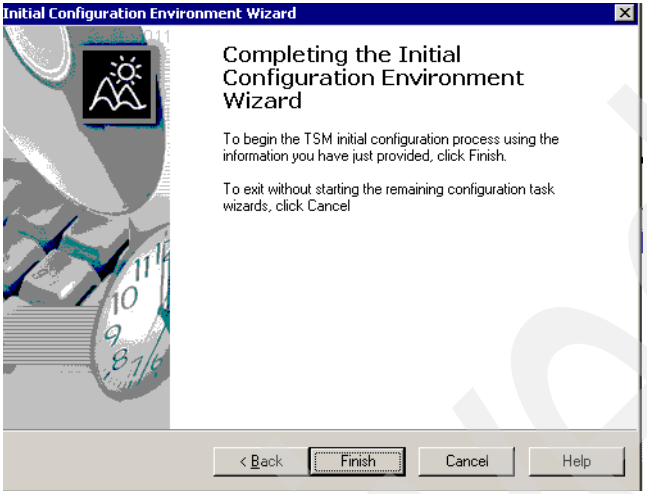

Figure 5-118 Initial configuration

7. The next task is to complete the Performance Configuration Wizard. We click **Next** (Figure 5-119).

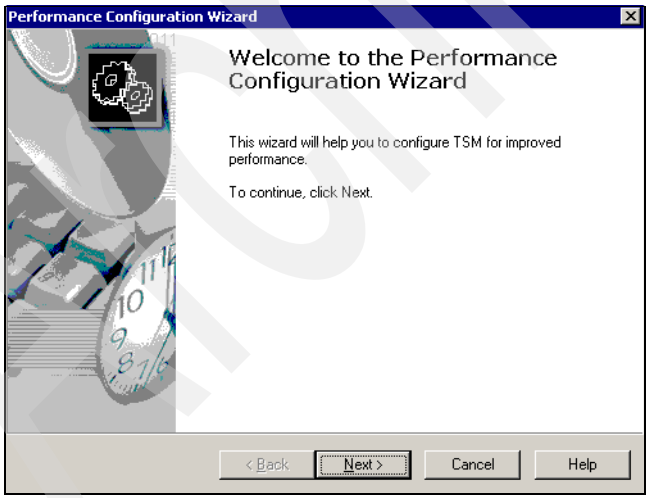

Figure 5-119 Welcome Performance Environment wizard

Archived

8. In Figure 5-120 we provide information about our own environment. Tivoli Storage Manager will use this information for tuning. For our lab, we used the defaults. In a real installation, it is necessary to select the values that best fit that environment. We click **Next**.

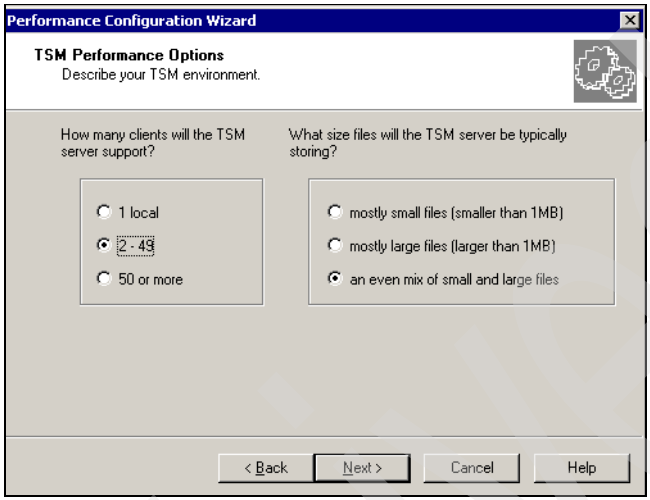

Figure 5-120 Performance options

9. The wizard starts to analyze the hard drives as shown in Figure 5-121. When the process ends, we click **Finish**.

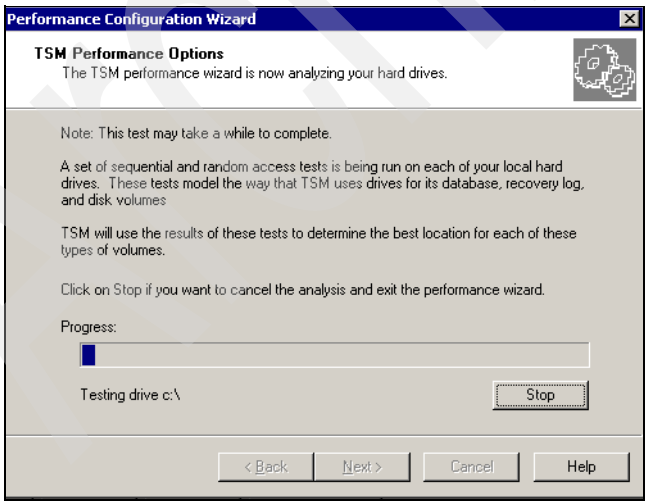

Figure 5-121 Drive analysis

10. The Performance Configuration task is completed (Figure 5-122).

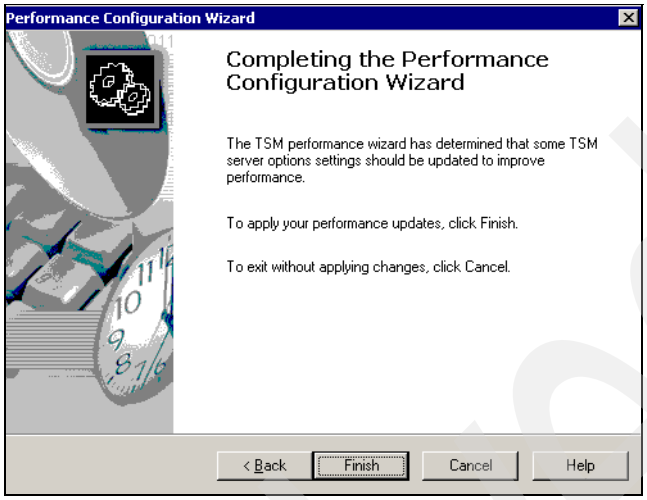

Figure 5-122 Performance wizard

11. Next step is the initialization of the Tivoli Storage Manager server instance. We click **Next** (Figure 5-123).

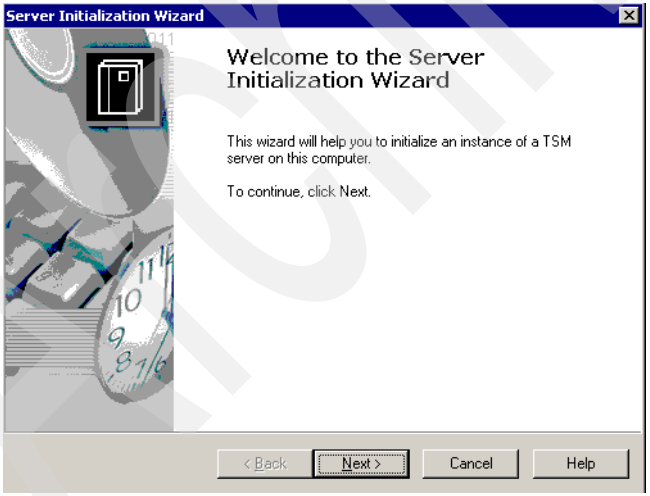

Figure 5-123 Server instance initialization wizard

Archived

12. The initialization process detects that there is a cluster installed. The option **Yes** is already selected. We leave this default in Figure 5-124 and we click **Next** so that Tivoli Storage Manager server instance is installed correctly.

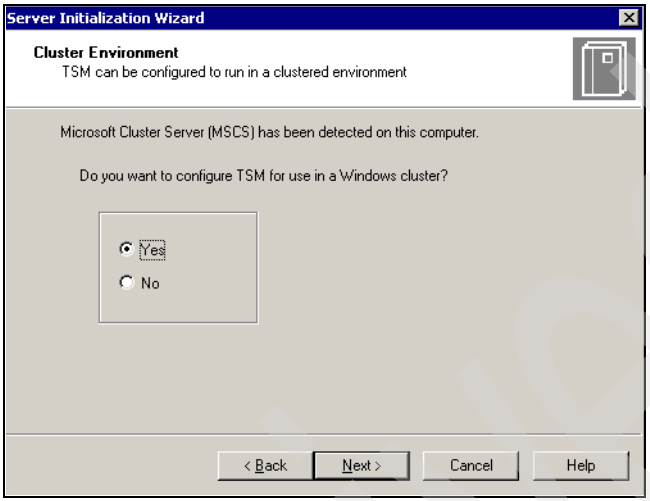

Figure 5-124 Cluster environment detection

13. We select the cluster group where Tivoli Storage Manager server instance will be created. This cluster group initially must contain only disk resources. For our environment this is TSM Group. Then we click **Next** (Figure 5-125).

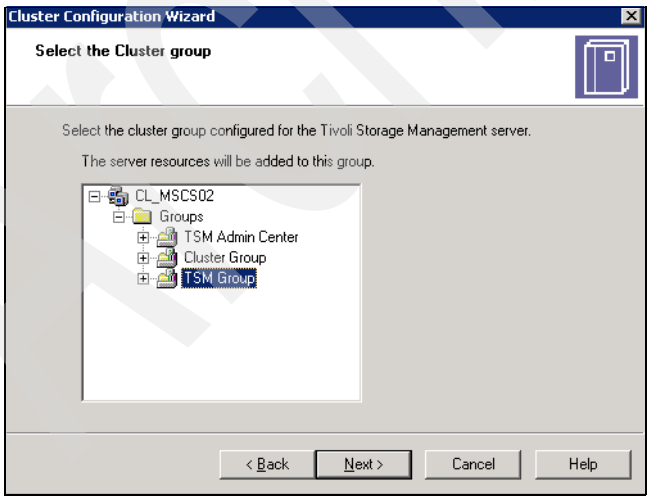

Figure 5-125 Cluster group selection

Important: The cluster group we choose here must match the cluster group used when configuring the cluster in Figure 5-134 on page 198.

14. In Figure 5-126 we select the directory where the files used by Tivoli Storage Manager server will be placed. It is possible to choose any disk on the Tivoli Storage Manager cluster group. We change the drive letter to use e: and click **Next** (Figure 5-126).

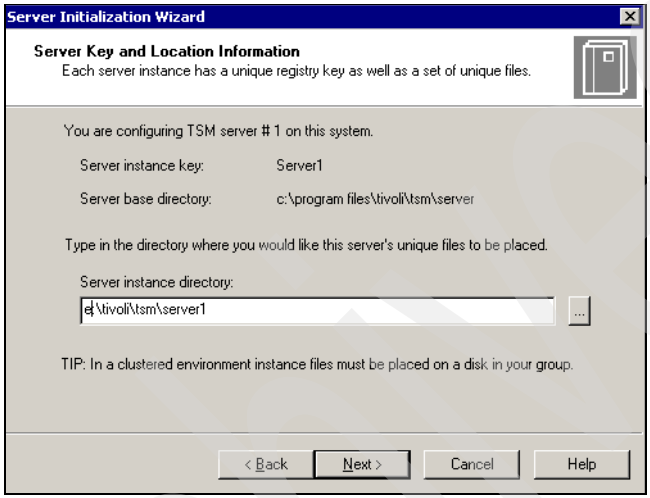

Figure 5-126 Server initialization wizard

15. In Figure 5-127 we type the complete paths and sizes of the initial volumes to be used for database, recovery log and disk storage pools. Refer to Table 5-5 on page 181 where we planned the use of the disk drives.

    A specific installation should choose its own values.

    We also check the two boxes on the two bottom lines to let Tivoli Storage Manager create additional volumes as needed.

    With the selected values, we will initially have a 1000 MB size database volume with name db1.dsm, a 500 MB size recovery log volume called log1.dsm, and a 5 GB size storage pool volume of name disk1.dsm. If we need, we can create additional volumes later.

    We input our values and click **Next**.

Archived

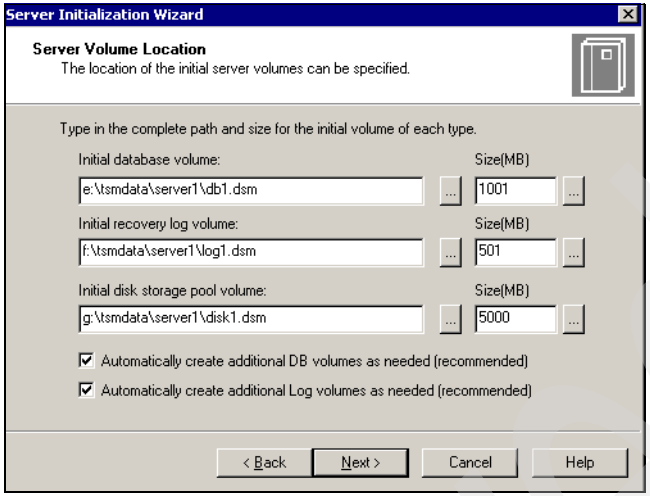

Figure 5-127 Server volume location

16. On the server service logon parameters shown in Figure 5-128, we select the Windows **account** and user ID that Tivoli Storage Manager server instance will use when logging onto Windows. We recommend to leave the defaults and click **Next**.

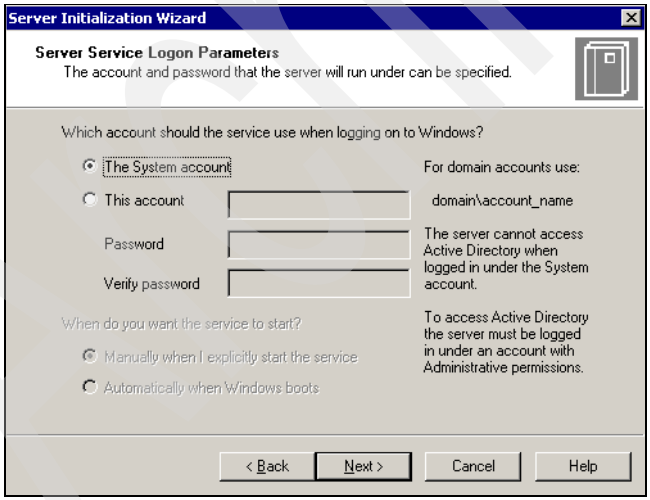

Figure 5-128 Server service logon parameters

17. In Figure 5-129, we specify the server name that Tivoli Storage Manager will use as well as its password. The server password is used for server-to-server communications. We will need it later on with the Storage Agent. This password can also be set later using the administrator interface. We click **Next**.

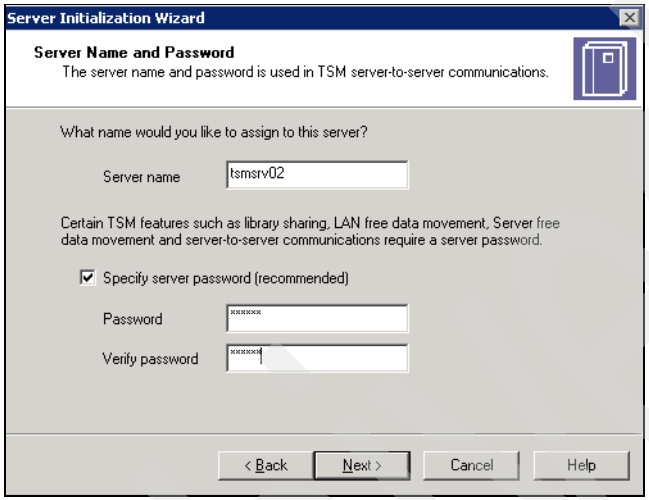

Figure 5-129 Server name and password

Important: The server name we select here must be the same name that we will use when configuring Tivoli Storage Manager on the other node of the MSCS.

Archived

18. We click **Finish** in Figure 5-130 to start the process of creating the server instance.

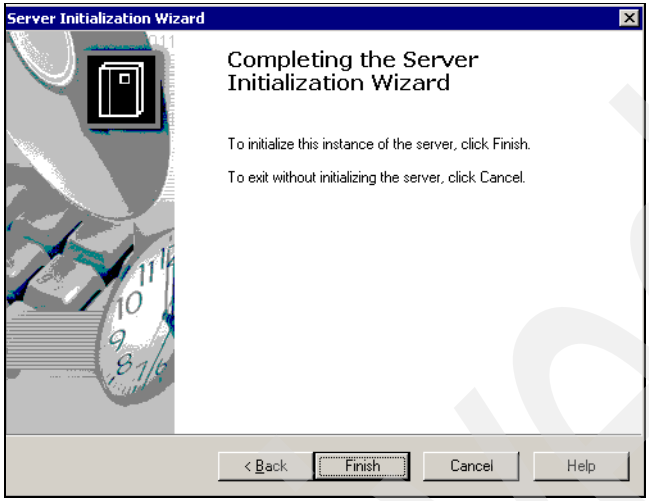

Figure 5-130 Completing the Server Initialization wizard

19. The wizard starts the process of the server initialization and shows a progress bar (Figure 5-131).

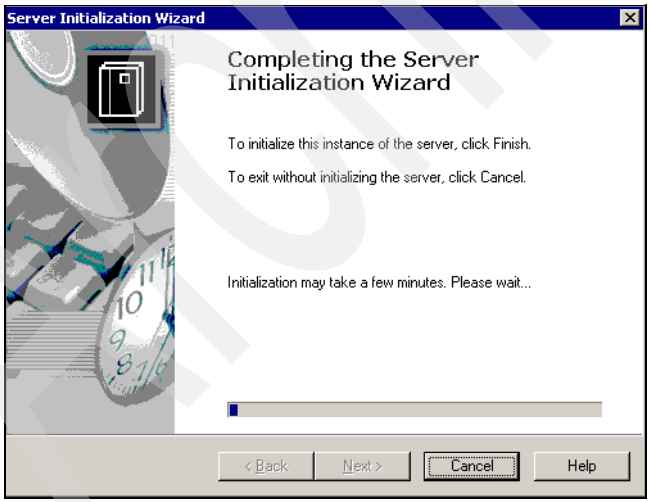

Figure 5-131 Completing the server installation wizard

Archived

20. If the initialization ends without any errors we receive the following informational message. We click **OK** (Figure 5-132).

Figure 5-132 Tivoli Storage Manager Server has been initialized

21. The next task performed by the wizard if the Cluster Configuration. We click **Next** on the welcome page (Figure 5-133).

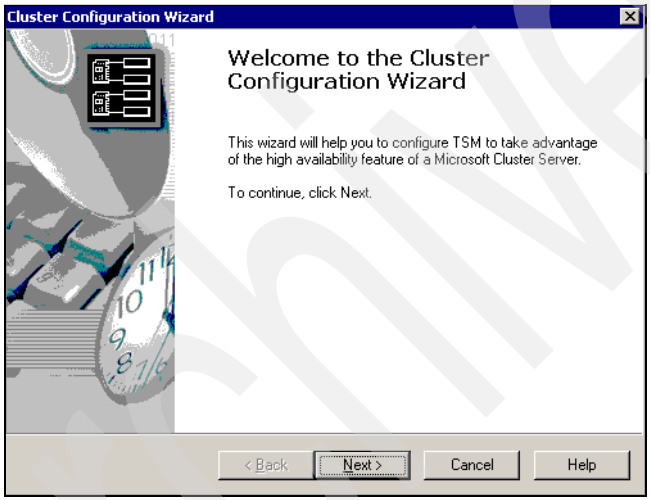

Figure 5-133 Cluster configuration wizard

22. We select the cluster group where Tivoli Storage Manager server will be configured and click **Next** (Figure 5-134).

> **Important:** Do not forget that the cluster group we select here must match the cluster group used during the server initialization wizard process in Figure 5-125 on page 192.

Archived

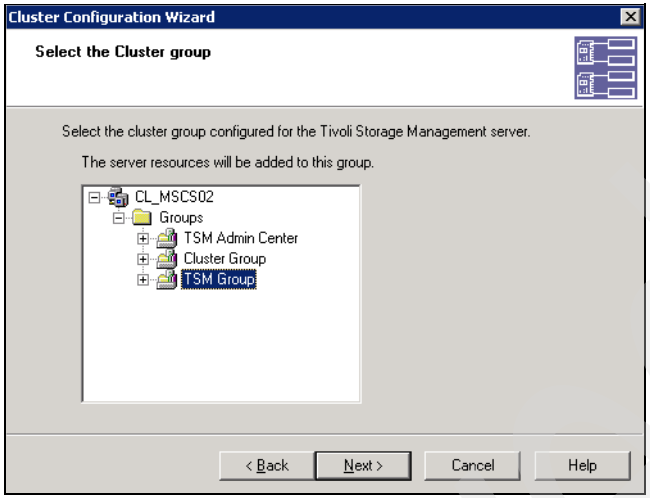

Figure 5-134 Select the cluster group

23. In Figure 5-135 we can configure Tivoli Storage Manager to manage tape failover in the cluster.

Note: MSCS does not support the failover of tape devices. However, Tivoli Storage Manager can manage this type of failover using a shared SCSI bus for the tape devices. Each node in the cluster must contain an additional SCSI adapter card. The hardware and software requirements for tape failover to work are described on Tivoli Storage Manager documentation.

Archived

Our lab environment does not meet the requirements for tape failover support so we select **Do not configure TSM to manage tape failover** and click **Next** (Figure 5-136).

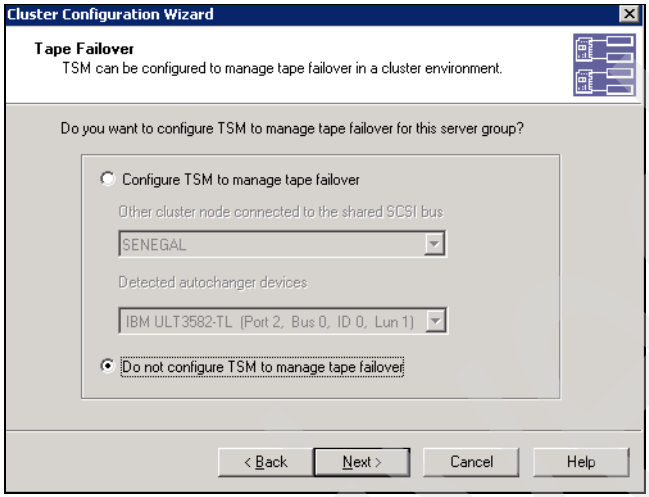

Figure 5-135 Tape failover configuration

24. In Figure 5-136 we enter the IP address and Subnest Mask that Tivoli Storage Manager virtual server will use in the cluster. This IP address must match the IP address selected in our planning and design worksheets (see Table 5-5 on page 181).

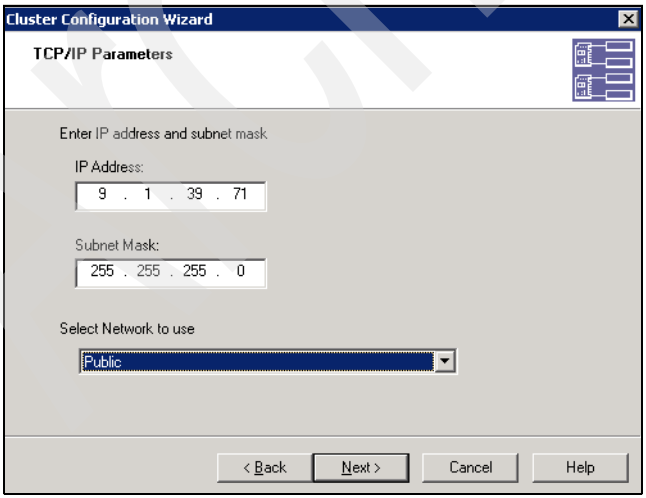

Figure 5-136 IP address

25. In Figure 5-137 we enter the Network name. This must match the network name we selected in our planning and design worksheets (see Table 5-5 on page 181). We enter TSMSRV02 and click **Next**.

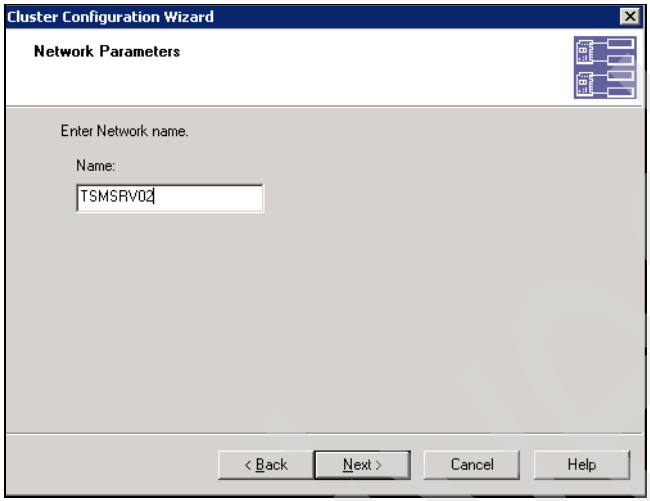

Figure 5-137 Network Name

26. On the next menu we check that everything is correct and we click **Finish**. This completes the cluster configuration on TONGA (Figure 5-138).

Figure 5-138 Completing the Cluster configuration wizard

27. We receive the following informational message and we click **OK** (Figure 5-139).

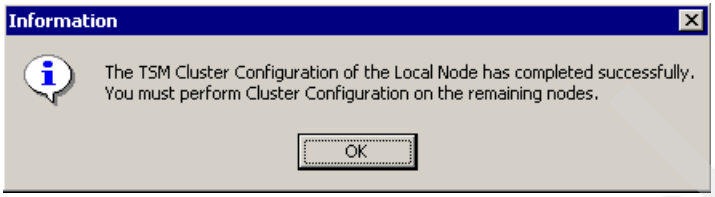

Figure 5-139 End of Tivoli Storage Manager Cluster configuration

At this time, we can continue with the initial configuration wizard, to set up devices, nodes and media. However, for the purpose of this book we will stop here. These tasks are the same we would follow in a regular Tivoli Storage Manager server. So, we click **Cancel** when the Device Configuration welcome menu displays.

So far Tivoli Storage Manager server instance is installed and started on TONGA. If we open the Tivoli Storage Manager console we can check that the service is running as shown in Figure 5-140.

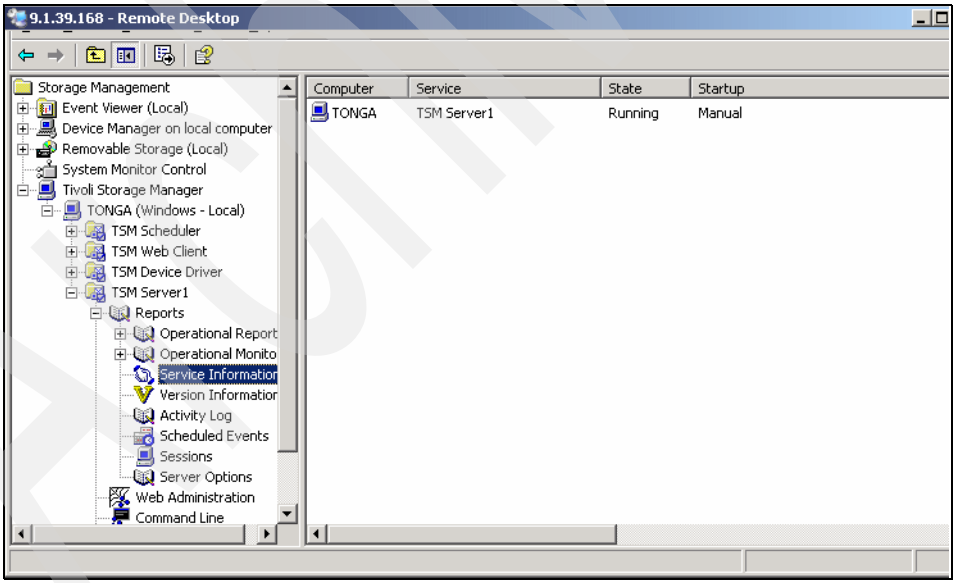

Figure 5-140 Tivoli Storage Manager console

Important: before starting the initial configuration for Tivoli Storage Manager on the second node, we must stop the instance on the first node.

28. We stop the Tivoli Storage Manager server instance on TONGA before going on with the configuration on SENEGAL.

Configuring the second node

In this section we describe how to configure Tivoli Storage Manager on the second node of the MSCS. We follow the same process as for the first node. The only difference is that the Tivoli Storage Manager server instance was already created on the first node. Now the installation will allow the second node to host that server instance.

1. First of all we move the Tivoli Storage Manager cluster group to the second node using the Cluster Administrator. Once moved, the resources should be hosted by SENEGAL, as shown in Figure 5-141.

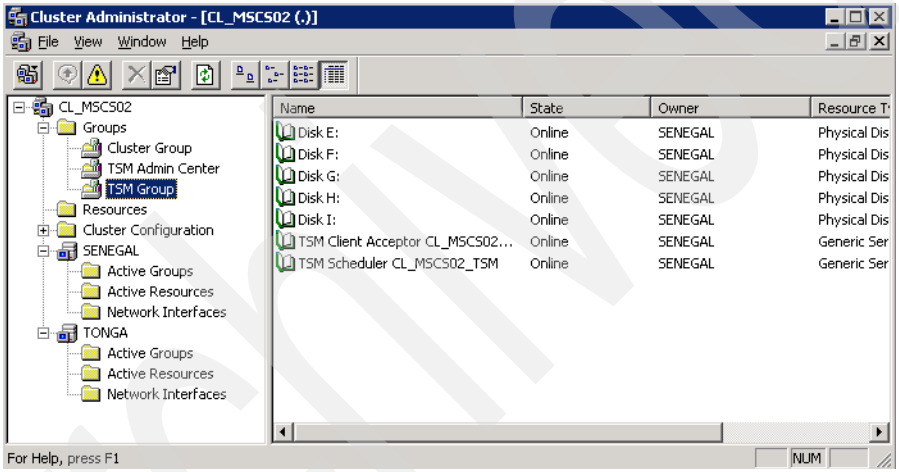

Figure 5-141 Cluster resources

Note: As we can see in Figure 5-141 the IP address and network name resources are not created yet. We still have only disk resources in the TSM resource group. When the configuration ends in SENEGAL, the process will create those resources for us.

2. We open the Tivoli Storage Manager console to start the initial configuration on the second node and follow the same steps (1 to 18) from section “Configuring the first node” on page 185, until we get into the **Cluster Configuration** Wizard in Figure 5-142. We click **Next**.

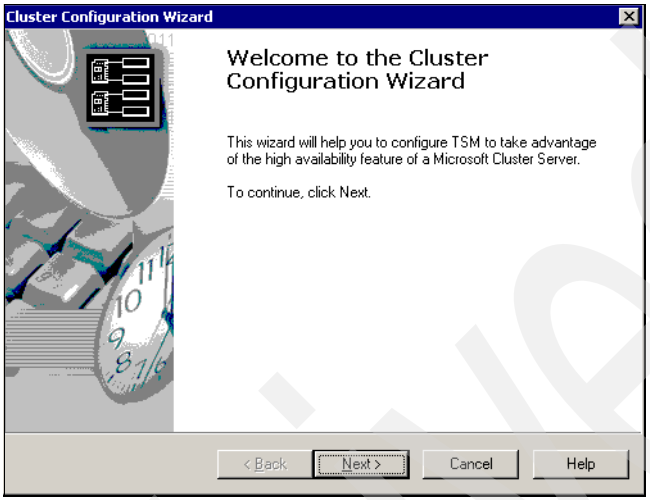

Figure 5-142 Cluster configuration wizard

3. On the Select **Cluster Group** menu in Figure 5-143 we select the same group, the TSM Group, and then click **Next** (Figure 5-143).

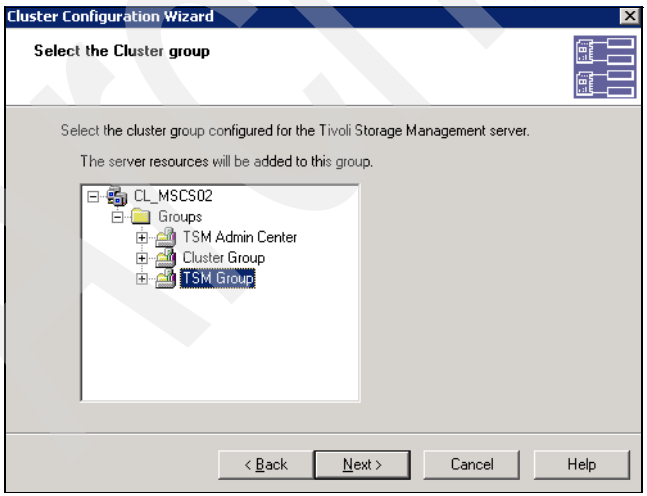

Figure 5-143 Selecting the cluster group

Archived

4. In Figure 5-144 we check that the information reported is correct and then we click **Finish** (Figure 5-144).

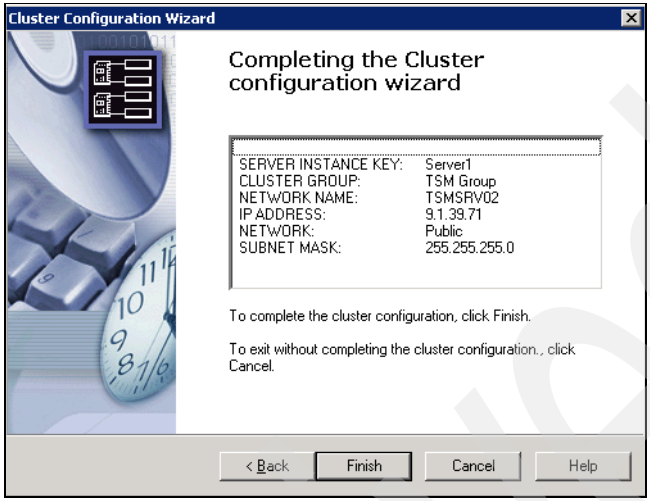

Figure 5-144 Completing the Cluster Configuration wizard

5. The wizard starts the configuration for the server as shown in Figure 5-145.

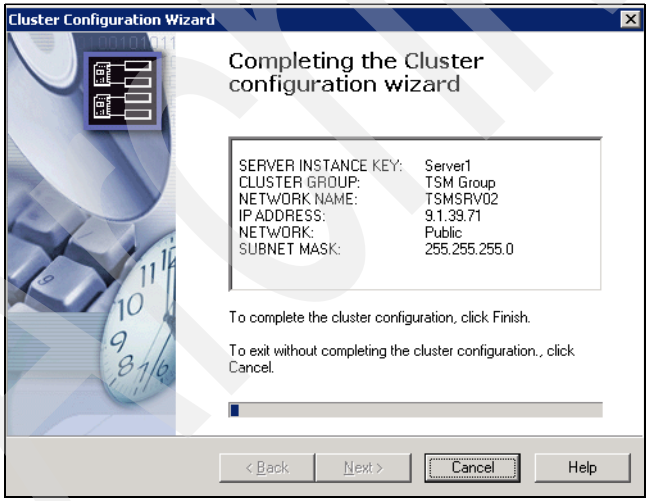

Figure 5-145 The wizard starts the cluster configuration

Archived

6. When the configuration is successfully completed the following message is displayed. We click **OK** (Figure 5-146).

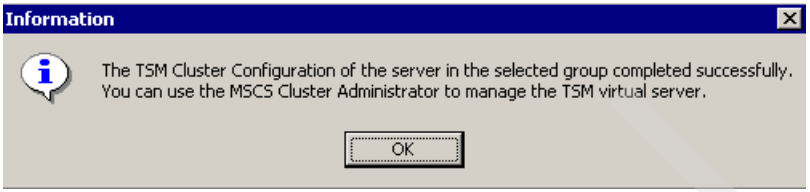

Figure 5-146 Successful installation

So far the Tivoli Storage Manager is correctly configured on the second node. To manage the virtual server, we have to use the MSCS Cluster Administrator.

When we open the MSCS Cluster Administrator to check the results of the process followed on this node. As we can see in Figure 5-147, the cluster configuration process itself creates the following resources on the TSM cluster group:

- TSM Group IP Address: the one we specified in Figure 5-136 on page 199.
- TSM Group Network name: the specified in Figure 5-137 on page 200.
- TSM Group Server: the Tivoli Storage Manager server instance.

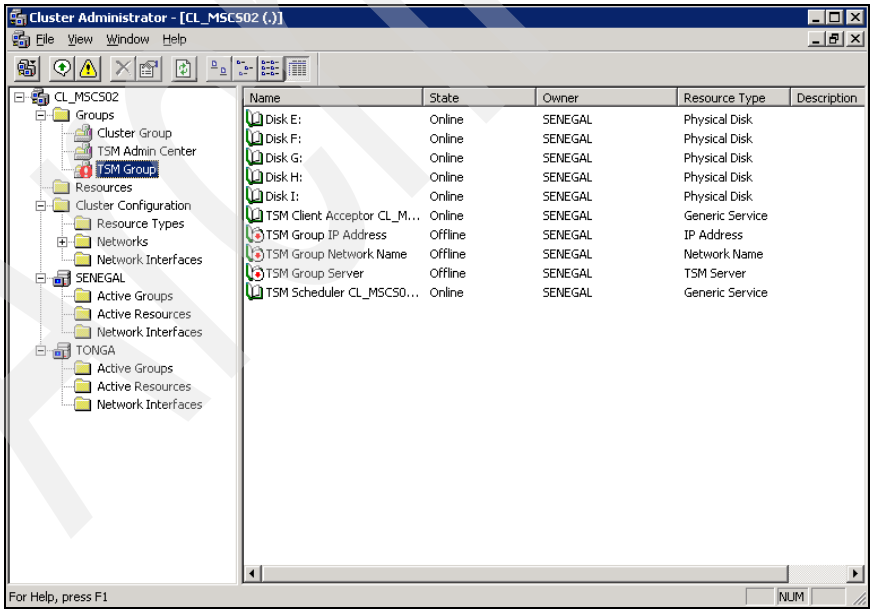

Figure 5-147 TSM Group resources

The TSM Group cluster group is offline because the new resources are offline. Now we must bring online every resource on this group as shown in Figure 5-148.

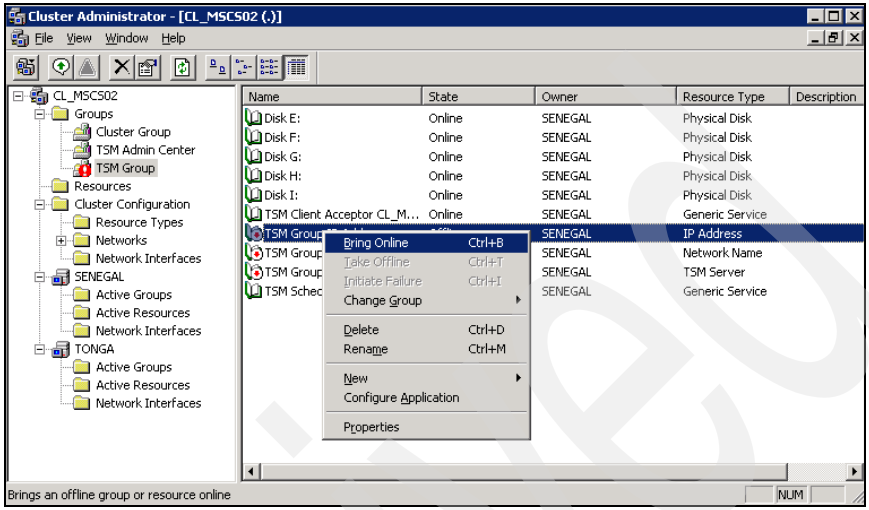

Figure 5-148 Bringing resources online

In this figure we show how to bring online the TSM Group IP Address. The same process should be done for the remaining resources.

The final menu should display as shown in Figure 5-149.

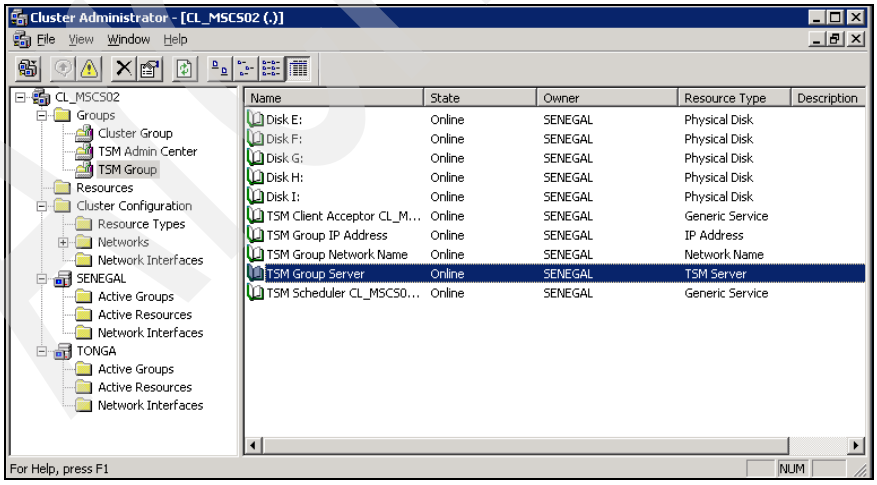

Figure 5-149 TSM Group resources online

Now the TSM server instance is running on SENEGAL, which is the node which hosts the resources. If we go into the Windows services menu, Tivoli Storage Manager server instance is started, as shown in Figure 5-150.

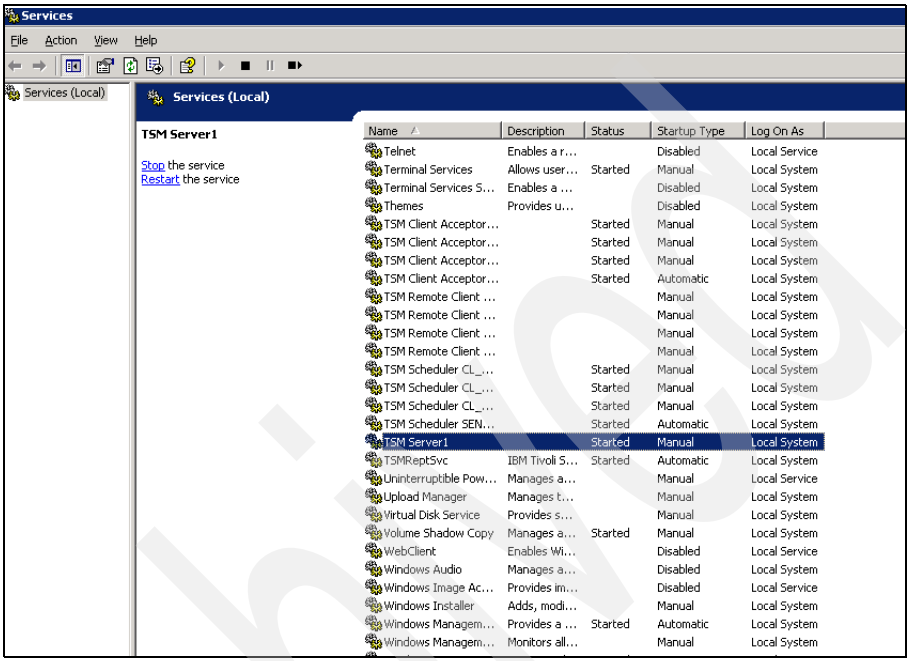

Figure 5-150 Services

Important: Do not forget to manage always the Tivoli Storage Manager server instance using the Cluster Administrator menu, to bring it online or offline.

We are now ready to test the cluster.

Archived

5.6.3 Testing the server on Windows 2003

In order to check the high availability of Tivoli Storage Manager server on our lab environment, we must do some testing.

Our objective with these tests is to show how Tivoli Storage Manager in a clustered environment manages its own resources to achieve high availability and how it can respond after certain kinds of failures that affect the shared resources.

Testing client incremental backup using the GUI

Our first test uses the Tivoli Storage Manager GUI to start an incremental backup.

Objective

The objective of this test is to show what happens when a client incremental backup starts using the Tivoli Storage Manager GUI and suddenly the node which hosts the Tivoli Storage Manager server fails.

Activities

To do this test, we perform these tasks:

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager cluster group as shown in Figure 5-151.

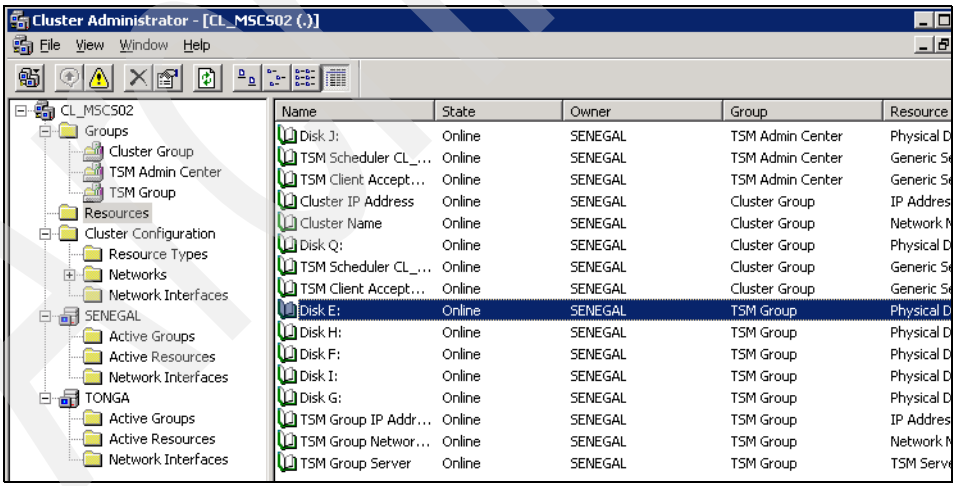

Figure 5-151 Cluster Administrator shows resources on SENEGAL

2. We start an incremental backup from the second node, TONGA, using the Tivoli Storage Manager backup/archive GUI client, which is also installed on each node of the cluster. We select the local drives, the System State, and the System Services as shown in Figure 5-152.

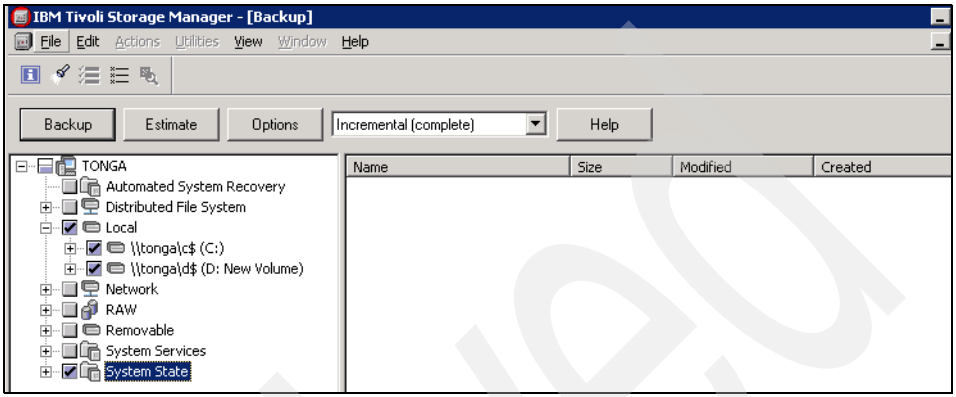

Figure 5-152 Selecting a client backup using the GUI

3. The transfer of files starts, as we can see in Figure 5-153.

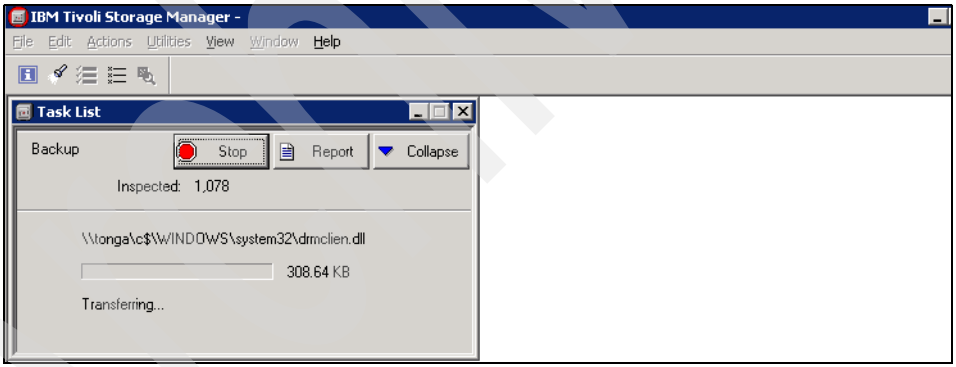

Figure 5-153 Transferring files to the server

4. While the client is transferring files to the server we force a failure on SENEGAL, the node that hosts the Tivoli Storage Manager server. When Tivoli Storage Manager restarts on the second node, we can see in the GUI client that backup is held and a reopening session message is received, as shown in Figure 5-154.

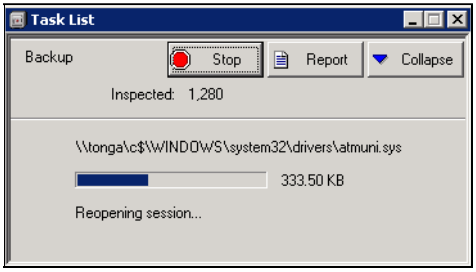

Figure 5-154 Reopening the session

5. When the connection is re-established, the client continues sending files to the server, as shown in Figure 5-155.

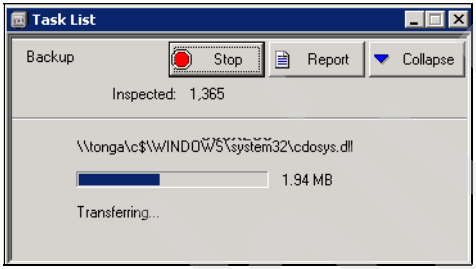

Figure 5-155 Transfer of data goes on when the server is restarted

6. The client backup ends successfully.

Results summary

The result of the test shows that when we start a backup from a client and there is a failure that forces Tivoli Storage Manager server to fail, backup is held and when the server is up again, the client reopens a session with the server and continues transferring data.

> **Note:** In the test we have just described, we used a disk storage pool as the destination storage pool. We also tested using a tape storage pool as destination and we got the same results. The only difference is that when the Tivoli Storage Manager server is again up, the tape volume it was using on the first node is unloaded from the drive and loaded again into the second drive, and the client receives a “media wait” message while this process takes place. After the tape volume is mounted, the backup continues and ends successfully.

Archived

Testing a scheduled client backup

The second test consists of a scheduled backup.

Objective

The objective of this test is to show what happens when a scheduled client backup is running and suddenly the node which hosts the Tivoli Storage Manager server fails.

Activities

We perform these tasks:

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager cluster group: TONGA.
2. We schedule a client incremental backup operation using the Tivoli Storage Manager server scheduler and this time we associate the schedule to the Tivoli Storage Manager client installed on SENEGAL.
3. A client session starts from SENEGAL as shown in Example 5-19.

Example 5-19 Activity log when the client starts a scheduled backup

```
02/07/2005 14:45:01 ANR2561I Schedule prompter contacting SENEGAL (session 16)
                     to start a scheduled operation. (SESSION: 16)
02/07/2005 14:45:03 ANR0403I Session 16 ended for node SENEGAL (). (SESSION:
                     16)
02/07/2005 14:45:03 ANR0406I Session 17 started for node SENEGAL (WinNT)
                     (Tcp/Ip senegal.tsmw2003.com(1491)). (SESSION: 17)
```

4. The client starts sending files to the server as shown in Example 5-20.

Example 5-20 Schedule log file shows the start of the backup on the client

```
02/07/2005 14:45:03 --- SCHEDULEREC QUERY BEGIN
02/07/2005 14:45:03 --- SCHEDULEREC QUERY END
02/07/2005 14:45:03 Next operation scheduled:
02/07/2005 14:45:03
------------------------------------------------------------
02/07/2005 14:45:03 Schedule Name:         DAILY_INCR
02/07/2005 14:45:03 Action:                Incremental
02/07/2005 14:45:03 Objects:
02/07/2005 14:45:03 Options:
02/07/2005 14:45:03 Server Window Start:   14:45:00 on 02/07/2005
02/07/2005 14:45:03
------------------------------------------------------------
02/07/2005 14:45:03
Executing scheduled command now.
02/07/2005 14:45:03 --- SCHEDULEREC OBJECT BEGIN DAILY_INCR 02/07/2005 14:45:00
02/07/2005 14:45:03 Incremental backup of volume ‘\\senegal\c$’
```

```
02/07/2005 14:45:03 Incremental backup of volume ‘\\senegal\d$’
02/07/2005 14:45:03 Incremental backup of volume ‘SYSTEMSTATE’
02/07/2005 14:45:03 Backup System State using shadow copy...
02/07/2005 14:45:05 Backup System State: ‘System Files’.

02/07/2005 14:45:05 Backup System State: ‘System Volume’.

02/07/2005 14:45:05 Backup System State: ‘Active Directory’.

02/07/2005 14:45:05 Backup System State: ‘Registry’.

02/07/2005 14:45:05 Backup System State: ‘COM+ Database’.

02/07/2005 14:45:05 Incremental backup of volume ‘SYSTEMSERVICES’
02/07/2005 14:45:05 Backup System Services using shadow copy...
02/07/2005 14:45:05 Backup System Service: ‘Event Log’.

02/07/2005 14:45:05 Backup System Service: ‘RSM Database’.

02/07/2005 14:45:05 Backup System Service: ‘WMI Database’.

02/07/2005 14:45:05 Backup System Service: ‘Cluster DB’.

02/07/2005 14:45:07 ANS1898I ***** Processed     1,000 files *****
02/07/2005 14:45:07 Directory-->                 0 \\senegal\c$\ [Sent]
02/07/2005 14:45:07 Directory-->                 0 \\senegal\c$\Documents and
Settings [Sent]
02/07/2005 14:45:07 Directory-->                 0 \\senegal\c$\IBMTOOLS
[Sent]
02/07/2005 14:45:07 Directory-->                 0 \\senegal\c$\Program Files
[Sent]
02/07/2005 14:45:07 Directory-->                 0 \\senegal\c$\RECYCLER
[Sent]
02/07/2005 14:45:07 Directory-->                 0 \\senegal\c$\sdwork [Sent]
02/07/2005 14:45:07 Directory-->                 0 \\senegal\c$\swd [Sent]
02/07/2005 14:45:07 Directory-->                 0 \\senegal\c$\System Volume
Information [Sent]
02/07/2005 14:45:07 Directory-->                 0 \\senegal\c$\temp [Sent
```

5. While the client continues sending files to the server, we force TONGA to fail. The following sequence occurs:

   a. In the client, backup is held and an error is received as shown in Example 5-21.

Example 5-21 Error log when the client lost the session

```
02/07/2005 14:49:38 sessSendVerb: Error sending Verb, rc: -50
02/07/2005 14:49:38 ANS1809W Session is lost; initializing session reopen
procedure.
02/07/2005 14:49:38 ANS1809W Session is lost; initializing session reopen
procedure.
02/07/2005 14:50:35 ANS5216E Could not establish a TCP/IP connection with
address ‘9.1.39.71:1500’. The TCP/IP error is ‘Unknown error’ (errno = 10060).
02/07/2005 14:50:35 ANS4039E Could not establish a session with a TSM server or
client agent.  The TSM return code is -50.
```

b. In the Cluster Administrator, TONGA goes down and SENEGAL begins to bring the resources online.

c. When the Tivoli Storage Manager server instance resource is online (now hosted by SENEGAL), the client backup restarts again as shown on the schedule log file in Example 5-22.

Example 5-22 Schedule log file when backup is restarted on the client

```
02/07/2005 14:49:38 ANS1809W Session is lost; initializing session reopen
procedure.
02/07/2005 14:58:49 ... successful
02/07/2005 14:58:49 Retry # 1  Normal File-->    549,376
\\senegal\c$\WINDOWS\system32\printui.dll [Sent]
02/07/2005 14:58:49 Retry # 1  Normal File-->     55,340
\\senegal\c$\WINDOWS\system32\prncnfg.vbs [Sent]
02/07/2005 14:58:49 Retry # 1  Normal File-->     25,510
\\senegal\c$\WINDOWS\system32\prndrvr.vbs [Sent]
02/07/2005 14:58:49 Retry # 1  Normal File-->     35,558
\\senegal\c$\WINDOWS\system32\prnjobs.vbs [Sent]
02/07/2005 14:58:49 Retry # 1  Normal File-->     43,784
\\senegal\c$\WINDOWS\system32\prnmngr.vbs [Sent]
```

d. The following messages in Example 5-23 are received on the Tivoli Storage Manager server activity log after restarting.

Example 5-23 Activity log after the server is restarted

```
02/07/2005 14:58:48   ANR4726I The NAS-NDMP support module has been loaded.
02/07/2005 14:58:48   ANR4726I The Centera support module has been loaded.
02/07/2005 14:58:48   ANR4726I The ServerFree support module has been loaded.
02/07/2005 14:58:48   ANR2803I License manager started.
02/07/2005 14:58:48   ANR8260I Named Pipes driver ready for connection with
                       clients.
02/07/2005 14:58:48   ANR8200I TCP/IP driver ready for connection with clients
                       on port 1500.
02/07/2005 14:58:48   ANR8280I HTTP driver ready for connection with clients on
                       port 1580.
```

```
02/07/2005 14:58:48   ANR0984I Process 1 for EXPIRATION started in the
                       BACKGROUND at 14:58:48. (PROCESS: 1)
02/07/2005 14:58:48   ANR0993I Server initialization complete.
02/07/2005 14:58:48   ANR2560I Schedule manager started.
02/07/2005 14:58:48   ANR4747W The web administrative interface is no longer
                       supported. Begin using the Integrated Solutions Console
                       instead.
02/07/2005 14:58:48   ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK1.DSM varied
                       online.
02/07/2005 14:58:48   ANR0811I Inventory client file expiration started as
                       process 1. (PROCESS: 1)
02/07/2005 14:58:48   ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli is
                       now ready for use.
02/07/2005 14:58:48   ANR2828I Server is licensed to support Tivoli Storage
                       Manager Basic Edition.
02/07/2005 14:58:48   ANR0984I Process 2 for AUDIT LICENSE started in the
                       BACKGROUND at 14:58:48. (PROCESS: 2)
02/07/2005 14:58:48   ANR2820I Automatic license audit started as process 2.
                       (PROCESS: 2)
02/07/2005 14:58:48   ANR0812I Inventory file expiration process 1 completed:
                       examined 1 objects, deleting 0 backup objects, 0 archive
                       objects, 0 DB backup volumes, and 0 recovery plan files.
                       0 errors were encountered. (PROCESS: 1)
02/07/2005 14:58:48   ANR0985I Process 1 for EXPIRATION running in the
                       BACKGROUND completed with completion state SUCCESS at
                       14:58:48. (PROCESS: 1)
02/07/2005 14:58:48   ANR2825I License audit process 2 completed successfully -
                       2 nodes audited.  (PROCESS: 2)
02/07/2005 14:58:48   ANR0987I Process 2 for AUDIT LICENSE running in the
                       BACKGROUND processed 2 items with a completion state of
                       SUCCESS at 14:58:48. (PROCESS: 2)
02/07/2005 14:58:49   ANR0406I Session 1 started for node SENEGAL (WinNT)
```

6. When the backup ends the client sends the statistics messages we show on the schedule log file in Example 5-24.

Example 5-24 Schedule log file shows backup statistics on the client

```
02/07/2005 15:05:47 Successful incremental backup of 'System Services'

02/07/2005 15:05:47 --- SCHEDULEREC STATUS BEGIN
02/07/2005 15:05:47 Total number of objects inspected:   15,797
02/07/2005 15:05:47 Total number of objects backed up:    2,709
02/07/2005 15:05:47 Total number of objects updated:          4
02/07/2005 15:05:47 Total number of objects rebound:          0
02/07/2005 15:05:47 Total number of objects deleted:          0
02/07/2005 15:05:47 Total number of objects expired:          4
02/07/2005 15:05:47 Total number of objects failed:           0
02/07/2005 15:05:47 Total number of bytes transferred:    879.32 MB
```

```
02/07/2005 15:05:47 Data transfer time:                    72.08 sec
02/07/2005 15:05:47 Network data transfer rate:        12,490.88 KB/sec
02/07/2005 15:05:47 Aggregate data transfer rate:      4,616.12 KB/sec
02/07/2005 15:05:47 Objects compressed by:                    0%
02/07/2005 15:05:47 Elapsed processing time:           00:03:15
02/07/2005 15:05:47 --- SCHEDULEREC STATUS END
02/07/2005 15:05:47 --- SCHEDULEREC OBJECT END DAILY_INCR 02/07/2005 14:45:00
02/07/2005 15:05:47 Scheduled event 'DAILY_INCR' completed successfully.
02/07/2005 15:05:47 Sending results for scheduled event 'DAILY_INCR'.
02/07/2005 15:05:47 Results sent to server for scheduled event 'DAILY_INCR'
```

Results summary

The test results show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a scheduled backup started from one client is restarted after the failover on the other node of the MSCS.

On the server event report, the schedule is shown as completed with a return code 8, as shown in Figure 5-156. This is due to the communication loss, but the backup ends successfully.

```
tsm: TSMSRV02>q event * * begind=-2 f=d

Policy Domain Name: STANDARD
     Schedule Name: DAILY_INCR
         Node Name: SENEGAL
   Scheduled Start: 02/07/2005 14:45:00
      Actual Start: 02/07/2005 14:45:03
         Completed: 02/07/2005 15:05:47
            Status: Completed
            Result: 8
            Reason: The operation completed with at least one warning
message.
```

Figure 5-156 Schedule result

Note: In the test we have just described, we used a disk storage pool as the destination storage pool. We also tested using a tape storage pool as destination and we got the same results. The only difference is that when the Tivoli Storage Manager server is again up, the tape volume it was using on the first node is unloaded from the tape drive and loaded again into the second drive, and the client receives a "media wait" message while this process takes place. After the tape volume is mounted the backup continues and ends successfully.

Archived

Testing a scheduled client restore

Our third test consists of a scheduled restore.

Objective

OUr objective here is to show what happens when a scheduled client restore is running and the node which hosts the Tivoli Storage Manager server fails.

Activities

We perform these tasks:

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager cluster group: TONGA.
2. We schedule a client restore operation using the Tivoli Storage Manager server scheduler and we associate the schedule to CL_MSCS02_SA, one of the virtual clients installed on this Windows 2003 MSCS.
3. When it is the scheduled time, the client starts a session for the restore operation, as we see on the activity log in Example 5-25.

Example 5-25 Restore starts in the event log

```
tsm: TSMSRV02>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name     Status
-------------------- -------------------- ------------- ------------- ---------
02/24/2005 16:27:08  02/24/2005 16:27:19  RESTORE       CL_MSCS02_SA  Started
```

4. The client starts restoring files as shown in its schedule log file in Example 5-26.

Example 5-26 Restore starts in the schedule log file of the client

```
Executing scheduled command now.
02/24/2005 16:27:19 --- SCHEDULEREC OBJECT BEGIN RESTORE 02/24/2005 16:27:08
02/24/2005 16:27:19 Restore function invoked.
02/24/2005 16:27:20 ANS1247I Waiting for files from the server...Restoring  0
\\cl_mscs02\j$\code\adminc [Done]
02/24/2005 16:27:21 Restoring  0 \\cl_mscs02\j$\code\drivers_lto [Done]
02/24/2005 16:27:21 Restoring  0 \\cl_mscs02\j$\code\isc [Done]
02/24/2005 16:27:21 Restoring  0 \\cl_mscs02\j$\code\lto2k3 [Done]
02/24/2005 16:27:21 Restoring  0 \\cl_mscs02\j$\code\storageagent [Done]
02/24/2005 16:27:21 Restoring  0 \\cl_mscs02\j$\code\drivers_lto\checked [Done]
02/24/2005 16:27:21 Restoring  0 \\cl_mscs02\j$\code\isc\RuntimeExt [Done]
02/24/2005 16:27:21 Restoring  0 \\cl_mscs02\j$\code\isc\tutorial [Done]
02/24/2005 16:27:21 Restoring  0 \\cl_mscs02\j$\code\isc\wps [Done]
02/24/2005 16:27:21 Restoring  0 \\cl_mscs02\j$\code\isc\RuntimeExt\eclipse
[Done]
02/24/2005 16:27:21 Restoring  0 \\cl_mscs02\j$\code\isc\RuntimeExt\ewase
[Done]
```

```
02/24/2005 16:27:21 Restoring  0
\\cl_mscs02\j$\code\isc\RuntimeExt\ewase_efixes [Done]
02/24/2005 16:27:21 Restoring  0
\\cl_mscs02\j$\code\isc\RuntimeExt\ewase_modification [Done]
02/24/2005 16:27:21 Restoring  0 \\cl_mscs02\j$\code\isc\RuntimeExt\misc [Done]
02/24/2005 16:27:21 Restoring  0 \\cl_mscs02\j$\code\isc\RuntimeExt\pzn [Done]
02/24/2005 16:27:21 Restoring  0 \\cl_mscs02\j$\code\isc\RuntimeExt\uninstall
[Done]
02/24/2005 16:27:21 Restoring  0
\\cl_mscs02\j$\code\isc\RuntimeExt\eclipse\windows [Done]
```

5. While the client continues receiving files from the server, we force TONGA to fail. The following sequence occurs:

   a. In the client, the session is lost temporarily and it starts the procedure to reopen a session with the server. We see this in its schedule log file in Example 5-27.

Example 5-27 The session is lost in the client

```
02/24/2005 16:27:31 Restoring 527,360
\\cl_mscs02\j$\code\drivers_lto\checked\ibmtp2k3.pdb [Done]
02/24/2005 16:27:31 Restoring 285,696
\\cl_mscs02\j$\code\drivers_lto\checked\ibmtp2k3.sys [Done]
02/24/2005 16:28:01 ANS1809W Session is lost; initializing session reopen
procedure.
```

   b. In the Cluster Administrator, SENEGAL begins to bring the resources online.

   c. When Tivoli Storage Manager server instance resource is online (now hosted by SENEGAL), the client reopens its session and the restore restarts from the point of the last committed transaction in the server database. We can see this in its schedule log file in Example 5-28.

Example 5-28 The client reopens a session with the server

```
02/24/2005 16:27:31 Restoring 285,696
\\cl_mscs02\j$\code\drivers_lto\checked\ibmtp2k3.sys [Done]
02/24/2005 16:28:01 ANS1809W Session is lost; initializing session reopen
procedure.
02/24/2005 16:28:36 ... successful
02/24/2005 16:28:36 ANS1247I Waiting for files from the server...Restoring
327,709,515 \\cl_mscs02\j$\code\isc\C8241ML.exe [Done]
02/24/2005 16:29:05 Restoring 20,763 \\cl_mscs02\j$\code\isc\dsminstall.jar
[Done]
02/24/2005 16:29:06 Restoring 6,484,490 \\cl_mscs02\j$\code\isc\ISCAction.jar
[Done]
```

Archived

d. The activity log shows the event as restarted as shown in Example 5-29.

Example 5-29 The schedule is restarted in the activity log

```
tsm: TSMSRV02>q ev * *
Session established with server TSMSRV02: Windows
  Server Version 5, Release 3, Level 0.0
  Server date/time: 02/24/2005 16:27:58  Last access: 02/24/2005 16:23:35

Scheduled Start      Actual Start         Schedule Name Node Name     Status
-------------------- -------------------- ------------- ------------- ---------
02/24/2005 16:27:08  02/24/2005 16:27:19  RESTORE       CL_MSCS02_SA  Restarted
```

6. The client ends the restore, it reports the restore statistics to the server, and it writes those statistics in its schedule log file as we can see in Example 5-30.

Example 5-30 Restore final statistics

```
02/24/2005 16:29:55 Restoring     111,755,569
\\cl_mscs02\j$\code\storageagent\c8117ml.exe [Done]
02/24/2005 16:29:55
Restore processing finished.
02/24/2005 16:29:57 --- SCHEDULEREC STATUS BEGIN
02/24/2005 16:29:57 Total number of objects restored:     1,864
02/24/2005 16:29:57 Total number of objects failed:           0
02/24/2005 16:29:57 Total number of bytes transferred:     1.31 GB
02/24/2005 16:29:57 Data transfer time:                  104.70 sec
02/24/2005 16:29:57 Network data transfer rate:        13,142.61 KB/sec
02/24/2005 16:29:57 Aggregate data transfer rate:      8,752.74 KB/sec
02/24/2005 16:29:57 Elapsed processing time:           00:02:37
02/24/2005 16:29:57 --- SCHEDULEREC STATUS END
02/24/2005 16:29:57 --- SCHEDULEREC OBJECT END RESTORE 02/24/2005 16:27:08
02/24/2005 16:29:57 --- SCHEDULEREC STATUS BEGIN
02/24/2005 16:29:57 --- SCHEDULEREC STATUS END
02/24/2005 16:29:57 ANS1512E Scheduled event ‘RESTORE’ failed.  Return code =
12.
02/24/2005 16:29:57 Sending results for scheduled event ‘RESTORE’.
02/24/2005 16:29:57 Results sent to server for scheduled event ‘RESTORE’.
```

7. In the activity log, the event figures as “failed” with return code = 12 as shown in Example 5-31.

Example 5-31 The activity log shows the event failed

```
tsm: TSMSRV02>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name     Status
-------------------- -------------------- ------------- ------------- ---------
02/24/2005 16:27:08  02/24/2005 16:27:19  RESTORE       CL_MSCS02_SA  Failed
```

Results summary

The test results show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a scheduled restore started from one client is restarted after the server is again up in the second node of the MSCS.

Depending on the amount of data being restored before the failure of the Tivoli Storage Manager server, the schedule ends as “failed” or it can also end as “completed”.

If the Tivoli Storage Manager server committed the transaction for the files already restored to the client, when the server starts again in the second node of the MSCS, the client restarts the restore from the point of failure. However, since there was a failure and the session was lost by the client, the event shows “failed” and it reports a return code 12. However, the restore worked correctly and there were no files missing.

If the Tivoli Storage Manager server did not commit the transaction for the files already restored to the client, when the server starts again in the second node of the MSCS, the session for the restore operation is not reopened by the client and the schedule log file does not report any information after the failure. The restore session is marked as restartable on the Tivoli Storage Manager server, and it is necessary to restart the scheduler in the client. When the scheduler starts, if the startup window is not elapsed, the client restores the files from the beginning. If the scheduler starts when the startup window elapsed, the restore is still in a restartable state.

If the client starts a manual session with the server (using the command line or the GUI) while the restore is in a restartable state, it can restore the rest of the files. If the timeout for the restartable restore session expires, the restore cannot be restarted.

Testing migration from disk storage pool to tape storage pool

This time we test a server process: migration from disk storage pool to tape storage pool.

Objective

The objective of this test is to show what happens when a disk storage pool migration process starts on the Tivoli Storage Manager server and the node that hosts the server instance fails.

Activities

For this test, we perform these tasks:

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager cluster group: TONGA.

2. We update the disk storage pool (SPD_BCK) high threshold migration to 0. This forces migration of backup versions to its next storage pool, a tape storage pool (SPT_BCK).
3. A process starts for the migration and Tivoli Storage Manager prompts the tape library to mount a tape volume as shown in Example 5-32.

Example 5-32 Disk storage pool migration started on server

```
02/08/2005 17:07:19 ANR1000I Migration process 3 started for storage pool
                     SPD_BCK automatically, highMig=0, lowMig=0, duration=No.
                     (PROCESS: 3)
02/08/2005 17:07:19 ANR0513I Process 3 opened output volume 026AKKL2.
                     (PROCESS: 3)
02/08/2005 17:07:21 ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                     (SESSION: 1)
```

4. While migration is running we force a failure on TONGA. When the Tivoli Storage Manager Server instance resource is online (hosted by SENEGAL), the tape volume is unloaded from the drive. Since the high threshold is still 0, a new migration process is started and the server prompts to mount the same tape volume as shown in Example 5-33.

Example 5-33 Disk storage pool migration started again on the server

```
02/08/2005 17:08:30  ANR0984I Process 2 for MIGRATION started in the BACKGROUND
                     at 17:08:30. (PROCESS: 2)
02/08/2005 17:08:30  ANR1000I Migration process 2 started for storage pool
                     SPT_BCK automatically, highMig=0, lowMig=0, duration=No.
                     (PROCESS: 2)
02/08/2005 17:09:17  ANR8439I SCSI library LIBLTO is ready for operations.
02/08/2005 17:09:42  ANR8337I LTO volume 026AKKL2 mounted in drive
DRIVE1
                     (mt0.0.0.2). (PROCESS: 2)
02/08/2005 17:09:42  ANR0513I Process 2 opened output volume 026AKKL2.
                     (PROCESS: 2)
02/08/2005 17:09:51  ANR2017I Administrator ADMIN issued command:
QUERY MOUNT
                     (SESSION: 1)
```

```
02/08/2005 17:09:51  ANR8330I LTO volume 026AKKL2 is mounted R/W in drive
                     DRIVE1 (mt0.0.0.2), status: IN USE. (SESSION: 1)
02/08/2005 17:09:51  ANR8334I        1 matches found. (SESSION: 1)
```

Attention: the migration process is not really restarted when the server failover occurs, as we can see comparing the process numbers for migration between Example 5-32 and Example 5-33. However, the tape volume is unloaded correctly after the failover and loaded again when the new migration process starts on the server.

5. The migration ends successfully as we show on the activity log taken from the server in Example 5-34.

Example 5-34 Disk storage pool migration ends successfully

```
02/08/2005 17:12:04  ANR1001I Migration process 2 ended for storage pool
                      SPT_BCK. (PROCESS: 2)
02/08/2005 17:12:04  ANR0986I Process 2 for MIGRATION running in the BACKGROUND
                      processed 1593 items for a total of 277,057,536 bytes
                      with a completion state of SUCCESS at 17:10:04. (PROCESS:
                      2)
```

Results summary

The results of our test show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a migration process started on the server before the failure, starts again using a new process number when the second node on the MSCS brings the Tivoli Storage Manager server instance online.

Testing backup from tape storage pool to copy storage pool

In this section we test another internal server process, backup from a tape storage pool to a copy storage pool.

Objective

The objective of this test is to show what happens when a backup storage pool process (from tape to tape) starts on the Tivoli Storage Manager server and the node that hosts the resource fails.

Activities

For this test, we perform these tasks:

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager cluster group: TONGA.
2. We run the following command to start an storage pool backup from our primary tape storage pool SPT_BCK to our copy storage pool SPCPT_BCK:

   ```
   ba stg spt_bck spcpt_bck
   ```

3. A process starts for the storage pool backup and Tivoli Storage Manager prompts to mount two tape volumes as shown in Example 5-35.

Example 5-35 Starting a backup storage pool process

```
02/09/2005 08:50:19  ANR2017I Administrator ADMIN issued command:
BACKUP

                STGPOOL spt_bck spcpt_bck  (SESSION: 1)

02/09/2005 08:50:19  ANR0984I Process 1 for BACKUP STORAGE POOL
started in the

                BACKGROUND at 08:50:19. (SESSION: 1, PROCESS: 1)

02/09/2005 08:50:19  ANR2110I BACKUP STGPOOL started as process 1.
(SESSION: 1,

                PROCESS: 1)

02/09/2005 08:50:19  ANR1210I Backup of primary storage pool SPT_BCK to
copy

                storage pool SPCPT_BCK started as process 1. (SESSION: 1,

                PROCESS: 1)

02/09/2005 08:50:19  ANR1228I Removable volume 026AKKL2 is required for
storage

                pool backup. (SESSION: 1, PROCESS: 1)

02/09/2005 08:50:31  ANR2017I Administrator ADMIN issued command:
QUERY MOUNT

                (SESSION: 1)

02/09/2005 08:50:31  ANR8379I Mount point in device class LTOCLASS1 is
waiting

                for the volume mount to complete, status: WAITING FOR
```

```
                    VOLUME. (SESSION: 1)

02/09/2005 08:50:31  ANR8379I Mount point in device class LTOCLASS1 is
waiting

                    for the volume mount to complete, status: WAITING FOR

                    VOLUME. (SESSION: 1)

02/09/2005 08:50:31  ANR8334I        2 matches found. (SESSION: 1)

02/09/2005 08:51:18  ANR8337I LTO volume 025AKKL2 mounted in drive
DRIVE1

                    (mt0.0.0.2). (SESSION: 1, PROCESS: 1)

02/09/2005 08:51:20  ANR8337I LTO volume 026AKKL2 mounted in drive
DRIVE2

                    (mt1.0.0.2). (SESSION: 1, PROCESS: 1)

02/09/2005 08:51:20  ANR1340I Scratch volume 025AKKL2 is now defined in
storage

                    pool SPCPT_BCK. (SESSION: 1, PROCESS: 1)

02/09/2005 08:51:20  ANR0513I Process 1 opened output volume 025AKKL2.

                    (SESSION: 1, PROCESS: 1)

02/09/2005 08:51:20  ANR0512I Process 1 opened input volume 026AKKL2.
(SESSION:

                    1, PROCESS: 1)
```

---

4. While the process is started and the two tape volumes are mounted on both drives, we force a failure on TONGA. When the Tivoli Storage Manager Server instance resource is online (hosted by SENEGAL), both tape volumes are unloaded from the drives and there is no process started in the activity log.
5. The backup storage pool process does not restart again unless we start it manually.
6. If the backup storage pool process sent enough data before the failure so that the server was able to commit the transaction in the database, when the Tivoli Storage Manager server starts again in the second node, those files already

copied in the copy storage pool tape volume and committed in the server database, are valid copied versions.

However, there are still files not copied from the primary tape storage pool. If we want to be sure that the server copies all the files from this primary storage pool, we need to repeat the command. Those files committed as copied in the database will not be copied again.

This happens both using roll-forward recovery log mode as well as normal recovery log mode.

7. If the backup storage pool task did not process enough data to commit the transaction into the database, when the Tivoli Storage Manager server starts again in the second node, those files copied in the copy storage pool tape volume before the failure are not recorded in the Tivoli Storage Manager server database. So, if we start a new backup storage pool task, they will be copied again.

   If the tape volume used for the copy storage pool before the failure was taken from the scratch pool in the tape library, (as in our case), it is given back to scratch status in the tape library.

   If the tape volume used for the copy storage pool before the failure had already data belonging to back up storage pool tasks from other days, the tape volume is kept in the copy storage pool but the new information written on it, is not valid.

   If we want to be sure that the server copies all the files from this primary storage pool, we need to repeat the command.

   This happens both using roll-forward recovery log mode as well as normal recovery log mode.

Results summary

The results of our test show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a backup storage pool process (from tape to tape) started on the server before the failure, does not restart when the second node on the MSCS brings the Tivoli Storage Manager server instance online.

Both tapes are correctly unloaded from the tape drives when the Tivoli Storage Manager server is again online, but the process is not restarted unless we run again the command.

Depending on the amount of data already sent when the task failed, (if it was committed to the database or not), the files backed up into the copy storage pool tape volume before the failure, will be reflected on the database or will be not.

If enough information was copied to the copy storage pool tape volume so that the transaction was committed before the failure, when the server restarts in the

second node the information is recorded in the database and the files figure as valid copies.

If the transaction was not committed to the database, there is no information in the database about the process and the files copied into the copy storage pool before the failure, will need to be copied again.

This situation happens either if the recovery log is set to roll-forward mode or it is set to normal mode.

In any of the cases to be sure that all information is copied from the primary storage pool to the copy storage pool, we should repeat the command.

There is no difference between a scheduled backup storage pool process or a manual process using the administrative interface. In our lab we tested both methods and the results were the same.

Testing server database backup

The following test consists of a server database backup.

Objective

The objective of this test is to show what happens when a Tivoli Storage Manager server database backup process is started on the Tivoli Storage Manager server and the node that hosts the resource fails.

Activities

For this test, we perform these tasks (see Example 5-36).

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager cluster group: SENEGAL.
2. We run the following command to start a full database backup:

   ```
   ba db t=full devc=cllto_1
   ```

3. A process starts for database backup and Tivoli Storage Manager mounts a tape.

Example 5-36 Starting a database backup on the server

```
02/08/2005 21:12:25     ANR2017I Administrator ADMIN issued command: BACKUP DB
                           devcl=cllto_2 type=f  (SESSION: 2)
02/08/2005 21:12:25     ANR0984I Process 1 for DATABASE BACKUP started in the
                        BACKGROUND at 21:12:25. (SESSION: 2, PROCESS: 1)
02/08/2005 21:12:25     ANR2280I Full database backup started as process 1.
                         (SESSION: 2, PROCESS: 1)
02/08/2005 21:12:53     ANR8337I LTO volume 027AKKL2 mounted in drive DRIVE1
                         (mt0.0.0.2). (SESSION: 2, PROCESS: 1)
02/08/2005 21:12:53     ANR0513I Process 1 opened output volume 027AKKL2.
```

```
                                    (SESSION: 2, PROCESS: 1)
```

4. While the backup is running we force a failure on SENEGAL. When the Tivoli Storage Manager Server is restarted in TONGA, the tape volume is unloaded from the drive, but the process is not restarted, as we can see in Example 5-37.

Example 5-37 After the server is restarted database backup does not restart

```
02/08/2005 21:13:19  ANR4726I The NAS-NDMP support module has been loaded.
02/08/2005 21:13:19  ANR4726I The Centera support module has been loaded.
02/08/2005 21:13:19  ANR4726I The ServerFree support module has been loaded.
02/08/2005 21:13:19  ANR2803I License manager started.
02/08/2005 21:13:19  ANR0993I Server initialization complete.
02/08/2005 21:13:19  ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli is
                      now ready for use.
02/08/2005 21:13:19  ANR2828I Server is licensed to support Tivoli Storage
                      Manager Basic Edition.
02/08/2005 21:13:19  ANR2560I Schedule manager started.
02/08/2005 21:13:19  ANR8260I Named Pipes driver ready for connection with
                      clients.
02/08/2005 21:13:19  ANR8280I HTTP driver ready for connection with clients on
                      port 1580.
02/08/2005 21:13:19  ANR8200I TCP/IP driver ready for connection with clients
                      on port 1500.
02/08/2005 21:13:19  ANR4747W The web administrative interface is no longer
                      supported. Begin using the Integrated Solutions Console
                      instead.
02/08/2005 21:13:19  ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK3.DSM varied
                      online.
02/08/2005 21:13:19  ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK1.DSM varied
                      online.
02/08/2005 21:13:19  ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK2.DSM varied
                      online.
02/08/2005 21:13:42  ANR0407I Session 1 started for administrator ADMIN (WinNT)
                      (Tcp/Ip tsmsrv02.tsmw2003.com(2233)). (SESSION: 1)
02/08/2005 21:13:46  ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                      (SESSION: 1)
02/08/2005 21:13:46  ANR0944E QUERY PROCESS: No active processes found.
                      (SESSION: 1)
```

5. If we want to do a database backup, we can start it now with the same command we used before.

6. If we query the volume history file, there is no record for that tape volume. However, if we query the library inventory the tape volume is in private status and it was last used for dbbackup.

Archived

7. We update the library inventory to change the status to scratch and then we run a new database backup.

Results summary

The results of our test show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a database backup process started on the server before the failure, does not restart when the second node on the MSCS brings the Tivoli Storage Manager server instance online.

The tape volume is correctly unloaded from the tape drive where it was mounted when the Tivoli Storage Manager server is again online, but the process is not restarted unless we run the command.

There is no difference between a scheduled process or a manual process using the administrative interface.

> **Important:** the tape volume used for the database backup before the failure is not useful. It is reported as a private volume in the library inventory but it is not recorded as valid backup in the volume history file. It is necessary to update the tape volume in the library inventory to scratch and start again a new database backup process.

Testing inventory expiration

In this section we test another server task: an inventory expiration process.

Objective

The objective of this test is to show what happens when Tivoli Storage Manager server is running the inventory expiration process and the node that hosts the server instance fails.

Activities

For this test, we perform these tasks:

1. We open the Cluster Administrator to check which node hosts the Tivoli Storage Manager cluster group: TONGA.
2. We to run the following command to start an inventory expiration process:

   ```
   expire inventory
   ```

3. A process starts for inventory expiration as shown in Example 5-38.

Example 5-38 Starting inventory expiration

```
02/09/2005 10:00:31    ANR2017I Administrator ADMIN issued command:
EXPIRE
```

```
                    INVENTORY  (SESSION: 20)

02/09/2005 10:00:31   ANR0984I Process 1 for EXPIRE INVENTORY started in
the

                    BACKGROUND at 10:00:31. (SESSION: 20, PROCESS: 1)

02/09/2005 10:00:31   ANR0811I Inventory client file expiration started as

                    process 1. (SESSION: 20, PROCESS: 1)

02/09/2005 10:00:31   ANR4391I Expiration processing node SENEGAL,
filespace

                    SYSTEM STATE, fsId 6, domain STANDARD, and management

                    class DEFAULT - for BACKUP type files. (SESSION: 20,

                    PROCESS: 1)

02/09/2005 10:00:31   ANR4391I Expiration processing node SENEGAL,
filespace

                    SYSTEM SERVICES, fsId 7, domain STANDARD, and
management

                    class DEFAULT - for BACKUP type files. (SESSION: 20,

                    PROCESS: 1)

02/09/2005 10:00:33   ANR4391I Expiration processing node SENEGAL,
filespace

                    \\senegal\c$, fsId 8, domain STANDARD, and management

                    class DEFAULT - for BACKUP type files. (SESSION: 20,

                    PROCESS: 1)
```

---

4. While Tivoli Storage Manager server is expiring objects, we force a failure on TONGA. When the Tivoli Storage Manager Server instance resource is online on SENEGAL, the inventory expiration process restarted. There are no errors in the activity log, just the process is not running, as shown in Example 5-39.

Archived

Example 5-39 No inventory expiration process after the failover

```
02/09/2005 10:01:07   ANR4726I The NAS-NDMP support module has been loaded.
02/09/2005 10:01:07   ANR4726I The Centera support module has been loaded.
02/09/2005 10:01:07   ANR4726I The ServerFree support module has been loaded.
02/09/2005 10:01:07   ANR8843E Initialization failed for SCSI library LIBLTO -
                       the library will be inaccessible.
02/09/2005 10:01:07   ANR8441E Initialization failed for SCSI library LIBLTO.
02/09/2005 10:01:07   ANR2803I License manager started.
02/09/2005 10:01:07   ANR2828I Server is licensed to support Tivoli Storage
                       Manager Basic Edition.
02/09/2005 10:01:07   ANR8280I HTTP driver ready for connection with clients on
                       port 1580.
02/09/2005 10:01:07   ANR4747W The web administrative interface is no longer
                      supported. Begin using the Integrated Solutions Console
                       instead.
02/09/2005 10:01:07   ANR0993I Server initialization complete.
02/09/2005 10:01:07   ANR2560I Schedule manager started.
02/09/2005 10:01:07   ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli is
                       now ready for use.
02/09/2005 10:01:07   ANR8200I TCP/IP driver ready for connection with clients
                       on port 1500.
02/09/2005 10:01:07   ANR8260I Named Pipes driver ready for connection with
                       clients.
02/09/2005 10:01:07   ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK1.DSM varied
                       online.
02/09/2005 10:01:07   ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK4.DSM varied
                       online.
02/09/2005 10:01:07   ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK2.DSM varied
                       online.
02/09/2005 10:01:07   ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK6.DSM varied
                       online.
02/09/2005 10:01:07   ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK3.DSM varied
                       online.
02/09/2005 10:01:07   ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK5.DSM varied
                       online.
02/09/2005 10:01:13   ANR0407I Session 1 started for administrator ADMIN
(WinNT)
                      (Tcp/Ip tsmsrv02.tsmw2003.com(3326)). (SESSION: 1)
02/09/2005 10:01:27   ANR0407I Session 2 started for administrator ADMIN
(WinNT)
                      (Tcp/Ip tsmsrv02.tsmw2003.com(3327)). (SESSION: 2)
02/09/2005 10:01:30   ANR2017I Administrator ADMIN issued command: QUERY
PROCESS
                       (SESSION: 2)
02/09/2005 10:01:30   ANR0944E QUERY PROCESS: No active processes found.
                       (SESSION: 2)
```

Archived

5. If we want to start the process again we just have to run the same command. Tivoli Storage Manager server runs the process and it ends successfully, such as shown in Example 5-40.

Example 5-40 Starting inventory expiration again

```
02/09/2005 10:01:33   ANR2017I Administrator ADMIN issued command: EXPIRE
                       INVENTORY  (SESSION: 2)
02/09/2005 10:01:33   ANR0984I Process 1 for EXPIRE INVENTORY started in the
                       BACKGROUND at 10:01:33. (SESSION: 2, PROCESS: 1)
02/09/2005 10:01:33   ANR0811I Inventory client file expiration started as
                       process 1. (SESSION: 2, PROCESS: 1)
02/09/2005 10:01:33   ANR4391I Expiration processing node SENEGAL, filespace
                       \\senegal\c$, fsId 8, domain STANDARD, and management
                       class DEFAULT - for BACKUP type files. (SESSION: 2,
                       PROCESS: 1)
02/09/2005 10:01:33  ANR4391I Expiration processing node SENEGAL, filespace
                       \\senegal\c$, fsId 8, domain STANDARD, and management
                       class DEFAULT - for BACKUP type files. (SESSION: 2,
                       PROCESS: 1)
02/09/2005 10:01:36  ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                       (SESSION: 2)
02/09/2005 10:01:46  ANR0407I Session 3 started for administrator ADMIN_CENTER
                       (DSMAPI) (Tcp/Ip 9.1.39.167(33681)). (SESSION: 3)
02/09/2005 10:01:46  ANR0418W Session 3 for administrator ADMIN_CENTER (DSMAPI)
                       is refused because an incorrect password was submitted.
                       (SESSION: 3)
02/09/2005 10:01:46  ANR0405I Session 3 ended for administrator ADMIN_CENTER
                       (DSMAPI). (SESSION: 3)
02/09/2005 10:01:56  ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                       (SESSION: 2)
02/09/2005 10:02:09  ANR4391I Expiration processing node SENEGAL, filespace
                       ASR, fsId 9, domain STANDARD, and management class
                       DEFAULT - for BACKUP type files. (SESSION: 2, PROCESS: 1)
02/09/2005 10:02:09  ANR4391I Expiration processing node SENEGAL, filespace
                       \\senegal\d$, fsId 10, domain STANDARD, and management
                       class DEFAULT - for BACKUP type files. (SESSION: 2,
                       PROCESS: 1)
02/09/2005 10:02:09  ANR4391I Expiration processing node TONGA, filespace
                       \\tonga\d$, fsId 5, domain STANDARD, and management class
                       DEFAULT - for BACKUP type files. (SESSION: 2, PROCESS: 1)
02/09/2005 10:02:14  ANR4391I Expiration processing node TONGA, filespace
                       \\tonga\c$, fsId 6, domain STANDARD, and management class
                       DEFAULT - for BACKUP type files. (SESSION: 2, PROCESS: 1)
02/09/2005 10:02:38  ANR4391I Expiration processing node KLCHV5D, filespace
                       \\klchv5d\c$, fsId 1, domain STANDARD, and management
                       class DEFAULT - for BACKUP type files. (SESSION: 2,
                       PROCESS: 1)
02/09/2005 10:02:38  ANR4391I Expiration processing node ROSANEG, filespace
```

```
                    \\rosaneg\c$, fsId 1, domain STANDARD, and management
                    class DEFAULT - for BACKUP type files. (SESSION: 2,
                    PROCESS: 1)
02/09/2005 10:02:38  ANR0812I Inventory file expiration process 1 completed:
                    examined 63442 objects, deleting 63429 backup objects, 0
                    Archive objects, 0 DB backup volumes, and 0 recovery plan
                    files. 0 errors were encountered. (SESSION: 2, PROCESS:
                    1)
02/09/2005 10:02:38  ANR0987I Process 1 for EXPIRE INVENTORY running in the
                    BACKGROUND processed 63429 items with a completion state
                    of SUCCESS at 10:02:38. (SESSION: 2, PROCESS: 1)
```

Results summary

The results of our test show that after a failure on the node that hosts the Tivoli Storage Manager server instance, an inventory expiration process started on the server before the failure, does not restart when the second node on the MSCS brings the Tivoli Storage Manager server instance online.

There is no error inside the Tivoli Storage Manager server database and we can restart the process again when the server is online.

5.7 Configuring ISC for clustering on Windows 2003

In 5.3.4, “Installation of the Administration Center” on page 92 we already described how we installed the Administration Center components on each node of the MSCS.

In this section we describe the method we use to configure the ISC as a clustered application on our MSCS Windows 2003. We need to create two new resources for the ISC services, in the cluster group where the shared disk used to install the code is located:

1. First we check that both nodes are again up and the two ISC services are stopped on them.
2. We open the Cluster Administrator menu and select the TSM Admin Center cluster group, the group that the shared disk j: belongs to. Then we select **New** →**Resource**, to create a new generic service resource as shown in Figure 5-157.

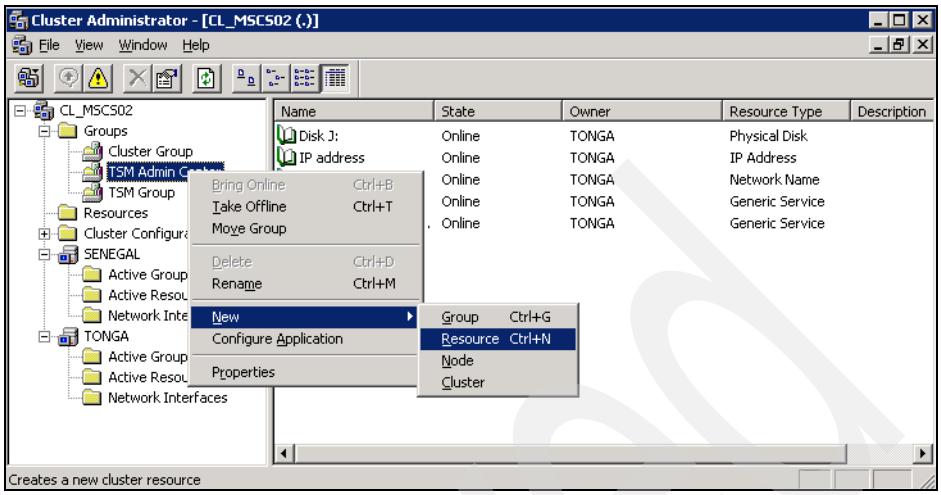

Figure 5-157 Defining a new resource for IBM WebSphere Application Server

3. We want to create a Generic Service resource related to the IBM WebSphere Application Server. We select a name for the resource and choose **Generic Service** as resource type in Figure 5-158, and we click **Next**:

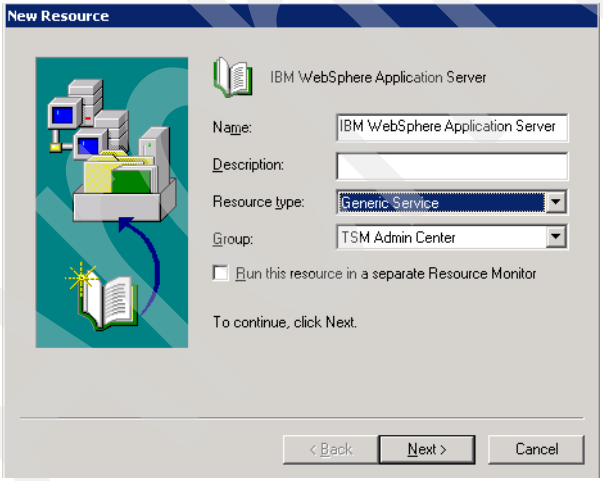

Figure 5-158 Specifying a resource name for IBM WebSphere application server

4. We leave both nodes as possible owners for the resource as shown in Figure 5-159 and we click **Next**.

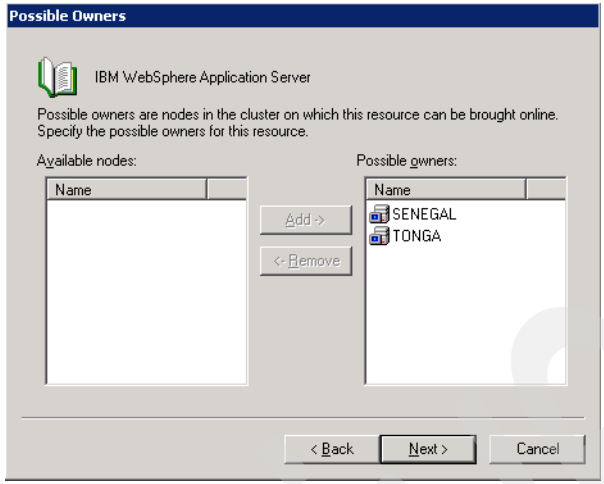

Figure 5-159 Possible owners for the IBM WebSphere application server resource

5. We select **Disk J** and **IP address** as dependencies for this resource and we click **Next** as shown in Figure 5-160.

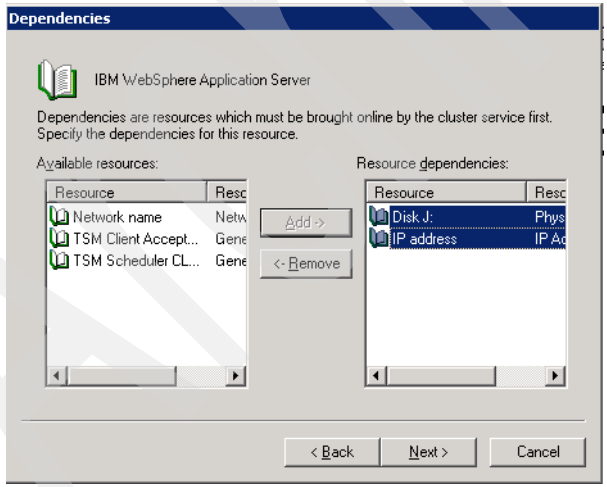

Figure 5-160 Dependencies for the IBM WebSphere application server resource

Important: the cluster group where the ISC services are defined must have an IP address resource. When the generic service is created using the Cluster Administrator menu, we use this IP address as dependency for the resource to be brought online. In this way when we start a Web browser to connect to the WebSphere Application server we use the IP for the cluster resource, instead of the local IP address for each node.

6. We type the real name of the IBM WebSphere Application Server service in Figure 5-161.

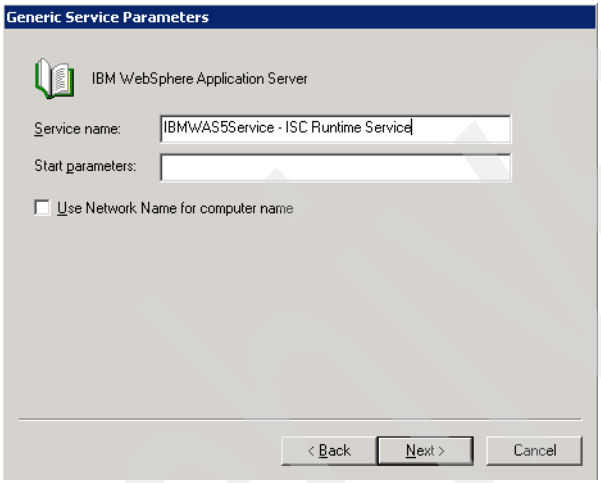

Figure 5-161 Specifying the same name for the service related to IBM WebSphere

Attention: make sure to specify the correct name in Figure 5-161. In the Windows services menu the name displayed for the service is not the real service name for it. Please, right-click the service and select **Properties** to check the service name for Windows.

7. We do not use any Registry key values to be replicated between nodes. We click **Next** in Figure 5-162.

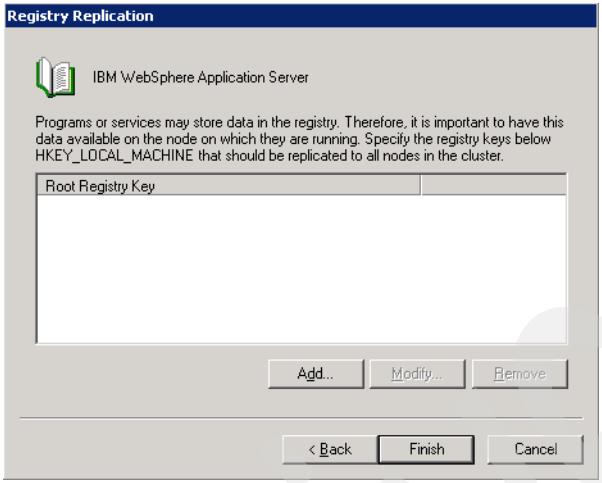

Figure 5-162 Registry replication values

8. The creation of the resource is successful as we can see in Figure 5-163. We click **OK** to finish.

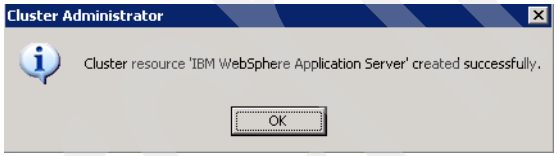

Figure 5-163 Successful creation of the generic resource

9. Now we bring this resource online.
10. The next task is the definition of a new Generic Service resource related to the ISC Help Service. We proceed using the same process as for the IBM WebSphere Application server.
11. We use ISC Help services as the name of the resource as shown in Figure 5-164.

Archived

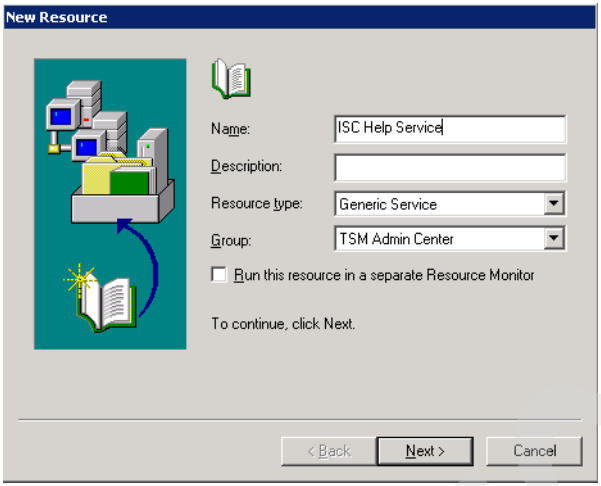

Figure 5-164 Selecting the resource name for ISC Help Service

12. As possible owners we select both nodes, in the dependencies menu we select the IBM WebSphere Application Server resource, and we do not use any Registry keys replication.
13. After the successful installation of the service, we bring it online using the Cluster Administrator menu.
14. At this moment both services are online in TONGA, the node that hosts the resources. To check that the configuration works correctly we proceed to move the resources to SENEGAL. Both services are now started in this node and stopped in TONGA.

5.7.1 Starting the Administration Center console

After the installation and configuration of ISC and administration center components in both nodes we are ready to start the Administration Center console to manage any Tivoli Storage Manager server.

We use the IP address related to the TSM Admin Center cluster group, which is the group where the ISC shared installation path is located.

1. In order to start an administrator Web session using the administrative client, we open a Web browser and type:

   `http://9.1.39.71:8421/ibm/console`

The login menu appears as shown in Figure 5-165.

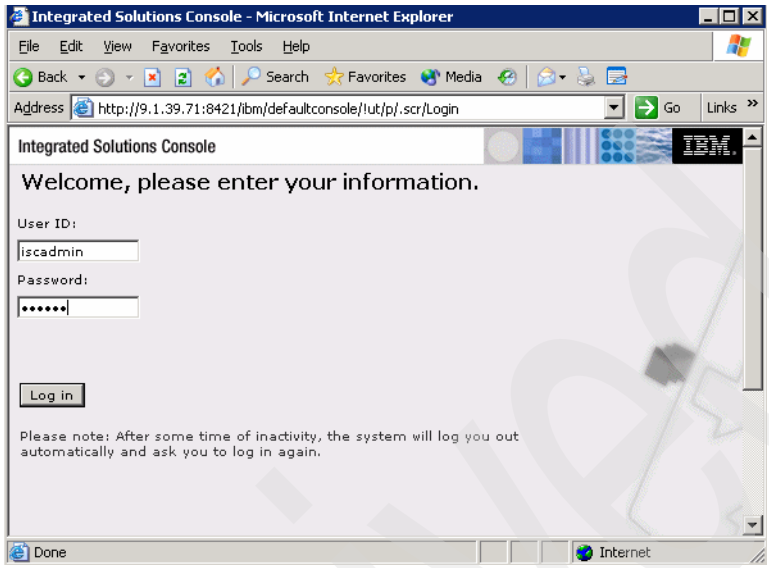

Figure 5-165 Login menu for the Administration Center

2. We type the user id and password we chose at ISC installation in Figure 5-26 and the following menu displays (Figure 5-166).

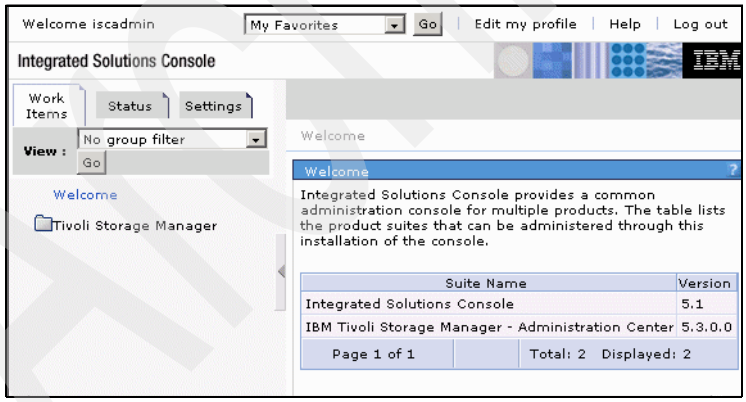

Figure 5-166 Administration Center

3. In Figure 5-166 we open the Tivoli Storage Manager folder on the right and the following menu displays (Figure 5-167).

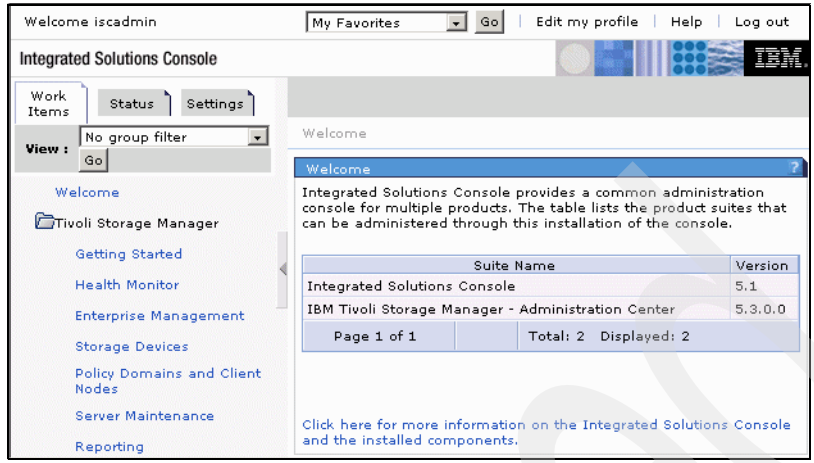

Figure 5-167 Options for Tivoli Storage Manager

4. We first need to create a new Tivoli Storage Manager server connection. To do this, we use Figure 5-167. We select Enterprise Management on that figure, and this takes us to the following menu (Figure 5-168).

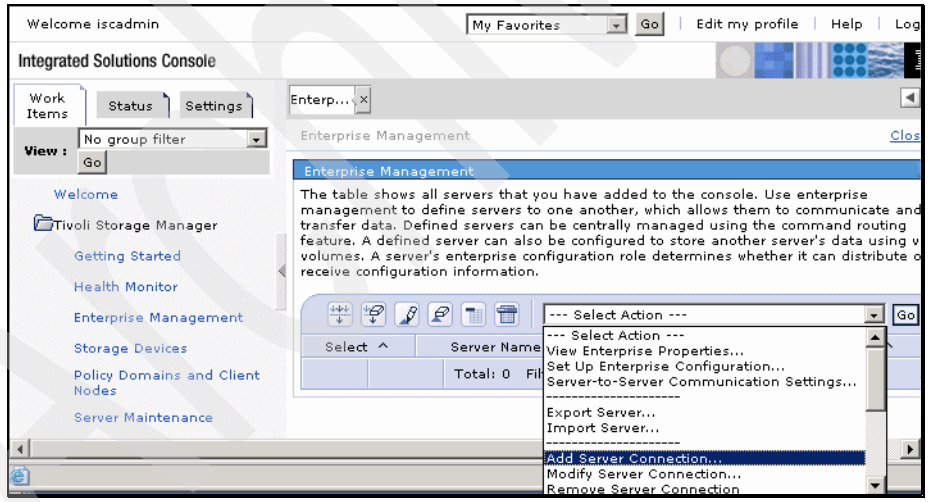

Figure 5-168 Selecting to create a new server connection

5. In Figure 5-168, if we open the pop-up menu such as we show, we have several options. To create a new server connection we select **Add Server Connection** and then we click **Go**.

The following menu displays (Figure 5-169).

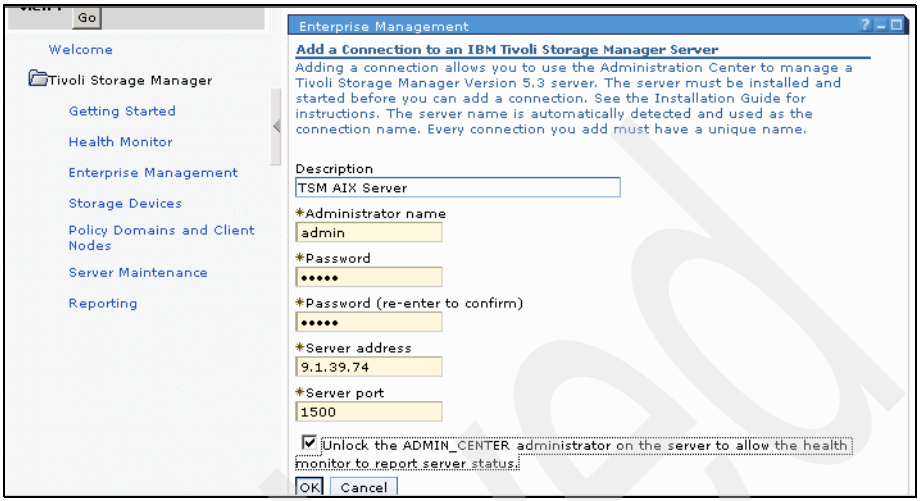

Figure 5-169 Specifying Tivoli Storage Manager server parameters

6. In Figure 5-169 we create a connection for a Tivoli Storage Manager server located in an AIX machine, whose name is TSMSRV03. We specify a Description (optional) as well as the Administrator name and Password to log into this server. We also specify the TCP/IP server address for our AIX server and its TCP port. Since we want to unlock the ADMIN_CENTER administrator to allow the health monitor to report server status, we check the box and then we click **OK**.

Archived

7. An information menu displays prompting to fill in the form below to configure the health monitor. We type the information such as shown in Figure 5-170.

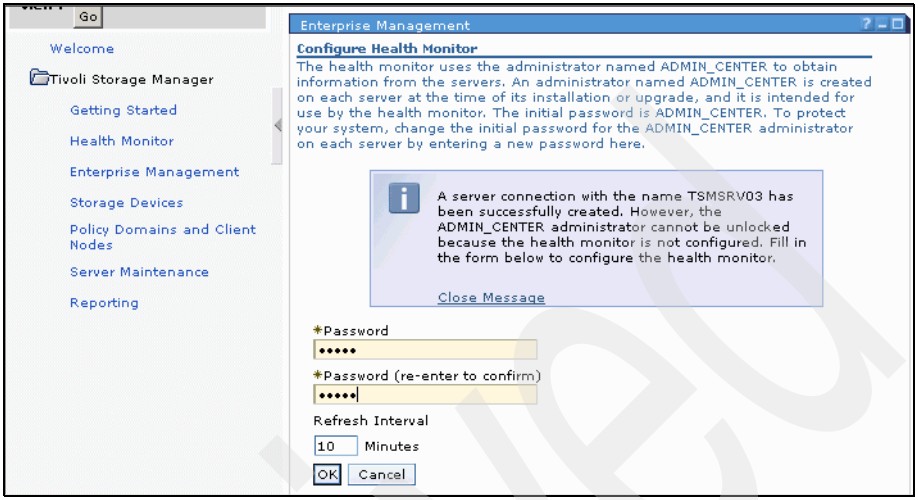

Figure 5-170 Filling a form to unlock ADMIN_CENTER

8. And finally, the panel shown in Figure 5-171 displays, where we can see the connection to TSMSRV03 server. We are ready to manage this server using the different options and commands provided by the Administration Center.

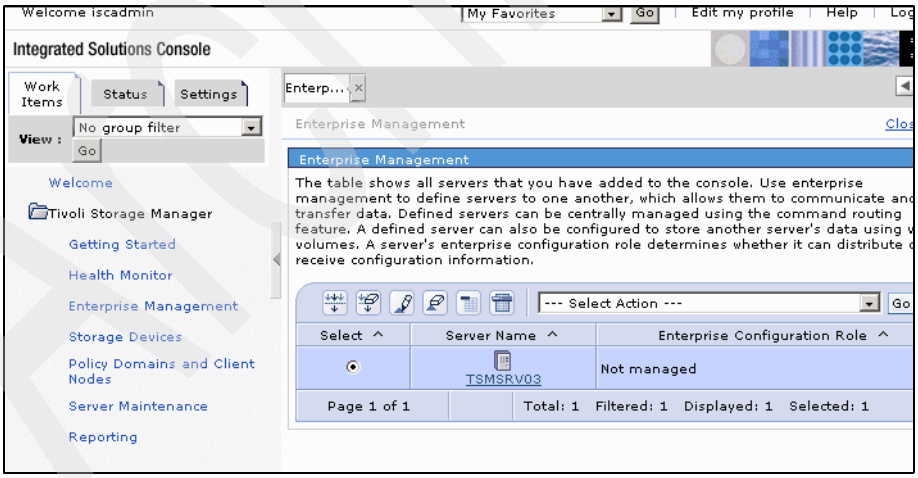

Figure 5-171 TSMSRV03 Tivoli Storage Manager server created

6

Microsoft Cluster Server and the IBM Tivoli Storage Manager Client

This chapter discusses how we set up Tivoli Storage Manager backup/archive client to work in a *Microsoft Cluster Services* (MSCS) for high availability.

We use two different environments:

- A Windows 2000 MSCS formed by two servers: POLONIUM and RADON
- A Windows 2003 MSCS formed by two servers: SENEGAL and TONGA.

© Copyright IBM Corp. 2005. All rights reserved.

6.1 Overview

When servers are set up in a cluster environment, applications can be active on different nodes at different times.

Tivoli Storage Manager backup/archive client is designed to support its implementation on an MSCS environment. However, it needs to be installed and configured following certain rules to run properly.

This chapter covers all the tasks we follow to achieve this goal.

6.2 Planning and design

We need to gather the following information to plan a backup strategy with Tivoli Storage Manager:

- Configuration of our cluster resource groups
- IP addresses and network names
- Shared disks that need to be backed up
- Tivoli Storage Manager nodenames used by each cluster group

Note:

- Service Pack 3 is required for backup and restore of SAN File Systems
- Windows 2000 hot fix 843198 is required to perform open file backup together with *Windows Encrypting File System* (EFS) files
- To back up the Windows 2003 system state or system services on local disks, Tivoli Storage Manager client must be connected to a Tivoli Storage Manager Version 5.2.0 or higher

We plan the names of the various services and resources so that they reflect our environment and ease our work.

6.3 Installing Tivoli Storage Manager client on MSCS

In order to implement Tivoli Storage Manager client to work correctly on a Windows 2000 MSCS or Windows 2003 MSCS environment to back up shared disk drives in the cluster, it is necessary to perform these tasks:

1. Installation of Tivoli Storage Manager client software components on each node of the MSCS, on local disk.

2. Configuration of Tivoli Storage Manager backup/archive client and Tivoli Storage Manager Web client for backup of local disks on each node.
3. Configuration of Tivoli Storage Manager backup/archive client and Tivoli Storage Manager Web client for backup of shared disks in the cluster.
4. Testing the Tivoli Storage Manager client clustering.

Some of these tasks are exactly the same for Windows 2000 or Windows 2003. For this reason, and to avoid duplicating the information, in this section we describe these common tasks. The specifics of each environment are described in sections “Tivoli Storage Manager client on Windows 2000” on page 248 and “Tivoli Storage Manager Client on Windows 2003” on page 289, also in this chapter.

6.3.1 Installation of Tivoli Storage Manager client components

The installation of Tivoli Storage Manager client on an MSCS Windows environment follows the same rules as in any single Windows machine. It is necessary to install the software on local disk in each node belonging to the same cluster.

In this section we describe this installation process. The same tasks apply to both Windows 2000 as well as Windows 2003 environments.

We use the same disk drive letter and installation path on each node:

```
c:\Program Files\Tivoli\tsm\baclient
```

To install the Tivoli Storage Manager client components we follow these steps:

1. On the first node of each MSCS, we run the **setup.exe** from the CD.
2. On the Choose Setup Language menu (Figure 6-1), we select the **English** language and click **OK**:

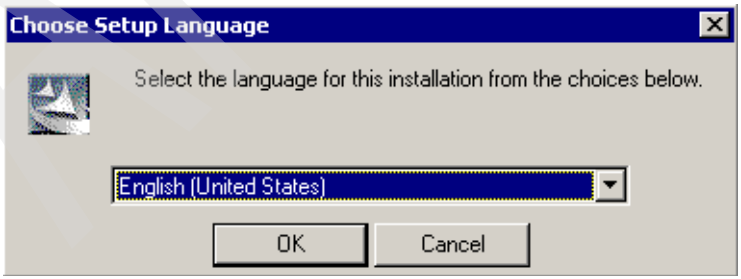

Figure 6-1 Setup language menu

3. The InstallShield Wizard for Tivoli Storage Manager Client displays (Figure 6-2). We click **Next**.

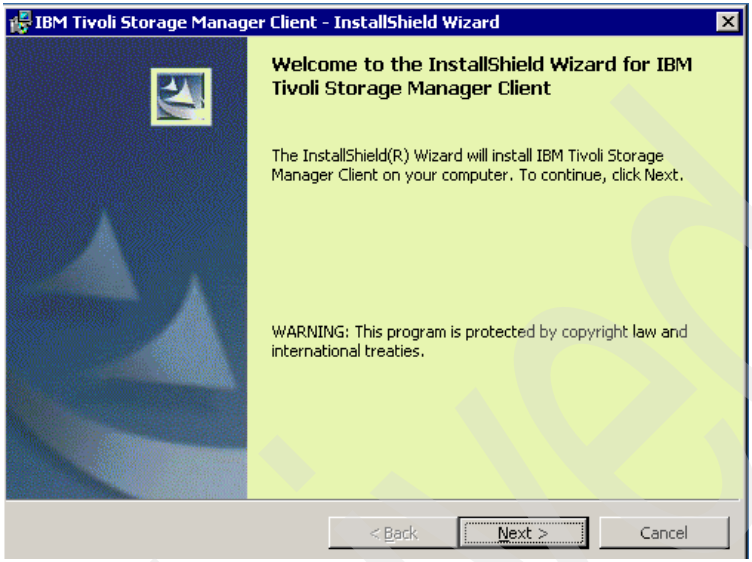

Figure 6-2 InstallShield Wizard for Tivoli Storage Manager Client

4. We choose the path where we want to install Tivoli Storage Manager backup/archive client. It is possible to select a local path or accept the default. We click **OK** (Figure 6-3).

Archived

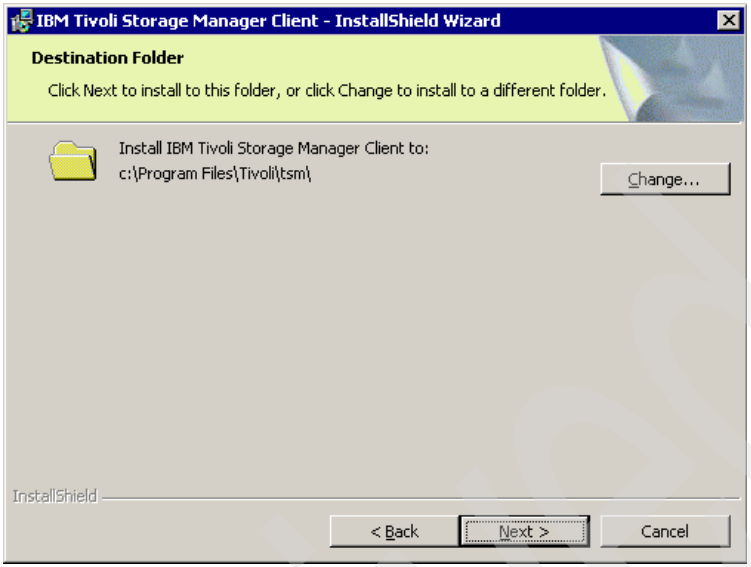

Figure 6-3 Installation path for Tivoli Storage Manager client

5. The next menu prompts for a Typical or Custom installation. Typical will install Tivoli Storage Manager GUI client, Tivoli Storage Manager command line client, and the API files. For our lab, we also want to install other components, so we select **Custom** and click **Next** (Figure 6-4).

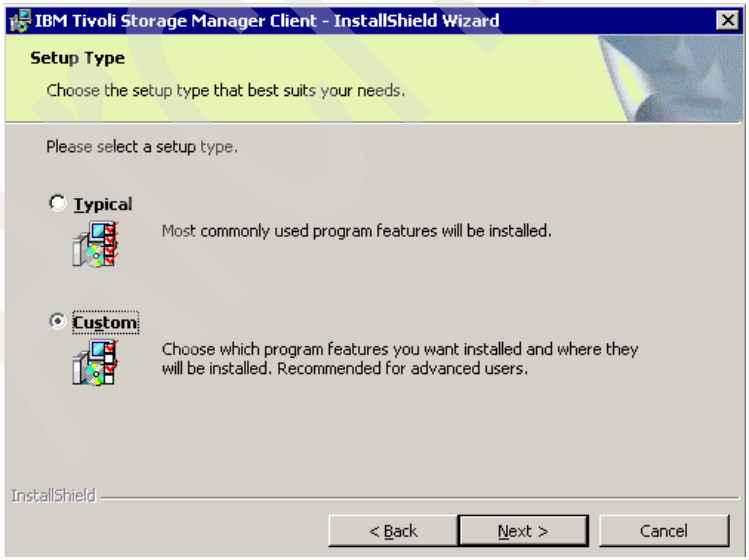

Figure 6-4 Custom installation

6. We select to install the **Administrative Client Command Line**, **Image Backup** and **Open File Support** packages. This choice depends on the actual environment (Figure 6-5).

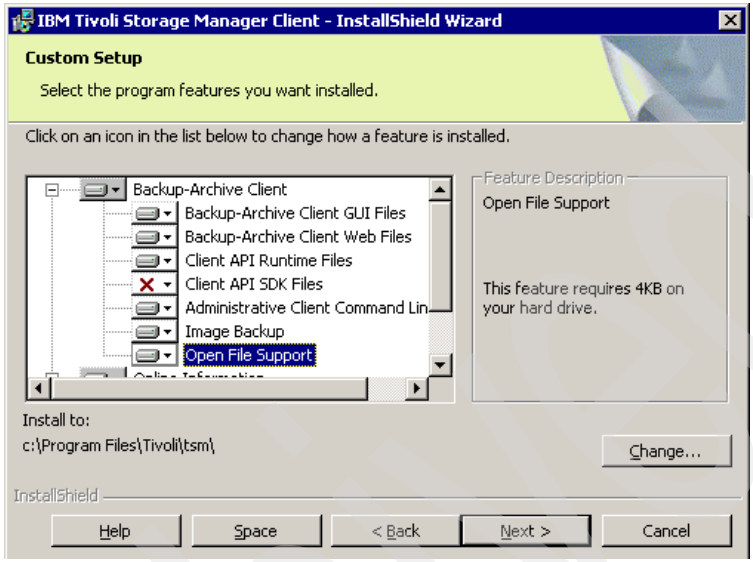

Figure 6-5 Custom setup

7. The system is now ready to install the software. We click **Install** (Figure 6-6).

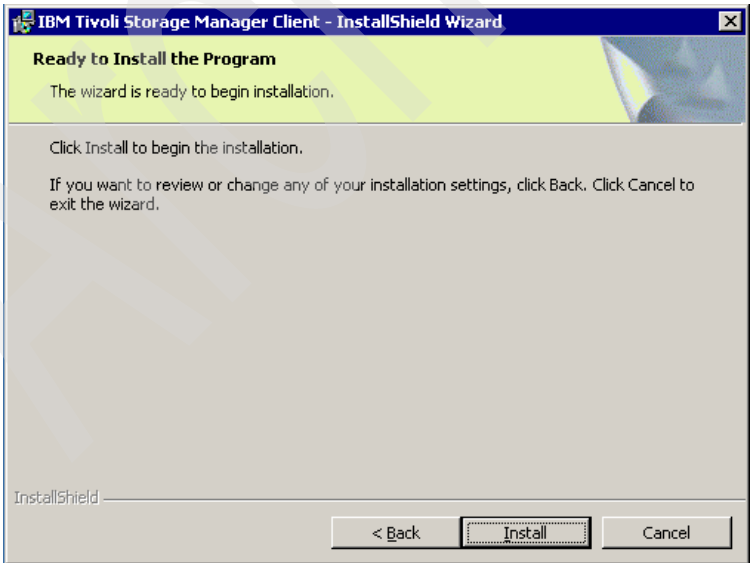

Figure 6-6 Start of installation of Tivoli Storage Manager client

8. The progress installation bar follows next (Figure 6-7).

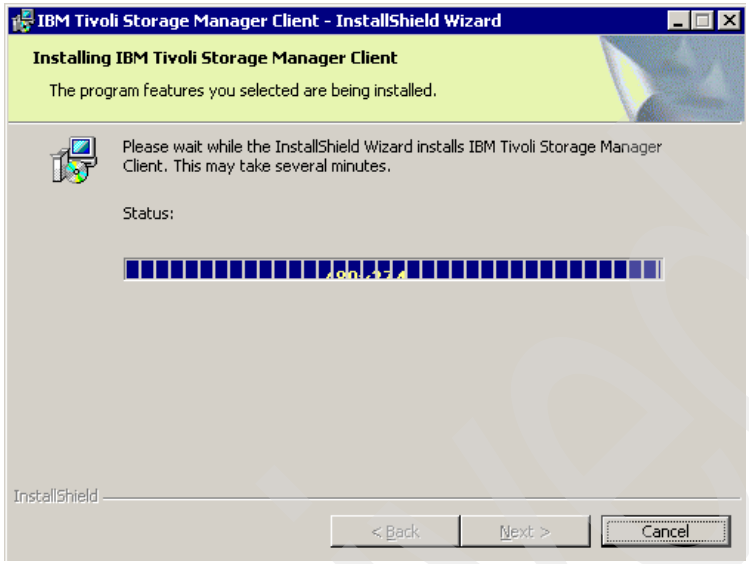

Figure 6-7 Status of the installation

9. When the installation ends we receive the following menu. We click **Finish** (Figure 6-8).

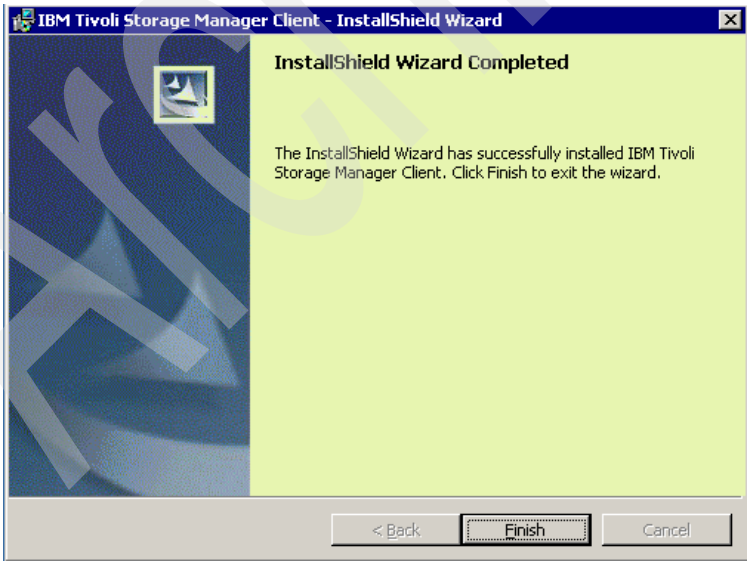

Figure 6-8 Installation completed

10. The system prompts to reboot the machine (Figure 6-9). If we can restart at this time, we should click **Yes**. If there are other applications running and it is not possible to restart the server now, we can do it later. We click **Yes**.

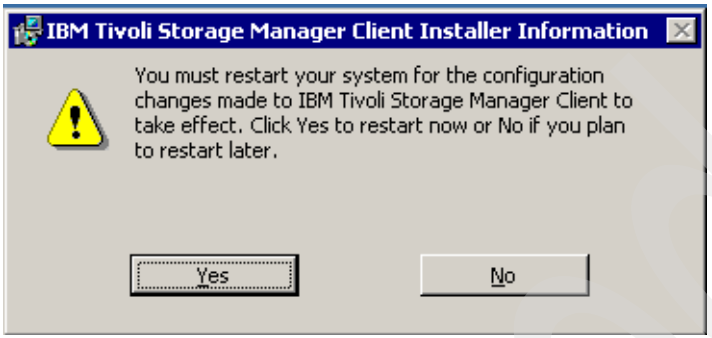

Figure 6-9 Installation prompts to restart the server

11. We repeat steps 1 to 10 for the second node of each MSCS, making sure to install Tivoli Storage Manager client on a local disk drive. We install it on the same path as the first node.

We follow all these tasks in our Windows 2000 MSCS (nodes POLONIUM and RADON), and also in our Windows 2003 MSCS (nodes SENEGAL and TONGA).

Refer to “Tivoli Storage Manager client on Windows 2000” on page 248 and “Tivoli Storage Manager Client on Windows 2003” on page 289 for the configuration tasks on each of this environments.

6.4 Tivoli Storage Manager client on Windows 2000

In this section we describe how we configure the Tivoli Storage Manager client software to be capable of running in our MSCS Windows 2000, the same cluster we installed and configured in 4.3, “Windows 2000 MSCS installation and configuration” on page 29.

Archived

6.4.1 Windows 2000 lab setup

Our clustered lab environment consists of two Windows 2000 Advanced Servers, RADON and POLONIUM.

The Windows 2000 Tivoli Storage Manager backup/archive client configuration for this cluster is shown in Figure 6-10.

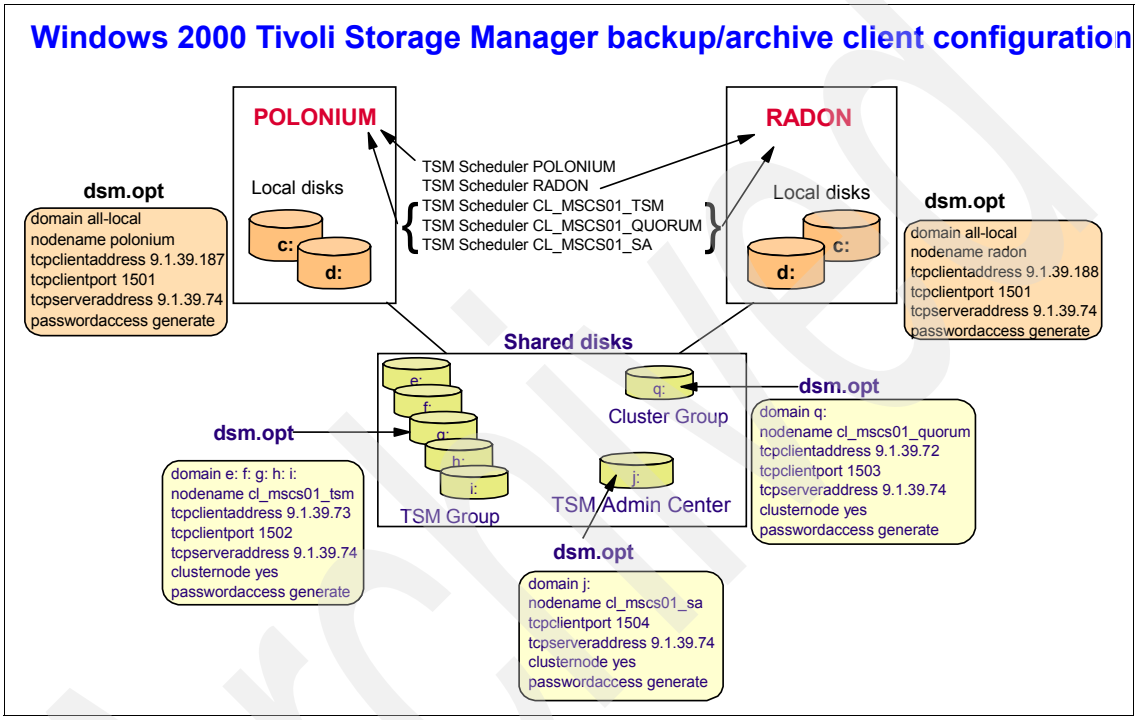

Figure 6-10 Tivoli Storage Manager backup/archive clustering client (Win.2000)

Refer to Table 4-1 on page 30, Table 4-2 on page 31, and Table 4-3 on page 31 for details of the MSCS cluster configuration used in our lab.

Archived

Table 6-1 and Table 6-2 show the specific Tivoli Storage Manager backup/archive client configuration we use for the purpose of this section.

Table 6-1 Tivoli Storage Manager backup/archive client for local nodes

Local node 1	
TSM nodename	POLONIUM
Backup domain	c: d: systemobject
Scheduler service name	TSM Scheduler POLONIUM
Client Acceptor service name	TSM Client Acceptor POLONIUM
Remote Client Agent service name	TSM Remote Client Agent POLONIUM
Local node 2	
TSM nodename	RADON
Backup domain	c: d: systemobject
Scheduler service name	TSM Scheduler RADON
Client Acceptor service name	TSM Client Acceptor RADON
Remote Client Agent service name	TSM Remote Client Agent RADON

Archived

Table 6-2 Tivoli Storage Manager backup/archive client for virtual nodes

Virtual node 1	
TSM nodename	CL_MSCS01_QUORUM
Backup domain	q:
Scheduler service name	TSM Scheduler CL_MSCS01_QUORUM
Client Acceptor service name	TSM Client Acceptor CL_MSCS01_QUORUM
Remote Client Agent service name	TSM Remote Client Agent CL_MSCS01_QUORUM
Cluster group name	Cluster Group
Virtual node 2	
TSM nodename	CL_MSCS01_SA
Backup domain	j:
Scheduler service name	TSM Scheduler CL_MSCS01_SA
Client Acceptor service name	TSM Client Acceptor CL_MSCS01_SA
Remote Client Agent service name	TSM Remote Client Agent CL_MSCS01_SA
Cluster group name	TSM Admin Center
Virtual node 3	
TSM nodename	CL_MSCS01_TSM
Backup domain	e: f: g: h: i:
Scheduler service name	TSM Scheduler CL_MSCS01_TSM
Client Acceptor service name	TSM Client Acceptor CL_MSCS01_TSM
Remote Client Agent service name	TSM Remote Client Agent CL_MSCS01_TSM
Cluster group name	TSM Group

Archived

6.4.2 Windows 2000 Tivoli Storage Manager Client configuration

We describe here how to configure the Tivoli Storage Manager backup/archive client in a Windows 2000 clustered environment. This is a two-step procedure:

1. Configuration to back up the local disk drives of each server.
2. Configuration to back up the shared disk drives of each group in the cluster.

Configuring the client to back up local disks

The configuration for the backup of the local disks is the same as for any standalone client:

1. We create a nodename for each server (POLONIUM and RADON) on the Tivoli Storage Manager server
2. We create the option file (dsm.opt) for each node on the local drive.

> **Important:** We should *only* use the `domain` option if not all local drives are going to be backed up. The default, if we do not specify anything, is backing up all local drives and system objects. We should not include any cluster drive in the domain parameter.

3. We generate the password locally by either opening the backup-archive GUI or issuing a query on the command prompt, such as **`dsmc q se`**.
4. We create the local Tivoli Storage Manager services as needed for each node, opening the backup-archive GUI client and selecting **Utilities** →**Setup Wizard**. The names we use for each service are:

- For RADON:
  - Tivoli Storage Manager Scheduler RADON
  - Tivoli Storage Manager Client Acceptor RADON
  - Tivoli Storage Manager Remote Client Agent RADON
- For POLONIUM:
  - Tivoli Storage Manager Scheduler POLONIUM
  - Tivoli Storage Manager Client Acceptor POLONIUM
  - Tivoli Storage Manager Remote Client Agent POLONIUM

5. After the configuration, the Windows services menu appears as shown in Figure 6-11. These are the Windows services for RADON. For POLONIUM we are presented with a very similar menu.

Archived

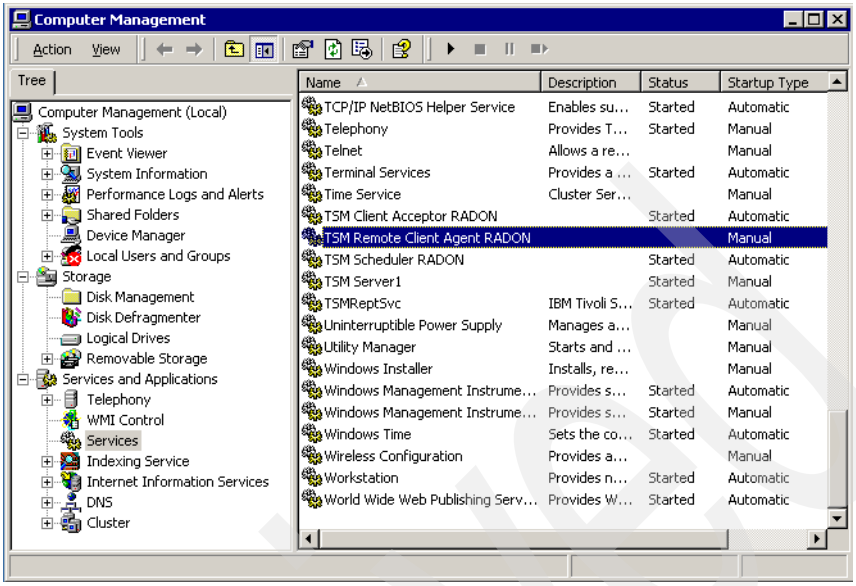

Figure 6-11 Tivoli Storage Manager client services

Configuring the client to back up shared disks

The configuration of Tivoli Storage Manager client to back up shared disks is slightly different for virtual nodes on MSCS.

For every resource group that has shared disks with backup requirements, we need to define an option file and an associated TSM scheduler service. If we want to use the Web client to access that virtual node from a browser, we also have to install the Web client services for that particular resource group.

For details of the nodenames, resources and services used for this part of the chapter, refer to Table 6-1 on page 250 and Table 6-2 on page 251.

Each resource group needs its own unique nodename. This ensures that Tivoli Storage Manager client correctly manages the disk resources in case of failure on any physical node, independently of the node who hosts the resources at that time.

As we can see in the tables mentioned above, we create three nodes in the Tivoli Storage Manager server database:

- CL_MSCS01_QUORUM: for the Cluster group
- CL_MSCS01_SA: for the TSM Admin Center group
- CL_MSCS01_TSM: for the TSM group

For each group, the configuration process consists of the following tasks:

1. Creation of the option files
2. Password generation
3. Installation (on each physical node on the MSCS) of the TSM scheduler service
4. Installation (on each physical node on the MSCS) of the TSM Web client services
5. Creation of a generic service resource for the TSM scheduler service using the Cluster Administrator application
6. Creation of a generic service resource for the TSM client acceptor service using the Cluster Administrator application

We describe each activity in the following sections.

Creation of the option files

For each group in the cluster we need to create an option file that will be used by the Tivoli Storage Manager nodename attached to that group.

The option file should be located on one of the shared disks hosted by this group. This ensures that both physical nodes have access to the file.

The dsm.opt file *must* contain at least the following options:

- nodename: Specifies the name that this group uses when it backs up data to the Tivoli Storage Manager server.
- domain: Specifies the disk drive letters managed by this group.
- passwordaccess *generate*: Specifies that the client generates a new password when the old one expires, and this new password is kept in the Windows registry.
- clusternode *yes*: To specify that it is a virtual node of a cluster. This is the main difference between the option file for a virtual node and the option file for a physical local node.

If we plan to use the schedmode prompted option to schedule backups, and we plan to use **the Web client interface for each virtual node**, we also should specify the following options:

- tcpclientaddress: Specifies **the unique IP address for this resource group**
- tcpclientport: Specifies a different TCP port for each node
- httpport: Specifies a different http port to contact with

There are other options we can specify, but the ones mentioned above are a requirement for a correct implementation of the client.

In our environment we create the dsm.opt files in the \tsm directory for the following drives:

- q: For the Cluster group
- j: For the Admin Center group
- g: For the TSM group

Option file for Cluster group

The dsm.opt file for this group contains the following options:

```
nodename cl_mscs01_quorum
passwordaccess generate
tcpserveraddress 9.1.39.73
errorlogretention 7
errorlogname q:\tsm\dsmerror.log
schedlogretention 7
schedlogname q:\tsm\dsmsched.log
domain q:
clusternode yes
schedmode prompted
tcpclientaddress 9.1.39.72
tcpclientport 1502
httpport 1582
```

Option file for TSM Admin Center group

The dsm.opt file for this group contains the following options:

```
nodename cl_mscs01_sa
passwordaccess generate
tcpserveraddress 9.1.39.73
errorlogretention 7
errorlogname j:\tsm\dsmerror.log
schedlogretention 7
schedlogname j:\tsm\dsmsched.log
domain j:
clusternode yes
tcpclientport 1503
httpport 1583
```

Archived

Option file for TSM Group

The dsm.opt file for this group contains the following options:

```
nodename cl_mscs_tsm
passwordaccess generate
tcpserveraddress 9.1.39.73
errorlogretention 7
errorlogname g:\tsm\dsmerror.log
schedlogretention 7
schedlogname g:\tsm\dsmsched.log
domain e: f: g: h: i:
clusternode yes
schedmode prompted
tcpclientaddress 9.1.39.73
tcpclientport 1504
httpport 1584
```

Password generation

The Windows registry of each server needs to be updated with the password used to register the nodenames for each resource group in the Tivoli Storage Manager server.

> **Important:** The steps below require that we run the following commands on both nodes while they own the resources. We recommend to move all resources to one of the nodes, complete the tasks for this node, and then move all resources to the other node and repeat the tasks.

Since the dsm.opt is located for each node in a different location, we need to specify the path for each, using the *-optfile* option of the **dsmc** command:

1. We run the following commands from a MS-DOS prompt in the Tivoli Storage Manager client directory (c:\program files\tivoli\tsm\baclient):

   ```
   dsmc q se -optfile=q:\tsm\dsm.opt
   ```

2. Tivoli Storage Manager prompts the nodename for the client (the specified in dsm.opt). If it is correct, press Enter.

Archived

3. Tivoli Storage Manager next asks for a password. We type the password and press Enter. Figure 6-12 shows the output of the command.

```
C:\WINNT\system32\cmd.exe

C:\Program Files\Tivoli\tsm\baclient>dsmc q se -optfile=g:\tsm\dsm.opt
IBM Tivoli Storage Manager
Command Line Backup/Archive Client Interface
  Client Version 5, Release 3, Level 0.0
  Client date/time: 01/27/2005 14:37:49
(c) Copyright by IBM Corporation and other(s) 1990, 2004. All Rights Reserved.

Node Name: CL_MSCS01_TSM
Please enter your user id <CL_MSCS01_TSM>:

Please enter password for user id "CL_MSCS01_TSM": ******

Session established with server TSMSRV01: Windows
  Server Version 5, Release 3, Level 0.0
  Server date/time: 01/27/2005 14:37:49  Last access: 01/27/2005 14:37:49

TSM Server Connection Information

Server Name.............: TSMSRV01
Server Type.............: Windows
Server Version..........: Ver. 5, Rel. 3, Lev. 0.0
Last Access Date........: 01/27/2005 14:37:49
Delete Backup Files.....: "No"
Delete Archive Files....: "Yes"

Node Name...............: CL_MSCS01_TSM
User Name...............:


C:\Program Files\Tivoli\tsm\baclient>_
```

Figure 6-12 Generating the password in the registry

Note: The password is kept in the Windows registry of this node and we do not need to type it any more. The client reads the password from the registry every time it opens a session with the Tivoli Storage Manager server.

4. We repeat the command for the other nodes

```
dsmc q se -optfile=j:\tsm\dsm.opt
dsmc q se -optfile=g:\tsm\dsm.opt
```

5. We move the resources to the other node and repeat steps 1 to 4.

Installing the TSM Scheduler service

For backup automation, using the Tivoli Storage Manager scheduler, we need to install and configure one scheduler service for each resource group.

Important: We must install the scheduler service for each cluster group exactly with the same name, which is case sensitive, on each of the physical nodes and on the MSCS Cluster Administrator, otherwise failover will not work.

1. We need to be sure we are located on the node that hosts all resources, in order to start with the Tivoli Storage Manager scheduler service installation.

2. We begin the installation of the scheduler service for each group on POLONIUM. This is the node that hosts the resources. We use the **dsmcutil program**. This utility is located on the Tivoli Storage Manager client installation path (c:\program files\tivoli\tsm\baclient).

   In our lab we installed three scheduler services, one for each cluster group.

3. We open an MS-DOS command line and, in the Tivoli Storage Manager client installation path, we issue the following command:

```
dsmcutil inst sched /name:"TSM Scheduler CL_MSCS01_QUORUM"
/clientdir:"c:\program files\tivoli\tsm\baclient" /optfile:q:\tsm\dsm.opt
/node:CL_MSCS01_QUORUM /password:itsosj /clustername:CL_MSCS01
/clusternode:yes /autostart:no
```

4. We show the result of executing the command in Figure 6-13.

```
C:\WINNT\system32\cmd.exe

Command: Install TSM Client Service
Machine: POLONIUM(Local Machine)


Locating the Cluster Services ...
MS cluster ...running


Installing TSM Client Service:

        Machine          : POLONIUM
        Service Name     : TSM Scheduler CL_MSCS01_QUORUM
        Client Directory : c:\program files\tivoli\tsm\baclient
        Automatic Start  : no
        Logon Account    : LocalSystem

The service was successfully installed.


Creating Registry Keys ...

Updated registry value 'ImagePath' .
Updated registry value 'EventMessageFile' .
Updated registry value 'TypesSupported' .
Updated registry value 'TSM Scheduler CL_MSCS01_QUORUM' .
Updated registry value 'ADSMClientKey' .
Updated registry value 'OptionsFile' .
Updated registry value 'EventLogging' .
Updated registry value 'ClientNodeName' .
Updated registry value 'ClusterNode' .
Updated registry value 'ClusterGroupName' .

Generating registry password ...
Authenticating TSM password for node CL_MSCS01_QUORUM ...

Connecting to TSM Server via client options file 'q:\tsm\dsm.opt' ...



Error 2 establishing TSM API session: The password couldn't be authenticated wit
h the TSM server.



Error 2 authenticating password: The password couldn't be authenticated with the
 TSM server.


Starting the 'TSM Scheduler CL_MSCS01_QUORUM' service ...

The service was successfully started.
```

Figure 6-13 Result of Tivoli Storage Manager scheduler service installation

5. We repeat this command to install the scheduler service for TSM Admin Center group, changing the information as needed. The command is:

```
dsmcutil inst sched /name:"TSM Scheduler CL_MSCS01_SA"
/clientdir:"c:\Program Files\Tivoli\tsm\baclient" /optfile:j:\tsm\dsm.opt
/node:CL_MSCS01_SA /password:itsosj /clusternode:yes /clustername:CL_MSCS01
/autostart:no
```

6. And again to install the scheduler service for TSM Group we use:

```
dsmcutil inst sched /name:"TSM Scheduler CL_MSCS01_TSM"
/clientdir:"c:\Program Files\Tivoli\tsm\baclient" /optfile:g:\tsm\dsm.opt
/node:CL_MSCS01_TSM /password:itsosj /clusternode:yes
/clustername:CL_MSCS01 /autostart:no
```

7. Be sure to stop all services using the Windows service menu before going on.
8. We move the resources to the second node, and run exactly the same commands as before (steps 1 to 7).

Attention: the Tivoli Storage Manager scheduler service names used on both nodes *must match*. Also remember to use the same parameters for the dsmcutil tool. *Do not forget* the clusternode yes and clustername options.

So far the Tivoli Storage Manager scheduler services are installed on both nodes of the cluster with exactly the same names for each resource group. The last task consists of the definition for a new resource on each cluster group.

Creating a generic service resource for TSM scheduler service

For a correct configuration of the Tivoli Storage Manager client we define, for each cluster group, a new generic service resource. This resource relates to the scheduler service name created for this group.

Important: Before continuing, we make sure to stop all services created in "Installing the TSM Scheduler service" on page 257 on all nodes. We also make sure all the resources are on one of the nodes.

1. We open the Cluster Administrator panel on the node that hosts all the resources and we select the first group (**Cluster Group**). We right-click the name and select **New** →**Resource** as shown in Figure 6-14.

Archived

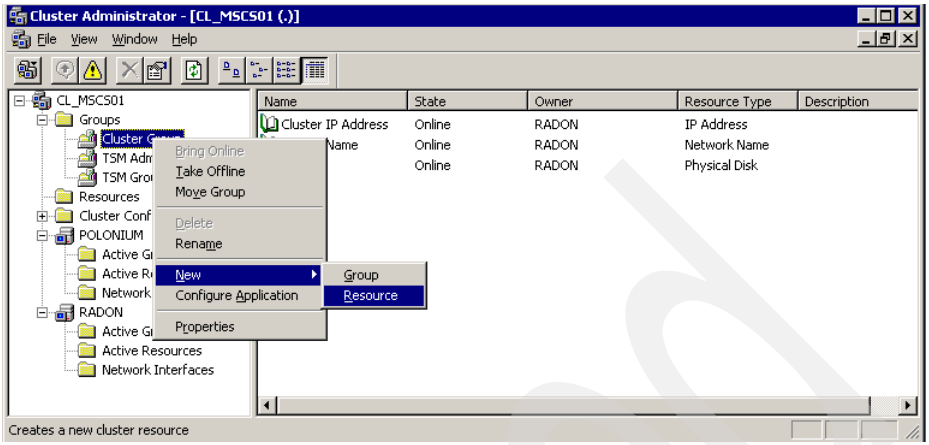

Figure 6-14 Creating new resource for Tivoli Storage Manager scheduler service

2. We type a Name for the resource (we recommend to use the same name as the scheduler service) and select Generic Service as resource type. We click **Next** as shown in Figure 6-15.

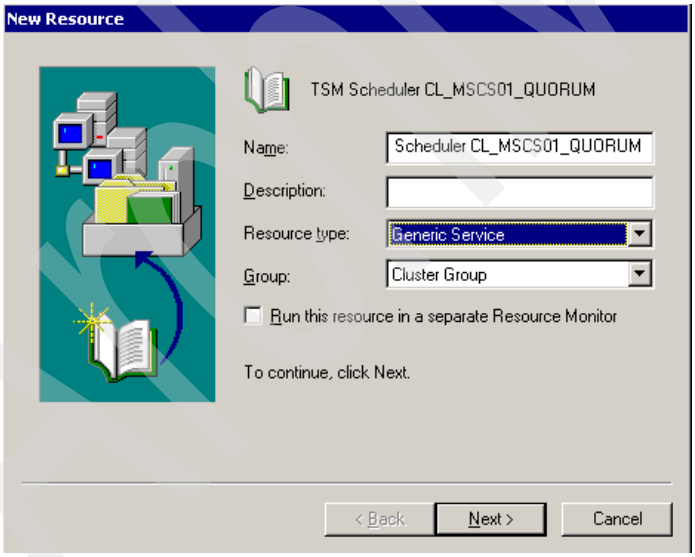

Figure 6-15 Definition of TSM Scheduler generic service resource

3. We leave both nodes as possible owners for the resource and click **Next** (Figure 6-16).

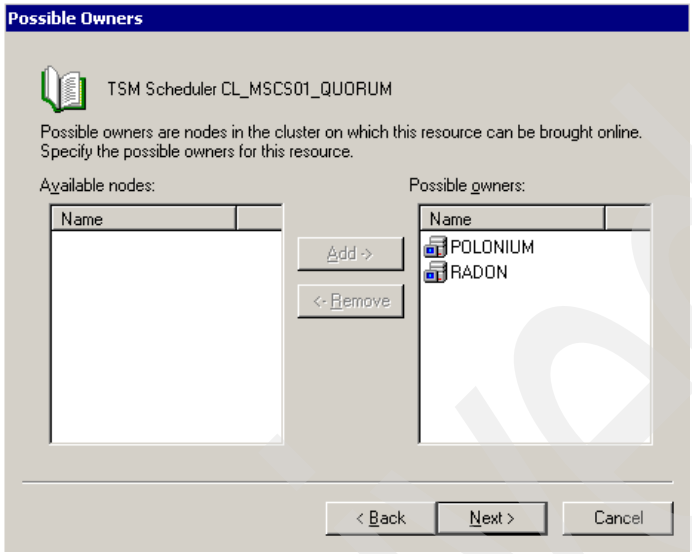

Figure 6-16 Possible owners of the resource

4. We **Add** the disk resource (q:) on Dependencies as shown in Figure 6-17. Then we click **Next** (Figure 6-17).

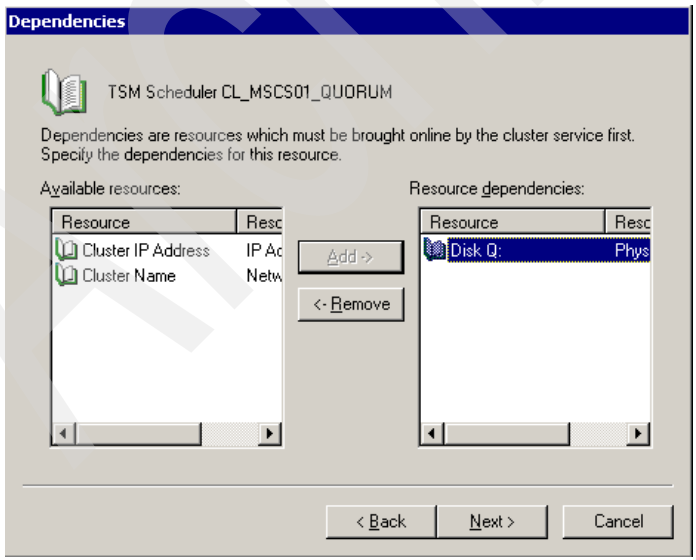

Figure 6-17 Dependencies

Archived

5. On the next menu we type a Service name. This must match the name used while installing the scheduler service on both nodes. Then we click **Next** (Figure 6-18).

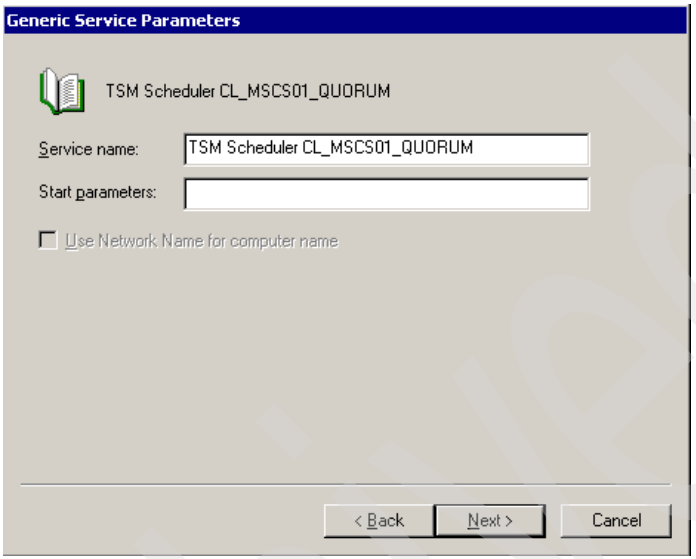

Figure 6-18 Generic service parameters

6. We click **Add** to type the Registry Key where Windows 2000 will save the generated password for the client. The registry key is:

```
SOFTWARE\IBM\ADSM\CurrentVersion\BackupClient\Nodes\<nodename>\<tsmservername>
```

Archived

We click **OK** (Figure 6-19).

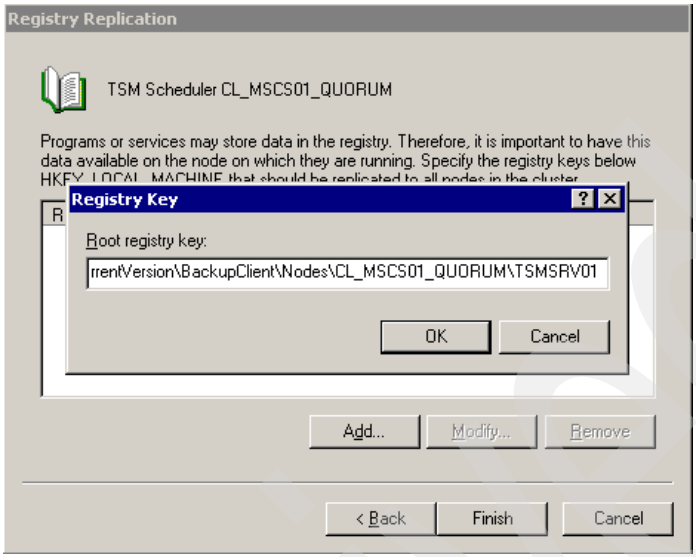

Figure 6-19 Registry key replication

7. If the resource creation is successful an information menu appears as shown in Figure 6-20. We click **OK**.

Figure 6-20 Successful cluster resource installation

8. As seen in Figure 6-21, the Cluster group is offline because the new resource is also offline. We bring it online.

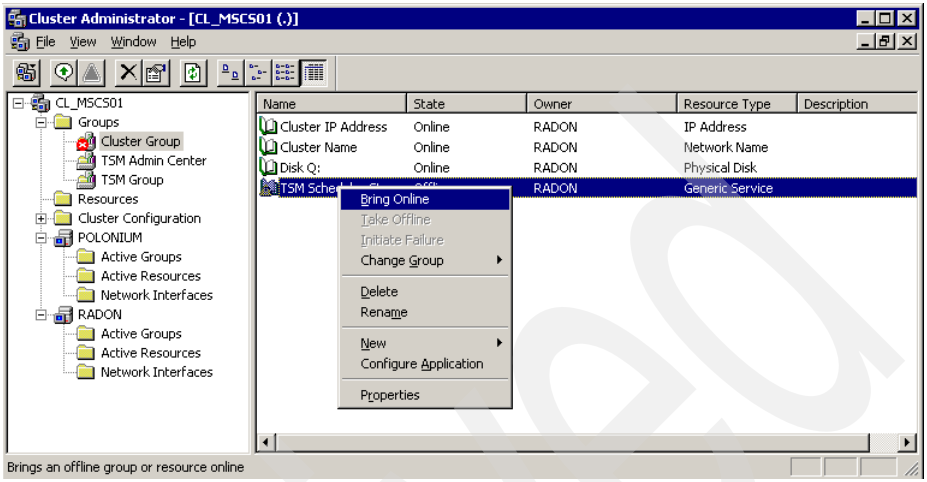

Figure 6-21 Bringing online the Tivoli Storage Manager scheduler service

9. The Cluster Administrator menu, after all resources are online, is shown in Figure 6-22.

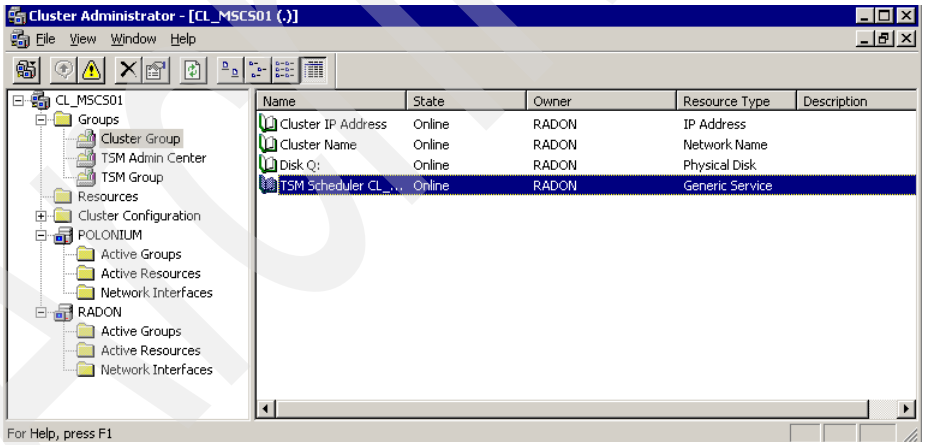

Figure 6-22 Cluster group resources online

Archived

10. If we go to the Windows service menu, Tivoli Storage Manager scheduler service is started on RADON, the node which now hosts this resource group (Figure 6-23).

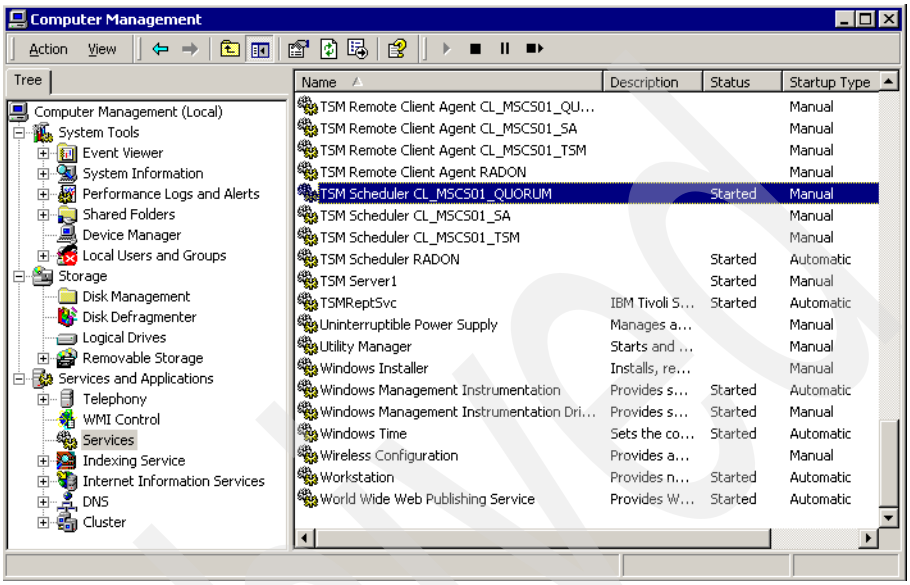

Figure 6-23 Windows service menu

11. We repeat steps 1-10 to create the Tivoli Storage Manager scheduler generic service resource for TSM Admin Center and TSM Group cluster groups. The resource names are:
   - TSM Scheduler CL_MSCS01_SA: for TSM Admin Center resource group
   - TSM Scheduler CL_MSCS01_TSM: for TSM Group resource group.

Important: To back up, archive, or retrieve data residing on MSCS, the Windows account used to start the Tivoli Storage Manager scheduler service on each local node must belong to the Administrators or Domain Administrators group or Backup Operators group.

12. We move the resources to check that Tivoli Storage Manager scheduler services successfully start on the second node while they are stopped on the first node.

Note: Use only the Cluster Administration menu to bring online/offline the Tivoli Storage Manager scheduler service for virtual nodes.

Archived

Installing the TSM Web client services

This task is not necessary if we do not want to use the Web client. However, if we want to be able to access virtual clients from a Web browser, we must follow the tasks explained in this section.

We install Tivoli Storage Manager Client Acceptor and Tivoli Storage Manager Remote Client Agent services on both physical nodes with the same service names and the same options.

1. We make sure we are in the cluster that hosts all resources in order to install the scheduler service.
2. We install the scheduler service for each group using the **`dsmcutil program`**. This utility is located on the Tivoli Storage Manager client installation path (c:\program files\tivoli\tsm\baclient).
3. In our lab we install three Client Acceptor services, one for each cluster group, and three Remote Client Agent services (one for each cluster group). When we start the installation the node that hosts the resources is POLONIUM.
4. We open a MS-DOS Windows command line and change to the Tivoli Storage Manager client installation path. We run the **`dsmcutil tool`** with the appropriate parameters to create the Tivoli Storage Manager client acceptor service for the Cluster group, as shown in Figure 6-24.

Archived

```
C:\WINNT\system32\cmd.exe
The service was successfully started.


C:\Program Files\Tivoli\tsm\baclient>dsmcutil inst cad /name:"TSM Client Accepto
r CL_MSCS01_QUORUM" /clientdir:"c:\program files\tivoli\tsm\baclient" /optfile:q
:\tsm\dsm.opt /node:CL_MSCS01_QUORUM /password:itsosj /clustername:CL_MSCS01 /cl
usternode:yes /autostart:no /httpport:1582

TSM Windows NT Client Service Configuration Utility
Command Line Interface - Version 5, Release 3, Level 0.0
(C) Copyright IBM Corporation, 1990, 2004, All Rights Reserved.
Last Updated Dec  8 2004
TSM Api Verison 5.3.0

Command: Install TSM Client Service
Machine: POLONIUM(Local Machine)



Installing TSM Client Service:

       Machine           : POLONIUM
       Service Name      : TSM Client Acceptor CL_MSCS01_QUORUM
       Client Directory  : c:\program files\tivoli\tsm\baclient
       Automatic Start   : no
       Logon Account     : LocalSystem

The service was successfully installed.


Creating Registry Keys ...

Updated registry value 'ImagePath' .
Updated registry value 'EventMessageFile' .
Updated registry value 'TypesSupported' .
Updated registry value 'TSM Client Acceptor CL_MSCS01_QUORUM' .
Updated registry value 'ADSMClientKey' .
Updated registry value 'OptionsFile' .
Updated registry value 'EventLogging' .
Updated registry value 'ClientNodeName' .
Updated registry value 'HttpPort' .

Generating registry password ...
Authenticating TSM password for node CL_MSCS01_QUORUM ...

Connecting to TSM Server via client options file 'q:\tsm\dsm.opt' ...

Password authentication successful.

The registry password for TSM node CL_MSCS01_QUORUM has been updated.
```

Figure 6-24 Installing the Client Acceptor service in the Cluster Group

5. After a successful installation of the client acceptor for this resource group, we run the **dsmcutil tool** again to create its remote client agent partner service typing the command:

   ```
   dsmcutil inst remoteagent /name:"TSM Remote Client Agent CL_MSCS01_QUORUM"
   /clientdir:"c:\Program Files\Tivoli\tsm\baclient" /optfile:q:\tsm\dsm.opt
   /node:CL_MSCS01_QUORUM /password:itsosj /clusternode:yes
   /clustername:CL_MSCS01 /startnow:no /partnername:"TSM Client Acceptor
   CL_MSCS01_QUORUM"
   ```

6. If the installation is successful, we receive the following sequence of messages as shown in Figure 6-25.

Archived

```
C:\WINNT\system32\cmd.exe

TSM Windows NT Client Service Configuration Utility
Command Line Interface - Version 5, Release 3, Level 0.0
(C) Copyright IBM Corporation, 1990, 2004, All Rights Reserved.
Last Updated Dec  8 2004
TSM Api Verison 5.3.0

Command: Install TSM Client Service
Machine: POLONIUM(Local Machine)


Locating partner client acceptor service 'TSM Client Acceptor CL_MSCS01_QUORUM'
...



Installing TSM Client Service:

        Machine          : POLONIUM
        Service Name     : TSM Remote Client Agent CL_MSCS01_QUORUM
        Client Directory : c:\program files\tivoli\tsm\baclient
        Automatic Start  : no
        Logon Account    : LocalSystem

The service was successfully installed.


Creating Registry Keys ...

Updated registry value 'ImagePath' .
Updated registry value 'EventMessageFile' .
Updated registry value 'TypesSupported' .
Updated registry value 'TSM Remote Client Agent CL_MSCS01_QUORUM' .
Updated registry value 'ADSMClientKey' .
Updated registry value 'OptionsFile' .
Updated registry value 'EventLogging' .
Updated registry value 'ClientNodeName' .
Updated registry value 'PartnerName' .
Updated Partner Client Acceptor registry value 'PartnerName' .

Generating registry password ...
Authenticating TSM password for node CL_MSCS01_QUORUM ...

Connecting to TSM Server via client options file 'q:\tsm\dsm.opt' ...

Password authentication successful.

The registry password for TSM node CL_MSCS01_QUORUM has been updated.
```

Figure 6-25 Successful installation, Tivoli Storage Manager Remote Client Agent

7. We follow the same process to install the services for the TSM Admin Center cluster group. We use the following commands:

```
dsmcutil inst cad /name:"TSM Client Acceptor CL_MSCS01_SA"
/clientdir:"c:\Program Files\Tivoli\tsm\baclient" /optfile:j:\tsm\dsm.opt
/node:CL_MSCS01_SA /password:itsosj /clusternode:yes /clustername:CL_MSCS01
/autostart:no /httpport:1583

dsmcutil inst remoteagent /name:"TSM Remote Client Agent CL_MSCS01_SA"
/clientdir:"c:\Program Files\Tivoli\tsm\baclient" /optfile:j:\tsm\dsm.opt
/node:CL_MSCS01_SA /password:itsosj /clusternode:yes /clustername:CL_MSCS01
/startnow:no /partnername:"TSM Client Acceptor CL_MSCS01_SA"
```

8. And finally we use the same process to install the services for the TSM Group, with the following commands:

```
dsmcutil inst cad /name:"TSM Client Acceptor CL_MSCS01_TSM"
/clientdir:"c:\Program Files\Tivoli\tsm\baclient" /optfile:g:\tsm\dsm.opt
/node:CL_MSCS01_TSM /password:itsosj /clusternode:yes
/clustername:CL_MSCS01 /autostart:no /httpport:1584
```

```
dsmcutil inst remoteagent /name:"TSM Remote Client Agent CL_MSCS01_TSM"
/clientdir:"c:\Program Files\Tivoli\tsm\baclient" /optfile:g:\tsm\dsm.opt
/node:CL_MSCS01_TSM /password:itsosj /clusternode:yes
/clustername:CL_MSCS01 /startnow:no /partnername:"TSM Client Acceptor
CL_MSCS01_TSM"
```

Important: The client acceptor and remote client agent services *must* be installed with the same name on each physical node on the MSCS, otherwise failover will not work. Also, do not forget the options clusternode yes and clustername as well as to specify the correct dsm.opt path file name in the optfile parameter of the **dsmcutil** command.

9. We move the resources to the second node (RADON) and repeat steps 1-8 with the same options for each resource group.

So far the Tivoli Storage Manager Web client services are installed on both nodes of the cluster with exactly the same names for each resource group. The last task consists of the definition for new resource on each cluster group. But first we go to the Windows Service menu and stop all the Web client services on RADON.

Creating a generic resource for TSM Client Acceptor service

For a correct configuration of the Tivoli Storage Manager Web client we define, for each cluster group, a new generic service resource. This resource will be related to the Client Acceptor service name created for this group.

Important: Before continuing, we make sure to stop all services created in "Installing the TSM Web client services" on page 266 on all nodes. We also make sure all resources are on one of the nodes.

Here are the steps we follow:

1. We open the Cluster Administrator menu on the node that hosts all resources and we select the first group (**Cluster Group**). We right-click the name and select **New** →**Resource** as shown in Figure 6-26.

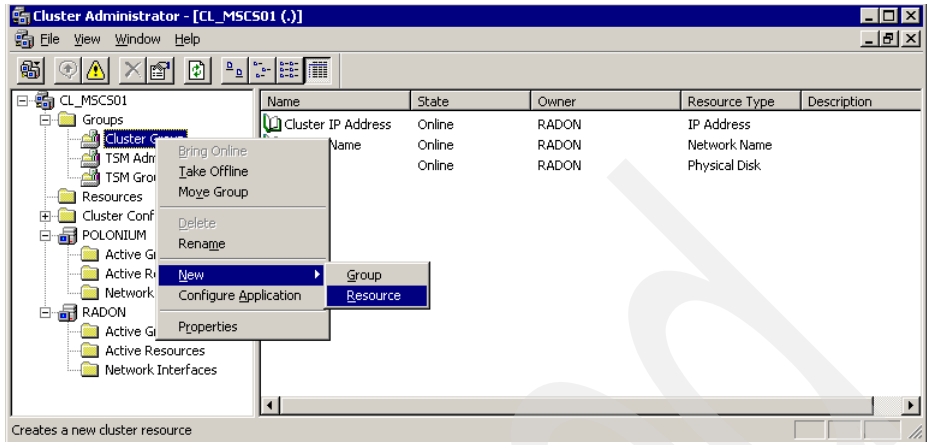

Figure 6-26 New resource for Tivoli Storage Manager Client Acceptor service

2. We type a Name for the resource (we recommend to use the same name as the scheduler service) and select Generic Service as resource type. We click Next as shown in Figure 6-27.

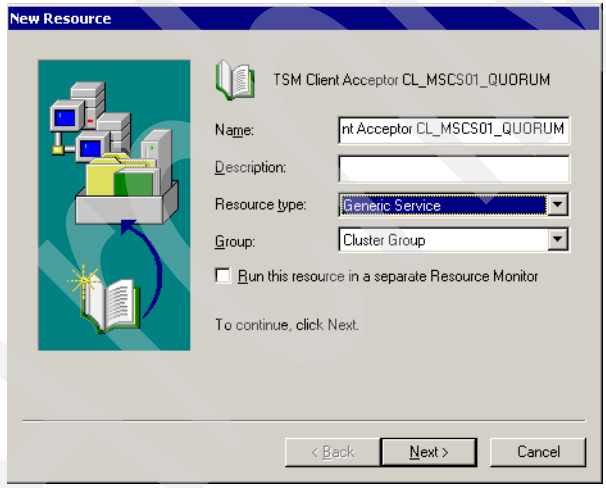

Figure 6-27 Definition of TSM Client Acceptor generic service resource

3. We leave both nodes as possible owners for the resource and we click **Next** (Figure 6-28).

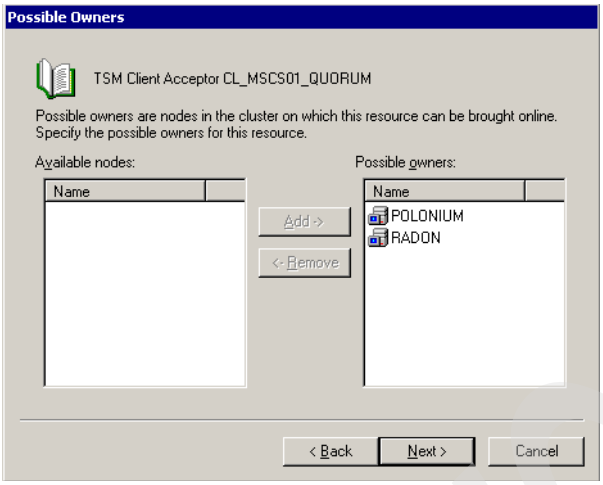

Figure 6-28 Possible owners of the TSM Client Acceptor generic service

4. We **Add** the disk resources (in this case q:) on Dependencies in Figure 6-29. We click **Next**.

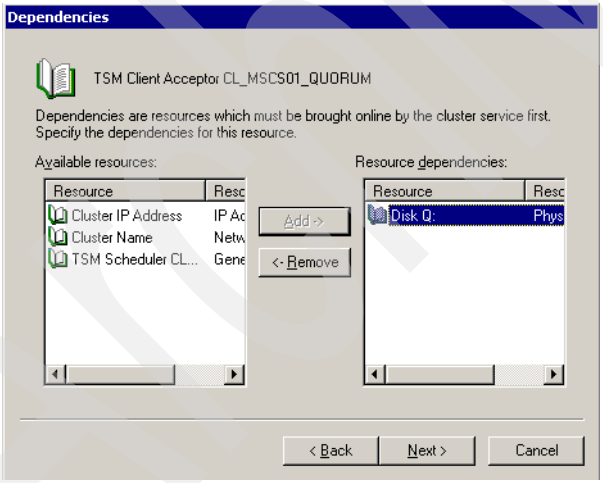

Figure 6-29 Dependencies for TSM Client Acceptor generic service

5. On the next menu (Figure 6-30), we type a Service name. This must match the name used while installing the client acceptor service on both nodes. We click **Next**.

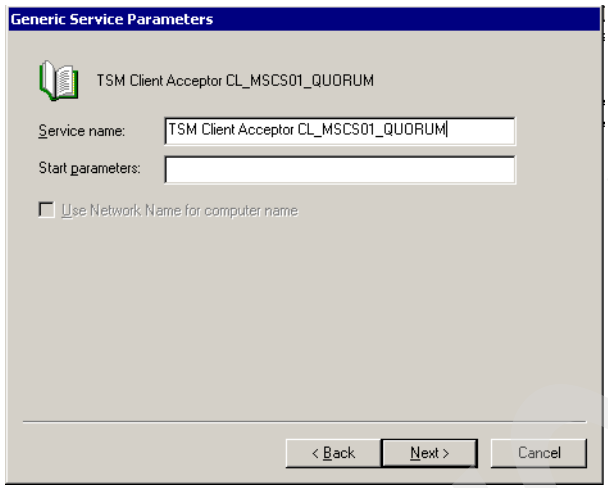

Figure 6-30 TSM Client Acceptor generic service parameters

6. Next we type the Registry Key where Windows 2000 will save the generated password for the client. It is the same path we typed in Figure 6-19 on page 263. We click **OK**.
7. If the resource creation is successful, we receive an information menu as shown in Figure 6-20 on page 263. We click **OK**.
8. As shown in the next figure, the Cluster Group is offline because the new resource is also offline. We bring it online (Figure 6-31).

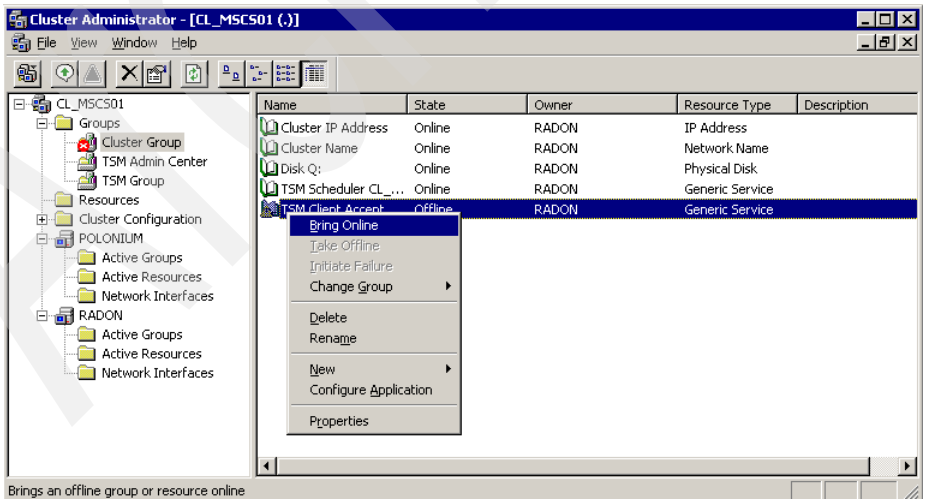

Figure 6-31 Bringing online the TSM Client Acceptor generic service

9. The Cluster Administrator menu displays next as shown in Figure 6-32.

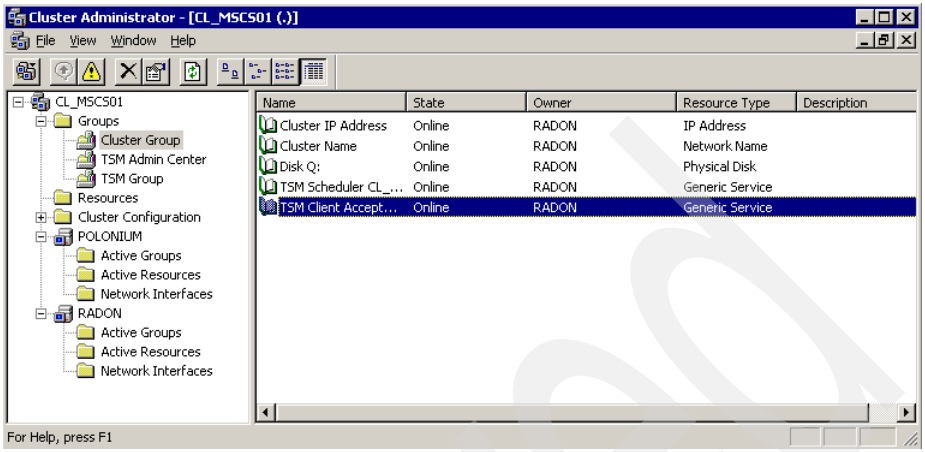

Figure 6-32 TSM Client Acceptor generic service online

10. If we go to the Windows service menu, Tivoli Storage Manager Client Acceptor service is started on RADON, the node which now hosts this resource group (Figure 6-33).

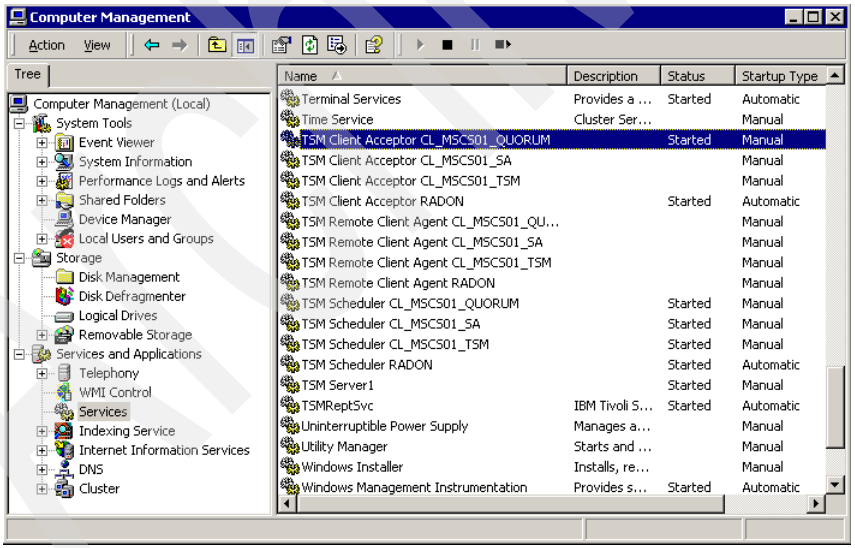

Figure 6-33 Windows service menu

Important: All Tivoli Storage Manager client services used by virtual nodes of the cluster must figure as *Manual* on the Startup Type column in Figure 6-33. They may only be started on the node that hosts the resource at that time.

11. We follow the same tasks to create the Tivoli Storage Manager client acceptor service resource for TSM Admin Center and TSM Group cluster groups. The resource names are:
    - TSM Client Acceptor CL_MSCS01_SA: for TSM Admin Center resource group
    - TSM Client Acceptor CL_MSCS01_TSM: for TSM Group resource group.
12. We move the resources to check that Tivoli Storage Manager client acceptor services successfully start on the second node, POLONIUM, while they are stopped on the first node.

Note: Use only the Cluster Administration menu to bring online/offline the Tivoli Storage Manager Client Acceptor service for virtual nodes.

Filespace names for local and virtual nodes

If the configuration of Tivoli Storage Manager client in our MSCS is correct, when the client backs up files against our Tivoli Storage Manager server, the filespace names for local (physical) nodes and virtual (shared) nodes are different. We show this in Figure 6-34.

Archived

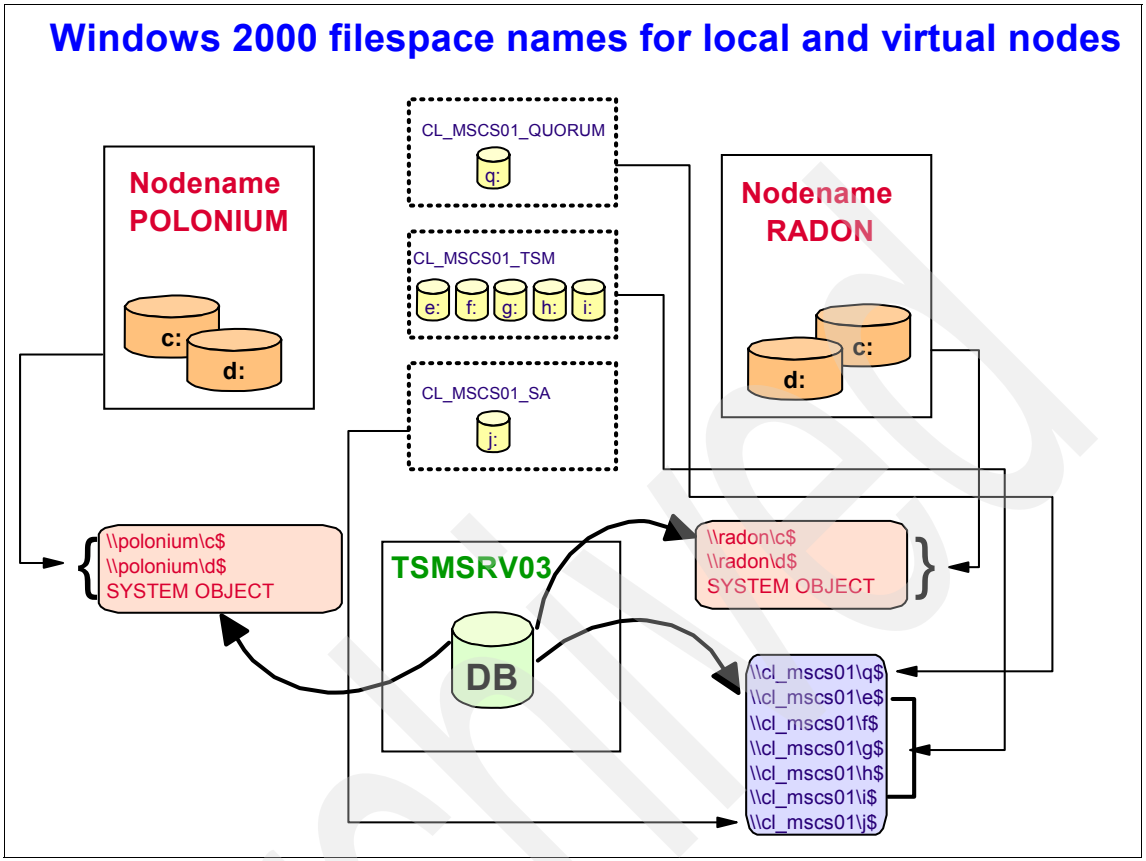

Figure 6-34 Windows 2000 filespace names for local and virtual nodes

When the local nodes back up files, their filespace names start with the physical nodename. However, when the virtual nodes back up files, their filespace names start with the cluster name, in our case, CL_MSCS01.

6.4.3 Testing Tivoli Storage Manager client on Windows 2000 MSCS

In order to check the high availability of Tivoli Storage Manager client on our lab environment, we must do some testing.

Our objective with these tests is to know how Tivoli Storage Manager client can respond, on a clustered environment, after certain kinds of failures that affect the shared resources.

For the purpose of this section, we use a Tivoli Storage Manager server installed on an AIX machine: TSMSRV03. For details of this server, refer to the AIX chapters in this book. Remember, our Tivoli Storage Manager virtual clients are:

- CL_MSCS01_QUORUM
- CL_MSCS01_TSM
- CL_MSCS01_SA

Testing client incremental backup

Our first test consists of an incremental backup started from the client.

Objective

The objective of this test is to show what happens when a client incremental backup is started for a virtual client in the cluster, and the node that hosts the resources at that moment suddenly fails.

Activities

To do this test, we perform these tasks:

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager client resource as shown in Figure 6-35.

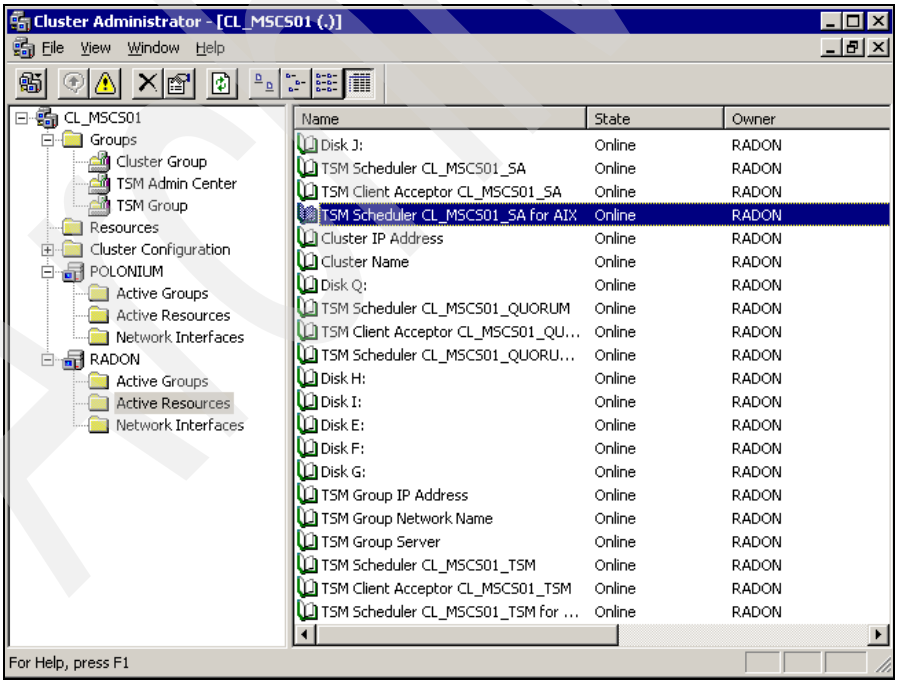

Figure 6-35 Resources hosted by RADON in the Cluster Administrator

As we can see in the figure, RADON hosts all the resources at this moment.

Note: TSM Scheduler CL_MSCS01_SA for AIX means the Tivoli Storage Manager scheduler service used by CL_MSCS01_SA when logs into the AIX server. We had to create this service on each node and then use the Cluster Administrator to define the generic service resource. To achieve this goal we followed the same tasks already explained for the rest of scheduler services.

2. We schedule a client incremental backup operation using the Tivoli Storage Manager server scheduler and we associate the schedule to CL_MSCS01_SA nodename.
3. A client session for CL_MSCS01_SA nodename starts on the server as shown in Example 6-1.

Example 6-1 Session started for CL_MSCS01_SA

```
02/01/2005 16:29:04      ANR0406I Session 70 started for node CL_MSCS01_SA (WinNT) (Tcp/Ip
9.1.39.188(2718)). (SESSION: 70)

02/01/2005 16:29:05      ANR0406I Session 71 started for node
CL_MSCS01_SA (WinNT) (Tcp/Ip 9.1.39.188(2719)). (SESSION: 71)
```

4. The client starts sending files to the server as we can see on the schedule log file in Example 6-2.

Example 6-2 Schedule log file shows the client sending files to the server

```
02/01/2005 16:36:17 --- SCHEDULEREC QUERY BEGIN
02/01/2005 16:36:17 --- SCHEDULEREC QUERY END
02/01/2005 16:36:17 Next operation scheduled:
02/01/2005 16:36:17
------------------------------------------------------------
02/01/2005 16:36:17 Schedule Name:         INCR_BACKUP
02/01/2005 16:36:17 Action:                Incremental
02/01/2005 16:36:17 Objects:
02/01/2005 16:36:17 Options:
02/01/2005 16:36:17 Server Window Start:   16:27:57 on 02/01/2005
02/01/2005 16:36:17
------------------------------------------------------------
02/01/2005 16:36:17
Executing scheduled command now.
02/01/2005 16:36:17 --- SCHEDULEREC OBJECT BEGIN INCR_BACKUP 02/01/2005
16:27:57
02/01/2005 16:36:17 Incremental backup of volume '\\cl_mscs01\j$'
02/01/2005 16:36:27 Directory-->                 0 \\cl_mscs01\j$\ [Sent]
```

```
02/01/2005 16:36:27 Directory-->                    0 \\cl_mscs01\j$\Program
Files [Sent]
02/01/2005 16:36:27 Directory-->                    0 \\cl_mscs01\j$\RECYCLER
[Sent]
02/01/2005 16:36:27 Directory-->                    0 \\cl_mscs01\j$\System
Volume Information [Sent]
02/01/2005 16:36:27 Directory-->                0 \\cl_mscs01\j$\TSM [Sent]
02/01/2005 16:36:27 Directory-->                    0 \\cl_mscs01\j$\TSM_Images
[Sent]
```

Note: Observe in Example 6-2 that the filespace name used by Tivoli Storage Manager to store the files in the server (`\\cl_mscs01\j$`). If the client is correctly configured to work on MSCS, the filespace name always starts with the cluster name. It does not use the local name of the physical node which hosts the resource at the time of backup.

5. While the client continues sending files to the server, we force RADON to fail. The following sequence takes place:
   a. The client loses its connection with the server temporarily, and the session terminates as we can see on the Tivoli Storage Manager server activity log shown in Example 6-3.

Example 6-3 The client loses its connection with the server

```
02/01/2005 16:29:54      ANR0480W Session 71 for node CL_MSCS01_SA (WinNT) terminated - connection with
client severed. (SESSION: 71)

02/01/2005 16:29:54      ANR0480W Session 70 for node CL_MSCS01_SA
(WinNT) terminated - connection with client severed. (SESSION: 70)
```

   b. In the Cluster Administrator menu, RADON is not in the cluster and POLONIUM begins to bring the resources online.
   c. After a while the resources are online on POLONIUM.
   d. When the TSM Scheduler CL_MSCS01_SA for AIX resource is online (hosted by POLONIUM), the client restarts the backup as we show on the schedule log file in Example 6-4.

Example 6-4 Schedule log file shows backup is restarted on the client

```
02/01/2005 16:37:07 Normal File-->             4,742 \\cl_mscs01\j$\Program
Files\IBM\ISC\AppServer\java\jre\lib\font.properties.te [Sent]
02/01/2005 16:37:07 Normal File-->             6,535 \\cl_mscs01\j$\Program
Files\IBM\ISC\AppServer\java\jre\lib\font.properties.th [Sent]
02/01/2005 16:38:39 Querying server for next scheduled event.
02/01/2005 16:38:39 Node Name: CL_MSCS01_SA
```

```
02/01/2005 16:38:39 Session established with server TSMSRV03: AIX-RS/6000
02/01/2005 16:38:39   Server Version 5, Release 3, Level 0.0
02/01/2005 16:38:39   Server date/time: 02/01/2005 16:31:26  Last access:
02/01/2005 16:29:57
02/01/2005 16:38:39 --- SCHEDULEREC QUERY BEGIN
02/01/2005 16:38:39 --- SCHEDULEREC QUERY END
02/01/2005 16:38:39 Next operation scheduled:
02/01/2005 16:38:39
------------------------------------------------------------
02/01/2005 16:38:39 Schedule Name:         INCR_BACKUP
02/01/2005 16:38:39 Action:                Incremental
02/01/2005 16:38:39 Objects:
02/01/2005 16:38:39 Options:
02/01/2005 16:38:39 Server Window Start:   16:27:57 on 02/01/2005
02/01/2005 16:38:39
------------------------------------------------------------
02/01/2005 16:38:39
Executing scheduled command now.
02/01/2005 16:38:39 --- SCHEDULEREC OBJECT BEGIN INCR_BACKUP 02/01/2005
16:27:57
02/01/2005 16:38:39 Incremental backup of volume '\\cl_mscs01\j$'
02/01/2005 16:38:50 ANS1898I ***** Processed       500 files *****
02/01/2005 16:38:52 ANS1898I ***** Processed     1,000 files *****
02/01/2005 16:38:54 ANS1898I ***** Processed     1,500 files *****
02/01/2005 16:38:56 ANS1898I ***** Processed     2,000 files *****
02/01/2005 16:38:57 ANS1898I ***** Processed     2,500 files *****
02/01/2005 16:38:59 ANS1898I ***** Processed     3,000 files *****
02/01/2005 16:38:59 Directory-->                 0 \\cl_mscs01\j$\ [Sent]
02/01/2005 16:38:59 Normal File-->       6,713,114 \\cl_mscs01\j$\Program
Files\IBM\ISC\AppServer\java\jre\lib\graphics.jar [Sent]
02/01/2005 16:38:59 Normal File-->         125,336 \\cl_mscs01\j$\Program
Files\IBM\ISC\AppServer\java\jre\lib\ibmcertpathprovider.jar [Sent]
02/01/2005 16:38:59 Normal File-->           9,210 \\cl_mscs01\j$\Program
Files\IBM\ISC\AppServer\java\jre\lib\ibmjaasactivelm.jar [Sent]
```

Here, the last file reported as sent to the server before the failure is:

\\cl_mscs01\j$\Program Files \IBM\ISC\AppServer\java\jre\lib\font.properties.th

When Tivoli Storage Manager scheduler is started on POLONIUM, it queries the server for a scheduled command, and since the schedule is still within the startup window, the incremental backup is restarted.

e. In the Tivoli Storage Manager server activity log, we can see how the connection was lost and a new session starts again for CL_MSCS01_SA as shown in Example 6-5.

Example 6-5 A new session is started for the client on the activity log

```
02/01/2005 16:29:54      ANR0480W Session 71 for node CL_MSCS01_SA (WinNT) terminated - connection with
client severed. (SESSION: 71)
```

02/01/2005 16:29:54 ANR0480W Session 70 for node CL_MSCS01_SA (WinNT) terminated - connection with client severed. (SESSION: 70)

02/01/2005 16:29:57 ANR0406I Session 72 started for node CL_MSCS01_SA (WinNT) (Tcp/Ip 9.1.39.187(2587)). (SESSION: 72)

02/01/2005 16:29:57 ANR1639I Attributes changed for node CL_MSCS01_SA: TCP Name from RADON to POLONIUM, TCP Address from 9.1.39.188 to 9.1.39.187, GUID from dd.41.76.e1.6e.59.11.d9.99.33.0-0.02.55.c6.fb.d0 to 77.24.3b.11.6e.5c.11.d9.86.b1.00.02.-55.c6.b9.07. (SESSION: 72)

02/01/2005 16:29:57 ANR0403I Session 72 ended for node CL_MSCS01_SA (WinNT). (SESSION: 72)

02/01/2005 16:31:26 ANR0406I Session 73 started for node CL_MSCS01_SA (WinNT) (Tcp/Ip 9.1.39.187(2590)). (SESSION: 73)

02/01/2005 16:31:28 ANR0406I Session 74 started for node CL_MSCS01_SA (WinNT) (Tcp/Ip 9.1.39.187(2592)). (SESSION: 74)

f. Also in the Tivoli Storage Manager server event log we see the scheduled event *restarted* as shown in Figure 6-36.

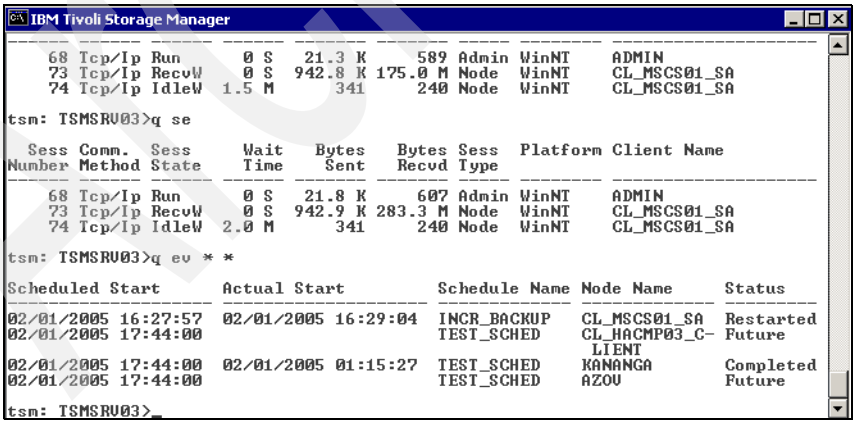

```
IBM Tivoli Storage Manager
------ ------ ------ ------ ------- ------- ----- -------- --------------------
    68 Tcp/Ip Run      0 S   21.3 K     589 Admin WinNT    ADMIN
    73 Tcp/Ip RecvW    0 S  942.8 K 175.0 M Node  WinNT    CL_MSCS01_SA
    74 Tcp/Ip IdleW  1.5 M     341     240 Node  WinNT    CL_MSCS01_SA

tsm: TSMSRV03>q se

  Sess Comm.  Sess      Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State     Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- --------------------
    68 Tcp/Ip Run      0 S   21.8 K     607 Admin WinNT    ADMIN
    73 Tcp/Ip RecvW    0 S  942.9 K 283.3 M Node  WinNT    CL_MSCS01_SA
    74 Tcp/Ip IdleW  2.0 M     341     240 Node  WinNT    CL_MSCS01_SA

tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name     Status
-------------------- -------------------- ------------- ------------- ---------
02/01/2005 16:27:57  02/01/2005 16:29:04  INCR_BACKUP   CL_MSCS01_SA  Restarted
02/01/2005 17:44:00                       TEST_SCHED    CL_HACMP03_C- Future
                                                         LIENT
02/01/2005 17:44:00  02/01/2005 01:15:27  TEST_SCHED    KANANGA       Completed
02/01/2005 17:44:00                       TEST_SCHED    AZOV          Future

tsm: TSMSRV03>_
```

Figure 6-36 Event log shows the schedule as restarted

6. The incremental backup ends without errors as we can see on the schedule log file in Example 6-6.

Example 6-6 Schedule log file shows the backup as completed

```
02/01/2005 16:43:30 Successful incremental backup of '\\cl_mscs01\j$'
02/01/2005 16:43:30 --- SCHEDULEREC STATUS BEGIN
02/01/2005 16:43:30 Total number of objects inspected:   17,878
02/01/2005 16:43:30 Total number of objects backed up:   15,084
02/01/2005 16:43:30 Total number of objects updated:         0
02/01/2005 16:43:30 Total number of objects rebound:         0
02/01/2005 16:43:30 Total number of objects deleted:         0
02/01/2005 16:43:30 Total number of objects expired:         0
02/01/2005 16:43:30 Total number of objects failed:          0
02/01/2005 16:43:30 Total number of bytes transferred:     1.10 GB
02/01/2005 16:43:30 Data transfer time:                   89.25 sec
02/01/2005 16:43:30 Network data transfer rate:       12,986.26 KB/sec
02/01/2005 16:43:30 Aggregate data transfer rate:      3,974.03 KB/sec
02/01/2005 16:43:30 Objects compressed by:                   0%
02/01/2005 16:43:30 Elapsed processing time:           00:04:51
02/01/2005 16:43:30 --- SCHEDULEREC STATUS END
02/01/2005 16:43:30 --- SCHEDULEREC OBJECT END INCR_BACKUP 02/01/2005 16:27:57
02/01/2005 16:43:30 Scheduled event 'INCR_BACKUP' completed successfully.
02/01/2005 16:43:30 Sending results for scheduled event 'INCR_BACKUP'.
02/01/2005 16:43:30 Results sent to server for scheduled event 'INCR_BACKUP'.
```

7. In the Tivoli Storage Manager server event log the schedule is *completed* as we see in Figure 6-37.

```
IBM Tivoli Storage Manager
    68 Tcp/Ip Run      0 S   24.3 K      705 Admin WinNT    ADMIN
    73 Tcp/Ip RecvW    0 S  943.3 K 883.7 M Node  WinNT    CL_MSCS01_SA
    74 Tcp/Ip IdleW  4.2 M      341      240 Node  WinNT    CL_MSCS01_SA

tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start          Schedule Name Node Name     Status
-------------------- --------------------- -------------- ------------ ---------
02/01/2005 16:27:57  02/01/2005 16:29:04   INCR_BACKUP   CL_MSCS01_SA  Restarted
02/01/2005 17:44:00                        TEST_SCHED    CL_HACMP03_C- Future
                                                          LIENT
02/01/2005 17:44:00  02/01/2005 01:15:27   TEST_SCHED    KANANGA       Completed
02/01/2005 17:44:00                        TEST_SCHED    AZOV          Future

tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start          Schedule Name Node Name     Status
-------------------- --------------------- -------------- ------------ ---------
02/01/2005 16:27:57  02/01/2005 16:29:04   INCR_BACKUP   CL_MSCS01_SA  Completed
02/01/2005 17:44:00                        TEST_SCHED    CL_HACMP03_C- Future
                                                          LIENT
02/01/2005 17:44:00  02/01/2005 01:15:27   TEST_SCHED    KANANGA       Completed
02/01/2005 17:44:00                        TEST_SCHED    AZOV          Future

tsm: TSMSRV03>_
```

Figure 6-37 Schedule completed on the event log

Checking that all files were correctly backed up

In this section we want to show a way of checking that the incremental backup did not miss any files while the failover process took place.

With this in mind, we perform these tasks:

1. In Example 6-4 on page 278, the last file reported as sent in the schedule log file is:

   \\cl_mscs01\j$\Program Files \IBM\ISC\AppServer\java\jre\lib\font.properties.th.

   And the first file sent after the failover is graphics.jar, also on the same path.

2. We open the explorer and go to this path, as we can see in Figure 6-38.

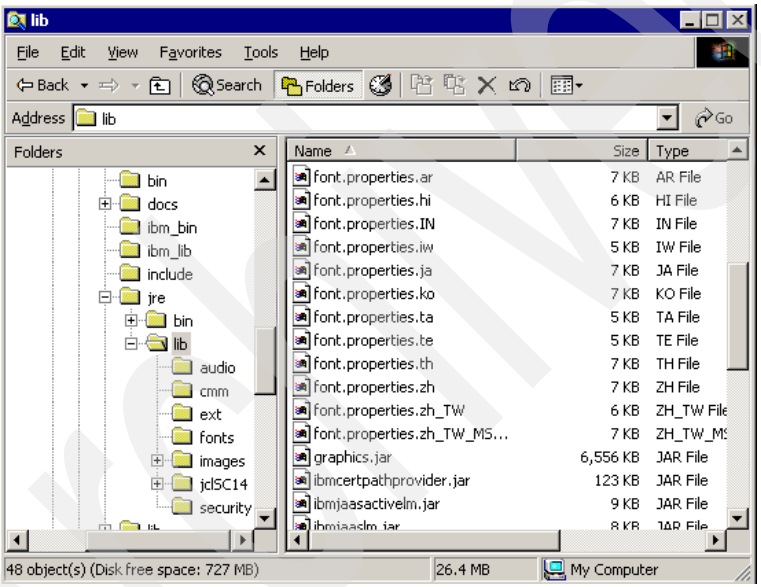

Figure 6-38 Windows explorer

3. If we have a look at last figure, between *font.properties.th* and *graphics.jar* files, there are three files not reported as backed up in the schedule log file.

4. We open a Tivoli Storage Manager GUI session to check, on the tree view of the Restore menu, whether these files were backed up (Figure 6-39).

Archived

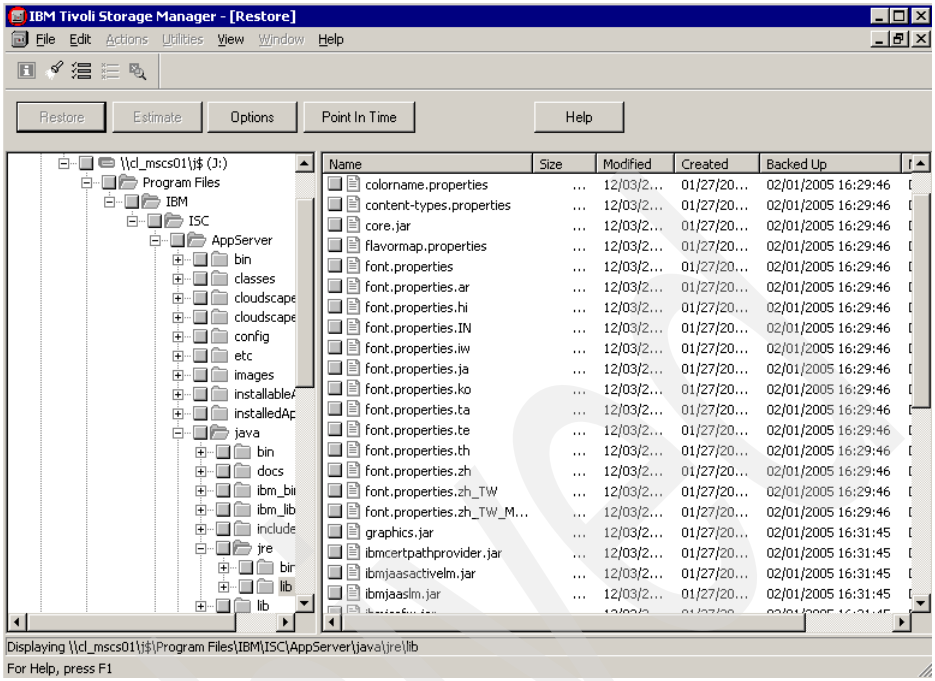

Figure 6-39 Checking backed up files using the TSM GUI

5. We see in Figure 6-39 that the client backed up the files correctly, even when they were not reported in the schedule log file. Since the session was lost, the client was not able of writing into the shared disk where the schedule log file is located.

Results summary

The test results show that, after a failure on the node that hosts the Tivoli Storage Manager scheduler service resource, a scheduled incremental backup started on one node is restarted and successfully completed on the other node that takes the failover.

This is true if the startup window used to define the schedule is not elapsed when the scheduler services restarts on the second node.

Testing client restore

Our second test consists of a scheduled restore of certain files under a directory.

Objective

The objective of this test is to show what happens when a client restore is started for a virtual node in the cluster, and the server that hosts the resources at that moment suddenly fails.

Activities

To do this test, we perform these tasks:

1. We open the Cluster Administrator to check which node hosts the Tivoli Storage Manager client resource: POLONIUM.
2. We schedule a client restore operation using the Tivoli Storage Manager server scheduler and we associate the schedule to CL_MSCS01_SA nodename.
3. A client session for CL_MSCS01_SA nodename starts on the server as shown in Figure 6-40.

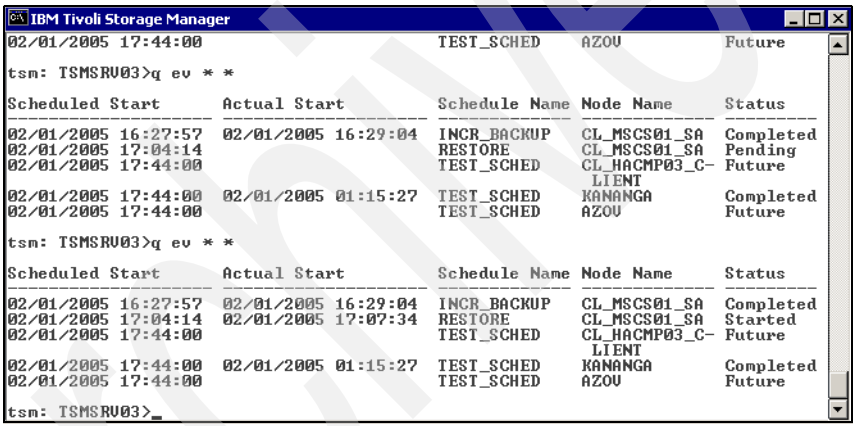

```
IBM Tivoli Storage Manager
02/01/2005 17:44:00                          TEST_SCHED     AZOV           Future

tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start          Schedule Name  Node Name      Status
-------------------- --------------------- -------------- -------------- ---------
02/01/2005 16:27:57  02/01/2005 16:29:04   INCR_BACKUP    CL_MSCS01_SA   Completed
02/01/2005 17:04:14                        RESTORE        CL_MSCS01_SA   Pending
02/01/2005 17:44:00                        TEST_SCHED     CL_HACMP03_C-  Future
                                                           LIENT
02/01/2005 17:44:00  02/01/2005 01:15:27   TEST_SCHED     KANANGA        Completed
02/01/2005 17:44:00                        TEST_SCHED     AZOV           Future

tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start          Schedule Name  Node Name      Status
-------------------- --------------------- -------------- -------------- ---------
02/01/2005 16:27:57  02/01/2005 16:29:04   INCR_BACKUP    CL_MSCS01_SA   Completed
02/01/2005 17:04:14  02/01/2005 17:07:34   RESTORE        CL_MSCS01_SA   Started
02/01/2005 17:44:00                        TEST_SCHED     CL_HACMP03_C-  Future
                                                           LIENT
02/01/2005 17:44:00  02/01/2005 01:15:27   TEST_SCHED     KANANGA        Completed
02/01/2005 17:44:00                        TEST_SCHED     AZOV           Future

tsm: TSMSRV03>_
```

Figure 6-40 Scheduled restore started for CL_MSCS01_SA

4. The client starts restoring files as we can see on the schedule log file in Example 6-7:

Example 6-7 Schedule log file shows the client restoring files

```
02/01/2005 17:23:38 Node Name: CL_MSCS01_SA
02/01/2005 17:23:38 Session established with server TSMSRV03: AIX-RS/6000
02/01/2005 17:23:38   Server Version 5, Release 3, Level 0.0
02/01/2005 17:23:38   Server date/time: 02/01/2005 17:16:25  Last access:
02/01/2005 17:15:40
02/01/2005 17:23:38 --- SCHEDULEREC QUERY BEGIN
02/01/2005 17:23:38 --- SCHEDULEREC QUERY END
02/01/2005 17:23:38 Next operation scheduled:
```

```
02/01/2005 17:23:38
------------------------------------------------------------
02/01/2005 17:23:38 Schedule Name:         RESTORE
02/01/2005 17:23:38 Action:                Restore
02/01/2005 17:23:38 Objects:               j:\tsm_images\tsmsrv5300_win\tsm64\*
02/01/2005 17:23:38 Options:               -subdir=yes -replace=yes
02/01/2005 17:23:38 Server Window Start:   17:15:17 on 02/01/2005
02/01/2005 17:23:38
------------------------------------------------------------
02/01/2005 17:23:38 Command will be executed in 2 minutes.
02/01/2005 17:25:38
Executing scheduled command now.
02/01/2005 17:25:38 Node Name: CL_MSCS01_SA
02/01/2005 17:25:38 Session established with server TSMSRV03: AIX-RS/6000
02/01/2005 17:25:38   Server Version 5, Release 3, Level 0.0
02/01/2005 17:25:38   Server date/time: 02/01/2005 17:18:25  Last access:
02/01/2005 17:16:25
02/01/2005 17:25:38 --- SCHEDULEREC OBJECT BEGIN RESTORE 02/01/2005 17:15:17
02/01/2005 17:25:38 Restore function invoked.
02/01/2005 17:25:39 ANS1247I Waiting for files from the server...Restoring
0 \\cl_mscs01\j$\TSM_Images\TSMSRV5300_WIN\TSM64\chs [Done]
02/01/2005 17:25:40 Restoring              0
\\cl_mscs01\j$\TSM_Images\TSMSRV5300_WIN\TSM64\cht [Done]
02/01/2005 17:25:40 Restoring              0
\\cl_mscs01\j$\TSM_Images\TSMSRV5300_WIN\TSM64\deu [Done]
02/01/2005 17:25:40 Restoring              0
\\cl_mscs01\j$\TSM_Images\TSMSRV5300_WIN\TSM64\driver [Done]
02/01/2005 17:25:40 Restoring              0
\\cl_mscs01\j$\TSM_Images\TSMSRV5300_WIN\TSM64\esp [Done]
...............................
02/01/2005 17:25:49 Restoring            729
\\cl_mscs01\j$\TSM_Images\TSMSRV5300_WIN\TSM64\cht\program
files\Tivoli\TSM\console\working_cht.htm [Done]
```

5. While the client is restoring the files, we force POLONIUM to fail. The following sequence takes place:

   a. The client loses temporarily its connection with the server, and the session is terminated as we can see on the Tivoli Storage Manager server activity log in Example 6-8.

Example 6-8 Connection is lost on the server

```
02/01/2005 17:18:38      ANR0480W Session 84 for node CL_MSCS01_SA (WinNT) terminated - connection with
client severed. (SESSION: 84)
```

Archived

b. In the Cluster Administrator, POLONIUM is not in the cluster and RADON begins to bring online the resources.

c. After a while the resources are online on RADON.

d. When the Tivoli Storage Manager scheduler service resource is again online on RADON, and queries the server for a schedule, if the startup window for the scheduled operation is not elapsed, the restore process restarts from the beginning, as we can see on the schedule log file in Example 6-9.

Example 6-9 Schedule log for the client starting the restore again

```
02/01/2005 17:27:24 Querying server for next scheduled event.
02/01/2005 17:27:24 Node Name: CL_MSCS01_SA
02/01/2005 17:27:24 Session established with server TSMSRV03: AIX-RS/6000
02/01/2005 17:27:24   Server Version 5, Release 3, Level 0.0
02/01/2005 17:27:24   Server date/time: 02/01/2005 17:20:11  Last access:
02/01/2005 17:18:42
02/01/2005 17:27:24 --- SCHEDULEREC QUERY BEGIN
02/01/2005 17:27:24 --- SCHEDULEREC QUERY END
02/01/2005 17:27:24 Next operation scheduled:
02/01/2005 17:27:24
------------------------------------------------------------
02/01/2005 17:27:24 Schedule Name:         RESTORE
02/01/2005 17:27:24 Action:                Restore
02/01/2005 17:27:24 Objects:               j:\tsm_images\tsmsrv5300_win\tsm64\*
02/01/2005 17:27:24 Options:               -subdir=yes -replace=yes
02/01/2005 17:27:24 Server Window Start:   17:15:17 on 02/01/2005
02/01/2005 17:27:24
------------------------------------------------------------
02/01/2005 17:27:24 Command will be executed in 1 minute.
02/01/2005 17:28:24
Executing scheduled command now.
02/01/2005 17:28:24 Node Name: CL_MSCS01_SA
02/01/2005 17:28:24 Session established with server TSMSRV03: AIX-RS/6000
02/01/2005 17:28:24   Server Version 5, Release 3, Level 0.0
02/01/2005 17:28:24   Server date/time: 02/01/2005 17:21:11  Last access:
02/01/2005 17:20:11
02/01/2005 17:28:24 --- SCHEDULEREC OBJECT BEGIN RESTORE 02/01/2005 17:15:17
02/01/2005 17:28:24 Restore function invoked.
02/01/2005 17:28:25 ANS1247I Waiting for files from the server...Restoring
0 \\cl_mscs01\j$\TSM_Images\TSMSRV5300_WIN\TSM64\chs [Done]
02/01/2005 17:28:26 Restoring               0
\\cl_mscs01\j$\TSM_Images\TSMSRV5300_WIN\TSM64\cht [Done]
02/01/2005 17:28:26 Restoring               0
\\cl_mscs01\j$\TSM_Images\TSMSRV5300_WIN\TSM64\deu [Done]
02/01/2005 17:28:26 Restoring               0
\\cl_mscs01\j$\TSM_Images\TSMSRV5300_WIN\TSM64\driver [Done]
```

e. In the activity log of Tivoli Storage Manager server we see that a new session is started for CL_MSCS01_SA as shown in Example 6-10.

Example 6-10 New session started on the activity log for CL_MSCS01_SA

```
02/01/2005 17:18:38      ANR0480W Session 84 for node CL_MSCS01_SA
(WinNT) terminated - connection with client severed. (SESSION: 84)

02/01/2005 17:18:42      ANR0406I Session 85 started for node
CL_MSCS01_SA (WinNT) (Tcp/Ip 9.1.39.188(2895)). (SESSION: 85)

02/01/2005 17:18:42      ANR1639I Attributes changed for node
CL_MSCS01_SA: TCP Name from POLONIUM to RADON, TCP Address from
9.1.39.187 to 9.1.39.188, GUID from
77.24.3b.11.6e.5c.11.d9.86.b1.0-0.02.55.c6.b9.07 to
dd.41.76.e1.6e.59.11.d9.99.33.00.02.-55.c6.fb.d0. (SESSION: 85)

02/01/2005 17:18:42      ANR0403I Session 85 ended for node CL_MSCS01_SA
(WinNT). (SESSION: 85)

02/01/2005 17:20:11      ANR0406I Session 86 started for node
CL_MSCS01_SA (WinNT) (Tcp/Ip 9.1.39.188(2905)). (SESSION: 86)

02/01/2005 17:20:11      ANR0403I Session 86 ended for node CL_MSCS01_SA
(WinNT).  (SESSION: 86)

02/01/2005 17:21:11      ANR0406I Session 87 started for node
CL_MSCS01_SA (WinNT) (Tcp/Ip 9.1.39.188(2906)). (SESSION: 87)
```

f. And the event log of Tivoli Storage Manager server shows the schedule as restarted (Figure 6-41).

Archived

```
IBM Tivoli Storage Manager
  Sess Comm.  Sess      Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State     Time    Sent   Recvd Type
------ ------ ------ ------- ------- ------ ----- -------- ----------------
    77 Tcp/Ip Run       0 S   25.0 K   1.9 K Admin WinNT    ADMIN

tsm: TSMSRV03>q se

  Sess Comm.  Sess      Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State     Time    Sent   Recvd Type
------ ------ ------ ------- ------- ------ ----- -------- ----------------
    77 Tcp/Ip Run       0 S   25.3 K   1.9 K Admin WinNT    ADMIN
    87 Tcp/Ip SendW     0 S   48.8 M     751 Node  WinNT    CL_MSCS01_SA

tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name     Status
-------------------- -------------------- ------------- ------------- ---------
02/01/2005 16:27:57  02/01/2005 16:29:04  INCR_BACKUP   CL_MSCS01_SA  Completed
02/01/2005 17:15:17  02/01/2005 17:18:25  RESTORE       CL_MSCS01_SA  Restarted
02/01/2005 17:44:00                       TEST_SCHED    CL_HACMP03_C- Future
                                                         LIENT
02/01/2005 17:44:00  02/01/2005 01:15:27  TEST_SCHED    KANANGA       Completed
02/01/2005 17:44:00                       TEST_SCHED    AZOV          Future

tsm: TSMSRV03>_
```

Figure 6-41 Schedule restarted on the event log for CL_MSCS01_SA

6. When the restore completes we can see the final statistics in the schedule log file of the client for a successful operation as shown in Example 6-11.

Example 6-11 Schedule log file on client shows statistics for the restore operation

```
Restore processing finished.
02/01/2005 17:29:42 --- SCHEDULEREC STATUS BEGIN
02/01/2005 17:29:42 Total number of objects restored:       675
02/01/2005 17:29:42 Total number of objects failed:           0
02/01/2005 17:29:42 Total number of bytes transferred:    221.68 MB
02/01/2005 17:29:42 Data transfer time:                    38.85 sec
02/01/2005 17:29:42 Network data transfer rate:        5,842.88 KB/sec
02/01/2005 17:29:42 Aggregate data transfer rate:      2,908.60 KB/sec
02/01/2005 17:29:42 Elapsed processing time:           00:01:18
02/01/2005 17:29:42 --- SCHEDULEREC STATUS END
02/01/2005 17:29:42 --- SCHEDULEREC OBJECT END RESTORE 02/01/2005 17:15:17
02/01/2005 17:29:42 --- SCHEDULEREC STATUS BEGIN
02/01/2005 17:29:42 --- SCHEDULEREC STATUS END
02/01/2005 17:29:42 Scheduled event ‘RESTORE’ completed successfully.
02/01/2005 17:29:42 Sending results for scheduled event ‘RESTORE’.
02/01/2005 17:29:42 Results sent to server for scheduled event ‘RESTORE’.
```

7. And the event log of Tivoli Storage Manager server shows the scheduled operation as completed (Figure 6-42).

```
IBM Tivoli Storage Manager
02/01/2005 17:44:00                          TEST_SCHED     AZOV           Future

tsm: TSMSRV03>  q se
ANR2000E Unknown command - ?Q.
ANS8001I Return code 2.

tsm: TSMSRV03>q se

  Sess Comm.  Sess      Wait    Bytes    Bytes Sess  Platform Client Name
Number Method State     Time     Sent    Recvd Type
------ ------ ------ -------- -------- -------- ------ -------- --------------------
    77 Tcp/Ip Run       0 S    26.2 K    2.0 K Admin WinNT    ADMIN

tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name      Status
-------------------- -------------------- -------------- -------------- ---------
02/01/2005 16:27:57  02/01/2005 16:29:04  INCR_BACKUP    CL_MSCS01_SA   Completed
02/01/2005 17:15:17  02/01/2005 17:18:25  RESTORE        CL_MSCS01_SA   Completed
02/01/2005 17:44:00                       TEST_SCHED     CL_HACMP03_C-  Future
                                                          LIENT
02/01/2005 17:44:00  02/01/2005 01:15:27  TEST_SCHED     KANANGA        Completed
02/01/2005 17:44:00                       TEST_SCHED     AZOV           Future

tsm: TSMSRV03>_
```

Figure 6-42 Event completed for schedule name RESTORE

Results summary

The test results show that, after a failure on the node that hosts the Tivoli Storage Manager client scheduler instance, a scheduled restore operation started on this node is started again on the second node of the cluster when the service is online.

This is true if the startup window for the scheduled restore operation is not elapsed when the scheduler client is online again on the second node.

Also notice that the restore is not restarted from the point of failure, but started from the beginning. The scheduler queries the Tivoli Storage Manager server for a scheduled operation and a new session is opened for the client after the failover.

6.5 Tivoli Storage Manager Client on Windows 2003

In this section we describe how we configure the Tivoli Storage Manager client software to be capable of running in our MSCS Windows 2003, the same cluster we installed and configured in 4.4, “Windows 2003 MSCS installation and configuration” on page 44.

6.5.1 Windows 2003 lab setup

Our lab environment consists of a Microsoft Windows 2003 Enterprise Server Cluster with two nodes, SENEGAL and TONGA, as we can see in Figure 6-43.

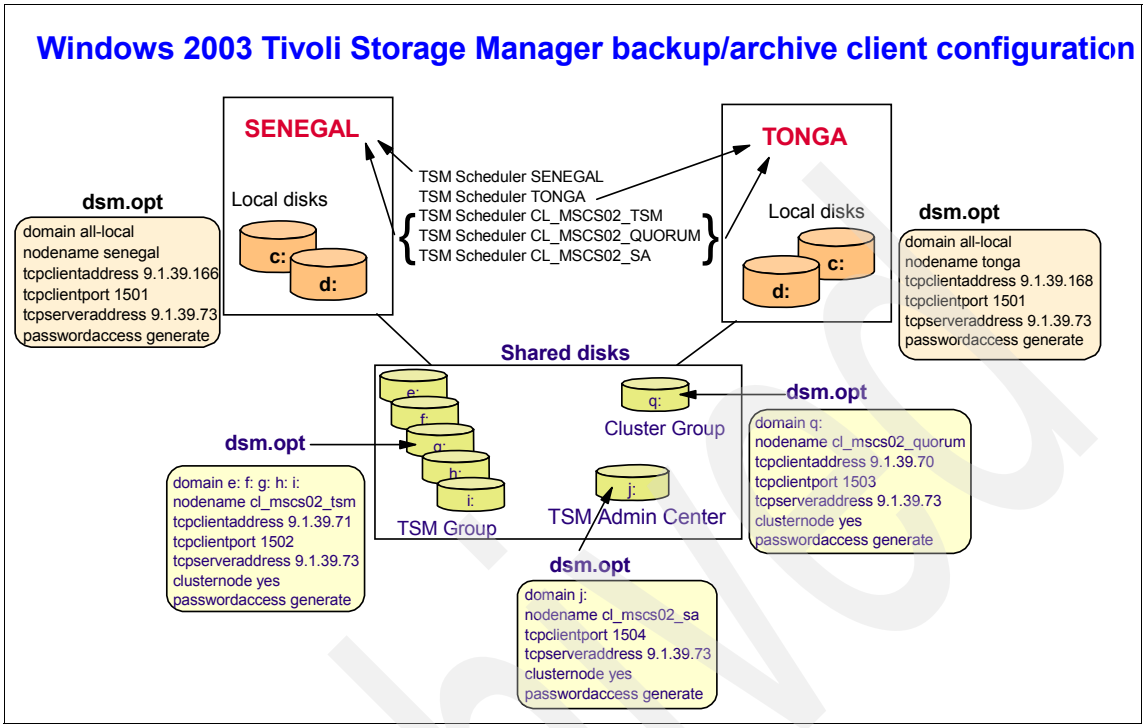

Figure 6-43 Tivoli Storage Manager backup/archive clustering client (Win.2003)

Refer to Table 4-4 on page 46, Table 4-5 on page 47 and Table 4-6 on page 47 for details of the MSCS cluster configuration used in our lab.

Table 6-3 and Table 6-4 show the specific Tivoli Storage Manager backup/archive client configuration we use for the purpose of this section.

Table 6-3 Windows 2003 TSM backup/archive configuration for local nodes

Local node 1	
TSM nodename	SENEGAL
Backup domain	c: d: systemstate systemservices
Scheduler service name	TSM Scheduler SENEGAL
Client Acceptor service name	TSM Client Acceptor SENEGAL
Remote Client Agent service name	TSM Remote Client Agent SENEGAL

Local node 2	
TSM nodename	TONGA
Backup domain	c: d: systemstate systemservices
Scheduler service name	TSM Scheduler TONGA
Client Acceptor service name	TSM Client Acceptor TONGA
Remote Client Agent service name	TSM Remote Client Agent TONGA

Table 6-4 Windows 2003 TSM backup/archive client for virtual nodes

Virtual node 1	
TSM nodename	CL_MSCS02_QUORUM
Backup domain	q:
Scheduler service name	TSM Scheduler CL_MSCS02_QUORUM
Client Acceptor service name	TSM Client Acceptor CL_MSCS02_QUORUM
Remote Client Agent service name	TSM Remote Client Agent CL_MSCS02_QUORUM
Cluster group name	Cluster Group
Virtual node 2	
TSM nodename	CL_MSCS02_SA
Backup domain	j:
Scheduler service name	TSM Scheduler CL_MSCS02_SA
Client Acceptor service name	TSM Client Acceptor CL_MSCS02_SA
Remote Client Agent service name	TSM Remote Client Agent CL_MSCS02_SA
Cluster group name	TSM Admin Center
Virtual node 3	
TSM nodename	CL_MSCS02_TSM
Backup domain	e: f: g: h: i:
Scheduler service name	TSM Scheduler CL_MSCS02_TSM

Archived

Virtual node 1	
TSM nodename	CL_MSCS02_QUORUM
Backup domain	q:
Scheduler service name	TSM Scheduler CL_MSCS02_QUORUM
Client Acceptor service name	TSM Client Acceptor CL_MSCS02_QUORUM
Remote Client Agent service name	TSM Remote Client Agent CL_MSCS02_QUORUM
Cluster group name	Cluster Group
Virtual node 2	
TSM nodename	CL_MSCS02_SA
Backup domain	j:
Scheduler service name	TSM Scheduler CL_MSCS02_SA
Client Acceptor service name	TSM Client Acceptor CL_MSCS02_SA
Remote Client Agent service name	TSM Remote Client Agent CL_MSCS02_SA
Cluster group name	TSM Admin Center
Virtual node 3	
Client Acceptor service name	TSM Client Acceptor CL_MSCS02_TSM
Remote Client Agent service name	TSM Remote Client Agent CL_MSCS02_TSM
Cluster group name	TSM Group

6.5.2 Windows 2003 Tivoli Storage Manager Client configurations

In this section we describe how to configure the Tivoli Storage Manager backup/archive client in our Windows 2000 MSCS environment. This is a two-step procedure:

1. Configuration to back up the local disk drives of each server
2. Configuration to back up shared disk drives of each group in the cluster

Archived

Configuring the client to back up local disks

The configuration for the backup of the local disks is the same as for any standalone client:

1. We create a nodename for each server (TONGA and SENEGAL) on the Tivoli Storage Manager server
2. We create the option file (dsm.opt) for each node on the local drive.

> **Important:** We should only use the `domain` option if not all local drives are going to be backed up. The default, if we do not specify anything, is backup all local drives and system objects. We should not include any cluster drive in the domain parameter.

3. We generate the password locally by either opening the backup-archive GUI or issuing a query on the command prompt, such as **`dsmc q se`**.
4. We create the local Tivoli Storage Manager services as needed for each node, opening the backup-archive GUI client and selecting **Utilities** →**Setup Wizard**. The names we use for each service are:

- For SENEGAL:
  - Tivoli Storage Manager Scheduler SENEGAL
  - Tivoli Storage Manager Client Acceptor SENEGAL
  - Tivoli Storage Manager Remote Client Agent SENEGAL
- For TONGA:
  - Tivoli Storage Manager Scheduler TONGA
  - Tivoli Storage Manager Client Acceptor TONGA
  - Tivoli Storage Manager Remote Client Agent TONGA

5. After the configuration, the Windows services menu appears as shown in Figure 6-44. These are the Windows services for TONGA. For SENEGAL we are presented with a very similar menu.

Archived

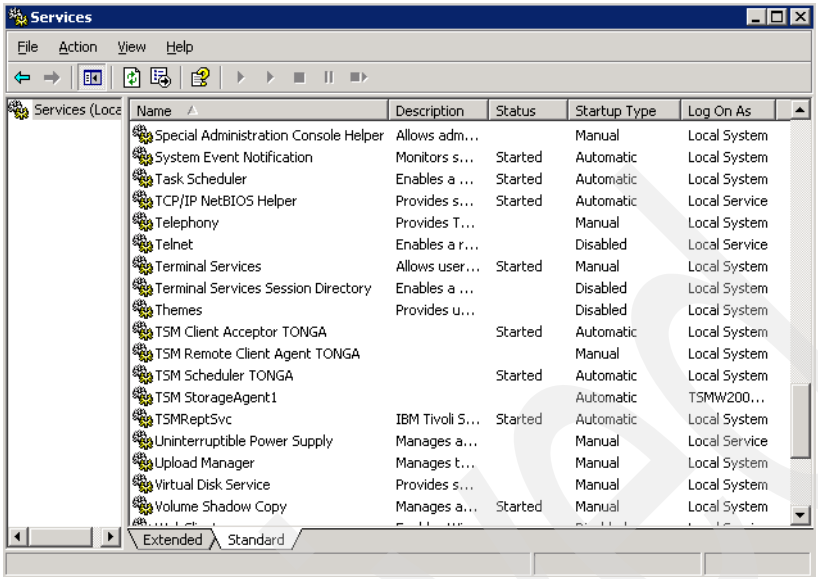

Figure 6-44 Tivoli Storage Manager client services

Configuring the client to back up shared disks

The configuration of Tivoli Storage Manager client to back up shared disks is slightly different for virtual nodes on MSCS.

For every resource group that has shared disks with backup requirements, we need to define an options file and an associated TSM scheduler service. If we want to use the Web client to access that virtual node from a browser, we also have to install the Web client services for that particular resource group.

The cluster environment for this section, formed by TONGA and SENEGAL, has the following resource groups:

- Cluster Group: Contains the quorum physical disk q:
- TSM Admin Center: Contains physical disk j:
- TSM Group: Contains physical disks e: f: g: h: i:

Each resource group needs its own unique nodename. This ensures that Tivoli Storage Manager client correctly manages the disk resources in case of failure on any physical node, independently of the node who hosts the resources at that time.

We created the following nodes on the Tivoli Storage Manager server:

- CL_MSCS02_QUORUM: for Cluster Group
- CL_MSCS02_SA: for TSM Admin Center
- CL_MSCS02_TSM: for TSM Group

For each group, the configuration process consists of the following tasks:

1. Creation of the option files
2. Password generation
3. Installation (on each physical node on the MSCS) of the TSM Scheduler service
4. Installation (on each physical node on the MSCS) of the TSM Web client services
5. Creation of a generic service resource for the TSM Scheduler service using the Cluster Administration
6. Creation of a generic service resource for the TSM Client Acceptor service using the Cluster Administration

We describe each activity in the following sections.

Creation of the option files

For each group in the cluster we need to create an option file that will be used by the Tivoli Storage Manager nodename attached to that group.

The option file must be located on one of the shared disks hosted by this group. This ensures that both physical nodes have access to the file.

The dsm.opt file *must* contain at least the following options:

- nodename: Specifies the name that this group uses when it backs up data to the Tivoli Storage Manager server
- domain: Specifies the disk drive letters managed by this group
- passwordaccess *generate*: Specifies that the client generates a new password when the old one expires, and this new password is kept in the Windows registry.
- clusternode yes: To specify that it is a virtual node of a cluster. This is the main difference between the option file for a virtual node and the option file for a physical node.

If we plan to use the schedmode promted option to schedule backups, and we plan to use **the Web client interface for each virtual node**, we also should specify the following options:

- tcpclientaddress: Specifies **the unique IP address for this resource group**
- tcpclientport: Specifies a different TCP port for each node
- httpport: Specifies a different http port to contact with.

There are other options we can specify but the ones mentioned above are a requirement for a correct implementation of the client.

In our environment we create the dsm.opt files in a directory called \tsm in the following drives:

- For the Cluster Group: drive q:
- For the Admin Center Group: drive j:
- For the TSM Group: drive g:

Option file for Cluster Group

The dsm.opt file for this group contains the following options:

```
nodename cl_mscs02_quorum
passwordaccess generate
tcpserveraddress 9.1.39.74
errorlogretention 7
errorlogname q:\tsm\dsmerror.log
schedlogretention 7
schedlogname q:\tsm\dsmsched.log
domain q:
clusternode yes
schedmode prompted
tcpclientaddress 9.1.39.70
tcpclientport 1502
httpport 1582
```

Option file for TSM Admin Center group

The **dsm.opt** file for this group contains the following options:

```
nodename cl_mscs02_sa
passwordaccess generate
tcpserveraddress 9.1.39.74
errorlogretention 7
errorlogname j:\tsm\dsmerror.log
schedlogretention 7
schedlogname j:\tsm\dsmsched.log
domain j:
clusternode yes
tcpclientport 1503
httpport 1583
```

Archived

Option file for TSM Group

The **dsm.opt** file for this group contains the following options:

```
nodename cl_mscs02_tsm
passwordaccess generate
tcpserveraddress 9.1.39.74
errorlogretention 7
errorlogname g:\tsm\dsmerror.log
schedlogretention 7
schedlogname g:\tsm\dsmsched.log
domain e: f: g: h: i:
clusternode yes
schedmode prompted
tcpclientaddress 9.1.39.71
tcpclientport 1504
httpport 1584
```

Password generation

The Windows registry of each server needs to be updated with the password used to register, in the Tivoli Storage Manager server, the nodenames for each resource group.

> **Important:** The following steps require that the commands shown below are run on both nodes while they own the resources. We recommend to move all resources to one of the nodes, complete the tasks below, and then move all resources to the other node and repeat the tasks.

Since the dsm.opt is located for each node in a different location, we need to specify the path for each using the *-optfile* option of the **dsmc** command.

1. We run the following command on a MS-DOS prompt on the Tivoli Storage Manager client directory (c:\program files\tivoli\tsm\baclient):

   ```
   dsmc q se -optfile=q:\tsm\dsm.opt
   ```

2. Tivoli Storage Manager prompts the nodename for the client (the specified in dsm.opt). If it is correct, press Enter.
3. Tivoli Storage Manager next asks for a password. We type the password and press Enter. Figure 6-45 shows the output of the command.

Archived

```
C:\WINDOWS\system32\cmd.exe

Please enter password for user id "CL_MSCS02_QUORUM": ******

Session established with server TSMSRV03: AIX-RS/6000
  Server Version 5, Release 3, Level 0.0
  Server date/time: 02/03/2005 14:18:58  Last access: 02/03/2005 14:18:58

TSM Server Connection Information

Server Name.............: TSMSRV03
Server Type.............: AIX-RS/6000
Server Version..........: Ver. 5, Rel. 3, Lev. 0.0
Last Access Date........: 02/03/2005 14:18:58
Delete Backup Files.....: "No"
Delete Archive Files....: "Yes"

Node Name...............: CL_MSCS02_QUORUM
User Name...............:


C:\Program Files\Tivoli\TSM\baclient>_
```

Figure 6-45 Generating the password in the registry

Note: The password is kept in the Windows registry of this node and we do not need to type it any more. The client reads the password from the registry every time it opens a session with the Tivoli Storage Manager server.

4. We repeat the command for the other nodes:

```
dsmc q se -optfile=j:\tsm\dsm.opt
dsmc q se -optfile=g:\tsm\dsm.opt
```

5. We move the resources to the other node and repeat steps 1 to 4.

Installing the TSM Scheduler service

For backup automation, using the Tivoli Storage Manager scheduler, we need to install and configure one scheduler service for each resource group.

Important: We must install the scheduler service for each cluster group exactly with the same name, which is case sensitive, on each of the physical nodes and on the MSCS Cluster Administrator, otherwise failover will not work.

1. We need to be sure we are located on the node that hosts all resources, in order to start with the Tivoli Storage Manager scheduler service installation.
2. We begin the installation of the scheduler service for each group on TONGA. This is the node that hosts the resources. We use the **`dsmcutil program`**. This utility is located on the Tivoli Storage Manager client installation path (c:\program files\tivoli\tsm\baclient).

   In our lab we installed three scheduler services, one for each cluster group.

3. We open an MS-DOS command line and, in the Tivoli Storage Manager client installation path, we issue the following command:

```
dsmcutil inst sched /name:"TSM Scheduler CL_MSCS02_QUORUM"
/clientdir:"c:\program files\tivoli\tsm\baclient" /optfile:q:\tsm\dsm.opt
/node:CL_MSCS02_QUORUM /password:itsosj /clustername:CL_MSCS02
/clusternode:yes /autostart:no
```

4. The result is shown in Figure 6-46.

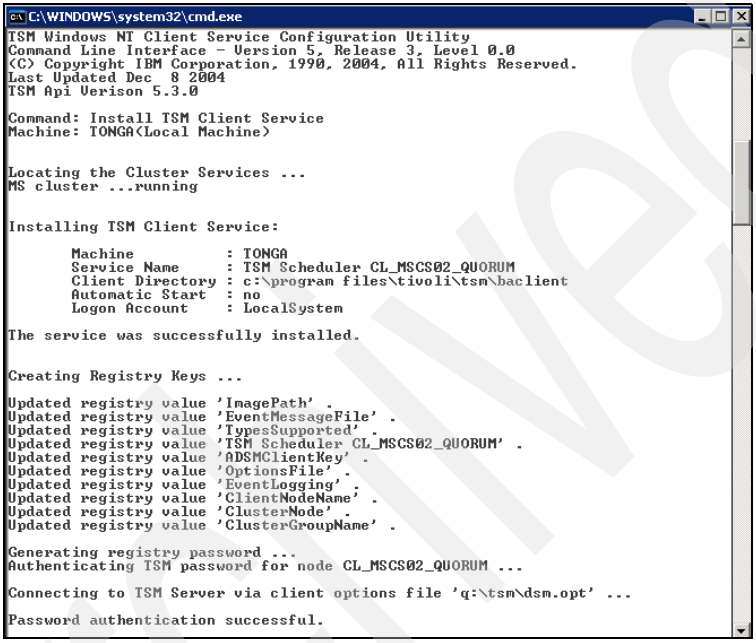

```
C:\WINDOWS\system32\cmd.exe
TSM Windows NT Client Service Configuration Utility
Command Line Interface - - Version 5, Release 3, Level 0.0
(C) Copyright IBM Corporation, 1990, 2004, All Rights Reserved.
Last Updated Dec  8 2004
TSM Api Verison 5.3.0

Command: Install TSM Client Service
Machine: TONGA(Local Machine)


Locating the Cluster Services ...
MS cluster ...running


Installing TSM Client Service:

        Machine          : TONGA
        Service Name     : TSM Scheduler CL_MSCS02_QUORUM
        Client Directory : c:\program files\tivoli\tsm\baclient
        Automatic Start  : no
        Logon Account    : LocalSystem

The service was successfully installed.


Creating Registry Keys ...

Updated registry value 'ImagePath' .
Updated registry value 'EventMessageFile' .
Updated registry value 'TypesSupported' .
Updated registry value 'TSM Scheduler CL_MSCS02_QUORUM' .
Updated registry value 'ADSMClientKey' .
Updated registry value 'OptionsFile' .
Updated registry value 'EventLogging' .
Updated registry value 'ClientNodeName' .
Updated registry value 'ClusterNode' .
Updated registry value 'ClusterGroupName' .

Generating registry password ...
Authenticating TSM password for node CL_MSCS02_QUORUM ...

Connecting to TSM Server via client options file 'q:\tsm\dsm.opt' ...

Password authentication successful.
```

Figure 6-46 Result of Tivoli Storage Manager scheduler service installation

5. We repeat this command to install the scheduler service for TSM Admin Center Group, changing the information as needed. The command is:

```
dsmcutil inst sched /name:"TSM Scheduler CL_MSCS02_SA"
/clientdir:"c:\Program Files\Tivoli\tsm\baclient" /optfile:j:\tsm\dsm.opt
/node:CL_MSCS02_SA /password:itsosj /clusternode:yes /clustername:CL_MSCS02
/autostart:no
```

6. And we do this again to install the scheduler service for TSM Group we use:

```
dsmcutil inst sched /name:"TSM Scheduler CL_MSCS02_TSM"
/clientdir:"c:\Program Files\Tivoli\tsm\baclient" /optfile:g:\tsm\dsm.opt
/node:CL_MSCS02_TSM /password:itsosj /clusternode:yes
/clustername:CL_MSCS02 /autostart:no
```

7. Be sure to stop all services using the Windows service menu before continuing.

Archived

8. We move the resources to the second node, SENEGAL, and run exactly the same commands as before (steps 1 to 7).

> **Attention:** the Tivoli Storage Manager scheduler service names used on both nodes must match. Also remember to use the same parameters for the dsmcutil tool. Do not forget the clusternode yes and clustername options.

So far the Tivoli Storage Manager scheduler services are installed on both nodes of the cluster with exactly the same names for each resource group. The last task consists of the definition for a new resource on each cluster group.

Creating a generic service resource for TSM scheduler service

For a correct configuration of the Tivoli Storage Manager client we define, for each cluster group, a new generic service resource. This resource will be related to the scheduler service name created for this group.

> **Important:** Before continuing, we make sure to stop all services created in "Installing the TSM Scheduler service" on page 298 on all nodes. We also make sure all resources are on one of the nodes.

1. We open the Cluster Administrator menu on the node that hosts all resources and select the first group (Cluster Group). We right-click the name and select **New** →**Resource** as shown in Figure 6-47.

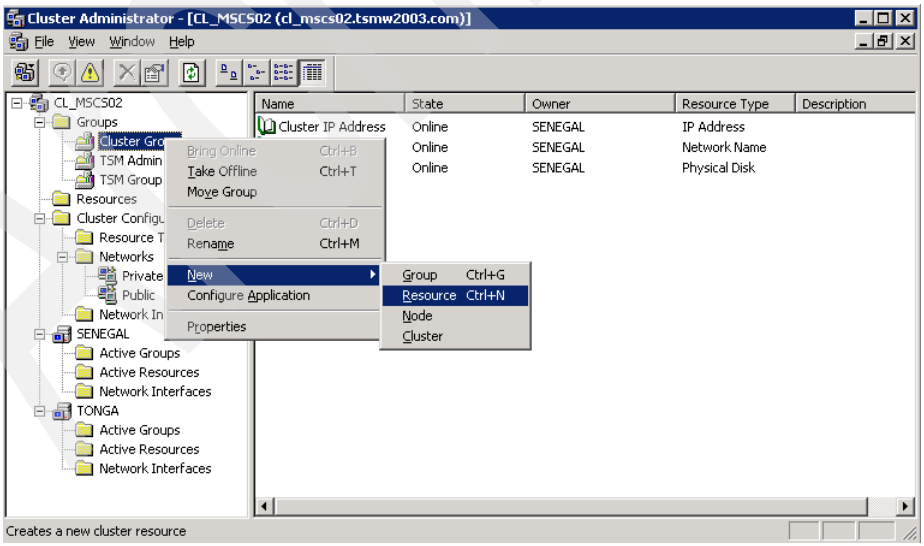

Figure 6-47 Creating new resource for Tivoli Storage Manager scheduler service

2. We type a Name for the resource (we recommend to use the same name as the scheduler service) and select **Generic Service** as resource type. We click **Next** as shown in Figure 6-48.

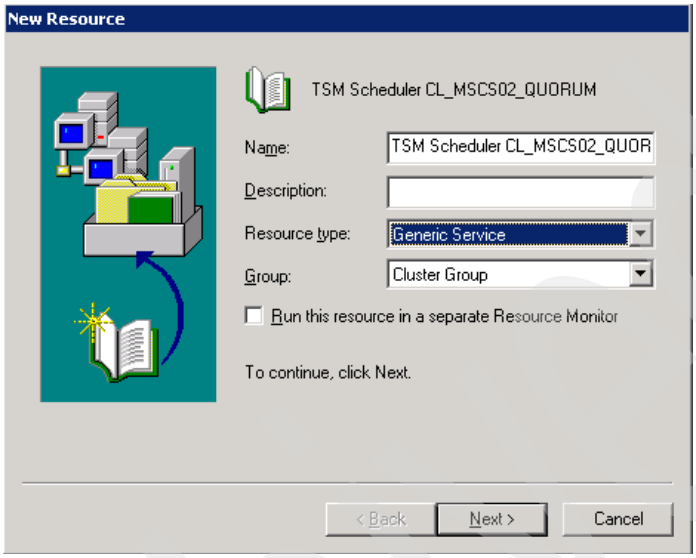

Figure 6-48 Definition of TSM Scheduler generic service resource

3. We leave both nodes as possible owners for the resource and click **Next** (Figure 6-49).

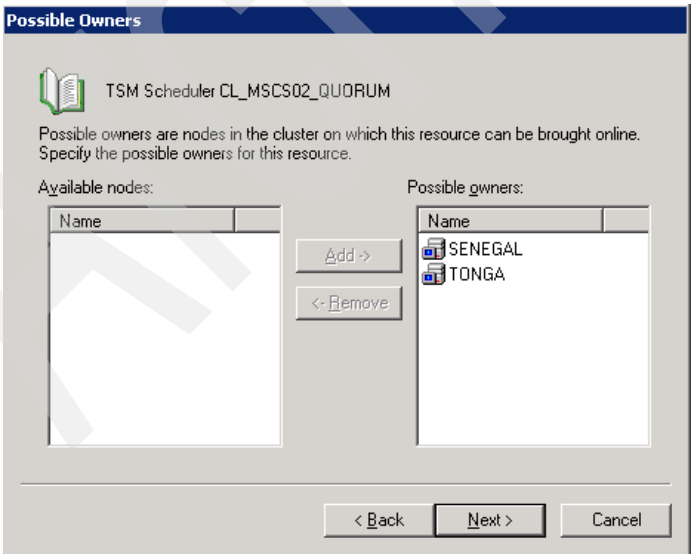

Figure 6-49 Possible owners of the resource

4. We **Add** the disk resource (q:) on Dependencies as shown in Figure 6-50. We click **Next**.

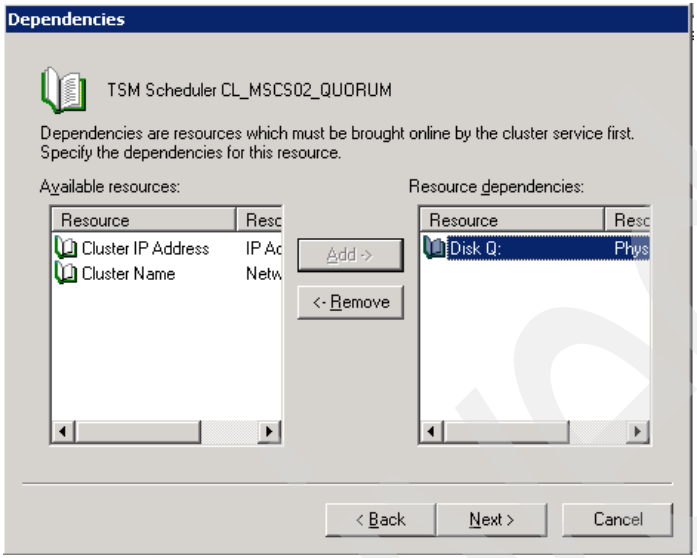

Figure 6-50 Dependencies

5. Next (see Figure 6-51) we type a Service name. This must match the name used while installing the scheduler service on both nodes. We click **Next**:

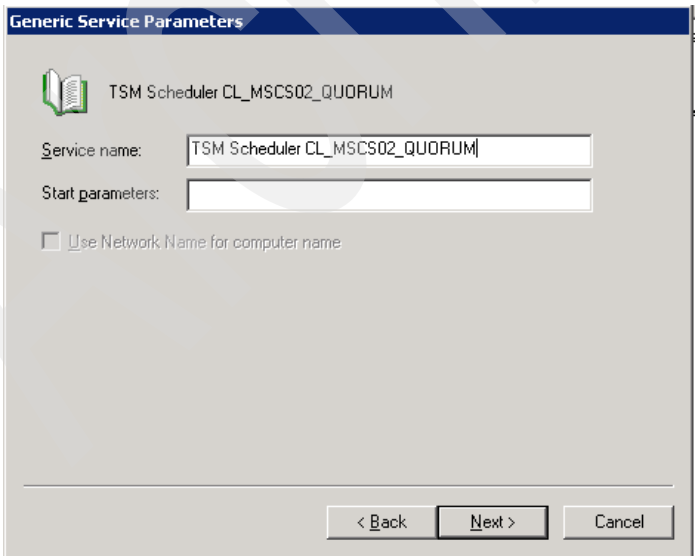

Figure 6-51 Generic service parameters

6. We click **Add** to type the Registry Key where Windows 2003 will save the generated password for the client. The registry key is

   `SOFTWARE\IBM\ADSM\CurrentVersion\BackupClient\Nodes\<nodename>\<tsmservername>`

   We click **OK** (Figure 6-52).

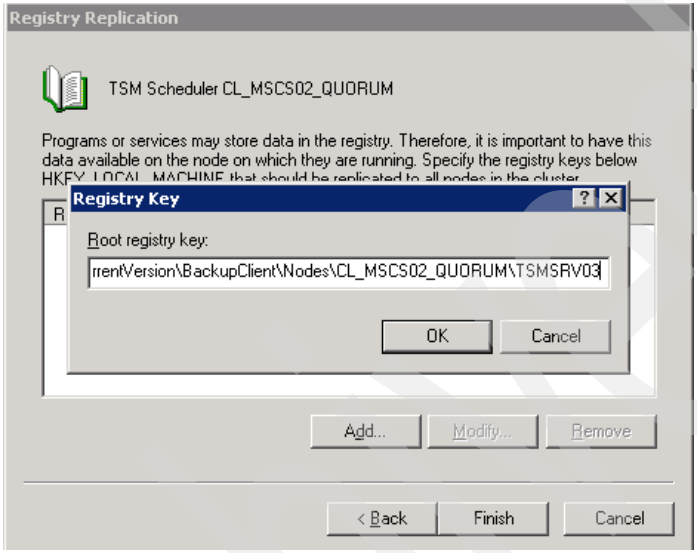

Figure 6-52 Registry key replication

7. If the resource creation is successful, an information menu appears as shown in Figure 6-53. We click **OK**.

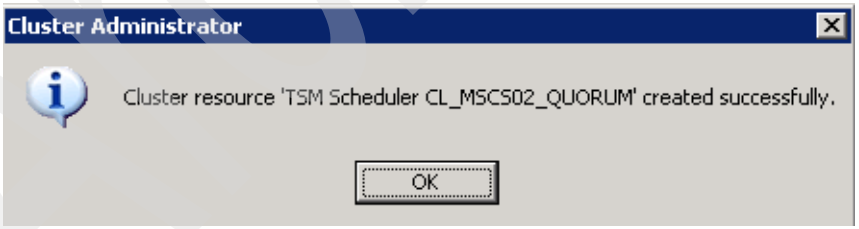

Figure 6-53 Successful cluster resource installation

8. As seen in Figure 6-54, the Cluster Group is offline because the new resource is also offline. We bring it online.

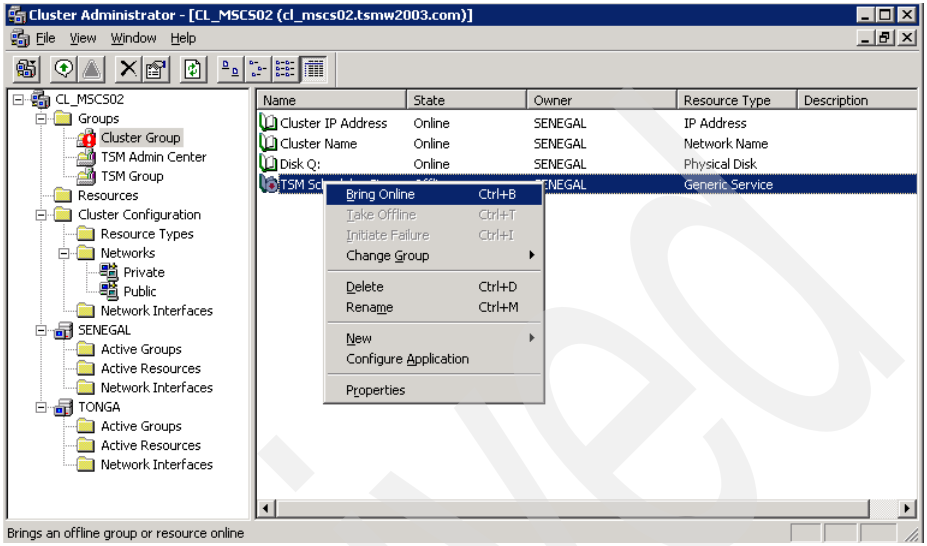

Figure 6-54 Bringing online the Tivoli Storage Manager scheduler service

9. The Cluster Administrator menu after all resources are online is shown in Figure 6-55.

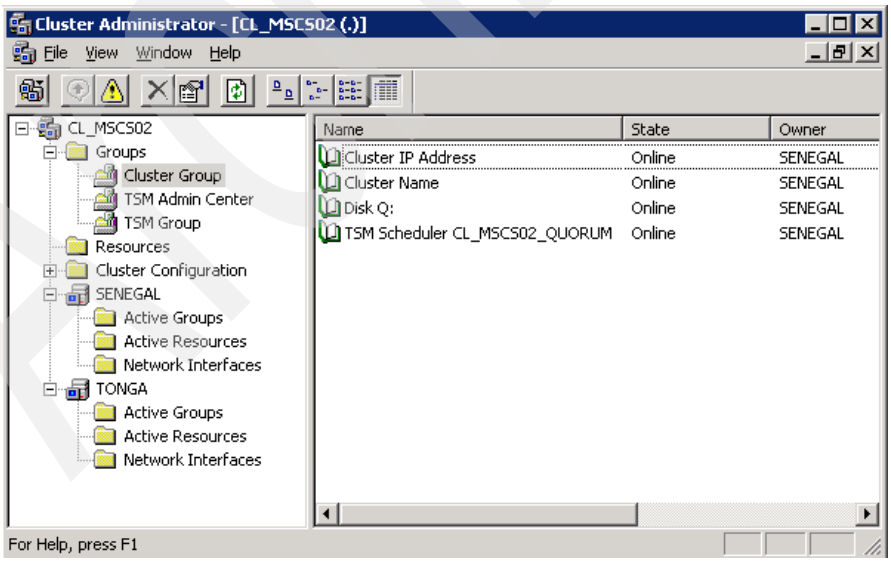

Figure 6-55 Cluster group resources online

10. If we go to the Windows service menu, Tivoli Storage Manager scheduler service is started on SENEGAL, the node which now hosts this resource group (Figure 6-56).

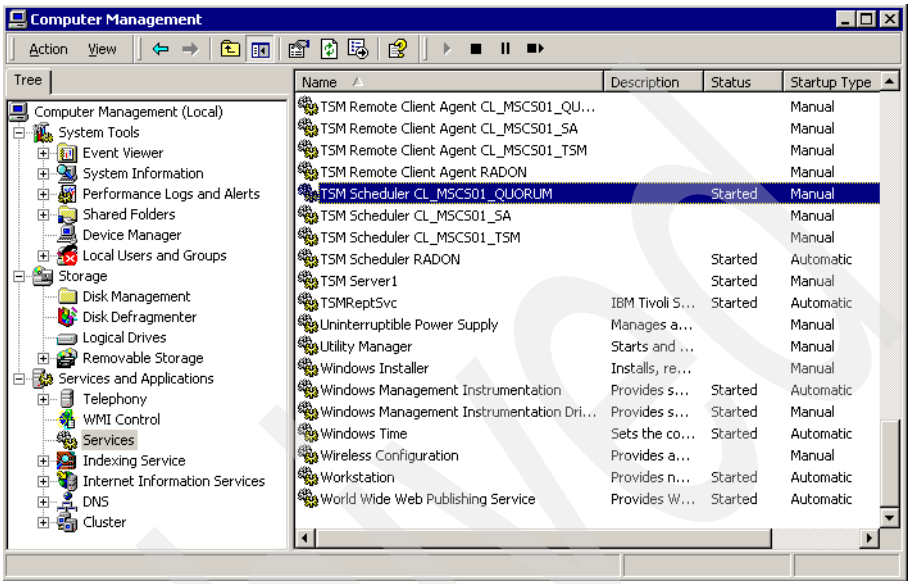

Figure 6-56 Windows service menu

11. We repeat steps 1-10 to create the Tivoli Storage Manager scheduler generic service resource for TSM Admin Center and TSM Group cluster groups. The resource names are:
   - TSM Scheduler CL_MSCS02_SA: for TSM Admin Center resource group
   - TSM Scheduler CL_MSCS02_TSM: for TSM Group resource group.

Important: To back up, archive, or retrieve data residing on MSCS, the Windows account used to start the Tivoli Storage Manager scheduler service on each local node must belong to the Administrators or Domain Administrators group or Backup Operators group.

12. We move the resources to check that Tivoli Storage Manager scheduler services successfully start on TONGA while they are stopped on SENEGAL.

Note: Use only the Cluster Administration menu to bring online/offline the Tivoli Storage Manager scheduler service for virtual nodes.

Installing the TSM Web client services

This task is not necessary if we do not want to use the Web client. However, if we want to be able to access virtual clients from a Web browser, we must follow the tasks explained in this section.

We need to install Tivoli Storage Manager Client Acceptor and Tivoli Storage Manager Remote Client Agent services on both physical nodes with the same service names and the same options:

1. We need to be sure we are located on the node that hosts all resources, in order to start with the Tivoli Storage Manager Web services installation.
2. We begin the installation of the Tivoli Storage Manager Client Acceptor service for each group on TONGA. This is the node that hosts the resources. We use the **`dsmcutil program`**. This utility is located on the Tivoli Storage Manager client installation path (c:\program files\tivoli\tsm\baclient).
3. In our lab we installed three Client Acceptor services, one for each cluster group, and three Remote Client Agent services (one for each cluster group). When we start the installation the node that hosts the resources is TONGA.
4. We open a MS-DOS Windows command line and change to the Tivoli Storage Manager client installation path. We run the **`dsmcutil tool`** with the appropriate parameters to create the Tivoli Storage Manager Client Acceptor service for the Cluster Group, as shown in Figure 6-57.

Archived

```
C:\WINDOWS\system32\cmd.exe

C:\Program Files\Tivoli\TSM\baclient>dsmcutil inst cad /name:"TSM Client Accepto
r CL_MSCS02_QUORUM" /clientdir:"c:\program files\tivoli\tsm\baclient" /optfile:q
:\tsm\dsm.opt /node:CL_MSCS02_QUORUM /password:itsosj /clustername:CL_MSCS02 /cl
usternode:yes /autostart:no /httpport:1582

TSM Windows NT Client Service Configuration Utility
Command Line Interface - Version 5, Release 3, Level 0.0
(C) Copyright IBM Corporation, 1990, 2004, All Rights Reserved.
Last Updated Dec  8 2004
TSM Api Verison 5.3.0

Command: Install TSM Client Service
Machine: TONGA(Local Machine)


Installing TSM Client Service:

        Machine           : TONGA
        Service Name      : TSM Client Acceptor CL_MSCS02_QUORUM
        Client Directory  : c:\program files\tivoli\tsm\baclient
        Automatic Start   : no
        Logon Account     : LocalSystem

The service was successfully installed.


Creating Registry Keys ...

Updated registry value 'ImagePath' .
Updated registry value 'EventMessageFile' .
Updated registry value 'TypesSupported' .
Updated registry value 'TSM Client Acceptor CL_MSCS02_QUORUM' .
Updated registry value 'ADSMClientKey' .
Updated registry value 'OptionsFile' .
Updated registry value 'EventLogging' .
Updated registry value 'ClientNodeName' .
Updated registry value 'HttpPort' .

Generating registry password ...
Authenticating TSM password for node CL_MSCS02_QUORUM ...

Connecting to TSM Server via client options file 'q:\tsm\dsm.opt' ...

Password authentication successful.

The registry password for TSM node CL_MSCS02_QUORUM has been updated.
```

Figure 6-57 Installing the Client Acceptor service in the Cluster Group

5. After a successful installation of the Client Acceptor for this resource group, we run the **dsmcutil tool** again to create its Remote Client Agent partner service, typing the command:

   ```
   dsmcutil inst remoteagent /name:"TSM Remote Client Agent CL_MSCS02_QUORUM"
   /clientdir:"c:\Program Files\Tivoli\tsm\baclient" /optfile:q:\tsm\dsm.opt
   /node:CL_MSCS02_QUORUM /password:itsosj /clusternode:yes
   /clustername:CL_MSCS02 /startnow:no /partnername:"TSM Client Acceptor
   CL_MSCS02_QUORUM".
   ```

6. If the installation is successful we receive the following sequence of messages (Figure 6-58).

Archived

```
C:\WINDOWS\system32\cmd.exe
TSM Windows NT Client Service Configuration Utility
Command Line Interface - Version 5, Release 3, Level 0.0
(C) Copyright IBM Corporation, 1990, 2004, All Rights Reserved.
Last Updated Dec  8 2004
TSM Api Verison 5.3.0

Command: Install TSM Client Service
Machine: TONGA(Local Machine)


Locating partner client acceptor service 'TSM Client Acceptor CL_MSCS02_QUORUM'
...



Installing TSM Client Service:

        Machine          : TONGA
        Service Name     : TSM Remote Client Agent CL_MSCS02_QUORUM
        Client Directory : c:\program files\tivoli\tsm\baclient
        Automatic Start  : no
        Logon Account    : LocalSystem

The service was successfully installed.


Creating Registry Keys ...

Updated registry value 'ImagePath' .
Updated registry value 'EventMessageFile' .
Updated registry value 'TypesSupported' .
Updated registry value 'TSM Remote Client Agent CL_MSCS02_QUORUM' .
Updated registry value 'ADSMClientKey' .
Updated registry value 'OptionsFile' .
Updated registry value 'EventLogging' .
Updated registry value 'ClientNodeName' .
Updated registry value 'PartnerName' .
Updated Partner Client Acceptor registry value 'PartnerName' .

Generating registry password ...
Authenticating TSM password for node CL_MSCS02_QUORUM ...

Connecting to TSM Server via client options file 'q:\tsm\dsm.opt' ...

Password authentication successful.

The registry password for TSM node CL_MSCS02_QUORUM has been updated.
```

Figure 6-58 Successful installation, Tivoli Storage Manager Remote Client Agent

7. We follow the same process to install the services for the TSM Admin Center cluster group. We use the following commands:

   dsmcutil inst cad /name:"TSM **Client Acceptor** CL_MSCS02_SA" /clientdir:"c:\Program Files\Tivoli\tsm\baclient" /optfile:j:\tsm\dsm.opt /node:CL_MSCS02_SA /password:itsosj /clusternode:yes /clustername:CL_MSCS02 /autostart:no /httpport:1584

   dsmcutil inst remoteagent /name:**"TSM Remote Client Agent CL_MSCS02_SA"** /clientdir:"c:\Program Files\Tivoli\tsm\baclient" /optfile:j:\tsm\dsm.opt /node:CL_MSCS02_SA /password:itsosj /clusternode:yes /clustername:CL_MSCS02 /startnow:no /partnername:"TSM Client Acceptor CL_MSCS02_SA"

8. And finally we use the same process to install the services for the TSM Group, with the following commands:

   dsmcutil inst cad /name:"TSM Client Acceptor CL_MSCS02_TSM" /clientdir:"c:\Program Files\Tivoli\tsm\baclient" /optfile:g:\tsm\dsm.opt /node:CL_MSCS02_TSM /password:itsosj /clusternode:yes /clustername:CL_MSCS02 /autostart:no /httpport:1583

Archived

```
dsmcutil inst remoteagent /name:"TSM Remote Client Agent CL_MSCS02_TSM"
/clientdir:"c:\Program Files\Tivoli\tsm\baclient" /optfile:g:\tsm\dsm.opt
/node:CL_MSCS02_TSM /password:itsosj /clusternode:yes
/clustername:CL_MSCS02 /startnow:no /partnername:"TSM Client Acceptor
CL_MSCS02_TSM"
```

Important: The client acceptor and remote client agent services must be installed with the same name on each physical node on the MSCS, otherwise failover will not work. Also do not forget the options clusternode yes and clustername as well as to specify the correct dsm.opt path file name in the optfile parameter of the **dsmcutil** command.

9. We move the resources to the second node (SENEGAL) and repeat steps 1-8 with the same options for each resource group.

So far the Tivoli Storage Manager Web client services are installed on both nodes of the cluster with exactly the same names for each resource group. The last task consists of the definition for new resource on each cluster group. But first we go to the Windows Service menu and stop all the Web client services on SENEGAL.

Creating a generic resource for TSM Client Acceptor service

For a correct configuration of the Tivoli Storage Manager Web client we define, for each cluster group, a new generic service resource. This resource will be related to the Client Acceptor service name created for this group.

Important: Before continuing, we make sure to stop all services created in "Installing the TSM Web client services" on page 306 on all nodes. We also make sure all resources are on one of the nodes.

Archived

Here are the steps we follow:

1. We open the Cluster Administrator menu on the node that hosts all resources and select the first group (Cluster Group). We right-click the name and select **New →Resource** as shown in Figure 6-59.

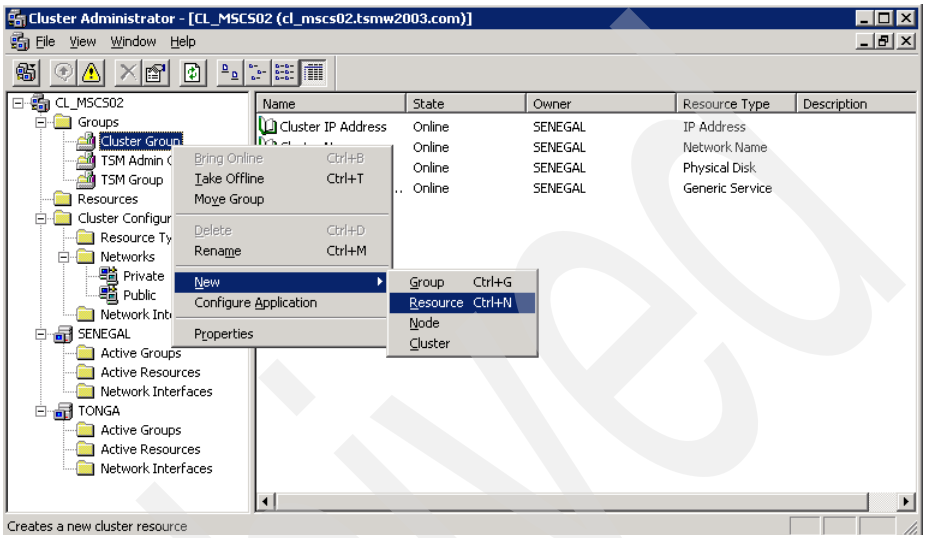

Figure 6-59 New resource for Tivoli Storage Manager Client Acceptor service

2. We type a Name for the resource (we recommend to use the same name as the scheduler service) and select Generic Service as resource type. We click **Next** as shown in Figure 6-60.

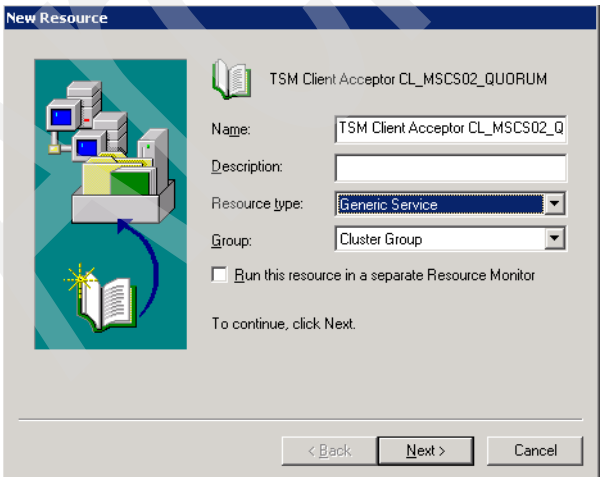

Figure 6-60 Definition of TSM Client Acceptor generic service resource

3. We leave both nodes as possible owners for the resource and click **Next** (Figure 6-61).

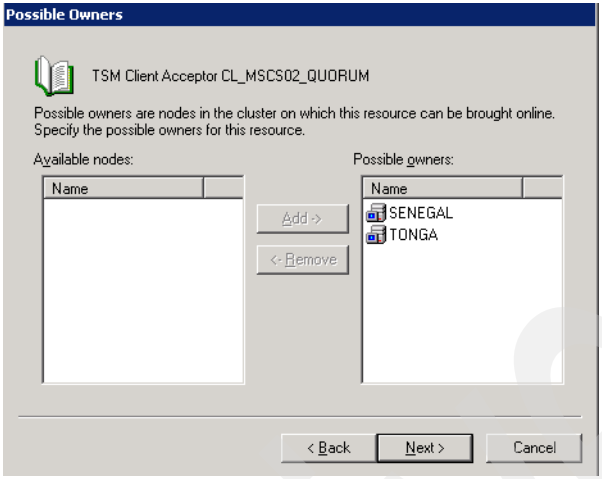

Figure 6-61 Possible owners of the TSM Client Acceptor generic service

4. We **Add** the disk resources (in this case q:) on Dependencies in Figure 6-62. We click **Next**.

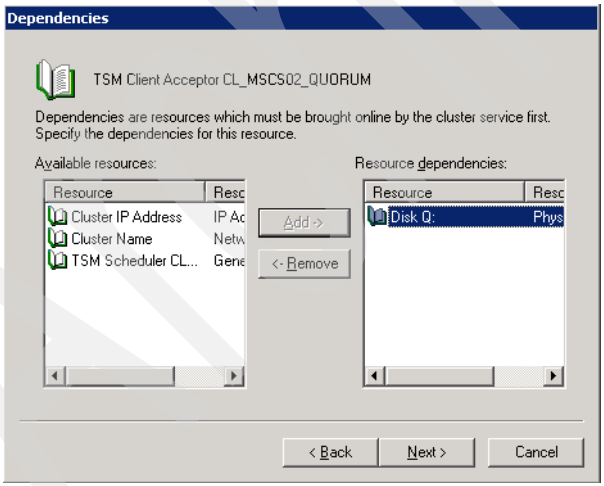

Figure 6-62 Dependencies for TSM Client Acceptor generic service

Archived

5. On the next menu we type a Service name. This must match the name used while installing the Client Acceptor service on both nodes. We click **Next** (Figure 6-63).

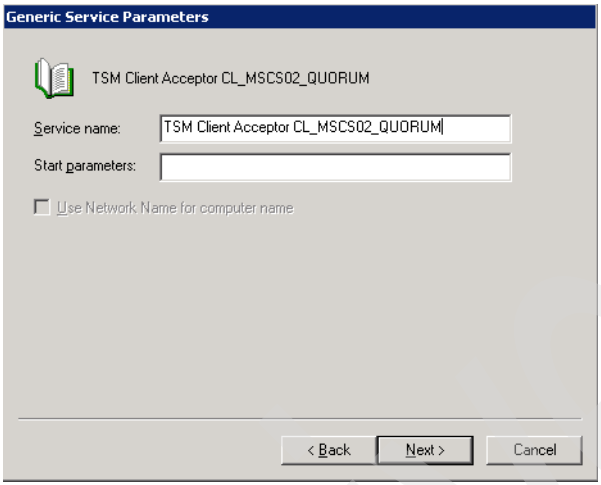

Figure 6-63 TSM Client Acceptor generic service parameters

6. Next we type the Registry Key where Windows 2003 will save the generated password for the client. It is the same path we typed in Figure 6-52 on page 303. We click **OK**.
7. If the resource creation is successful we receive an information menu as was shown in Figure 6-53 on page 303. We click **OK**.

Archived

8. Now, as shown in Figure 6-64 below, the Cluster Group is offline because the new resource is also offline. We bring it online.

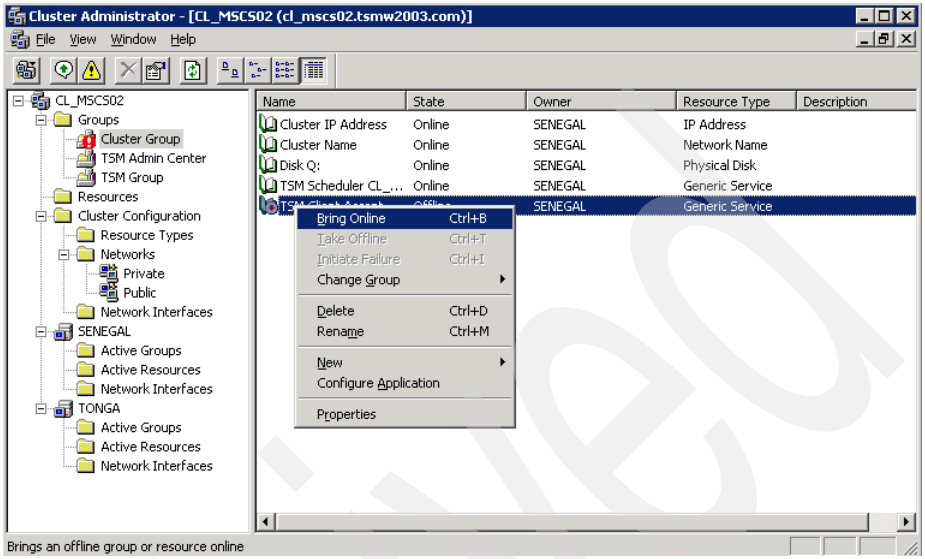

Figure 6-64 Bringing online the TSM Client Acceptor generic service

9. The Cluster Administrator menu displays next as shown in Figure 6-65.

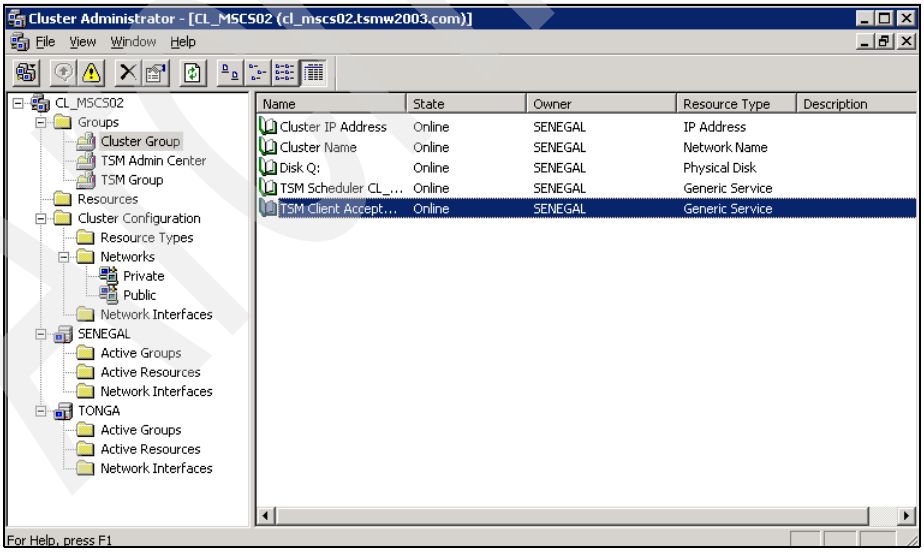

Figure 6-65 TSM Client Acceptor generic service online

10. If we go to the Windows service menu, Tivoli Storage Manager Client Acceptor service is started on SENEGAL, the node which now hosts this resource group:

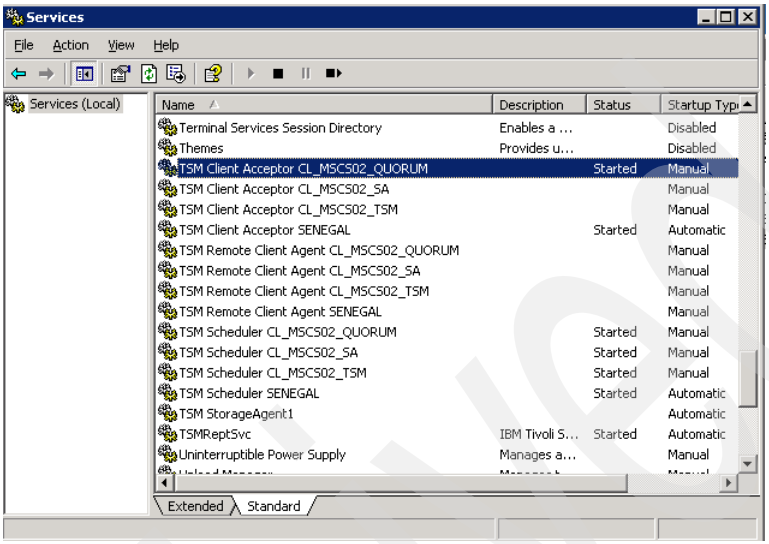

Figure 6-66 Windows service menu

Important: all Tivoli Storage Manager client services used by virtual nodes of the cluster must figure as *Manual* on the Startup Type column in Figure 6-66. They may only be started on the node that hosts the resource at that time.

11. We follow the same tasks to create the Tivoli Storage Manager Client Acceptor service resource for TSM Admin Center and TSM Group cluster groups. The resource names are:
    - TSM Client Acceptor CL_MSCS02_SA: for TSM Admin Center resource group
    - TSM Client Acceptor CL_MSCS02_TSM: for TSM Group resource group.
12. We move the resources to check that Tivoli Storage Manager Client Acceptor services successfully start on the second node, TONGA, while they are stopped on the first node.

Filespace names for local and virtual nodes

If the configuration of Tivoli Storage Manager client in our MSCS is correct, when the client backs up files against our Tivoli Storage Manager server, the filespace names for local (physical) nodes and virtual (shared) nodes are different. We show this in Figure 6-67.

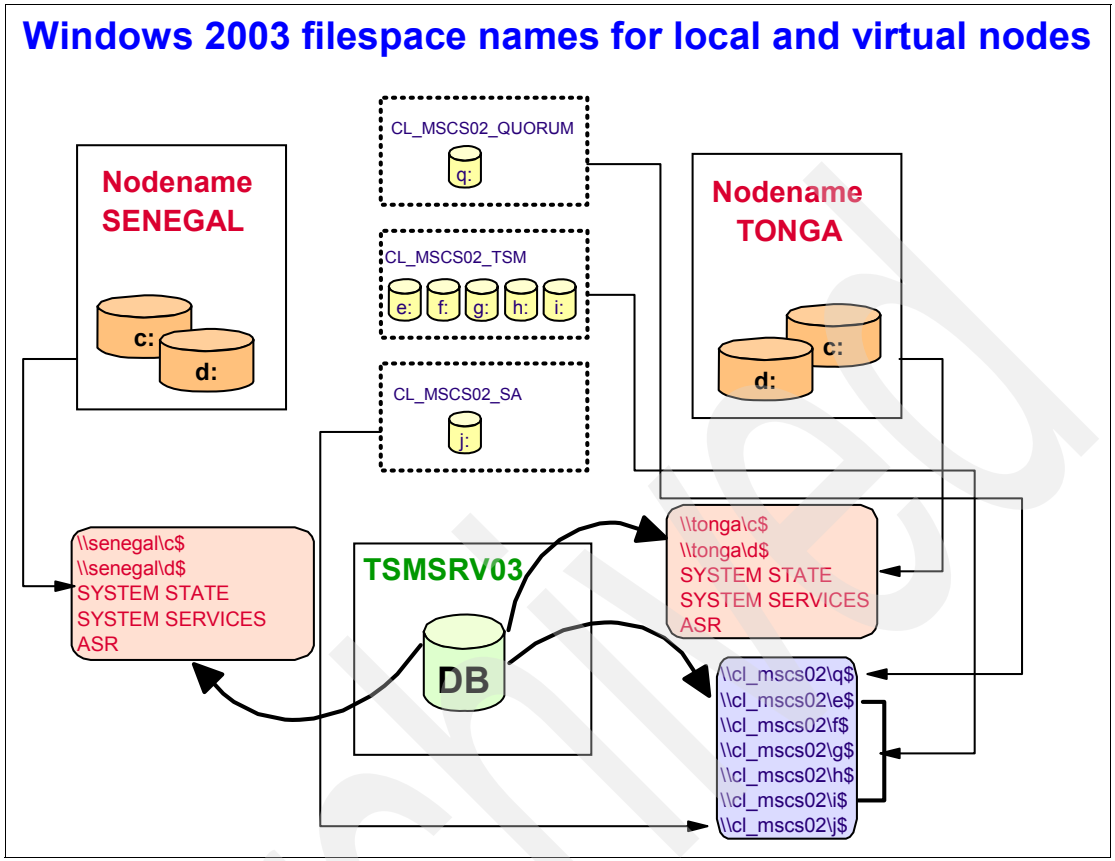

Figure 6-67 Windows 2003 filespace names for local and virtual nodes

When the local nodes back up files, their filespace names start with the physical nodename. However, when the virtual nodes back up files, their filespace names start with the cluster name, in our case, CL_MSCS02.

6.5.3 Testing Tivoli Storage Manager client on Windows 2003

In order to check the high availability of Tivoli Storage Manager client on our lab environment, we must do some testing.

Our objective with these tests is to know how Tivoli Storage Manager can respond, on a clustered environment, after certain kinds of failures that affect the shared resources.

For the purpose of this section, we will use a Tivoli Storage Manager server installed on an AIX machine: TSMSRV03. For details of this server, refer to the AIX chapters in this book.

Remember, our Tivoli Storage Manager clients are:

- CL_MSCS02_QUORUM
- CL_MSCS02_TSM
- CL_MSCS02_SA

Testing client incremental backup

Our first test consists of an incremental backup started from the client.

Objective

The objective of this test is to show what happens when a client incremental backup is started for a virtual client in the cluster, and the node that hosts the resources at that moment suddenly fails.

Activities

To do this test, we perform these tasks:

1. We open the Cluster Administrator to check which node hosts the Tivoli Storage Manager client resource as shown in Figure 6-68.

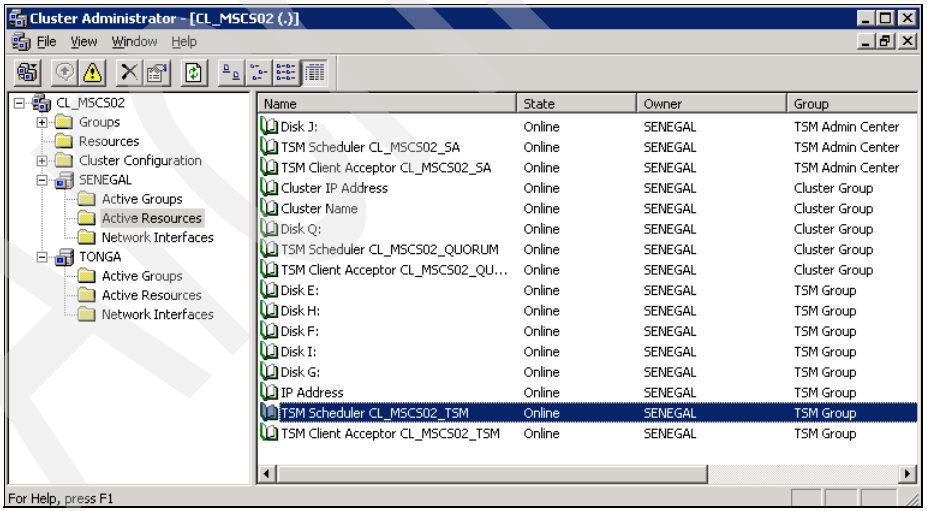

Figure 6-68 Resources hosted by SENEGAL in the Cluster Administrator

As we can see in the figure, SENEGAL hosts all the resources at this moment.

2. We schedule a client incremental backup operation using the Tivoli Storage Manager server scheduler and we associate the schedule to CL_MSCS02_TSM nodename.
3. A client session for CL_MSCS02_TSM nodename starts on the server as shown in Figure 6-69.

```
IBM Tivoli Storage Manager
tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name     Status
-------------------- -------------------- ------------- ------------- ---------
02/03/2005 17:44:00                       TEST_SCHED    CL_HACMP03_C- Pending
                                                         LIENT
02/03/2005 17:44:00  02/03/2005 17:55:03  TEST_SCHED    KANANGA       Completed
02/03/2005 17:44:00  02/03/2005 17:57:40  TEST_SCHED    AZOV          Completed
02/03/2005 17:59:35  02/03/2005 18:00:58  INCR_BCK      CL_MSCS02_TSM Started

tsm: TSMSRV03>q se

  Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type  -------- --------------------
------ ------ ------ ------ ------- ------- ----- -------- --------------------
     1 Tcp/Ip Run      0 S  192.9 K     139 Admin AIX      ADMIN
   368 Tcp/Ip Run      0 S   88.4 K   1.5 K Admin WinNT    ADMIN
   420 Tcp/Ip IdleW   25 S    1.2 K   1.4 K Node  WinNT    CL_MSCS02_TSM
   421 Tcp/Ip RecvW    0 S      484 239.9 M Node  WinNT    CL_MSCS02_TSM
```

Figure 6-69 Scheduled incremental backup started for CL_MSCS02_TSM

4. The client starts sending files to the server as we can see on the schedule log file shown in Figure 6-70.

```
dsmsched.log - Notepad
File  Edit  Format  View  Help
02/03/2005 18:00:57 --- SCHEDULEREC QUERY BEGIN
02/03/2005 18:00:57 --- SCHEDULEREC QUERY END
02/03/2005 18:00:57 Next operation scheduled:
02/03/2005 18:00:57 ------------------------------------------------------------
02/03/2005 18:00:57 Schedule Name:         INCR_BCK
02/03/2005 18:00:57 Action:                Incremental
02/03/2005 18:00:57 Objects:
02/03/2005 18:00:57 Options:
02/03/2005 18:00:57 Server Window Start:   17:59:35 on 02/03/2005
02/03/2005 18:00:57 ------------------------------------------------------------
02/03/2005 18:00:57 Command will be executed in 1 minute.
02/03/2005 18:01:57
Executing scheduled command now.
02/03/2005 18:01:57 Node Name: CL_MSCS02_TSM
02/03/2005 18:01:57 Session established with server TSMSRV03: AIX-RS/6000
02/03/2005 18:01:57   Server Version 5, Release 3, Level 0.0
02/03/2005 18:01:57   Server date/time: 02/03/2005 18:00:58  Last access: 02/03/2005 17:59:5

02/03/2005 18:01:57 --- SCHEDULEREC OBJECT BEGIN INCR_BCK 02/03/2005 17:59:35
02/03/2005 18:01:57 Incremental backup of volume '\\cl_mscs02\e$'
02/03/2005 18:01:57 Incremental backup of volume '\\cl_mscs02\f$'
02/03/2005 18:01:57 Incremental backup of volume '\\cl_mscs02\g$'
02/03/2005 18:01:57 Incremental backup of volume '\\cl_mscs02\h$'
02/03/2005 18:01:57 Incremental backup of volume '\\cl_mscs02\i$'
02/03/2005 18:01:57 Directory-->                  0 \\cl_mscs02\e$\ [Sent]
02/03/2005 18:01:59 Directory-->                  0 \\cl_mscs02\e$\Recycled [Sent]
02/03/2005 18:01:59 Directory-->                  0 \\cl_mscs02\e$\RECYCLER [Sent]
02/03/2005 18:01:59 Directory-->                  0 \\cl_mscs02\e$\System Volume Informatio
02/03/2005 18:01:59 Retry # 1  Directory-->                 0 \\cl_mscs02\e$\ [Sent]
02/03/2005 18:01:59 Retry # 1  Directory-->                 0 \\cl_mscs02\e$\Recycled [Sen
```

Figure 6-70 Schedule log file: incremental backup starting for CL_MSCS02_TSM

Note: Observe that, in Figure 6-70, the filespace name used by Tivoli Storage Manager to store the files in the server (`\\cl_mscs02\e$`). If the client is correctly configured to work on MSCS, the filespace name always starts with the cluster name. It does not use the local name of the physical node which hosts the resource at the time of backup.

5. While the client continues sending files to the server, we force SENEGAL to fail. The following sequence takes place:

   a. The client loses its connection with the server temporarily, and the session is terminated as we can see on the Tivoli Storage Manager server activity log shown in Figure 6-71.

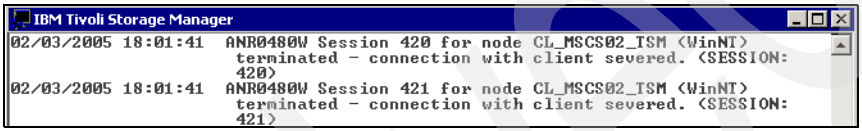

Figure 6-71 CL_MSCS02_TSM loss its connection with the server

   b. In the Cluster Administrator, SENEGAL is not in the cluster and TONGA begins to take the failover for the resources.

   c. In the schedule log file for CL_MSCS02_TSM, there is an interruption message (Figure 6-72).

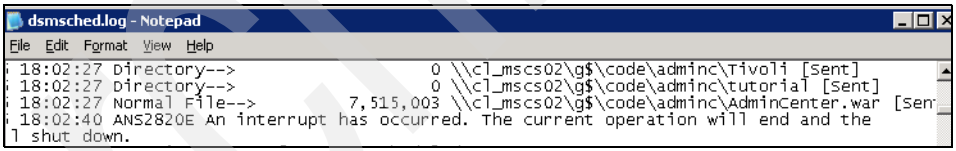

Figure 6-72 The schedule log file shows an interruption of the session

   d. After a short period of time the resources are online on TONGA.

   e. When the TSM Scheduler CL_MSCS02_TSM resource is online (hosted by TONGA), the client restarts the backup as we show on the schedule log file in Figure 6-73.

```
dsmsched.log - Notepad
File  Edit  Format  View  Help
02/03/2005 18:04:10 --- SCHEDULEREC QUERY BEGIN
02/03/2005 18:04:10 --- SCHEDULEREC QUERY END
02/03/2005 18:04:10 Next operation scheduled:
02/03/2005 18:04:10 ------------------------------------------------------------
02/03/2005 18:04:10 Schedule Name:         INCR_BCK
02/03/2005 18:04:10 Action:                Incremental
02/03/2005 18:04:10 Objects:
02/03/2005 18:04:10 Options:
02/03/2005 18:04:10 Server Window Start:   17:59:35 on 02/03/2005
02/03/2005 18:04:10 ------------------------------------------------------------
02/03/2005 18:04:10 Command will be executed in 2 minutes.
02/03/2005 18:06:10
Executing scheduled command now.
02/03/2005 18:06:10 Node Name: CL_MSCS02_TSM
02/03/2005 18:06:10 Session established with server TSMSRV03: AIX-RS/6000
02/03/2005 18:06:10   Server Version 5, Release 3, Level 0.0
02/03/2005 18:06:10   Server date/time: 02/03/2005 18:05:11  Last access: 02/03/2005 18:03:1

02/03/2005 18:06:10 --- SCHEDULEREC OBJECT BEGIN INCR_BCK 02/03/2005 17:59:35
02/03/2005 18:06:10 Incremental backup of volume '\\cl_mscs02\e$'
02/03/2005 18:06:10 Incremental backup of volume '\\cl_mscs02\f$'
02/03/2005 18:06:10 Incremental backup of volume '\\cl_mscs02\g$'
02/03/2005 18:06:10 Incremental backup of volume '\\cl_mscs02\h$'
02/03/2005 18:06:10 Incremental backup of volume '\\cl_mscs02\i$'
02/03/2005 18:06:12 Successful incremental backup of '\\cl_mscs02\e$'

02/03/2005 18:06:12 Successful incremental backup of '\\cl_mscs02\f$'

02/03/2005 18:06:12 Normal File-->           20,955 \\cl_mscs02\g$\code\adminc\dsminstall.j
02/03/2005 18:06:14 Normal File-->       23,846,959 \\cl_mscs02\g$\code\adminc\help.jar [se
02/03/2005 18:06:14 Normal File-->        6,484,490 \\cl_mscs02\g$\code\adminc\ISCAction.ja
02/03/2005 18:06:14 Normal File-->           79,853 \\cl_mscs02\g$\code\adminc\LICENSE.TXT
```

Figure 6-73 Schedule log shows how the incremental backup restarts

In Figure 6-73, we see how Tivoli Storage Manager client scheduler queries the server for a scheduled command, and since the schedule is still within the startup window, the incremental backup starts sending files for the g: drive. The files belonging to e: and f: shared disks are not sent again because the client already backed up them before the interruption.

f. In the Tivoli Storage Manager server activity log in Figure 6-74 we can see how the resource for CL_MSCS02_TSM moves from SENEGAL to TONGA and a new session is started again for this client (Figure 6-74).

```
IBM Tivoli Storage Manager
02/03/2005 18:03:05  ANR0406I Session 432 started for node CL_MSCS02_TSM
                      (WinNT) (Tcp/Ip 9.1.39.168(1435)). (SESSION: 432)
02/03/2005 18:03:05  ANR1639I Attributes changed for node CL_MSCS02_TSM: TCP
                      Name from SENEGAL to TONGA, TCP Address from 9.1.39.166
                      to 9.1.39.168, GUID from 46.6e.f4.d0.75.3a.11.d9.bc.7b.0-
                      0.0d.60.49.41.dd to 18.72.62.50.75.3b.11.d9.8f.b7.00.a0.-
                      cc.d2.58.ac. (SESSION: 432)
```

Figure 6-74 Attributes changed for node CL_MSCS02_TSM

Archived

g. Also, in the Tivoli Storage Manager server event log, we see the scheduled event *restarted* as shown in Figure 6-75.

```
IBM Tivoli Storage Manager
02/03/2005 17:44:00  02/03/2005 17:55:03  TEST_SCHED    KANANGA        Completed
02/03/2005 17:44:00  02/03/2005 17:57:40  TEST_SCHED    AZOU           Completed
02/03/2005 17:59:35  02/03/2005 18:00:58  INCR_BCK      CL_MSCS02_TSM  Started

tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name      Status
-------------------- -------------------- ------------- ------------   ---------
02/03/2005 17:44:00                       TEST_SCHED    CL_HACMP03_C-  Pending
                                                         LIENT
02/03/2005 17:44:00  02/03/2005 17:55:03  TEST_SCHED    KANANGA        Completed
02/03/2005 17:44:00  02/03/2005 17:57:40  TEST_SCHED    AZOU           Completed
02/03/2005 17:59:35  02/03/2005 18:00:58  INCR_BCK      CL_MSCS02_TSM  Restarted
```

Figure 6-75 Event log shows the incremental backup schedule as restarted

6. The incremental backup ends successfully as we see on the activity log in Figure 6-76.

```
IBM Tivoli Storage Manager
02/03/2005 18:06:38  ANR2507I Schedule INCR_BCK for domain STANDARD started at
                      02/03/05 17:59:35 for node CL_MSCS02_TSM completed
                      successfully at 02/03/05 18:06:38. (SESSION: 440)
02/03/2005 18:06:38  ANR0403I Session 440 ended for node CL_MSCS02_TSM (WinNT).
                      (SESSION: 440)
02/03/2005 18:06:38  ANR0406I Session 444 started for node CL_MSCS02_TSM
                      (WinNT) (Tcp/Ip 9.1.39.168(1630)). (SESSION: 444)
02/03/2005 18:06:38  ANR0403I Session 444 ended for node CL_MSCS02_TSM (WinNT).
                      (SESSION: 444)
```

Figure 6-76 Schedule INCR_BCK completed successfully

7. In the Tivoli Storage Manager server event log, the schedule is *completed* (Figure 6-77).

```
IBM Tivoli Storage Manager
02/03/2005 17:44:00  02/03/2005 17:55:03  TEST_SCHED    KANANGA        Completed
02/03/2005 17:44:00  02/03/2005 17:57:40  TEST_SCHED    AZOU           Completed
02/03/2005 17:59:35  02/03/2005 18:00:58  INCR_BCK      CL_MSCS02_TSM  Restarted

tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name      Status
-------------------- -------------------- ------------- ------------   ---------
02/03/2005 17:44:00                       TEST_SCHED    CL_HACMP03_C-  Pending
                                                         LIENT
02/03/2005 17:44:00  02/03/2005 17:55:03  TEST_SCHED    KANANGA        Completed
02/03/2005 17:44:00  02/03/2005 17:57:40  TEST_SCHED    AZOU           Completed
02/03/2005 17:59:35  02/03/2005 18:00:58  INCR_BCK      CL_MSCS02_TSM  Restarted

tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name      Status
-------------------- -------------------- ------------- ------------   ---------
02/03/2005 17:44:00                       TEST_SCHED    CL_HACMP03_C-  Pending
                                                         LIENT
02/03/2005 17:44:00  02/03/2005 17:55:03  TEST_SCHED    KANANGA        Completed
02/03/2005 17:44:00  02/03/2005 17:57:40  TEST_SCHED    AZOU           Completed
02/03/2005 17:59:35  02/03/2005 18:00:58  INCR_BCK      CL_MSCS02_TSM  Completed

tsm: TSMSRV03>_
```

Figure 6-77 Schedule completed on the event log

Checking that all files were correctly backed up

In this section we want to show a way of checking that the incremental backup did not miss any file while the failover process took place.

With this in mind, we perform these tasks:

1. In Figure 6-72 on page 318, the last file reported as sent in the schedule log file is \\cl_mscs02\g$\code\adminc\AdminCenter.war. And the first file sent after the failover is dsminstall.jar, also on the same path.
2. We open the explorer and go to this path (Figure 6-78).

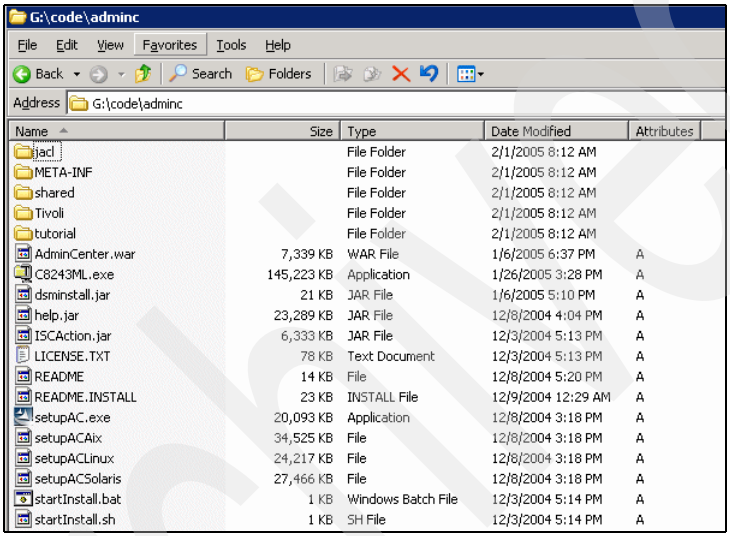

Name	Size	Type	Date Modified	Attributes
jacl		File Folder	2/1/2005 8:12 AM	
META-INF		File Folder	2/1/2005 8:12 AM	
shared		File Folder	2/1/2005 8:12 AM	
Tivoli		File Folder	2/1/2005 8:12 AM	
tutorial		File Folder	2/1/2005 8:12 AM	
AdminCenter.war	7,339 KB	WAR File	1/6/2005 6:37 PM	A
C8243ML.exe	145,223 KB	Application	1/26/2005 3:28 PM	A
dsminstall.jar	21 KB	JAR File	1/6/2005 5:10 PM	A
help.jar	23,289 KB	JAR File	12/8/2004 4:04 PM	A
ISCAction.jar	6,333 KB	JAR File	12/3/2004 5:13 PM	A
LICENSE.TXT	78 KB	Text Document	12/3/2004 5:13 PM	A
README	14 KB	File	12/8/2004 5:20 PM	A
README.INSTALL	23 KB	INSTALL File	12/9/2004 12:29 AM	A
setupAC.exe	20,093 KB	Application	12/8/2004 3:18 PM	A
setupACAix	34,525 KB	File	12/8/2004 3:18 PM	A
setupACLinux	24,217 KB	File	12/8/2004 3:18 PM	A
setupACSolaris	27,466 KB	File	12/8/2004 3:18 PM	A
startInstall.bat	1 KB	Windows Batch File	12/3/2004 5:14 PM	A
startInstall.sh	1 KB	SH File	12/3/2004 5:14 PM	A

Figure 6-78 Windows explorer

3. If we have a look at Figure 6-78 between Admincenter.war and dsminstall.jar, there is one file not reported as backed up in the schedule log file.
4. We open a Tivoli Storage Manager GUI session to check, on the tree view of the *Restore* menu, whether these files were backed up (Figure 6-79).

Archived

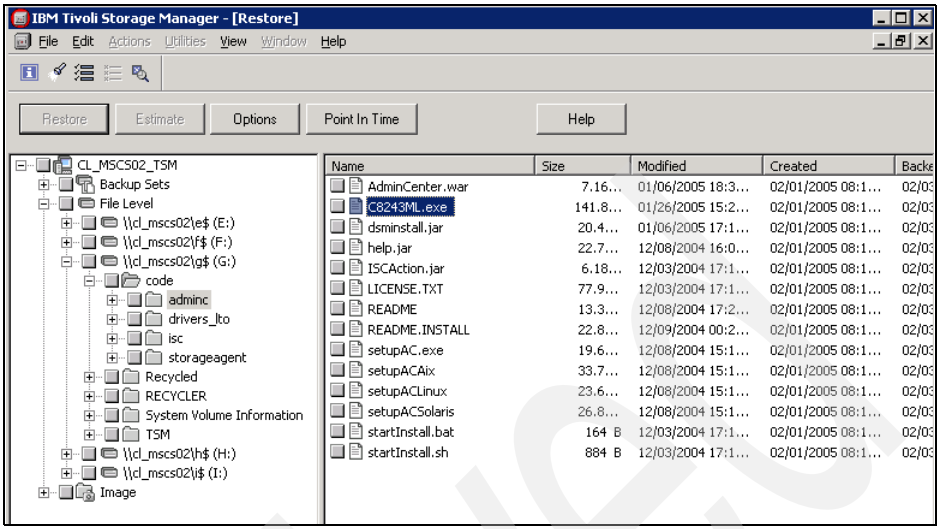

Figure 6-79 Checking backed up files using the TSM GUI

5. We see in Figure 6-79 that the client backed up the files correctly, even when they were not reported in the schedule log file. Since the session was lost, the client was not able of writing into the shared disk where the schedule log file is located.

Results summary

The test results show that, after a failure on the node that hosts the Tivoli Storage Manager scheduler service resource, a scheduled incremental backup started on one node is restarted and successfully completed on the other node that takes the failover.

This is true if the startup window used to define the schedule is not elapsed when the scheduler services restarts on the second node.

Testing client restore

Our second test consists of a scheduled restore of certain files under a directory.

Objective

The objective of this test is to show what happens when a client restore is started for a virtual client in the cluster, and the node that hosts the resources at that moment suddenly fails.

Activities

To do this test, we perform these tasks:

1. We open the Cluster Administrator to check which node hosts the Tivoli Storage Manager client resource: TONGA.
2. We schedule a client restore operation using the Tivoli Storage Manager server scheduler and we associate the schedule to CL_MSCS02_TSM nodename.
3. A client session for CL_MSCS02_TSM nodename starts on the server as shown in Figure 6-80.

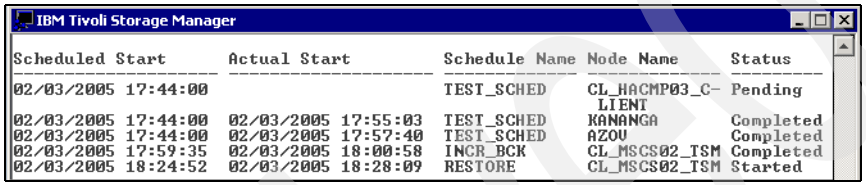

IBM Tivoli Storage Manager

Scheduled Start	Actual Start	Schedule Name	Node Name	Status
02/03/2005 17:44:00		TEST_SCHED	CL_HACMP03_C-LIENT	Pending
02/03/2005 17:44:00	02/03/2005 17:55:03	TEST_SCHED	KANANGA	Completed
02/03/2005 17:44:00	02/03/2005 17:57:40	TEST_SCHED	AZOU	Completed
02/03/2005 17:59:35	02/03/2005 18:00:58	INCR_BCK	CL_MSCS02_TSM	Completed
02/03/2005 18:24:52	02/03/2005 18:28:09	RESTORE	CL_MSCS02_TSM	Started

Figure 6-80 Scheduled restore started for CL_MSCS02_TSM

4. The client starts restoring files as we see on the schedule log file in Figure 6-81.

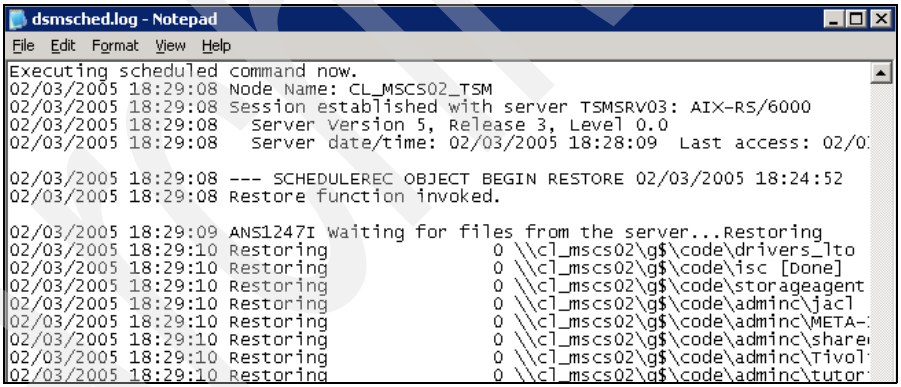

dsmsched.log - Notepad

File Edit Format View Help

```
Executing scheduled command now.
02/03/2005 18:29:08 Node Name: CL_MSCS02_TSM
02/03/2005 18:29:08 Session established with server TSMSRV03: AIX-RS/6000
02/03/2005 18:29:08   Server Version 5, Release 3, Level 0.0
02/03/2005 18:29:08   Server date/time: 02/03/2005 18:28:09  Last access: 02/0

02/03/2005 18:29:08 --- SCHEDULEREC OBJECT BEGIN RESTORE 02/03/2005 18:24:52
02/03/2005 18:29:08 Restore function invoked.

02/03/2005 18:29:09 ANS1247I Waiting for files from the server...Restoring
02/03/2005 18:29:10 Restoring              0 \\cl_mscs02\g$\code\drivers_lto
02/03/2005 18:29:10 Restoring              0 \\cl_mscs02\g$\code\isc [Done]
02/03/2005 18:29:10 Restoring              0 \\cl_mscs02\g$\code\storageagent
02/03/2005 18:29:10 Restoring              0 \\cl_mscs02\g$\code\adminc\jacl
02/03/2005 18:29:10 Restoring              0 \\cl_mscs02\g$\code\adminc\META-
02/03/2005 18:29:10 Restoring              0 \\cl_mscs02\g$\code\adminc\share
02/03/2005 18:29:10 Restoring              0 \\cl_mscs02\g$\code\adminc\Tivol
02/03/2005 18:29:10 Restoring              0 \\cl_mscs02\g$\code\adminc\tutor
```

Figure 6-81 Restore starts in the schedule log file for CL_MSCS02_TSM

5. While the client is restoring the files, we force TONGA to fail. The following sequence takes place:
   a. The client loses temporarily its connection with the server, and the session is terminated as we can see on the Tivoli Storage Manager server activity log shown in Figure 6-82.

```
IBM Tivoli Storage Manager
02/03/2005 18:28:34  ANR8216W Error sending data on socket 11.  Reason 32.
                      (SESSION: 453)
02/03/2005 18:28:34  ANR0480W Session 453 for node CL_MSCS02_TSM (WinNT)
                      terminated - connection with client severed. (SESSION:
                      453)
02/03/2005 18:28:34  ANR8216W Error sending data on socket 11.  Reason 32.
                      (SESSION: 453)
```

Figure 6-82 Restore session is lost for CL_MSCS02_TSM

b. In the Cluster Administrator, TONGA is not in the cluster and SENEGAL begins to bring the resources online.

c. In the schedule log file for CL_MSCS02_TSM we also see a message informing us about a connection lost (Figure 6-83).

```
dsmsched.log - Notepad
File  Edit  Format  View  Help
02/03/2005 18:29:29 Restoring     181,664,860 \\cl_mscs02\g$\code\TSMSRV5300_WIN.exe [Done]
02/03/2005 18:29:32 Restoring      29,205,340 \\cl_mscs02\g$\code\TSMSTA5300_WIN.exe [Done]
02/03/2005 18:29:33 ANS2820E An interrupt has occurred. The current operation will end and
client will shut down.
```

Figure 6-83 Schedule log file shows an interruption for the restore operation

d. After some minutes, the resources are online on SENEGAL. The Tivoli Storage Manager server activity log shows the resource for CL_MSCS02_TSM moving from TONGA to SENEGAL (Figure 6-84).

```
IBM Tivoli Storage Manager
02/03/2005 18:29:46  ANR0406I Session 464 started for node CL_MSCS02_TSM
                      (WinNT) (Tcp/Ip 9.1.39.166(1905)). (SESSION: 464)
02/03/2005 18:29:46  ANR1639I Attributes changed for node CL_MSCS02_TSM: TCP
                      Name from TONGA to SENEGAL, TCP Address from 9.1.39.168
                      to 9.1.39.166, GUID from 18.72.62.50.75.3b.11.d9.8f.b7.0-
                      0.a0.cc.d2.58.ac to 46.6e.f4.d0.75.3a.11.d9.bc.7b.00.0d.-
                      60.49.41.dd. (SESSION: 464)
```

Figure 6-84 Attributes changed from node CL_MSCS02_TSM to SENEGAL

e. When the Tivoli Storage Manager scheduler service resource is again online on SENEGAL and queries the server for a schedule, if the startup window for the scheduled operation is not elapsed, the restore process restarts from the beginning, as we can see on the schedule log file in Figure 6-85.

Archived

```
dsmsched.log - Notepad
File  Edit  Format  View  Help
02/03/2005 18:30:50 --- SCHEDULEREC QUERY BEGIN
02/03/2005 18:30:50 --- SCHEDULEREC QUERY END
02/03/2005 18:30:50 Next operation scheduled:
02/03/2005 18:30:50 ------------------------------------------------------------
02/03/2005 18:30:50 Schedule Name:         RESTORE
02/03/2005 18:30:50 Action:                Restore
02/03/2005 18:30:50 objects:               g:\code\*
02/03/2005 18:30:50 options:               -subdir=yes -replace=yes
02/03/2005 18:30:50 Server Window Start:   18:24:52 on 02/03/2005
02/03/2005 18:30:50 ------------------------------------------------------------
02/03/2005 18:30:50
Executing scheduled command now.
02/03/2005 18:30:50 --- SCHEDULEREC OBJECT BEGIN RESTORE 02/03/2005 18:24:52
02/03/2005 18:30:50 Restore function invoked.

02/03/2005 18:30:51 ANS1247I Waiting for files from the server...Restoring              0
02/03/2005 18:30:52 Restoring              0 \\cl_mscs02\g$\code\drivers_lto [Done]
02/03/2005 18:30:52 Restoring              0 \\cl_mscs02\g$\code\isc [Done]
02/03/2005 18:30:52 Restoring              0 \\cl_mscs02\g$\code\storageagent [Done]
02/03/2005 18:30:52 Restoring              0 \\cl_mscs02\g$\code\adminc\jacl [Done]
02/03/2005 18:30:52 Restoring              0 \\cl_mscs02\g$\code\adminc\META-INF [Done]
02/03/2005 18:30:52 Restoring              0 \\cl_mscs02\g$\code\adminc\shared [Done]
02/03/2005 18:30:52 Restoring              0 \\cl_mscs02\g$\code\adminc\Tivoli [Done]
```

Figure 6-85 Restore session starts from the beginning in the schedule log file

f. And the event log of Tivoli Storage Manager server shows the schedule as restarted (Figure 6-86).

```
IBM Tivoli Storage Manager
02/03/2005 18:24:52  02/03/2005 18:28:09  RESTORE       CL_MSCS02_TSM Started

tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name     Status
-------------------- -------------------- ------------- ------------- ---------
02/03/2005 17:44:00                       TEST_SCHED    CL_HACMP03_C- Pending
                                                         LIENT
02/03/2005 17:44:00  02/03/2005 17:55:03  TEST_SCHED    KANANGA       Completed
02/03/2005 17:44:00  02/03/2005 17:57:40  TEST_SCHED    AZOV          Completed
02/03/2005 17:59:35  02/03/2005 18:00:58  INCR_BCK      CL_MSCS02_TSM Completed
02/03/2005 18:24:52  02/03/2005 18:28:09  RESTORE       CL_MSCS02_TSM Started

tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name     Status
-------------------- -------------------- ------------- ------------- ---------
02/03/2005 17:44:00                       TEST_SCHED    CL_HACMP03_C- Pending
                                                         LIENT
02/03/2005 17:44:00  02/03/2005 17:55:03  TEST_SCHED    KANANGA       Completed
02/03/2005 17:44:00  02/03/2005 17:57:40  TEST_SCHED    AZOV          Completed
02/03/2005 17:59:35  02/03/2005 18:00:58  INCR_BCK      CL_MSCS02_TSM Completed
02/03/2005 18:24:52  02/03/2005 18:28:09  RESTORE       CL_MSCS02_TSM Restarted

tsm: TSMSRV03>_
```

Figure 6-86 Schedule restarted on the event log for CL_MSCS02_TSM

Archived

6. When the restore is completed, we see in the schedule log file of the client the final statistics (Figure 6-87).

```
dsmsched.log - Notepad
File  Edit  Format  View  Help
02/03/2005 18:32:44 Restoring         110,861 \\cl_mscs02\g$\code\isc\wps\migration\templat
02/03/2005 18:32:44 Restoring           6,869 \\cl_mscs02\g$\code\isc\wps\migration\templat
02/03/2005 18:32:53 Restoring     111,755,569 \\cl_mscs02\g$\code\storageagent\c8117ml.exe
02/03/2005 18:32:54
Restore processing finished.
02/03/2005 18:32:56 --- SCHEDULEREC STATUS BEGIN
02/03/2005 18:32:56 Total number of objects restored:     1,902
02/03/2005 18:32:56 Total number of objects failed:           0
02/03/2005 18:32:56 Total number of bytes transferred:     1.28 GB
02/03/2005 18:32:56 Data transfer time:                  111.58 sec
02/03/2005 18:32:56 Network data transfer rate:        12,088.45 KB/sec
02/03/2005 18:32:56 Aggregate data transfer rate:      10,678.93 KB/sec
02/03/2005 18:32:56 Elapsed processing time:           00:02:06
02/03/2005 18:32:56 --- SCHEDULEREC STATUS END
02/03/2005 18:32:56 --- SCHEDULEREC OBJECT END RESTORE 02/03/2005 18:24:52
02/03/2005 18:32:56 --- SCHEDULEREC STATUS BEGIN
02/03/2005 18:32:56 --- SCHEDULEREC STATUS END
02/03/2005 18:32:56 Scheduled event 'RESTORE' completed successfully.
02/03/2005 18:32:56 Sending results for scheduled event 'RESTORE'.
02/03/2005 18:32:56 Results sent to server for scheduled event 'RESTORE'.
```

Figure 6-87 Statistics for the restore session

7. And the event log of Tivoli Storage Manager server shows the scheduled operation as completed (Figure 6-88).

```
IBM Tivoli Storage Manager
                    (SESSION: 368)

tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name     Status
-------------------- -------------------- ------------- ------------- ---------
02/03/2005 17:44:00                       TEST_SCHED    CL_HACMP03_C- Pending
                                                         LIENT
02/03/2005 17:44:00  02/03/2005 17:55:03  TEST_SCHED    KANANGA       Completed
02/03/2005 17:44:00  02/03/2005 17:57:40  TEST_SCHED    AZOV          Completed
02/03/2005 17:59:35  02/03/2005 18:00:58  INCR_BCK      CL_MSCS02_TSM Completed
02/03/2005 18:24:52  02/03/2005 18:28:09  RESTORE       CL_MSCS02_TSM Completed
```

Figure 6-88 Schedule name RESTORE completed for CL_MSCS02_TSM

Results summary

The test results show that after a failure on the node that hosts the Tivoli Storage Manager client scheduler instance, a scheduled restore operation started on this node is started again on the second node of the cluster when the service is online.

This is true if the startup window for the scheduled restore operation is not elapsed when the scheduler client is online again on the second node.

Also notice that the restore is not restarted from the point of failure, but started from the beginning. The scheduler queries the Tivoli Storage Manager server for a scheduled operation and a new session is opened for the client after the failover.

Archived

6.6 Protecting the quorum database

Although the MSCS database information is stored locally in the HKLM\Cluster registry hive, it is not sufficient to back up or restore the MSCS database simply by processing this registry hive.

The MSCS database is one of the several system objects available for backup via the Tivoli Storage Manager Backup-Archive client. Please refer to the *Backup-Archive Clients Installation and User's Guide* and to the IBM Redbook, *Deploying the Tivoli Storage Manager Client in a Windows 2000 Environment (SG25-6141)* for information about backup of Windows 2000 system objects.

The Tivoli Storage Manager Backup-Archive client uses the supported API function which creates a snapshot of the cluster configuration. The files are placed in c:\adsm.sys\clusterdb\<clustername> and then sent to Tivoli Storage Manager server. The backup is always full.

There are tools in Microsoft Resource kit that, together with Tivoli Storage Manager, should be used in case of a need to restore the cluster database. They are:

- Clustrest
- DumpConfig

Microsoft Knowledge Base has other materials concerning the backup/restore of the cluster database.

Archived

Archived

7

Microsoft Cluster Server and the IBM Tivoli Storage Manager Storage Agent

This chapter describes the use of Tivoli Storage Manager for Storage Area Network (also known as *Storage Agent*) to back up shared data of a Windows MSCS using the LAN-free path.

We use our two Windows MSCS environments described in Chapter 4:

- Windows 2000 MSCS formed by two servers: POLONIUM and RADON
- Windows 2003 MSCS formed by two servers: SENEGAL and TONGA.

© Copyright IBM Corp. 2005. All rights reserved.

7.1 Overview

The functionality of Tivoli Storage Manager for Storage Area Network (Storage Agent) has been described described under 2.1.2, “IBM Tivoli Storage Manager for Storage Area Networks V5.3” on page 14.

Through the current chapter, we focus in the use of this feature as applied to our Windows clustered environments.

7.2 Planning and design

There are different types of hardware configurations that take advantage of using the Storage Agent for LAN-free backup in a SAN.

An installation must carefully plan and design their own configuration, and they should also check the compatibility and support requirements for the Tivoli Storage Manager for Storage Area Network in order for it to work correctly.

In our lab we use IBM disk and tape Fibre Channel attached storage devices supported by LAN-free backup with Tivoli Storage Manager.

7.2.1 System requirements

Before implementing Tivoli Storage Manager for Storage Area Network, we should access the latest available software levels of all components and check supported hardware and software configurations. For information, see:

http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html

In order to use the Storage Agent for LAN-free backup, we need:

- A Tivoli Storage Manager server with LAN-free license.
- A Tivoli Storage Manager client or a Tivoli Storage Manager Data Protection application client
- A supported Storage Area Network configuration where storage devices and servers are attached for storage sharing purposes
- If we are sharing disk storage, Tivoli SANergy® must be installed. Tivoli SANergy Version 3.2.4 is included with the Storage Agent media.
- The Tivoli Storage Manager for Storage Area Network software.

Archived

7.2.2 System information

We gather all the information about our future client and server systems and use it to implement the LAN-free backup environment according to our needs.

We will need to plan and design carefully things such as:

- Name conventions for local nodes, virtual nodes and Storage Agents
- Number of Storage Agents to use depending upon the connections
- Number of tape drives to be shared and which servers will share them
- Segregate different types of data:
  - Large files and databases to use the LAN-free path
  - Small and numerous files to use the LAN path
- TCP/IP addresses and ports
- Device names used by Windows operating system for the storage devices

7.3 Installing the Storage Agent on Windows MSCS

In order to implement the Storage Agent to work correctly on a Windows 2000 MSCS or Windows 2003 MSCS environment, it is necessary to perform these tasks:

1. Installation of the Storage Agent software on each node of the MSCS, on local disk.
2. If necessary, installation of the correct tape drive device drivers on each node of the MSCS.
3. Configuration of the Storage Agent on each node for LAN-free backup of local disks and also LAN-free backup of shared disks in the cluster.
4. Testing the Storage Agent configuration.

Some of these tasks are exactly the same for Windows 2000 or Windows 2003. For this reason, and to avoid duplicating the information, in this section we describe these common tasks. The specifics of each environment are described later in this chapter, under 7.4, “Storage Agent on Windows 2000 MSCS” on page 333 and 7.5, “Storage Agent on Windows 2003 MSCS” on page 378.

For detailed information on Storage Agent and its implementation, refer to the *Tivoli Storage Manager for SAN for Windows Storage Agent User’s Guide*.

Archived

7.3.1 Installation of the Storage Agent

The installation of the Storage Agent in an MSCS Windows environment follows the same rules as in any single Windows server. It is necessary to install the software on local disk in each node belonging to the same cluster.

In this section we make a summary of this installation process. The same tasks apply to both Windows 2000 environments as well as Windows 2003.

We use the same disk drive letter and installation path on each node:

```
c:\Program Files\Tivoli\tsm\storageagent
```

We start the installation in the first node of each cluster, running **setup.exe** and selecting **Install Products** from the main menu. The Install Products menu appears (Figure 7-1). We first install the TSM Storage Agent and later the TSM Device Driver.

Figure 7-1 Install TSM Storage Agent

Note: Since the installation process is the same as for any other standalone server, we do not show all menus. We only describe a summary of the activities to follow.

TSM Storage Agent installation

To install the Storage Agent:

1. We select **TSM Storage Agent** as shown in Figure 7-1 on page 332.
2. We follow the sequence of panels providing the necessary information and we click **Next** when we are prompted to, accepting the license agreement and selecting the **Complete** installation.
3. After a successful installation, the process prompts for a reboot of the system. Since we are still going to install the TSM device driver, we reply **No**.

TSM device driver installation

To install the device driver:

1. We go back to the Install Products menu and we select **TSM Device Driver**.
2. We follow the sequence of panels providing the necessary information and we click **Next** when we are prompted to, accepting the license agreement and selecting the **Complete** installation.
3. After a successful installation, the process prompts for a reboot of the system. This time we reply **Yes** to reboot the server.

We follow the same tasks in the second node of each cluster.

7.4 Storage Agent on Windows 2000 MSCS

In this section we describe how we configure our Storage Agent software to run in our MSCS Windows 2000, the same cluster we installed and configured in 4.3, “Windows 2000 MSCS installation and configuration” on page 29.

7.4.1 Windows 2000 lab setup

Our Tivoli Storage Manager clients and Storage Agents for the purpose of this section are located on the same Microsoft Windows 2000 Advanced Server Cluster we introduce in Chapter 4, “Microsoft Cluster Server setup” on page 27.

Refer to Table 4-1 on page 30, Table 4-2 on page 31, and Table 4-3 on page 31 for details of the cluster configuration: local nodes, virtual nodes and cluster groups.

We use TSMSRV03, (an AIX machine), as the server, because Tivoli Storage Manager Version 5.3.0 for AIX is, so far, the only platform that supports high availability Library Manager functions for LAN-free backup.

Archived

Tivoli Storage Manager LAN-free configuration details

Figure 7-2 shows our LAN-free configuration:

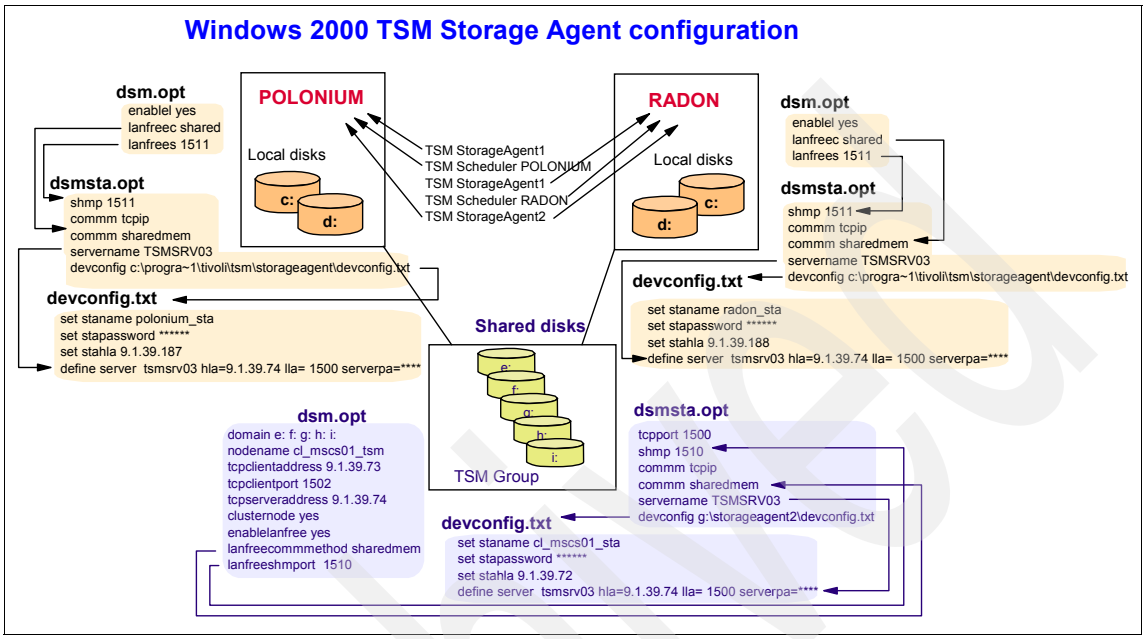

Figure 7-2 Windows 2000 TSM Storage Agent clustering configuration

Archived

For details of this configuration, refer to Table 7-1, Table 7-2, and Table 7-3.

Table 7-1 LAN-free configuration details

Node 1	
TSM nodename	POLONIUM
Storage Agent name	POLONIUM_STA
Storage Agent service name	TSM StorageAgent1
dsmsta.opt and devconfig.txt location	c:\program files\tivoli\tsm\storageagent
Storage Agent high level address	9.1.39.187
Storage Agent low level address	1502
Storage Agent shared memory port	1511
LAN-free communication method	sharedmem
Node 2	
TSM nodename	RADON
Storage Agent name	RADON_STA
Storage Agent service name	TSM StorageAgent1
dsmsta.opt and devconfig.txt location	c:\program files\tivoli\tsm\storageagent
Storage Agent high level address	9.1.39.188
Storage Agent low level address	1502
Storage Agent shared memory port	1511
LAN-free communication method	sharedmem
Virtual node	
TSM nodename	CL_MSCS01_TSM
Storage Agent name	CL_MSCS01_STA
Storage Agent service name	TSM StorageAgent2
dsmsta.opt and devconfig.txt location	g:\storageagent2
Storage Agent high level address	9.1.39.73
Storage Agent low level address	1500
Storage Agent shared memory port	1510

Archived

Node 1	
TSM nodename	POLONIUM
Storage Agent name	POLONIUM_STA
Storage Agent service name	TSM StorageAgent1
dsmsta.opt and devconfig.txt location	c:\program files\tivoli\tsm\storageagent
Storage Agent high level address	9.1.39.187
Storage Agent low level address	1502
Storage Agent shared memory port	1511
LAN-free communication method	sharedmem
Node 2	
TSM nodename	RADON
Storage Agent name	RADON_STA
Storage Agent service name	TSM StorageAgent1
dsmsta.opt and devconfig.txt location	c:\program files\tivoli\tsm\storageagent
Storage Agent high level address	9.1.39.188
Storage Agent low level address	1502
Storage Agent shared memory port	1511
LAN-free communication method	sharedmem
Virtual node	
TSM nodename	CL_MSCS01_TSM
Storage Agent name	CL_MSCS01_STA
Storage Agent service name	TSM StorageAgent2
dsmsta.opt and devconfig.txt location	g:\storageagent2
Storage Agent high level address	9.1.39.73
Storage Agent low level address	1500
LAN-free communication method	sharedmem

Archived

Table 7-2 TSM server details

TSM Server information	
Server name	TSMSRV03
High level address	9.1.39.74
Low level address	1500
Server password for server-to-server communication	password

Our SAN storage devices are described in Table 7-3:

Table 7-3 SAN devices details

SAN devices	
Disk	IBM DS4500 Disk Storage Subsystem
Tape Library	IBM LTO 3582 Tape Library
Tape drives	IBM 3580 Ultrium 2 tape drives
Tape drive device names for Storage Agents	drlto_1: mt0.0.0.4 drlto_2: mt1.0.0.4

Installing IBM 3580 tape drive drivers in Windows 2000

Before implementing the Storage Agent for LAN-free backup in our environment, we need to know that the Windows 2000 OS in each node recognizes the tape drives that will be shared with the Tivoli Storage Manager server.

In our Windows 2000 MSCS, both nodes, RADON and POLONIUM, are attached to the SAN. They recognize the two IBM 3580 tape drives of the IBM 3582 Tape Library managed by the Tivoli Storage Manager server for sharing.

However, when both nodes are started after connecting the devices, the IBM 3580 tape drives display as an interrogation mark under the Other Devices icon. This happens because we need to install the appropriate IBM device drivers for 3580 LTO tape drives.

Once installed, the device drivers must be updated in each local node of the cluster using the Device Manager wizard.

Archived

With this objective, we follow these steps:

1. We first download the latest available IBM TotalStorage® tape drivers from:

   http://www-1.ibm.com/servers/storage/support/allproducts/downloading.html

2. We open the Device Manager, right-click the tape drive, and select **Properties** →**Driver** →**Update Driver**, and the panel in Figure 7-3 displays.

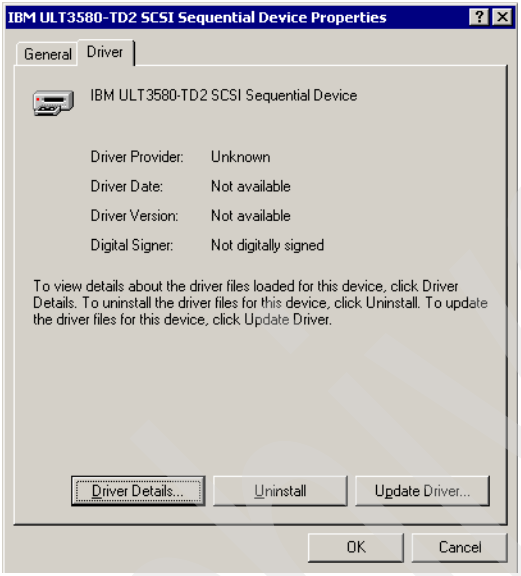

Figure 7-3 Updating the driver

Archived

3. The drivers installation process starts. We follow the sequence of menus, specifying (among other things) the path where the driver files were downloaded and, after a successful installation of the drivers, they should appear listed under the Tape drives icon, as shown in Figure 7-4.

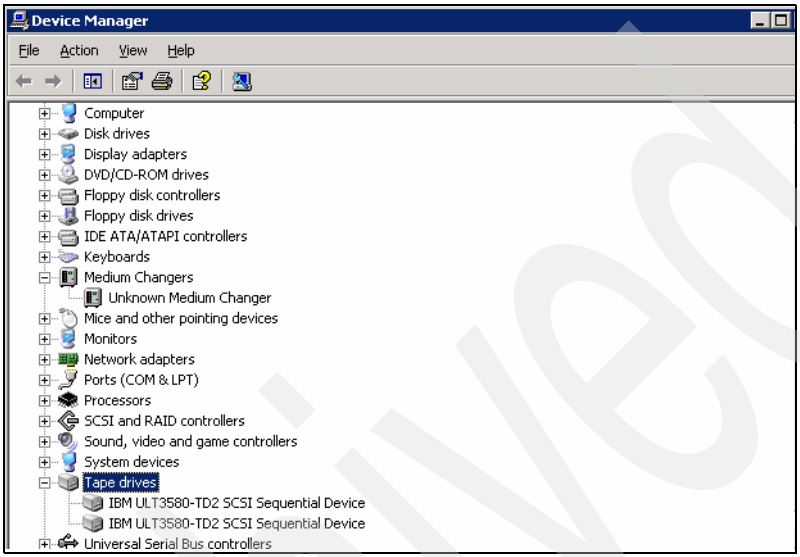

Figure 7-4 Device Manager menu after updating the drivers

Refer to the *IBM Ultrium device drivers Installation and User's Guide* for a detailed description of the installation procedure for the drivers.

7.4.2 Configuration of the Storage Agent on Windows 2000 MSCS

For the configuration of the Storage Agent to be capable of working in a cluster environment, this involves three steps:

1. Configuration of Tivoli Storage Manager server for LAN-free:
   - Establishment of server name, server password, server hladdress, and server lladdress
   - Definition of Storage Agents
   - Definition of the tape library as shared
   - Definition of paths from the Storage Agents to the tape drives
2. Installation of the Storage Agents in the client machines
3. Configuring the Storage Agent for local nodess to communicate with the client and the server for LAN-free purposes

Archived

Configuration of Tivoli Storage Manager server for LAN-free

The process of preparing the Tivoli Storage Manager server for LAN-free data movement is very complex, involving several phases.

Each Storage Agent must be defined as a server in the TSM server. For our lab we use one Storage Agent for each local node and one Storage Agent for the TSM cluster group for high-availability. The naming conventions for these are given in Table 7-1 on page 335.

Setting up parameters for the server

The first task we must do is establishing the server name, server password, server hladdress, and server lladdress on the Tivoli Storage Manager server for the server itself.

Only by setting up these parameters will the Tivoli Storage Manager server be capable of communicating with other servers in the network for LAN-free backup.

From the administrative command line, we run the following commands for our Tivoli Storage Manager AIX server:

```
set servername tsmsrv03
set serverpassword password
set serverhladdress 9.1.39.74
set hladdress 1500
```

LAN-free tasks

These are the activities we follow in our Tivoli Storage Manager server for each Storage Agent:

- Update of the tape library definition as *shared yes*
- Definition of the Storage Agent as a server
- Definition of paths from the Storage Agent to each drive on the tape library
- Setup of a storage pool for LAN-free backup
- Definition of the policy (management class) that points to the LAN-free storage pool
- Validation of the LAN-free environment

Archived

Using the administration console wizard

To set up server-to-server communications, we use the new Administrative Center console of Tivoli Storage Manager Version 5.3.0. This console helps us to cover all the LAN-free tasks.

For details about the Administrative Center installation and how to start a session using this new Web interface, refer to 5.5.1, “Starting the Administration Center console” on page 173.

In this section we only describe, with more detail, the process of enabling LAN-free data movement for one client. We do not show all menus, just the panels we need to achieve this goal.

As an example, we show the activities to define RADON_STA as the Storage Agent used by RADON for LAN-free data movement. We follow the same steps to define POLONIUM_STA (as Storage Agent for POLONIUM) and CL_MSCS01_STA (as Storage Agent for CL_MSCS01_TSM).

1. We open the administration console using a Web browser and we authenticate with a user id (iscadmin) and a password. These are the user id and password we defined in 5.3.4, “Installation of the Administration Center” on page 92.
2. We select the folder **Policy Domains and Client Nodes**.
3. We choose the **TSMSRV03** server, which is the Tivoli Storage Manager server whose policy domain we wish to administer.
4. We select the Domain Name that we want to use for LAN-Free operations, **Standard** in our case. This will take us to open the ***domainname* Properties** portlet.
5. We expand the **Client Nodes** item of the portlet to show a list of clients.

Archived

6. We select the client node for which we want to use LAN-Free data movement, RADON, using the Select radio button. We open the drop down menu, scroll down to **Enable LAN-free Data Movement...** as shown in Figure 7-5 and we click **Go**.

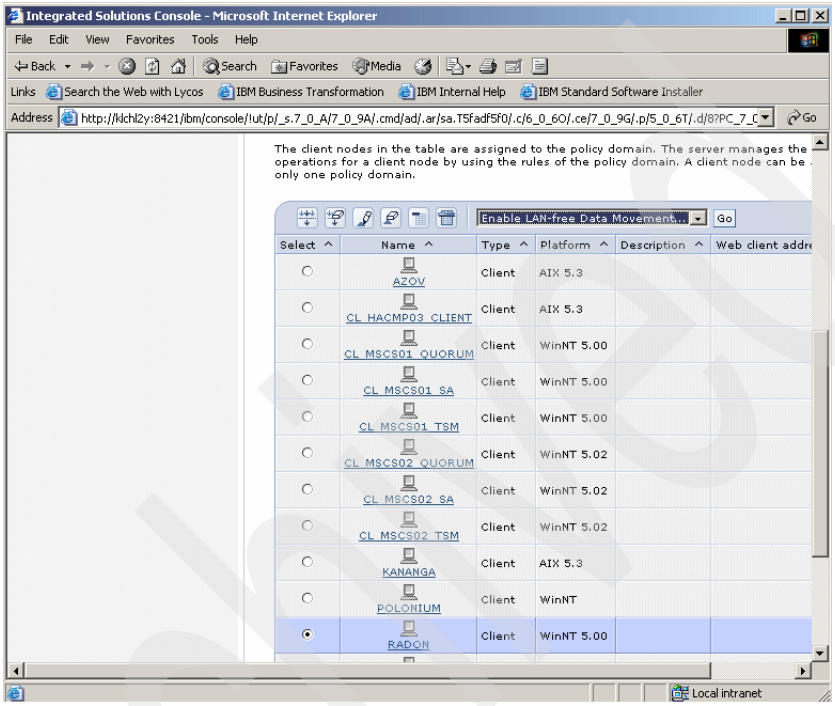

Figure 7-5 Choosing RADON for LAN-free backup

Archived

7. This launches the Enable LAN-free Data Movement wizard as shown in Figure 7-6. We click **Next** in this panel.

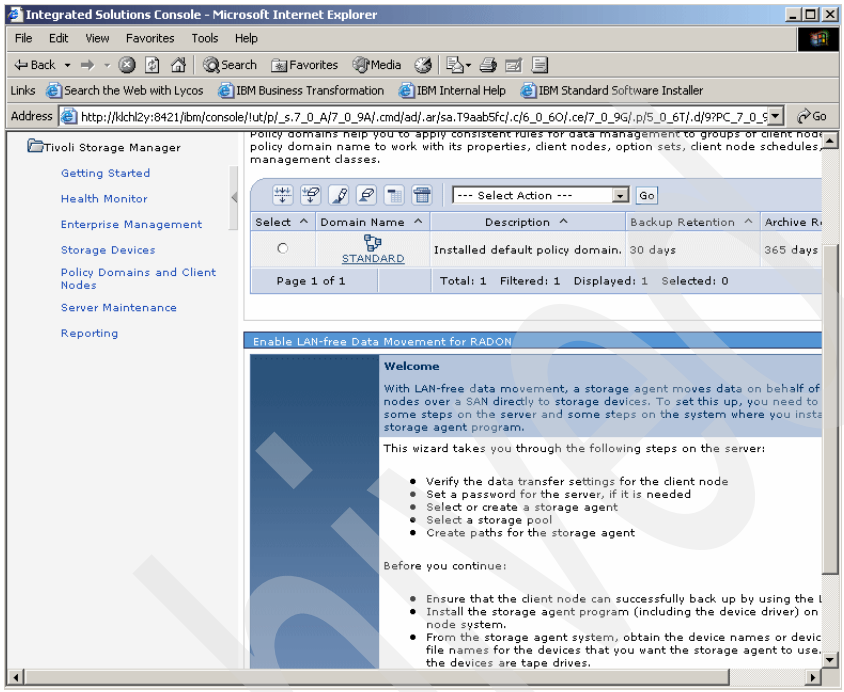

Figure 7-6 Enable LAN-free Data Movement wizard for RADON

Archived

8. In Figure 7-7 we select to allow both LAN as well as LAN-free data transfer and we click **Next**. In this way, if the SAN path fails, the client can use the LAN path.

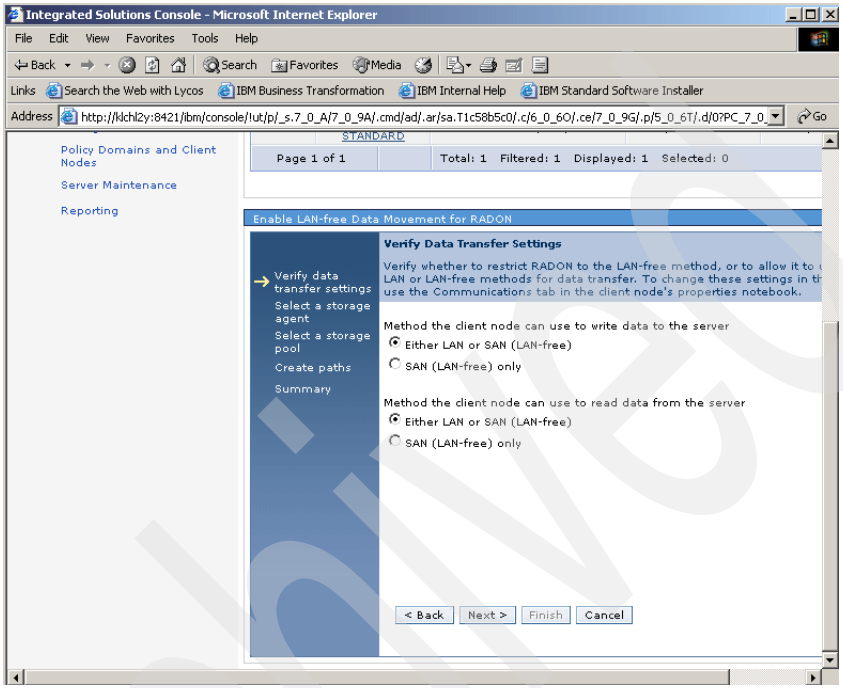

Figure 7-7 Allowing LAN and LAN-free operations for RADON

Archived

9. In Figure 7-8 we choose to **Create a Storage Agent** and we click **Next**.

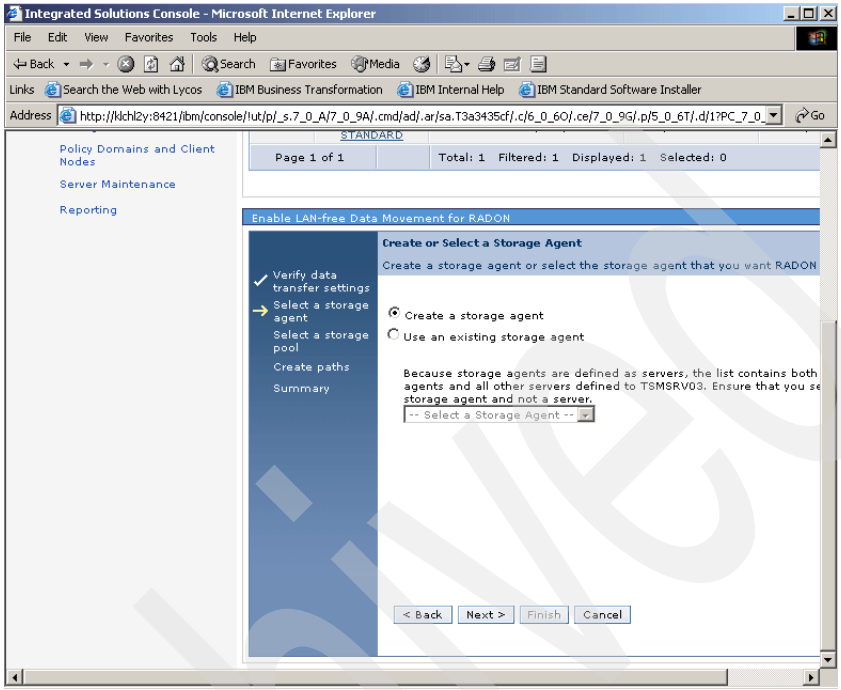

Figure 7-8 Creating a new Storage Agent

Archived

10. We type the name, password, TCP/IP address and port number for the Storage Agent being defined as shown in Figure 7-9 and we click **Next**. Filling in this information in this menu is the same as using the **`define server`** command in the administrative command line.

> **Important:** We must be sure to use the same name, password, TCP/IP address, and port number in Figure 7-8 as when we configure the Storage Agent on the client machine that will use LAN-free backup.

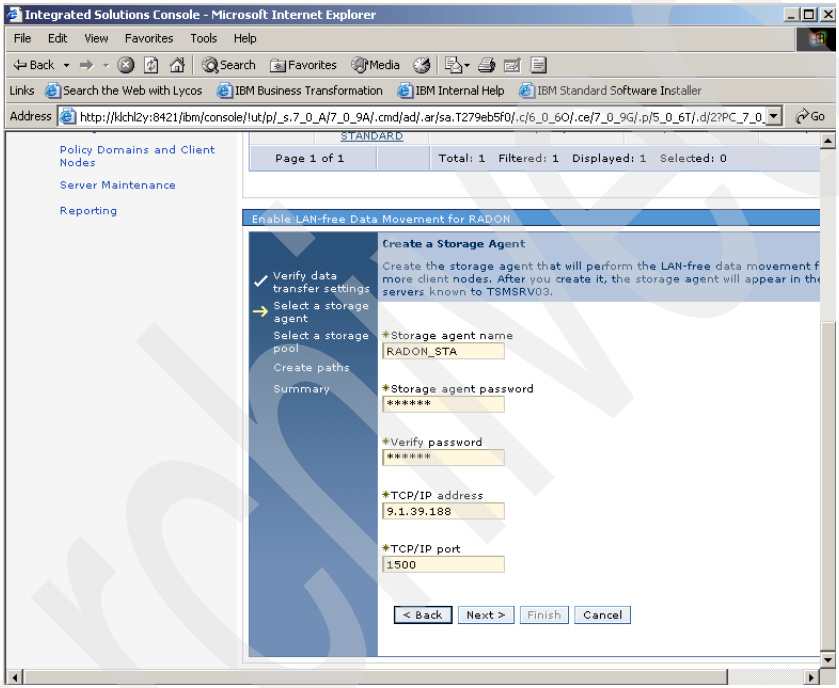

Figure 7-9 Storage agent parameters for RADON

Archived

11. We select which storage pool we want to use for LAN-free backups as shown in Figure 7-10 and we click **Next**. This storage pool had to be defined first.

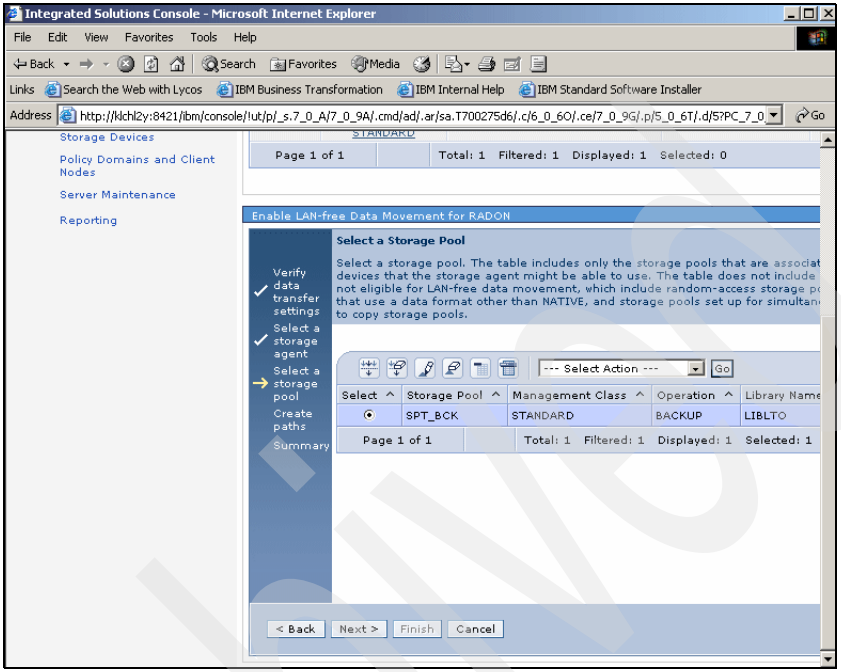

Figure 7-10 Storage pool selection for LAN-free backup

Archived

12. Now we create the paths between the Storage Agent and the tape drives as shown in Figure 7-11. We first choose one drive, select **Modify drive path** and we click **Go**.

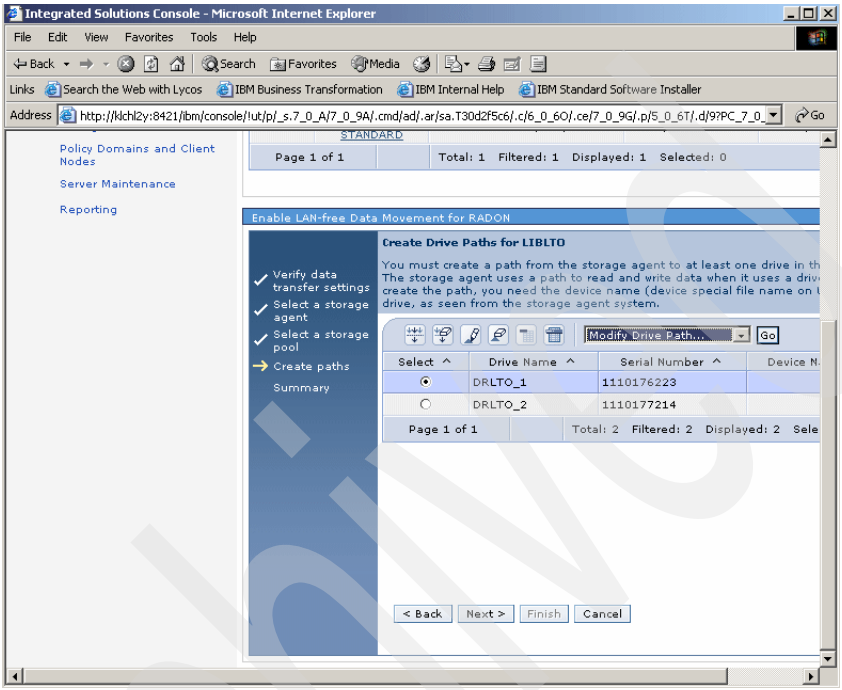

Figure 7-11 Modify drive paths for Storage Agent RADON_STA

Archived

13. In Figure 7-12 we type the device name such as Windows 2000 operating system sees the first drive and we click **Next**.

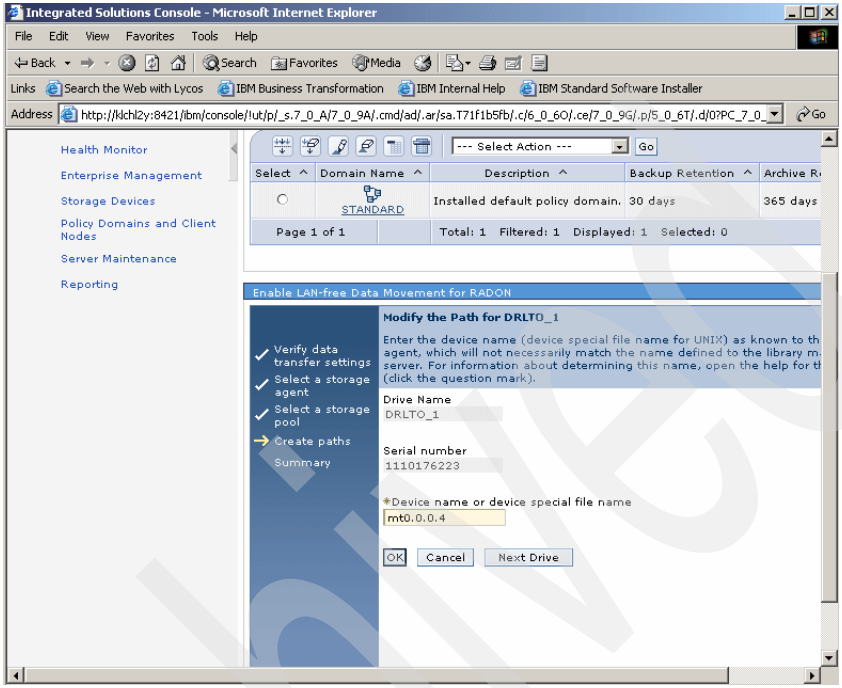

Figure 7-12 Specifying the device name from the operating system view

Archived

The information provided in Figure 7-12 is the same as we would use in the **`define path`** command if we run the administrative command line interface instead.

To know which is the device name for Windows we need to open **Tivoli** Storage Manager management console, in RADON, and go to **Tivoli Storage Manager →TSM Device Driver →Reports →Device Information** as we show in Figure 7-13.

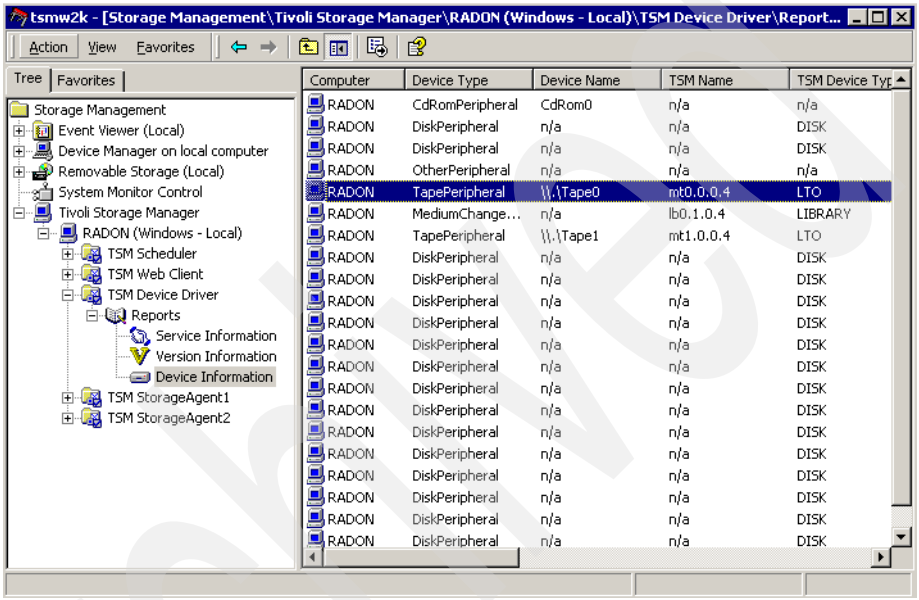

Figure 7-13 Device names for 3580 tape drives attached to RADON

14. Since there is **a second drive in the tape library**, the **configuration process** will ask next for the **device name** of this second drive. We also define the device name for the second drive, and finally the wizard ends. A summary menu displays, informing us about the completion of the LAN-free setup. This menu also advises us about the rest of the tasks we should follow to use LAN-free backup on the client side. We cover these activities in the following sections (Figure 7-14).

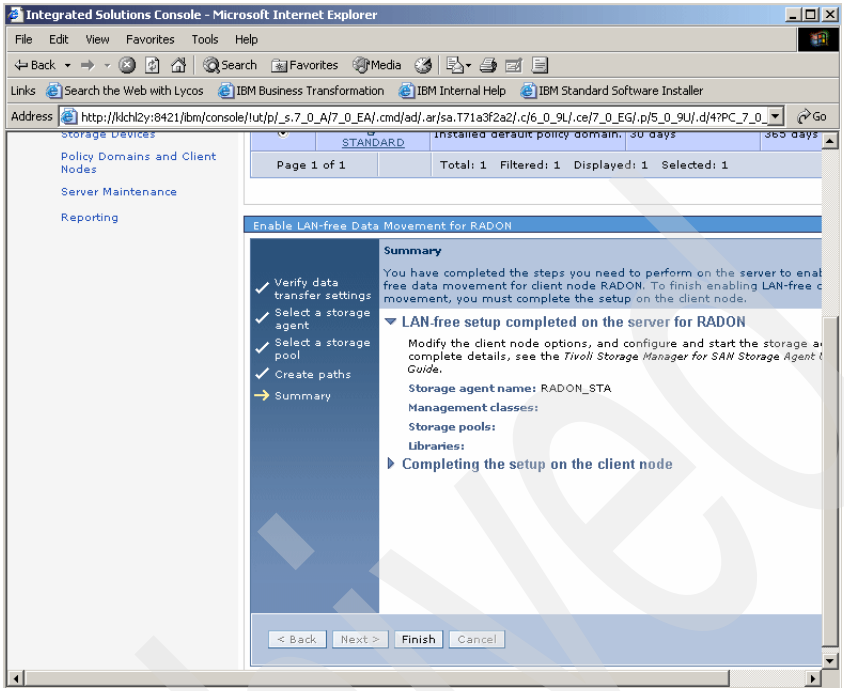

Figure 7-14 LAN-free configuration summary

Configuring the Storage Agent for local nodes

In our lab we use three Storage Agents: one local in each node and one for the TSM Group in the cluster. The configuration process differs between them. Here we describe the configuration tasks for local nodes.

To back up local disk drives on each node using the LAN-free path, we follow the same process we would follow for any single node.

Updating the dsmsta.opt

Before starting to use the management console to initialize an Storage Agent, we change the dsmsta.opt file which is located in the installation path.

We update the option *devconfig* to make sure that points to the whole path where the device configuration file is located:

```
devconfig c:\progra~1\tivoli\tsm\storageagent\devconfig.txt
```

Note: We need to update the dsmsta.opt because the service used to start the Storage Agent uses as default the path where the command is run, not the installation path.

Using the management console to initialize a Storage Agent

We open the management console using **Start →Programs →Tivoli Storage Manager →Management Console**.

1. We start the configuring process with RADON. The initialization wizard starts as shown in Figure 7-15. We click **Next**:

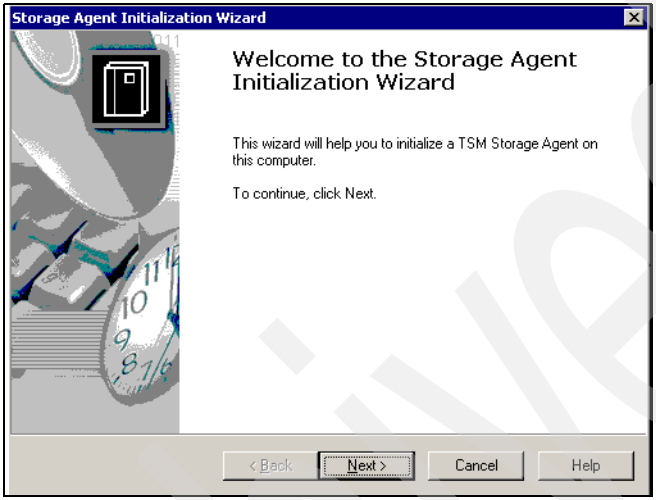

Figure 7-15 Initialization of a local Storage Agent

2. We provide the appropriate information for this Storage Agent: its name, password and high level address and we click **Next** (Figure 7-16).

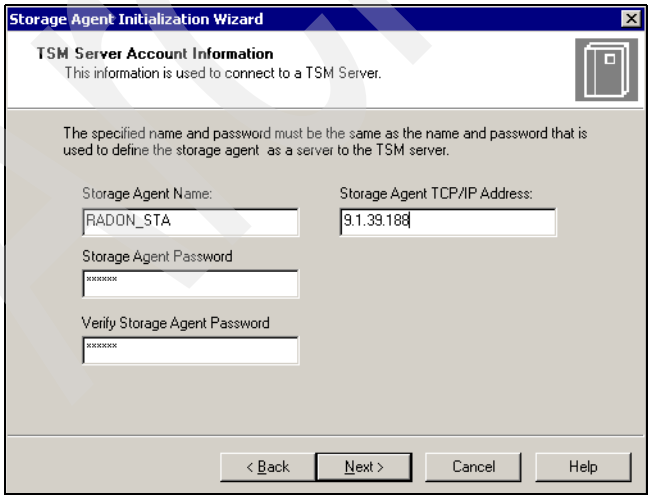

Figure 7-16 Specifying parameters for Storage Agent

Important: we must make sure the Storage Agent name and the rest of the information we provide in this menu matches the parameters used to define the Storage Agent in the Tivoli Storage Manager server in Figure 7-9 on page 346.

3. In the next menu we provide the Tivoli Storage Manager server information: its name, password, TCP/IP address and TCP port. Then we click **Next** (Figure 7-17).

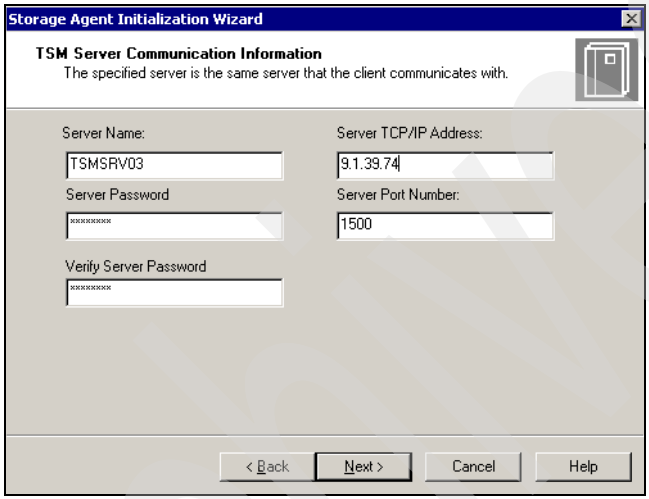

Figure 7-17 Specifying parameters for the Tivoli Storage Manager server

Important: The information provided in Figure 7-17 must match the information provided in the `set servername`, `set serverpassword`, **set serverhladdress** and `set serverlladdress` commands in the Tivoli Storage Manager server.

Archived

4. We select the account under which the service will be started and we also choose **Automatically when Windows boots**. We click **Next** (Figure 7-18).

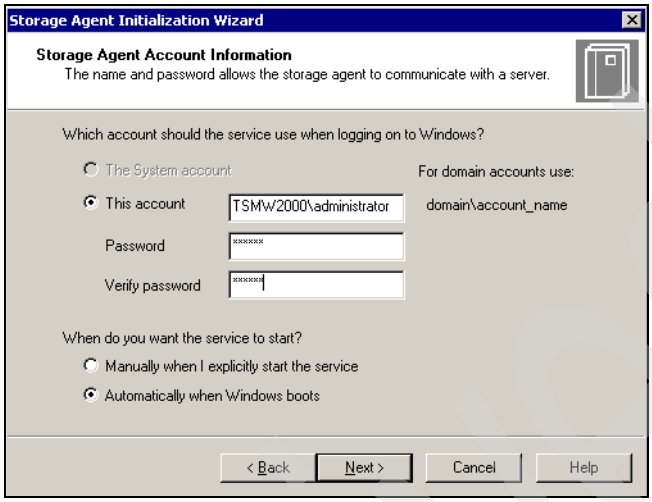

Figure 7-18 Specifying the account information

5. The Completing the Storage Agent Initialization Wizard displays. We click **Finish** in Figure 7-19.

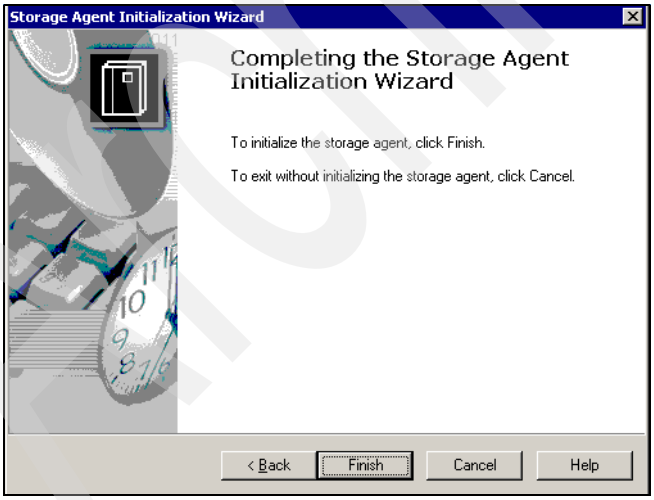

Figure 7-19 Completing the initialization wizard

6. We receive an information menu showing that the account has been granted the right to start the service. We click **OK** (Figure 7-20).

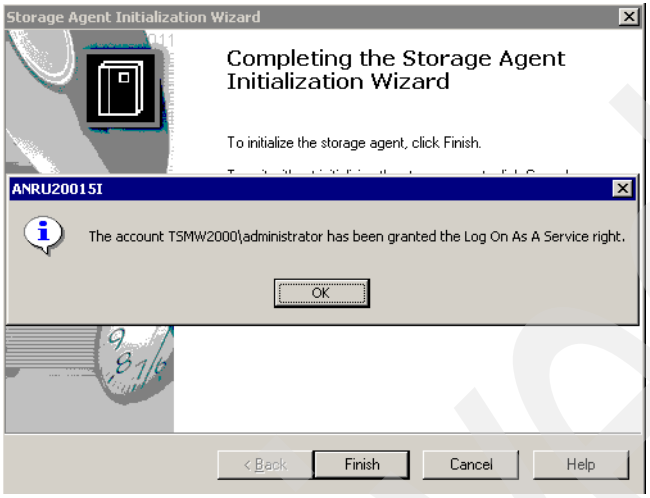

Figure 7-20 Granted access for the account

7. Finally we receive the message that the Storage Agent has been initialized. We click **OK** in Figure 7-21 to end the wizard.

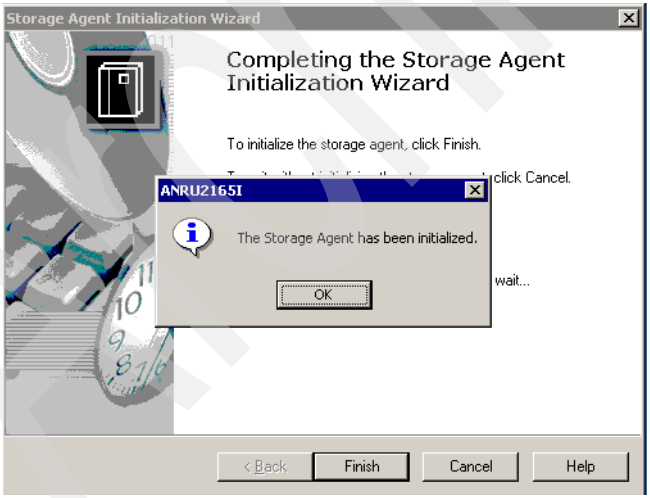

Figure 7-21 Storage agent is successfully initialized

Archived

8. In RADON, after the successful initialization of its Storage Agent the management console displays as shown in Figure 7-22.

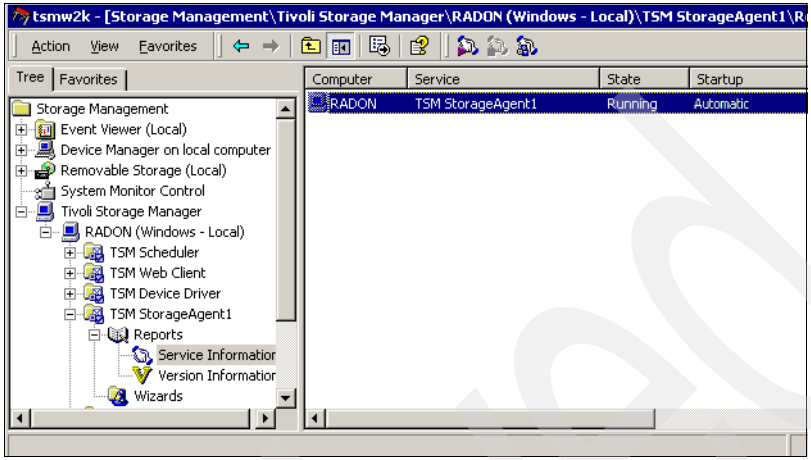

Figure 7-22 TSM StorageAgent1 is started on RADON

For POLONIUM we get a similar menu.

Updating the client option file

To be capable of using LAN-free backup for each local node we must specify certain special options in the client option file.

We edit c:\program files\tivoli\tsm\baclient\dsm.opt and we include the following options:

```
ENABLELANFREE yes
LANFREECOMMMETHOD SHAREDMEM
LANFREESHMPORT 1511
```

We specify the 1511 port for Shared Memory instead of 1510 (the default), because we will use this default port to communicate with the Storage Agent related to the cluster. Port 1511 will be used by the local nodes when communicating to the local Storage Agents.

Instead of the options specified above we also can use:

```
ENABLELANFREE yes
LANFREECOMMMETHOD TCPIP
LANFREETCPPORT 1502
```

Archived

Restarting the Tivoli Storage Manager scheduler service

To use the LAN-free path is necessary, after including the LAN-free options in dsm.opt, to restart the Tivoli Storage Manager scheduler service. If we do not restart the service, the new options will not be read by the client.

Configuring the Storage Agent for virtual nodes

In order to back up shared disk drives in the cluster using the LAN-free path, we can use the Storage Agent instance created for the local nodes. Depending upon the node that hosts the resources at that time, it will be used one local Storage Agent or another one.

This is the technically supported way of configuring LAN-free backup for clustered configurations. Each virtual node in the cluster should use the local Storage Agent in the local node that hosts the resource at that time.

However, in order to also have high-availability for the Storage Agent, we configure a new Storage Agent instance that will be used for the cluster.

> **Attention:** This is not a technically supported configuration but, in our lab tests, it worked.

In the following sections we describe the process for our TSM Group, where a TSM Scheduler generic service resource is located for backup of e: f: g: h: and i: shared disk drives.

Using the dsmsta setstorageserver utility

Instead of using the management console to create the new instance we use the **`dsmsta`** utility from an MS-DOS prompt. The reason to use this tool is because we have to create a new registry key for this Storage Agent. If we start the management console we would use the default key, StorageAgent1, and we need a different one.

With that objective, we perform these tasks:

1. We begin the configuration in the node that hosts the shared disk drives, POLONIUM.
2. We start copying the “storageagent” folder (created at installation time) from c:\program files\tivoli\tsm onto a shared disk drive (g:) with the name “storageagent2”.
3. We open an MS-DOS prompt and change to g:\storageagent2.

4. From this path we run the command we see in Figure 7-23 to create another instance for a Storage Agent called StorageAgent2. For this instance, the option (dsmsta.opt) and device configuration (devconfig.txt) files will be located on this path.

```
Command Prompt

G:\storageagent2>dsmsta -k StorageAgent2 setstorageserver myname=CL_MSCS01_STA m
ypassword=itsosj myhla=9.1.39.73 servername=tsmsrv03 serverpa=password hla=9.1.3
9.74 lla=1500
ANR0900I Processing options file G:\storageagent2\dsmsta.opt.
ANR7800I DSMSERV generated at 12:45:40 on Jan 28 2005.

Tivoli Storage Manager for Windows
Version 5, Release 3, Level 0.1

*****************************************************************
* Patch 5.3.0.1 TSM Server (Driver 0000, Build 05028.1245) *
*                                                               *
*****************************************************************
Licensed Materials - Property of IBM

(C) Copyright IBM Corporation 1990, 2004.
All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

ANR4726I The ICC support module has been loaded.
ANR1432I Updating device configuration information to defined files.
ANR1433I Device configuration information successfully written to g:\storageag-
ent2\devconfig.txt.
ANR2119I The SERVERNAME option has been changed in the options file.
ANR0467I The SETSTORAGESERVER command completed successfully.

G:\storageagent2>_
```

Figure 7-23 Installing Storage Agent for LAN-free backup of shared disk drives

Attention: Notice in Figure 7-23 the new registry key used for this Storage Agent, StorageAgent2, as well as the name and IP address specified in the “myname” and “myhla” parameters. The Storage Agent name is CL_MSCS01_STA, and its IP address is the IP address of the TSM Group. Also notice that executing the command from g:\storageagent2 we make sure that the dsmsta.opt and devconfig.txt updated files are the ones in this path.

5. Now, from the same path, we run a command to install a service called “TSM StorageAgent2” related to the StorageAgent2 instance created in step 4. The command and the result of its execution are shown in Figure 7-24.

Figure 7-24 Installing the service related to StorageAgent2

6. If we open the Tivoli Storage Manager management console in this node, we now can see two instances for two Storage Agents: the one we created for the local node, TSM StorageAgent1, and a new one, TSM Storage Agent2, which is set to *Manual*. This last instance is stopped, as we can see in Figure 7-25.

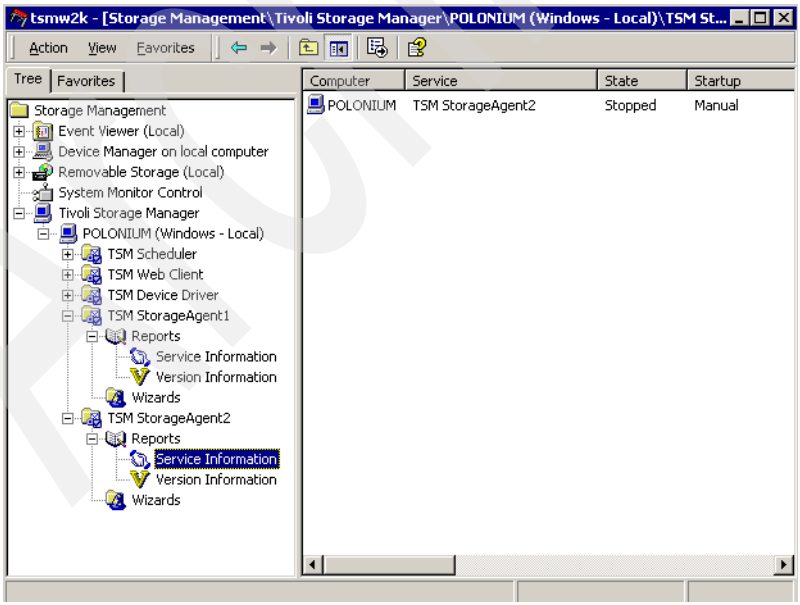

Figure 7-25 Management console displays two Storage Agents

Archived

7. We start the TSM StorageAgent2 instance right-clicking and selecting **Start** as we show in Figure 7-26.

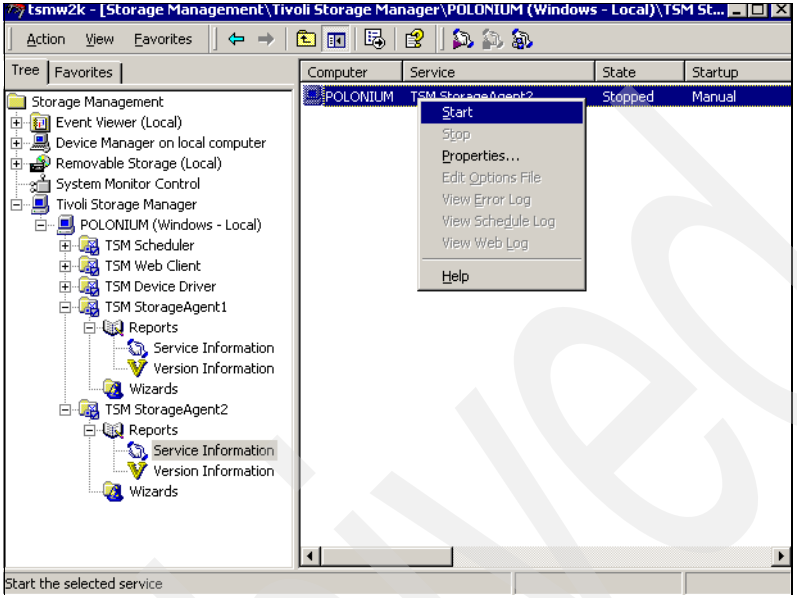

Figure 7-26 Starting the TSM StorageAgent2 service in POLONIUM

8. Now we have two Storage Agent instances running in POLONIUM:
   - TSM StorageAgent1: Related to the local node, that uses the dsmsta.opt and devconfig.txt files located in c:\program files\tivoli\tsm\storageagent
   - TSM StorageAgent2: Related to the virtual node, which uses the dsmsta.opt and devconfig.txt files located in g:\storageagent2
9. We stop the TSM StorageAgent2 and move the resources to RADON.

Archived

10. In RADON, we follow steps 3 to 5. Then, we open the Tivoli Storage Manager management console and we again find two Storage Agent instances: TSM StorageAgent1 (for the local node) and TSM StorageAgent2 (for the virtual node). This last instance is stopped and set to manual as shown in Figure 7-27.

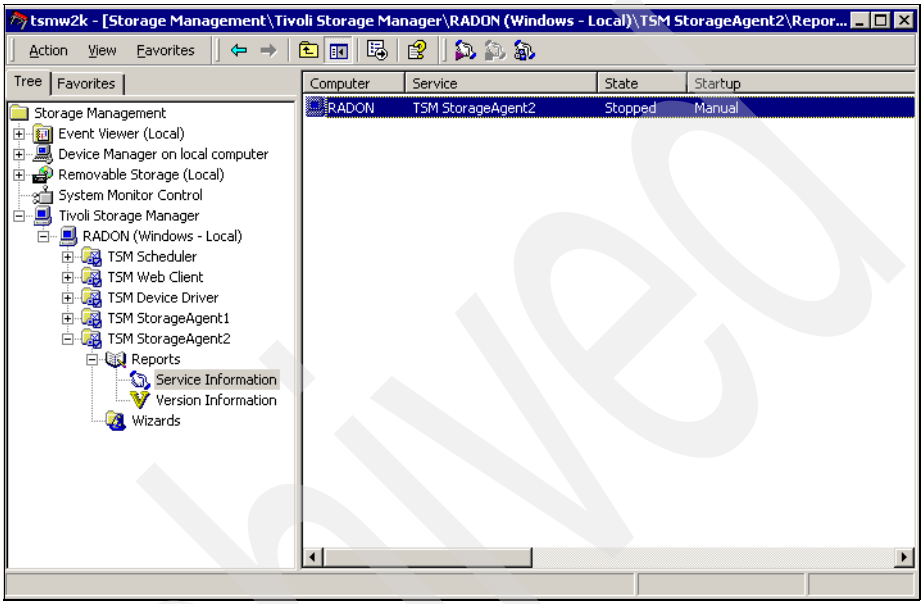

Figure 7-27 TSM StorageAgent2 installed in RADON

11. We start the instance right-clicking and selecting **Start**. After a successful start, we stop it again.

12. Finally, the last task consists of the definition of TSM StorageAgent2 as a cluster resource. To do this, we open the Cluster Administrator, we right-click the resource group where Tivoli Storage Manager scheduler service is defined, TSM Group, and we select to define a new resource as shown in Figure 7-28.

Archived

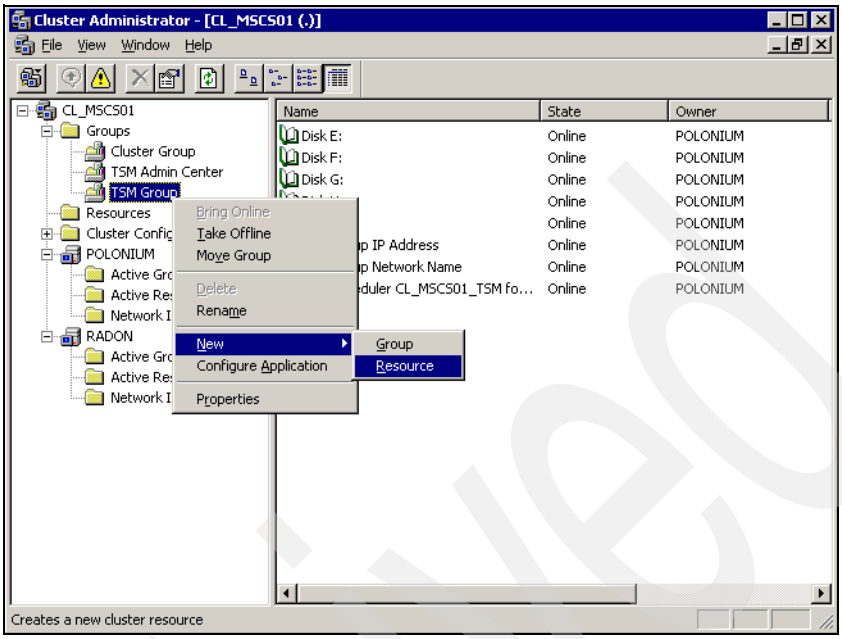

Figure 7-28 Use cluster administrator to create resource for TSM StorageAgent2

13. We type a name for the resource and we select Generic Service as the resource type. Then we click **Next** as we see in Figure 7-29.

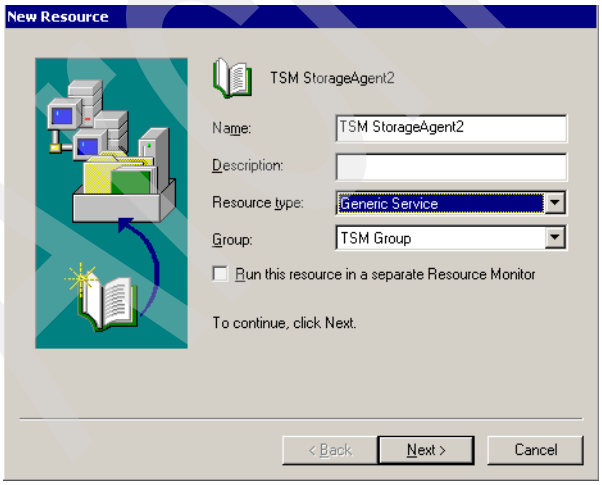

Figure 7-29 Defining a generic service resource for TSM StorageAgent2

14. In Figure 7-30 we leave both nodes as possible owners and we click **Next**.

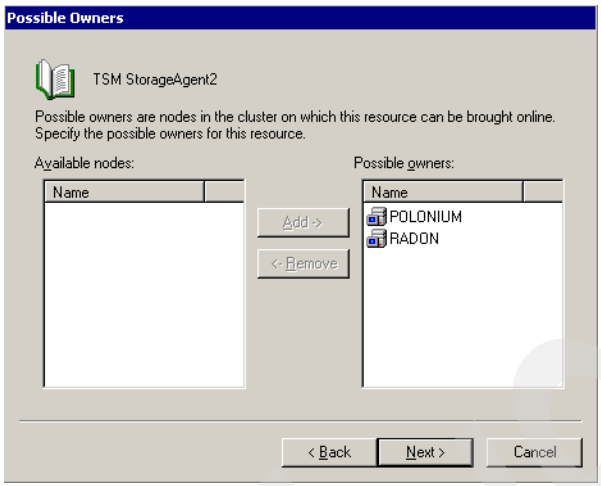

Figure 7-30 Possible owners for TSM StorageAgent2

15. As TSM StorageAgent2 dependencies we select the **Disk G:** drive which is where the configuration files are located for this instance. After adding the disk, we click **Next** in Figure 7-31.

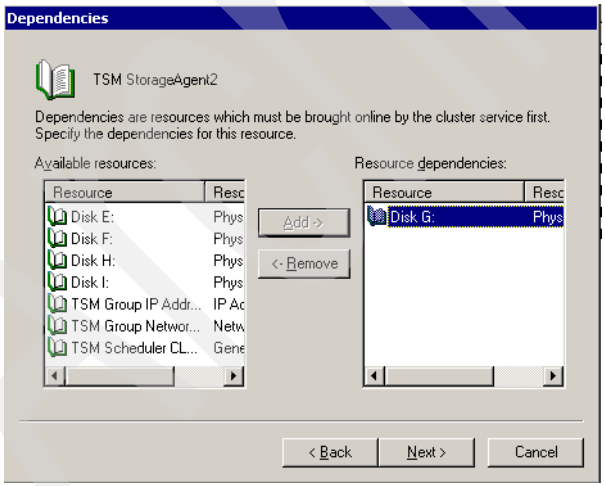

Figure 7-31 Dependencies for TSM StorageAgent2

16. We provide the name of the service, **TSM StorageAgent2** and then we click **Next** in Figure 7-32.

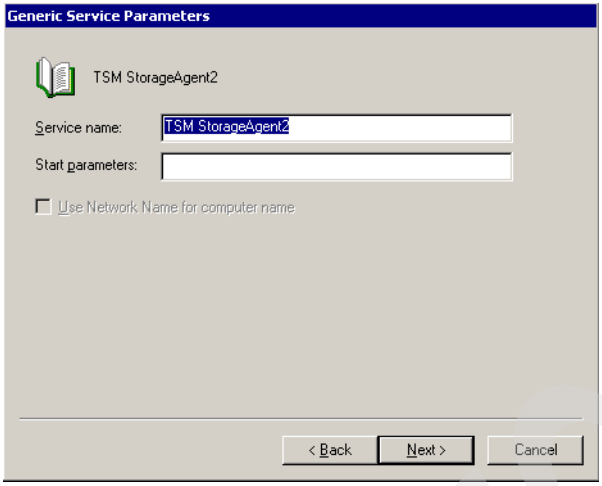

Figure 7-32 Service name for TSM StorageAgent2

Important: The name of the service in Figure 7-32 must match exactly the name we used to install the instance in both nodes.

17. We do not use any registry key replication for this resource. We click **Finish** in Figure 7-33.

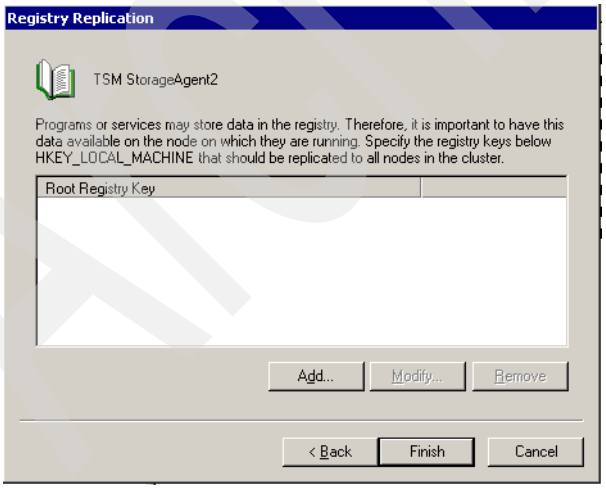

Figure 7-33 Registry key for TSM StorageAgent2

18. The new resource is successfully created as Figure 7-34 displays. We click **OK**.

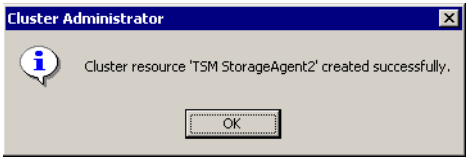

Figure 7-34 Generic service resource created successfully:TSM StorageAgent2

19. The last task is bringing online the new resource as we show in Figure 7-35.

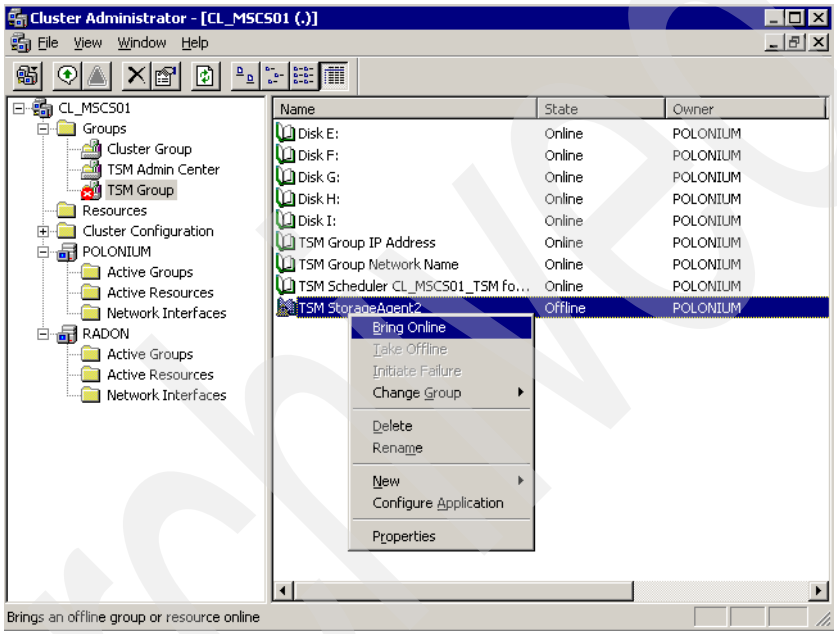

Figure 7-35 Bringing the TSM StorageAgent2 resource online

20. At this time the service is started in the node that hosts the resource group. To check the successful implementation of this Storage Agent, we move the resources to the second node and we check that TSM StorageAgent2 is now started in this second node and stopped in the first one.

> **Important:** be sure to use only the Cluster Administrator to start and stop the StorageAgent2 instance at any time.

Changing the dependencies for the TSM Scheduler resource

Since we want the Tivoli Storage Manager scheduler always to use the LAN-free path when it starts, it is necessary to update its associated resource in Cluster Administrator to add TSM StorageAgent2 as a dependency to bring it online.

For this reason, we open the Cluster Administrator, select the TSM Scheduler resource for CL_MSCS01_TSM and go to **Properties** → **Dependencies** → **Modify**. Once there, we add TSM StorageAgent2 as a dependency, as we show in Figure 7-36.

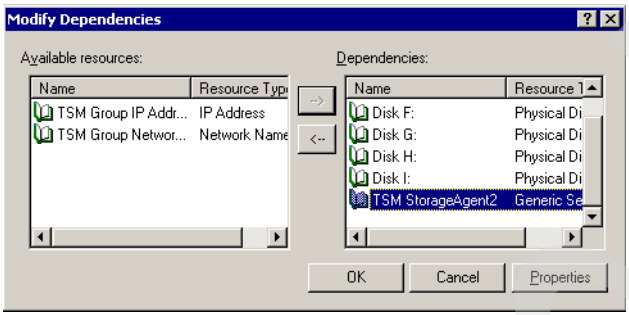

Figure 7-36 Adding Storage Agent resource as dependency for TSM Scheduler

We click **OK** and bring the resource online again. With this dependency we make sure the Tivoli Storage Manager scheduler is not started for this cluster group before the Storage Agent does.

Updating the client option file

To be capable of using LAN-free backup for the virtual node, we must specify certain special options in the client option file for the virtual node.

We open g:\tsm\dsm.opt and we include the following options:

```
ENABLELANFREE yes
LANFREECOMMMETHOD SHAREDMEM
LANFREESHMPORT 1510
```

For the virtual node we use the default shared memory port, 1510.

Instead of the options above, we also can use:

```
ENABLELANFREE yes
LANFREECOMMMETHOD TCPIP
LANFREETCPPORT 1500
```

Restarting the Tivoli Storage Manager scheduler service

After including the LAN-free options in dsm.opt, we restart the Tivoli Storage Manager scheduler service for the TSM Group using the Cluster Administrator. If we do not restart the service, the new options will not be read by the client.

Archived

7.4.3 Testing Storage Agent high availability on Windows 2000 MSCS

The purpose of this section is to test our LAN-free setup for the clustering.

We use the TSM Group (nodename CL_MSCS01_TSM) to test LAN-free backup/restore of shared data in our Windows 2000 cluster.

Our objective with these tasks is to know how the Storage Agent and the Tivoli Storage Manager Library Manager work together to respond, on a client clustered environment, after certain kinds of failures that affect the shared resources.

Again, for details of our LAN-free configuration, refer to Table 7-1 on page 335 and Table 7-2 on page 337.

Testing LAN-free client incremental backup

First we test a scheduled client incremental backup using the SAN path.

Objective

The objective of this test is to show what happens when a LAN-free client incremental backup is started for a virtual node in the cluster using the Storage Agent created for this group (CL_MSCS01_STA), and the node that hosts the resources at that moment suddenly fails.

Activities

To do this test, we perform these tasks:

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager scheduler service for TSM Group. At this time RADON does.
2. We schedule a client incremental backup operation using the Tivoli Storage Manager server scheduler and we associate the schedule to CL_MSCS01_TSM nodename.
3. We make sure that TSM StorageAgent2 and TSM Scheduler for CL_MSCS01_TSM are online resources on RADON.
4. When it is the scheduled time, a client session for CL_MSCS01_TSM nodename starts on the server. At the same time, several sessions are also started for CL_MSCS01_STA for Tape Library Sharing and the Storage Agent prompts the Tivoli Storage Manager server to mount a tape volume, as we can see in Figure 7-37.

Archived

```
IBM Tivoli Storage Manager
02/07/2005 10:21:50  ANR0406I Session 1050 started for node CL_MSCS01_TSM
                      (WinNT) (Tcp/Ip 9.1.39.188(4846)). (SESSION: 1050)
02/07/2005 10:21:56  ANR2017I Administrator ADMIN issued command: QUERY ACTLOG
                      begint=now-0:03  (SESSION: 1017)
02/07/2005 10:21:56  ANR0406I (Session: 1046, Origin: CL_MSCS01_STA)  Session 8
                      started for node CL_MSCS01_TSM (WinNT) (Tcp/Ip
                      radon.tsmw2000.com(4849)). (SESSION: 1046)
02/07/2005 10:21:56  ANR0408I Session 1051 started for server CL_MSCS01_STA
                      (Windows) (Tcp/Ip) for storage agent.  (SESSION: 1051)
02/07/2005 10:21:56  ANR0408I (Session: 1046, Origin: CL_MSCS01_STA)  Session 9
more...   <<ENTER> to continue, 'C' to cancel>

                      started for server CL_MSCS01_STA (AIX-RS/6000) (Tcp/Ip)
                      for storage agent.  (SESSION: 1046)
02/07/2005 10:21:56  ANR0415I Session 1051 proxied by CL_MSCS01_STA started for
                      node CL_MSCS01_TSM. (SESSION: 1051)
02/07/2005 10:21:56  ANR0408I Session 1052 started for server CL_MSCS01_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 1052)
02/07/2005 10:21:56  ANR0408I (Session: 1046, Origin: CL_MSCS01_STA)  Session
                      10 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 1046)
02/07/2005 10:21:56  ANR0409I (Session: 1046, Origin: CL_MSCS01_STA)  Session
                      10 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                      1046)
02/07/2005 10:21:56  ANR0409I Session 1052 ended for server CL_MSCS01_STA
                      (Windows). (SESSION: 1052)
02/07/2005 10:21:56  ANR0408I Session 1053 started for server CL_MSCS01_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 1053)
02/07/2005 10:21:56  ANR0408I (Session: 1046, Origin: CL_MSCS01_STA)  Session
                      11 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 1046)
02/07/2005 10:22:11  ANR2017I Administrator ADMIN issued command: QUERY MO
                      (SESSION: 1017)
02/07/2005 10:22:11  ANR8379I Mount point in device class LTO is waiting for
                      the volume mount to complete, status: WAITING FOR VOLUME.
```

Figure 7-37 Storage agent CL_MSCS01_STA session for tape library sharing

5. After a few seconds, the Tivoli Storage Manager server mounts the tape volume 028AKK in drive DRLTO_2, and it informs the Storage Agent about the drive where the volume is mounted. The Storage Agent CL_MSCS01_STA opens then the tape volume as an output volume and starts sending data to the DRLTO_2 as shown in Figure 7-38.

```
IBM Tivoli Storage Manager
02/07/2005 10:22:23  ANR8337I LTO volume 028AKK mounted in drive DRLTO_2
                      (/dev/rmt1). (SESSION: 1053)
02/07/2005 10:22:23  ANR0409I (Session: 1046, Origin: CL_MSCS01_STA)  Session
                      11 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                      1046)
02/07/2005 10:22:23  ANR0409I Session 1053 ended for server CL_MSCS01_STA
                      (Windows). (SESSION: 1053)
02/07/2005 10:22:23  ANR8337I (Session: 1046, Origin: CL_MSCS01_STA)  LTO
                      volume 028AKK mounted in drive DRLTO_2 (mt1.0.0.4).
                      (SESSION: 1046)
more...   <<ENTER> to continue, 'C' to cancel>

02/07/2005 10:22:23  ANR0511I Session 1051 opened output volume 028AKK.
                      (SESSION: 1051)
02/07/2005 10:22:23  ANR0511I (Session: 1046, Origin: CL_MSCS01_STA)  Session 8
                      opened output volume 028AKK. (SESSION: 1046)
```

Figure 7-38 A tape volume is mounted and the Storage Agent starts sending data

6. The client, by means of the Storage Agent, starts sending files to the drive using the SAN path, as we see on its schedule log file in Figure 7-39.

```
dsmsched_aix - Notepad
File  Edit  Format  Help
Executing scheduled command now.
02/07/2005 10:21:24 Node Name: CL_MSCS01_TSM
02/07/2005 10:21:24 Session established with server TSMSRV03: AIX-RS/6000
02/07/2005 10:21:24   Server Version 5, Release 3, Level 0.0
02/07/2005 10:21:24   Server date/time: 02/07/2005 10:14:02  Last access: 02/0

02/07/2005 10:21:24 --- SCHEDULEREC OBJECT BEGIN INCR_BACKUP 02/07/2005 10:11:
02/07/2005 10:21:24 Incremental backup of volume '\\cl_mscs01\e$'
02/07/2005 10:21:24 Incremental backup of volume '\\cl_mscs01\f$'
02/07/2005 10:21:24 Incremental backup of volume '\\cl_mscs01\g$'
02/07/2005 10:21:24 Incremental backup of volume '\\cl_mscs01\h$'
02/07/2005 10:21:24 Incremental backup of volume '\\cl_mscs01\i$'
02/07/2005 10:21:39 Directory-->                    0 \\cl_mscs01\e$\ [Sent]
02/07/2005 10:21:39 Directory-->                    0 \\cl_mscs01\e$\RECYCLER [
02/07/2005 10:21:39 Directory-->                    0 \\cl_mscs01\e$\System vol
02/07/2005 10:21:39 Directory-->                    0 \\cl_mscs01\e$\tivoli [Se
02/07/2005 10:21:39 Directory-->                    0 \\cl_mscs01\e$\tsmdata [S
02/07/2005 10:21:39 Directory-->                    0 \\cl_mscs01\e$\RECYCLER\S
02/07/2005 10:21:39 Directory-->                    0 \\cl_mscs01\e$\RECYCLER\S
02/07/2005 10:21:39 Normal File-->                 65 \\cl_mscs01\e$\RECYCLER\S
02/07/2005 10:21:39 Normal File-->                820 \\cl_mscs01\e$\RECYCLER\S
02/07/2005 10:21:39 Directory-->                    0 \\cl_mscs01\e$\RECYCLER\S
```

Figure 7-39 Client starts sending files to the TSM server in the schedule log file

7. While the client continues sending files to the server, we force RADON to fail. The following sequence takes place:

   a. The client and also the Storage Agent lose their connections with the server temporarily, and both sessions are terminated as we can see on the Tivoli Storage Manager server activity log shown in Figure 7-40.

```
IBM Tivoli Storage Manager
02/07/2005 10:26:00  ANR7835I (Session: 1046, Origin: CL_MSCS01_STA)  The
                      server thread 7 (tid 3044) terminated in response to
                      server shutdown. (SESSION: 1046)
02/07/2005 10:26:00  ANR7835I (Session: 1046, Origin: CL_MSCS01_STA)  The
                      server thread 8 (tid 3224) terminated in response to
                      server shutdown. (SESSION: 1046)
02/07/2005 10:26:00  ANR7835I (Session: 1046, Origin: CL_MSCS01_STA)  The
                      server thread 9 (tid 2532) terminated in response to
                      server shutdown. (SESSION: 1046)
02/07/2005 10:26:00  ANR7835I (Session: 1046, Origin: CL_MSCS01_STA)  The
                      server thread 10 (tid 412) terminated in response to
                      server shutdown. (SESSION: 1046)
02/07/2005 10:26:00  ANR7835I (Session: 1046, Origin: CL_MSCS01_STA)  The
                      server thread 11 (tid 2892) terminated in response to
more...    (<ENTER> to continue, 'C' to cancel)

                      server shutdown. (SESSION: 1046)
02/07/2005 10:26:00  ANR7835I (Session: 1046, Origin: CL_MSCS01_STA)  The
                      server thread 12 (tid 3036) terminated in response to
                      server shutdown. (SESSION: 1046)
02/07/2005 10:26:00  ANR7835I (Session: 1046, Origin: CL_MSCS01_STA)  The
                      server thread 13 (tid 3008) terminated in response to
                      server shutdown. (SESSION: 1046)
02/07/2005 10:26:00  ANR7835I (Session: 1046, Origin: CL_MSCS01_STA)  The
                      server thread 14 (tid 2900) terminated in response to
                      server shutdown. (SESSION: 1046)
02/07/2005 10:26:00  ANR7835I (Session: 1046, Origin: CL_MSCS01_STA)  The
                      server thread 15 (tid 3172) terminated in response to
                      server shutdown. (SESSION: 1046)
02/07/2005 10:26:00  ANR0991I (Session: 1046, Origin: CL_MSCS01_STA)  Storage
                      agent shutdown complete. (SESSION: 1046)
02/07/2005 10:26:00  ANR0480W Session 1049 for node CL_MSCS01_TSM (WinNT)
                      terminated - connection with client severed. (SESSION:
                      1049)
02/07/2005 10:26:00  ANR0480W Session 1050 for node CL_MSCS01_TSM (WinNT)
                      terminated - connection with client severed. (SESSION:
                      1050)
02/07/2005 10:26:02  ANR0480W Session 1051 for node CL_MSCS01_TSM (WinNT)
                      terminated - connection with client severed. (SESSION:
```

Figure 7-40 Sessions for TSM client and Storage Agent are lost in the activity log

b. In the Cluster Administrator menu, RADON is not in the cluster and POLONIUM begins to bring the resources online.

c. The tape volume is still mounted on the same drive.

d. After a short period of time the resources are online on POLONIUM.

e. When the Storage Agent CL_MSCS01_STA is again online (in POLONIUM), the TSM Scheduler service also is started (because of the dependency between these two resources). We can see this on the activity log in Figure 7-41.

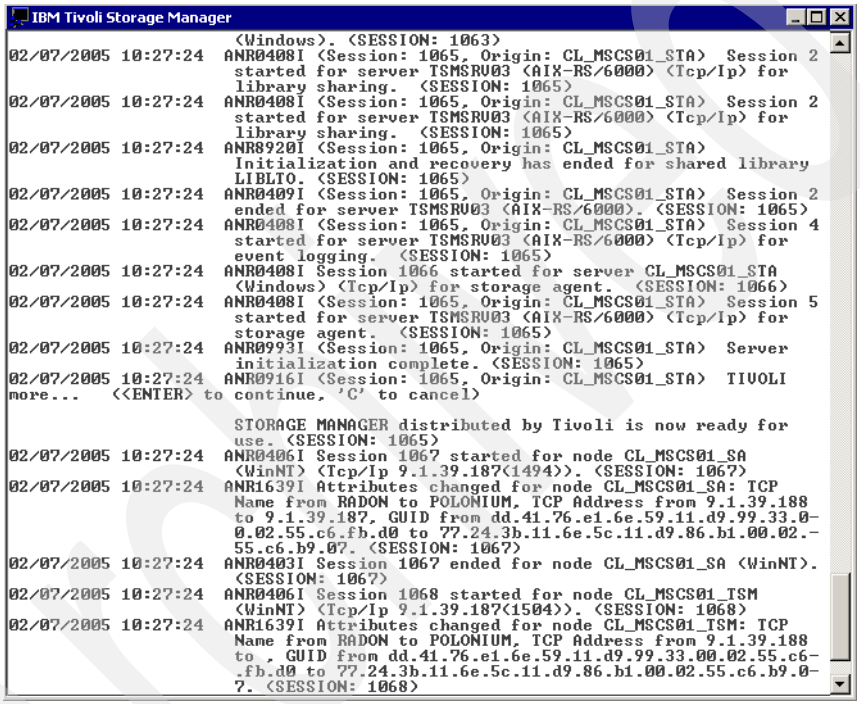

Figure 7-41 Both Storage Agent and TSM client restart sessions in second node

f. The Tivoli Storage Manager server resets the SCSI bus, dismounting the tape volume from the drive for the Storage Agent CL_MSCS01_STA, as we can see in Figure 7-42.

```
IBM Tivoli Storage Manager
02/07/2005 10:27:31  ANR8331I LTO volume 028AKK is mounted R/W in drive DRLTO_2
more...   <<ENTER> to continue, 'C' to cancel>

                      (/dev/rmt1), status: DISMOUNTING. (SESSION: 1056)
02/07/2005 10:27:31  ANR8334I          1 matches found. (SESSION: 1056)
02/07/2005 10:27:35  ANR0408I (Session: 1065, Origin: CL_MSCS01_STA)  Session 3
                      started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 1065)
02/07/2005 10:27:35  ANR0408I Session 1064 started for server CL_MSCS01_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 1063)
02/07/2005 10:27:35  ANR0409I Session 1064 ended for server CL_MSCS01_STA
                      (Windows). (SESSION: 1063)
02/07/2005 10:27:35  ANR0409I (Session: 1065, Origin: CL_MSCS01_STA)  Session 3
                      ended for server TSMSRV03 (AIX-RS/6000). (SESSION: 1065)
02/07/2005 10:28:05  ANR2017I Administrator ADMIN issued command: QUERY ACTLOG
                      begint=now-0:05  (SESSION: 1056)
02/07/2005 10:28:17  ANR8336I Verifying label of LTO volume 028AKK in drive
                      DRLTO_2 (/dev/rmt1). (SESSION: 1053)
02/07/2005 10:28:53  ANR0406I Session 1069 started for node CL_MSCS01_SA
                      (WinNT) (Tcp/Ip 9.1.39.187(1515)). (SESSION: 1069)
02/07/2005 10:28:53  ANR0403I Session 1069 ended for node CL_MSCS01_SA (WinNT).
                      (SESSION: 1069)
02/07/2005 10:28:54  ANR0406I Session 1070 started for node CL_MSCS01_TSM
                      (WinNT) (Tcp/Ip 9.1.39.187(1517)). (SESSION: 1070)
02/07/2005 10:28:55  ANR8468I LTO volume 028AKK dismounted from drive DRLTO_2
                      (/dev/rmt1) in library LIBLTO. (SESSION: 1053)
```

Figure 7-42 Tape volume is dismounted by the Storage Agent

g. Finally, the client restarts its scheduled incremental backup using the SAN path and the tape volume is mounted again by the Tivoli Storage Manager server for use of the Storage Agent, as we can see in Figure 7-43.

```
IBM Tivoli Storage Manager

02/07/2005 10:29:02  ANR0406I Session 1071 started for node CL_MSCS01_TSM
                      (WinNT) (Tcp/Ip 9.1.39.187(1519)). (SESSION: 1071)
02/07/2005 10:29:08  ANR0406I (Session: 1065, Origin: CL_MSCS01_STA)  Session
                      10 started for node CL_MSCS01_TSM (WinNT) (Tcp/Ip
                      polonium.tsmw2000.com(1523)). (SESSION: 1065)
02/07/2005 10:29:08  ANR0408I Session 1072 started for server CL_MSCS01_STA
                      (Windows) (Tcp/Ip) for storage agent.  (SESSION: 1072)
02/07/2005 10:29:08  ANR0408I (Session: 1065, Origin: CL_MSCS01_STA)  Session
                      11 started for server CL_MSCS01_STA (AIX-RS/6000)
                      (Tcp/Ip) for storage agent.  (SESSION: 1065)
02/07/2005 10:29:08  ANR0415I Session 1072 proxied by CL_MSCS01_STA started for
                      node CL_MSCS01_TSM. (SESSION: 1072)
02/07/2005 10:29:08  ANR0408I Session 1073 started for server CL_MSCS01_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 1073)
02/07/2005 10:29:08  ANR0408I (Session: 1065, Origin: CL_MSCS01_STA)  Session
                      12 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 1065)
02/07/2005 10:29:08  ANR0409I (Session: 1065, Origin: CL_MSCS01_STA)  Session
                      12 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                      1065)
02/07/2005 10:29:08  ANR0409I Session 1073 ended for server CL_MSCS01_STA
                      (Windows). (SESSION: 1073)
02/07/2005 10:29:08  ANR0408I Session 1074 started for server CL_MSCS01_STA
more...   <<ENTER> to continue, 'C' to cancel>

                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 1074)
02/07/2005 10:29:08  ANR0408I (Session: 1065, Origin: CL_MSCS01_STA)  Session
                      13 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 1065)
02/07/2005 10:29:40  ANR8337I LTO volume 028AKK mounted in drive DRLTO_1
                      (/dev/rmt0). (SESSION: 1074)
02/07/2005 10:29:40  ANR0409I (Session: 1065, Origin: CL_MSCS01_STA)  Session
                      13 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                      1065)
02/07/2005 10:29:40  ANR0409I Session 1074 ended for server CL_MSCS01_STA
                      (Windows). (SESSION: 1074)
02/07/2005 10:29:40  ANR8337I (Session: 1065, Origin: CL_MSCS01_STA)  LTO
                      volume 028AKK mounted in drive DRLTO_1 (mt0.0.0.4).
                      (SESSION: 1065)
02/07/2005 10:29:40  ANR0511I Session 1072 opened output volume 028AKK.
                      (SESSION: 1072)
02/07/2005 10:29:40  ANR0511I (Session: 1065, Origin: CL_MSCS01_STA)  Session
                      10 opened output volume 028AKK. (SESSION: 1065)
```

Figure 7-43 The scheduled is restarted and the tape volume mounted again

8. The incremental backup ends successfully, as we can see on the final statistics recorded by the client in its schedule log file in Figure 7-44.

```
dsmsched_aix - Notepad
File  Edit  Format  Help

02/07/2005 10:39:31 --- SCHEDULEREC STATUS BEGIN
02/07/2005 10:39:31 Total number of objects inspected:    3,711
02/07/2005 10:39:31 Total number of objects backed up:    1,301
02/07/2005 10:39:31 Total number of objects updated:          0
02/07/2005 10:39:31 Total number of objects rebound:          0
02/07/2005 10:39:31 Total number of objects deleted:          0
02/07/2005 10:39:31 Total number of objects expired:          0
02/07/2005 10:39:31 Total number of objects failed:           0
02/07/2005 10:39:31 Total number of bytes transferred:    434.87 MB
02/07/2005 10:39:31 LanFree data bytes:                   434.81 MB
02/07/2005 10:39:31 Data transfer time:                     5.02 sec
02/07/2005 10:39:31 Network data transfer rate:        88,636.56 KB/sec
02/07/2005 10:39:31 Aggregate data transfer rate:      2,275.79 KB/sec
02/07/2005 10:39:31 Objects compressed by:                    0%
02/07/2005 10:39:31 Elapsed processing time:           00:03:15
02/07/2005 10:39:31 --- SCHEDULEREC STATUS END
02/07/2005 10:39:31 --- SCHEDULEREC OBJECT END INCR_BACKUP 02/07/2005 10:20:39
02/07/2005 10:39:31 Scheduled event 'INCR_BACKUP' completed successfully.
02/07/2005 10:39:31 Sending results for scheduled event 'INCR_BACKUP'.
02/07/2005 10:39:32 Results sent to server for scheduled event 'INCR_BACKUP'.
```

Figure 7-44 Final statistics for LAN-free backup

Results summary

The test results show that, after a failure on the node that hosts both the Tivoli Storage Manager scheduler as well as the Storage Agent shared resources, a scheduled incremental backup started on one node for LAN-free is restarted and successfully completed on the other node, also using the SAN path.

This is true if the startup window used to define the schedule is not elapsed when the scheduler service restarts on the second node.

The Tivoli Storage Manager server on AIX resets the SCSI bus when the Storage Agent is restarted on the second node. This permits us to dismount the tape volume from the drive where it was mounted before the failure. When the client restarts the LAN-free operation, the same Storage Agent commands the server to mount again the tape volume to continue the backup.

Restriction: This configuration, with two Storage Agents installed on the same node, is not technically supported by Tivoli Storage Manager for SAN. However, in our lab environment it worked.

Note: In other tests we made using the local Storage Agent on each node for communication to the virtual client for LAN-free, the SCSI bus reset did not work. The reason is that when Tivoli Storage Manager server on AIX acts as a Library Manager, can handle the SCSI bus reset only when the Storage Agent name is the same for the failing and recovering Storage Agent.

In other words, if we use local Storage Agents for LAN-free backup of the virtual client (CL_MSCS01_TSM), the following conditions must be taken into account: The failure of the node RADON means that all local services will also fail, including RADON_STA (the local Storage Agent). MSCS will cause a failover to the second node where the local Storage Agent will be started again, but with a different name (POLONIUM_STA). It is this discrepancy in naming which will cause the LAN-free backup to fail, as clearly, the virtual client will be unable to connect to RADON_STA.

Tivoli Storage Manager server does not know what happened to the first Storage Agent, because it does not receive any alert from it until the node that failed is again up, so that the tape drive is in a RESERVED status until the default timeout (10 minutes) elapses. If the scheduler for CL_MSCS01_TSM starts a new session before the ten minutes timeout elapses, it tries to communicate to the local Storage Agent of this second node, POLONIUM_STA, and this prompts the Tivoli Storage Manager server to mount the same tape volume.

Since this tape volume is still mounted on the first drive by RADON_STA (even when the node failed) and the drive is RESERVED, the only option for the Tivoli Storage Manager server is to mount a new tape volume in the second drive. If either there are not enough tape volumes in the tape storage pool, or the second drive is busy at that time with another operation, or if the client node has its maximum mount points limited to 1, the backup is cancelled.

Testing client restore

Our second test is a scheduled restore using the SAN path.

Objective

The objective of this test is to show what happens when a LAN-free restore is started for a virtual node in the cluster, and the node that hosts the resources at that moment suddenly fails.

Activities

To do this test, we perform these tasks:

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager scheduler resource: POLONIUM.
2. We schedule a client restore operation using the Tivoli Storage Manager server scheduler and we associate the schedule to CL_MSCS01_TSM nodename.
3. We make sure that TSM StorageAgent2 and TSM Scheduler for CL_MSCS01_TSM are online resources on POLONIUM.

4. When it is the scheduled time, a client session for CL_MSCS01_TSM nodename starts on the server. At the same time, several sessions are also started for CL_MSCS01_STA for Tape Library Sharing, and the Storage Agent prompts the Tivoli Storage Manager server to mount a tape volume. The tape volume is mounted in drive DRLTO_2. All of these events are shown in Figure 7-45.

```
IBM Tivoli Storage Manager
02/07/2005 11:20:19  ANR0406I Session 1086 started for node CL_MSCS01_TSM
                      (WinNT) (Tcp/Ip 9.1.39.187(1792)). (SESSION: 1086)
02/07/2005 11:20:21  ANR0408I Session 1087 started for server CL_MSCS01_STA
                      (Windows) (Tcp/Ip) for storage agent.  (SESSION: 1087)
02/07/2005 11:20:21  ANR0415I Session 1087 proxied by CL_MSCS01_STA started for
                      node CL_MSCS01_TSM. (SESSION: 1087)
02/07/2005 11:20:22  ANR0408I Session 1088 started for server CL_MSCS01_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 1088)
02/07/2005 11:20:22  ANR0409I Session 1088 ended for server CL_MSCS01_STA
                      (Windows). (SESSION: 1088)
02/07/2005 11:20:22  ANR0408I Session 1089 started for server CL_MSCS01_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 1089)
02/07/2005 11:20:48  ANR8337I LTO volume 028AKK mounted in drive DRLTO_2
                      (/dev/rmt1). (SESSION: 1089)
02/07/2005 11:20:48  ANR0409I Session 1089 ended for server CL_MSCS01_STA
                      (Windows). (SESSION: 1089)
```

Figure 7-45 Starting restore session for LAN-free

5. The client starts restoring files as we can see on the schedule log file in Figure 7-46.

```
dsmsched_aix - Notepad
File  Edit  Format  Help
Executing scheduled command now.
02/07/2005 11:27:41 Node Name: CL_MSCS01_TSM
02/07/2005 11:27:41 Session established with server TSMSRV03: AIX-RS/6000
02/07/2005 11:27:41   Server Version 5, Release 3, Level 0.0
02/07/2005 11:27:41   Server date/time: 02/07/2005 11:20:19  Last access: 02/0

02/07/2005 11:27:41 --- SCHEDULEREC OBJECT BEGIN RESTORE 02/07/2005 11:16:05
02/07/2005 11:27:41 Restore function invoked.

02/07/2005 11:27:42 ANS1247I Waiting for files from the server... ** Interrupt
02/07/2005 11:28:16 ANS1114I Waiting for mount of offline media.
02/07/2005 11:28:16 Restoring              0 \\cl_mscs01\g$\code\adminc [Done
02/07/2005 11:28:16 Restoring              0 \\cl_mscs01\g$\code\isc [Done]
02/07/2005 11:28:16 Restoring            263 \\cl_mscs01\g$\code\GSKit-README
02/07/2005 11:28:16 Restoring      1,205,657 \\cl_mscs01\g$\code\GSKit.zip [D
02/07/2005 11:28:17 Restoring     25,388,669 \\cl_mscs01\g$\code\TSM530C_WIN3
02/07/2005 11:28:18 Restoring     10,762,389 \\cl_mscs01\g$\code\TSMCON5300_W
02/07/2005 11:28:19 Restoring     12,426,660 \\cl_mscs01\g$\code\TSMLNG5300_C
02/07/2005 11:28:20 Restoring     12,444,265 \\cl_mscs01\g$\code\TSMLNG5300_C
02/07/2005 11:28:20 Restoring     12,865,309 \\cl_mscs01\g$\code\TSMLNG5300_D
02/07/2005 11:28:21 Restoring     12,826,564 \\cl_mscs01\g$\code\TSMLNG5300_E
02/07/2005 11:28:22 Restoring     12,881,222 \\cl_mscs01\g$\code\TSMLNG5300_F
02/07/2005 11:28:22 Restoring     12,770,448 \\cl_mscs01\g$\code\TSMLNG5300_I
02/07/2005 11:28:23 Restoring     12,885,559 \\cl_mscs01\g$\code\TSMLNG5300_J
02/07/2005 11:28:23 Restoring     12,543,008 \\cl_mscs01\g$\code\TSMLNG5300_K
02/07/2005 11:28:24 Restoring     12,718,897 \\cl_mscs01\g$\code\TSMLNG5300_F
```

Figure 7-46 Restore starts on the schedule log file

6. While the client is restoring the files, we force POLONIUM to fail. The following sequence takes place:

   a. The client CL_MSCS01_TSM and the Storage Agent CL_MSCS01_STA temporarily lose both of their connections with the server, as shown in Figure 7-47.

```
IBM Tivoli Storage Manager
02/07/2005 11:21:20  ANR0481W Session 1086 for node CL_MSCS01_TSM (WinNT)
                      terminated - client did not respond within 60 seconds.
                      (SESSION: 1086)
02/07/2005 11:21:22  ANR2017I Administrator ADMIN issued command: QUERY EVENT *
                      *  (SESSION: 1085)
02/07/2005 11:21:41  ANR0568W Session 1082 for admin ADMIN (WinNT) terminated -
                      connection with client severed. (SESSION: 1082)
02/07/2005 11:21:45  ANR2017I Administrator ADMIN issued command: QUERY MO
                      (SESSION: 1085)
02/07/2005 11:21:45  ANR8330I LTO volume 028AKK is mounted R/W in drive DRLTO_2
                      (/dev/rmt1), status: IN USE. (SESSION: 1085)
02/07/2005 11:21:45  ANR8334I         1 matches found. (SESSION: 1085)
02/07/2005 11:21:46  ANR0409I Session 1062 ended for server CL_MSCS01_STA
                      (Windows). (SESSION: 1062)
02/07/2005 11:21:47  ANR3605E Unable to communicate with storage agent.
                      (SESSION: 1066)
02/07/2005 11:21:47  ANR0480W Session 1087 for node CL_MSCS01_TSM (WinNT)
                      terminated - connection with client severed. (SESSION:
                      1087)
```

Figure 7-47 Both sessions for the Storage Agent and the client lost in the server

   b. The tape volume is still mounted on the same drive.

   c. After a short period of time the resources are online on RADON.

   d. When the Storage Agent CL_MSCS01_STA is again online (in RADON), the TSM Scheduler service also is started (because of the dependency between these two resources). We can see this on the activity log in Figure 7-48.

```
IBM Tivoli Storage Manager
02/07/2005 11:23:04  ANR0408I Session 1093 started for server CL_MSCS01_STA
                      (Windows) (Tcp/Ip) for storage agent.  (SESSION: 1093)
02/07/2005 11:23:04  ANR0408I Session 1094 started for server CL_MSCS01_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 1094)
02/07/2005 11:23:04  ANR0409I Session 1094 ended for server CL_MSCS01_STA
                      (Windows). (SESSION: 1094)
02/07/2005 11:23:04  ANR0408I Session 1096 started for server CL_MSCS01_STA
                      (Windows) (Tcp/Ip) for event logging.  (SESSION: 1096)
02/07/2005 11:23:04  ANR0408I Session 1097 started for server CL_MSCS01_STA
                      (Windows) (Tcp/Ip) for storage agent.  (SESSION: 1097)
02/07/2005 11:23:05  ANR0408I (Session: 1096, Origin: CL_MSCS01_STA)  Session 2
                      started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 1096)
02/07/2005 11:23:05  ANR0408I (Session: 1096, Origin: CL_MSCS01_STA)  Session 2
                      started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 1096)
02/07/2005 11:23:05  ANR8920I (Session: 1096, Origin: CL_MSCS01_STA)
                      Initialization and recovery has ended for shared library
                      LIBLTO. (SESSION: 1096)
02/07/2005 11:23:05  ANR0409I (Session: 1096, Origin: CL_MSCS01_STA)  Session 2
                      ended for server TSMSRV03 (AIX-RS/6000). (SESSION: 1096)
02/07/2005 11:23:05  ANR0408I (Session: 1096, Origin: CL_MSCS01_STA)  Session 3
                      started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      event logging.  (SESSION: 1096)
02/07/2005 11:23:05  ANR0408I (Session: 1096, Origin: CL_MSCS01_STA)  Session 5
                      started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      storage agent.  (SESSION: 1096)
02/07/2005 11:23:05  ANR0406I Session 1098 started for node CL_MSCS01_TSM
                      (WinNT) (Tcp/Ip 9.1.39.188(1537)). (SESSION: 1098)
02/07/2005 11:23:05  ANR1639I Attributes changed for node CL_MSCS01_TSM: TCP
                      Name from POLONIUM to RADON, GUID from
                      77.24.3b.11.6e.5c.11.d9.86.b1.00.02.55.c6.b9.07 to
```

Figure 7-48 Resources are started again in the second node

e. The Tivoli Storage Manager resets the SCSI bus and dismounts the tape volume such as we can see in Figure 7-49.

```
IBM Tivoli Storage Manager
02/07/2005 11:23:30  ANR0409I (Session: 1096, Origin: CL_MSCS01_STA)  Session 4
                      ended for server TSMSRV03 (AIX-RS/6000). (SESSION: 1096)
02/07/2005 11:23:41  ANR2017I Administrator ADMIN issued command: QUERY SESSION
                       (SESSION: 1085)
02/07/2005 11:23:44  ANR2017I Administrator ADMIN issued command: QUERY MO
                      (SESSION: 1085)
02/07/2005 11:23:44  ANR8331I LTO volume 028AKK is mounted R/W in drive DRLTO_2
                      (/dev/rmt1), status: DISMOUNTING. (SESSION: 1085)
02/07/2005 11:23:44  ANR8334I          1 matches found. (SESSION: 1085)
02/07/2005 11:23:53  ANR8336I Verifying label of LTO volume 028AKK in drive
                      DRLTO_2 (/dev/rmt1). (SESSION: 1089)
02/07/2005 11:24:31  ANR8468I LTO volume 028AKK dismounted from drive DRLTO_2
                      (/dev/rmt1) in library LIBLTO. (SESSION: 1089)
```

Figure 7-49 Tape volume is dismounted by the Storage Agent

f. Finally, the client restarts its scheduled restore and the tape volume is mounted again by the Tivoli Storage Manager server for use of the Storage Agent as we can see in Figure 7-50.

```
IBM Tivoli Storage Manager
                      (Tcp/Ip) for storage agent.  (SESSION: 1096)
02/07/2005 11:24:37  ANR0415I Session 1101 proxied by CL_MSCS01_STA started for
                      node CL_MSCS01_TSM. (SESSION: 1101)
02/07/2005 11:24:37  ANR0408I Session 1102 started for server CL_MSCS01_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 1102)
02/07/2005 11:24:37  ANR0408I (Session: 1096, Origin: CL_MSCS01_STA)  Session
                      11 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 1096)
02/07/2005 11:24:37  ANR0409I (Session: 1096, Origin: CL_MSCS01_STA)  Session
                      11 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                      1096)
02/07/2005 11:24:37  ANR0409I Session 1102 ended for server CL_MSCS01_STA
                      (Windows). (SESSION: 1102)
02/07/2005 11:24:37  ANR0408I Session 1103 started for server CL_MSCS01_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 1103)
02/07/2005 11:24:37  ANR0408I (Session: 1096, Origin: CL_MSCS01_STA)  Session
                      12 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 1096)
02/07/2005 11:25:02  ANR8337I LTO volume 028AKK mounted in drive DRLTO_1
                      (/dev/rmt0). (SESSION: 1103)
02/07/2005 11:25:02  ANR0409I (Session: 1096, Origin: CL_MSCS01_STA)  Session
                      12 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                      1096)
02/07/2005 11:25:02  ANR0409I Session 1103 ended for server CL_MSCS01_STA
                      (Windows). (SESSION: 1103)
02/07/2005 11:25:02  ANR8337I (Session: 1096, Origin: CL_MSCS01_STA)  LTO
                      volume 028AKK mounted in drive DRLTO_1 (mt0.0.0.4).
                      (SESSION: 1096)
```

Figure 7-50 The tape volume is mounted again by the Storage Agent

7. When the restore is completed we can see the final statistics in the schedule log file of the client for a successful operation as shown in Figure 7-51.

```
dsmsched_aix - Notepad
File  Edit  Format  Help
02/07/2005 11:34:33 Restoring         110,861 \\cl_mscs01\g$\code\isc\wps\migr
02/07/2005 11:34:33 Restoring           6,869 \\cl_mscs01\g$\code\isc\wps\migr
02/07/2005 11:34:33
Restore processing finished.
02/07/2005 11:34:35 --- SCHEDULEREC STATUS BEGIN
02/07/2005 11:34:35 Total number of objects restored:     1,884
02/07/2005 11:34:35 Total number of objects failed:           0
02/07/2005 11:34:35 Total number of bytes transferred:     1.28 GB
02/07/2005 11:34:35 LanFree data bytes:                    1.28 GB
02/07/2005 11:34:35 Data transfer time:                   40.94 sec
02/07/2005 11:34:35 Network data transfer rate:        32,818.81 KB/sec
02/07/2005 11:34:35 Aggregate data transfer rate:      8,477.99 KB/sec
02/07/2005 11:34:35 Elapsed processing time:           00:02:38
02/07/2005 11:34:35 --- SCHEDULEREC STATUS END
02/07/2005 11:34:35 --- SCHEDULEREC OBJECT END RESTORE 02/07/2005 11:16:05
02/07/2005 11:34:35 --- SCHEDULEREC STATUS BEGIN
02/07/2005 11:34:35 --- SCHEDULEREC STATUS END
02/07/2005 11:34:35 Scheduled event 'RESTORE' completed successfully.
02/07/2005 11:34:35 Sending results for scheduled event 'RESTORE'.
02/07/2005 11:34:35 Results sent to server for scheduled event 'RESTORE'.
```

Figure 7-51 Final statistics for the restore on the schedule log file

Attention: Notice that the restore process is started from the beginning. It is not restarted.

Results summary

The test results show that after a failure on the node that hosts the Tivoli Storage Manager client scheduler instance, a scheduled restore operation started on this node using the LAN-free path is started again from the beginning on the second node of the cluster when the service is online.

This is true if the startup window for the scheduled restore operation is not elapsed when the scheduler client is online again on the second node.

Also notice that the restore is not restarted from the point of failure, but started from the beginning. The scheduler queries the Tivoli Storage Manager server for a scheduled operation and a new session is opened for the client after the failover.

Restriction: Notice again that this configuration, with two Storage Agents in the same machine, is not technically supported by Tivoli Storage Manager for SAN. However, in our lab environment it worked. In other tests we made using the local Storage Agents for communication to the virtual client for LAN-free, the SCSI bus reset did not work and the restore process failed.

Archived

7.5 Storage Agent on Windows 2003 MSCS

In this section we describe how we configure our Storage Agent software to be capable of running in our MSCS Windows 2003, the same cluster we installed and configured in 4.4, “Windows 2003 MSCS installation and configuration” on page 44.

7.5.1 Windows 2003 lab setup

Refer to Table 4-4 on page 46, Table 4-5 on page 47, and Table 4-6 on page 47 for details of the cluster configuration: local nodes, virtual nodes, and cluster groups.

We use TSMSRV03, (an AIX machine), as the server because Tivoli Storage Manager Version 5.3.0 for AIX is, so far, the only platform that supports high availability Library Manager functions for LAN-free backup.

Tivoli Storage Manager LAN-free configuration details

Figure 7-52 shows the Storage Agent configuration we use in this chapter.

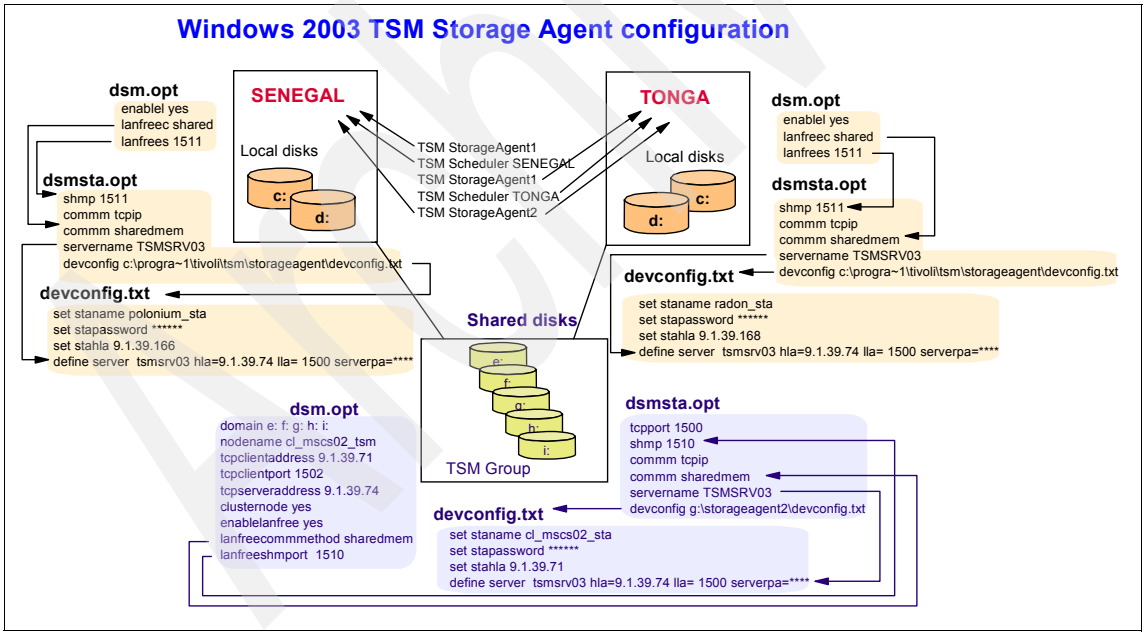

Figure 7-52 Windows 2003 Storage Agent configuration

Table 7-4 and Table 7-5 below give details about the client and server systems we use to install and configure the Storage Agent in our environment.

Table 7-4 Windows 2003 LAN-free configuration of our lab

Node 1	
TSM nodename	SENEGAL
Storage Agent name	SENEGAL_STA
Storage Agent service name	TSM StorageAgent1
dsmsta.opt and devconfig.txt location	c:\program files\tivoli\tsm\storageagent
Storage Agent high level address	9.1.39.166
Storage Agent low level address	1502
Storage Agent shared memory port	1511
LAN-free communication method	SharedMemory
Node 2	
TSM nodename	TONGA
Storage Agent name	TONGA_STA
Storage Agent service name	TSM StorageAgent1
dsmsta.opt and devconfig.txt location	c:\program files\tivoli\tsm\storageagent
Storage Agent high level address	9.1.39.168
Storage Agent low level address	1502
Storage Agent shared memory port	1511
LAN-free communication method	SharedMemory
Virtual node	
TSM nodename	CL_MSCS02_TSM
Storage Agent name	CL_MSCS02_STA
Storage Agent service name	TSM StorageAgent2
dsmsta.opt and devconfig.txt location	g:\storageagent2
Storage Agent high level address	9.1.39.71
Storage Agent low level address	1500
Storage Agent shared memory port	1510

Archived

Node 1	
TSM nodename	SENEGAL
Storage Agent name	SENEGAL_STA
Storage Agent service name	TSM StorageAgent1
dsmsta.opt and devconfig.txt location	c:\program files\tivoli\tsm\storageagent
Storage Agent high level address	9.1.39.166
Storage Agent low level address	1502
Storage Agent shared memory port	1511
LAN-free communication method	SharedMemory
Node 2	
TSM nodename	TONGA
Storage Agent name	TONGA_STA
Storage Agent service name	TSM StorageAgent1
dsmsta.opt and devconfig.txt location	c:\program files\tivoli\tsm\storageagent
Storage Agent high level address	9.1.39.168
Storage Agent low level address	1502
Storage Agent shared memory port	1511
LAN-free communication method	SharedMemory
Virtual node	
TSM nodename	CL_MSCS02_TSM
Storage Agent name	CL_MSCS02_STA
Storage Agent service name	TSM StorageAgent2
dsmsta.opt and devconfig.txt location	g:\storageagent2
Storage Agent high level address	9.1.39.71
Storage Agent low level address	1500
LAN-free communication method	SharedMemory

Archived

Table 7-5 Server information

Server information	
Servername	TSMSRV03
High level address	9.1.39.74
Low level address	1500
Server password for server-to-server communication	password

Our Storage Area Network devices are shown in Table 7-6.

Table 7-6 Storage devices used in the SAN

SAN devices	
Disk	IBM DS4500 Disk Storage Subsystem
Library	IBM LTO 3582 Tape Library
Tape drives	3580 Ultrium 2
Tape drive device name	drlto_1: mt0.0.0.2 drlto_2: mt1.0.0.2

Installing IBM 3580 tape drive drivers in Windows 2003

Before implementing the Storage Agent for LAN-free backup in our environment, we need to make sure that Windows 2003 OS in each node recognizes the tape drives that will be shared with the Tivoli Storage Manager server.

When we started our two servers, SENEGAL and TONGA, after connecting the devices, the IBM 3580 tape drives displayed as an interrogation mark under the **Other devices** icon. This happens because we need to install the appropriate IBM device drivers for 3580 LTO tape drives.

Once installed, the device drivers must be updated in each local node of the cluster using the Device Manager wizard.

We do not show in this section the whole installation process for the drivers. We only describe the main tasks to achieve this goal. For a detailed description of the tasks we follow refer to *IBM Ultrium device drivers Installation and User's Guide*.

To accomplish this requirement, we follow these steps:

1. We download the latest available IBM TotalStorage tape drivers from:

   http://www-1.ibm.com/servers/storage/support/allproducts/downloading.html

2. We open the device manager, right-click the tape drive, and choose **Update Driver** as shown in Figure 7-53. We follow the wizard process informing us of the path where the file was downloaded.

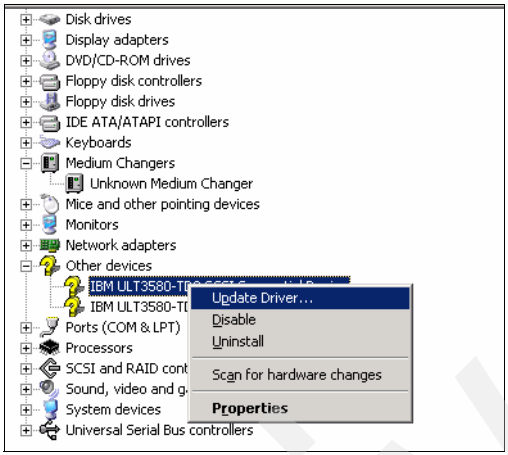

Figure 7-53 Tape devices in device manager page

3. After a successful installation, the drives are listed under Tape drives as shown in Figure 7-54.

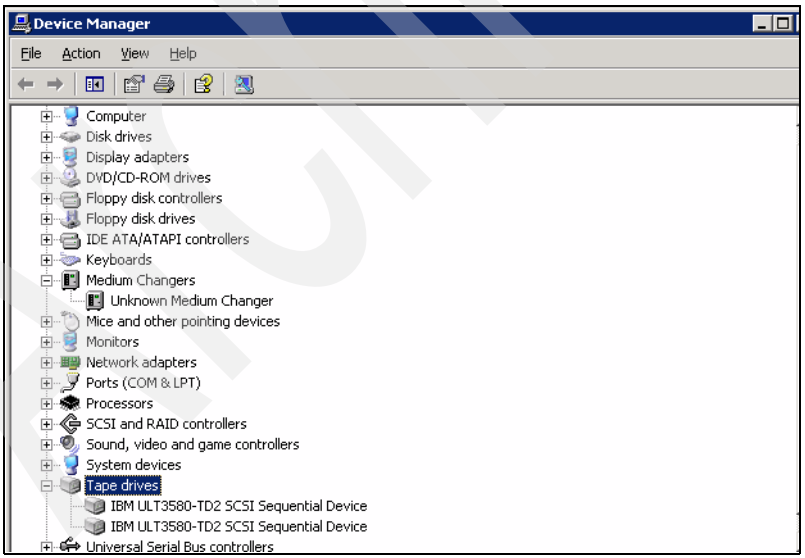

Figure 7-54 Device Manager page after updating the drivers

Archived

7.5.2 Configuration of the Storage Agent on Windows 2003 MSCS

The installation and configuration of the Storage Agent involves three steps:

1. Configuration of Tivoli Storage Manager server for LAN-free operation.
2. “Installation of the Storage Agent” on page 332.
3. Configuring the Storage Agent for local nodes.

Configuration of Tivoli Storage Manager server for LAN-free

The process of preparing a server for LAN-free data movement is very complex, involving several phases.

Each Storage Agent must be defined as a server in the Tivoli Storage Manager server. For our lab we define one Storage Agent for each local node and another one for the cluster node.

In 7.4.2, “Configuration of the Storage Agent on Windows 2000 MSCS” on page 339, we show how to set up server-to-server communications and path definitions using the new Administrative Center console. In this section we use instead the administrative command line interface.

1. Preparation of the server for enterprise management. We use the following commands:

```
set servername tsmsrv03
set serverpassword password
set serverhladress 9.1.39.74
set serverlladdress 1500
```

2. Definition of the Storage Agents as servers. We use the following commands:

```
define server senegal_sta serverpa=itsosj hla=9.1.39.166 lla=1500
define server tonga_sta serverpa=itsosj hla=9.1.39.168 lla=1500
define server cl_mscs02_sta serverpa=itsosj hla=9.1.39.71 lla=1500
```

3. Change of the node’s properties to allow either LAN or LAN-free movement of data:

```
update node senegal datawritepath=any datareadpath=any
update node tonga datawritepath=any datareadpath=any
update node cl_mscs02_tsm datawritepath=any datareadpath=any
```

4. Definition of tape library as shared (if this was not done when the library was first defined):

```
update library liblto shared=yes
```

5. Definition of paths from the Storage Agents to each tape drive in the Tivoli Storage Manager server. We use the following commands:

```
define path senegal_sta drlto_1 srctype=server desttype=drive
library=liblto device=mt0.0.0.2
```

```
define path senegal_sta drlto_2 srctype=server desttype=drive
library=liblto device=mt1.0.0.2
define path tonga_sta drlto_1 srctype=server desttype=drive library=liblto
device=mt0.0.0.2
define path tonga_sta drlto_2 srctype=server desttype=drive library=liblto
device=mt1.0.0.2
define path cl_mscs02_sta drlto_1 srctype=server desttype=drive
library=liblto device=mt0.0.0.2
define path cl_mscs02_sta drlto_2 srctype=server desttype=drive
library=liblto device=mt1.0.0.2
```

6. Definition of the storage pool for LAN-free backup:

```
define stgpool spt_bck lto pooltype=PRIMARY maxscratch=4
```

7. Definition/update of the policies to point to the storage pool above and activation of the policy set to refresh the changes. In our case we update the backup copygroup in the standard domain:

```
update copygroup standard standard standard type=backup dest=spt_bck
validate policyset standard standard
activate policyset standard standard
```

Configuring the Storage Agent for local nodes

As mentioned before, we set up three Storage Agents: one local for each node (SENEGAL_STA and TONGA_STA) and one for the TSM Group of the cluster (CL_MSCS02_STA).

The configuration process differs whether it is local or cluster. Here we describe the tasks we follow to configure the Storage Agent for local nodes.

Updating dsmsta.opt

Before we start configuring the Storage Agent we need to edit the dsmsta.opt file located in c:\program files\tivoli\tsm\storageagent.

We change the following line, to make sure it points to the whole path where the device configuration file is located:

```
DEVCONFIG C:\PROGRA~1\TIVOLI\TSM\STORAGEAGENT\DEVCONFIG.TXT
```

Figure 7-55 Modifying the devconfig option to point to devconfig file in dsmsta.opt

Note: We need to update the dsmsta.opt because the service used to start the Storage Agent uses as default the path where the command is run, not the installation path.

Using the management console to initialize the Storage Agent

To initialize the Storage Agent:

1. We open the Management Console (**Start** →**Programs** → **Tivoli Storage Manager** → **Management Console**) and we click **Next** on the welcome menu of the wizard.
2. We provide the Storage Agent information: name, password and TCP/IP address (high level address) as shown in Figure 7-56.

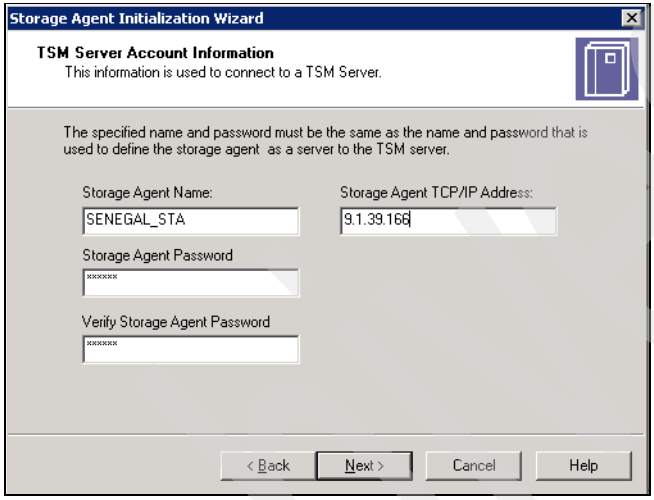

Figure 7-56 Specifying parameters for the Storage Agent

Important: We make sure that the Storage Agent name, and the rest of the information we provide in this menu, match the parameters used to define the Storage Agent in the Tivoli Storage Manager server in step 2 on page 383.

Archived

3. We provide all the server information: name, password, TCP/IP, and TCP port information as shown in Figure 7-57, and we click **Next**.

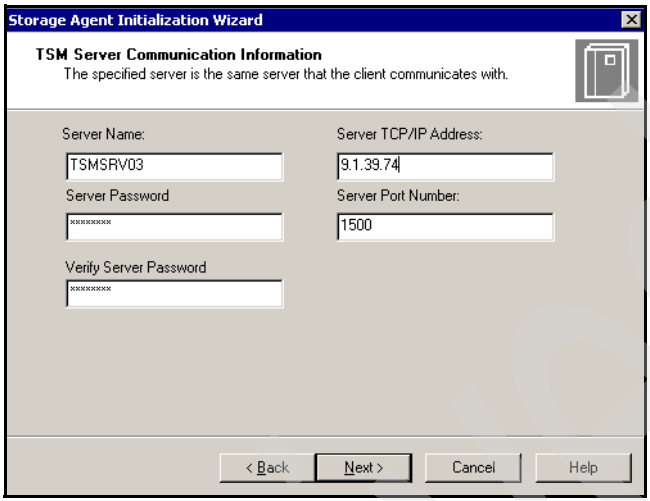

Figure 7-57 Specifying parameters for the Tivoli Storage Manager server

Important: The information provided in Figure 7-57 must match the information provided in the `set servername`, `set serverpassword`, **set serverhladdress** and `set serverlladdress` commands in the Tivoli Storage Manager server in step 1 on page 383.

4. We select the account that the service will use to start. We specify here the administrator account, but we could also have created a specific account to be used. This account should be in the administrator's group. We type the password and accept the service to start automatically when the server is started, we then click **Next** (Figure 7-58).

Archived

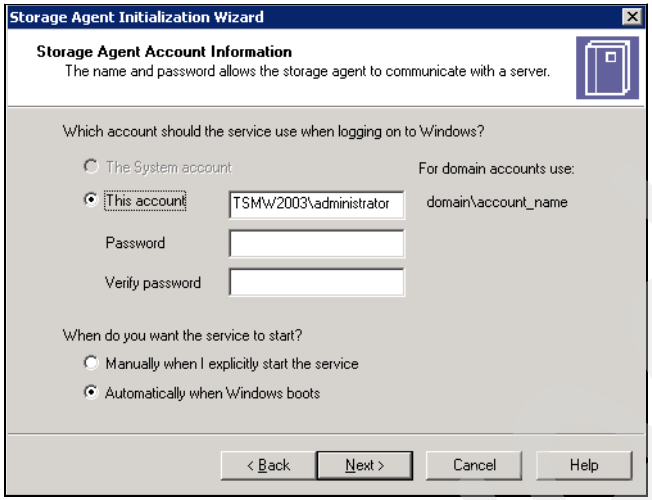

Figure 7-58 Specifying the account information

5. We click **Finish** when the wizard is complete.
6. We click **OK** on the message that says that the user has been granted rights to log on as a service.
7. The wizard finishes, informing us that the Storage Agent has been initialized. We click **OK** (Figure 7-59).

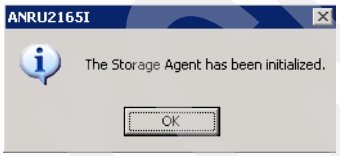

Figure 7-59 Storage agent initialized

8. The Management Console now displays the TSM StorageAgent1 service running, as shown in Figure 7-60.

Archived

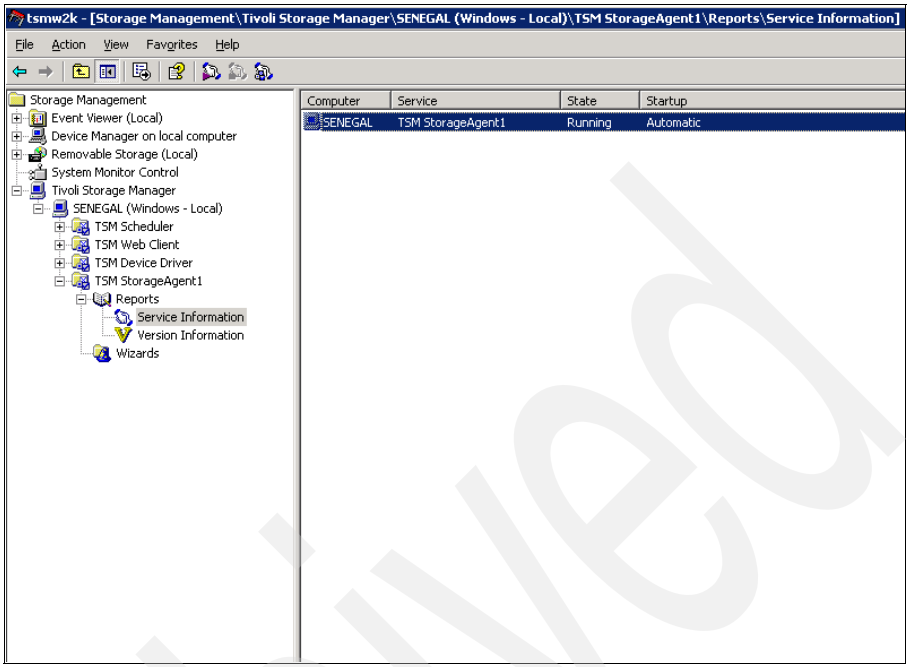

Figure 7-60 TSM StorageAgent1 is started

9. We repeat the same steps in the other server (TONGA).

This wizard can be re-run at any time if needed, from the Management Console, under **TSM StorageAgent1 →Wizards**.

Updating the client option file

To be capable of using LAN-free backup for each local node, we include the following options in the dsm.opt client file:

```
ENABLELANFREE yes
LANFREECOMMMETHOD sharedmem
LANFREESHMPORT 1511
```

We specify the 1511 port for Shared Memory instead of 1510 (the default), because we will use this default port to communicate with the Storage Agent associated to the cluster. Port 1511 will be used by the local nodes when communicating to the local Storage Agents.

Instead of the options specified above, we also can use:

```
ENABLELANFREE yes
LANFREECOMMMETHOD TCPIP
LANFREETCPPORT 1502
```

Archived

Restarting the Tivoli Storage Manager scheduler service

To use the LAN-free path, it is necessary, after including the LAN-free options in dsm.opt, to restart the Tivoli Storage Manager scheduler service. If we do not restart the service, the new options will not be read by the client.

Configuring Storage Agent for virtual nodes

In order to back up shared disk drives in the cluster using the LAN-free path, we can use the Storage Agent instance created for the local nodes. Depending upon the node that hosts the resources at that time, it will be used one local Storage Agent or another one.

This is the technically supported way of configuring LAN-free backup for clustered configurations. Each virtual node in the cluster should use the local Storage Agent in the local node that hosts the resource at that time.

However, in order to also have high-availability for the Storage Agent, we configure a new Storage Agent instance that will be used for the cluster.

> **Attention:** This is not a technically supported configuration but, in our lab tests, it worked.

In the following sections we describe the process for our TSM Group, where a TSM Scheduler generic service resource is located for backup of e: f: g: h: and i: shared disk drives.

Using the dsmsta setstorageserver utility

Instead of using the management console to create the new instance, we use the **`dsmsta`** utility from an MS-DOS prompt. The reason to use this tool is because we have to create a new registry key for this Storage Agent. If we start the management console we would use the default key, StorageAgent1, and we need a different one.

To achieve this goal, we perform these tasks:

1. We begin the configuration in the node that hosts the shared disk drives.
2. We copy the storageagent folder (created at installation time) from c:\program files\tivoli\tsm onto a shared disk drive (g:) with the name “storageagent2”.
3. We open a Windows MS-DOS prompt and change to g:\storageagent2.
4. We change the line *devconfig* in the dsmsta.opt file to point to g:\storageagent2\devconfig.txt.
5. From this path, we run the command we see in Figure 7-61 to create another instance for a Storage Agent called StorageAgent2. For this instance, the option (dsmsta.opt) and device configuration (devconfig.txt) files will be located on this path.

```
Command Prompt

G:\storageagent2>dsmsta -k StorageAgent2 setstorageserver myname=CL_MSCS02_STA m
ypa=itsosj myhla=9.1.39.71 servername=TSMSRV03 serverpassword=password hla=9.1.3
9.74 lla=1500
ANR0900I Processing options file G:\storageagent2\dsmsta.opt.
ANR7800I DSMSERV generated at 12:45:40 on Jan 28 2005.

Tivoli Storage Manager for Windows
Version 5, Release 3, Level 0.1

*****************************************************************
* Patch 5.3.0.1 TSM Server (Driver 0000, Build 05028.1245) *
*                                                               *
*****************************************************************
Licensed Materials - Property of IBM

(C) Copyright IBM Corporation 1990, 2004.
All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

ANR4726I The ICC support module has been loaded.
ANR1432I Updating device configuration information to defined files.
ANR1433I Device configuration information successfully written to g:\storageag-
ent2\devconfig.txt.
ANR2119I The SERVERNAME option has been changed in the options file.
ANR0467I The SETSTORAGESERVER command completed successfully.

G:\storageagent2>_
```

Figure 7-61 Installing Storage Agent for LAN-free backup of shared disk drives

Attention: Notice, in Figure 7-61, the new registry key that is used for this Storage Agent, StorageAgent2, as well as the name and IP address specified in the “myname” and “myhla” parameters. The Storage Agent name is CL_MSCS02_STA, and its IP address is the IP address of the TSM Group. Also notice that, when executing the command from g:\storageagent2, we make sure that the dsmsta.opt and devconfig.txt updated files are the ones in this path.

6. Now, from the same path, we run a command to install a service called “TSM StorageAgent2” related to the StorageAgent2 instance created in step 5. The command and the result of its execution is shown in Figure 7-62.

```
C:\WINDOWS\system32\cmd.exe

G:\storageagent2>install "TSM StorageAgent2" g:\storageagent2\dstasvc.exe tsmsto
rageagent itsosj
Service installed

G:\storageagent2>_
```

Figure 7-62 Installing the service attached to StorageAgent2

7. If we open the Tivoli Storage Manager management console in this node, we now can see two instances for two Storage Agents: the one we created for the local node, TSM StorageAgent1, and a new one, TSM Storage Agent2, which is set to *Manual*. This last instance is stopped, as we can see in Figure 7-63.

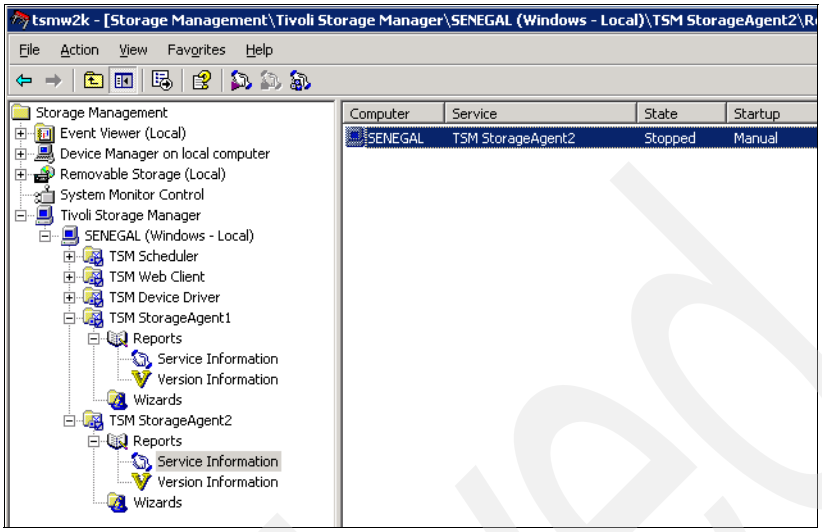

Figure 7-63 Management console displays two Storage Agents

8. We start the TSM StorageAgent2 instance right-clicking and selecting **Start** as we show in Figure 7-64.

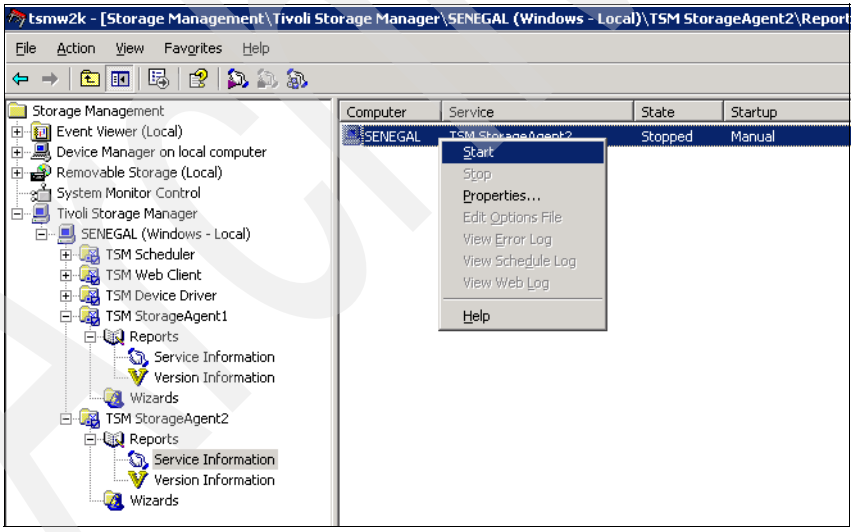

Figure 7-64 Starting the TSM StorageAgent2 service in SENEGAL

Archived

9. Now we have two Storage Agent instances running in SENEGAL:
   - TSM StorageAgent1: related to the local node and using the dsmsta.opt and devconfig.txt files located in c:\program files\tivoli\tsm\storageagent
   - TSM StorageAgent2: related to the virtual node and using the dsmsta.opt and devconfig.txt files located in g:\storageagent2.
10. We stop the TSM StorageAgent2 and move the resources to TONGA.
11. In TONGA, we follow steps 3 to 6. After that, we open the Tivoli Storage Manager management console and we again find two Storage Agent instances: TSM StorageAgent1 (for the local node) and TSM StorageAgent2 (for the virtual node). This last instance is stopped and set to manual as shown in Figure 7-65.

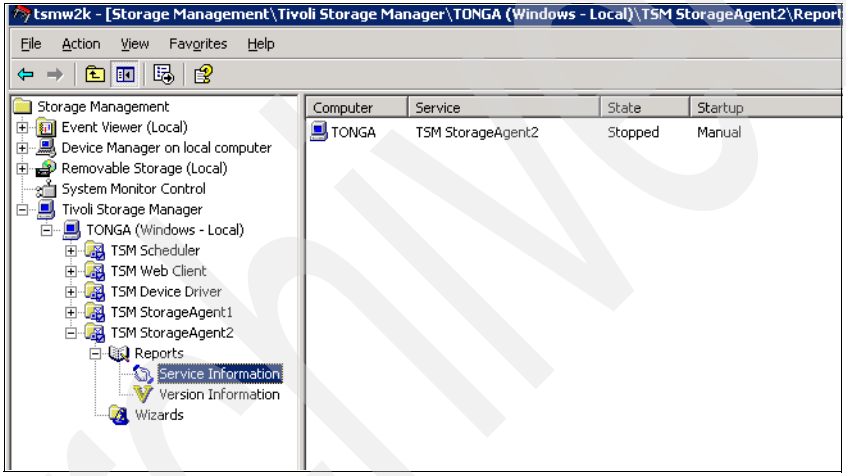

Figure 7-65 TSM StorageAgent2 installed in TONGA

12. We start the instance right-clicking and selecting **Start**. After a successful start, we stop it again.
13. Finally, the last task consists of the definition of TSM StorageAgent2 service as a cluster resource. To do this we open the Cluster Administrator menu, we right-click the resource group where Tivoli Storage Manager scheduler service is defined, TSM Group, and select to define a new resource as shown in Figure 7-66.

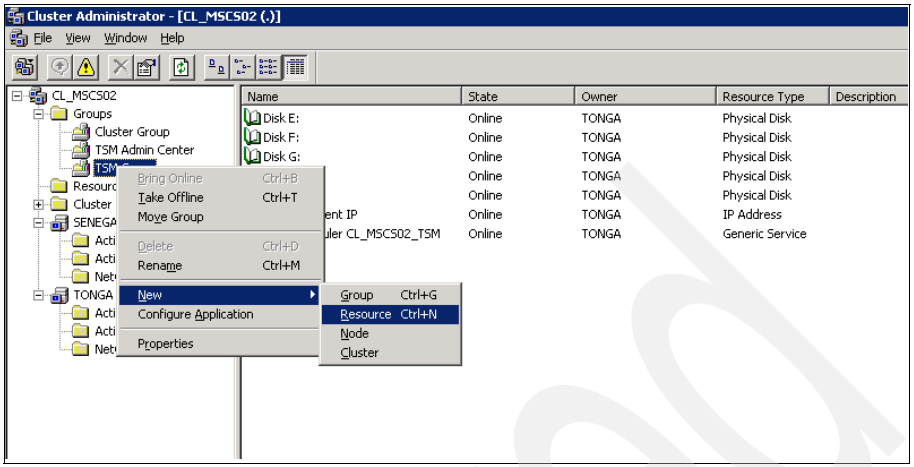

Figure 7-66 Use cluster administrator to create a resource: TSM StorageAgent2

14. We type a name for the resource and select Generic Service as the resource type and click **Next** as we see in Figure 7-67.

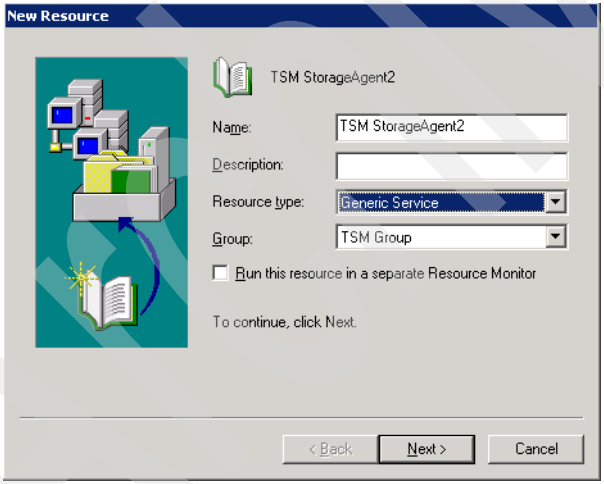

Figure 7-67 Defining a generic service resource for TSM StorageAgent2

15. We leave both nodes as possible owners and click **Next** in Figure 7-68.

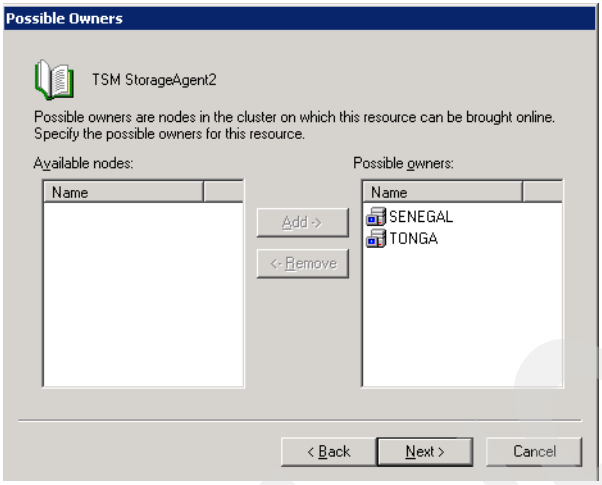

Figure 7-68 Possible owners for TSM StorageAgent2

16. As TSM StorageAgent2 dependencies, we select **Disk G:** which is where the configuration files are located for this instance. We click **Next** in Figure 7-69.

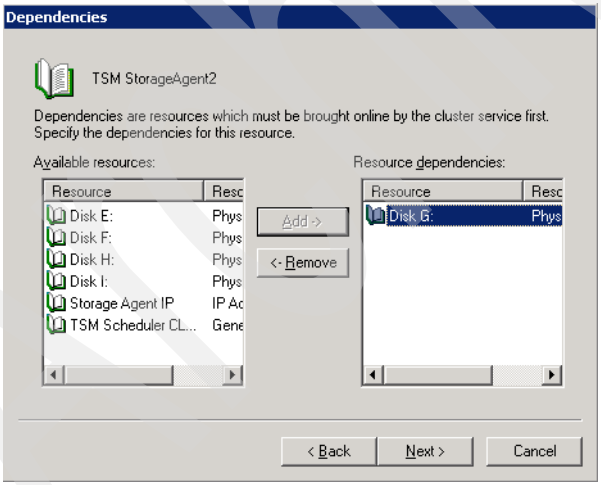

Figure 7-69 Dependencies for TSM StorageAgent2

17. We type the name of the service, **TSM StorageAgent2**. We click **Next** in Figure 7-70.

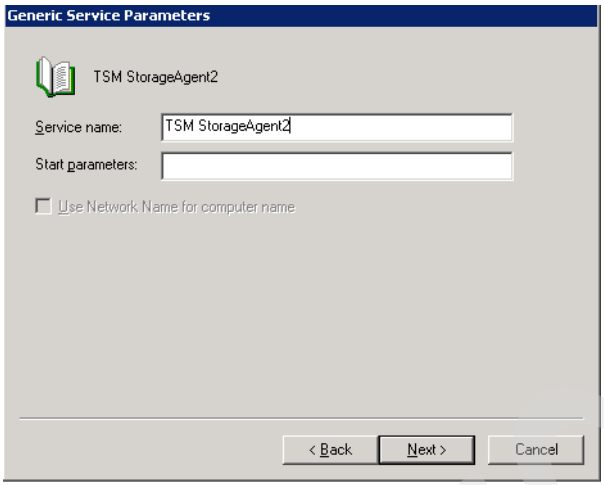

Figure 7-70 Service name for TSM StorageAgent2

Important: The name of the service in Figure 7-70 must match the name we used to install the instance in both nodes.

18. We do not use any registry key replication for this resource. We click **Finish** in Figure 7-71.

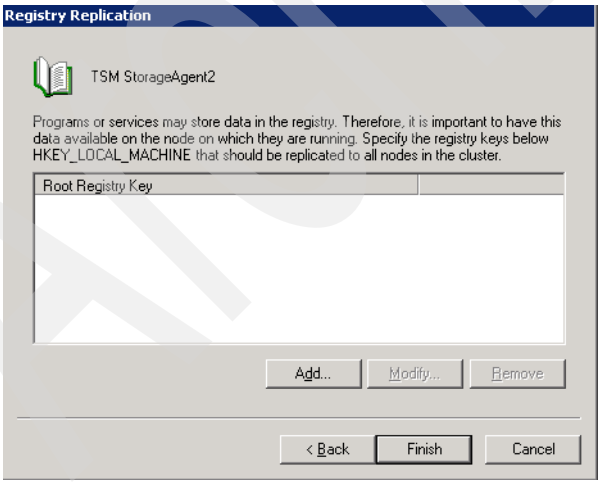

Figure 7-71 Registry key for TSM StorageAgent2

19. The new resource is successfully created as Figure 7-72 displays. We click **OK**.

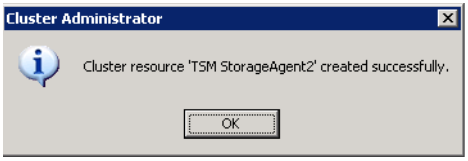

Figure 7-72 Generic service resource created successfully: TSM StorageAgent2

20. The last task is bringing online the new resource, as we show in Figure 7-73.

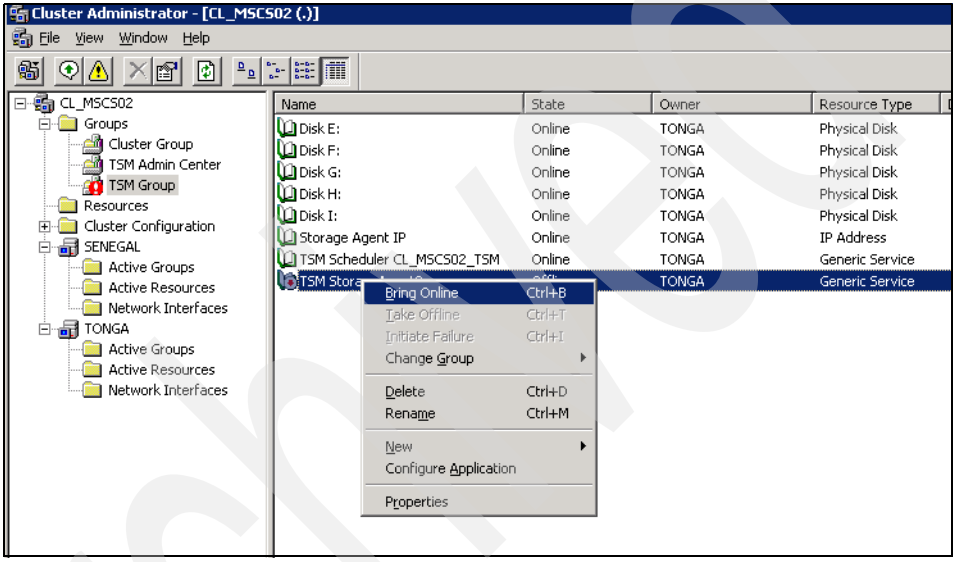

Figure 7-73 Bringing the TSM StorageAgent2 resource online

21. At this time the service is started in the node that hosts the resource group. To check the successful implementation of this Storage Agent, we move the resources to the second node and we check that TSM StorageAgent2 is now started in this second node and stopped in the first one.

> **Important:** Be sure to use only the Cluster Administrator to start and stop the StorageAgent2 instance at any time.

Changing the dependencies for the TSM Scheduler resource

Since we want the Tivoli Storage Manager scheduler always to use the LAN-free path when it starts, it is necessary to update its associated resource in the Cluster Administrator to add TSM StorageAgent2 as a dependency to bring it online.

For this reason, we open the Cluster Administrator menu, select the TSM Scheduler resource for CL_MSCS02_TSM and go to **Properties** → **Dependencies** → **Modify**. Once there, we add TSM StorageAgent2 as a dependency such as we show in Figure 7-74.

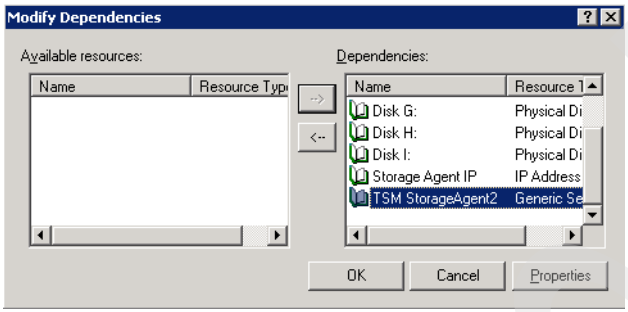

Figure 7-74 Adding Storage Agent resource as dependency for TSM Scheduler

We click **OK** and bring the resource online again. With this dependency we make sure the Tivoli Storage Manager scheduler is not started for this cluster group before the Storage Agent does.

Updating the client option file

To be capable of using LAN-free backup for the virtual node, we must specify certain special options in the client option file for the virtual node.

We open g:\tsm\dsm.opt and we include the following options:

```
ENABLELANFREE yes
LANFREECOMMMETHOD SHAREDMEM
LANFREESHMPORT 1510
```

For the virtual node we use the default shared memory port, 1510.

Instead of the options above we also can use:

```
ENABLELANFREE yes
LANFREECOMMMETHOD TCPIP
LANFREETCPPORT 1500
```

Restarting the Tivoli Storage Manager scheduler service

After including the LAN-free options in dsm.opt, we restart the Tivoli Storage Manager scheduler service for the TSM Group using the Cluster Administrator. If we do not restart the service, the new options will not be read by the client.

Archived

7.5.3 Testing the Storage Agent high availability

The purpose of this section is to test our LAN-free setup for the clustering.

We use the TSM Group (nodename CL_MSCS02_TSM) to test LAN-free backup/restore of shared data in our Windows 2003 cluster.

Our objective with these tasks is to know how the Storage Agent and the Tivoli Storage Manager Library Manager work together to respond, on a client clustered environment, after certain kinds of failures that affect the shared resources.

Again, for details of our LAN-free configuration, refer to Table 7-4 on page 379 and Table 7-5 on page 381.

Testing LAN-free client incremental backup on Windows 2003

In this section we test LAN-free incremental backup.

Objective

The objective of this test is to show what happens when a LAN-free client incremental backup is started for a virtual node in the cluster using the Storage Agent created for this group (CL_MSCS02_STA), and the node that hosts the resources at that moment suddenly fails.

Activities

To do this test, we perform these tasks:

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager scheduler service for TSM Group. At this time SENEGAL does.
2. We schedule a client incremental backup operation using the Tivoli Storage Manager server scheduler and we associate the schedule to CL_MSCS02_TSM nodename.
3. We make sure that TSM StorageAgent2 and TSM Scheduler for CL_MSCS02_TSM are online resources on SENEGAL.
4. When it is the scheduled time, a client session for CL_MSCS02_TSM nodename starts on the server. At the same time, several sessions are also started for CL_MSCS02_STA for Tape Library Sharing and the Storage Agent prompts the Tivoli Storage Manager server to mount a tape volume. The tape volume is mounted in drive DRLTO_2 as we can see in Figure 7-75:

Archived

```
IBM Tivoli Storage Manager
                        (SESSION: 21)
02/17/2005 10:20:12  ANR8379I Mount point in device class LTO is waiting for
                      the volume mount to complete, status: WAITING FOR VOLUME.
                      (SESSION: 21)
02/17/2005 10:20:12  ANR8334I          1 matches found. (SESSION: 21)
02/17/2005 10:20:13  ANR2017I Administrator ADMIN issued command: QUERY SESSION
                        (SESSION: 21)
02/17/2005 10:20:35  ANR8337I LTO volume 030AKK mounted in drive DRLTO_2
                      (/dev/rmt1). (SESSION: 90)
02/17/2005 10:20:35  ANR0409I Session 90 ended for server CL_MSCS02_STA
                      (Windows). (SESSION: 90)
02/17/2005 10:20:36  ANR0511I Session 88 opened output volume 030AKK. (SESSION:
                      88)
02/17/2005 10:20:46  ANR2017I Administrator ADMIN issued command: QUERY SESSION
                        (SESSION: 21)
more...   (<ENTER> to continue, 'C' to cancel)

02/17/2005 10:20:47  ANR2017I Administrator ADMIN issued command: QUERY MO
                      (SESSION: 21)
02/17/2005 10:20:47  ANR8330I LTO volume 030AKK is mounted R/W in drive DRLTO_2
                      (/dev/rmt1), status: IN USE. (SESSION: 21)
```

Figure 7-75 Storage agent CL_MSCS02_STA mounts tape for LAN-free backup

5. The client, by means of the Storage Agent, starts sending files to the drive using the SAN path as we see on its schedule log file in Figure 7-76.

```
dsmsched.log - Notepad
File  Edit  Format  View  Help
Executing scheduled command now.
02/17/2005 10:20:45 --- SCHEDULEREC OBJECT BEGIN INCR_BACKUP 02/17/2005 10:15:28
02/17/2005 10:20:45 Incremental backup of volume 'g:\code\*'
02/17/2005 10:20:45 Directory-->                  0 \\cl_mscs02\g$\code [Sent]
02/17/2005 10:20:45 Directory-->                  0 \\cl_mscs02\g$\code\adminc [Se
02/17/2005 10:20:45 Directory-->                  0 \\cl_mscs02\g$\code\drivers_lt
02/17/2005 10:20:45 Directory-->                  0 \\cl_mscs02\g$\code\isc [Sent]
02/17/2005 10:20:45 Directory-->                  0 \\cl_mscs02\g$\code\lto2k3 [Se
02/17/2005 10:20:45 Directory-->                  0 \\cl_mscs02\g$\code\storageage
02/17/2005 10:20:45 Normal File-->        1,862,528 \\cl_mscs02\g$\code\IBMTape.W3
02/17/2005 10:20:45 ANS1114I Waiting for mount of offline media.
02/17/2005 10:20:46 Retry # 1  Directory-->                  0 \\cl_mscs02\g$\code
02/17/2005 10:20:46 Retry # 1  Directory-->                  0 \\cl_mscs02\g$\code
02/17/2005 10:20:46 Retry # 1  Directory-->                  0 \\cl_mscs02\g$\code
02/17/2005 10:21:44 Retry # 1  Directory-->                  0 \\cl_mscs02\g$\code
02/17/2005 10:21:44 Retry # 1  Directory-->                  0 \\cl_mscs02\g$\code
02/17/2005 10:21:44 Retry # 1  Directory-->                  0 \\cl_mscs02\g$\code
02/17/2005 10:21:44 Retry # 1  Normal File-->        1,862,528 \\cl_mscs02\g$\code
02/17/2005 10:21:51 Normal File-->       25,388,669 \\cl_mscs02\g$\code\TSM530C_WI
02/17/2005 10:21:53 Normal File-->       10,762,389 \\cl_mscs02\g$\code\TSMCON5300
02/17/2005 10:22:06 Normal File-->      181,664,860 \\cl_mscs02\g$\code\TSMSRV5300
02/17/2005 10:22:10 Normal File-->       29,205,340 \\cl_mscs02\g$\code\TSMSTA5300
02/17/2005 10:22:14 Normal File-->       29,065,149 \\cl_mscs02\g$\code\TSMSTA5301
02/17/2005 10:22:14 Directory-->                  0 \\cl_mscs02\g$\code\adminc\jac
```

Figure 7-76 Client starts sending files to the TSM server in the schedule log file

6. While the client continues sending files to the server, we force SENEGAL to fail. The following sequence takes place:

   a. The client and also the Storage Agent lose their connections with the server temporarily, and both sessions are terminated as we can see on the Tivoli Storage Manager server activity log shown in Figure 7-77.

```
IBM Tivoli Storage Manager
02/17/2005 10:20:47  ANR8334I          1 matches found. (SESSION: 21)
02/17/2005 10:21:33  ANR0409I Session 85 ended for server CL_MSCS02_STA
                      (Windows). (SESSION: 85)
02/17/2005 10:21:33  ANR0409I Session 82 ended for server CL_MSCS02_STA
                      (Windows). (SESSION: 82)
02/17/2005 10:21:34  ANR0480W Session 87 for node CL_MSCS02_TSM (WinNT)
                      terminated - connection with client severed. (SESSION:
                      87)
02/17/2005 10:21:35  ANR0480W Session 88 for node CL_MSCS02_TSM (WinNT)
                      terminated - connection with client severed. (SESSION:
                      88)
```

Figure 7-77 Sessions for TSM client and Storage Agent are lost in the activity log

   b. We can also see that the connection is lost on the schedule log client file in Figure 7-78.

```
dsmsched.log - Notepad
File  Edit  Format  View  Help
02/17/2005 10:22:34 Normal File-->        6,484,490 \\cl_mscs02\g$\code\adminc\ISC
02/17/2005 10:22:34 Normal File-->           79,853 \\cl_mscs02\g$\code\adminc\LIC
02/17/2005 10:22:34 Normal File-->           13,661 \\cl_mscs02\g$\code\adminc\REA
02/17/2005 10:22:34 Normal File-->           23,408 \\cl_mscs02\g$\code\adminc\REA
02/17/2005 10:22:34 ANS2820E An interrupt has occurred. The current operation will
client will shut down.
```

Figure 7-78 Connection is lost in the client while the backup is running

   c. In the Cluster Administrator menu SENEGAL is not in the cluster and TONGA begins to bring the resources online.

   d. The tape volume is still mounted on the same drive.

   e. After a while the resources are online on TONGA.

   f. When the Storage Agent CL_MSCS02_STA is again online (in TONGA), the TSM Scheduler service also is started (because of the dependency between these two resources). We can see this on the activity log in Figure 7-79.

```
IBM Tivoli Storage Manager
02/17/2005 10:22:02  ANR0406I Session 92 started for node CL_MSCS02_TSM (WinNT)
                      (Tcp/Ip 9.1.39.168(3983)). (SESSION: 92)
02/17/2005 10:22:02  ANR1639I Attributes changed for node CL_MSCS02_TSM: TCP
                      Name from SENEGAL to TONGA, TCP Address from 9.1.39.166
                      to 9.1.39.168, GUID from 46.6e.f4.d0.75.3a.11.d9.bc.7b.0-
                      0.0d.60.49.41.dd to 18.72.62.50.75.3b.11.d9.8f.b7.00.a0.-
                      cc.d2.58.ac. (SESSION: 92)
02/17/2005 10:22:02  ANR0403I Session 92 ended for node CL_MSCS02_TSM (WinNT).
                      (SESSION: 92)
02/17/2005 10:22:03  ANR0408I Session 93 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for storage agent.  (SESSION: 93)
02/17/2005 10:22:03  ANR0408I Session 94 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 94)
02/17/2005 10:22:03  ANR0408I Session 95 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for event logging.  (SESSION: 95)
02/17/2005 10:22:03  ANR0409I Session 94 ended for server CL_MSCS02_STA
                      (Windows). (SESSION: 94)
02/17/2005 10:22:03  ANR0408I Session 97 started for server CL_MSCS02_STA
more...   <<ENTER> to continue, 'C' to cancel>

                      (Windows) (Tcp/Ip) for storage agent.  (SESSION: 97)
02/17/2005 10:22:05  ANR0406I Session 98 started for node CL_MSCS02_TSM (WinNT)
                      (Tcp/Ip 9.1.39.168(3988)). (SESSION: 98)
02/17/2005 10:22:05  ANR1639I Attributes changed for node CL_MSCS02_TSM: TCP
                      Address from 9.1.39.168 to . (SESSION: 98)
02/17/2005 10:22:05  ANR0408I Session 99 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for storage agent.  (SESSION: 99)
02/17/2005 10:22:05  ANR0415I Session 99 proxied by CL_MSCS02_STA started for
                      node CL_MSCS02_TSM. (SESSION: 99)
02/17/2005 10:22:05  ANR0408I Session 100 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 100)
02/17/2005 10:22:05  ANR0409I Session 100 ended for server CL_MSCS02_STA
                      (Windows). (SESSION: 100)
02/17/2005 10:22:06  ANR0408I Session 101 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 101)
02/17/2005 10:22:08  ANR0408I Session 96 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 94)
02/17/2005 10:22:08  ANR0409I Session 96 ended for server CL_MSCS02_STA
                      (Windows). (SESSION: 94)
02/17/2005 10:22:08  ANR2017I Administrator ADMIN issued command: QUERY EVENT *
                      *  (SESSION: 21)
02/17/2005 10:22:26  ANR8336I Verifying label of LTO volume 030AKK in drive
                      DRLTO_2 (/dev/rmt1). (SESSION: 90)
more...   <<ENTER> to continue, 'C' to cancel>
```

Figure 7-79 Both Storage Agent and TSM client restart sessions in second node

g. The Tivoli Storage Manager server resets the SCSI bus, dismounting the tape volume from one drive and it mounts the tape volume on the other drive for the Storage Agent CL_MSCS02_STA to use as we can see in Figure 7-80.

```
IBM Tivoli Storage Manager
                      *  (SESSION: 21)
02/17/2005 10:22:26  ANR8336I Verifying label of LTO volume 030AKK in drive
                      DRLTO_2 (/dev/rmt1). (SESSION: 90)
more...   <<ENTER> to continue, 'C' to cancel>

02/17/2005 10:22:57  ANR2017I Administrator ADMIN issued command: QUERY MO
                      (SESSION: 21)
02/17/2005 10:22:57  ANR8379I Mount point in device class LTO is waiting for
                      the volume mount to complete, status: WAITING FOR VOLUME.
                      (SESSION: 21)
02/17/2005 10:22:57  ANR8331I LTO volume 030AKK is mounted R/W in drive DRLTO_2
                      (/dev/rmt1), status: DISMOUNTING. (SESSION: 21)
02/17/2005 10:22:57  ANR8334I          2 matches found. (SESSION: 21)
02/17/2005 10:23:03  ANR8468I LTO volume 030AKK dismounted from drive DRLTO_2
                      (/dev/rmt1) in library LIBLTO. (SESSION: 90)
02/17/2005 10:23:29  ANR8337I LTO volume 030AKK mounted in drive DRLTO_1
                      (/dev/rmt0). (SESSION: 101)
02/17/2005 10:23:30  ANR0409I Session 101 ended for server CL_MSCS02_STA
                      (Windows). (SESSION: 101)
02/17/2005 10:23:30  ANR0511I Session 99 opened output volume 030AKK. (SESSION:
                      99)
02/17/2005 10:24:46  ANR2017I Administrator ADMIN issued command: QUERY MO
                      (SESSION: 21)
02/17/2005 10:24:46  ANR8330I LTO volume 030AKK is mounted R/W in drive DRLTO_1
                      (/dev/rmt0), status: IN USE. (SESSION: 21)
02/17/2005 10:24:46  ANR8334I          1 matches found. (SESSION: 21)
```

Figure 7-80 Tape volume is dismounted and mounted again by the server

h. The client restarts its scheduled incremental backup using the SAN path as we can see on the schedule log file in Figure 7-81.

```
dsmsched.log - Notepad
File  Edit  Format  View  Help
02/17/2005 10:22:34 ANS2820E An interrupt has occurred. The current operation will
client will shut down.
02/17/2005 10:23:06 Querying server for next scheduled event.
02/17/2005 10:23:06 Node Name: CL_MSCS02_TSM
02/17/2005 10:23:06 Session established with server TSMSRV03: AIX-RS/6000
02/17/2005 10:23:06   Server Version 5, Release 3, Level 0.0
02/17/2005 10:23:06   Server date/time: 02/17/2005 10:22:05  Last access: 02/17/200

02/17/2005 10:23:06 --- SCHEDULEREC QUERY BEGIN
02/17/2005 10:23:06 --- SCHEDULEREC QUERY END
02/17/2005 10:23:06 Next operation scheduled:
02/17/2005 10:23:06 ------------------------------------------------------------
02/17/2005 10:23:06 Schedule Name:         INCR_BACKUP
02/17/2005 10:23:06 Action:                Incremental
02/17/2005 10:23:06 Objects:               g:\code\*
02/17/2005 10:23:06 Options:               -subdir=yes
02/17/2005 10:23:06 Server Window Start:   10:15:28 on 02/17/2005
02/17/2005 10:23:06 ------------------------------------------------------------
02/17/2005 10:23:06
Executing scheduled command now.
02/17/2005 10:23:06 --- SCHEDULEREC OBJECT BEGIN INCR_BACKUP 02/17/2005 10:15:28
02/17/2005 10:23:06 Incremental backup of volume 'g:\code\*'
02/17/2005 10:23:06 Normal File-->         6,484,490 \\cl_mscs02\g$\code\adminc\ISC
02/17/2005 10:23:06 Normal File-->            79,853 \\cl_mscs02\g$\code\adminc\LIC
02/17/2005 10:23:06 Normal File-->            13,661 \\cl_mscs02\g$\code\adminc\REA
02/17/2005 10:23:06 Normal File-->            23,408 \\cl_mscs02\g$\code\adminc\REA
02/17/2005 10:23:06 ANS1114I Waiting for mount of offline media.
02/17/2005 10:24:49 Retry # 1  Normal File-->         6,484,490 \\cl_mscs02\g$\code
02/17/2005 10:24:49 Retry # 1  Normal File-->            79,853 \\cl_mscs02\g$\code
02/17/2005 10:24:49 Retry # 1  Normal File-->            13,661 \\cl_mscs02\g$\code
02/17/2005 10:24:49 Retry # 1  Normal File-->            23,408 \\cl_mscs02\g$\code
02/17/2005 10:24:56 Normal File-->        20,574,949 \\cl_mscs02\g$\code\adminc\set
02/17/2005 10:25:00 Normal File-->        35,352,876 \\cl_mscs02\g$\code\adminc\set
02/17/2005 10:25:04 Normal File-->        24,797,258 \\cl_mscs02\g$\code\adminc\set
02/17/2005 10:25:08 Normal File-->        28,124,203 \\cl_mscs02\g$\code\adminc\set
```

Figure 7-81 The scheduled is restarted and the tape volume mounted again

Archived

7. The incremental backup ends successfully as we can see on the final statistics recorded by the client in its schedule log file in Figure 7-82.

```
dsmsched.log - Notepad
File  Edit  Format  View  Help
02/17/2005 10:26:30 Normal File-->            8,032 \\cl_mscs02\g$\code\lto2k3\mag
02/17/2005 10:26:30 Normal File-->          233,472 \\cl_mscs02\g$\code\lto2k3\ntu
02/17/2005 10:26:30 Normal File-->          143,360 \\cl_mscs02\g$\code\lto2k3\uni
02/17/2005 10:26:30 Normal File-->          527,360 \\cl_mscs02\g$\code\lto2k3\che
02/17/2005 10:26:30 Normal File-->          279,552 \\cl_mscs02\g$\code\lto2k3\che
02/17/2005 10:26:30 Normal File-->          543,744 \\cl_mscs02\g$\code\lto2k3\che
02/17/2005 10:26:30 Normal File-->          282,112 \\cl_mscs02\g$\code\lto2k3\che
02/17/2005 10:26:30 Normal File-->           46,592 \\cl_mscs02\g$\code\lto2k3\che
02/17/2005 10:26:30 Normal File-->          510,976 \\cl_mscs02\g$\code\lto2k3\che
02/17/2005 10:26:30 Normal File-->          282,368 \\cl_mscs02\g$\code\lto2k3\che
02/17/2005 10:26:30 Normal File-->          527,360 \\cl_mscs02\g$\code\lto2k3\che
02/17/2005 10:26:30 Normal File-->          285,696 \\cl_mscs02\g$\code\lto2k3\che
02/17/2005 10:26:38 Normal File-->      111,755,569 \\cl_mscs02\g$\code\storageage
02/17/2005 10:26:38 Successful incremental backup of '\\cl_mscs02\g$\code\*'

02/17/2005 10:26:38 --- SCHEDULEREC STATUS BEGIN
02/17/2005 10:26:38 Total number of objects inspected:    1,928
02/17/2005 10:26:38 Total number of objects backed up:    1,907
02/17/2005 10:26:38 Total number of objects updated:          0
02/17/2005 10:26:38 Total number of objects rebound:          0
02/17/2005 10:26:38 Total number of objects deleted:          0
02/17/2005 10:26:38 Total number of objects expired:          0
02/17/2005 10:26:38 Total number of objects failed:           0
02/17/2005 10:26:38 Total number of bytes transferred:    891.74 MB
02/17/2005 10:26:38 LanFree data bytes:                   891.62 MB
02/17/2005 10:26:38 Data transfer time:                   103.13 sec
02/17/2005 10:26:38 Network data transfer rate:        8,854.20 KB/sec
02/17/2005 10:26:38 Aggregate data transfer rate:      4,300.56 KB/sec
02/17/2005 10:26:38 Objects compressed by:                    0%
02/17/2005 10:26:38 Elapsed processing time:           00:03:32
02/17/2005 10:26:38 --- SCHEDULEREC STATUS END
02/17/2005 10:26:38 --- SCHEDULEREC OBJECT END INCR_BACKUP 02/17/2005 10:15:28
02/17/2005 10:26:38 Scheduled event 'INCR_BACKUP' completed successfully.
02/17/2005 10:26:38 Sending results for scheduled event 'INCR_BACKUP'.
02/17/2005 10:26:38 Results sent to server for scheduled event 'INCR_BACKUP'.
```

Figure 7-82 Final statistics for LAN-free backup

Archived

8. In the activity log there are messages reporting the end of the LAN-free backup, and the tape volume is correctly dismounted by the server. We see all these events in Figure 7-83.

```
IBM Tivoli Storage Manager
                        (SESSION: 21)
02/17/2005 10:25:38  ANR2507I Schedule INCR_BACKUP for domain STANDARD started
                      at 02/17/05 10:15:28 for node CL_MSCS02_TSM completed
                      successfully at 02/17/05 10:25:38. (SESSION: 98)
02/17/2005 10:25:38  ANR0403I Session 98 ended for node CL_MSCS02_TSM (WinNT).
                      (SESSION: 98)
02/17/2005 10:25:38  ANR0408I Session 102 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 102)
02/17/2005 10:25:38  ANR0514I Session 99 closed volume 030AKK. (SESSION: 99)
02/17/2005 10:25:38  ANR0409I Session 102 ended for server CL_MSCS02_STA
                      (Windows). (SESSION: 102)
02/17/2005 10:25:38  ANR0408I Session 103 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 103)
02/17/2005 10:25:38  ANR0409I Session 103 ended for server CL_MSCS02_STA
                      (Windows). (SESSION: 103)
more...   <<ENTER> to continue, 'C' to cancel>

02/17/2005 10:25:38  ANR0406I Session 104 started for node CL_MSCS02_TSM
                      (WinNT) (Tcp/Ip 9.1.39.168(4042)). (SESSION: 104)
02/17/2005 10:25:38  ANR0403I Session 99 ended for node CL_MSCS02_TSM (WinNT).
                      (SESSION: 99)
02/17/2005 10:25:38  ANR0403I Session 104 ended for node CL_MSCS02_TSM (WinNT).
                      (SESSION: 104)
02/17/2005 10:26:38  ANR8325I Dismounting volume 030AKK - 1 minute mount
                      retention expired.
02/17/2005 10:26:42  ANR0408I Session 105 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 105)
02/17/2005 10:27:27  ANR0409I Session 105 ended for server CL_MSCS02_STA
                      (Windows). (SESSION: 105)
02/17/2005 10:27:35  ANR8336I Verifying label of LTO volume 030AKK in drive
                      DRLTO_1 (/dev/rmt0). (SESSION: 101)
02/17/2005 10:27:44  ANR2017I Administrator ADMIN issued command: QUERY MO
                      (SESSION: 21)
02/17/2005 10:27:44  ANR8331I LTO volume 030AKK is mounted R/W in drive DRLTO_1
                      (/dev/rmt0), status: DISMOUNTING. (SESSION: 21)
02/17/2005 10:27:44  ANR8334I         1 matches found. (SESSION: 21)
02/17/2005 10:28:10  ANR8468I LTO volume 030AKK dismounted from drive DRLTO_1
                      (/dev/rmt0) in library LIBLTO. (SESSION: 101)
```

Figure 7-83 Activity log shows tape volume is dismounted when backup ends

Results summary

The test results show that, after a failure on the node that hosts both the Tivoli Storage Manager scheduler as well as the Storage Agent shared resources, a scheduled incremental backup started on one node for LAN-free is restarted and successfully completed on the other node, also using the SAN path.

This is true if the startup window used to define the schedule is not elapsed when the scheduler service restarts on the second node.

The Tivoli Storage Manager server on AIX resets the SCSI bus when the Storage Agent is restarted on the second node. This permits us to dismount the tape volume from the drive where it was mounted before the failure. When the client restarts the LAN-free operation, the same Storage Agent commands the server to mount again the tape volume to continue the backup.

Restriction: This configuration, with two Storage Agents on the same machine, is not technically supported by Tivoli Storage Manager for SAN. However, in our lab environment it worked.

Note: In other tests we made using the local Storage Agent on each node for communication to the virtual client for LAN-free, the SCSI bus reset did not work. The reason is that when the Tivoli Storage Manager server on AIX acts as a Library Manager, it can handle the SCSI bus reset only when the Storage Agent name is the same for the failing and recovering Storage Agent.

In other words, if we use local Storage Agents for LAN-free backup of the virtual client (CL_MSCS02_TSM), the following conditions must be taken into account: The failure of the node SENEGAL means that all local services will also fail, including SENEGAL_STA (the local Storage Agent). MSCS will cause a failover to the second node where the local Storage Agent will be started again, but with a different name (TONGA_STA). It is this discrepancy in naming which will cause the LAN-free backup to fail, as clearly, the virtual client will be unable to connect to SENEGAL_STA.

Tivoli Storage Manager server does not know what happened to the first Storage Agent because it does not receive any alert from it, until the node that failed is up again, so that the tape drive is in a RESERVED status until the default timeout (10 minutes) elapses. If the scheduler for CL_MSCS02_TSM starts a new session before the ten minutes timeout elapses, it tries to communicate to the local Storage Agent of this second node, TONGA_STA, and this prompts the Tivoli Storage Manager server to mount the same tape volume.

Since this tape volume is still mounted on the first drive by SENEGAL_STA (even when the node failed) and the drive is RESERVED, the only option for the Tivoli Storage Manager server is to mount a new tape volume in the second drive. If either there are not enough tape volumes in the tape storage pool, or the second drive is busy at that time with another operation, or if the client node has its maximum mount points limited to 1, the backup is cancelled.

Testing LAN-free client restore

In this section we test LAN-free client restore.

Objective

The objective of this test is to show what happens when a LAN-free restore is started for a virtual node in the cluster, and the node that hosts the resources at that moment suddenly fails.

Activities

To do this test, we perform these tasks:

1. We open the Cluster Administrator menu to check which node hosts the Tivoli Storage Manager scheduler resource: SENEGAL.
2. We schedule a client restore operation using the Tivoli Storage Manager server scheduler and we associate the schedule to CL_MSCS02_TSM nodename.
3. We make sure that TSM StorageAgent2 and TSM Scheduler for CL_MSCS02_TSM are online resources on SENEGAL.
4. When it is the scheduled time, a client session for CL_MSCS02_TSM nodename starts on the server. At the same time several sessions are also started for CL_MSCS02_STA for Tape Library Sharing and the Storage Agent prompts the Tivoli Storage Manager server to mount a tape volume. The tape volume is mounted in drive DRLTO_1. All of these events are shown in Figure 7-84.

```
IBM Tivoli Storage Manager
02/17/2005 10:46:03  ANR0406I Session 155 started for node CL_MSCS02_TSM
                      (WinNT) (Tcp/Ip 9.1.39.166(1822)). (SESSION: 155)
02/17/2005 10:46:05  ANR0408I Session 156 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for storage agent.  (SESSION: 156)
02/17/2005 10:46:05  ANR0415I Session 156 proxied by CL_MSCS02_STA started for
                      node CL_MSCS02_TSM. (SESSION: 156)
02/17/2005 10:46:05  ANR0408I Session 157 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for storage agent.  (SESSION: 157)
02/17/2005 10:46:05  ANR0408I Session 158 started for server CL_MSCS02_STA
more...   (<ENTER> to continue, 'C' to cancel)

                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 158)
02/17/2005 10:46:05  ANR0409I Session 158 ended for server CL_MSCS02_STA
                      (Windows). (SESSION: 158)
02/17/2005 10:46:05  ANR0408I Session 159 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 159)
02/17/2005 10:46:05  ANR0409I Session 159 ended for server CL_MSCS02_STA
                      (Windows). (SESSION: 159)
02/17/2005 10:46:05  ANR0406I Session 160 started for node TONGA (WinNT)
                      (Tcp/Ip 9.1.39.168(1179)). (SESSION: 160)
02/17/2005 10:46:05  ANR0403I Session 160 ended for node TONGA (WinNT).
                      (SESSION: 160)
02/17/2005 10:46:25  ANR8337I LTO volume 030AKK mounted in drive DRLTO_1
                      (/dev/rmt0). (SESSION: 153)
02/17/2005 10:46:25  ANR0409I Session 153 ended for server CL_MSCS02_STA
```

Figure 7-84 Starting restore session for LAN-free

5. The client starts restoring files using the CL_MSCS02_STA Storage Agent as we can see on the schedule log file in Figure 7-85.

```
dsmsched.log - Notepad
File  Edit  Format  View  Help
Executing scheduled command now.
02/17/2005 10:47:04 --- SCHEDULEREC OBJECT BEGIN RESTORE 02/17/2005 10:45:39
02/17/2005 10:47:04 Restore function invoked.

02/17/2005 10:47:05 ANS1247I Waiting for files from the server... ** Interrupted **
02/17/2005 10:47:08 ANS1114I Waiting for mount of offline media.
02/17/2005 10:48:47 Restoring               0 \\cl_mscs02\g$\code\adminc [Done]
02/17/2005 10:48:47 Restoring               0 \\cl_mscs02\g$\code\drivers_lto [Done
02/17/2005 10:48:47 Restoring               0 \\cl_mscs02\g$\code\isc [Done]
02/17/2005 10:48:47 Restoring               0 \\cl_mscs02\g$\code\lto2k3 [Done]
02/17/2005 10:48:47 Restoring               0 \\cl_mscs02\g$\code\storageagent [Don
02/17/2005 10:48:47 Restoring       1,862,528 \\cl_mscs02\g$\code\IBMTape.W32_6068.
02/17/2005 10:48:48 Restoring      25,388,669 \\cl_mscs02\g$\code\TSM530C_WIN32.exe
02/17/2005 10:48:48 Restoring      10,762,389 \\cl_mscs02\g$\code\TSMCON5300_WIN.ex
02/17/2005 10:48:54 Restoring     181,664,860 \\cl_mscs02\g$\code\TSMSRV5300_WIN.ex
02/17/2005 10:48:55 Restoring      29,205,340 \\cl_mscs02\g$\code\TSMSTA5300_WIN.ex
02/17/2005 10:48:55 Restoring      29,065,149 \\cl_mscs02\g$\code\TSMSTA5301_WIN.ex
02/17/2005 10:48:55 Restoring               0 \\cl_mscs02\g$\code\adminc\jacl [Done
02/17/2005 10:48:55 Restoring               0 \\cl_mscs02\g$\code\adminc\META-INF [
02/17/2005 10:48:55 Restoring               0 \\cl_mscs02\g$\code\adminc\shared [Do
02/17/2005 10:48:55 Restoring               0 \\cl_mscs02\g$\code\adminc\Tivoli [Do
02/17/2005 10:48:55 Restoring               0 \\cl_mscs02\g$\code\adminc\tutorial [
02/17/2005 10:48:56 Restoring       7,515,003 \\cl_mscs02\g$\code\adminc\AdminCente
02/17/2005 10:49:00 Restoring     148,707,743 \\cl_mscs02\g$\code\adminc\C8243ML.ex
02/17/2005 10:49:00 Restoring          20,955 \\cl_mscs02\g$\code\adminc\dsminstall
```

Figure 7-85 Restore starts on the schedule log file

6. In Figure 7-86 we see that the Storage Agent has an opened session with the virtual client, CL_MSCS02_TSM, as well as Tivoli Storage Manager, TSMSRV03, and the tape volume is mounted for its use.

```
IBM Tivoli Storage Manager

  Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------   ------ ------ ------ -----  -------- ----------------
     1 Tcp/Ip Run      0 S    7.0 K  40.1 K Serv-  AIX-RS/- TSMSRV03
                                             er     6000
     3 Tcp/Ip IdleW    0 S    6.5 K    253 Serv-  AIX-RS/- TSMSRV03
                                             er     6000
     4 Tcp/Ip IdleW   12 S    4.6 K  37.2 K Serv-  AIX-RS/- TSMSRV03
                                             er     6000
    14 Shared SendW    0 S  381.8 M    551 Node   WinNT    CL_MSCS02_TSM
       Memo-
       ry
    15 Tcp/Ip Run      0 S    4.2 K   6.9 K Serv-  AIX-RS/- TSMSRV03
                                             er     6000
    16 Tcp/Ip IdleW    0 S   24.1 K  24.7 K Serv-  AIX-RS/- TSMSRV03
                                             er     6000
    19 Tcp/Ip Run      0 S    1.7 K    160 Admin  WinNT    ADMIN

tsm: CL_MSCS02_STA>q mo
ANR8330I LTO volume 030AKK is mounted R/W in drive DRLTO_1 (mt0.0.0.2), status:
IN USE.
ANR8334I         1 matches found.

tsm: CL_MSCS02_STA>_
```

Figure 7-86 Storage agent shows sessions for the server and the client

7. While the client is restoring the files, we force SENEGAL to fail. The following sequence takes place:

   a. The client CL_MSCS02_TSM and the Storage Agent CL_MSCS02_STA lose both temporarily their connections with the server, as shown in Figure 7-87.

```
IBM Tivoli Storage Manager
02/10/2005 18:47:58  ANR2017I Administrator ADMIN issued command: QUERY ACTLOG
                      begint=now-0:01  (SESSION: 412)
02/10/2005 18:48:20  ANR0568W Session 412 for admin ADMIN (WinNT) terminated -
                      connection with client severed. (SESSION: 412)
02/10/2005 18:48:23  ANR0409I Session 362 ended for server SENEGAL_STA
                      (Windows). (SESSION: 362)
more...   (<ENTER> to continue, 'C' to cancel)

02/10/2005 18:48:23  ANR0409I Session 415 ended for server CL_MSCS02_STA
                      (Windows). (SESSION: 415)
02/10/2005 18:48:23  ANR0409I (Session: 417, Origin: CL_MSCS02_STA)  Session 1
                      ended for server TSMSRV01 (Windows). (SESSION: 417)
02/10/2005 18:48:23  ANR7835I (Session: 417, Origin: CL_MSCS02_STA)  The server
                      thread 1 (tid 2200) terminated in response to server
                      shutdown. (SESSION: 417)
02/10/2005 18:48:23  ANR7835I (Session: 417, Origin: CL_MSCS02_STA)  The server
                      thread 3 (tid 2140) terminated in response to server
                      shutdown. (SESSION: 417)
02/10/2005 18:48:23  ANR7835I (Session: 417, Origin: CL_MSCS02_STA)  The server
                      thread 4 (tid 2356) terminated in response to server
                      shutdown. (SESSION: 417)
02/10/2005 18:48:23  ANR7835I (Session: 417, Origin: CL_MSCS02_STA)  The server
                      thread 5 (tid 2908) terminated in response to server
                      shutdown. (SESSION: 417)
02/10/2005 18:48:23  ANR7835I (Session: 417, Origin: CL_MSCS02_STA)  The server
                      thread 6 (tid 2924) terminated in response to server
                      shutdown. (SESSION: 417)
02/10/2005 18:48:23  ANR7835I (Session: 417, Origin: CL_MSCS02_STA)  The server
                      thread 7 (tid 3068) terminated in response to server
                      shutdown. (SESSION: 417)
02/10/2005 18:48:25  ANR3605E Unable to communicate with storage agent.
more...   (<ENTER> to continue, 'C' to cancel)

                      (SESSION: 418)
02/10/2005 18:48:25  ANR0480W Session 422 for node CL_MSCS02_TSM (WinNT)
                      terminated - connection with client severed. (SESSION:
                      422)
```

Figure 7-87 Both sessions for the Storage Agent and the client lost in the server

   b. The tape volume is still mounted on the same drive.

   c. After a short period of time the resources are online on TONGA.

   d. When the Storage Agent CL_MSCS02_STA is again online (in SENEGAL), the TSM Scheduler service also is started (because of the dependency between these two resources). The Tivoli Storage Manager resets the SCSI bus when the Storage Agent starts, and it dismounts the tape volume. We show this on the activity log for the server in Figure 7-88.

```
IBM Tivoli Storage Manager
02/17/2005 10:48:53  ANR0406I Session 181 started for node CL_MSCS02_TSM
                      (WinNT) (Tcp/Ip 9.1.39.168(1390)). (SESSION: 181)
02/17/2005 10:48:53  ANR1639I Attributes changed for node CL_MSCS02_TSM: TCP
                      Name from SENEGAL to TONGA, TCP Address from 9.1.39.166
                      to , GUID from 46.6e.f4.d0.75.3a.11.d9.bc.7b.00.0d.60.49-
                      .41.dd to 18.72.62.50.75.3b.11.d9.8f.b7.00.a0.cc.d2.58.a-
                      c. (SESSION: 181)
02/17/2005 10:48:54  ANR0403I Session 181 ended for node CL_MSCS02_TSM (WinNT).
                      (SESSION: 181)
02/17/2005 10:48:54  ANR0406I Session 182 started for node CL_MSCS02_SA (WinNT)
                      (Tcp/Ip 9.1.39.168(1397)). (SESSION: 182)
02/17/2005 10:48:54  ANR0403I Session 182 ended for node CL_MSCS02_SA (WinNT).
                      (SESSION: 182)
02/17/2005 10:48:56  ANR0406I Session 183 started for node CL_MSCS02_QUORUM
                      (WinNT) (Tcp/Ip 9.1.39.168(1401)). (SESSION: 183)
02/17/2005 10:48:56  ANR0403I Session 183 ended for node CL_MSCS02_QUORUM
                      (WinNT). (SESSION: 183)
02/17/2005 10:48:57  ANR0406I Session 184 started for node CL_MSCS02_SA (WinNT)
                      (Tcp/Ip 9.1.39.168(1402)). (SESSION: 184)
02/17/2005 10:48:57  ANR0403I Session 184 ended for node CL_MSCS02_SA (WinNT).
                      (SESSION: 184)
02/17/2005 10:48:58  ANR0406I Session 185 started for node CL_MSCS02_TSM
                      (WinNT) (Tcp/Ip 9.1.39.168(1403)). (SESSION: 185)
02/17/2005 10:49:01  ANR0408I Session 186 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for storage agent.  (SESSION: 186)
02/17/2005 10:49:01  ANR0415I Session 186 proxied by CL_MSCS02_STA started for
                      node CL_MSCS02_TSM. (SESSION: 186)
02/17/2005 10:49:01  ANR0408I Session 187 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 187)
02/17/2005 10:49:01  ANR0409I Session 187 ended for server CL_MSCS02_STA
                      (Windows). (SESSION: 187)
02/17/2005 10:49:01  ANR0408I Session 188 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 188)
02/17/2005 10:49:28  ANR0409I Session 188 ended for server CL_MSCS02_STA
                      (Windows). (SESSION: 188)
02/17/2005 10:49:31  ANR0407I Session 189 started for administrator ADMIN
                      (WinNT) (Tcp/Ip 9.1.39.168(1409)). (SESSION: 189)
02/17/2005 10:49:33  ANR2017I Administrator ADMIN issued command: QUERY SESSION
                      (SESSION: 189)
02/17/2005 10:49:34  ANR2017I Administrator ADMIN issued command: QUERY MO
more...   <<ENTER> to continue, 'C' to cancel>

                      (SESSION: 189)
02/17/2005 10:49:34  ANR8330I LTO volume 030AKK is mounted R/W in drive DRLTO_1
                      (/dev/rmt0), status: IN USE. (SESSION: 189)
```

Figure 7-88 Resources are started again in the second node

e. For the Storage Agent, at the same time, the tape volume is idle because there is no session with the client yet, and the tape volume is dismounted (Figure 7-89).

```
IBM Tivoli Storage Manager

  Sess Comm.  Sess      Wait    Bytes    Bytes Sess  Platform Client Name
Number Method State     Time     Sent    Recvd Type
------ ------ ------ -------  -------  ------ ----- -------- --------------------
     1 Tcp/Ip IdleW    10 S   110.7 K  116.6 K Serv- AIX-RS/- TSMSRV03
                                                 er    6000
     3 Tcp/Ip IdleW     0 S     5.3 K      253 Serv- AIX-RS/- TSMSRV03
                                                 er    6000
     4 Tcp/Ip IdleW    10 S     8.8 K  1,021.3 Serv- AIX-RS/- TSMSRV03
                                           K  er    6000
    11 Tcp/Ip Run       0 S     3.8 K      210 Admin WinNT    ADMIN
    12 Tcp/Ip IdleW     0 S    59.9 K   63.6 K Serv- AIX-RS/- TSMSRV03
                                                 er    6000

tsm: CL_MSCS02_STA>q mo
ANR8329I LTO volume 030AKK is mounted R/W in drive DRLTO_1 (mt0.0.0.2), status:
IDLE.
ANR8334I          1 matches found.

tsm: CL_MSCS02_STA>q mo
ANR8331I LTO volume 030AKK is mounted R/W in drive DRLTO_1 (mt0.0.0.2), status:
DISMOUNTING.
ANR8334I          1 matches found.

tsm: CL_MSCS02_STA>_
```

Figure 7-89 Storage agent commands the server to dismount the tape volume

f. When the client restarts the session, the Storage Agent commands the server to mount the tape volume and it starts sending data directly to the client, as we see in Figure 7-90.

```
IBM Tivoli Storage Manager
tsm: CL_MSCS02_STA>q se

  Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- --------------------
     1 Tcp/Ip IdleW    0 S  110.7 K 116.6 K Serv- AIX-RS/- TSMSRV03
                                          er    6000
     3 Tcp/Ip IdleW    0 S    3.2 K     253 Serv- AIX-RS/- TSMSRV03
                                          er    6000
     4 Tcp/Ip Run      0 S    6.7 K 576.7 K Serv- AIX-RS/- TSMSRV03
                                          er    6000
     7 Shared SendW    0 S  998.8 M     821 Node  WinNT    CL_MSCS02_TSM
       Memo-
       ry
     8 Tcp/Ip Run      0 S  183.8 K 300.7 K Serv- AIX-RS/- TSMSRV03
                                          er    6000
    11 Tcp/Ip Run      0 S    1.1 K     138 Admin WinNT    ADMIN
    12 Tcp/Ip IdleW    0 S      341     304 Serv- AIX-RS/- TSMSRV03
                                          er    6000

tsm: CL_MSCS02_STA>q mo
ANR8330I LTO volume 030AKK is mounted R/W in drive DRLTO_1 (mt0.0.0.2), status:
IN USE.
ANR8334I         1 matches found.
```

Figure 7-90 Storage agent writes to the volume again

g. When the tape volume is mounted again, the client restarts its scheduled restore from the beginning such as we can see in Figure 7-91.

```
dsmsched.log - Notepad
File  Edit  Format  View  Help
02/17/2005 10:49:07 Restoring              0 \\cl_mscs02\g$\code\isc\wps [Done]
02/17/2005 10:49:16 ANS2820E An interrupt has occurred. The current operation will
client will shut down.
02/17/2005 10:49:59 Querying server for next scheduled event.
02/17/2005 10:49:59 Node Name: CL_MSCS02_TSM
02/17/2005 10:49:59 Session established with server TSMSRV03: AIX-RS/6000
02/17/2005 10:49:59   Server Version 5, Release 3, Level 0.0
02/17/2005 10:49:59   Server date/time: 02/17/2005 10:48:58  Last access: 02/17/200

02/17/2005 10:49:59 --- SCHEDULEREC QUERY BEGIN
02/17/2005 10:49:59 --- SCHEDULEREC QUERY END
02/17/2005 10:49:59 Next operation scheduled:
02/17/2005 10:49:59 ------------------------------------------------------------
02/17/2005 10:49:59 Schedule Name:         RESTORE
02/17/2005 10:49:59 Action:                Restore
02/17/2005 10:49:59 Objects:               g:\code\*
02/17/2005 10:49:59 Options:               -subdir=yes -replace=yes
02/17/2005 10:49:59 Server Window Start:   10:45:39 on 02/17/2005
02/17/2005 10:49:59 ------------------------------------------------------------
02/17/2005 10:49:59
Executing scheduled command now.
02/17/2005 10:49:59 --- SCHEDULEREC OBJECT BEGIN RESTORE 02/17/2005 10:45:39
02/17/2005 10:49:59 Restore function invoked.

02/17/2005 10:50:00 ANS1247I Waiting for files from the server... ** Interrupted **
02/17/2005 10:50:04 ANS1114I Waiting for mount of offline media.
02/17/2005 10:50:38 Restoring              0 \\cl_mscs02\g$\code\adminc [Done]
02/17/2005 10:50:38 Restoring              0 \\cl_mscs02\g$\code\drivers_lto [Done
02/17/2005 10:50:38 Restoring              0 \\cl_mscs02\g$\code\isc [Done]
```

Figure 7-91 The client restarts the restore

8. When the restore is completed, we look at the final statistics in the schedule log file of the client as shown in Figure 7-92.

```
dsmsched.log - Notepad
File  Edit  Format  View  Help
02/17/2005 10:51:33 Restoring         143,360 \\cl_mscs02\g$\code\lto2k3\uninst.exe
02/17/2005 10:51:33 Restoring         527,360 \\cl_mscs02\g$\code\lto2k3\checked\ib
02/17/2005 10:51:33 Restoring         279,552 \\cl_mscs02\g$\code\lto2k3\checked\ib
02/17/2005 10:51:33 Restoring         543,744 \\cl_mscs02\g$\code\lto2k3\checked\ib
02/17/2005 10:51:33 Restoring         282,112 \\cl_mscs02\g$\code\lto2k3\checked\ib
02/17/2005 10:51:33 Restoring          46,592 \\cl_mscs02\g$\code\lto2k3\checked\ib
02/17/2005 10:51:33 Restoring         510,976 \\cl_mscs02\g$\code\lto2k3\checked\ib
02/17/2005 10:51:33 Restoring         282,368 \\cl_mscs02\g$\code\lto2k3\checked\ib
02/17/2005 10:51:33 Restoring         527,360 \\cl_mscs02\g$\code\lto2k3\checked\ib
02/17/2005 10:51:33 Restoring         285,696 \\cl_mscs02\g$\code\lto2k3\checked\ib
02/17/2005 10:51:39 Restoring     111,755,569 \\cl_mscs02\g$\code\storageagent\c811
02/17/2005 10:51:40
Restore processing finished.
02/17/2005 10:51:42 --- SCHEDULEREC STATUS BEGIN
02/17/2005 10:51:42 Total number of objects restored:     1,927
02/17/2005 10:51:42 Total number of objects failed:           0
02/17/2005 10:51:42 Total number of bytes transferred:     1.29 GB
02/17/2005 10:51:42 LanFree data bytes:                    1.29 GB
02/17/2005 10:51:42 Data transfer time:                   37.62 sec
02/17/2005 10:51:42 Network data transfer rate:       35,993.08 KB/sec
02/17/2005 10:51:42 Aggregate data transfer rate:     13,170.25 KB/sec
02/17/2005 10:51:42 Elapsed processing time:          00:01:42
02/17/2005 10:51:42 --- SCHEDULEREC STATUS END
02/17/2005 10:51:42 --- SCHEDULEREC OBJECT END RESTORE 02/17/2005 10:45:39
02/17/2005 10:51:42 --- SCHEDULEREC STATUS BEGIN
02/17/2005 10:51:42 --- SCHEDULEREC STATUS END
02/17/2005 10:51:42 Scheduled event 'RESTORE' completed successfully.
02/17/2005 10:51:42 Sending results for scheduled event 'RESTORE'.
02/17/2005 10:51:42 Results sent to server for scheduled event 'RESTORE'.
```

Figure 7-92 Final statistics for the restore on the schedule log file

Note: Notice that the restore process is started from the beginning, it is not restarted.

Archived

9. In the activity log the restore ends successfully and the tape volume is dismounted correctly as we see in Figure 7-93.

```
IBM Tivoli Storage Manager
02/17/2005 10:50:41  ANR2507I Schedule RESTORE for domain STANDARD started at
                      02/17/05 10:45:39 for node CL_MSCS02_TSM completed
                      successfully at 02/17/05 10:50:41. (SESSION: 185)
02/17/2005 10:50:41  ANR0403I Session 185 ended for node CL_MSCS02_TSM (WinNT).
                      (SESSION: 185)
02/17/2005 10:50:41  ANR0406I Session 193 started for node CL_MSCS02_TSM
                      (WinNT) (Tcp/Ip 9.1.39.168(1429)). (SESSION: 193)
02/17/2005 10:50:41  ANR0403I Session 186 ended for node CL_MSCS02_TSM (WinNT).
                      (SESSION: 186)
02/17/2005 10:50:41  ANR0403I Session 193 ended for node CL_MSCS02_TSM (WinNT).
                      (SESSION: 193)
02/17/2005 10:51:00  ANR0406I Session 194 started for node SENEGAL (WinNT)
                      (Tcp/Ip 9.1.39.166(1315)). (SESSION: 194)
02/17/2005 10:51:00  ANR0403I Session 194 ended for node SENEGAL (WinNT).
                      (SESSION: 194)
more...   (<ENTER> to continue, 'C' to cancel)

02/17/2005 10:51:34  ANR2017I Administrator ADMIN issued command: QUERY MO
                      (SESSION: 189)
02/17/2005 10:51:34  ANR8329I LTO volume 030AKK is mounted R/W in drive DRLTO_1
                      (/dev/rmt0), status: IDLE. (SESSION: 189)
02/17/2005 10:51:34  ANR8334I         1 matches found. (SESSION: 189)
02/17/2005 10:51:36  ANR2017I Administrator ADMIN issued command: QUERY SESSION
                      (SESSION: 189)
02/17/2005 10:51:37  ANR8325I Dismounting volume 030AKK - 1 minute mount
                      retention expired.
02/17/2005 10:51:40  ANR2017I Administrator ADMIN issued command: QUERY MO
                      (SESSION: 189)
02/17/2005 10:51:40  ANR8331I LTO volume 030AKK is mounted R/W in drive DRLTO_1
                      (/dev/rmt0), status: DISMOUNTING. (SESSION: 189)
02/17/2005 10:51:40  ANR8334I         1 matches found. (SESSION: 189)
02/17/2005 10:51:42  ANR0408I Session 195 started for server CL_MSCS02_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 195)
02/17/2005 10:52:20  ANR0409I Session 195 ended for server CL_MSCS02_STA
                      (Windows). (SESSION: 195)
02/17/2005 10:52:28  ANR8336I Verifying label of LTO volume 030AKK in drive
                      DRLTO_1 (/dev/rmt0). (SESSION: 153)
02/17/2005 10:52:43  ANR2017I Administrator ADMIN issued command: QUERY MO
                      (SESSION: 189)
02/17/2005 10:52:43  ANR8331I LTO volume 030AKK is mounted R/W in drive DRLTO_1
                      (/dev/rmt0), status: DISMOUNTING. (SESSION: 189)
02/17/2005 10:52:43  ANR8334I         1 matches found. (SESSION: 189)
02/17/2005 10:53:02  ANR8468I LTO volume 030AKK dismounted from drive DRLTO_1
                      (/dev/rmt0) in library LIBLTO. (SESSION: 153)
02/17/2005 10:54:07  ANR2017I Administrator ADMIN issued command: UPDATE
```

Figure 7-93 Restore completed and volume dismounted by the server in actlog

Results summary

The test results show that after a failure on the node that hosts the Tivoli Storage Manager client scheduler instance, a scheduled restore operation started on this node using the LAN-free path is started again from the beginning on the second node of the cluster when the service is online.

This is true if the startup window for the scheduled restore operation is not elapsed when the scheduler client is online again on the second node.

Also notice that the restore is not restarted from the point of failure, but started from the beginning. The scheduler queries the Tivoli Storage Manager server for a scheduled operation and a new session is opened for the client after the failover.

Restriction: Notice again that this configuration, with two Storage Agents in the same machine, is not officially supported by Tivoli Storage Manager for SAN. However, in our lab environment it worked. In other tests we made using the local Storage Agents for communication to the virtual client for LAN-free, the SCSI bus reset did not work and the restore process failed.

Archived

Archived

Part 3

AIX V5.3 with HACMP V5.2 environments and IBM Tivoli Storage Manager Version 5.3

In this part of the book, we discuss highly available clustering, using the AIX operating system. There are many different configurations possible; however, we will document the configurations we believe will provide a balance between availability and cost effective computing. We will cover two clustering products, High Availability Cluster Multi-Processing (HACMP) and VERITAS Cluster Services (VCS).

© Copyright IBM Corp. 2005. All rights reserved.

Archived

8

Establishing an HACMP infrastructure on AIX

This chapter describes the planning and installation of HACMP Version 5.2 on AIX Version 5.3. We establish an HACMP cluster infrastructure, in which we will then build our application environment in the chapters that follow.

© Copyright IBM Corp. 2005. All rights reserved.

8.1 Overview

In this overview we discuss topics which our team reviewed, and believe the reader would also want to review and fully understand, prior to advancing to later chapters.

8.1.1 AIX overview

There are many AIX V5.3 enhancements. In this overview we list items that are most relative to a large number of Tivoli Storage Manager and HACMP environments. We recommend reviewing the details, which are available in the IBM Redbook, *AIX 5L Differences Guide Version 5.3 Edition*, SG24-7463-00.

Storage management

AIX 5L™ introduces several new features for the current and emerging storage requirements.

These enhancements include:

- LVM enhancements
- Performance improvement of LVM commands
- Removal of classical concurrent mode support
- Scalable volume groups
- Striped column support for logical volumes
- Volume group pbuf pools
- Variable logical track group
- JFS2 enhancements
- Disk quotas support for JFS2
- JFS2 file system shrink
- JFS2 extended attributes Version 2 support
- JFS2 ACL support for NFS V4
- ACL inheritance support
- JFS2 logredo scalability
- JFS2 file system check scalability

Reliability, availability, and serviceability

This section includes descriptions of the following enhancements for AIX 5L.

Error logging, core files, and system dumps

These enhancements include:

- Error log RAS
- Enhancements for a large number of devices
- Core file creation and compression

- System dump enhancements
- DVD support for system dumps
- snap command enhancements

Trace enhancements

These enhancements include:

- Administrative control of the user trace buffers
- Single thread trace

System management

AIX 5L provides many enhancements in the area of system management and utilities. This section discusses these enhancements. Topics include:

- InfoCenter for AIX 5L Version 5.3
- Multiple desktop selection from BOS menus
- Erasing hard drive during BOS install
- Service Update Management Assistant
- Long user and group name support
- Dynamic reconfiguration usability
- Paging space garbage collection
- Dynamic support for large page pools
- Interim Fix Management
- List installed filesets by bundle
- Configuration file modification surveillance
- DVD backup using the mkdvd command
- NIM security
- High Available NIM (HA NIM)
- General NIM enhancements

8.2 HACMP overview

This overview contains an introduction to IBM High Availability Cluster Multi-Processing (HACMP) for AIX product line, and the concepts on which IBM's high availability products are based. It is essential that the reader fully understand how HACMP works, and what HACMP is designed to deliver with regard to availability.

We discuss the following topics:

- What is HACMP?
- HACMP concepts
- HACMP terminology

Archived

8.2.1 What is HACMP?

IBM's high availability solution for AIX, High Availability Cluster Multi Processing, based on IBM's well-proven clustering technology, consists of two components:

- **High availability:** The process of ensuring that an application is available for use through the use of duplicated and/or shared resources
- **Cluster multi-processing:** Multiple applications running on the same nodes with shared or concurrent access to the data

A high availability solution based on HACMP provides automated failure detection, diagnosis, application recovery, and node reintegration. With an appropriate application, HACMP can also provide concurrent access to the data for parallel processing applications, thus offering excellent horizontal scalability.

A typical HACMP environment is shown in Figure 8-1.

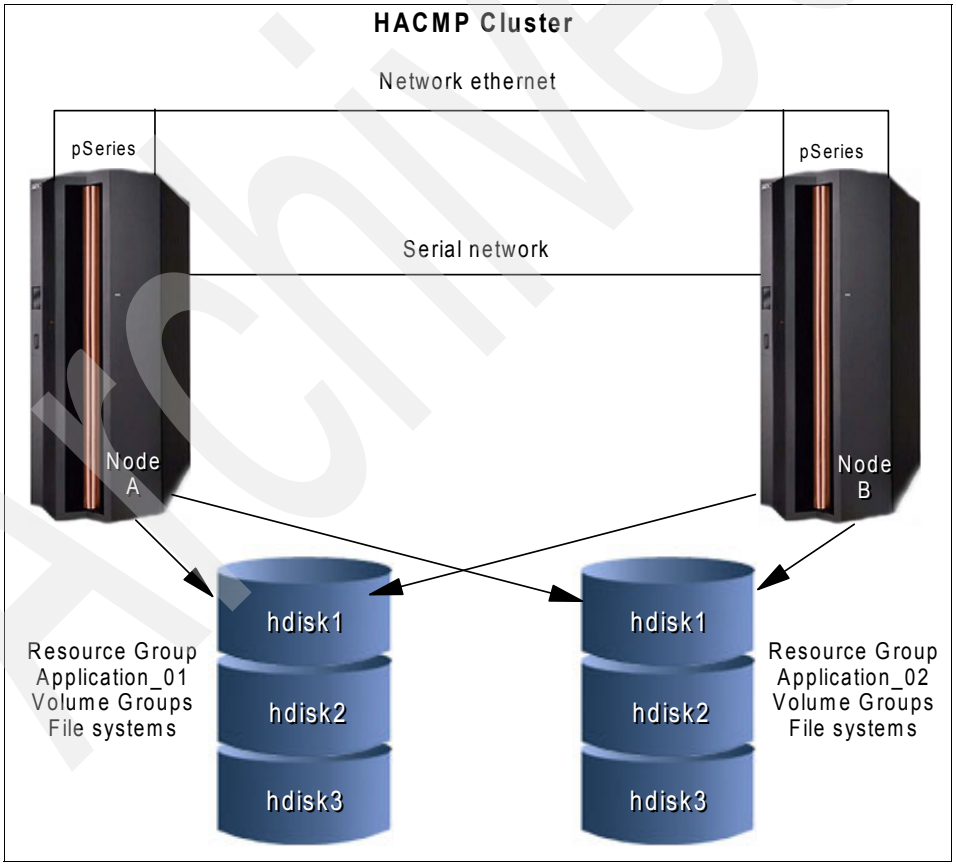

Figure 8-1 HACMP cluster

Archived

8.3 HACMP concepts

The basic concepts of HACMP can be classified as follows:

- Cluster topology:

  Contains basic cluster components nodes, networks, communication interfaces, communication devices, and communication adapters.

- Cluster resources:

  Entities that are being made highly available (for example, file systems, raw devices, service IP labels, and applications). Resources are grouped together in resource groups (RGs), which HACMP keeps highly available as a single entity. Resource groups can be available from a single node or, in the case of concurrent applications, available simultaneously from multiple nodes.

- Fallover:

  Represents the movement of a resource group from one active node to another node (backup node) in response to a failure on that active node.

- Fallback:

  Represents the movement of a resource group back from the backup node to the previous node, when it becomes available. This movement is typically in response to the reintegration of the previously failed node.

8.3.1 HACMP terminology

To understand the correct functionality and utilization of HACMP, it is necessary to know some important terms:

- Cluster:

  Loosely-coupled collection of independent systems (nodes) or LPARs organized into a network for the purpose of sharing resources and communicating with each other.

  HACMP defines relationships among cooperating systems where peer cluster nodes provide the services offered by a cluster node should that node be unable to do so. These individual nodes are together responsible for maintaining the functionality of one or more applications in case of a failure of any cluster component.

- Node:

  An IBM @server® pSeries® machine (or LPAR) running AIX and HACMP that is defined as part of a cluster. Each node has a collection of resources (disks, file systems, IP address(es), and applications) that can be transferred to another node in the cluster in case the node fails.

Archived

- Resource:

  Resources are logical components of the cluster configuration that can be moved from one node to another. All the logical resources necessary to provide a Highly Available application or service are grouped together in a resource group (RG).

  The components in a resource group move together from one node to another in the event of a node failure. A cluster may have more than one resource group, thus allowing for efficient use of the cluster nodes (thus the “Multi-Processing” in HACMP).

- Takeover:

  This is the operation of transferring resources between nodes inside the cluster. If one node fails due to a hardware problem or crash of AIX, its resources application will be moved to the another node.

- Client:

  A client is a system that can access the application running on the cluster nodes over a local area network. Clients run a client application that connects to the server (node) where the application runs.

- Heartbeat:

  In order for an HACMP cluster to recognize and respond to failures, it must continually check the health of the cluster. Some of these checks are provided by the heartbeat function. Each cluster node sends heartbeat messages at specific intervals to other cluster nodes, and expects to receive heartbeat messages from the nodes at specific intervals. If messages stop being received, HACMP recognizes that a failure has occurred. Heartbeats can be sent over:

  - TCP/IP networks
  - Point-to-point networks
  - Shared disks.

8.4 Planning and design

In this section we talk about planning and design considerations for the HACMP environment.

8.4.1 Supported hardware and software

We will first ensure that our system meets the hardware and software requirements established for HACMP and SAN connectivity. For up to date information about the required and supported hardware for HACMP, see the sales guide for the product.

To locate the sales guide:

1. Go to the following URL:

   http://www.ibm.com/common/ssi

2. Select your country and language.
3. Select HW and SW Description (SalesManual, RPQ) for a Specific Information Search.

Next, we review up to date information about compatibility of devices and adapter over your SAN. Check the appropriate *Interoperability Matrix* from Storage Support Home page.

1. Go to the following URL:

   http://www-1.ibm.com/servers/storage/support/

2. Select your Product Family: Storage area network (SAN).
3. Select Your switch type and model (our case SAN32B-2).
4. Click either the **Plan** or **Upgrade** folder tab.
5. Click **Interoperability Matrix** link to open the document or right-click to save.

Tip: We must take note of required firmware levels, as we may require this information later in the process.

8.4.2 Planning for networking

Here we list some HACMP networking features we are going to exploit, with the planning for our lab.

Point-to-point networks

We can increase availability by configuring non-IP point-to-point connections that directly link cluster nodes. These connections provide:

- An alternate heartbeat path for a cluster that uses a single TCP/IP-based network, and prevent the TCP/IP software from being a single point of failure
- Protection against cluster partitioning. For more information, see the section, “Cluster Partitioning” in the *HACMP Planning and Installation Guide*.

We can configure heartbeat paths over the following types of networks:

- Serial (RS232)
- Disk heartbeat (over an enhanced concurrent mode disk)
- Target Mode SSA
- Target Mode SCSI

Archived

In our implementation example

We will configure:

- Serial heartbeat
- Disk heartbeat

IP Address Takeover via IP aliases

We can configure IP Address Takeover on certain types of networks using the IP aliases network capabilities supported in AIX. Assigning IP aliases to NICs allows you to create more than one IP label on the same network interface.

HACMP allows the use of IPAT via IP aliases with the following network types that support gratuitous ARP in AIX:

- Ethernet
- Token Ring
- FDDI
- SP Switch1 and SP Switch2.

During IP Address Takeover via IP aliases, when an IP label moves from one NIC to another, the target NIC receives the new IP label as an IP alias and keeps the original IP label and hardware address.

To enable IP Address Takeover via IP aliases, configure NICs to meet the following requirements:

- At least one boot-time IP label must be assigned to the service interface on each cluster node.
- Hardware Address Takeover can not be configured for any interface that has an IP alias configured.
- Subnet requirements:
  - Multiple boot-time addresses configured on a node should be defined on different subnets.
  - Service addresses must be on a different subnet from all non-service addresses defined for that network on the cluster node. This requirement enables HACMP to comply with the IP route striping functionality of AIX 5L 5.1, which allows multiple routes to the same subnet.
- Service address labels configured for IP Address Takeover via IP aliases can be included in all non-concurrent resource groups.
- Multiple service labels can coexist as aliases on a given interface.
- The netmask for all IP labels in an HACMP network must be the same.
- You cannot mix aliased and non-aliased service IP labels in the same resource group.

Archived

HACMP non-service labels are defined on the nodes as the boot-time address assigned by AIX after a system reboot and before the HACMP software is started.

When the HACMP software is started on a node, the node's service IP label is added as an alias onto one of the NICs that has a non-service label.

In our implementation example

We will configure:

- 2 *non-service* subnets
- 2 adapters with a *boot IP label* for each cluster node
- 1 *service* address to be included in the Tivoli Storage Manager resource group
- 1 *service* address to be included in the ISC resource group.

Persistent node IP label

A persistent node IP label is an IP alias that can be assigned to a specific node on a cluster network. A persistent node IP label:

- Always stays on the same node (is node-bound)
- Coexists on a NIC that already has a service or non-service IP label defined
- Does not require installing an additional physical NIC on that node
- Is not part of any resource group.

Assigning a persistent node IP label provides a node-bound address that you can use for administrative purposes, because a connection to a persistent node IP label always goes to a specific node in the cluster. You can have one persistent node IP label per network per node.

After a persistent node IP label is configured on a specified network node, it becomes available at boot time and remains configured even if HACMP is shut down on that node.

In our implementation example

We will configure:

- A *persistent address* for each cluster node

Archived

8.4.3 Plan for cascading versus rotating

A cascading resource group defines a list of all the nodes that can control the resource group and then, by assigning a takeover priority to each node, specifies a preference for which cluster node controls the resource group. When a fallover occurs, the active node with the highest priority acquires the resource group. If that node is unavailable, the node with the next-highest priority acquires the resource group, and so on.

The list of participating nodes establishes the resource chain for that resource group. When a node with a higher priority for that resource group joins or reintegrates into the cluster, it takes control of the resource group, that is, the resource group falls back from nodes with lesser priorities to the higher priority node.

Special cascading resource group attributes

Cascading resource groups support the following attributes:

- Cascading without fallback
- Inactive takeover
- Dynamic node priority

Cascading without fallback (CWOF) is a cascading resource group attribute that allows you to refine fall-back behavior. When the Cascading Without Fallback flag is set to false, this indicates traditional cascading resource group behavior: When a node of higher priority than that on which the resource group currently resides joins or reintegrates into the cluster, and interfaces are available, the resource group falls back to the higher priority node. When the flag is set to true, the resource group will not fall back to any node joining or reintegrating into the cluster, even if that node is a higher priority node. A resource group with CWOF configured does not require IP Address Takeover.

Inactive takeover is a cascading resource group attribute that allows you to fine tune the initial acquisition of a resource group by a node. If inactive takeover is true, then the first node in the resource group to join the cluster acquires the resource group, regardless of the node's designated priority. If Inactive Takeover is false, each node to join the cluster acquires only those resource groups for which it has been designated the highest priority node. The default is false.

Dynamic node priority lets you use the state of the cluster at the time of the event to determine the order of the takeover node list.

In our implementation example

We will configure:

- Two *cascading* resource groups having the following features:
  - *Policies*:
    ONLINE ON HOME NODE ONLY
    FALLOVER TO NEXT PRIORITY NODE
    NEVER FALLBACK
  - Nodes and *priority*:
    AVOV, KANAGA for the Tivoli Storage Manager server
    KANAGA, AVOZ for the *ISC* with *Administration Center*.

8.5 Lab setup

In Figure 8-2, we show the Storage Area Network and the IP network we implemented in our lab from a “physical” point of view.

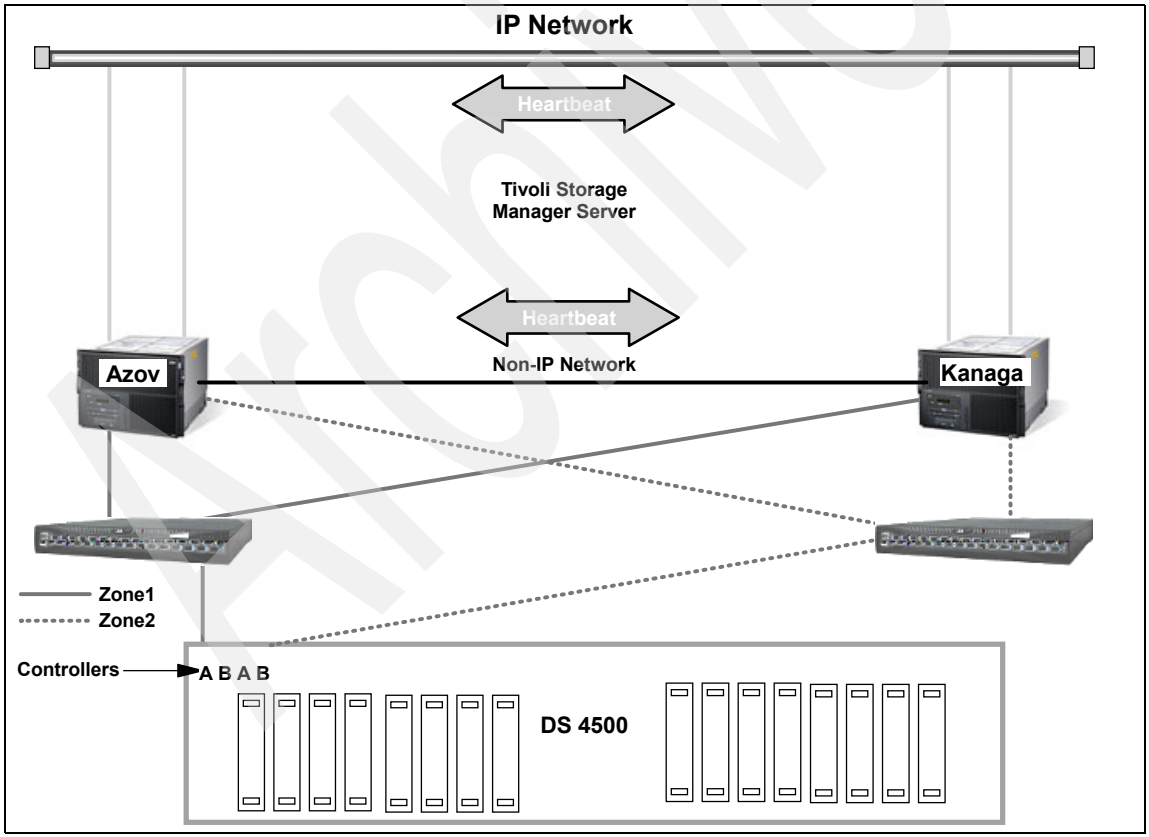

Figure 8-2 AIX Clusters - SAN (Two fabrics) and network

Archived

In Figure 8-3 we provide a logical view of our lab, showing the layout for AIX and Tivoli Storage Manager filesystems, devices, and network.

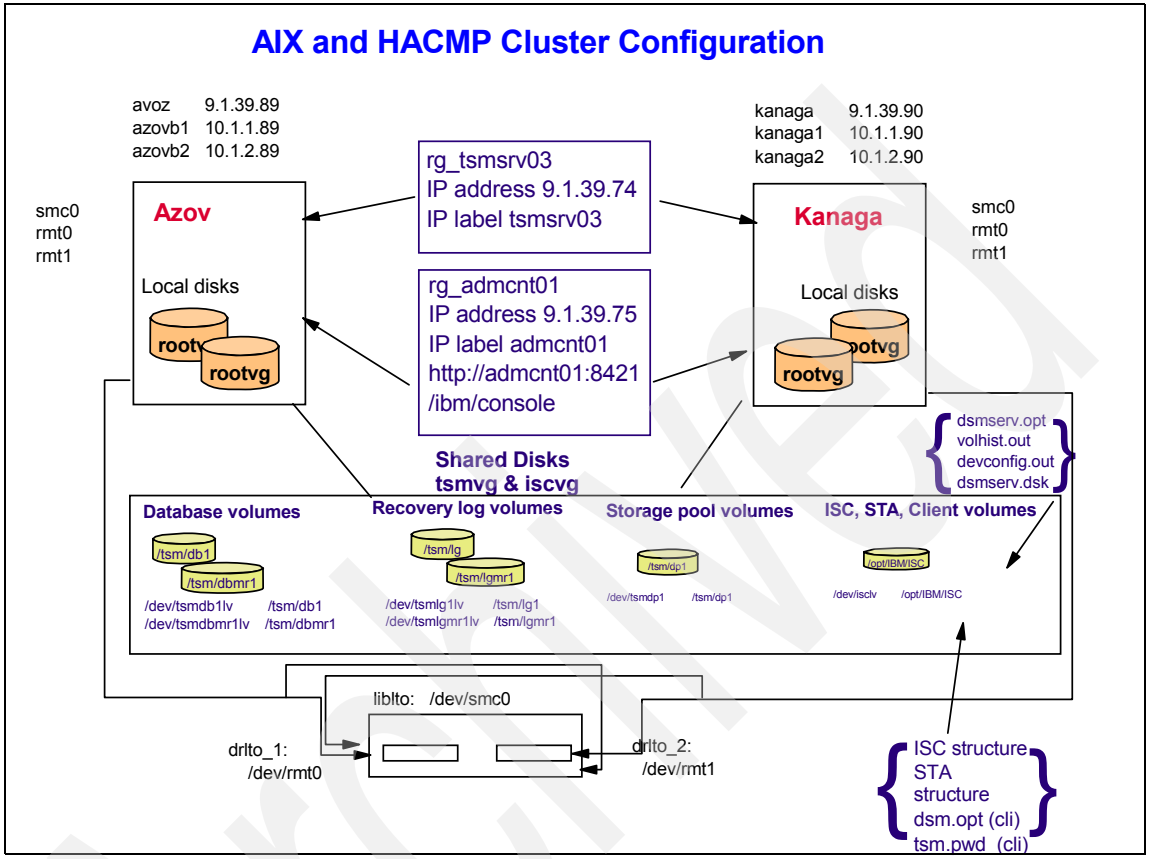

Figure 8-3 Logical layout for AIX and TSM filesystems, devices, and network

Archived

Table 8-1 and Table 8-2 provide some more details about our configuration.

Table 8-1 HACMP cluster topology

HACMP Cluster	
Cluster name	CL_HACMP01
IP network	net_ether_01
IP network / Boot subnet 1	net_ether_01 / 10.1.1.0/24
IP network / Boot subnet 2	net_ether_01 / 10.1.2.0/24
IP network / Service subnet	net_ether_01 / 9.1.39.0/24
Point to point network 1	net_rs232_01
Point to point network 2	net_diskhb_01
Node 1	
Name	AZOV
Boot IP address / IP label 1	10.1.1.89 / azovb1
Boot IP address / IP label 2	10.1.2.89 / azovb1
Persistent address / IP label	9.1.39.89 / azov
Point to point network 1 device	/dev/tty0
Point to point network 2 device	/dev/hdisk3
Node 2	
Name	KANAGA
Boot IP address / IP label 1	10.1.1.90 / kanagab1
Boot IP address / IP label 2	10.1.2.90 / kanagab2
Persistent address / IP label	9.1.39.90 / kanaga
Point to point network 1	/dev/tty0
Point to point network 2	/dev/hdisk3

Archived

Table 8-2 HACMP resources groups

Resource Group 1	
Name	RG_TSMSRV03
Participating Nodes and priority order	AZOV, KANAGA
Policy	ONLINE ON HOME NODE ONLY, FALLOVER TO NEXT PRIORITY NODE and NEVER FALLBACK
IP address / IP label	9.1.39.74 tsmsrv03
Network name	net_ether_01
Volume group	tsmvg
Applications	TSM Server
Resource Group 2	
Name	RG_ADMCNT01
Participating Nodes and priority order	KANAGA, AZOV
Policy	ONLINE ON HOME NODE ONLY, FALLOVER TO NEXT PRIORITY NODE and NEVER FALLBACK
Volume group	iscvg
IP address	9.1.39.75 admcnt01
Applications	IBM WebSphere Application Server, ISC Help Service, TSM Storage Agent and Client

8.5.1 Pre-installation tasks

Here we do the first configuration steps.

Name resolution and remote connection permissions

Note: We execute all following task on both cluster nodes.

1. At first, we insert all planned entries in the local **`/etc/hosts`** file (Example 8-1).

Note: We prefer local resolution for cluster addresses.

Example 8-1 /etc/hosts file after the changes

```
127.0.0.1       loopback localhost

# Boot network 1
10.1.1.89       azovb1
10.1.1.90       kanagab1

# Boot network 2
10.1.2.89       azovb2
10.1.2.90       kanagab2

# Persistent addresses
9.1.39.89       azov
9.1.39.90       kanaga

# Service addresses
9.1.39.74       tsmsrv03
9.1.39.75       admcnt01
```

2. Next, we inserted the first boot network adapters addresses to enable clcomd communication for initial resource discovery, and cluster configuration into the /usr/es/sbin/etc/cluster/rhosts file. /.rhosts can be used, with *host user* entries, but is suggested to remove it as soon as possible (Example 8-2).

Example 8-2 The edited /usr/es/sbin/etc/cluster/rhosts file

```
azovb1
kanagab1
```

Note: Full resolved iplabels are to be used, or use IP addresses instead.

Software requirement

For up-to-date information, always refer to the readme file that comes with the latest maintenance or patches you are going to install.

We have a prerequisite for HACMP and Tivoli Storage Manager to be installed.

1. The base operating system filesets listed in Example 8-3 are required to be installed prior to HACMP installation.

Example 8-3 The AIX bos filesets that must be installed prior to installing HACMP

```
bos.adt.lib
bos.adt.libm
bos.adt.syscalls
bos.clvm.enh (if you going to use disk hb)
bos.net.tcp.client
```

Archived

```
bos.net.tcp.server
bos.rte.SRC
bos.rte.libc
bos.rte.libcfg
bos.rte.libcur
bos.rte.libpthreads
bos.rte.odm
```

Tip: Only bos.adt.libm, bos.adt.syscalls, and bos.clvm.enh are not installed by default at OS installation time.

2. The AIX command **lslpp** is to be used to verify for filesets installed as in Example 8-4.

Example 8-4 The lslpp -L command

```
lslpp -L bos.adt.lib
azov/: lslpp -L bos.adt.lib
Fileset                      Level  State  Type  Description (Uninstaller)
----------------------------------------------------------------------------
bos.adt.lib               5.3.0.10    A     F    Base Application Development
Libraries
```

3. The RSCT filesets needed for HACMP installation are listed in Example 8-5.

Example 8-5 The RSCT filesets required prior to HACMP installation

```
rsct.basic.hacmp 2.4.0.1
rsct.compat.clients.hacmp 2.4.0.1
rsct.msg.en_US.basic.rte 2.4.0.1
```

Tip: The following versions of RSCT filesets are required:

RSCT 2.2.1.36 or higher is required for AIX 5L V5.1.
RSCT 2.3.3.1 or higher is required for AIX 5L V5.2.
RSCT 2.4.0.0 or higher is required for AIX 5L V5.3.

4. Then the devices.common.IBM.fc.hba-api AIX fileset is required to enable the Tivoli Storage Manager *SAN environment support* functions (Example 8-6).

Example 8-6 The AIX fileset that must be installed for the SAN discovery function

```
devices.common.IBM.fc.hba-api
```

Archived

5. We then install the needed AIX filesets listed above from the AIX installation CD using **smitty installp** fast path. An example of the **installp** usage is shown in “Installation” on page 455.

Device driver installation

We install now the device drivers required for our storage subsystems following subsystems documentation and reboot the systems for the changes to take effect.

Devices will be connected and configured later on, setting up external storage.

snmpd configuration

> **Important:** The following change is not necessary for HACMP Version 5.2 or HACMP Version 5.1 with APAR IY56122 because HACMP Version 5.2 now supports SNMP Version 3.

The SNMP Version 3 (the default on AIX 5.3) will not work with older HACMP versions; you need to run the **fix_snmpdv3_conf** script on each node to add the necessary entries to the /etc/snmpdv3.conf file. This is shown in Example 8-7.

Example 8-7 SNMPD script to switch from v3 to v2 support

```
/usr/es/sbin/cluster/samples/snmp/fix_snmpdv3_conf
```

8.5.2 Serial network setup

> **Note:** When using integrated serial ports, be aware that not all native ports are supported with HACMP serial networks. For example, sa0 could be in use by the service processor. Check the server model announcement letter or search:
>
> http://www.ibm.com

We now configure the RS232 serial line by doing the following activities.

1. Initially, we ensure that we have physically installed the RS232 serial line between the two nodes before configuring it; this should be a “cross” or “null-modem” cable, which is usually ordered with the servers (Example 8-8).

Example 8-8 HACMP serial cable features

```
3124 Serial to Serial Port Cable for Drawer/Drawer
or 3125 Serial to Serial Port Cable for Rack/Rack
```

Or you can use a 9-pin cross cable as shown in Figure 8-4.

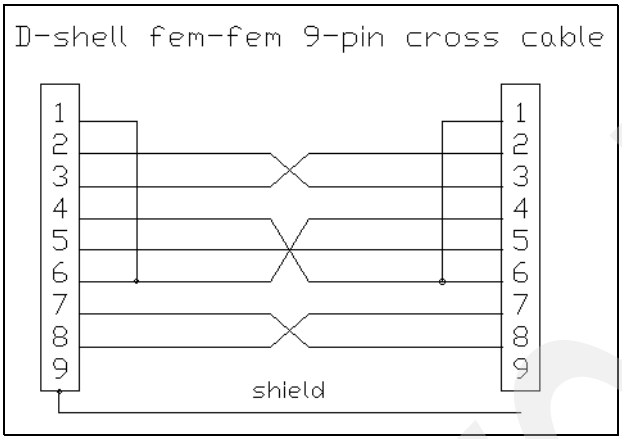

Figure 8-4 9-pin D shell cross cable example

2. We then use the AIX **`smitty tty`** fast path to define the device on each node that will be connected to the RS232 line
3. Next, we select **Add a TTY**.
4. We then select the option, **tty rs232 Asynchronous Terminal**.
5. SMIT prompts you to identify the parent adapter. We use *sa1 Available 01-S2 Standard I/O Serial Port* (on our server serial ports 2 and 3 are supported with *RECEIVE trigger level* set to 0).
6. We then select the appropriate port number and press Enter. The port that you select is the port to which the RS232 cable is connected; we select port 0.
7. We set the login field to DISABLE to prevent **`getty`** processes from spawning on this device.

> **Tip:** In the field, *Flow Control*, leave the default of **xon**, as Topology Services will disable the xon setting when it begins using the device. If xon is not available, then use **none**. Topology Services cannot disable rts, and that setting has (in rare instances) caused problems with the use of the adapter by Topology Services.

8. We will type 0 in *RECEIVE trigger level* as for suggestions found searching `http://www.ibm.com` for the server model.

Archived

9. Then we press Enter (Figure 8-5).

```
                                   Add a TTY

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[TOP]                                                   [Entry Fields]
  TTY type                                             tty
  TTY interface                                        rs232
  Description                                          Asynchronous Terminal
  Parent adapter                                       sa1
* PORT number                                          [0]                    +
  Enable LOGIN                                         disable                +
  BAUD rate                                            []                     +
  PARITY                                               [none]                 +
  BITS per character                                   [8]                    +
  Number of STOP BITS                                  [1]                    +
  TIME before advancing to next port setting           [0]                    +#
  TERMINAL type                                        [dumb]
  FLOW CONTROL to be used                              [xon]                  +
[MORE...31]

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 8-5 tty configuration

Note: Regardless of the baud rate setting of the tty when it is created, all RS232 networks used by HACMP are brought up by RSCT with a default baud rate of 38400. Some RS232 networks that are extended to longer distances and some CPU load conditions will require the baud rate to be lowered from the default of 38400.

For more information, see 8.7.5, “Further cluster customization tasks” on page 448 of this book, and refer to the section “Changing an RS232 Network Module Baud Rate” in *Managing the Cluster Topology,* included in the *Administration and Troubleshooting Guide*.

Test communication over the serial line

To test communication over the serial line after creating the tty device on both nodes:

1. On the first node, we enter the AIX command **`stty < /dev/ttyx`** where /dev/ttyx is the newly added tty device.
2. Then the command line on the first node should hang until the second node receives a return code.
3. Now, on the second node, we enter the AIX command **`stty < /dev/ttyx`** where /dev/ttyx is the newly added tty device.
4. Then if the nodes are able to communicate over the serial line, both nodes display their tty settings and return to the prompt.

Note: This is a valid communication test of a newly added serial connection before the HACMP for AIX /usr/es/sbin/cluster/clstrmgr daemon has been started. This test is not valid when the HACMP daemon is running. The original settings are restored when the HACMP for AIX software exits.

8.5.3 External storage setup

Next we configure external storage resources (devices) used for Tivoli Storage Manager server, *Integrated Solutions Console*, Administration Center, and disk heartbeat functions.

Tape drive names

We need to ensure that the removable media storage devices are configured with the same names on the production and standby nodes.

We may have to define dummy devices on one of the nodes to accomplish this, such as the case of having an internal tape drive on one node only.

To define a dummy device, we can follow these steps:

1. Issue the command **`smit devices`** and go through the **`smit`** panels to define the device.
2. Choose an unused SCSI address for the device.
3. Rather than pressing Enter on the last panel to define the device, press F6 instead to obtain the command that **`smit`** is about to execute.
4. Exit from **`smit`** and enter the same command on the command line, adding the **`-d`** flag to the command. If you attempt to define the device using **`smit`**, the attempt will fail because there is no device at the unused SCSI address you have chosen.

Provide volumes access

Next we perform the following tasks to verify and configure the resources and devices, but we do not go into fine detail with the hardware related tasks. Rather, we just mention the higher level topics:

1. We verify servers adapter cards, storage and tape subsystems and SAN switches for planned firmware levels or update as needed.
2. Then we connect fibre connections from servers adapters and storage subsystems to the SAN switches.
3. We configure zoning as planned to give storage and tape subsystems to servers.

Archived

4. Then we run **cfgmgr** on both nodes to configure tape storage subsystem and make the disk storage subsystem recognize the host adapters.
5. Tape storage devices are now available on both servers; **lsdev** output in Example 8-9.

Example 8-9 lsdev command for tape subsystems

```
azov:/# lsdev -Cctape
rmt0 Available 1Z-08-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt1 Available 1D-08-02 IBM 3580 Ultrium Tape Drive (FCP)
smc0 Available 1Z-08-02 IBM 3582 Library Medium Changer (FCP)

kanaga:/# lsdev -Cctape
rmt1 Available 1Z-08-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt0 Available 1D-08-02 IBM 3580 Ultrium Tape Drive (FCP)
smc0 Available 1Z-08-02 IBM 3582 Library Medium Changer (FCP)
```

6. On the disk storage subsystem, we can configure servers host adapters and assign planned LUNs to them now.

In Figure 8-6 we show the configuration of the DS4500 we used in our lab.

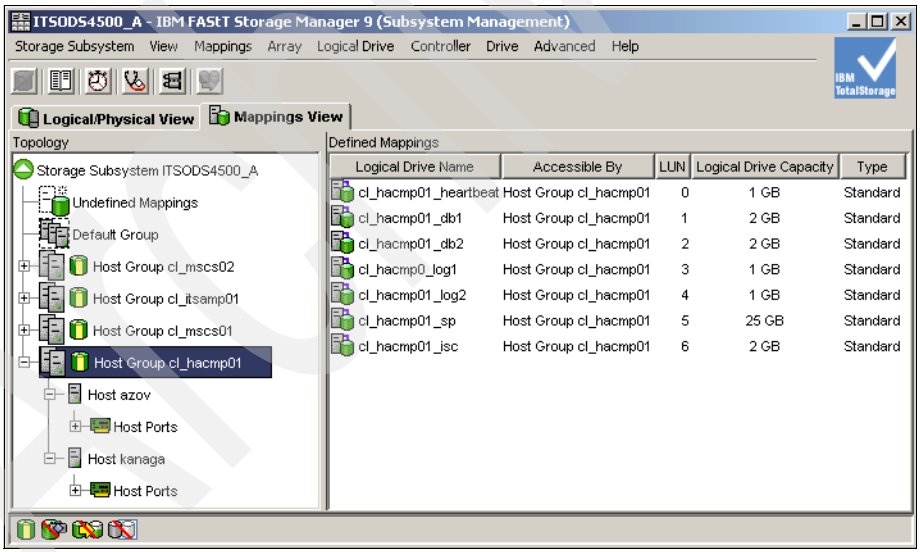

Figure 8-6 DS4500 configuration layout.

7. Now we run **cfgmgr -S** on the first server.

8. We verify the volumes' availability with the **lspv** command (Example 8-10).

Example 8-10 The lspv command output

```
hdisk0          0009cd9aea9f4324                    rootvg          active
hdisk1          0009cd9af71db2c1                    rootvg          active
hdisk2          0009cd9ab922cb5c                    None
hdisk3          none                                None
hdisk4          none                                None
hdisk5          none                                None
hdisk6          none                                None
hdisk7          none                                None
hdisk8          none                                None
```

9. We identify storage subsystems configured LUNs to operating systems physical volumes using the **lscfg** command (Example 8-11).

Example 8-11 The lscfg command

```
azov/: lscfg -vpl hdisk4
  hdisk4          U0.1-P2-I4/Q1-W200400A0B8174432-L1000000000000 1742-900
(900) Disk Array Device
```

Create a non-concurrent shared volume group

We now create a shared volume and the shared filesystems required for the *Tivoli Storage Manager Server*. This same procedure will also be used for setting up *the storage resources for the Integrated Solutions Console* and *Administration Center*.

1. We will create the non-concurrent shared volume group on a node, using the **mkvg** command (Example 8-12).

Example 8-12 mkvg command to create the volume group

```
mkvg -n -y tsmvg -V 50 hdisk4 hdisk5 hdisk6 hdisk7 hdisk8
```

Important:

- Do not activate the volume group AUTOMATICALLY at system restart. Set to **no** (-n flag) so that the volume group can be activated as appropriate by the cluster event scripts.
- Use the **lvlstmajor** command on each node to determine a free major number common to all nodes.
- If using SMIT, **smitty vg** fast path, use the default fields that are already populated wherever possible, unless the site has specific requirements.

Archived

2. Then we create the logical volumes using the **mklv** command. This will create the logical volumes for the jfs2log, Tivoli Storage Manager disk storage pools, and configuration files on the RAID1 volume (Example 8-13).

Example 8-13 mklv commands to create logical volumes

```
/usr/sbin/mklv -y tsmvglg -t jfs2log tsmvg 1 hdisk8
/usr/sbin/mklv -y tsmlv -t jfs2 tsmvg 1 hdisk8
/usr/sbin/mklv -y tsmdp1lv -t jfs2 tsmvg 790 hdisk8
```

3. Next, we create the logical volumes for Tivoli Storage Manager database and log files on the RAID0 volumes (Example 8-14).

Example 8-14 mklv commands used to create the logical volumes

```
/usr/sbin/mklv -y tsmdb1lv -t jfs2 tsmvg 63 hdisk4
/usr/sbin/mklv -y tsmdbmr1lv -t jfs2 tsmvg 63 hdisk5
/usr/sbin/mklv -y tsmlg1lv -t jfs2 tsmvg 31 hdisk6
/usr/sbin/mklv -y tsmlgmr1lv -t jfs2 tsmvg 31 hdisk7
```

4. We then format the jfs2log device, to be used when we create the filesystems (Example 8-15).

Example 8-15 The logform command

```
logform /dev/tsmvglg
logform: destroy /dev/rtsmvglg (y)?y
```

5. Then, we create the filesystems on the previously defined logical volumes using the **crfs** command (Example 8-16).

Example 8-16 The crfs commands used to create the filesystems

```
/usr/sbin/crfs -v jfs2 -d tsmlv -m /tsm/files -A no -p rw -a agblksize=4096
/usr/sbin/crfs -v jfs2 -d tsmdb1lv -m /tsm/db1 -A no -p rw -a agblksize=4096
/usr/sbin/crfs -v jfs2 -d tsmdbmr1lv -m /tsm/dbmr1 -A no -p rw -a
agblksize=4096
/usr/sbin/crfs -v jfs2 -d tsmlg1lv -m /tsm/lg1 -A no -p rw -a agblksize=4096
/usr/sbin/crfs -v jfs2 -d tsmlgmr1lv -m /tsm/lgmr1 -A no -p rw -a
agblksize=4096
/usr/sbin/crfs -v jfs2 -d tsmdp1lv -m /tsm/dp1 -A no -p rw -a agblksize=4096
```

6. We then vary offline the shared volume group (Example 8-17).

Example 8-17 The varyoffvg command

```
varyoffvg tsmvg
```

7. We then run **cfgmgr -S** on second node, and check for the presence of tsmvg's PVIDs on the second node.

Important: If PVIDs are not present, we issue the **chdev -l** *hdiskname* **-a pv=yes** for the required physical volumes:

```
chdev -l hdisk4 -a pv=yes
```

8. We then import the volume group *tsmvg* on the second node (Example 8-18).

Example 8-18 The importvg command

```
importvg -y tsmvg -V 50 hdisk4
```

9. Then, we change the *tsmvg* volume group, so it will not varyon (activate) at boot time (Example 8-19).

Example 8-19 The chvg command

```
chvg -a n tsmvg
```

10. Lastly, we varyoff the tsmvg volume group on the second node (Example 8-20).

Example 8-20 The varyoffvg command

```
varyoffvg tsmvg
```

Creating an enhanced concurrent capable volume group

We will now create a non-concurrent shared volume group on a node, using the AIX command line.

This volume group is to be used for the disk heartbeat.

Important: Use the **lvlstmajor command** on each node to determine a unique major number common to all nodes.

1. We create the volume group using the **mkvg** command (Example 8-21).

Example 8-21 The mkvg command

```
azov:/# mkvg -n -y diskhbvg -V 55 hdisk3
```

2. Then, we change the *diskhbvg* volume group into an Enhanced Concurrent Capable volume group using the **chvg** command (Example 8-22).

Example 8-22 The chvg command

```
azov:/# chvg -C diskhbvg
```


Archived

3. Next, we vary offline the diskhbvg volume from the first node using the **varyoffvg** command (Example 8-23).

Example 8-23 The varyoffvg command

```
varyoffvg diskhbvg
```

4. Lastly, we import the diskhbvg volume group on the second node using the **importvg** command (Example 8-24).

Example 8-24 The importvg command

```
kanaga/: importvg -y diskhbvg -V 55 hdisk3
synclvodm: No logical volumes in volume group diskhbvg.
diskhbvg
0516-783 importvg: This imported volume group is concurrent capable.
Therefore, the volume group must be varied on manually.
```

8.6 Installation

Here we will install the HACMP code.

For **installp** usage examples, see: “Installation” on page 455.

8.6.1 Install the cluster code

For up-to-date information, always refer to the readme file that comes with latest maintenance or patches you are going to install.

With the standard AIX filesets installation method (**installp**), install on both nodes the required HACMP V5.2 filesets at the latest level:

```
cluster.es.client.lib
cluster.es.client.rte
cluster.es.client.utils
cluster.es.clvm.rte
cluster.es.cspoc.cmds
cluster.es.cspoc.dsh
cluster.es.cspoc.rte
cluster.es.server.diag
cluster.es.server.events
cluster.es.server.rte
cluster.es.server.utils
cluster.license
```

Note: AIX 5L V5.3 (5765-G03); HACMP V5.2 requires **IY58496**.

Archived

Once you have installed HACMP, check to make sure you have the required APAR applied with the **instfix** command.

Example 8-25 shows the output on a system having APAR IY58496 installed.

Example 8-25 APAR installation check with instfix command.

```
instfix -ick IY58496
#Keyword:Fileset:ReqLevel:InstLevel:Status:Abstract
IY58496:cluster.es.client.lib:5.2.0.1:5.2.0.1:=:Base fixes for hacmp 5.2.0
IY58496:cluster.es.client.rte:5.2.0.1:5.2.0.1:=:Base fixes for hacmp 5.2.0
IY58496:cluster.es.client.utils:5.2.0.1:5.2.0.1:=:Base fixes for hacmp 5.2.
IY58496:cluster.es.cspoc.cmds:5.2.0.1:5.2.0.1:=:Base fixes for hacmp 5.2.0
IY58496:cluster.es.cspoc.rte:5.2.0.1:5.2.0.1:=:Base fixes for hacmp 5.2.0
IY58496:cluster.es.server.diag:5.2.0.1:5.2.0.1:=:Base fixes for hacmp 5.2.0
IY58496:cluster.es.server.events:5.2.0.1:5.2.0.1:=:Base fixes for hacmp 5.2
IY58496:cluster.es.server.rte:5.2.0.1:5.2.0.1:=:Base fixes for hacmp 5.2.0
IY58496:cluster.es.server.utils:5.2.0.1:5.2.0.1:=:Base fixes for hacmp 5.2.
```

8.7 HACMP configuration

Cluster information will be entered on **one node only** between synchronizations.

> **Tip:** We suggest choosing a “primary” node and then using this node to enter all the cluster information. This will help you avoid losing configuration data or incurring inconsistencies.

Network adapters configuration

We will now configure our network adapters with the boot addresses, from Table 8-1 on page 429.

During these steps, we will require an alternative network connection to telnet to the servers or to login from a local console, as our network connection will be severed.

> **Attention:** Make a note of the default router address and other routing table entries. This is due to the ip address changes deleting the routing information, which will have to be added back later.

1. We use the **smitty chinet** fast path.
2. Then, on the **Available Network Interfaces** panel, we select our first targeted network adapter.

3. We fill in the required fields and press Enter (Figure 8-7).

```
                  Change / Show a Standard Ethernet Interface

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
  Network Interface Name                                 en0
  INTERNET ADDRESS (dotted decimal)                     [10.1.1.89]
  Network MASK (hexadecimal or dotted decimal)          [255.255.255.0]
  Current STATE                                          up                      +
  Use Address Resolution Protocol (ARP)?                 yes                     +
  BROADCAST ADDRESS (dotted decimal)                    []
  Interface Specific Network Options
     ('NULL' will unset the option)
    rfc1323                                             []
    tcp_mssdflt                                         []
    tcp_nodelay                                         []
    tcp_recvspace                                       []
    tcp_sendspace                                       []
  Apply change to DATABASE only                          no                      +

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 8-7 boot address configuration

4. We repeat the above steps for the two adapters of both servers.

8.7.1 Initial configuration of nodes

We will now configure the Cluster Name, Cluster Node Name, and the initial communication paths between the nodes:

1. On the AIX command line, enter the command **`smitty hacmp`**.
2. Within SMIT, select the **Extended Configuration** option.
3. Next, select the **Extended Topology Configuration** option.
4. Then, select the **Configure an HACMP Cluster** option.
5. Lastly, select the **Add/Change/Show an HACMP Cluster** option.
6. Then we enter our Cluster Name, which is *cl_hacmp01*.
7. Press Enter to complete the configuration (Figure 8-8).

Archived

```
                        Add/Change/Show an HACMP Cluster

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* Cluster Name                                        [cl_hacmp01]

  NOTE: HACMP must be RESTARTED
  on all nodes in order for change to take effect

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 8-8 Define cluster example.

8. Then we go back to the **Extended Topology Configuration** panel (3 layers back).
9. We select the **Configure HACMP Nodes** option.
10. Then we select the **Add a Node to the HACMP Cluster** option.
11. We fill in the **Node Name** field.
12. For the next field below, we press the F4 key to select from a list of available communication paths to the node.
13. Press Enter to complete the change (Figure 8-9).

Archived

```
                        Add a Node to the HACMP Cluster

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* Node Name                                            [azov]
  Communication Path to Node                           [azovb1]                +



F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 8-9 An add cluster node example

14. We now go back thorough the SMIT menus using the F3 key, and then repeat the process for the second node.

8.7.2 Resource discovery

Now we will use the cluster software discovery function to have HACMP locating the available hardware resources which are available to the nodes.

1. From the AIX command line, we enter the **`smitty hacmp`** command.
2. We select **Extended Configuration** option.
3. Then, we select the **Discover HACMP-related Information from Configured Nodes** option.

> **Note:** The Discover utility runs for few seconds (depending on the configuration) and ends, showing an OK status.

8.7.3 Defining HACMP interfaces and devices

The cluster should have more than one network path to avoid a single point of failure for the high available service.

Network paths configured to the cluster are used for heartbeat also. To improve HACMP problem determination and fault isolation, we use both IP and non-IP based networks as heartbeat paths.

Archived

Now we are going to configure planned communication devices and interfaces.

> **Note:** Configuring the first network interface or communication device for a point to point network makes a corresponding cluster network object configured too.

Configuring the non-IP communication devices

We now configure the HACMP discovered serial and disk devices for the cluster heartbeat using SMIT.

1. Enter the AIX **`smitty hacmp`** command.
2. Select the **Extended Configuration** option.
3. Then, select the **Extended Topology Configuration** option.
4. Next, select **Configure HACMP Communication Interfaces/Devices** option.
5. Select the **Discovered** option (Figure 8-10).
6. Select the **Communications Devices type** from the selection options.
7. The screen **Select Point-to-Point Pair of Discovered Communication Devices to Add** appears; devices that are already added to the cluster are filtered from the pick list.
8. Now select both devices for the same network at once and press Enter.
9. We then repeat this process for the second serial network type. In our cluster we configure two point-to-point network types, rs232 and disk.

```
                 Configure HACMP Communication Interfaces/Devices

Move cursor to desired item and press Enter.
  Add Communication Interfaces/Devices
  Change/Show Communication Interfaces/Devices
  Remove Communication Interfaces/Devices
  +-------------------------------------------------------------------------+
  |  Select Point-to-Point Pair of Discovered Communication Devices to Add  |
  |                                                                         |
  | Move cursor to desired item and press F7. Use arrow keys to scroll.     |
  |     ONE OR MORE items can be selected.                                  |
  | Press Enter AFTER making all selections.                                |
  |                                                                         |
  |   # Node                                Device   Device Path   Pvid     |
  | >   azov                                hdisk3   /dev/hdisk3   0009cddaa6d |
  | >   kanaga                              hdisk3   /dev/hdisk3   0009cddaa6d |
  |     azov                                tty0     /dev/tty0             |
  |     kanaga                              tty0     /dev/tty0             |
  |                                                                         |
  | F1=Help                 F2=Refresh               F3=Cancel              |
  | F7=Select               F8=Image                 F10=Exit               |
F1| Enter=Do                /=Find                   n=Find Next            |
F9+-------------------------------------------------------------------------+^L
```

Figure 8-10 Configure HACMP Communication Interfaces/Devices panel

Configuring the IP interfaces

We will now configure the HACMP IP-based discovered interfaces for the cluster using SMIT:

1. Enter the AIX command **`smitty hacmp`**.
2. Then select the **Extended Topology Configuration** option.
3. Next, select the **Configure HACMP Communication Interfaces/Devices** option.
4. We then select the **Add Discovered Communication Interface and Devices** option.
5. Now, select **Add Communication Interfaces/Devices** option panel.
6. We then select **Discovered Communication Interface and Devices** panel.
7. Then, we select the **Communication Interfaces** option.
8. Lastly, we select **ALL**.
9. We mark all the planned network interfaces (see Table 8-1 on page 429).
10. Press Enter to complete the selection processing (Figure 8-11).

```
Mo+-------------------------------------------------------------------------+
  |      Select one or more Discovered Communication Interfaces to Add      |
  |                                                                         |
  | Move cursor to desired item and press F7.                               |
  |     ONE OR MORE items can be selected.                                  |
  | Press Enter AFTER making all selections.                                |
  |                                                                         |
  |   # Network / Node                                                      |
  |   #       Interface          IP Label                        IP Address |
  |                                                                         |
  |   # net_ether_01 / azov                                                 |
  | >         en0                azovb1                           10.1.1.89 |
  | >         en1                azovb2                           10.1.2.89 |
  |                                                                         |
  |   # net_ether_01 / kanaga                                               |
  | >         en0                kanagab1                         10.1.1.90 |
  | >         en1                kanagab2                         10.1.2.90 |
  |                                                                         |
  | F1=Help                 F2=Refresh              F3=Cancel               |
  | F7=Select               F8=Image                F10=Exit                |
F1| Enter=Do                /=Find                  n=Find Next             |
F9+-------------------------------------------------------------------------+
```

Figure 8-11 Selecting communication interfaces.

8.7.4 Persistent addresses

Next we implement persistent addressing to enable network connectivity for a cluster node regardless of the service state or a single adapter failure situation.

Archived

We accomplish this by entering **smitty hacmp** on an AIX command line:

1. Then, select **Extended Topology Configuration**.
2. We then select **Configure HACMP Persistent Node IP Label/Addresses**.
3. Then we select the **Add a Persistent Node IP Label/Address** option.
4. We then select the first node.
5. Then we pick from list the network name.
6. And then we pick the plannednode persistent IP Label/Address (see Table 8-1 on page 429).
7. We then press Enter to complete the selection process (Figure 8-12).
8. Lastly, we repeat the process for the second node.

```
                    Add a Persistent Node IP Label/Address

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                      [Entry Fields]
* Node Name                                           azov
* Network Name                                      [net_ether_01]          +
* Node IP Label/Address                             [azov]                  +



F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 8-12 The Add a Persistent Node IP Label/Address panel

8.7.5 Further cluster customization tasks

Here we go on to other tasks that are highly dependent on the solution design and available HW.

Refer to the *HACMP Planning and Installation Guide* for further explanation about these tasks.

Configure network modules

As we are not interested, due to the nature or our application, in an extremely sensitive cluster, we chose to lower the *Failure Detection Rate* for utilized network modules, avoiding unwanted takeovers in case of particular events such as high CPU load.

1. We enter **smitty cm_config_networks** from the AIX command line.
2. Then, we choose **Change a Network Module using Predefined Values**.

3. We then select **diskhb**.
4. Next, we change **Failure Detection Rate** to **Slow**.
5. We then press Enter to complete the processing.
6. We then repeat the process for the **ether** and **rs232** networks.

Lower RS232 speed

In situations when the CPU load is high, the default RSCT baud rate is too high (for serial networks this is 38400 bps). In the case of some integrated adapters or long distance connections, this can lead to problems. We choose to lower this rate to 9600 bps.

1. We enter `smitty cm_config_networks` from the AIX command line.
2. Then we choose **Change a Network Module using Custom Values**.
3. Next, we select **rs232**.
4. Then we type in the value of “9600” in **Parameter** field.
5. Lastly, we press Enter to complete the processing.

Change/Show syncd frequency

Here we change the frequency with which I/O buffers are flushed. For nodes in HACMP clusters, the recommended frequency is 10.

1. We enter `smitty cm_tuning_parms_chsyncd` on the AIX command line.
2. Then we change the **syncd frequency (in seconds)** field value to 10.

Configure automatic error notification

The Automatic Error Notification utility will discover single points of failure. For these single points of failure, this utility will create an *Error Notify Method* that is used to react to errors with a takeover.

1. We enter `smitty hacmp` on the AIX command line.
2. Then, we select **Problem Determination Tools**.
3. And then select the **HACMP Error Notification** option.
4. Next, we select the **Configure Automatic Error Notification** option.
5. We then select the **Add Error Notify Methods for** Cluster Resources option.
6. We then press Enter and the processing completes.
7. Once this completes, we go back up to the **Configure Automatic Error Notification** option.
8. We then use **List Error Notify Methods for Cluster Resources** to verify the configured Notify methods.

Archived

Note: If a non-mirrored logical volume exists, Takeover Notify methods are configured for the used physical volumes. Take, for example, the dump logical volume that has to be **not mirrored**; in this case the simplest way to exit is to have it mirrored **only while** the automatic error notification utility runs.

Here we have completed the base cluster infrastructure. The next steps are resources configuration and cluster testing. Those steps are described in Chapter 9, “AIX and HACMP with IBM Tivoli Storage Manager Server” on page 451, where we install the Tivoli Storage Manager server, configure storage and network resources, and make it an HACMP highly available application.

Archived

9

AIX and HACMP with IBM Tivoli Storage Manager Server

In this chapter we provide detailed coverage, including an overview, planning, installing, configuring, testing, and troubleshooting of Tivoli Storage Manager V5.3, as an application resource controlled by HACMP.

© Copyright IBM Corp. 2005. All rights reserved.

9.1 Overview

Here is a brief overview of IBM Tivoli Storage Manager 5.3 enhancements.

9.1.1 Tivoli Storage Manager Version 5.3 new features overview

IBM Tivoli Storage Manager V5.3 is designed to provide some significant improvements to ease of use as well as ease of administration and serviceability. These enhancements help you improve the productivity of personnel who are administering and using IBM Tivoli Storage Manager. Additionally, the product is easier to use for new administrators and users.

- Improved application availability:
  - IBM Tivoli Storage Manager for Space Management: HSM for AIX JFS2,enhancements to HSM for AIX and Linux GPFS
  - IBM Tivoli Storage Manager for application products update
- Optimized storage resource utilization:
  - Improved device management, SAN attached device dynamic mapping, native STK ACSLS drive sharing and LAN-free operations, improved tape checkin, checkout, and label operations, and new device support
  - Disk storage pool enhancements, collocation groups, proxy node support, improved defaults, reduced LAN-free CPU utilization, parallel reclamation, and migration
- Enhanced storage personnel productivity:
  - New Administrator Web GUI
  - Task-oriented interface with wizards to simplify tasks such as scheduling, managing server maintenance operations (storage pool backup, migration, reclamation), and configuring devices
  - Health monitor which shows status of scheduled events, the database and recovery log, storage devices, and activity log messages
  - Calendar-based scheduling for increased flexibility of client and administrative schedules
  - Operational customization for increased ability to control and schedule server operations

Archived

Server enhancements, additions, and changes

This section lists all the functional enhancements, additions, and changes for the IBM Tivoli Storage Manager Server introduced in Version 5.3.

Here are the latest changes:

- ACSLS Library Support Enhancements
- Accurate SAN Device mapping for UNIX Servers
- ACSLS Library Support Enhancements
- Activity Log Management
- Check-In and Check-Out Enhancements
- Collocation by Group
- Communications Options
- Database Reorganization
- Disk-only Backup
- Enhancements for Server Migration and Reclamation Processes
- IBM 3592 WORM Support
- Improved Defaults
- Increased Block Size for Writing to Tape
- LAN-free Environment Configuration
- NDMP Operations
- Net Appliance SnapLock Support
- New Interface to Manage Servers: Administration Center
- Server Processing Control in Scripts
- Simultaneous Write Inheritance Improvements
- Space Triggers for Mirrored Volumes
- Storage Agent and Library Sharing Fallover
- Support for Multiple IBM Tivoli Storage Manager Client Nodes
- IBM Tivoli Storage Manager Scheduling Flexibility

Client enhancements, additions and changes

This section lists all the functional enhancements, additions, and changes for the IBM Tivoli Storage Manager Backup Archive Client introduced in Version 5.3.

Here are the latest changes:

- Include-exclude enhancements
- Enhancements to query schedule command
- IBM Tivoli Storage Manager Administration Center
- Support for deleting individual backups from a server file space
- Optimized option default values
- New links from the backup-archive client Java GUI to the IBM Tivoli Storage Manager and Tivoli Home Pages
- New options, Errorlogmax and Schedlogmax, and DSM_LOG environment variable changes
- Enhanced encryption

Archived

- Dynamic client tracing
- Web client enhancements
- Client node proxy support [asnodename]
- Java GUI and Web client enhancements
- IBM Tivoli Storage Manager backup-archive client for HP-UX Itanium 2
- Linux for zSeries offline image backup
- Journal based backup enhancements
- Single drive support for Open File Support (OFS) or online image backups.

9.1.2 Planning for storage and database protection

In this section we will give some considerations how to plan the storage and Tivoli Storage Manager database protection. For more details, please refer to: “Protecting and Recovering Your Server” in the *Administrators Guide.*

For this configuration example we chose to have:

Tivoli Storage Manager server:

- Code installed under rootvg filesystems /usr on both nodes
- Tivoli Storage Manager mirroring for database and log volumes
- RAID0 shared disks volumes configured on separate storage subsystem arrays for database and log volumes copies
  - /tsm/db1
  - /tsm/db1mr
  - /tsm/lg1
  - /tsm/lg1mr
- Database and log writes set to sequential (which disables DBPAGESHADOW)
- Log mode set to RollForward
- RAID1 shared disk volumes for configuration files and disk storage pools.
  - /tsm/files
  - /tsm/dp1

Tivoli Storage Manager Administration Center

Note: The Administration Center can be a critical application for environments where administrator and operators are not confident with the *IBM Tivoli Storage Manager Command Line Administrative Interface*. So we decided to experiment with a clustered installation, even if it is not currently supported.

Archived

- RAID1 shared disk volume for both code and data (server connections and ISC user definitions) under a shared filesystem that we are going to create and activate before going on to ISC code installation.
  - /opt/IBM/ISC

The physical layout is shown in 8.5, “Lab setup” on page 427.

9.2 Lab setup

Here we use the lab setup described in Chapter 8, “Establishing an HACMP infrastructure on AIX” on page 417.

9.3 Installation

Next we install Tivoli Storage Manager server and client code.

9.3.1 Tivoli Storage Manager Server AIX filesets

For up-to-date information, always refer to the Tivoli Storage Manager Web pages under `http://www.ibm.com/tivoli` orsee the readme file that comes with the latest maintenance or patches you are going to install.

Server code

Use normal AIX filesets install procedures (**`installp`**) to install server code filesets according to your environment at the latest level on both cluster nodes.

32-bit hardware, 32-bit AIX kernel

tivoli.tsm.server.com
tivoli.tsm.server.rte
tivoli.tsm.msg.en_US.server
tivoli.tsm.license.cert
tivoli.tsm.license.rte
tivoli.tsm.webcon
tivoli.tsm.msg.en-US.devices
tivoli.tsm.devices.aix5.rte

Archived

64-bit hardware, 64-bit AIX kernel

tivoli.tsm.server.com
tivoli.tsm.server.aix5.rte64
tivoli.tsm.msg.en_US.server
tivoli.tsm.license.cert
tivoli.tsm.license.rte
tivoli.tsm.webcon
tivoli.tsm.msg.en-US.devices
tivoli.tsm.devices.aix5.rte
tivoli.tsm.devices.aix5.rte

64-bit hardware, 32-bit AIX kernel

tivoli.tsm.server.com
tivoli.tsm.server.rte
tivoli.tsm.msg.en_US.server
tivoli.tsm.license.cert
tivoli.tsm.license.rte
tivoli.tsm.webcon
tivoli.tsm.msg.en-US.devices
tivoli.tsm.devices.aix5.rte

9.3.2 Tivoli Storage Manager Client AIX filesets

Important: It is necessary to install the *Command Line Administrative Interface* during this process (**dsmadmc** command).

Even if we have no plans to use the Tivoli Storage Manager client, we still need to have ,these components installed on both servers, as the scripts to be configured within HACMP for starting, stopping and eventually monitoring the server require the **dsmadmc** command.

tivoli.tsm.client.api.32bit
tivoli.tsm.client.ba.32bit.base
tivoli.tsm.client.ba.32bit.common
tivoli.tsm.client.ba.32bit.web

9.3.3 Tivoli Storage Manager Client Installation

We will install the Tivoli Storage Manager client into the default location of /usr/tivoli/tsm/client/ba/bin and the API into /usr/tivoli/tsm/client/api/bin on all systems in the cluster.

Archived

1. First we change into the directory which holds our installation images, and issue the **`smitty installp`** AIX command as shown in Figure 9-1.

```
                         Install and Update Software

Move cursor to desired item and press Enter.

  Install Software
  Update Installed Software to Latest Level (Update All)
  Install Software Bundle
  Update Software by Fix (APAR)
  Install and Update from ALL Available Software


F1=Help             F2=Refresh          F3=Cancel           F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 9-1 The smit install and update panel

2. Then, for the input device, we used a **dot**, implying the current directory, as shown in Figure 9-2.

```
                Install and Update from ALL Available Software

Type or select a value for the entry field.
Press Enter AFTER making all desired changes.

                                                      [Entry Fields]
* INPUT device / directory for software               [.]                +


F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 9-2 Launching SMIT from the source directory, only dot (.) is required

Archived

3. For the next **smit** panel, we select a LIST using the F4 key.
4. We then select the required filesets to install using the F7 key, as seen in Figure 9-3.

```
                         SOFTWARE to install

Move cursor to desired item and press F7. Use arrow keys to scroll.
    ONE OR MORE items can be selected.
Press Enter AFTER making all selections.


  tivoli.tsm.books.en_US.client                                      ALL
   + 5.3.0.0  TSM BOOKS - Using UNIX Clients - HTML Format
   + 5.3.0.0  TSM BOOKS - Using UNIX Clients - PDF Format

  tivoli.tsm.client.api.32bit                                        ALL
>  + 5.3.0.0  TSM Client - Application Programming Interface 32bit

  tivoli.tsm.client.api.64bit                                        ALL
   + 5.3.0.0  TSM Client - 64bit Application Programming Interface

  tivoli.tsm.client.ba                                               ALL
>  + 5.3.0.0  TSM Client - Backup/Archive Base Files
>  + 5.3.0.0  TSM Client - Backup/Archive Common Files
>  + 5.3.0.0  TSM Client - Backup/Archive WEB Client
   + 5.3.0.0  TSM Client - IMAGE Backup Client
   + 5.3.0.0  TSM Client - NAS Backup Client

F1=Help                 F2=Refresh              F3=Cancel
F7=Select               F8=Image                F10=Exit
Enter=Do                /=Find                  n=Find Next
```

Figure 9-3 AIX installp filesets chosen: Tivoli Storage Manager client installation

Archived

5. After the selection and pressing enter, we change the default **smit** panel options to allow for a detailed preview first, as shown in Figure 9-4.

```
                 Install and Update from ALL Available Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* INPUT device / directory for software                  .
* SOFTWARE to install                                   [tivoli.tsm.client.api> +
  PREVIEW only? (install operation will NOT occur)       yes                    +
  COMMIT software updates?                               no                     +
  SAVE replaced files?                                   no                     +
  AUTOMATICALLY install requisite software?              yes                    +
  EXTEND file systems if space needed?                   yes                    +
  OVERWRITE same or newer versions?                      no                     +
  VERIFY install and check file sizes?                   no                     +
  DETAILED output?                                       yes                    +
  Process multiple volumes?                              yes                    +
  ACCEPT new license agreements?                         no                     +
  Preview new LICENSE agreements?                        no                     +

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 9-4 Changing the defaults to preview with detail first prior to installing

6. Following a successful preview, we change the **smit** panel configuration to reflect a detailed and committed installation as shown in Figure 9-5.

```
                 Install and Update from ALL Available Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* INPUT device / directory for software                  .
* SOFTWARE to install                                   [tivoli.tsm.client.api> +
  PREVIEW only? (install operation will NOT occur)       no                     +
  COMMIT software updates?                               yes                    +
  SAVE replaced files?                                   no                     +
  AUTOMATICALLY install requisite software?              yes                    +
  EXTEND file systems if space needed?                   yes                    +
  OVERWRITE same or newer versions?                      no                     +
  VERIFY install and check file sizes?                   no                     +
  DETAILED output?                                       yes                    +
  Process multiple volumes?                              yes                    +
  ACCEPT new license agreements?                         no                     +
  Preview new LICENSE agreements?                        no                     +

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 9-5 The smit panel demonstrating a detailed and committed installation

Archived

7. Finally, we review the installed filesets using the **AIX command lslpp** as shown in Figure 9-6.

```
kanaga/images/TSM/client: lslpp -h |grep -p tsm.client
  tivoli.tsm.client.api.32bit
                  5.3.0.0   COMMIT       COMPLETE     01/24/05     17:16:11

  tivoli.tsm.client.ba.32bit.base
                  5.3.0.0   COMMIT       COMPLETE     01/24/05     17:16:18

  tivoli.tsm.client.ba.32bit.common
                  5.3.0.0   COMMIT       COMPLETE     01/24/05     17:16:18

  tivoli.tsm.client.ba.32bit.web
                  5.3.0.0   COMMIT       COMPLETE     01/24/05     17:16:18
```

Figure 9-6 AIX lslpp command to review the installed filesets

9.3.4 Installing the Tivoli Storage Manager Server software

We will install the Tivoli Storage Manager server into the default location of /usr/tivoli/tsm/server/bin on all systems in the cluster which could host the Tivoli Storage Manager server if a failover were to occur.

1. First we change into the directory which holds our installation images, and issue the **smitty installp** AIX command, which presents the first install panel, as shown in Figure 9-7.

```
                         Install and Update Software

Move cursor to desired item and press Enter.

  Install Software
  Update Installed Software to Latest Level (Update All)
  Install Software Bundle
  Update Software by Fix (APAR)
  Install and Update from ALL Available Software

F1=Help             F2=Refresh          F3=Cancel           F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 9-7 The smit software installation panel

2. Then, for the input device, we used a **dot**, implying the current directory, as shown in Figure 9-8.

```
                 Install and Update from ALL Available Software

Type or select a value for the entry field.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* INPUT device / directory for software                 [.]                +



F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 9-8 The smit input device panel

3. Next, we select the filesets which will be required for our clustered environment, using the F7 key. Our selection is shown in Figure 9-9.

Archived

```
                  Install and Update from ALL Available Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* INPUT device / directory for software                 .
*
                              SOFTWARE to install

    Move cursor to desired item and press F7. Use arrow keys to scroll.
        ONE OR MORE items can be selected.
    Press Enter AFTER making all selections.


    > tivoli.tsm.devices.aix5.rte                                      ALL
      + 5.3.0.0  IBM Tivoli Storage Manager Device Support runtime

    > tivoli.tsm.license.aix5.rte64                                    ALL
      + 5.3.0.0  IBM Tivoli Storage Manager 64 bit  License Registration

    > tivoli.tsm.license.cert                                          ALL
      + 5.3.0.0  IBM Tivoli Storage Manager License Certificates

    > tivoli.tsm.loc.server.ela                                        ALL
      + 5.3.0.0  IBM Tivoli Storage Manager Server E-Lic

    > tivoli.tsm.msg.EN_US.devices                                     ALL
      + 5.3.0.0  IBM Tivoli Storage Manager Devices SMIT Menus, US English

    > tivoli.tsm.msg.EN_US.server                                      ALL
      + 5.3.0.0  IBM Tivoli Storage Manager Server Msgs, US English utf8

    > tivoli.tsm.msg.en_US.devices                                     ALL
      + 5.3.0.0  IBM Tivoli Storage Manager Devices SMIT Menus, US English

    > tivoli.tsm.msg.en_US.server                                      ALL
      + 5.3.0.0  IBM Tivoli Storage Manager Server Msgs, US English

    > tivoli.tsm.server.aix5.rte64                                     ALL
      + 5.3.0.0  IBM Tivoli Storage Manager 64 bit Server Runtime

    > tivoli.tsm.server.com                                            ALL
      + 5.3.0.0  IBM Tivoli Storage Manager Server common services

    > tivoli.tsm.server.webcon                                         ALL
      + 5.3.0.0  IBM Tivoli Storage Manager Web Console intfc

    F1=Help             F2=Refresh            F3=Cancel
F1  F7=Select           F8=Image              F10=Exit
F5  Enter=Do            /=Find                n=Find Next
```

Figure 9-9 The smit selection screen for Tivoli Storage Manager filesets

4. We then press Enter after the selection has been made.
5. On this next panel presented, we change the default values for preview, commit, detailed, accept. This allows us to verify that we have all the prerequisites installed prior to running a commit installation. The changes to these default options are shown in Figure 9-10.

Archived

```
                 Install and Update from ALL Available Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* INPUT device / directory for software                  .
* SOFTWARE to install                                   []                    +
  PREVIEW only? (install operation will NOT occur)       yes                  +
  COMMIT software updates?                               no                   +
  SAVE replaced files?                                   no                   +
  AUTOMATICALLY install requisite software?              yes                  +
  EXTEND file systems if space needed?                   yes                  +
  OVERWRITE same or newer versions?                      no                   +
  VERIFY install and check file sizes?                   no                   +
  DETAILED output?                                       yes                  +
  Process multiple volumes?                              yes                  +
  ACCEPT new license agreements?                         yes                  +
  Preview new LICENSE agreements?                        no                   +

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 9-10 The smit screen showing non-default values for a detailed preview

6. After we successfully complete the preview, we change the installation panel to reflect a detailed, committed installation and accept the new license agreements. This is shown in Figure 9-11.

```
                 Install and Update from ALL Available Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* INPUT device / directory for software                  .
* SOFTWARE to install                                   [tivoli.tsm.devices.ai> +
  PREVIEW only? (install operation will NOT occur)       no                   +
  COMMIT software updates?                               yes                  +
  SAVE replaced files?                                   no                   +
  AUTOMATICALLY install requisite software?              yes                  +
  EXTEND file systems if space needed?                   yes                  +
  OVERWRITE same or newer versions?                      no                   +
  VERIFY install and check file sizes?                   no                   +
  DETAILED output?                                       yes                  +
  Process multiple volumes?                              yes                  +
  ACCEPT new license agreements?                         yes                  +
  Preview new LICENSE agreements?                        no                   +

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 9-11 The final smit install screen with selections and a commit installation

Archived

7. After the installation has been successfully completed, we review the installed filesets from the AIX command line with the **lslpp** command, as shown in Figure 9-12.

```
  tivoli.tsm.loc.server.ela
                 5.3.0.0   COMMIT       COMPLETE     01/24/05     15:22:45

  tivoli.tsm.msg.EN_US.server
                 5.3.0.0   COMMIT       COMPLETE     01/24/05     15:22:42

  tivoli.tsm.msg.en_US.server
                 5.3.0.0   COMMIT       COMPLETE     01/24/05     15:22:44

  tivoli.tsm.server.aix5.rte64
                 5.3.0.0   COMMIT       COMPLETE     01/24/05     15:22:34

  tivoli.tsm.server.com
                 5.3.0.0   COMMIT       COMPLETE     01/24/05     15:22:30

  tivoli.tsm.server.webcon
                 5.3.0.0   COMMIT       COMPLETE     01/24/05     15:22:54

  tivoli.tsm.server.aix5.rte64
                 5.3.0.0   COMMIT       COMPLETE     01/24/05     15:22:42

  tivoli.tsm.server.com
                 5.3.0.0   COMMIT       COMPLETE     01/24/05     15:22:33

kanaga/images/TSM/server:
```

Figure 9-12 AIX lslpp command listing of the server installp images

8. Lastly, we repeat all of these processes on the other cluster node.

9.3.5 Installing the ISC and the Administration Center

The installation of Tivoli Storage Manager Administration Center is a two-step install. First install the Integrated Solutions Console. Then deploy the Tivoli Storage Manager Administration Center into the Integrated Solutions Console. Once both pieces are installed, you will be able to administer Tivoli Storage Manager from a browser anywhere in your network.

In addition, these two software components will be a resource within our HACMP cluster. To achieve this, these software packages will be installed onto shared disk, and on the second node in the Tivoli Storage Manager cluster. This will make this cluster configuration an active/active configuration.

Archived

Shared installation

As planned in “Planning for storage and database protection” on page 454, we are going to install the code on a shared filesystem.

We set up a /opt/IBM/ISC filesystem, as we do for the Tivoli Storage Manager server ones in “External storage setup” on page 436.

Then we can:

- Activate it temporarily by hand with **varyonvg iscvg** and **mount /opt/IBM/ISC** commands o the n primary node, run the code installation, and then deactivate it with **umount /opt/IBM/ISC** and **varyoffvg iscvg** (otherwise the following cluster activities will fail).

Or we can:

- Run the ISC code installation later on, after the /opt/IBM/ISC filesystems have been made available through HACMP and before configuring ISC start and stop scripts as an application server.

9.3.6 Installing Integrated Solutions Console Runtime

Here we install the ISC:

1. First we extract the contents of the file TSM_ISC_5300_AIX.tar (Example 9-1).

Example 9-1 The tar command extraction

```
tar xvf TSM_ISC_5300_AIX.tar
```

2. Then we change directory into iscinstall and run the **setupISC** InstallShield command (Example 9-2).

Example 9-2 setupISC usage

```
setupISC
```

Archived

Note: Depending on what the screen and graphics requirements would be, the following options exist for this installation.

Run one of the following commands to install the runtime:

- For InstallShield wizard install, run: setupISC.
- For console wizard install, run: setupISC -console.
- For silent install, run the following command on a single line:

```
setupISC -silent -W ConfigInput.adminName="<user name>"
```

- Flags:
  - W ConfigInput.adminPass="<user password>"
  - W ConfigInput.verifyPass="<user password>"
  - W PortInput.webAdminPort="<web administration port>"
  - W PortInput.secureAdminPort="<secure administration port>"
  - W MediaLocationInput.installMediaLocation="<media location>"
  - P ISCProduct.installLocation="<install location>"

Note: The installation process can take anywhere from 30 minutes to 2 hours to complete. The time to install depends on the speed of your processor and memory.

The following screen captures are for the Java based installation process:

1. We click **Next** on the Welcome message panel (Figure 9-13).

Archived

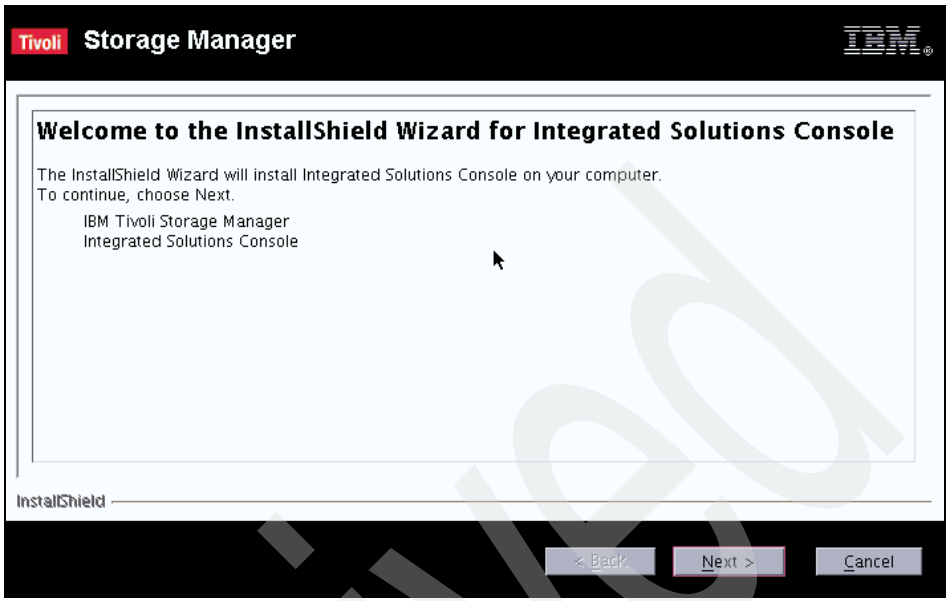

Figure 9-13 ISC installation screen

2. We accept the license agreement and click **Next** on **License Agreement** pane (Figure 9-14).

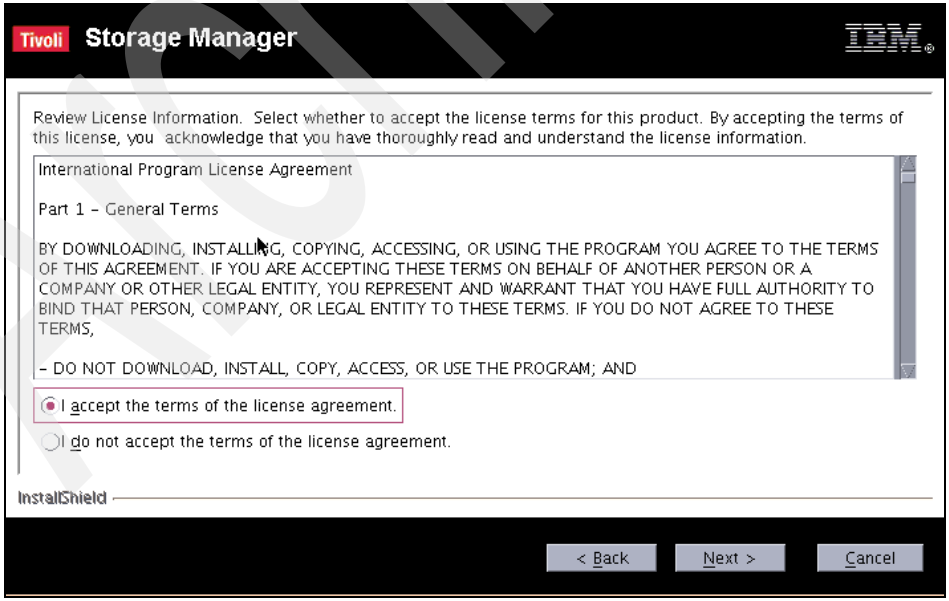

Figure 9-14 ISC installation screen, license agreement

Archived

3. We accept the proposed location for install files and click **Next** on **Source path** panel (Figure 9-15).

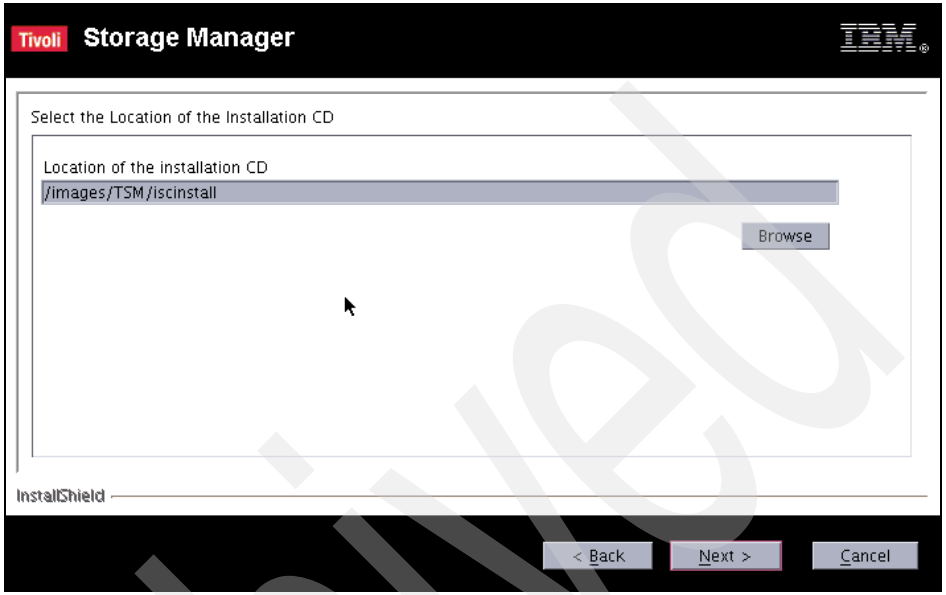

Figure 9-15 ISC installation screen, source path

4. We verify proposed installation path and click **Next** on the install location panel (Figure 9-16).

Archived

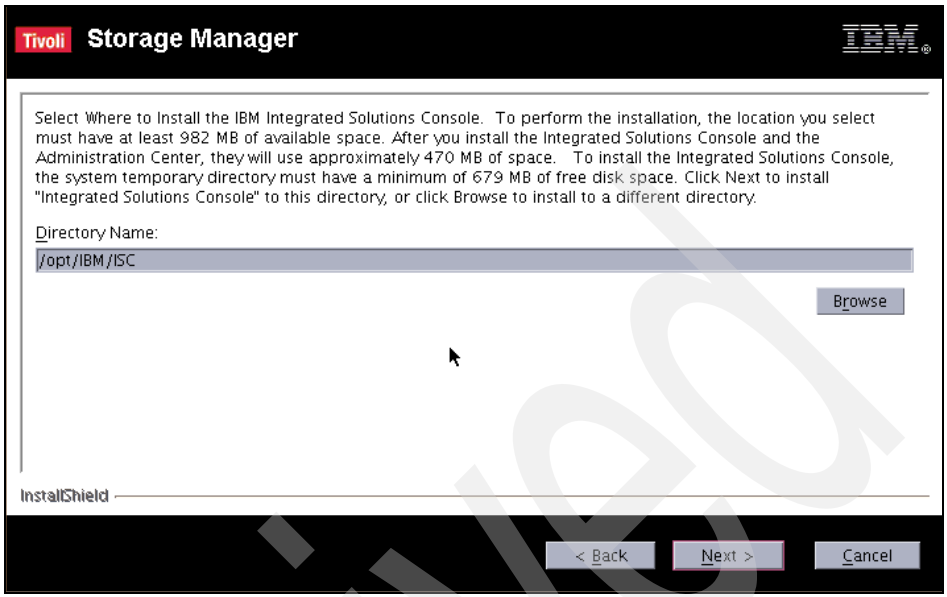

Figure 9-16 ISC installation screen, target path - our shared disk for this node

5. We accept the default name (iscadmin) for the ISC user ID, choose and type type in password and verify password and click **Next** on **Create a User ID and Password** panel (Figure 9-17).

Archived

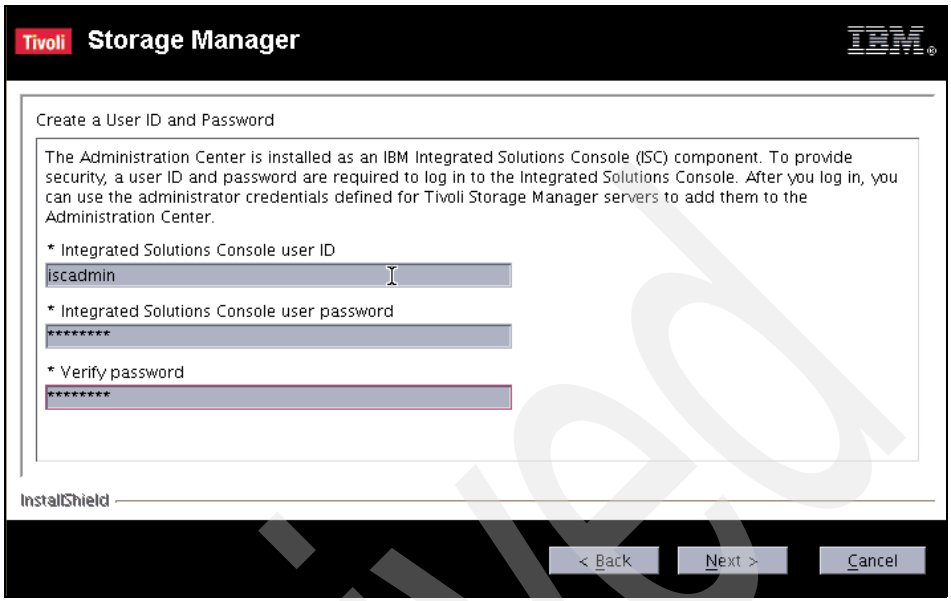

Figure 9-17 ISC installation screen, establishing a login and password

6. We accept the default port numbers for http and https and click **Next** on the **Select the Ports the IBM ISC Can use** panel (Figure 9-18).

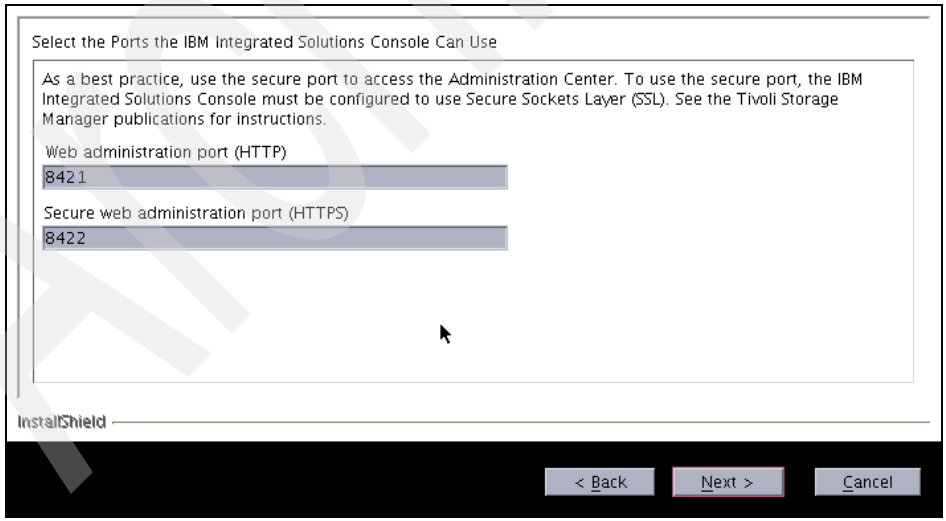

Figure 9-18 ISC installation screen establishing the ports which will be used

7. We verify entered options and click **Next** on **Review** panel (Figure 9-19).

Archived

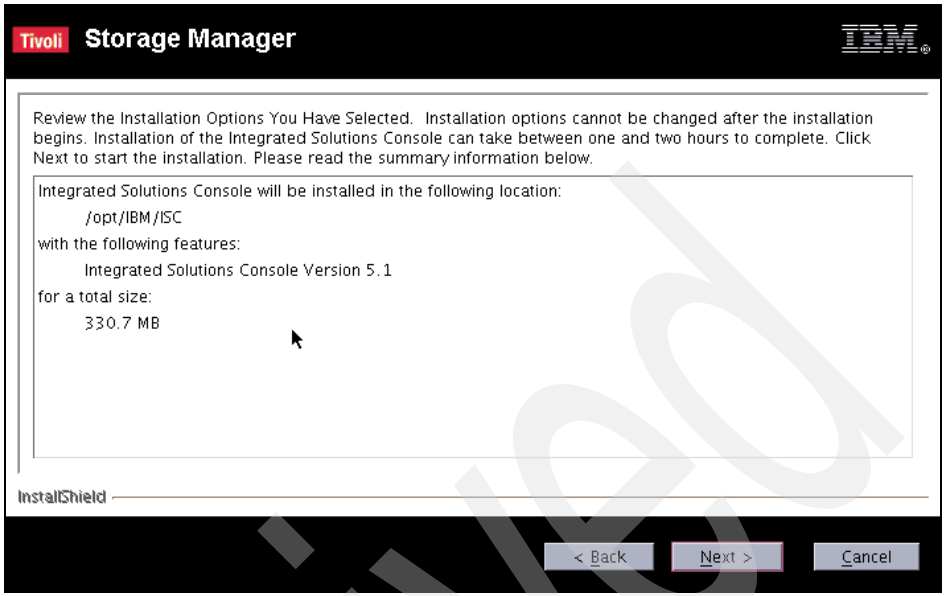

Figure 9-19 ISC installation screen, reviewing selections and disk space required

8. Then we wait for the completion panel and click **Next** on it (Figure 9-20).

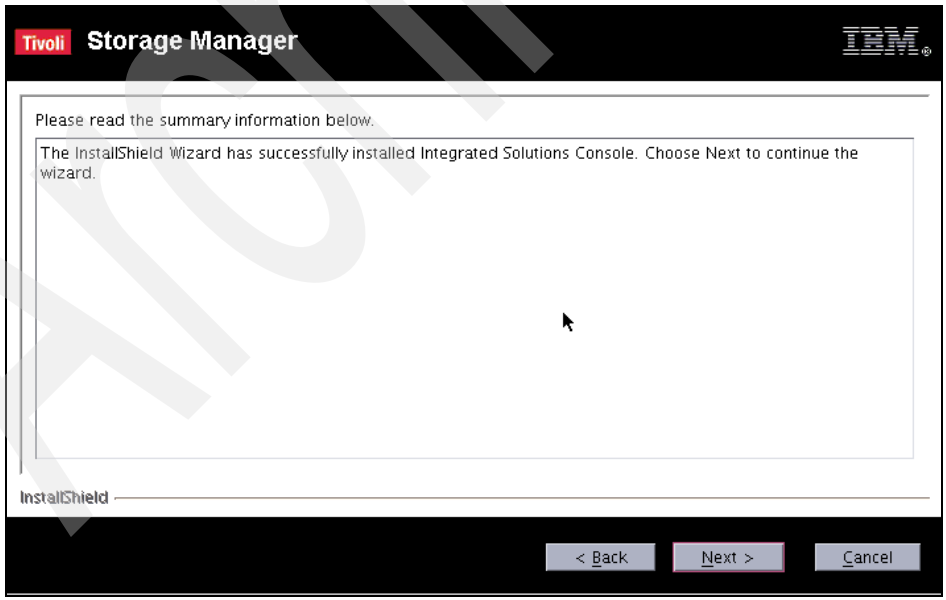

Figure 9-20 ISC installation screen showing completion

Archived

9. Now we make note of the ISC address onthe **Installation Summary** panel and click **Next** on it (Figure 9-21).

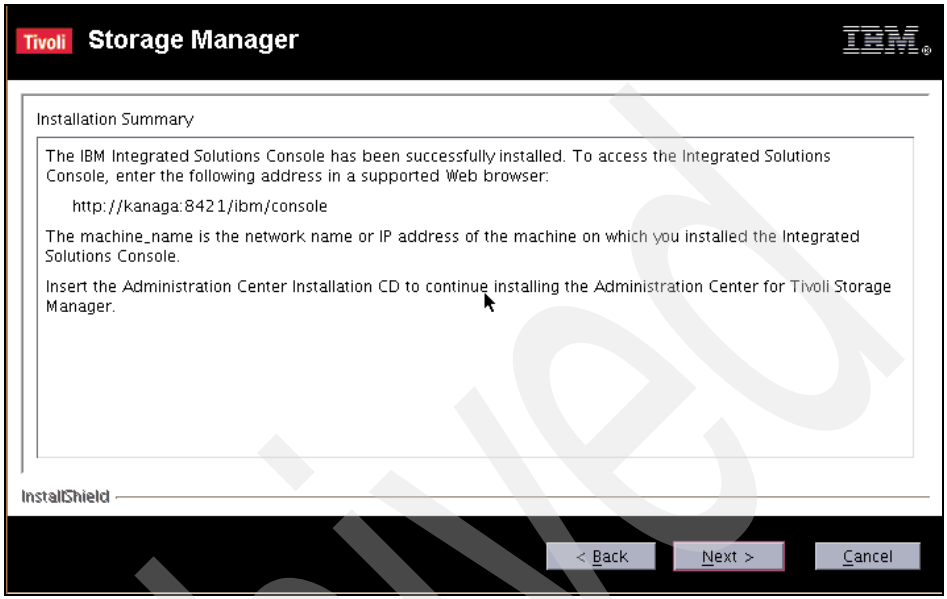

Figure 9-21 ISC installation screen, final summary providing URL for connection

9.3.7 Installing the Tivoli Storage Manager Administration Center

Here we install the Tivoli Storage Manager Administration center.

1. First we extract the contents of the file TSMAdminCenter5300_AIX.tar (Example 9-3).

Example 9-3 The tar command extraction

```
tar xvf TSMAdminCenter5300_AIX.tar
```

2. Then we change directory into acinstall and run the startInstall.sh InstallShield command script (Example 9-4).

Example 9-4 startInstall.sh usage

```
startInstall.sh
```

Archived

Note: Depending on what are screen and graphics requirements would be, the following options exist for this installation.

Run one of the following commands to install the Administration Center:

- For Installshield wizard install, run: startInstall.sh
- For console wizard install, run: startInstall.sh -console
- For silent install, run the following command on a single line:
  `startInstall.sh -silent -W AdminNamePanel.adminName="<user name>"`
- Flags:
  - W PasswordInput.adminPass="<user password>"
  - W PasswordInput.verifyPass="<user password>"
  - W MediaLocationInput.installMediaLocation="<media location>"
  - W PortInput.webAdminPort="<web administration port>"
  - P AdminCenterDeploy.installLocation="<install location>"

Note: The installation process can take anywhere from 30 minutes to 2 hours to complete. The time to install depends on the speed of your processor and memory.

3. We choose to use the console install method for Administration Center, so we launch **`startInstall.sh -console`**. Example 9-5 shows how we did this.

Example 9-5 Command line installation for the Administration Center

```
azov:/# cd /install/acinstall
azov:/install/acinstall# ./startInstall.sh -console

InstallShield Wizard

Initializing InstallShield Wizard...

Preparing Java(tm) Virtual Machine...
....................................
....................................
....................................
....................................
....................................
....................................
....................................
....................................
....................................
....................................
....................................
....................................
....................................
```

Archived

```
.................................
.................................
.................................
.................................
.................................
.................................
.................................
.................................
.................................
.................................
........
Welcome to the InstallShield Wizard for Administration Center

The InstallShield Wizard will install Administration Center on your computer.
To continue, choose Next.

IBM Tivoli Storage Manager
Administration Center
Version 5.3


Press 1 for Next, 3 to Cancel or 4 to Redisplay [1]
Welcome

The Administration Center is a Web-based interface that can be used to
centrally configure and manage IBM Tivoli Storage Manager Version 5.3 servers.
The Administration Center is installed as an IBM Integrated Solutions Console
component. The Integrated Solutions Console allows you to create custom
solutions by installing components provided by one or more IBM applications.

Version 5.1 of the Integrated Solutions Console is required to use the
Administration Center. If an earlier version of the Integrated Solutions
Console is already installed, use the Integrated Solutions Console CD in this
package to upgrade to version 5.1

For the latest product information, see the readme file on the installation CD
or the Tivoli Storage Manager technical support website


(http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.h
tml).

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1] 1

Review License Information. Select whether to accept the license terms for this
product. By accepting the terms of this license, you acknowledge that you have
thoroughly read and understand the license information.

International Program License Agreement
```

```
Part 1 - General Terms

BY DOWNLOADING, INSTALLING, COPYING, ACCESSING, OR USING THE PROGRAM YOU AGREE
TO THE TERMS OF THIS AGREEMENT. IF YOU ARE ACCEPTING THESE TERMS ON BEHALF OF
ANOTHER PERSON OR A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT AND WARRANT
THAT YOU HAVE FULL AUTHORITY TO BIND THAT PERSON, COMPANY, OR LEGAL ENTITY TO
THESE TERMS. IF YOU DO NOT AGREE TO THESE TERMS,

- DO NOT DOWNLOAD, INSTALL, COPY, ACCESS, OR USE THE PROGRAM; AND

- PROMPTLY RETURN THE PROGRAM AND PROOF OF ENTITLEMENT TO THE PARTY FROM WHOM
YOU ACQUIRED IT TO OBTAIN A REFUND OF THE AMOUNT YOU PAID. IF YOU DOWNLOADED
THE PROGRAM, CONTACT THE PARTY FROM WHOM YOU ACQUIRED IT.

"IBM" is International Business Machines Corporation or one of its
subsidiaries.

"License Information" ("LI") is a document that provides information specific

Press ENTER to read the text [Type q to quit] q


Please choose from the following options:

[ ] 1 - I accept the terms of the license agreement.
[X] 2 - I do not accept the terms of the license agreement.

To select an item enter its number, or 0 when you are finished: [0]1

Enter 0 to continue or 1 to make another selection: [0]


Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]

Review Integrated Solutions Console Configuration Information

To deploy the Administration Center component to the IBM Integrated Solutions
Console, the information listed here for the Integrated Solutions Console must
be correct. Verify the following information.

IBM Integrated Solutions Console installation path:

/opt/IBM/ISC

IBM Integrated Solutions Console Web Administration Port:

8421
```

Archived

```
IBM Integrated Solutions Console user ID:

iscadmin

[X] 1 - The information is correct.
[ ] 2 - I would like to update the information.

To select an item enter its number, or 0 when you are finished: [0]

To select an item enter its number, or 0 when you are finished: [0]


Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]

Enter the Integrated Solutions Console Password

Enter the password for user ID iscadmin

* Integrated Solutions Console user password
Please press Enter to Continue
Password: scadmin


* Verify password
Please press Enter to Continue
Password: scadmin


Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]

Select the Location of the Installation CD



Location of the installation CD [/install/acinstall]


Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]

Administration Center will be installed in the following location:

/opt/IBM/ISC

with the following features:

Administration Center Deployment

for a total size:
```

Archived

```
 305 MB

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]


Installing Administration Center. Please wait...


Installing Administration Center. Please wait... - Extracting...


Installing Administration Center. Please wait...


Installing the Administration Center
Install Log location /opt/IBM/ISC/Tivoli/dsm/logs/ac_install.log



Creating uninstaller...

The InstallShield Wizard has successfully installed Administration Center.
Choose Next to continue the wizard.

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1] 1

Installation Summary

The Administration Center has been successfully installed. To access the
Administration Center, enter the following address in a supported Web browser:

http://azov.almaden.ibm.com:8421/ibm/console

The machine_name is the network name or IP address of the machine on which you
installed the Administration Center

To get started, log in using the Integrated Solutions Console user ID and
password you specified during the installation. When you successfully log in,
the Integrated Solutions Console welcome page is displayed. Expand the Tivoli
Storage Manager folder in the Work Items list and click Getting Started to
display the Tivoli Storage Manager welcome page. This page provides
instructions for using the Administration Center.

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]

The wizard requires that you logout and log back in.

Press 3 to Finish or 4 to Redisplay [3]
```

Archived

4. Then we can access the Administration Center via http://azov.almaden.ibm.com:8421/ibm/console

9.3.8 Configure resources and resource groups

Resource groups are collections of resources which are managed as group during cluster operations.

In this section we are showing how we configure resources prepared in Chapter 8, “Establishing an HACMP infrastructure on AIX” on page 417 and the resource group to be used with **Tivoli Storage Manager server**.

Then we will use the same procedure configuring the **ISC** and **Admin Center** resources and resource group; only the names and network/storage objects will change.

Configure service addresses

Network addresses that are included in the */etc/hosts* file prior to the *HACMP resource discovery* run (see “Resource discovery” on page 445) can be picked from the list when configuring service addresses, as we are doing here:

1. We enter `smitty hacmp` on the AIX command line.
2. Then, we select **Extended Configuration**.
3. Next, we select the **Extended Resource Configuration** option.
4. Then, we choose the **HACMP Extended Resources Configuration** option.
5. We then select the **Configure HACMP Service IP Labels/Addresses** panel.
6. We choose the **Add a Service IP Label/Address** option.
7. Then, we select **Configurable on Multiple Nodes**.
8. We then choose the applicable network.
9. Choose the IP Label/Address to be used with Tivoli Storage Manager server.
10. We then press Enter to complete the processing (Figure 9-22).

Archived

```
    Add a Service IP Label/Address configurable on Multiple Nodes (extended)

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* IP Label/Address                                     tsmsrv03                +
* Network Name                                         net_ether_01
  Alternate HW Address to accompany IP Label/Address []



F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 9-22 Service address configuration

Create resource groups

Creating a resource group for managing Tivoli Storage Manager server:

1. We go back to **Extended Resource Configuration**.
2. Then we select **HACMP Extended Resource Group Configuration**.
3. And then we **Add a Resource Group**.
4. We type in resource group name, *rg_tsmsrv03*.
5. We pick from the list the participating nodes name.
6. We check that the nodes name order matches the **nodes priority order** for cascading resource groups; we write, don' t pick, if the order differs.

   Note: Nodes priority is determined by the order in which the node names appear.

7. We select **Startup/Fallover/Fallback** policies, and we choose:
   a. Online On Home Node only
   b. Fallover To Next Priority Node
   c. Never Fallback

   See “Planning and design” on page 422 and “Plan for cascading versus rotating” on page 426. Using F1 on the parameter line gives exhaustive help or you can refer to “Resource Groups and Their Behavior During Startup, Fallover, and Fallback” in the HACMP 5.2 *Planning and Installation Guide*).

Archived

8. We press Enter (Figure 9-23).

```
                    Add a Resource Group (extended)

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* Resource Group Name                                  [rg_tsmsrv03]
* Participating Nodes (Default Node Priority)          [azov kanaga]           +

  Startup Policy                                        Online On Home Node O> +
  Fallover Policy                                       Fallover To Next Prio> +
  Fallback Policy                                       Fallback To Higher Pr> +


F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 9-23 Add a resource group

Add resources to the resource group

Adding resources to the Tivoli Storage Manager server resource group:

1. We go back to **HACMP Extended Resource Group Configuration**.
2. Then we select **Change/Show Resources and Attributes for a Resource Group**.
3. And then we select our resource group name.
4. The **Change/Show Resources and Attributes for a Resource Group** shows, and we pick from the list **Service IP Labels** and **Volume Groups**.
5. We leave empty the Filesystems field, that means all filesystems in selected **VGs** are to be managed.
6. We check node priority and policies.

Archived

7. We press Enter (Figure 9-24).

```
          Change/Show All Resources and Attributes for a Resource Group

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[TOP]                                                   [Entry Fields]
  Resource Group Name                                  rg_tsmsrv03
  Participating Nodes (Default Node Priority)          azov kanaga
  [                                                    _
  Startup Policy                                       Online On Home Node O>
  Fallover Policy                                      Fallover To Next Prio>
  Fallback Policy                                      Never Fallback

  Service IP Labels/Addresses                         [tsmsrv03]              +
  Application Servers                                 [as_tsmsrv03]           +

  Volume Groups                                       [tsmvg]                 +
  Use forced varyon of volume groups, if necessary     false                  +
  Automatically Import Volume Groups                   false                  +

  Filesystems (empty is ALL for VGs specified)        []                      +
  Filesystems Consistency Check                        fsck                   +
  Filesystems Recovery Method                          sequential             +
  Filesystems mounted before IP configured             false                  +
  Filesystems/Directories to Export                   []                      +
                                                                              +
  Filesystems/Directories to NFS Mount                []
  Network For NFS Mount                               []                      +

  Tape Resources                                      [ltolib1 ltolib1dr1 lt> +
  Raw Disk PVIDs                                      []                      +
[MORE...8]

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 9-24 Add resources to the resource group

9.3.9 Synchronize cluster configuration and make resource available

Here we are synchronizing and starting up the cluster resources.

Before synchronizing the cluster configuration, we should verify that the clcomd daemon is added to /etc/inittab and started by init on all nodes in the cluster.

Synchronize cluster configuration

A copy of the cluster configuration is stored on each node; now we are going to synchronize them.

> **Note:** Remember to do that from the node you where you are inserting cluster data.

1. We use **`smitty hacmp fast path`**.

2. We select the **Extended Configuration** menu.
3. Then we select **Extended Verification and Synchronization**.
4. We leave the defaults and press Enter.
5. We look at the result and take appropriate action for errors and warnings if needed (we ignore warnings about netmon.cf missing for point-to-point networks) (Figure 9-25).

```
                                COMMAND STATUS

Command: OK            stdout: yes           stderr: no

Before command completion, additional instructions may appear below.

[TOP]
_
Verification to be performed on the following:
        Cluster Topology
        Cluster Resources

Retrieving data from available cluster nodes.  This could take a few minutes....
..

Verifying Cluster Topology...

WARNING: File 'netmon.cf' is missing or empty on node azov. This file is needed
for a cluster with the single-adapter network net_rs232_01. Please create 'netmo
[MORE...80]

F1=Help             F2=Refresh          F3=Cancel           F6=Command
F8=Image            F9=Shell            F10=Exit            /=Find
n=Find Next
```

Figure 9-25 Cluster resources synchronization

Reconfigure default gateway

Once the first synchronization has run, persistent addresses are available, so these addresses are available for network connections, and a default gateway configuration, which has been deleted configuring boot addresses, can be restored now:

1. We use **`smitty route`** fast path.
2. We select **Add a Static Route**.
3. And we fill in as required and press Enter.

Start cluster services to make resource available

Now we make available cluster resources needed for Tivoli Storage Manager server configuration.

Start and stop scripts for the Tivoli Storage Manager sever will be customized and added to cluster resource later on.

Archived

We can start the cluster services by using the SMIT fast path smitty clstart. From there, we can select the nodes on which we want cluster services to start. We choose to **don't** start the **cluster lock services** (not needed in our configuration) and to **start** the **cluster information daemon**.

1. First, we issue the **`smitty clstart`** fast path command.
2. Next, we configure as shown in Figure 9-26 (using F1 on parameter lines gives exhaustive help).
3. To complete the process, press Enter.

```
                          Start Cluster Services

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                      [Entry Fields]
* Start now, on system restart or both                 now                 +
  Start Cluster Services on these nodes                [azov]              +
  BROADCAST message at startup?                         false              +
  Startup Cluster Information Daemon?                   true               +
  Reacquire resources after forced down ?               false              +

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 9-26 Starting cluster services.

4. Monitor the status of the cluster services using the command **`lssrc -g cluster`** (Example 9-6).

Example 9-6 lssrc -g cluster

```
azov:/# lssrc -g cluster
Subsystem         Group            PID          Status
 clstrmgrES       cluster          213458       active
 clsmuxpdES       cluster          233940       active
 clinfoES         cluster          238040       active
```

Note: After having the cluster services started, resources are being taken online. You can view the /tmp/hacmp.log log file for operations progress monitor (tail -f /tmp/hacmp.out).

Archived

5. Overall cluster status monitor is available thought /usr/es/sbin/cluster/clstat.
   - It comes up with an X11 interface if a graphical environment is available (Figure 9-27).

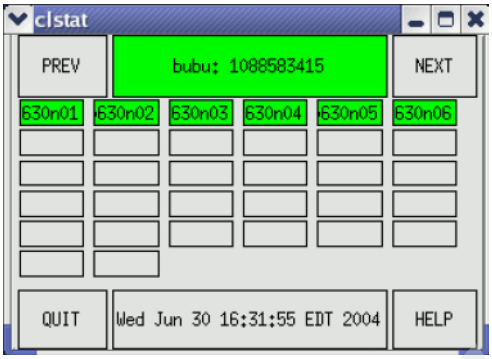

Figure 9-27 X11 clstat example

   - Otherwise a character based interface is shown as in Figure 9-28, where we can monitor state in our cluster for:
     - Cluster
     - Nodes
     - Interfaces
     - Resource groups

```
                clstat - HACMP Cluster Status Monitor
                -------------------------------------

Cluster: cl_hacmp01     (1107294149)
Fri Jan 28 10:05:18 PST 2005
                State: UP               Nodes: 2
                SubState: STABLE

        Node: azov              State: UP
           Interface: azovb1 (1)                Address: 10.1.1.89
                                                State:   UP
           Interface: azovb2 (1)                Address: 10.1.2.89
                                                State:   UP
           Interface: azov_hdisk3_01 (0)        Address: 0.0.0.0
                                                State:   UP
           Interface: azov_tty0_01 (2)          Address: 0.0.0.0
                                                State:   UP
           Interface: tsmsrv03 (1)              Address: 9.1.39.74
                                                State:   UP
           Resource Group: rg_tsmsrv03                  State:  On line

************************ f/forward, b/back, r/refresh, q/quit *****************
```

Figure 9-28 clstat output

   - Starting with HACMP 5.2, you can use the WebSMIT version of **`clstat`** (**`wsm_clstat.cgi`**) (Figure 9-29).

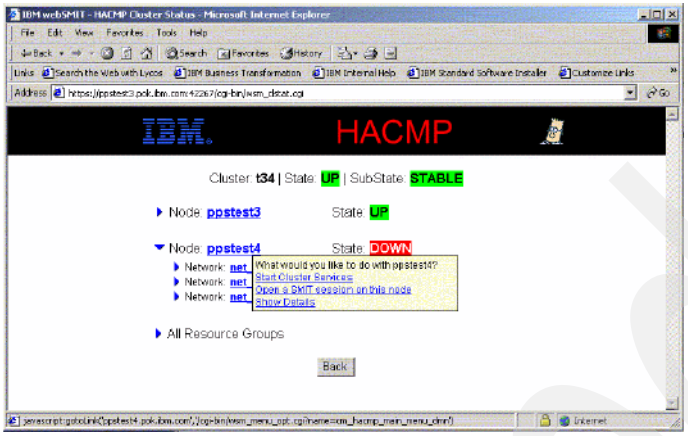

Figure 9-29 WebSMIT version of clstat example

See “Monitoring Clusters with clstat” on the HACMP *Administration and Troubleshooting Guide* for more details about clstat and the WebSMIT version of clstat setup.

6. Finally, we check for resources with operating system commands (Figure 9-30).

```
azov/: df
Filesystem    512-blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4           65536     29776   55%     1977    36% /
/dev/hd2         3997696    219536   95%    32587    54% /usr
/dev/hd9var       131072     58056   56%      552     8% /var
/dev/hd3         2621440   2595248    1%      213     1% /tmp
/dev/hd1           65536     64832    2%        5     1% /home
/proc                  -         -    -        -     -  /proc
/dev/hd10opt     2424832   2231952    8%     2196     1% /opt
/dev/fslv00      2097152   1185688   44%     1501     2% /install
/dev/tsmdb1lv    4128768     29432  100%        5     1% /tsm/db1
/dev/tsmdbmr1lv    4128768     29432  100%        5     1% /tsm/dbmr1
/dev/tsmdp1lv   51773440    564880   99%        5     1% /tsm/dp1
/dev/tsmlv        131072    130224    1%        9     1% /tsm/files
/dev/tsmlg1lv    2031616     78904   97%        5     1% /tsm/lg1
/dev/tsmlgmr1lv    2031616     78904   97%        5     1% /tsm/lgmr1
azov/: netstat -i
Name  Mtu   Network     Address             Ipkts Ierrs    Opkts Oerrs  Coll
en0   1500  link#2      0.2.55.4f.46.b2     1316131     0   740498     0     0
en0   1500  10.1.1      azovb1              1316131     0   740498     0     0
en0   1500  9.1.39      azov                1316131     0   740498     0     0
en1   1500  link#3      0.6.29.6b.83.e4    14087227     0  7147709     3     0
en1   1500  10.1.2      azovb2             14087227     0  7147709     3     0
en1   1500  9.1.39      tsmsrv03           14087227     0  7147709     3     0
lo0   16896 link#1                          1101844     0  1100371     0     0
lo0   16896 127         loopback            1101844     0  1100371     0     0
lo0   16896 ::1                             1101844     0  1100371     0     0
azov/: _
```

Figure 9-30 Check for available resources

Core testing

At this point, we recommend testing at least the main cluster operation, and we do so. Basic tasks such as putting resources online and offline, or moving them across the cluster node, to verify basic cluster operation and set a check point, are shown in “Core HACMP cluster testing” on page 496.

9.4 Tivoli Storage Manager Server configuration

Now that the needed storage and network resources are available, it is possible to configure the Tivoli Storage Manager server and set up start and stop scripts to be used by the HACMP cluster.

Default installation cleanup

Since we are going to create a new instance on shared disks, we can clean up the installation-created one.

These steps are to be executed on both nodes:

1. We remove the entry from /etc/inittab that starts the IBM Tivoli Storage Manager server, using the **rmitab autosrvr** command.
2. We stop the default server installation instance, if running (Example 9-7).

Example 9-7 Stop the initial server installation instance

```
# ps -ef|grep dsmserv
root  41304 176212 0 09:52:48  pts/3  0:00 grep dsmserv
root 229768      1 0 07:39:36 - 0:56 /usr/tivoli/tsm/server/bin/dsmserv quiet
# kill 229768
```

3. We clean up the default server installation files which are not required, we remove the default created database, recovery log, space management, archive, and backup files created. We also remove the dsmserv.dsk and the dsmserv.opt files (Example 9-8).

Example 9-8 Files to remove after the initial server installation

```
# cd /usr/tivoli/tsm/server/bin
# rm dsmserv.opt
# rm dsmserv.dsk
# rm db.dsm
# rm spcmgmt.dsm
# rm log.dsm
# rm backup.dsm
# rm archive.dsm
# rm archive.dsm
```

Server instance installation and mirroring

Here we create the shared disk installed instance and execute the main customization tasks; further customization can be done as with any other installation.

1. We configure IBM Tivoli Storage Manager to use the TCP/IP communication method. See the HACMP *Installation Guide* for more information on specifying server and client communications. TCP/IP is the default in dsmserv.opt.smp. Copy dsmserv.opt.smp to /tsm/files/dsmserv.opt.
2. Then we configure the local client to communicate with the server (only basic communication parameters in dsm.sys found in the /usr/tivoli/tsm/client/ba/bin directory) (Example 9-9). We will use this server stanza for the *Command Line Administrative Interface communication.*

Example 9-9 The server stanza for the client dsm.sys file

```
* Server stanza for admin connection purpose
SErvername tsmsrv03_admin
COMMMethod TCPip
TCPPor 1500
TCPServeraddress 127.0.0.1
ERRORLOGRETENTION 7
ERRORLOGname /usr/tivoli/tsm/client/ba/bin/dsmerror.log
```

Note: We used loopback address because we want to be sure that the stop script that we are going to set up later on, connects only when server is local.

3. We set up the appropriate IBM Tivoli Storage Manager server directory environment setting for the current shell issuing the following commands (Example 9-10).

Example 9-10 The variables which must be exported in our environment

```
# export DSMSERV_CONFIG=/tsm/files/dsmserv.opt
# export DSMSERV_DIR=/usr/tivoli/tsm/server/bin
```

Tip: For information about running the server from a directory different from the default database that was created during the server installation, also see the *Installation Guide*.

4. Then we allocate the IBM Tivoli Storage Manager database, recovery log, and storage pools on the shared IBM Tivoli Storage Manager volume group. To accomplish this, we will use the **`dsmfmt command`** to format database, log and disk storage pools files on the shared filesystems (Example 9-11).

Example 9-11 dsmfmt command to create database, recovery log, storage pool files

```
# cd /tsm/files
# dsmfmt -m -db /tsm/db1/vol1 2000
# dsmfmt -m -db /tsm/dbmr1/vol1 2000
# dsmfmt -m -log /tsm/lg1/vol1 1000
# dsmfmt -m -log /tsm/lgmr1/vol1 1000
# dsmfmt -m -data /tsm/dp1/bckvol1 25000
```

5. We change the current directory to the new server directory and we then issue the **`dsmserv format command`** to initialize the database and recovery log and create the **dsmserv.dsk file**, which point to the database and log files (Example 9-12).

Example 9-12 The dsmserv format prepares db & log files and the dsmserv.dsk file

```
# cd /tsm/files
# dsmserv format 1 /tsm/lg1/vol1 1 /tsm/db1/vol1
```

6. And then we start the Tivoli Storage Manager Server in the foreground by issuing the command dsmserv from the installation directory and with the proper environment variables set within the running shell (Example 9-13).

Example 9-13 Starting the server in the foreground

```
# pwd
/tsm/files
# dsmserv
```

7. Once the Tivoli Storage Manager Server has completed the startup, we run the Tivoli Storage Manager server commands: **`set servername`** to name the new server, **`define dbcopy`** and **`define logcopy`** to mirror database and log, and then we set the **log mode** to **Roll forward** as planned in “Planning for storage and database protection” on page 454 (Example 9-14).

Example 9-14 Our server naming and mirroring.

```
TSM:SERVER03> set servername tsmsrv03
TSM:TSMSRV03> define dbcopy /tsm/db1/vol1 /tsm/dbmr1/vol1
TSM:TSMSRV03> define logcopy /tsm/lg1/vol1 /tsm/lgmr1/vol1
TSM:TSMSRV03> set logmode rollforward
```

Former customization

1. We then define a DISK storage pool with a volume on the shared filesystem /tsm/dp1 which is configured on a RAID1 protected storage device (Example 9-15).

Example 9-15 The define commands for the diskpool

```
TSM:TSMSRV03> define stgpool spd_bck disk
TSM:TSMSRV03> define volume spd_bck /tsm/dp1/bckvol1
```

2. We now define the tape library and tape drive configurations using the define library, define drive and define path commands (Example 9-16).

Example 9-16 An example of define library, define drive and define path commands

```
TSM:TSMSRV03> define library liblto libtype=scsi
TSM:TSMSRV03> define path tsmsrv03 liblto srctype=server desttype=libr
device=/dev/smc0
TSM:TSMSRV03> define drive liblto drlto_1
TSM:TSMSRV03> define drive liblto drlto_2
TSM:TSMSRV03> define path tsmsrv03 drlto_1 srctype=server desttype=drive
libr=liblto device=/dev/rmt0
TSM:TSMSRV03> define path tsmsrv03 drlto_2 srctype=server desttype=drive
libr=liblto device=/dev/rmt1
```

3. We set library parameter **resetdrives=yes**, this enables a new Tivoli Storage Manager 5.3 server for AIX function that resets SCSI reserved tape drives on server or Storage Agent restart. If we use a older version we still need a SCSI reset from HACMP tape resources management and/or older TSM server startup samples scripts (Example 9-17).

Note: In a library client/server or LAN-free environment, this is function is available only if a Tivoli Storage Manager for AIX server, 5.3 or later, acts as library server.

Example 9-17 Library parameter RESETDRIVES set to YES

```
TSM:TSMSRV03> update library liblto RESETDRIVES=YES
```

4. We will now register the **admin** administrator with the **system** authority with the register admin and grant authority commands to enable further server customization and server administration, though the ISC and command line (Example 9-18).

Example 9-18 The register admin and grant authority commands

```
TSM:TSMSRV03> register admin admin admin
   TSM:TSMSRV03> grant authority admin classes=system
```

5. Now we register a **script_operator** administrator with the operator authority with the register admin and grant authority commands to be used in the server stop script (Example 9-19).

Archived

Example 9-19 The register admin and grant authority commands

```
TSM:TSMSRV03> register admin script_operator password
   TSM:TSMSRV03> grant authority script_operator classes=operator
```

Start and stop scripts setup

Here we set up application start and stop scripts to be configured as application server objects in HACMP.

Tivoli Storage Manager server

We chose to use the standard HACMP application scripts directory for start and stop scripts.

1. At first we create the **/usr/es/sbin/cluster/local/tsmsrv** directory on both nodes.
2. Then from **/usr/tivoli/tsm/server/bin/** we copy the two sample scripts to our scripts directory on the first node (Example 9-20).

Example 9-20 Copy the example scripts on the first node

```
cd /usr/tivoli/tsm/server/bin/
cp startserver /usr/es/sbin/cluster/local/tsmsrv/starttsmsrv03.sh
cp stopserver /usr/es/sbin/cluster/local/tsmsrv/stoptsmsrv03.sh
```

3. Now we adapt the start script to our environment, setting the correct running directory for **dsmserv** and other operating system related environment variables, crosschecking them with the latest /usr/tivoli/tsm/server/bin/rc.adsmserv file (Example 9-21).

Example 9-21 Setting running environment in the start script

```
#!/bin/ksh
##############################################################################
#                                                                            #
# Shell script to start a TSM server.                                        #
#                                                                            #
# Please note commentary below indicating the places where this shell script  #
# may need to be modified in order to tailor it for your environment.        #
#                                                                            #
##############################################################################
#                                                                            #
# Update the cd command below to change to the directory that contains the   #
# dsmserv.dsk file and change the export commands to point to the dsmserv.opt #
# file and /usr/tivoli/tsm/server/bin directory for the TSM server being     #
# started.  The export commands are currently set to the defaults.           #
#                                                                            #
##############################################################################
echo “Starting TSM now...”
```

```
cd /tsm/files

export DSMSERV_CONFIG=/tsm/files/dsmserv.opt
export DSMSERV_DIR=/usr/tivoli/tsm/server/bin

# Allow the server to pack shared memory segments
export EXTSHM=ON

# max out size of data area
ulimit -d unlimited

# Make sure we run in the correct threading environment
export AIXTHREAD_MNRATIO=1:1
export AIXTHREAD_SCOPE=S

##############################################################################
#                                                                            #
# set the server language.  These two statements need to be modified by the  #
# user to set the appropriate language.                                      #
#                                                                            #
##############################################################################
export LC_ALL=en_US
export LANG=en_US

#OK, now fire-up the server in quiet mode.
$DSMSERV_DIR/dsmserv quiet &
```

4. Then we modify the stop script following header inserted instructions (Example 9-22).

Example 9-22 Stop script setup instructions

```
[...]
# Please note that changes must be made to the dsmadmc command below in   #
# order to tailor it for your environment:                                #
#                                                                         #
#    1. Set -servername= to the TSM server name on the SErvername option  #
#       in the /usr/tivoli/tsm/client/ba/bin/dsm.sys file.                #
#                                                                         #
#    2. Set -id= and -password= to a TSM userid that has been granted     #
#       operator authority, as described in the section:                  #
#       “Chapter 3. Customizing Your Tivoli Storage Manager System -      #
#       Adding Administrators”, in the Quick Start manual.                #
#                                                                         #
#    3. Edit the path in the LOCKFILE= statement to the directory where   #
#       your dsmserv.dsk file exists for this server.                     #
[...]
```

5. We modify the lock file path (Example 9-23).

Example 9-23 Modifying the lock file path

```
[...]
# TSM lock file
LOCKFILE=”/tsm/files/adsmserv.lock”
[...]
```

6. We set server stanza name, user id, and password (Example 9-24).

Example 9-24 dsmadmc command setup

```
[...]
/usr/tivoli/tsm/client/ba/bin/dsmadmc -servername=tsmsrv03_admin
-id=script_operator -password=password -noconfirm << EOF
[...]
```

7. Then now we can test the start and stop scripts and, as this works fine, we copy all directory content to the second cluster node.

Integrated Solution Console

The installation procedure has set an **inittab** entry for starting the ISC at boot time. We copy the command from that line, before removing it with **rmitab** command, and create a script with only that command within it. Example 9-25 shows our startisc.sh script.

Example 9-25 ISC startup command

```
#!/bin/ksh

# Startup the ISC_Portal to make the TSM Admin Center available
/opt/IBM/ISC/PortalServer/bin/startISC.sh ISC_Portal iscadmin iscadmin
```

Then we found, in the product readme files, instructions, and a sample script for stopping the ISC that we are going to use, named stopisc.sh (Example 9-26).

Example 9-26 ISC stop sample script

```
#!/bin/ksh

# Stop The Portal
/opt/IBM/ISC/PortalServer/bin/stopISC.sh ISC_Portal iscadmin iscadmin

# killing all AppServer related java processes left running

JAVAASPIDS=`ps -ef | egrep "java|AppServer" | awk '{ print $2 }'`

for PID in $JAVAASPIDS
```


Archived

```
 do
  kill  $PID
 done

exit 0
```

Application servers configuration and activation

Application server is an HACMP object that identifies start and stop scripts for and application to be made high available.

Here we show how we configure that object for the Tivoli Storage Manager server, then we use the same procedure for the ISC.

1. We use the **`smitty hacmp fast path`**.
2. Then, we select **Extended Configuration**.
3. Then we select **Extended Resource Configuration** option.
4. We then select **HACMP Extended Resources Configuration** option.
5. Then we select **Configure HACMP Applications**.
6. Then we select **Configure HACMP Application Servers**.
7. And then we **Add an Application Server**.
8. We type in Server Name (we type *as_tsmsrv03*), Start Script, Stop Script, and press Enter.
9. Then we go back to **Extended Resource Configuration** and select **HACMP Extended Resource Group Configuration**.
10. We elect **Change/Show Resources and Attributes for a Resource Group** and pick the resource group name to which to add the application server.
11. In the **Application Servers** field, we chose **as_tsmsrv03** from the list.
12. We press Enter and, after the command result, we go back to the **Extended Configuration** panel.
13. Here we select **Extended Verification and Synchronization**, leave the defaults, and press Enter.
14. The cluster verification and synchronization utility runs, and after a successful completion, the application server start script is executed, making the Tivoli Storage Manager server instance running.
15. We repeat the above steps, creating *as_admcnt01* application server with the *startisc.sh* and *stopisc.sh* scripts.

Archived

Application server monitor configuration (optional)

HACMP can monitor specified applications and automatically take action to restart them upon detecting process termination or other application failures. In HACMP 5.2, you can configure multiple application monitors and associate them with one or more application servers.

You can select either of two application monitoring methods: process application monitoring, which detects the termination of one or more processes of an application; or custom application monitoring, which checks the health of an application with a custom monitor method at user-specified polling intervals. Process monitoring is easier to set up, as it uses the built-in monitoring capability provided by RSCT and requires no custom scripts; custom monitoring can monitor more subtle aspects of an application's performance and is more customizable, but it takes more planning, as you must create the custom scripts.

> **Note:** For more detailed information, see:"Configuring HACMP Application Servers" in the HACMP *Administration and Troubleshooting Guide.*

16. We write a monitor script that checks the return code from a query session command issued through the administrative command line interface (**dsmadmc**) as shown in Example 9-27. At least the session for that query has to be found if server is running and accessible, allowing the **dsmadmc** console to exit with RC=0.

Example 9-27 Monitor script example

```
#!/bin/ksh
#######################################################
#
# Module:       monitortsmsrv03.sh
#
# Function:     Simple query to ensure TSM is running and responsive
#
# Author:       Dan Edwards (IBM Canada Ltd.)
#
# Date:         February 09, 2005
#
#######################################################

# Define some variables for use throughout the script
export ID=script_operator          # TSM admin ID
export PASS=password               # TSM admin password

#
#Query tsmsrv looking for a response
#
```

```
/usr/tivoli/tsm/client/ba/bin/dsmadmc -es=tsmsrv03_admin -id=${ID} -pa=${PASS}
“
q session” >/dev/console 2>&1
#
if [ $? -gt 0 ]
        then exit 1
fi
```

17. And then we configure the application custom monitor using the **smitty cm_cfg_custom_appmon** fast path.
18. We select **Add a Custom Application Monitor**.
19. We fill in our choice and press Enter (Figure 9-31).

In this example we choose just to have cluster notification, no restart on failure, and a long monitor interval to avoid having the actlog filled by query messages. We can use any other notification method such as signaling a Tivoli Management product or sending an snmp trap, e-mail, or other notifications of choice.

Note: To have or not to have HACMP restarting the Tivoli Storage Manager server is a highly *solution dependent* choice.

```
                        Add a Custom Application Monitor

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                      [Entry Fields]
* Monitor Name                                       [am_tsmsrv03]
* Application Server(s) to Monitor                    as_tsmsrv03            +
* Monitor Mode                                       [Long-running monitori> +
* Monitor Method                                     <rv/monitortsmsrv03.sh
  Monitor Interval                                   [300]                   #
  Hung Monitor Signal                                []                      #
* Stabilization Interval                              [120]                   #
* Restart Count                                      [0]                     #
  Restart Interval                                   []                      #
* Action on Application Failure                      [notify]               +
  Notify Method                                      []
  Cleanup Method                                     []
  Restart Method                                     []


F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 9-31 The Add a Custom Application Monitor panel

9.5 Testing

Now we can start testing our configuration.

9.5.1 Core HACMP cluster testing

Here we are testing basic cluster functions.This is a checkpoint that can help in problem determination if something goes wrong later on

Here, tests are run with only storage and network resource configured. We suggest running further testing after server code installation and configuration.

We start cluster services, if not already running, via the **`smitty clstart`** fast path

Before every test, we check the status for cluster services, resource groups, and resources on both nodes; In Example 9-28 we are verifying on the primary node.

Example 9-28 Verify available cluster resources

```
azov:/# lssrc -g cluster
Subsystem        Group           PID          Status
 clstrmgrES      cluster         213458       active
 clsmuxpdES      cluster         233940       active
 clinfoES        cluster         238040       active


azov:/# /usr/es/sbin/cluster/utilities/clRGinfo -p
-----------------------------------------------------------------------------
Group Name     Type         State      Location      Priority Override
-----------------------------------------------------------------------------
rg_tsmsrv03    non-concurrent ONLINE     azov
                              OFFLINE    kanaga


azov:/# lsvg -o
tsmvg
rootvg
```

```
azov:/# lsvg -l tsmvg
tsmvg:
LV NAME          TYPE      LPs   PPs   PVs  LV STATE      MOUNT POINT
tsmvglg          jfs2log   1     1     1    open/syncd    N/A
tsmdb1lv         jfs2      63    63    1    open/syncd    /tsm/db1
tsmdbmr1lv       jfs2      63    63    1    open/syncd    /tsm/dbmr1
tsmlg1lv         jfs2      31    31    1    open/syncd    /tsm/lg1
tsmlgmr1lv       jfs2      31    31    1    open/syncd    /tsm/lgmr1
tsmdp1lv         jfs2      790   790   1    open/syncd    /tsm/dp1
tsmlv            jfs2      3     3     1    open/syncd    /tsm/files

azov:/# df
Filesystem    512-blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4          65536     29392   56%     1963    36% /
/dev/hd2        3997696    173024   96%    32673    59% /usr
/dev/hd9var      131072     62984   52%      569     8% /var
/dev/hd3        2621440   2589064    2%      292     1% /tmp
/dev/hd1          65536     64832    2%        5     1% /home
/proc                -         -    -        -     -  /proc
/dev/hd10opt     2424832   2244272    8%     2196     1% /opt
/dev/tsmdb1lv    4128768     29432  100%        5     1% /tsm/db1
/dev/tsmdbmr1lv  4128768     29432  100%        5     1% /tsm/dbmr1
/dev/tsmdp1lv   51773440    564792   99%       11     1% /tsm/dp1
/dev/tsmlv        196608    195848    1%       12     1% /tsm/files
/dev/tsmlg1lv    2031616     78904   97%        5     1% /tsm/lg1
```

```
/dev/tsmlgmr1lv    2031616     78904   97%       5     1% /tsm/lgmr1


azov:/# netstat -i

Name  Mtu   Network     Address           Ipkts Ierrs    Opkts Oerrs  Coll

en0   1500  link#2      0.2.55.4f.46.b2    1149378     0    33173     0     0

en0   1500  10.1.1      azovb1             1149378     0    33173     0     0

en0   1500  9.1.39      azov               1149378     0    33173     0     0

en1   1500  link#3      0.6.29.6b.83.e4      34578     0   531503     3     0

en1   1500  10.1.2      azovb2               34578     0   531503     3     0

en1   1500  9.1.39      tsmsrv03             34578     0   531503     3     0

lo0   16896 link#1                          48941     0    49725     0     0

lo0   16896 127         loopback             48941     0    49725     0     0

lo0   16896 ::1                            48941     0    49725     0     0
```

Manual Fallover (clstop with takeover)

Here we move a resource group from primary to secondary node.

1. To manually takeover the resource group to the secondary node, we enter the `smitty clstop` fast path on the primary node.
2. Then we change **BROADCAST cluster shutdown?** to **false** and **Shutdown mode** to **takeover** (Figure 9-32).

Archived

```
                              Stop Cluster Services

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* Stop now, on system restart or both                   now                    +
  Stop Cluster Services on these nodes                  [azov]                 +
  BROADCAST cluster shutdown?                            false                 +
* Shutdown mode                                          takeover              +

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 9-32 Clstop with takeover

3. We press Enter and wait for the command status result.
4. After the command result shows the cluster services stopping, we can monitor the progress of operation looking at the hacmp.log file using **`tail -f /tmp/hacmp.out`** on the target node (Example 9-29).

Example 9-29 Takeover progress monitor

```
:get_local_nodename[51] [[ azov = kanaga ]]
:get_local_nodename[51] [[ kanaga = kanaga ]]
:get_local_nodename[54] print kanaga
:get_local_nodename[55] exit 0
LOCALNODENAME=kanaga
:cl_hb_alias_network[82] STATUS=0
:cl_hb_alias_network[85] cllsnw -Scn net_rs232_01
:cl_hb_alias_network[85] grep -q hb_over_alias
:cl_hb_alias_network[85] cut -d: -f4
:cl_hb_alias_network[85] exit 0
:network_down_complete[120] exit 0
Feb  2 09:15:02 EVENT COMPLETED: network_down_complete -1 net_rs232_01

                        HACMP Event Summary
Event: network_down_complete -1 net_rs232_01
Start time: Wed Feb  2 09:15:02 2005

End time: Wed Feb  2 09:15:02 2005

Action:         Resource:                       Script Name:
------------------------------------------------------------------------------
```

```
No resources changed as a result of this event
-------------------------------------------------------------------------------
```

5. Once the takeover operation has completed we check the status of resources on both nodes; Example 9-30 shows some check results on the target node.

Example 9-30 Post takeover resource checking

```
kanaga:/# /usr/es/sbin/cluster/utilities/clRGinfo -p
-----------------------------------------------------------------------------
Group Name     Type           State      Location      Priority Override
-----------------------------------------------------------------------------
rg_tsmsrv03    non-concurrent OFFLINE    azov
                              ONLINE     kanaga

kanaga:/# lsvg -o
tsmvg
rootvg

kanaga:/# lsvg -l tsmvg
tsmvg:
LV NAME             TYPE       LPs   PPs   PVs  LV STATE      MOUNT POINT
tsmvglg             jfs2log    1     1     1    open/syncd    N/A
tsmdb1lv            jfs2       63    63    1    open/syncd    /tsm/db1
tsmdbmr1lv          jfs2       63    63    1    open/syncd    /tsm/dbmr1
tsmlg1lv            jfs2       31    31    1    open/syncd    /tsm/lg1
tsmlgmr1lv          jfs2       31    31    1    open/syncd    /tsm/lgmr1
tsmdp1lv            jfs2       790   790   1    open/syncd    /tsm/dp1
tsmlv               jfs2       2     2     1    open/syncd    /tsm/files

kanaga:/# netstat -i
Name  Mtu   Network     Address            Ipkts Ierrs    Opkts Oerrs  Coll
en0   1500  link#2      0.2.55.4f.5c.a1    1056887     0  1231419     0     0
en0   1500  10.1.1      kanagab1           1056887     0  1231419     0     0
en0   1500  9.1.39      admcnt01           1056887     0  1231419     0     0
en0   1500  9.1.39      tsmsrv03           1056887     0  1231419     0     0
en1   1500  link#3      0.6.29.6b.69.91    3256868     0  5771540     5     0
en1   1500  10.1.2      kanagab2           3256868     0  5771540     5     0
en1   1500  9.1.39      kanaga             3256868     0  5771540     5     0
lo0   16896 link#1                          542020     0   536418     0     0
lo0   16896 127         loopback            542020     0   536418     0     0
lo0   16896 ::1                             542020     0   536418     0     0
```

Manual fallback (resource group moving)

We restart cluster services on the primary node and move back the resource group to it.

1. To move the resource group backup to the to the primary node, we at first have to restart cluster services on it via the **`smitty clstart`** fast path.
2. Once the cluster services are started, we check with the **`lssrc -g cluster command`**, we go to the **`smitty hacmp panel`**.
3. Then we select **System Management (C-SPOC)**.
4. Next we select **HACMP Resource Group and Application Management**.
5. Then we select **Move a Resource Group to Another Node**.
6. At **Select a Resource Group**, we select the resource group to be moved.
7. At **Select a Destination Node**, we chose **Restore_Node_Priority_Order**.

> **Important: Restore_Node_Priority_Order** selection has to be used when restoring a resource group to the high priority node, otherwise the **Fallback Policy** will be overridden.

8. We leave the defaults and press Enter.
9. While waiting for the command result, we can monitor the progress of operation looking at the hacmp.log file using **`tail -f /tmp/hacmp.out on`** the target node (Example 9-31).

Example 9-31 Monitor resource group moving

```
rg_tsmsrv03:rg_move_complete[218] [ 0 -ne 0 ]
rg_tsmsrv03:rg_move_complete[227] [ 0 = 1 ]
rg_tsmsrv03:rg_move_complete[251] [ 0 = 1 ]
rg_tsmsrv03:rg_move_complete[307] exit 0
Feb  2 09:36:52 EVENT COMPLETED: rg_move_complete azov 1

                        HACMP Event Summary
Event: rg_move_complete azov 1
Start time: Wed Feb  2 09:36:52 2005

End time: Wed Feb  2 09:36:52 2005

Action:         Resource:                       Script Name:
------------------------------------------------------------------------------
Acquiring resource:     All_servers     start_server
Search on: Wed.Feb.2.09:36:52.PST.2005.start_server.All_servers.rg_tsmsrv03.ref
Resource online:        All_nonerror_servers    start_server
Search on:
Wed.Feb.2.09:36:52.PST.2005.start_server.All_nonerror_servers.rg_tsmsrv03.ref
Resource group online:  rg_tsmsrv03     node_up_local_complete
Search on: Wed.Feb.2.09:36:52.PST.2005.node_up_local_complete.rg_tsmsrv03.ref
------------------------------------------------------------------------------
```

10. Once the move operation has terminated, we check the status of resources on both nodes as before, especially for **Priority Override** (Example 9-32).

Example 9-32 Resource group state check

```
azov:/# /usr/es/sbin/cluster/utilities/clRGinfo -p
-----------------------------------------------------------------------------
Group Name     Type           State      Location     Priority Override
-----------------------------------------------------------------------------
rg_tsmsrv03    non-concurrent ONLINE     azov
                              OFFLINE    kanaga
```

Stop resource group (bring offline)

Here we are checking the cluster ability to put a resource group offline:

1. To put a resource group to the offline state, we go to the **`smitty hacmp panel`**.
2. Then we select **System Management (C-SPOC)**.
3. And then select **HACMP Resource Group and Application Management**.
4. And then we select **Bring a Resource Group Offline**.
5. At **Select a Resource Group**, we select the resource group to be put offline.
6. At **Select an Online Node**, we choose the node where our resource group is online.
7. We leave default **Persist Across Cluster Reboot?** set to **false** and press Enter.
8. While waiting for the command result, we can monitor the progress of the operation looking at the hacmp.log file by using **`tail -f /tmp/hacmp.out`** on the target node (Example 9-33).

Example 9-33 Monitor resource group moving

```
tail -f /tmp/hacmp.out

rg_admcnt01:node_up_remote_complete[204] [ 0 -ne 0 ]
rg_admcnt01:node_up_remote_complete[208] exit 0
Feb  3 11:11:37 EVENT COMPLETED: node_up_remote_complete kanaga

rg_admcnt01:rg_move_complete[206] [ 0 -ne 0 ]
rg_admcnt01:rg_move_complete[212] [ RELEASE = ACQUIRE ]
rg_admcnt01:rg_move_complete[218] [ 0 -ne 0 ]
rg_admcnt01:rg_move_complete[227] [ 0 = 1 ]
rg_admcnt01:rg_move_complete[251] [ 0 = 1 ]
rg_admcnt01:rg_move_complete[307] exit 0
Feb  3 11:11:37 EVENT COMPLETED: rg_move_complete kanaga 2

                         HACMP Event Summary
```

```
Event: rg_move_complete kanaga 2
Start time: Thu Feb  3 11:11:36 2005

End time: Thu Feb  3 11:11:37 2005

Action:         Resource:                       Script Name:
------------------------------------------------------------------------------
Resource group offline: rg_admcnt01     node_up_remote_complete
Search on: Thu.Feb.3.11:11:37.PST.2005.node_up_remote_complete.rg_admcnt01.ref
------------------------------------------------------------------------------
```

9. Once the bring offline operation has terminated, we check the status of resources on both nodes as before, especially for **Priority Override** (Example 9-34).

Example 9-34 Resource group state check

```
kanaga:/# /usr/es/sbin/cluster/utilities/clRGinfo -p
------------------------------------------------------------------------------
Group Name     Type          State      Location      Priority Override
------------------------------------------------------------------------------
rg_admcnt01    non-concurrent OFFLINE    kanaga       OFFLINE
                              OFFLINE    azov         OFFLINE

kanaga:/# lsvg -o
rootvg
kanaga:/# netstat -i
Name  Mtu   Network     Address            Ipkts Ierrs    Opkts Oerrs  Coll
en0   1500  link#2      0.2.55.4f.5c.a1      17759     0    11880     0     0
en0   1500  10.1.1      kanagab1             17759     0    11880     0     0
en1   1500  link#3      0.6.29.6b.69.91      28152     0    21425     5     0
en1   1500  10.1.2      kanagab2             28152     0    21425     5     0
en1   1500  9.1.39      kanaga               28152     0    21425     5     0
lo0   16896 link#1                           17775     0    17810     0     0
lo0   16896 127         loopback             17775     0    17810     0     0
lo0   16896 ::1                              17775     0    17810     0     0
```

Start resource group (bring online)

Here we are checking the cluster ability to bring a resource group online:

1. To put a resource group to the to the offline state, we go to the **`smitty hacmp panel`**.
2. Then we select **System Management (C-SPOC)**.
3. And then select **HACMP Resource Group and Application Management**.
4. And then we select **Bring a Resource Group Online**.
5. At **Select a Resource Group**, we select the resource group to be put online.

6. At **Select a Destination Node**, we choose the node where we want to bring our resource group online.

> **Attention:** Unless our intention is to put the resource group online on a node different from the primary one, we have to select **Restore_Node_Priority_Order** to avoid a resource group Startup/Failback policy override.

7. We leave default **Persist Across Cluster Reboot?** set to **false** and press Enter.
8. While waiting for the command result, we can monitor the progress of operation looking at the hacmp.log file using **`tail -f /tmp/hacmp.out`** on the target node (Example 9-35).

Example 9-35 Monitor resource group moving

```
# Tail -f /tmp/hacmp.out

End time: Thu Feb  3 11:43:48 2005

Action:         Resource:                          Script Name:
------------------------------------------------------------------------------
Acquiring resource:     All_servers     start_server
Search on: Thu.Feb.3.11:43:48.PST.2005.start_server.All_servers.rg_admcnt01.ref
Resource online:        All_nonerror_servers    start_server
Search on:
Thu.Feb.3.11:43:48.PST.2005.start_server.All_nonerror_servers.rg_admcnt01.ref
Resource group online:  rg_admcnt01     node_up_local_complete
Search on: Thu.Feb.3.11:43:48.PST.2005.node_up_local_complete.rg_admcnt01.ref
------------------------------------------------------------------------------
ADMU0116I: Tool information is being logged in file
          /opt/IBM/ISC/AppServer/logs/ISC_Portal/startServer.log
ADMU3100I: Reading configuration for server: ISC_Portal
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server ISC_Portal open for e-business; process id is 454774
+ [[ high = high ]]
+ version=1.2
+ + cl_get_path
HA_DIR=es
+ STATUS=0
+ set +u
+ [ ]
+ exit 0
```

9. Once the bring online operation has terminated, we check the status of resources on both nodes as before, especially for **Priority Override** (Example 9-36).

Example 9-36 Resource group state check

```
kanaga:/# /usr/es/sbin/cluster/utilities/clRGinfo -p
-----------------------------------------------------------------------------
Group Name     Type           State      Location      Priority Override
-----------------------------------------------------------------------------
rg_admcnt01    non-concurrent ONLINE     kanaga
                              OFFLINE    azov

kanaga:/# lsvg -o
iscvg
rootvg
kanaga:/# lsvg -l iscvg
iscvg:
LV NAME             TYPE       LPs   PPs   PVs  LV STATE      MOUNT POINT
iscvglg             jfs2log    1     1     1    open/syncd    N/A
ibmisclv            jfs2       500   500   1    open/syncd    /opt/IBM/ISC
kanaga:/# netstat -i
Name  Mtu   Network     Address            Ipkts Ierrs    Opkts Oerrs  Coll
en0   1500  link#2      0.2.55.4f.5c.a1      20385     0    13678     0     0
en0   1500  10.1.1      kanagab1             20385     0    13678     0     0
en0   1500  9.1.39      admcnt01             20385     0    13678     0     0
en1   1500  link#3      0.6.29.6b.69.91      31094     0    23501     5     0
en1   1500  10.1.2      kanagab2             31094     0    23501     5     0
en1   1500  9.1.39      kanaga               31094     0    23501     5     0
lo0   16896 link#1                           22925     0    22966     0     0
lo0   16896 127         loopback             22925     0    22966     0     0
lo0   16896 ::1                              22925     0    22966     0     0
```

Further testing on infrastructure and resources

So far we have showed the checking for cluster base functionality before Tivoli Storage Manager installation and configuration; other tests we need to do are adapter related tests such as pulling out SAN and Ethernet cables.

The SAN failures are successfully recovered by the storage subsystem device driver once the operating system declares as failed the test involved adapter, accessing the DASDs through the surviving one; a “freeze” in storage access is noted for a few seconds.

Network adapter failures are recovered for the HACMP cluster software moving the involved IP addresses (alias configured) to the other adapter.

Archived

Refer to HACMP and Storage Subsystem documentation for more in depth testing on network and storage resources.

We are going to do further testing once the installation and configuration tasks are complete.

9.5.2 Failure during Tivoli Storage Manager client backup

Our first test with failure and recovery during a client backup is described here.

Objective

In this test we are verifying client operation surviving a server takeover.

Preparation

Here we prepare test environment:

1. We verify that the cluster services are running with the **lssrc -g cluster command** on both nodes.
2. On resource group secondary node we use **tail -f /tmp/hacmp.out** to monitor cluster operation.
3. Then we start a client incremental backup with the command line and look for metadata and data sessions starting on the server (Example 9-37).

Example 9-37 Client sessions starting

```
01/31/05   16:13:57      ANR0406I Session 19 started for node CL_HACMP03_CLIENT
(AIX) (Tcp/Ip 9.1.39.90(46686)). (SESSION: 19)
01/31/05   16:14:02      ANR0406I Session 20 started for node CL_HACMP03_CLIENT
(AIX) (Tcp/Ip 9.1.39.90(46687)). (SESSION: 20)
```

4. On the server, we verify that data is being transferred via the **query session** command (Example 9-38).

Example 9-38 Query sessions for data transfer

```
tsm: TSMSRV03>q se

  Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- --------------------
    19 Tcp/Ip IdleW    0 S    3.5 M     432 Node  AIX      CL_HACMP03_CLIENT
    20 Tcp/Ip Run      0 S      285  87.6 M Node  AIX      CL_HACMP03_CLIENT
```

Failure

Now we simulate a server crash:

1. Being sure that client backup is running, we issue **halt -q** on the AIX server running the Tivoli Storage Manager server; the **halt -q** command stops any activity immediately and powers off the server.
2. The client stops sending data to the server; it keeps retrying (Example 9-39).

Example 9-39 client stops sending data

```
Normal File-->             6,820
/opt/IBM/ISC/AppServer/config/cells/DefaultNode/applications/Tracing_PA_1_0_3B.
ear/deployments/Tracing_PA_1_0_3B/Tracing.war/WEB-INF/portlet.xml [Sent]
Normal File-->               627
/opt/IBM/ISC/AppServer/config/cells/DefaultNode/applications/Tracing_PA_1_0_3B.
ear/deployments/Tracing_PA_1_0_3B/Tracing.war/WEB-INF/web.xml [Sent]
Directory-->                 256
/opt/IBM/ISC/AppServer/config/cells/DefaultNode/applications/favorites_PA_1_0_3
8.ear/deployments [Sent]
Normal File-->         3,352,904
/opt/IBM/ISC/AppServer/config/cells/DefaultNode/applications/favorites_PA_1_0_3
8.ear/favorites_PA_1_0_38.ear  ** Unsuccessful **

ANS1809W Session is lost; initializing session reopen procedure.
A Reconnection attempt will be made in 00:00:14
[...]
A Reconnection attempt will be made in 00:00:00
A Reconnection attempt will be made in 00:00:14
```

Recovery

Now we see how recovery is managed:

1. The secondary cluster nodes take over the resources and restart the Tivoli Storage Manager server.
2. Once the server is restarted, the client is able to reconnect and continue the incremental backup (Example 9-40 and Example 9-41).

Example 9-40 The restarted Tivoli Storage Manager accept client rejoin

```
01/31/05   16:16:25      ANR2100I Activity log process has started.

01/31/05   16:16:25      ANR4726I The NAS-NDMP support module has been
loaded.

01/31/05   16:16:25      ANR1794W TSM SAN discovery is disabled by options.

01/31/05   16:16:25      ANR2803I License manager started.

01/31/05   16:16:25      ANR8200I TCP/IP driver ready for connection with clients
on port 1500.
```

```
01/31/05   16:16:25      ANR2560I Schedule manager started.

01/31/05   16:16:25      ANR0993I Server initialization complete.

01/31/05   16:16:25      ANR0916I TIVOLI STORAGE MANAGER distributed by
Tivoli is now ready for use.

01/31/05   16:16:25      ANR1305I Disk volume /tsm/dp1/bckvol1 varied online.

01/31/05   16:16:25      ANR0984I Process 1 for AUDIT LICENSE started in the
BACKGROUND.                    BACKGROUND at 16:16:25. (PROCESS: 1)

01/31/05   16:16:25      ANR2820I Automatic license audit started as process 1.
(PROCESS: 1)

01/31/05   16:16:26      ANR2825I License audit process 1 completed
successfully - 3 nodes audited.  (PROCESS: 1)

01/31/05   16:16:26      ANR0987I Process 1 for AUDIT LICENSE running in the
BACKGROUND processed 3 items with a completion state of SUCCESS at
16:16:26. (PROCESS: 1)

01/31/05   16:16:26      ANR0406I Session 1 started for node
CL_HACMP03_CLIENT (AIX) (Tcp/Ip 9.1.39.90(46698)). (SESSION: 1)

01/31/05   16:16:47      ANR0406I Session 2 started for node
CL_HACMP03_CLIENT (AIX) (Tcp/Ip 9.1.39.90(46699)). (SESSION: 2)
```

Example 9-41 The client reconnect and continue operations

```
A Reconnection attempt will be made in 00:00:00 ... successful

Retry # 1  Directory-->            4,096 /opt/IBM/ISC/ [Sent]

Retry # 1  Directory-->            4,096 /opt/IBM/ISC/backups [Sent]

Retry # 1  Normal File-->            482 /opt/IBM/ISC/isc.properties [Sent]

Retry # 1  Normal File-->             68 /opt/IBM/ISC/product.reg [Sent]

Retry # 1  Normal File-->         14,556
/opt/IBM/ISC/AppServer/WEB-INF/portlet.xml [Sent]
```

Archived

Scheduled backup

We repeat the same test using a scheduled backup operation.

Also in this case, the client operation restarts and then completes incremental backup, but instead of a successful operation reports RC=12 even if all files are backed up (Example 9-42).

Example 9-42 Scheduled backup case

```
01/31/05   17:55:42 Normal File-->               207
/opt/IBM/ISC/backups/backups/PortalServer/odc/editors/rt/DocEditor/images/undo_
rtl.gif [Sent]
01/31/05   17:56:34 Normal File-->         2,002,443
/opt/IBM/ISC/backups/backups/PortalServer/odc/editors/ss/SpreadsheetBlox.ear
** Unsuccessful **
01/31/05   17:56:34 ANS1809W Session is lost; initializing session reopen
procedure.
01/31/05   17:57:35 ... successful
01/31/05   17:57:35 Retry # 1  Normal File-->         5,700,745
/opt/IBM/ISC/backups/backups/PortalServer/odc/editors/pr/Presentation.war
[Sent]
01/31/05   17:57:35 Retry # 1  Directory-->               4,096
/opt/IBM/ISC/backups/backups/PortalServer/odc/editors/rt/DocEditor [Sent]
[...]
01/31/05   17:57:56 Successful incremental backup of '/opt/IBM/ISC'


01/31/05   17:57:56 --- SCHEDULEREC STATUS BEGIN

01/31/05   17:57:56 Total number of objects inspected:   37,081

01/31/05   17:57:56 Total number of objects backed up:    5,835

01/31/05   17:57:56 Total number of objects updated:         0

01/31/05   17:57:56 Total number of objects rebound:         0

01/31/05   17:57:56 Total number of objects deleted:         0

01/31/05   17:57:56 Total number of objects expired:         1

01/31/05   17:57:56 Total number of objects failed:          0

01/31/05   17:57:56 Total number of bytes transferred:    371.74 MB

01/31/05   17:57:56 Data transfer time:                 10.55 sec

01/31/05   17:57:56 Network data transfer rate:       36,064.77 KB/sec
```

Archived

01/31/05 17:57:56 Aggregate data transfer rate: 2,321.44 KB/sec

01/31/05 17:57:56 Objects compressed by: 0%

01/31/05 17:57:56 Elapsed processing time: 00:02:43

01/31/05 17:57:56 --- SCHEDULEREC STATUS END

01/31/05 17:57:56 --- SCHEDULEREC OBJECT END TEST_SCHED 01/31/05 17:44:00

01/31/05 17:57:56 ANS1512E Scheduled event 'TEST_SCHED' failed. R.C. = 12.

This results also from the event query (Example 9-43).

Example 9-43 Query event result

```
tsm: TSMSRV03>q ev * TEST_SCHED

Policy Domain Name: STANDARD
     Schedule Name: TEST_SCHED
         Node Name: CL_HACMP03_CLIENT
   Scheduled Start: 01/31/05   17:44:00
      Actual Start: 01/31/05   17:55:16
         Completed: 01/31/05   17:57:56
            Status: Failed
            Result: 12
            Reason: The operation completed with at least one error message
(except for error messages for skipped files).
```

3. We turn back to the primary node of our resource group as described in "Manual fallback (resource group moving)" on page 500.

Result summary

In both cases, the cluster is able to manage server failure and make the Tivoli Storage Manager available to the client in about 1 minute, and the client is able to continue its operations successfully to the end.

With the scheduled operation we get RC=12, but by checking the logs, we are aware of the successful backup completion.

9.5.3 Tivoli Storage Manager server failure during LAN-free restore

Now we test the recovery of a LAN-free operation.

Objective

In this test we are verifying that client LAN-free operation is able to be restarted immediately after a Tivoli Storage Manager server takeover.

Setup

In this test, we use a LAN-free enabled node setup as described in 11.4.3, “Tivoli Storage Manager Storage Agent configuration” on page 562.

1. We register on our server the node with the **register node** command: (Example 9-44).

Example 9-44 Register node command

```
register node atlantic atlantic
```

2. Then we add the related Storage Agent server to our server with define server command (Example 9-45).

Example 9-45 Define server using the command line.

```
TSMSRV03> define server atalntic_sta serverpassword=password hladdress=atlantic
lladdress=1502
```

3. Then we use the define path commands (Example 9-46).

Example 9-46 Define path commands

```
def path atlantic_sta drlto_1 srct=server destt=dri libr=liblto1
device=/dev/rmt2
def path atlantic_sta drlto_2 srct=server destt=dri libr=liblto1
device=/dev/rmt3
```

Preparation

We prepare to test LAN-free backup failure and recovery:

1. We verify that the cluster services are running with the **lssrc -g cluster command** on both nodes.
2. On the resource group secondary node, we use **tail -f /tmp/hacmp.out** to monitor cluster operation.
3. Then we start a LAN-free client restore using the command line (Example 9-47).

Example 9-47 Client sessions starting

```
Node Name: ATLANTIC
Session established with server TSMSRV03: AIX-RS/6000
  Server Version 5, Release 3, Level 0.0
```

Archived

```
  Server date/time: 02/15/05   18:12:09  Last access: 02/15/05   17:41:22

tsm> restore -subdir=yes "/install/backups/*"
Restore function invoked.

ANS1247I Waiting for files from the server...
Restoring             256 /install/backups [Done]
 ** Interrupted **]
ANS1114I Waiting for mount of offline media.
Restoring   1,034,141,696 /install/backups/520005.tar [Done]
<    1.27 GB> [ - ]
```

4. On the server, we wait for the Storage Agent tape mount messages (Example 9-48).

Example 9-48 Tape mount for LAN-free messages

```
ANR8337I LTO volume ABA924 mounted in drive DRLTO_1 (/dev/rmt2).
ANR0510I Session 13 opened input volume ABA924.
```

5. On the Storage Agent, we verify that data is being transferred, routing to it the **query session** command (Example 9-49).

Example 9-49 Query session for data transfer

```
tsm: TSMSRV03>ATLANTIC_STA:q se

  Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- --------------------
    10 Tcp/Ip IdleW    0 S    5.5 K     257 Serv- AIX-RS/- TSMSRV03
                                             er    6000
    13 Tcp/Ip SendW    0 S    1.6 G     383 Node  AIX      ATLANTIC
    14 Tcp/Ip Run      0 S    1.2 K   1.9 K Serv- AIX-RS/- TSMSRV03
                                             er    6000
```

Failure

Now we make the server fail:

1. Being sure that client is restoring using the LAN-free method, we issue **halt -q** on the AIX server running the Tivoli Storage Manager server; the **halt -q** command stops any activity immediately and powers off the server.
2. The Storage Agent gets errors for the dropped server connection and unmounts the tape (Example 9-50).

Example 9-50 Storage unmount the tapes for the dropped server connection

```
ANR8214E Session open with 9.1.39.74 failed due to connection refusal.
```

```
ANR0454E Session rejected by server TSMSRV03, reason: Communication Failure.
ANR3602E Unable to communicate with database server.
ANR3602E Unable to communicate with database server.
ANR0107W bfrtrv.c(668): Transaction  was not committed due to an internal
error.
ANR8216W Error sending data on socket 12.  Reason 32.
ANR0479W Session 10 for server TSMSRV03 (AIX-RS/6000) terminated - connection
with server severed.
ANR8216W Error sending data on socket 12.  Reason 32.
ANR0546W Retrieve or restore failed for session 13 for node ATLANTIC (AIX)
internal server error detected.
[...]
ANR0514I Session 13 closed volume ABA924.
[...]
ANR8214E Session open with 9.1.39.74 failed due to connection refusal.
[...]
ANR8336I Verifying label of LTO volume ABA924 in drive DRLTO_1 (/dev/rmt2).
[...]
ANR8938E Initialization failed for Shared library LIBLTO1; will retry within 5
minute(s).
[...]
ANR8468I LTO volume ABA924 dismounted from drive DRLTO_1 (/dev/rmt2) in library
LIBLTO1.
```

3. Then the client interrupts the restore operation (Example 9-51).

Example 9-51 client stops receiving data

```
<    1.92 GB> [ - ]ANS9201W LAN-free path failed.
Node Name: ATLANTIC

Total number of objects restored:         2
Total number of objects failed:           0
Total number of bytes transferred:     1.92 GB
LanFree data bytes:                    1.92 GB
Data transfer time:                  194.97 sec
Network data transfer rate:        10,360.53 KB/sec
Aggregate data transfer rate:      4,908.31 KB/sec
Elapsed processing time:           00:06:51
ANS1301E Server detected system error

tsm>
```

Recovery

Here is how the failure is managed:

1. The secondary cluster node takes over the resources and restarts the Tivoli Storage Manager server.

2. Once the server is restarted, it reconnects to the Storage Agent (Example 9-52).

Example 9-52 The restarted Tivoli Storage Manager rejoin the Storage Agent.

```
ANR8439I SCSI library LIBLTO1 is ready for operations.
ANR0408I Session 1 started for server ATLANTIC_STA (AIX-RS/6000) (Tcp/Ip) for
storage agent.  (SESSION: 1)
ANR0408I Session 2 started for server ATLANTIC_STA (AIX-RS/6000) (Tcp/Ip) for
library sharing.  (SESSION: 2)
ANR0409I Session 2 ended for server ATLANTIC_STA (AIX-RS/6000). (SESSION: 2)
ANR0408I Session 3 started for server ATLANTIC_STA (AIX-RS/6000) (Tcp/Ip) for
library sharing.  (SESSION: 2)
ANR0409I Session 3 ended for server ATLANTIC_STA (AIX-RS/6000). (SESSION: 2)
ANR0408I Session 4 started for server ATLANTIC_STA (AIX-RS/6000) (Tcp/Ip) for
event logging.  (SESSION: 4)
```

3. Library recovery is successful for the Storage Agent (Example 9-53).

Example 9-53 Library recovery for Storage Agent

```
ANR8439I SCSI library LIBLTO1 is ready for operations.
ANR0408I Session 1 started for server ATLANTIC_STA (AIX-RS/6000) (Tcp/Ip) for
storage agent.  (SESSION: 1)
ANR0408I Session 2 started for server ATLANTIC_STA (AIX-RS/6000) (Tcp/Ip) for
library sharing.  (SESSION: 2)
ANR0409I Session 2 ended for server ATLANTIC_STA (AIX-RS/6000). (SESSION: 2)
ANR0408I Session 3 started for server ATLANTIC_STA (AIX-RS/6000) (Tcp/Ip) for
library sharing.  (SESSION: 2)
ANR0409I Session 3 ended for server ATLANTIC_STA (AIX-RS/6000). (SESSION: 2)
ANR0408I Session 4 started for server ATLANTIC_STA (AIX-RS/6000) (Tcp/Ip) for
event logging.  (SESSION: 4)
```

4. The client restore command is re-issued with the **replace=all** option (Example 9-54) and the volume is mounted (Example 9-55).

Example 9-54 New restore operation

```
tsm> restore -subdir=yes -replace=all "/install/backups/*"
Restore function invoked.

ANS1247I Waiting for files from the server...
 ** Interrupted **]
ANS1114I Waiting for mount of offline media.
Restoring   1,034,141,696 /install/backups/520005.tar [Done]
Restoring   1,034,141,696 /install/backups/tarfile.tar [Done]
Restoring     809,472,000 /install/backups/VCS_TSM_package.tar [Done]

Restore processing finished.
```

```
Total number of objects restored:         3
Total number of objects failed:           0
Total number of bytes transferred:     2.68 GB
Data transfer time:                  248.37 sec
Network data transfer rate:        11,316.33 KB/sec
Aggregate data transfer rate:      7,018.05 KB/sec
Elapsed processing time:           00:06:40
```

Example 9-55 Volume mounted for restore after the recovery

```
ANR8337I LTO volume ABA924 mounted in drive DRLTO_1 (/dev/rmt2).
ANR0510I Session 9 opened input volume ABA924.
ANR0514I Session 9 closed volume ABA924.
```

Result summary

Once restarted on the secondary node, the Tivoli Storage Manager server reconnects to the Storage Agent for the shared library recovery and takes control of the removable storage resources.

Then we are able to restart our restore operation without any problem.

9.5.4 Failure during disk to tape migration operation

Now we start testing failure during server operations, at first, for a migration.

Objectives

We are testing the recovery of a failure during a disk to tape migration operation and checking to see if the operation continues.

Preparation

Here we prepare for a failure during the migration test:

1. We verify that the cluster services are running with the **`lssrc -g cluster command`** on both nodes.
2. On the resource group secondary node, we use **`tail -f /tmp/hacmp.out`** to monitor cluster operation.
3. We have a disk storage pool used at 87%, with a tape storage pool as next.
4. Lowering **`highMig`** below the used percentage, we make the migration begin.
5. We wait for a tape cartridge mount: Example 9-56 before “crash and restart”.
6. Then we check for data being transferred form disk to tape using the **`query process`** command.

Archived

Failure

We use the **halt -q** command to stop AIX immediately and power off the server.

Recovery

Now we see how the failure is managed:

1. The secondary cluster nodes take over the resources.
2. The Tivoli Storage Manager server is restarted.
3. The tape is unloaded by the reset issued from the TSM server at its restart.
4. Once the server is restarted, the migration restarts because of the used percentage still above the **highMig** percentage (Example 9-56).

Example 9-56 Migration restarts after a takeover

```
02/01/05   07:57:46      ANR0984I Process 1 for MIGRATION started in the
BACKGROUND at 07:57:46. (PROCESS: 1)
02/01/05   07:57:46      ANR1000I Migration process 1 started for storage pool
SPD_BCK automatically, highMig=20, lowMig=10, duration=No. (PROCESS: 1)
02/01/05   07:58:14      ANR8337I LTO volume 029AKK mounted in drive DRLTO_1
(/dev/rmt0). (PROCESS: 1)
02/01/05   07:58:14      ANR1340I Scratch volume 029AKK is now defined in
storage pool TAPEPOOL. (PROCESS: 1)
02/01/05   07:58:14      ANR0513I Process 1 opened output volume 029AKK.
(PROCESS: 1)

[crash and restart]

02/01/05   08:00:09      ANR4726I The NAS-NDMP support module has been loaded.
02/01/05   08:00:09      ANR1794W TSM SAN discovery is disabled by options.
02/01/05   08:00:18      ANR2803I License manager started.
02/01/05   08:00:18      ANR8200I TCP/IP driver ready for connection with
clients on port 1500.
02/01/05   08:00:18      ANR2560I Schedule manager started.
02/01/05   08:00:18      ANR0993I Server initialization complete.
02/01/05   08:00:18      ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli
is now ready for use.
02/01/05   08:00:18      ANR2828I Server is licensed to support Tivoli Storage
Manager Basic Edition.
02/01/05   08:00:18      ANR2828I Server is licensed to support Tivoli Storage
Manager Extended Edition.
02/01/05   08:00:19      ANR1305I Disk volume /tsm/dp1/bckvol1 varied online.
02/01/05   08:00:20      ANR0984I Process 1 for MIGRATION started in the
BACKGROUND at 08:00:20. (PROCESS: 1)
02/01/05   08:00:20      ANR1000I Migration process 1 started for storage pool
SPD_BCK automatically, highMig=20, lowMig=10, duration=No. (PROCESS: 1)
02/01/05   08:00:30      ANR8358E Audit operation is required for library
LIBLTO.
```

```
02/01/05   08:00:31      ANR8439I SCSI library LIBLTO is ready for operations.
02/01/05   08:00:58      ANR8337I LTO volume 029AKK mounted in drive DRLTO_1
(/dev/rmt0). (PROCESS: 1)
02/01/05   08:00:58      ANR0513I Process 1 opened output volume 029AKK.
(PROCESS: 1)
```

5. In Example 9-56 we saw that the same tape volume used before is used also.
6. The process terminate successfully (Example 9-57).

Example 9-57 Migration process ending

```
02/01/05   08:11:11      ANR0986I Process 1 for MIGRATION running in the
BACKGROUND processed 48979 items for a total of 18,520,035,328 bytes with a
completion state of SUCCESS at 08:11:11. (PROCESS: 1)
```

7. We turn back to primary node our resource group as described in “Manual fallback (resource group moving)” on page 500.

Result summary

Also in this case, the cluster is able to manage server failure and make Tivoli Storage Manager available in a somewhat longer time, because of the reset and unload of the tape drive.

A new migration process is started because of the **highMig** setting.

The tape volume involved in the failure is still in a read/write state and is reused.

9.5.5 Failure during backup storage pool operation

Now we describe failure during backup storage pool operation.

Objectives

Here we are testing the recovery of a failure during a tape storage pool backup operation and checking to see if we are able to restart the process without any particular intervention.

Preparation

We first prepare the test environment:

1. We verify that the cluster services are running with the **lssrc -g cluster command** on both nodes.
2. On resource group secondary node we use **tail -f /tmp/hacmp.out** to monitor cluster operation.

3. We have a primary sequential storage pool called SPT_BCK containing an amount of backup data and a copy storage pool called SPC_BCK.
4. The backup storage pool SPT_BCK PSC_BCK command is issued.
5. We wait for tape cartridges to mount: see Example 9-58 before “crash and recovery”.
6. Then we check for data being transferred form disk to tape using the **query process** command.

Failure

We use the **halt -q** command to stop AIX and immediately power off the server.

Recovery

1. The secondary cluster nodes take over the resources.
2. The tapes are unloaded by reset issued during cluster takeover operations.
3. The Tivoli Storage Manager server is restarted (Example 9-58).

Example 9-58 Tivoli Storage Manager restarts after a takeover

```
02/01/05   08:43:51      ANR1210I Backup of primary storage pool SPT_BCK to
copy storage pool SPC_BCK started as process 5. (SESSION: 1, PROCESS: 5)
02/01/05   08:43:51      ANR1228I Removable volume 028AKK is required for
storage pool backup. (SESSION: 1, PROCESS: 5)
02/01/05   08:43:52      ANR0512I Process 5 opened input volume 028AKK.
(SESSION: 1, PROCESS: 5)
02/01/05   08:44:19      ANR8337I LTO volume 029AKK mounted in drive DRLTO_2
(/dev/rmt1). (SESSION: 1, PROCESS: 5)
02/01/05   08:44:19      ANR1340I Scratch volume 029AKK is now defined in
storage pool SPC_BCK. (SESSION: 1, PROCESS: 5)
02/01/05   08:44:19      ANR0513I Process 5 opened output volume 029AKK.
(SESSION: 1, PROCESS: 5)

[crash and restart]

02/01/05   08:49:19      ANR4726I The NAS-NDMP support module has been loaded.
02/01/05   08:49:19      ANR1794W TSM SAN discovery is disabled by options.
02/01/05   08:49:28      ANR2803I License manager started.
02/01/05   08:49:28      ANR8200I TCP/IP driver ready for connection with
clients on port 1500.
02/01/05   08:49:28      ANR2560I Schedule manager started.
02/01/05   08:49:28      ANR0993I Server initialization complete.
02/01/05   08:49:28      ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli
is now ready for use.
02/01/05   08:49:28      ANR1305I Disk volume /tsm/dp1/bckvol1 varied online.
02/01/05   08:49:28      ANR2828I Server is licensed to support Tivoli Storage
Manager Basic Edition.
```

```
02/01/05   08:49:28      ANR2828I Server is licensed to support Tivoli Storage
Manager Extended Edition.
02/01/05   08:51:11      ANR8439I SCSI library LIBLTO is ready for operations.
02/01/05   08:51:38      ANR0407I Session 1 started for administrator ADMIN
(AIX) (Tcp/Ip 9.1.39.89(32793)). (SESSION: 1)
02/01/05   08:51:57      ANR2017I Administrator ADMIN issued command: BACKUP
STGPOOL SPT_BCK SPC_BCK  (SESSION: 1)
02/01/05   08:51:57      ANR0984I Process 1 for BACKUP STORAGE POOL started in
the BACKGROUND at 08:51:57. (SESSION: 1, PROCESS: 1)
02/01/05   08:51:57      ANR2110I BACKUP STGPOOL started as process 1.
(SESSION: 1, PROCESS: 1)
02/01/05   08:51:57      ANR1210I Backup of primary storage pool SPT_BCK to
copy storage pool SPC_BCK started as process 1. (SESSION: 1, PROCESS: 1)
02/01/05   08:51:58      ANR1228I Removable volume 028AKK is required for
storage pool backup. (SESSION: 1, PROCESS: 1)
02/01/05   08:52:25      ANR8337I LTO volume 029AKK mounted in drive DRLTO_1
(/dev/rmt0). (SESSION: 1, PROCESS: 1)
02/01/05   08:52:25      ANR0513I Process 1 opened output volume 029AKK.
(SESSION: 1, PROCESS: 1)
02/01/05   08:52:56      ANR8337I LTO volume 028AKK mounted in drive DRLTO_2
(/dev/rmt1). (SESSION: 1, PROCESS: 1)
02/01/05   08:52:56      ANR0512I Process 1 opened input volume 028AKK.
(SESSION: 1, PROCESS: 1)
02/01/05   09:01:43      ANR1212I Backup process 1 ended for storage pool
SPT_BCK.  (SESSION: 1, PROCESS: 1)
02/01/05   09:01:43      ANR0986I Process 1 for BACKUP STORAGE POOL running in
the BACKGROUND processed 20932 items for a total of 16,500,420,858 bytes with a
completion state of SUCCESS at 09:01:43. (SESSION: 1, PROCESS: 1)
```

4. And then we restart the backup storage pool by reissuing the command:
5. The same output tape volume is mounted and used as before: Example 9-58.
6. The process terminate successfully.
7. We turn back to the primary node for our resource group as described in “Manual fallback (resource group moving)” on page 500.

Result summary

Also in this case, the cluster is able to manage server failure and make Tivoli Storage Manager available in a short time; now it has taken 5 minutes total, because of the two tape drives to be reset/unload.

The backup storage pool process has to be restarted, and completed with a consistent state.

The Tivoli Storage Manager database survives the crash with all volumes synchronized.

Archived

The tape volumes involved in the failure have remained in a read/write state and reused.

9.5.6 Failure during database backup operation

Now we describe failure during database backup operation.

Objectives

Here we test the recovery of a failure during database backup.

Preparation

First we prepare the test environment:

1. We verify that the cluster services are running with the **`lssrc -g cluster command`** on both nodes.
2. On the resource group secondary node, we use **`tail -f /tmp/hacmp.out`** to monitor cluster operation.
3. We issue a **`backup db type=full devc=lto`** command.
4. Then we wait for a tape mount and for the first ANR4554I message.

Failure

We use the **`halt -q`** command to stop AIX immediately and power off the server.

Recovery

Here we see how the failure is managed:

1. The secondary cluster nodes take over the resources.
2. The tape is unloaded by reset issued during cluster takeover operations.
3. The Tivoli Storage Manager server is restarted (Example 9-59).

Example 9-59 Tivoli Storage Manager restarts after a takeover

```
02/01/05   09:12:07      ANR2280I Full database backup started as process 2.  (SESSION: 1, PROCESS: 2)

02/01/05   09:13:04      ANR8337I LTO volume 030AKK mounted in drive
DRLTO_1 (/dev/rmt0). (SESSION: 1, PROCESS: 2)

02/01/05   09:13:04      ANR0513I Process 2 opened output volume 030AKK.
(SESSION: 1, PROCESS: 2)

02/01/05   09:13:07      ANR1360I Output volume 030AKK opened (sequence
number 1).  (SESSION: 1, PROCESS: 2)
```

02/01/05 09:13:08 ANR4554I Backed up 6720 of 13555 database pages. (SESSION: 1, PROCESS: 2)

[crash and recovery]

02/01/05 09:15:42 ANR2100I Activity log process has started.

02/01/05 09:19:21 ANR4726I The NAS-NDMP support module has been loaded.

02/01/05 09:19:21 ANR1794W TSM SAN discovery is disabled by options.

02/01/05 09:19:30 ANR8200I TCP/IP driver ready for connection with clients on port 1500.

02/01/05 09:19:30 ANR2803I License manager started.

02/01/05 09:19:30 ANR2560I Schedule manager started.

02/01/05 09:19:30 ANR0993I Server initialization complete.

02/01/05 09:19:30 ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli is now ready for use.

02/01/05 09:19:30 ANR1305I Disk volume /tsm/dp1/bckvol1 varied online.

02/01/05 09:19:30 ANR2828I Server is licensed to support Tivoli Storage Manager Basic Edition.

02/01/05 09:19:30 ANR2828I Server is licensed to support Tivoli Storage Manager Extended Edition.

02/01/05 09:19:31 ANR0407I Session 1 started for administrator ADMIN (AIX) (Tcp/Ip 9.1.39.75(32794)). (SESSION: 1)

02/01/05 09:21:13 ANR8439I SCSI library LIBLTO is ready for operations.

02/01/05 09:21:36 ANR2017I Administrator ADMIN issued command: QUERY VOLHISTORY t=dbb (SESSION: 2)

02/01/05 09:21:36 ANR2034E QUERY VOLHISTORY: No match found using this criteria. (SESSION: 2)

Archived

02/01/05 09:21:36 ANR2017I Administrator ADMIN issued command: ROLLBACK (SESSION: 2)

02/01/05 09:21:39 ANR2017I Administrator ADMIN issued command: QUERY LIBV (SESSION: 2)

02/01/05 09:22:13 ANR2017I Administrator ADMIN issued command: BACKUP DB t=f devc=lto (SESSION: 2)

02/01/05 09:22:13 ANR0984I Process 1 for DATABASE BACKUP started in the BACKGROUND at 09:22:13. (SESSION: 2, PROCESS: 1)

02/01/05 09:22:13 ANR2280I Full database backup started as process 1. (SESSION: 2, PROCESS: 1)

02/01/05 09:22:40 ANR8337I LTO volume 031AKK mounted in drive DRLTO_1 (/dev/rmt0). (SESSION: 2, PROCESS: 1)

02/01/05 09:22:40 ANR0513I Process 1 opened output volume 031AKK. (SESSION: 2, PROCESS: 1)

02/01/05 09:22:43 ANR1360I Output volume 031AKK opened (sequence number 1). (SESSION: 2, PROCESS: 1)

02/01/05 09:22:43 ANR4554I Backed up 6720 of 13556 database pages. (SESSION: 2, PROCESS: 1)

02/01/05 09:22:43 ANR4554I Backed up 13440 of 13556 database pages. (SESSION: 2, PROCESS: 1)

02/01/05 09:22:46 ANR1361I Output volume 031AKK closed. (SESSION: 2, PROCESS: 1)

02/01/05 09:22:46 ANR0515I Process 1 closed volume 031AKK. (SESSION: 2, PROCESS: 1)

02/01/05 09:22:46 ANR4550I Full database backup (process 1) complete, 13556 pages copied. (SESSION: 2, PROCESS: 1)

---

4. Then we check the state of database backup in execution at halt time with **q vol** and **q libv** commands (Example 9-60).

Example 9-60 Search for database backup volumes

```
tsm: TSMSRV03>q volh t=dbb
ANR2034E QUERY VOLHISTORY: No match found using this criteria.
```

```
ANS8001I Return code 11.

tsm: TSMSRV03>q libv

Library Name   Volume Name   Status    Owner      Last Use    Home      Device
                                                              Element   Type
------------   -----------   -------   --------   ---------   -------   ------

LIBLTO         028AKK        Private   TSMSRV03   Data        4,104     LTO
LIBLTO         029AKK        Private   TSMSRV03   Data        4,105     LTO
LIBLTO         030AKK        Private   TSMSRV03   DbBackup    4,106     LTO
LIBLTO         031AKK        Scratch   TSMSRV03               4,107     LTO
```

5. For Example 9-60 we see that the volume state has been reserved for database backup but the operation has not finished.
6. We used **BACKUP DB t=f devc=lto** to start a new database backup process.
7. The new process skips the previous volume, takes a new one, and completes as can be seen in the final portion of actlog in Example 9-59.
8. Then we have to return to scratch the volume 030AKK with the command, **upd libv LIBLTO 030AKK status=scr**.
9. At the end of testing, we turn backup to the primary node for our resource group as in “Manual fallback (resource group moving)” on page 500.

Result summary

Also in this case, the cluster is able to manage server failure and make Tivoli Storage Manager available in a short time.

Database backup has to be restarted.

The tape volume use in the database backup process running at failure time has remained in a non-scratch status to which has to be returned using a command.

9.5.7 Failure during expire inventory process

Now we describe failure during the expire inventory process.

Objectives

Now we to test the recovery of a Tivoli Storage Manager server failure while expire inventory is running.

Preparation

Here we prepare the test environment.

1. We verify that the cluster services are running with the **lssrc -g cluster command** on both nodes.
2. On the resource group secondary node, we use **tail -f /tmp/hacmp.out** to monitor cluster operation.
3. We issue the **expire inventory** command.
4. Then we wait for the first ANR0811I and ANR4391I messages (Example 9-61).

Example 9-61 Expire inventory process starting

```
ANR2017I Administrator ADMIN issued command: EXPIRE INVENTORY  (SESSION: 1)
ANR0984I Process 2 for EXPIRE INVENTORY started in the BACKGROUND at 11:18:00.
(SESSION: 1, PROCESS: 2)
ANR0811I Inventory client file expiration started as process 2. (SESSION: 1,
PROCESS: 2)
ANR4391I Expiration processing node CL_HACMP03_CLIENT, filespace
/opt/IBM/ISC_old, fsId 1, domain STANDARD, and management class DEFAULT - for
BACKUP type files. (SESSION: 1, PROCESS: 2)
```

Failure

We use the **halt -q** command to stop AIX immediately and power off the server.

Recovery

1. The secondary cluster nodes take over the resources.
2. The Tivoli Storage Manager is restarted (Example 9-62).

Example 9-62 Tivoli Storage Manager restarts

```
ANR4726I The NAS-NDMP support module has been loaded.
ANR1794W TSM SAN discovery is disabled by options.
ANR2803I License manager started.
ANR8200I TCP/IP driver ready for connection with clients on port 1500.
ANR2560I Schedule manager started.
ANR0993I Server initialization complete.
ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli is now ready for use.
ANR1305I Disk volume /tsm/dp1/bckvol1 varied online.
ANR2828I Server is licensed to support Tivoli Storage Manager Basic Edition.
ANR2828I Server is licensed to support Tivoli Storage Manager Extended Edition.
ANR8439I SCSI library LIBLTO1 is ready for operations.
```

3. We check the database and log volumes with and find all of them in a synchronized state (Example 9-63).

Example 9-63 Database and log volumes state

```
sm: TSMSRV03>q dbv

Volume Name       Copy    Volume Name       Copy    Volume Name       Copy
(Copy 1)          Status  (Copy 2)          Status  (Copy 3)          Status
----------------  ------  ----------------  ------  ----------------  ------
/tsm/db1/vol1     Sync'd  /tsm/dbmr1/vol1   Sync'd                    Undef-
                                                                        ined

tsm: TSMSRV03>q logv

Volume Name       Copy    Volume Name       Copy    Volume Name       Copy
(Copy 1)          Status  (Copy 2)          Status  (Copy 3)          Status
----------------  ------  ----------------  ------  ----------------  ------
/tsm/lg1/vol1     Sync'd  /tsm/lgmr1/vol1   Sync'd                    Undef-
                                                                        ined
```

4. We issue the **`expire inventory`** command for a second time to start a new expire process; the new process runs successfully to the end (Example 9-64).

Example 9-64 New expire inventory execution

```
ANR2017I Administrator ADMIN issued command: EXPIRE INVENTORY
ANR0984I Process 1 for EXPIRE INVENTORY started in the BACKGROUND at 11:27:38.
ANR0811I Inventory client file expiration started as process 1.
ANR4391I Expiration processing node CL_HACMP03_CLIENT, filespace
/opt/IBM/ISC_old, fsId 1, domain STANDARD, and management class DEFAULT - for
BACKUP type files.
ANR4391I Expiration processing node CL_HACMP03_CLIENT, filespace
/opt/IBM/ISC_old, fsId 1, domain STANDARD, and management class DEFAULT - for
BACKUP type files.
ANR4391I Expiration processing node CL_HACMP03_CLIENT, filespace /opt/IBM/ISC,
fsId 4, domain STANDARD, and management class DEFAULT - for BACKUP type files.
ANR4391I Expiration processing node KANANGA, filespace /, fsId 1, domain
STANDARD, and management class DEFAULT - for BACKUP type files.
ANR4391I Expiration processing node KANANGA, filespace /usr, fsId 2, domain
STANDARD, and management class DEFAULT - for BACKUP type files.
ANR4391I Expiration processing node KANANGA, filespace /var, fsId 3, domain
STANDARD, and management class DEFAULT - for BACKUP type files.
ANR4391I Expiration processing node AZOV, filespace /, fsId 1, domain STANDARD,
and management class DEFAULT - for BACKUP type files.
ANR4391I Expiration processing node AZOV, filespace /usr, fsId 2, domain
STANDARD, and management class DEFAULT - for BACKUP type files.
ANR4391I Expiration processing node AZOV, filespace /var, fsId 3, domain
STANDARD, and management class DEFAULT - for BACKUP type files.
ANR4391I Expiration processing node AZOV, filespace /opt, fsId 5, domain
STANDARD, and management class STANDARD - for BACKUP type files.
ANR2369I Database backup volume and recovery plan file expiration starting
under process 1.
```

```
ANR0812I Inventory file expiration process 1 completed: examined 88167 objects,
deleting 88139 backup objects, 0 archive objects, 0 DB backup volumes, and 0
recovery plan files. 0 errors were encountered.
ANR0987I Process 1 for EXPIRE INVENTORY running in the BACKGROUND processed
88139 items with a completion state of SUCCESS at 11:29:46.
```

Result summary

Tivoli Storage Manager server restarted with all datafiles synchronized even if an intensive update activity was running.

The process is to be restarted just like any other server interrupted activity.

The new expire inventory process completes to the end without any errors.

Archived

10

AIX and HACMP with IBM Tivoli Storage Manager Client

In this chapter we discuss the details related to the installation and configuration of the Tivoli Storage Manager client V5.3, installed on AIX V5.3, and running as a highly available application under the control of HACMP V5.2.

© Copyright IBM Corp. 2005. All rights reserved.

10.1 Overview

An application that has been made highly available needs a backup program with the same high availability. High Availability Cluster Multi Processing (HACMP) allows scheduled Tivoli Storage Manager client operations to continue processing during a failover situation.

Tivoli Storage Manager in an HACMP environment can back up anything that Tivoli Storage Manager can normally back up. However, we must be careful when backing up non-clustered resources due to the *after failover* effects. Local resources should never be backed up or archived from clustered Tivoli Storage Manager client nodes. Local Tivoli Storage Manager client nodes should be used for local resources.

In our lab, Tivoli Storage Manager client code will be installed on both cluster nodes, and three client nodes will be defined, one clustered and two locals. One dsm.sys file will be used for all Tivoli Storage Manager clients, and located within the default directory /usr/tivoli/tsm/client/ba/bin and hold a unique stanza for each client. We maintain a unique dsm.sys, copied on both nodes, containing all of the three nodes stanzas for an easier synchronizing.

All cluster resource groups which are highly available will have its own Tivoli Storage Manager client. In our lab environment, the ISC with Tivoli Storage Manager Administration Center will be an application within a resource group, and will have the HACMP Tivoli Storage Manager client node included.

For the clustered client nodes, the dsm.opt file, password file, and inclexcl.lst files will be highly available, and located on the application shared disk. The Tivoli Storage Manager client environment variables which reference these option files will be placed in the startup script configured within HACMP.

10.2 Clustering Tivoli Data Protection

Generally, as we configure a Tivoli Storage Manager client to be able to access a Tivoli Storage Manager server across cluster nodes, a clustered API connection can be enabled for Tivoli Data Protection too.

This can be accomplished using the same server stanza the clustered client is using in dsm.sys, or through a dedicated one pointed out by the dsm.opt referenced with the DSMI_CONFIG variable.

Password encryption files and processes that can be required by some Tivoli Data Protection applications will be managed in a different way.

In most cases, the Tivoli Data Protection product manuals have a cluster related section. Refer to these documents if you are interested in clustering Tivoli Data Protection.

10.3 Planning and design

The HACMP planning, installation, and configuration is the same as documented in the previous chapters: Chapter 8, “Establishing an HACMP infrastructure on AIX” on page 417 and Chapter 9, “AIX and HACMP with IBM Tivoli Storage Manager Server” on page 451.

In addition to the documented environment setup for HACMP and the SAN, understanding the Tivoli Storage Manager client requirements is essential.

There must be a requirement to configure an HACMP Tivoli Storage Manager client. The most common requirement would be an application, such as a database product that has been configured and running under HACMP control. In such cases, the Tivoli Storage Manager client will be configured within the same resource group as this application, as an application server. This ensures that the Tivoli Storage Manager client is tightly coupled with the application which requires backup and recovery services.

Our case application is the ISC with the Tivoli Storage Manager Administration Console, which we set up as highly available in Chapter 8, “Establishing an HACMP infrastructure on AIX” on page 417 and Chapter 9, “AIX and HACMP with IBM Tivoli Storage Manager Server” on page 451.

Now we are testing the configuration and clustering for one or more Tivoli Storage Manager client node instances and demonstrating the possibility of restarting a client operation just after the takeover of a crashed node.

Our design considers a 2-node cluster, with 2 local Tivoli Storage Manager client nodes to be used with local storage resources and a clustered client node to manage shared storage resources backup and archive.

To distinguish the 3 client nodes we use different paths for configuration files and running directory, different TCP/IP addresses and different TCP/IP ports (Table 10-1).

Table 10-1 Tivoli Storage Manager client distinguished configuration

Node name	Node directory	TCP/IP addr	TCP/IP port
kanaga	/usr/tivoli/tsm/client/ba/bin	kanaga	1501
azov	/usr/tivoli/tsm/client/ba/bin	azov	1501

Node name	Node directory	TCP/IP addr	TCP/IP port
cl_hacmp03_client	/opt/IBM/ISC/tsm/client/ba/bin	admcnt01	1503

- We use default local paths for the local client nodes instances and a path on a shared filesystem for the clustered one.
- Default port 1501 is used for the local client nodes agent instances while 1503 is used for the clustered one.
- Persistent addresses are used for local Tivoli Storage Manager resources.
- After reviewing the *Backup-Archive Clients Installation and User's Guide*, we then proceed to complete our environment configuration in Table 10-2.

Table 10-2 .Client nodes configuration of our lab

Node 1	
TSM nodename	AZOV
dsm.opt location	/usr/tivoli/tsm/client/ba/bin
Backup domain	/, /usr, /var, /home, /opt
Client Node high level address	azov
Client Node low level address	1501
Node 2	
TSM nodename	KANAGA
dsm.opt location	/usr/tivoli/tsm/client/ba/bin
Backup domain	/, /usr, /var, /home, /opt
Client Node high level address	kanaga
Client Node low level address	1501
Virtual node	
TSM nodename	CL_HACMP03_CLIENT
dsm.opt location	/opt/IBM/ISC/tsm/client/ba/bin
Backup domain	/opt/IBM/ISC
Client Node high level address	admcnt01
Client Node low level address	1503

Archived

10.4 Lab setup

We use the lab already set up for clustered client testing in Chapter 9, “AIX and HACMP with IBM Tivoli Storage Manager Server” on page 451.

10.5 Installation

Our team has already installed all of the needed code now. In the following sections we provide installation details.

10.5.1 HACMP V5.2 installation

We have installed, configured, and tested HACMP prior to this point, and will utilize this infrastructure to hold our highly available application, and our highly available Tivoli Storage Manager client. To reference the HACMP installation, see 8.5, “Lab setup” on page 427.

10.5.2 Tivoli Storage Manager Client Version 5.3 installation

We have installed the Tivoli Storage Manager Client Version 5.3 prior to this point, and will focus our efforts on the configuration in this chapter. To reference the client installation, refer to 9.3.3, “Tivoli Storage Manager Client Installation” on page 456

10.5.3 Tivoli Storage Manager Server Version 5.3 installation

We have installed the Tivoli Storage Manager Server Version 5.3 prior to this point. To reference the server installation, refer to 9.3.4, “Installing the Tivoli Storage Manager Server software” on page 460.

10.5.4 Integrated Solution Console and Administration Center

We have installed the Integrated Solution Console (ISC) and Administration Center prior to this point, and will utilize this function for configuration tasks throughout this chapter, and future chapters. To reference the ISC and Administration Center installation, see 9.3.5, “Installing the ISC and the Administration Center” on page 464.

Archived

10.6 Configuration

Here we configure a highly available node, tied to a highly available application.

1. We have already defined a basic client configuration for use with both the local clients and the administrative command line interface, shown in 9.3.1, “Tivoli Storage Manager Server AIX filesets” on page 455.
2. We then start a Tivoli Storage Manager administration command line client by using the **dsmadmc** command in AIX.
3. Next, we issue the **register node cl_hacmp03_client password passexp=0** Tivoli Storage Manager command.
4. Then, on the primary HACMP node in which the cluster application resides, we create a directory on the application resource shared disk to hold the Tivoli Storage Manager configuration files. In our case, the path is /opt/IBM/ISC/tsm/client/ba/bin, with the mount point for the filesystem being /opt/IBM/ISC.
5. Now, we copy the default dsm.opt.smp to shared disk directory as dsm.opt and edit the file with the servername to be used by this client (Example 10-1).

Example 10-1 dsm.opt file contents located in the application shared disk

```
kanaga/opt/IBM/ISC/tsm/client/ba/bin: more dsm.opt
************************************************
* Tivoli Storage Manager                       *
*                                              *
************************************************
*                                              *
*  This servername is the reference for the    *
*  highly available TSM client.                *
*                                              *
************************************************

SErvername              tsmsrv03_ha
```

6. And then we add a new stanza into dsm.sys for the high available Tivoli Storage Manager client nodes, as shown in Example 10-2, with:
   a. *clusternode* parameter set to yes.
      Clusternode set to yes makes the password encryption not affected by the hostname, so we are able to use the same password file on both nodes.
   b. *passworddir* parameter points to a shared directory.
   c. *managedservices* set to “*schedule webclient*”, to have the *dsmc sched* waked up by the client acceptor daemon at schedule start time as from the example script as suggested in the *UNIX and Linux Backup-Archive Clients Installation and User’s Guide*.

d. Last but most important, we add a domain statement for our shared filesystems. Domain statements are required to tie each filesystem to the corresponding Tivoli Storage Manager client node. Without that, each node will save all of the local mounted filesystems during incremental backups.

> **Important:** When domain statements, one or more, are used in a client configuration, **only** those domains (filesystems) will be backed up during incremental backup.

Example 10-2 dsm.sys file contents located in the default directory

```
kanaga/usr/tivoli/tsm/client/ba/bin: more dsm.sys
************************************************************************
* Tivoli Storage Manager                                               *
*                                                                      *
* Client System Options file for AIX                                   *
************************************************************************

* Server stanza for admin connection purpose
SErvername              tsmsrv03_admin
   COMMMethod           TCPip
   TCPPort              1500
   TCPServeraddress     9.1.39.75
   ERRORLOGRETENTION    7
   ERRORLOGname         /usr/tivoli/tsm/client/ba/bin/dsmerror.log

* Server stanza for the HACMP highly available client connection purpose
SErvername              tsmsrv03_ha
nodename                cl_hacmp03_client
   COMMMethod           TCPip
   TCPPort              1500
   TCPServeraddress     9.1.39.74
   HTTPPORT             1582
   ERRORLOGRETENTION    7
   ERRORLOGname         /opt/IBM/ISC/tsm/client/ba/bin/dsm_error.log
   passwordaccess       generate
   clusternode          yes
   passworddir          /opt/IBM/ISC/tsm/client/ba/bin
   managedservices      schedule webclient
   domain               /opt/IBM/ISC
```

Archived

7. We then connect to the Tivoli Storage Manager server using **dsmc -server=tsmsrv03_ha set password <*old_password*> <*new_password*>** from the AIX command line. This will generate the TSM.PWD file as shown in Example 10-3.

Example 10-3 Current contents of the shared disk directory for the client

```
kanaga/opt/IBM/ISC/tsm/client/ba/bin: ls -l
total 16
-rw-------   1 root     system          151 Jan 26 09:58 TSM.PWD
-rw-r--r--   1 root     system          470 Jan 27 14:25 dsm.opt
```

8. Next, we copy the Tivoli Storage Manager samples scripts (or create your own) for starting and stopping the Tivoli Storage Manager client with HACMP. We created the HACMP script directory '/usr/es/sbin/cluster/local/tsmcli' to hold these scripts, as shown in Example 10-4.

Example 10-4 The HACMP directory which holds the client start and stop scripts

```
kanaga/usr/es/sbin/cluster/local/tsmcli: ls
StartClusterTsmClient.sh StopClusterTsmClient.sh
```

9. Then we edit the sample files, and change the HADIR variable to the location on the shared disk that the Tivoli Storage Manager configuration files reside.
10. Now, the directory and files which have created or changed on the primary node must be copied to the other node. First we create the new hacmp script directory (identical to the primary node)
11. Then, we ftp the start and stop scripts into this new directory.
12. Next, we ftp the /usr/tivoli/tsm/client/ba/bin/dsm.sys.
13. Now, we switch back to the primary node for the application, configure an application server in HAMCP by following the **smit** panels as described in the following sequence.
    a. We select the **Extended Configuration** option.
    b. Then we select the **Extended Resource** Configuration option.
    c. Next we select the **HACMP Extended Resources Configuration** option.
    d. We then select the **Configure HACMP Applications** option.
    e. And then we select the **Configure HACMP Application** Servers option.
    f. Lastly, we select the **Add an Application Server** option, which is shown in Figure 10-1.

```
                              Add Application Server

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* Server Name                                          [as_hacmp03_client]
* Start Script                                         [/usr/es/sbin/cluster/]
* Stop Script                                          [/usr/es/sbin/cluster/]
  Application Monitor Name(s)                                            +

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 10-1 HACMP application server configuration for the client's start and stop

g. Type in the application **Server Name** (we type *as_hacmp03_client*), Start Script, Stop Script, and press Enter.

h. Then we go back to the **Extended Resource Configuration** and select **HACMP Extended Resource Group Configuration**.

i. We select **Change/Show Resources and Attributes for a Resource Group** and pick the resource group name to which to add the application server.

j. In the **Application Servers** field, we choose **as_hacmp03_client** from the list.

k. We press Enter and, after the command result, we go back to the **Extended Configuration** panel.

l. Here we select **Extended Verification and Synchronization**, leave the defaults, and press Enter.

m. The *cluster verification and synchronization utility* runs and after a successfully completion, executes the application server scripts, making the Tivoli Storage Manager cad start script begin running.

Archived

10.7 Testing server and client system failure scenarios

There are many client failure scenarios possible; however, we will test three client failure (failover) events while the clients are accessing the server, two with backup operation and one with restore.

10.7.1 Client system failover while the client is backing up to the disk storage pool

The first test is failover during a backup to disk storage pool.

Objective

In this test we are verifying a scheduled client selective backup operation restarting and completing after a takeover.

Preparation

Here we prepare our test environment:

1. We verify that the cluster services are running with the **`lssrc -g cluster command`** on both nodes.
2. On the resource group secondary node, we use **`tail -f /tmp/hacmp.out`** to monitor cluster operation.
3. Then we schedule a selective backup with client node CL_HACMP03_CLIENT associated to it (Example 10-5).

Example 10-5 Selective backup schedule

```
tsm: TSMSRV03>q sched * test_sched f=d

          Policy Domain Name: STANDARD
               Schedule Name: TEST_SCHED
                 Description:
                      Action: Selective
                     Options: -subdir=yes
                     Objects: /opt/IBM/ISC/
                    Priority: 5
           Start Date/Time: 01/31/05   17:03:14
                    Duration: 1 Hour(s)
             Schedule Style: Classic
                      Period: 1 Day(s)
                 Day of Week: Any
                       Month:
              Day of Month:
               Week of Month:
                  Expiration:
```

```
Last Update by (administrator): ADMIN
         Last Update Date/Time: 02/09/05   17:03:14
              Managing profile:
```

4. We wait for metadata and data sessions starting on server (Example 10-6).

Example 10-6 Client sessions starting

```
02/09/05   17:16:19      ANR0406I Session 452 started for node
CL_HACMP03_CLIENT (AIX) (Tcp/Ip 9.1.39.90(33177)). (SESSION: 452)
02/09/05   17:16:20      ANR0406I Session 453 started for node
CL_HACMP03_CLIENT (AIX) (Tcp/Ip 9.1.39.90(33178)). (SESSION: 453)
```

5. On the server, we verify that data is being transferred via the **query session** command.

Failure

Here we make the server fail:

1. Being sure that client backup is running, we issue **halt -q** on the AIX server running the Tivoli Storage Manager server; the **halt -q** command stops any activity immediately and powers off the client system.
2. The takeover takes more than 60 seconds, the server is not receiving data from the client and cancels a client session based on the *CommTimeOut* setting (Example 10-7).

Example 10-7 Client session cancelled due to the communication timeout.

```
02/09/05   17:20:35      ANR0481W Session 453 for node CL_HACMP03_CLIENT (AIX)
terminated - client did not respond within 60 seconds. (SESSION: 453)
```

Recovery

Here we see how recovery is managed:

1. The secondary cluster node takes over the resources and restarts the Tivoli Storage Manager Client Acceptor Daemon.
2. The scheduler is started and queries for schedules (Example 10-8 and Example 10-9).

Example 10-8 The restarted client scheduler queries for schedules (client log)

```
02/09/05   17:19:20 Directory-->                256 /opt/IBM/ISC/tsm/client/ba
[Sent]
02/09/05   17:19:20 Directory-->               4,096
/opt/IBM/ISC/tsm/client/ba/bin [Sent]
02/09/05   17:21:47 Scheduler has been started by Dsmcad.
02/09/05   17:21:47 Querying server for next scheduled event.
```

```
02/09/05   17:21:47 Node Name: CL_HACMP03_CLIENT
02/09/05   17:21:47 Session established with server TSMSRV03: AIX-RS/6000
02/09/05   17:21:47   Server Version 5, Release 3, Level 0.0
02/09/05   17:21:47   Server date/time: 02/09/05   17:21:47  Last access:
02/09/05   17:20:41

02/09/05   17:21:47 --- SCHEDULEREC QUERY BEGIN
[...]
02/09/05   17:30:51 Next operation scheduled:
02/09/05   17:30:51
------------------------------------------------------------
02/09/05   17:30:51 Schedule Name:         TEST_SCHED
02/09/05   17:30:51 Action:                Selective
02/09/05   17:30:51 Objects:               /opt/IBM/ISC/
02/09/05   17:30:51 Options:               -subdir=yes
02/09/05   17:30:51 Server Window Start:   17:03:14 on 02/09/05
02/09/05   17:30:51
------------------------------------------------------------
```

Example 10-9 The restarted client scheduler queries for schedules (server log)

```
02/09/05   17:20:41      ANR0406I Session 458 started for node
CL_HACMP03_CLIENT (AIX) (Tcp/Ip 9.1.39.89(37431)). (SESSION: 458)
02/09/05   17:20:41      ANR1639I Attributes changed for node
CL_HACMP03_CLIENT: TCP Name from kanaga to azov, TCP Address from 9.1.39.90 to
9.1.39.89, GUID from 00.00.00.00.6e.5c.11.d9.ae.7e.08.63.0a.01.01.5a to
00.00.00.00.6e.73.11.d9.98.cb.08.63.0a.01.01.59. (SESSION: 458)
02/09/05   17:20:41      ANR0403I Session 458 ended for node CL_HACMP03_CLIENT
(AIX). (SESSION: 458)
02/09/05   17:21:47      ANR0406I Session 459 started for node
CL_HACMP03_CLIENT (AIX) (Tcp/Ip 9.1.39.74(37441)). (SESSION: 459)
02/09/05   17:21:47      ANR1639I Attributes changed for node
CL_HACMP03_CLIENT: TCP Address from 9.1.39.89 to 9.1.39.74. (SESSION: 459)
02/09/05   17:21:47      ANR0403I Session 459 ended for node CL_HACMP03_CLIENT
(AIX). (SESSION: 459)
```

3. The backup operation restarts and goes through a successful completion (Example 10-10).

Example 10-10 The restarted backup operation

```
Executing scheduled command now.
02/09/05   17:30:51 --- SCHEDULEREC OBJECT BEGIN TEST_SCHED 02/09/05   17:03:14
02/09/05   17:30:51 Selective Backup function invoked.

02/09/05   17:30:52 ANS1898I ***** Processed     4,000 files *****
02/09/05   17:30:52 Directory-->               4,096 /opt/IBM/ISC/ [Sent]
02/09/05   17:30:52 Directory-->                 256
/opt/IBM/ISC/${SERVER_LOG_ROOT} [Sent]
```

```
02/09/05   17:30:52 Directory-->               4,096 /opt/IBM/ISC/AppServer
[Sent]
02/09/05   17:30:52 Directory-->               4,096 /opt/IBM/ISC/PortalServer
[Sent]
02/09/05   17:30:52 Directory-->                 256 /opt/IBM/ISC/Tivoli [Sent]
[...]
02/09/05   17:30:56 Normal File-->                96
/opt/IBM/ISC/AppServer/installedApps/DefaultNode/wps.ear/wps.war/doc/pt_BR/Info
Center/help/images/header_next.gif [Sent]
02/09/05   17:30:56 Normal File-->             1,890
/opt/IBM/ISC/AppServer/installedApps/DefaultNode/wps.ear/wps.war/doc/pt_BR/Info
Center/help/images/tabs.jpg [Sent]
02/09/05   17:30:56 Directory-->                 256
/opt/IBM/ISC/AppServer/installedApps/DefaultNode/wps.ear/wps.war/doc/ru/InfoCen
ter [Sent]

02/09/05   17:34:01 Selective Backup processing of '/opt/IBM/ISC/*' finished
without failure.

02/09/05   17:34:01 --- SCHEDULEREC STATUS BEGIN
02/09/05   17:34:01 Total number of objects inspected:   39,773
02/09/05   17:34:01 Total number of objects backed up:   39,773
02/09/05   17:34:01 Total number of objects updated:         0
02/09/05   17:34:01 Total number of objects rebound:         0
02/09/05   17:34:01 Total number of objects deleted:         0
02/09/05   17:34:01 Total number of objects expired:         0
02/09/05   17:34:01 Total number of objects failed:          0
02/09/05   17:34:01 Total number of bytes transferred:     1.73 GB
02/09/05   17:34:01 Data transfer time:                   10.29 sec
02/09/05   17:34:01 Network data transfer rate:        176,584.51 KB/sec
02/09/05   17:34:01 Aggregate data transfer rate:      9,595.09 KB/sec
02/09/05   17:34:01 Objects compressed by:                   0%
02/09/05   17:34:01 Elapsed processing time:           00:03:09
02/09/05   17:34:01 --- SCHEDULEREC STATUS END
02/09/05   17:34:01 --- SCHEDULEREC OBJECT END TEST_SCHED 02/09/05   17:03:14
02/09/05   17:34:01 Scheduled event 'TEST_SCHED' completed successfully.
02/09/05   17:34:01 Sending results for scheduled event 'TEST_SCHED'.
02/09/05   17:34:01 Results sent to server for scheduled event 'TEST_SCHED'.
```

Result summary

The cluster is able to manage server failure and make the Tivoli Storage Manager client available. The client is able to restart its operations successfully to the end. The schedule window is not expired and the backup is restarted.

In this example we use selective backup, so the entire operation is restarted from the beginning, and this can affect backup versioning, tape usage, and whole environment scheduling.

10.7.2 Client system failover while the client is backing up to tape

Our second test is failover during a backup to tape storage pool.

Objective

In this test we are verifying a scheduled client incremental backup to tape operation restarting after a client systems takeover.

Incremental backup of small files to tape storage pools is not a best practice, we are just testing it for differences from when a backup that sends data to disk.

Preparation

We follow these steps:

1. We verify that the cluster services are running with the **`lssrc -g cluster command`** on both nodes.
2. On resource group secondary node we use **`tail -f /tmp/hacmp.out`** to monitor cluster operation.
3. Then we schedule an incremental backup with client node CL_HACMP03_CLIENT association.
4. We wait for the metadata and data sessions starting on server and output volume being mounted and opened (Example 10-11).

Example 10-11 Client sessions starting

```
ANR0406I Session 677 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.90(32853)).
ANR0406I Session 678 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.90(32854)).
ANR8337I LTO volume ABA922 mounted in drive DRLTO_2 (/dev/rmt3).
ANR1340I Scratch volume ABA922 is now defined in storage pool SPT_BCK1.
ANR0511I Session 678 opened output volume ABA922.
```

5. On the server, we verify that data is being transferred via the **`query session`** command (Example 10-12).

Example 10-12 Monitoring data transfer through query session command

```
tsm: TSMSRV03>q se

  Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- --------------------
   677 Tcp/Ip IdleW    0 S    3.5 M     432 Node  AIX      CL_HACMP03_CLIENT
   678 Tcp/Ip Run      0 S      285  87.6 M Node  AIX      CL_HACMP03_CLIENT
```

Note: It can take several seconds to minutes from the volume mount completion to the real data writing because of the tape positioning operation.

Failure

6. Being sure that client backup is running, we issue **halt -q** on the AIX server running the Tivoli Storage Manager client; the **halt -q** command stops any activity immediately and powers off the server.
7. The server is not receiving data from the client, and sessions remain in idlew and recvw state (Example 10-13).

Example 10-13 Query sessions showing hanged client sessions

```
tsm: TSMSRV03>q se

  Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- --------------------
   677 Tcp/Ip IdleW   47 S    5.8 M     727 Node  AIX      CL_HACMP03_CLIENT
   678 Tcp/Ip RecvW   34 S      414 193.6 M Node  AIX      CL_HACMP03_CLIENT
```

Recovery

8. The secondary cluster node takes over the resources and restarts the Tivoli Storage Manager scheduler.
9. Then we see the scheduler querying the server for schedules and restarting the scheduled operation, while the server is cancelling old sessions for the expired communication timeout, and obtaining the same volume used before the crash (Example 10-14 and Example 10-15).

Example 10-14 The client reconnect and restarts incremental backup operations

```
02/10/05   08:50:05 Normal File-->            13,739
/opt/IBM/ISC/AppServer/java/jre/bin/libjsig.a [Sent]
02/10/05   08:50:05 Normal File-->           405,173
/opt/IBM/ISC/AppServer/java/jre/bin/libjsound.a [Sent]
02/10/05   08:50:05 Normal File-->           141,405
/opt/IBM/ISC/AppServer/java/jre/bin/libnet.a [Sent]
02/10/05   08:52:44 Scheduler has been started by Dsmcad.
02/10/05   08:52:44 Querying server for next scheduled event.
02/10/05   08:52:44 Node Name: CL_HACMP03_CLIENT
02/10/05   08:52:44 Session established with server TSMSRV03: AIX-RS/6000
02/10/05   08:52:44   Server Version 5, Release 3, Level 0.0
02/10/05   08:52:44   Server date/time: 02/10/05   08:52:44  Last access:
02/10/05   08:51:43
[...]
02/10/05   08:54:54 Next operation scheduled:
```

```
02/10/05   08:54:54
--------------------------------------------------------------
02/10/05   08:54:54 Schedule Name:         TEST_SCHED
02/10/05   08:54:54 Action:                Incremental
02/10/05   08:54:54 Objects:
02/10/05   08:54:54 Options:               -subdir=yes
02/10/05   08:54:54 Server Window Start:   08:47:14 on 02/10/05
02/10/05   08:54:54
--------------------------------------------------------------
02/10/05   08:54:54
Executing scheduled command now.
02/10/05   08:54:54 --- SCHEDULEREC OBJECT BEGIN TEST_SCHED 02/10/05   08:47:14
02/10/05   08:54:54 Incremental backup of volume '/opt/IBM/ISC'
02/10/05   08:54:56 ANS1898I ***** Processed     4,500 files *****
02/10/05   08:54:57 ANS1898I ***** Processed     8,000 files *****
02/10/05   08:54:57 ANS1898I ***** Processed    10,500 files *****
02/10/05   08:54:57 Normal File-->              336
/opt/IBM/ISC/AppServer/cloudscape/db2j.log [Sent]
02/10/05   08:54:57 Normal File-->          954,538
/opt/IBM/ISC/AppServer/logs/activity.log [Sent]
02/10/05   08:54:57 Normal File-->                6
/opt/IBM/ISC/AppServer/logs/ISC_Portal/ISC_Portal.pid [Sent]
02/10/05   08:54:57 Normal File-->           60,003
/opt/IBM/ISC/AppServer/logs/ISC_Portal/startServer.log [Sent]
```

Example 10-15 The Tivoli Storage Manager accept the client new sessions

```
ANR0406I Session 682 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.89(38386)).
ANR1639I Attributes changed for node CL_HACMP03_CLIENT: TCP Name from kanaga to
azov, TCP Address from 9.1.39.90 to 9.1.39.89, GUID from
00.00.00.00.6e.5c.11.d9.ae.7e.08.63.0a.01.01.5a to
00.00.00.00.6e.73.11.d9.98.cb.08.63.0a.01.01.59.
ANR0403I Session 682 ended for node CL_HACMP03_CLIENT (AIX).
ANR0514I Session 678 closed volume ABA922.
ANR0481W Session 678 for node CL_HACMP03_CLIENT (AIX) terminated - client did
not respond within 60 seconds.
ANR0406I Session 683 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.89(38395)).
ANR0403I Session 683 ended for node CL_HACMP03_CLIENT (AIX).
ANR0406I Session 685 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.89(38399)).
ANR0406I Session 686 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.89(38400)).
ANR0511I Session 686 opened output volume ABA922.
```

10. Then the new operation continues to the end and completes successfully (Example 10-16).

Example 10-16 Query event showing successful result.

```
tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name     Status
-------------------- -------------------- ------------- ------------- ---------
02/10/05   08:47:14  02/10/05   08:48:27  TEST_SCHED    CL_HACMP03_C- Completed
                                                        LIENT
```

Result summary

The cluster is able to manage client failure and make Tivoli Storage Manager client scheduler available on the secondary server, and the client is able to restart its operations successfully to the end.

Since this is an incremental backup, it backs up objects for which the backup operation has not taken place or has not been committed in the previous run and new created or modified files.

We see the server cancelling the tape holding session (Example 10-15 on page 542) for the communication timeout, so we want to check what happens if *CommTimeOut* is set to a higher value than usual for Tivoli Data Protection environments.

10.7.3 Client system failover while the client is backing up to tape with higher CommTimeOut

In this test we are verifying a scheduled client incremental backup to tape operation restarting after a client systems takeover with a greater commtimeout.

Objective

We suspect when something goes wrong in backup or archive operations that it used tapes with a commtimeout greater than the time needed for takeover.

Incremental backup of small files to tape storage pools is not a best practice, we are just testing it for differences from a backup that sends data to disk.

Preparation

Here we prepare the test environment:

1. We stop the Tivoli Storage Manager Server and insert the *CommTimeOut 600* parameter in the Tivoli Storage Manager server options file /tsm/files/dsmserv.opt.

2. Then we restart the server with the cluster script **/usr/es/sbin/cluster/local/tsmsrv/starttsmsrv03.sh**
3. We verify that the cluster services are running with the **lssrc -g cluster command** on both nodes.
4. On the resource group secondary node we use **tail -f /tmp/hacmp.out** to monitor cluster operation.
5. Then we schedule an incremental backup with client node CL_HACMP03_CLIENT association.
6. We wait for the metadata and data sessions starting on the server and output volume being mounted and opened (Example 10-17).

Example 10-17 Client sessions starting

```
ANR0406I Session 4 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.90(32799)).
ANR0406I Session 5 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.90(32800)).
ANR8337I LTO volume ABA922 mounted in drive DRLTO_1 (/dev/rmt2).
ANR0511I Session 5 opened output volume ABA922.
```

7. On the server, we verify that data is being transferred via **query session**.

Note: It takes some seconds from the volume mount completion to the real data writing because of tape positioning operation.

Failure

Now we make the server fail:

1. Being sure that client backup is transferring data, we issue **halt -q** on the AIX server running the Tivoli Storage Manager client; the **halt -q** command stops any activity immediately and powers off the server.
2. The server is not receiving data to server, and sessions remain in idlew and recvw state as for the previous test.

Recovery failure

Here we see how recovery is managed:

1. The secondary cluster nodes takes over the resources and restarts the Tivoli Storage Manager client acceptor daemon.
2. Then we can see the scheduler querying the server for schedules and restarting the scheduled operation, but the new session is not able to obtain a mount point because now the client node hits the *Maximum Mount Points Allowed* parameter: See the bottom part of Example 10-18.

Example 10-18 The client and restarts and hits MAXNUMMP

```
02/10/05   10:32:21 Normal File-->           100,262
/opt/IBM/ISC/AppServer/lib/txMsgs.jar [Sent]
02/10/05   10:32:21 Normal File-->             2,509
/opt/IBM/ISC/AppServer/lib/txRecoveryUtils.jar [Sent]
02/10/05   10:32:21 Normal File-->           111,133
/opt/IBM/ISC/AppServer/lib/uddi4j.jar [Sent]
02/10/05   10:35:09 Scheduler has been started by Dsmcad.
02/10/05   10:35:09 Querying server for next scheduled event.
02/10/05   10:35:09 Node Name: CL_HACMP03_CLIENT
02/10/05   10:35:09 Session established with server TSMSRV03: AIX-RS/6000
02/10/05   10:35:09   Server Version 5, Release 3, Level 0.0
02/10/05   10:35:09   Server date/time: 02/10/05   10:35:09  Last access:
02/10/05   10:34:09
02/10/05   10:35:09 --- SCHEDULEREC QUERY BEGIN
[...]
Executing scheduled command now.
02/10/05   10:35:09 --- SCHEDULEREC OBJECT BEGIN TEST_SCHED 02/10/05   10:17:02
02/10/05   10:35:10 Incremental backup of volume '/opt/IBM/ISC'
02/10/05   10:35:11 ANS1898I ***** Processed     4,000 files *****
02/10/05   10:35:12 ANS1898I ***** Processed     7,000 files *****
02/10/05   10:35:13 ANS1898I ***** Processed    13,000 files *****
02/10/05   10:35:13 Normal File-->               336
/opt/IBM/ISC/AppServer/cloudscape/db2j.log [Sent]
02/10/05   10:35:13 Normal File-->         1,002,478
/opt/IBM/ISC/AppServer/logs/activity.log [Sent]
02/10/05   10:35:13 Normal File-->                 6
/opt/IBM/ISC/AppServer/logs/ISC_Portal/ISC_Portal.pid [Sent]
[...]
02/10/05   10:35:18 ANS1228E Sending of object
'/opt/IBM/ISC/PortalServer/installedApps/taskmanager_PA_1_0_37.ear/taskmanager.
war/WEB-INF/classes/nls/taskmanager_zh.properties' failed
02/10/05   10:35:18 ANS0326E This node has exceeded its maximum number of mount
points.
02/10/05   10:35:18 ANS1228E Sending of object
'/opt/IBM/ISC/PortalServer/installedApps/taskmanager_PA_1_0_37.ear/taskmanager.
war/WEB-INF/classes/nls/taskmanager_zh_TW.properties' failed
02/10/05   10:35:18 ANS0326E This node has exceeded its maximum number of
mountpoints.
```

Troubleshooting

Using the parameter format=detail, we can see the previous data sending session still present and having a volume in output use (Example 10-19).

Example 10-19 Hanged client session with an output volume

```
 Sess Number: 5
             Comm. Method: Tcp/Ip
               Sess State: RecvW
                Wait Time: 58 S
               Bytes Sent: 139.8 M
              Bytes Recvd: 448.7 K
                Sess Type: Node
                 Platform: AIX
              Client Name: CL_HACMP03_CLIENT
      Media Access Status: Current output volume(s):  ABA922,(147 Seconds)
                User Name:
Date/Time First Data Sent:
   Proxy By Storage Agent:
```

That condition makes the number of mount points used to be already set at 1, that is, equal to the maximum allowed for our node, until the communication timeout expires and the session is cancelled.

Problem correction

Here we show how the team solved the problem:

1. We set up an administrator with operator privilege and modify the cad start script as follows
   a. To check about a *Client Acceptor Daemon* clean exit in the last run
   b. Then to search the Tivoli Storage Manager Server database for the CL_HACMP03_CLIENT's sessions that can be holding tape resources in case of a crash.
   c. Finally, a loop on cancelling any sessions found by the query above (we find a loop necessary because sometimes the session is not cancelled immediately at the first attempt)

Note: We are aware that in the client node failover case, all the existing sessions are to be cancelled by communication or idle timeout, so we are confident in what can be done with these client sessions.

In Example 10-20 we show the addition to the startup script.

Example 10-20 Old sessions cancelling work in startup script

```
[...]
# Set a temporary dir for output files
WORKDIR=/tmp
```

```
# Set up an appropriate administrator with operator (best) or system privileges
#  and an admin connection server stanza in dsm.sys.
TSM_ADMIN_CMD="dsmadmc -quiet -se=tsmsrv04_admin -id=script_operator
-pass=password"

# Set variable with node_name of the node being started by this script
tsmnode=CL_HACMP03_CLIENT
# Node name has to be uppercase to match TSM database entries
TSM_NODE=`echo $tsmnode | tr `[a-z] ` `[A-Z]``

#export DSM variables
export DSM_DIR=/usr/tivoli/tsm/client/ba/bin
export DSM_CONFIG=$HADIR/dsm.opt

##############################################
# Check for dsmcad clean exit last time.
##############################################
if [ -f $PIDFILE ]
 then
  # cad already running or not closed by stopscript
    PID=`cat $PIDFILE`
    ps  $PID
    if [ $? -ne 0 ]
     then
      # Old cad killed manually or a server crash has occoured
      # So search for hanged sessions in case of takeover
      COUNT=0
      while $TSM_ADMIN_CMD -outfile=$WORKDIR/SessionsQuery.out "select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='$TSM_NODE'"
        do
         let COUNT=$COUNT+1
         if [ $COUNT -gt 15 ]
          then
           echo "At least one session is not going away ... give up cancelling
it and start the CAD"
            break
         fi
         echo "If this node is restarting or on takeover, most likely now we
need to cancel its previous sessions."
        SESSIONS_TO_CANCEL=`cat $WORKDIR/SessionsQuery.out|grep $TSM_NODE|grep
-v ANS8000I|awk `{print $1}``
         echo $SESSIONS_TO_CANCEL
         for SESS in $SESSIONS_TO_CANCEL
          do
           $TSM_ADMIN_CMD cancel sess $SESS > /dev/null
           sleep 3
          done
        done
    fi
```

```
  echo "No hanged sessions have been left allocated to this node."
fi

# Remove tmp work file
if [ -f $WORKDIR/SessionsQuery.out ]
 then
  rm $WORKDIR/SessionsQuery.out
fi
[...]
```

New test

Here is the new execution of the test:

2. We repeat the above test and we can see what happens in the server activity log when the modified cad start script runs (Example 10-21).
   a. The select for searching a tape holding session.
   b. The cancel command for the above found session.
   c. A new select with no result because the first **cancel session** command is successful.
   d. The restarted client scheduler querying for schedules.
   e. The schedule is still in window, so a new incremental backup operation is started and it obtains the same output volume as before.

Example 10-21 Hanged tape holding sessions cancelling job

```
ANR0407I Session 54 started for administrator ADMIN (AIX)
(Tcp/Ip9.1.39.75(38721)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_HACMP03_CLIENT'
ANR0405I Session 54 ended for administrator ADMIN (AIX).
ANR0407I Session 55 started for administrator ADMIN (AIX) (Tcp/Ip
9.1.39.75(38722)).
ANR2017I Administrator ADMIN issued command: CANCEL SESSION 47
ANR0490I Canceling session 47 for node CL_HACMP03_CLIENT (AIX) .
ANR0524W Transaction failed for session 47 for node CL_HACMP03_CLIENT (AIX) -
data transfer interrupted.
ANR0405I Session 55 ended for administrator ADMIN (AIX).
ANR0514I Session 47 closed volume ABA922.
ANR0483W Session 47 for node CL_HACMP03_CLIENT (AIX) terminated - forced by
administrator.
ANR0407I Session 56 started for administrator ADMIN (AIX) (Tcp/Ip
9.1.39.75(38723)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_HACMP03_CLIENT'
ANR2034E SELECT: No match found using this criteria.
```

```
ANR2017I Administrator ADMIN issued command: ROLLBACK
ANR0405I Session 56 ended for administrator ADMIN (AIX).
ANR0406I Session 57 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.75(38725)).
ANR1639I Attributes changed for node CL_HACMP03_CLIENT: TCP Name from kanaga to
azov, TCP Address from 9.1.39.90 to 9.1.39.75, GUID from
00.00.00.00.6e.5c.11.d9.ae.7e.08.63.0a.01.01.5a to
00.00.00.00.6e.73.11.d9.98.cb.08.63.0a.01.01.59.
ANR0403I Session 57 ended for node CL_HACMP03_CLIENT (AIX).
ANR0406I Session 58 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.75(38727)).
ANR0403I Session 58 ended for node CL_HACMP03_CLIENT (AIX).
ANR0406I Session 60 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.75(38730)).
ANR0406I Session 61 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.75(38731)).
ANR0511I Session 61 opened output volume ABA922.
```

3. Now incremental backup runs successfully to the end as for the previous test and we can see the successful completion of the schedule (Example 10-22).

Example 10-22 Event result

```
tsm: TSMSRV03>q ev * * f=d


Policy Domain Name: STANDARD

   Schedule Name: TEST_SCHED

       Node Name: CL_HACMP03_CLIENT

  Scheduled Start: 02/10/05   14:44:33

    Actual Start: 02/10/05   14:49:53

       Completed: 02/10/05   14:56:24

          Status: Completed

          Result: 0

          Reason: The operation completed successfully.
```

Archived

Result summary

The cluster is able to manage client system failure and make the Tivoli Storage Manager client scheduler available on secondary server; the client is able to restart its operations successfully to the end.

We do some script work for freeing the Tivoli Storage Manager server in advance from hanged sessions that keep the mounted volumes number increased.

This can be avoided also with a higher *MAXUMMP* setting if the environment allows (more mount points and scratch volumes are needed).

10.7.4 Client system failure while the client is restoring

Now we do a scheduled restore scenario, such as the case of an application test environment having data refreshed daily using a production system backup run.

Objective

In this test we are verifying how a restore operation scenario is managed in a client takeover scenario.

In this test we use a scheduled operation with parameter **`replace=all`**, so the restore operation can be restarted from the beginning. In case of a manual restore, the *restartable restore* functionality can be exploited.

Preparation

Here we prepare the test environment.

1. We verify that the cluster services are running with the **`lssrc -g cluster command`** on both nodes.
2. On the resource group secondary node, we use **`tail -f /tmp/hacmp.out`** to monitor cluster operation.
3. Then we schedule a restore operation with client node CL_HACMP03_CLIENT (Example 10-23).

Example 10-23 Restore schedule

```
Policy Domain Name: STANDARD
                Schedule Name: RESTORE_SCHED
                  Description:
                       Action: Restore
                      Options: -subdir=yes -replace=all
                      Objects: /opt/IBM/ISC/backups/*
                     Priority: 5
            Start Date/Time: 01/31/05   19:48:55
                     Duration: 1 Hour(s)
```

```
                Schedule Style: Classic
                        Period: 1 Day(s)
                   Day of Week: Any
                         Month:
                  Day of Month:
                 Week of Month:
                    Expiration:
Last Update by (administrator): ADMIN
         Last Update Date/Time: 02/10/05   19:48:55
              Managing profile:
```

4. We wait for the client session starting on the server and an input volume being mounted and opened for it (Example 10-24).

Example 10-24 Client sessions starting

```
ANR0406I Session 6 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.90(32816)).
ANR8337I LTO volume ABA922 mounted in drive DRLTO_1 (/dev/rmt2).
ANR0510I Session 6 opened input volume ABA922.
```

5. On the server, we verify that data is being transferred via the **query session** command.

Failure

Now we make the server fail:

6. Being sure that client backup is running, we issue **halt -q** on the AIX server running the Tivoli Storage Manager client; the **halt -q** command stops any activity immediately and powers off the server.
7. The server is not receiving data to server, and sessions remain in idlew and recvw state.

Recovery

Here we see how recovery is managed:

8. The secondary cluster node takes over the resources and launches the Tivoli Storage Manager cad start script.
9. We can see in Example 10-25 the server activity log showing that the same events occurred in the backup test above:
   a. The select searching for a tape holding session.
   b. The **cancel** command for the session found above.
   c. A new select with no result because the first **cancel session** command is successful.

d. The restarted client scheduler querying for schedules.

e. The schedule is still in the window, so a new restore operation is started, and it obtains its input volume.

Example 10-25 The server log during restore restart

```
ANR0407I Session 7 started for administrator ADMIN (AIX) (Tcp/Ip
9.1.39.75(39399)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_HACMP03_CLIENT'
ANR0405I Session 7 ended for administrator ADMIN (AIX).
ANR0407I Session 8 started for administrator ADMIN (AIX) (Tcp/Ip
9.1.39.75(39400)).
ANR2017I Administrator ADMIN issued command: CANCEL SESSION 6
ANR0490I Canceling session 6 for node CL_HACMP03_CLIENT (AIX) .
ANR8216W Error sending data on socket 14.  Reason 32.
ANR0514I Session 6 closed volume ABA922.
ANR0483W Session 6 for node CL_HACMP03_CLIENT (AIX) terminated - forced by
administrator.
ANR0405I Session 8 ended for administrator ADMIN (AIX).
ANR0407I Session 9 started for administrator ADMIN (AIX) (Tcp/Ip
9.1.39.75(39401)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_HACMP03_CLIENT'
ANR2034E SELECT: No match found using this criteria.
ANR2017I Administrator ADMIN issued command: ROLLBACK
ANR0405I Session 9 ended for administrator ADMIN (AIX).
ANR0406I Session 10 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.75(39403)).
ANR1639I Attributes changed for node CL_HACMP03_CLIENT: TCP Name from kanaga to
azov, TCP Address from 9.1.39.90 to 9.1.39.75, GUID from
00.00.00.00.6e.5c.11.d9.ae.7e.08.63.0a.01.01.5a to
00.00.00.00.6e.73.11.d9.98.cb.08.63.0a.01.01.59.
ANR0403I Session 10 ended for node CL_HACMP03_CLIENT (AIX).
ANR2017I Administrator ADMIN issued command: QUERY SESSION f=d
ANR2017I Administrator ADMIN issued command: QUERY SESSION f=d
ANR0406I Session 11 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.75(39415)).
ANR0510I Session 11 opened input volume ABA922.
ANR0514I Session 11 closed volume ABA922.
ANR2507I Schedule RESTORE_SCHED for domain STANDARD started at 02/10/05
19:48:55 for node CL_HACMP03_CLIENT completed successfully at 02/10/05
19:59:21.
ANR0403I Session 11 ended for node CL_HACMP03_CLIENT (AIX).
ANR0406I Session 13 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.75(39419)).
ANR0403I Session 13 ended for node CL_HACMP03_CLIENT (AIX).
```

10. The new restore operation completes successfully.

11. In the client log we can see the restore interruption and restart (Example 10-26).

Example 10-26 The Tivoli Storage Manager client log

```
02/10/05   19:54:10 Restoring              47
/opt/IBM/ISC/backups/PortalServer/tmp/reuse18120.xml [Done]
02/10/05   19:54:10 Restoring              47
/opt/IBM/ISC/backups/PortalServer/tmp/reuse34520.xml [Done]
02/10/05   19:54:10 Restoring          37,341
/opt/IBM/ISC/backups/PortalServer/uninstall/wpscore/uninstall.dat [Don
e]
02/10/05   19:56:22 Scheduler has been started by Dsmcad.
02/10/05   19:56:22 Querying server for next scheduled event.
02/10/05   19:56:22 Node Name: CL_HACMP03_CLIENT
02/10/05   19:56:22 Session established with server TSMSRV03: AIX-RS/6000
02/10/05   19:56:22   Server Version 5, Release 3, Level 0.0
02/10/05   19:56:22   Server date/time: 02/10/05   19:56:22  Last access:
02/10/05   19:55:22

02/10/05   19:56:22 --- SCHEDULEREC QUERY BEGIN
02/10/05   19:56:22 --- SCHEDULEREC QUERY END
02/10/05   19:56:22 Next operation scheduled:
02/10/05   19:56:22
------------------------------------------------------------
02/10/05   19:56:22 Schedule Name:         RESTORE_SCHED
02/10/05   19:56:22 Action:                Restore
02/10/05   19:56:22 Objects:               /opt/IBM/ISC/backups/*
02/10/05   19:56:22 Options:               -subdir=yes -replace=all
02/10/05   19:56:22 Server Window Start:   19:48:55 on 02/10/05
02/10/05   19:56:22
------------------------------------------------------------
02/10/05   19:56:22
Executing scheduled command now.
02/10/05   19:56:22 --- SCHEDULEREC OBJECT BEGIN TEST_SCHED 02/10/05   19:48:55
02/10/05   19:56:22 Restore function invoked.

02/10/05   19:56:23 ANS1899I ***** Examined     1,000 files *****
[...]
02/10/05   19:56:24 ANS1899I ***** Examined    20,000 files *****
02/10/05   19:56:25 Restoring             256
/opt/IBM/ISC/backups/AppServer/config/.repository [Done]
02/10/05   19:56:25 Restoring             256
/opt/IBM/ISC/backups/AppServer/config/cells/DefaultNode/applications/AdminCente
r_PA_1_0_69.ear [Done]
02/10/05   19:56:25 Restoring             256
/opt/IBM/ISC/backups/AppServer/config/cells/DefaultNode/applications/Credential
_nistration_PA_1_0_3C.ear [Done]
```

```
[...]
02/10/05   19:59:19 Restoring          20,285
/opt/IBM/ISC/backups/backups/_uninst/uninstall.dat [Done]
02/10/05   19:59:19 Restoring       6,943,848
/opt/IBM/ISC/backups/backups/_uninst/uninstall.jar [Done]
02/10/05   19:59:19
Restore processing finished.
02/10/05   19:59:21 --- SCHEDULEREC STATUS BEGIN
02/10/05   19:59:21 Total number of objects restored:    20,338
02/10/05   19:59:21 Total number of objects failed:          0
02/10/05   19:59:21 Total number of bytes transferred:     1.00 GB
02/10/05   19:59:21 Data transfer time:                   47.16 sec
02/10/05   19:59:21 Network data transfer rate:        22,349.90 KB/sec
02/10/05   19:59:21 Aggregate data transfer rate:      5,877.97 KB/sec
02/10/05   19:59:21 Elapsed processing time:           00:02:59
02/10/05   19:59:21 --- SCHEDULEREC STATUS END
02/10/05   19:59:21 --- SCHEDULEREC OBJECT END RESTORE_SCHED 02/10/05
19:48:55
02/10/05   19:59:21 --- SCHEDULEREC STATUS BEGIN
02/10/05   19:59:21 --- SCHEDULEREC STATUS END
02/10/05   19:59:21 Scheduled event 'RESTORE_SCHED' completed successfully.
02/10/05   19:59:21 Sending results for scheduled event 'RESTORE_SCHED'.
02/10/05   19:59:21 Results sent to server for scheduled event 'RESTORE_SCHED'.
```

Result summary

The cluster is able to manage client failure and make Tivoli Storage Manager client scheduler available on the secondary server; the client is able to restart its operations successfully to the end.

Since this is a scheduled restore with **replace=all**, it is restarted from the beginning and completes successfully, overwriting the previously restored data. Otherwise, in a manual restore case, we can have a restartable one. Both client and server interfaces, in Example 10-27, can be used searching for restartable restores.

Example 10-27 Query server for restartable restores

```
tsm: TSMSRV03>q rest

  Sess Restore     Elapsed Node Name                Filespace         FSID
Number State       Minutes                          Name
------ ----------- ------- ------------------------ ----------- ----------
    -1 Restartable       8 CL_HACMP03_CLIENT        /opt/IBM/I-          1
                                                    SC
```

11

AIX and HACMP with the IBM Tivoli Storage Manager Storage Agent

This chapter describes our team's implementation of the IBM Tivoli Storage Manager Storage Agent under the control of the HACMP V5.2 product, which runs on AIX V5.3.

© Copyright IBM Corp. 2005. All rights reserved.

11.1 Overview

We can configure the Tivoli Storage Manager client and server so that the client, through a Storage Agent, can move its data directly to storage on a SAN. This function, called LAN-free data movement, is provided by IBM Tivoli Storage Manager for Storage Area Networks.

As part of the configuration, a Storage Agent is installed on the client system. Tivoli Storage Manager supports both tape libraries and FILE libraries. This feature supports SCSI, 349X, and ACSLS tape libraries.

For more information on configuring Tivoli Storage Manager for LAN-free data movement, see the *IBM Tivoli Storage Manager Storage Agent User's Guide*.

The configuration procedure we follow will depend on the type of environment we implement.

Tape drives: SCSI reserve concern

When a server running Tivoli Storage Manager server or Storage Agent crashes while using a tape drive, its SCSI reserve remains, preventing other servers from accessing the tape resources.

A new library parameter called ***resetdrives,*** which specifies whether the server performs a target reset when the server is restarted or when a library client or Storage Agent re-connection is established, has been made available in AIX Tivoli Storage Manager server for AIX 5.3. This parameter only applies to SCSI, 3494, Manual, and ACSLS type libraries.

An external SCSI reset is still needed to free up those resources if the library server is other than 5.3 or later running on AIX, or if the *resetdrives* parameter is set to no.

For those cases, we adapt a sample script, provided for starting the server in previous versions, to start up the Storage Agent.

We can't have HACMP do it using tape resources management, because it will reset all of the tape drives, even if they are in use from the server or other Storage Agents.

Archived

Advantage of clustering a Storage Agent

In a clustered client environment, Storage Agents can be a local or a clustered resource, for both backup/archive and API clients. They can be accessed, using shared memory communication with a specific port number or TCP/IP communication with loopback address and specific port number, or a TCP/IP address made *highly available*.

The advantage of clustering a Storage Agent, in a machine failover scenario, is to have Tivoli Storage Manager server reacting immediately when the Storage Agent restarts on a standby machine.

When the Tivoli Storage Manager server notices a Storage Agent restarting, it checks for resources previously allocated to that Storage Agent. If there are any, it tries to take them back, and issues SCSI resets if needed.

Otherwise, Tivoli Storage Manager reacts on a timeout only basis to Storage Agent failures.

11.2 Planning and design

Our design considers two AIX servers with one virtual Storage Agent to be used by a single virtual client. This design will simulate the most common configuration in production, which is an application such as a database product that has been configured as highly available. Now we will require a backup client and Storage Agent which will follow the application as it transitions though a cluster.

On our servers, local Storage Agents running with default environment settings are configured too. We can have more than one dsmsta running on a single machine as for servers and clients.

Clustered Tivoli Storage Manager resources are required for clustered application backups, so they have to be tied to the same resource group. In our example, we are using the ISC and Tivoli Storage Manager Administration Center as clustered applications, even if not much data is within them, but we are just demonstrating a configuration. Table 11-1 shows the location of our dsmsta.opt and devconfig.txt files.

Archived

A Storage Agent can be run on a directory other than the default one using the same environment setting as for a Tivoli Storage Manager server:

- To distinguish the two storage managers running on the same server, we use a different path for configuration files and running directory and different TCP/IP ports, as shown in Table 11-1.

Table 11-1 Storage Agents distinguished configuration

STA instance	Instance path	TCP/IP addr	TCP/IP port
kanaga_sta	/usr/tivoli/tsm/Storageagent/bin	kanaga	1502
azov_sta	/usr/tivoli/tsm/Storageagent/bin	azov	1502
cl_hacmp03_sta	/opt/IBM/ISC/tsm/Storageagent/bin	admcnt01	1504

- We use default local paths for the local Storage Agent instances and a path on a shared filesystem for the clustered one.
- Port 1502 is used for the local Storage Agent instances while 1504 is used for the clustered one.
- Persistent addresses are used for local Tivoli Storage Manager resources.
- Here we are using TCP/IP as a communication method, but shared memory also applies.
- After reviewing the User's Guide, we then proceed to fill out the Configuration Information Worksheet provided in the *User's Guide*.

Archived

Our complete environment configuration is shown in Table 11-2, Table 11-3, and Table 11-4.

Table 11-2 .LAN-free configuration of our lab

Node 1	
TSM nodename	AZOV
dsm.opt location	/usr/tivoli/tsm/client/ba/bin
Storage Agent name	AZOV_STA
dsmsta.opt and devconfig.txt location	/usr/tivoli/tsm/Storageagent/bin
Storage Agent high level address	azov
Storage Agent low level address	1502
LAN-free communication method	Tcpip
Node 2	
TSM nodename	KANAGA
dsm.opt location	/usr/tivoli/tsm/client/ba/bin
Storage Agent name	KANAGA_STA
dsmsta.opt and devconfig.txt location	/usr/tivoli/tsm/Storageagent/bin
Storage Agent high level address	kanaga
Storage Agent low level address	1502
LAN-free communication method	Tcpip
Virtual node	
TSM nodename	CL_HACMP03_CLIENT
dsm.opt location	/opt/IBM/ISC/tsm/client/ba/bin
Storage Agent name	CL_HACMP03_STA
dsmsta.opt and devconfig.txt location	/opt/IBM/ISC/tsm/Storageagent/bin
Storage Agent high level address	admcnt01
Storage Agent low level address	1504
LAN-free communication method	Tcpip

Archived

Table 11-3 Server information

Server information	
Servername	TSMSRV04
High level address	atlantic
Low level address	1500
Server password for server-to-server communication	*password*

Our Storage Area Network devices are listed inTable 11-4.

Table 11-4 Storage Area Network devices

SAN devices	
Disk	IBM DS4500 Disk Storage Subsystem
Library	IBM LTO 3583 Tape Library
Tape drives	3580 Ultrium 1
Tape drive device name	drlto_1: /dev/rmt2 drlto_2: /dev/rmt3

11.2.1 Lab setup

We use the lab already set up for clustered client testing in Chapter 10, “AIX and HACMP with IBM Tivoli Storage Manager Client” on page 527.

Once the installation and configuration of Tivoli Storage Manager Storage Agent has finished, we need to modify the existing clients’ configuration to make them use the LAN-free backup.

11.3 Installation

We will install the AIX Storage Agent V5.3 for LAN-free backup components on both nodes of the HACMP cluster. This installation will be a standard installation, following the product’s Storage Agent User’s Guide.

An appropriate tape device driver is also required to be installed.

For the above tasks, Chapter 9, “AIX and HACMP with IBM Tivoli Storage Manager Server” on page 451 can also be used as a reference.

Archived

At this point, our team has already installed the Tivoli Storage Manager Server and Tivoli Storage Manager Client, both configured for high availability.

1. We review the latest Storage Agent readme file and the *User's Guide*.
2. Using the AIX command **smitty installp**, we install the filesets for the Tivoli Storage Manager Storage Agent and tape subsystem device driver.

11.4 Configuration

We are using storage and network resources already managed by the cluster, so we configure the clustered Tivoli Storage Manager components relying on that resources, and local components on local disk and persistent addresses. We have configured and verified the communication paths between the client nodes and the server also. Then we set up start and stop scripts for Storage Agent and add it to the HACMP resource group configuration. After that we modify clients configuration for having it to use LAN-free.

11.4.1 Configure tape storage subsystems

Here we will configure external tape storage resources for Tivoli Storage Manager server. We will not go into fine detail regarding hardware related tasks, we will just mention the higher level topics.

1. We first verify server adapter cards, storage and tape subsystems, and SAN switches for planned firmware levels or update as needed.
2. Then we connect fibre connections from server adapters and tape storage subsystems to SAN switches.
3. We configure zoning as planned to give server access to tape subsystems.
4. Then we run **cfgmgr** on both nodes to configure the tape storage subsystem.
5. Tape storage devices are now available on both servers; see **lsdev** output in Example 11-1.

Example 11-1 lsdev command for tape subsystems

```
azov:/# lsdev -Cctape
rmt0 Available 1Z-08-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt1 Available 1D-08-02 IBM 3580 Ultrium Tape Drive (FCP)
smc0 Available 1Z-08-02 IBM 3582 Library Medium Changer (FCP)

kanaga:/# lsdev -Cctape
rmt1 Available 1Z-08-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt0 Available 1D-08-02 IBM 3580 Ultrium Tape Drive (FCP)
smc0 Available 1Z-08-02 IBM 3582 Library Medium Changer (FCP)
```

Archived

11.4.2 Configure resources and resource groups

The storage resource needs for Tivoli Storage Manager Storage Agents are a directory containing logs and configuration files, so we create the directory /opt/IBM/ISC/tsm/Storageagent/bin within the filesystem /opt/IBM/ISC which belongs to the resource group named *rg_admcnt01*. To the same resource group belongs the admcnt01 service address that are we going to use for Storage Agent communication with the server.

Once we have set up Storage Agent related start and stop scripts, they will be added to the main ISC start and stop scripts.

11.4.3 Tivoli Storage Manager Storage Agent configuration

Now we configure Tivoli Storage Manager server, server objects, Storage Agent instances, and Storage Agent tape paths for the LAN-free environment.

In Tivoli Storage Manager server, Storage Agent objects are to be configured as *Other Servers*.

> **Attention:** Take care when changing server settings like server name, address, port, and password in a currently running server, because it can impact whole Tivoli Storage Manager environment operations.

Set Tivoli Storage Manager server password

In order to enable the required server to server connection, a server password has to be set. If a server password has not been set yet, we need to do it now.

> **Note:** Check for server name, server password set, server address, and server port with the `query status` command on server administrator command line and use current values if applicable.

1. We select **Enterprise Administration under the administration center**.
2. Then we select our targeted Tivoli Storage Manager server, the **Server-to-Server Communication setting** wizard and click **Go** (Figure 11-1).

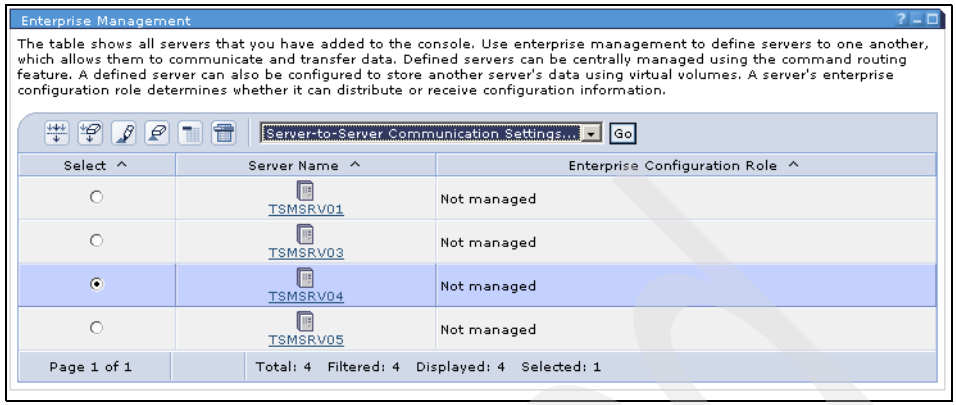

Figure 11-1 Start Server to Server Communication wizard

3. Then we make note of the server name and type in the fields for Server Password; Verify Password; TCP/IP Address; and TCP/IP Port for the server, if not yet set, and click **OK** (Figure 11-2).

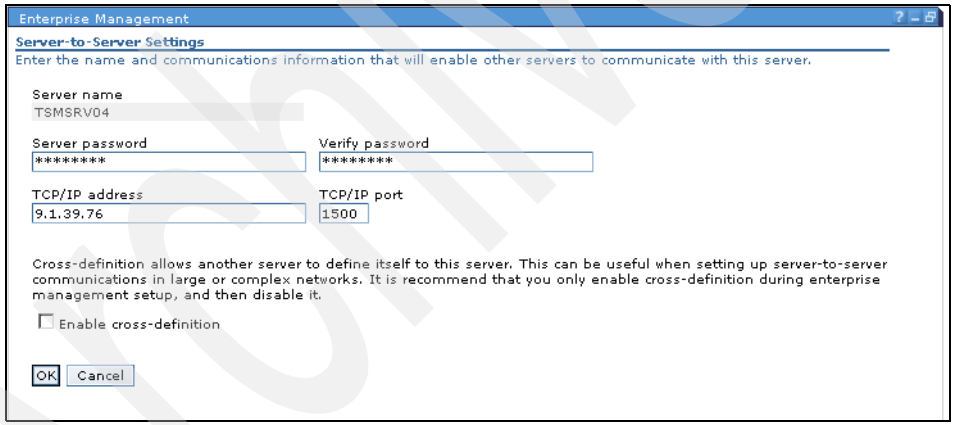

Figure 11-2 Setting Tivoli Storage Manager server password and address

From the administrator command line, the above tasks can be accomplished with these server commands (Example 11-2).

Example 11-2 Set server settings from command line

```
TSMSRV03> set serverpassword password
TSMSRV03> set serverhladdress atlantic
TSMSRV03> set serverlladdress 1500
```

Server object definitions for Storage Agents

Storage agents are configured to Tivoli Storage Manager server as other servers. Using data from Table 11-2 on page 559, we begin defining our Storage Agents on the targeted Tivoli Storage Manager server, by using the ISC administration interface.

1. We select **Enterprise Administration under the administration center**.
2. Then we select our targeted Tivoli Storage Manager server, **View Enterprise Properties** and click **Go** (Figure 11-3).

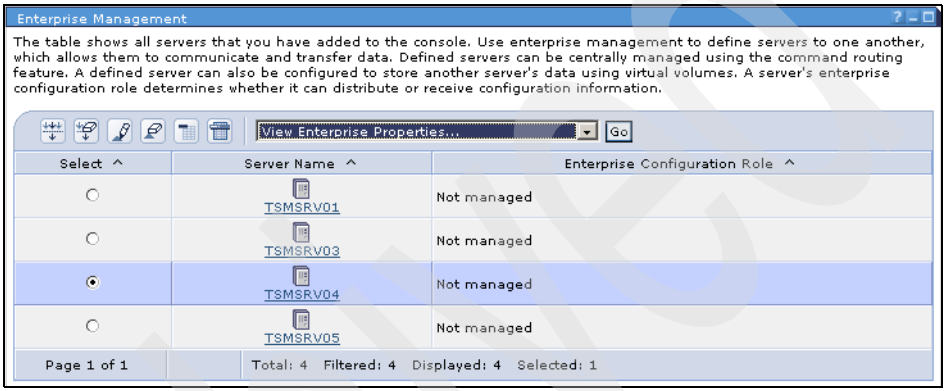

Figure 11-3 Select targeted server and View Enterprise Properties

3. We open the **Servers** section, choose **Define Server**, and click **Go** (Figure 11-4).

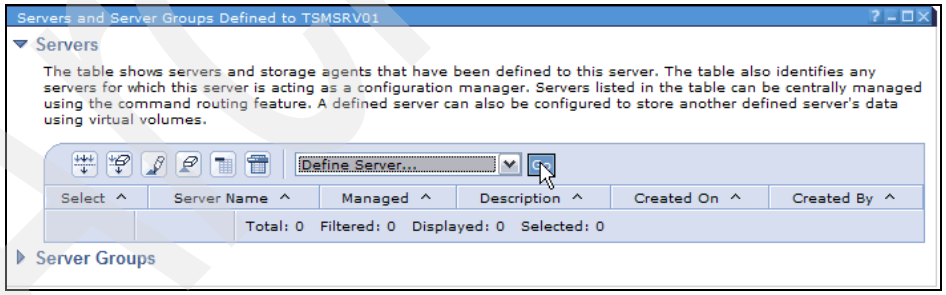

Figure 11-4 Define Server chose under Servers section

4. Then we click **Next** on the **Welcome** panel, and fill in the **General** panel fields with Tivoli Storage Manager Storage Agent name, password, description, and click **Next** (Figure 11-5).

Archived

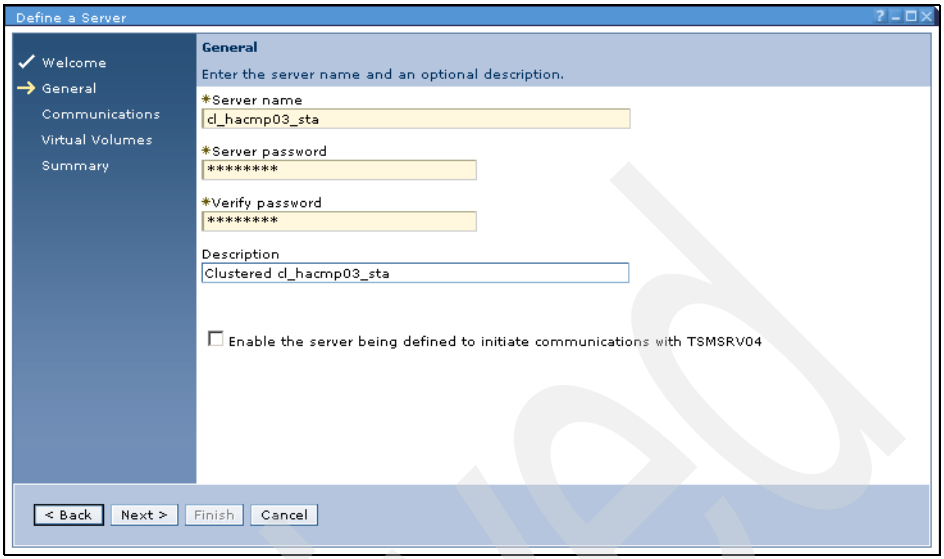

Figure 11-5 Entering Storage Agent name, password, and description

5. On the Communication panel we type in the fields for TCP/IP address (can be iplabel or dotted ip address) and TCP/IP port (Figure 11-6).

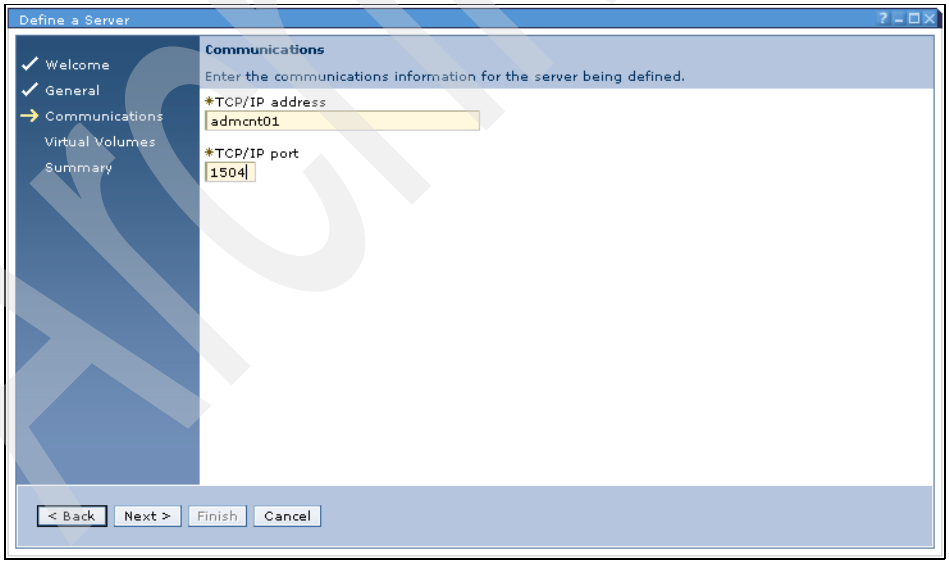

Figure 11-6 Insert communication data

Archived

6. We click **Next** on the **Virtual Volumes** panel (Figure 11-7).

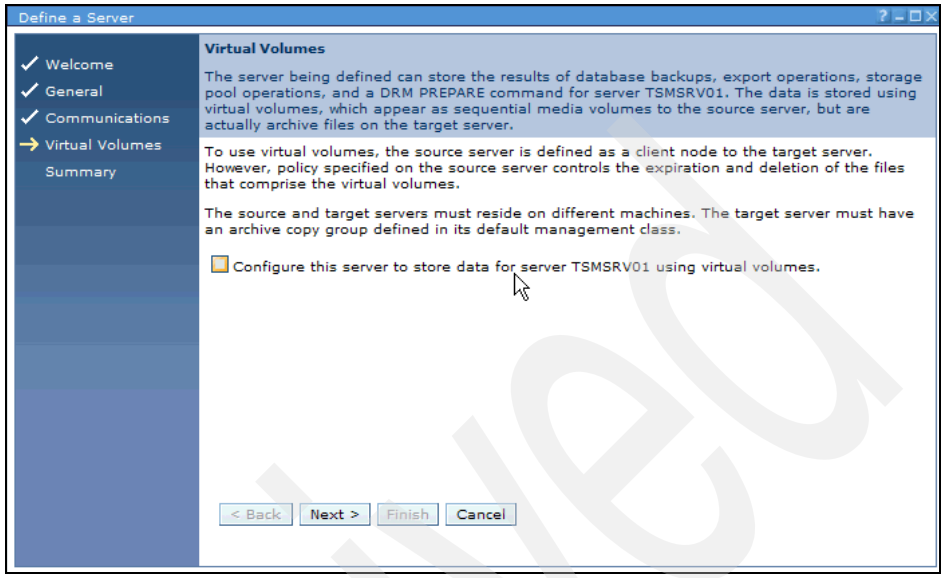

Figure 11-7 Click Next on Virtual Volumes panel

7. Then we verify entered data and click **Finish** on the **Summary** panel (Figure 11-8).

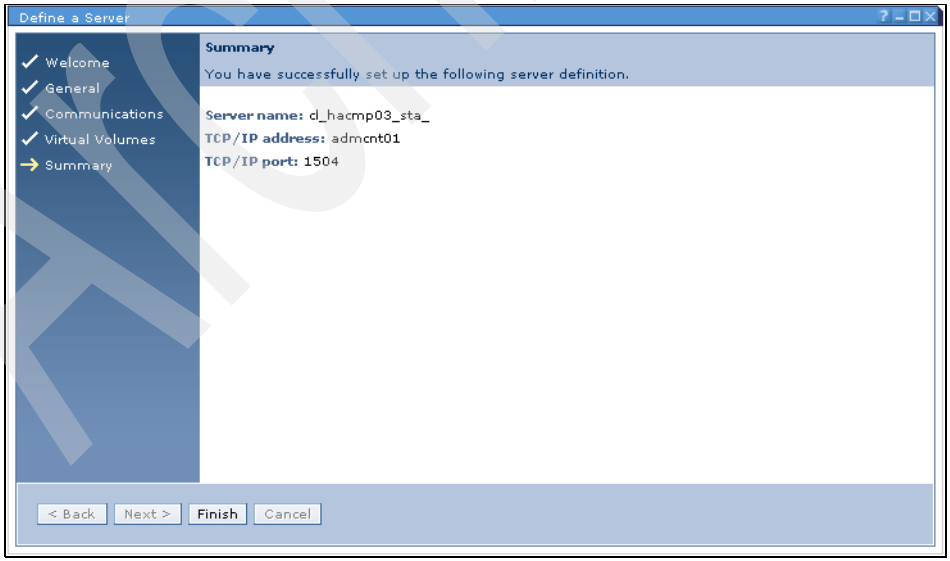

Figure 11-8 Summary panel

From the administrator command line, the above tasks can be accomplished with the server command shown in Example 11-3).

Example 11-3 Define server using the command line

```
TSMSRV03> define server cl_hacmp03_sta serverpassword=password
hladdress=admcnt01 lladdress=1504
```

Storage agent drive paths

Drive path definitions are needed in order to enable Storage Agents accessing the tape drives through the corresponding operating system device.

Using data fro mTable 11-4 on page 560, we configure all our Storage Agents device paths on the targeted Tivoli Storage Manager Server, by the ISC administration interface:

1. We select **Storage Devices** under the administration center.
2. Then, on the **Libraries for All Servers** panel, we select our targeted library for our targeted server, **Modify Library**, and click **Go**.
3. On the **Library_name Properties (Server_name)** panel, we check the boxes for **Share this library** and **Perform a target reset [...]** if not yet checked and click **Apply** (Figure 11-9).

Archived

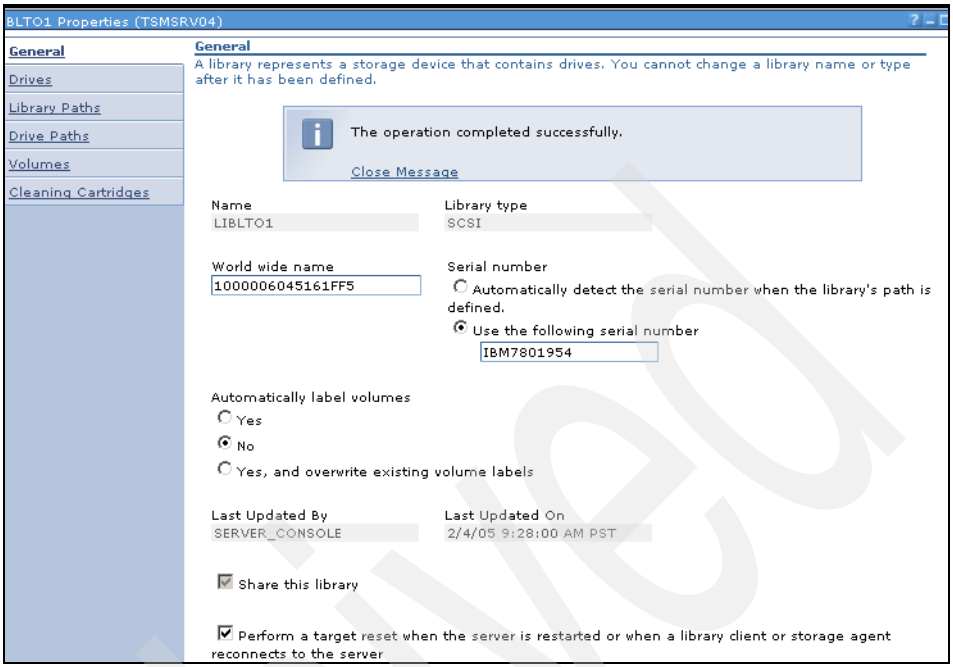

Figure 11-9 Share the library and set resetdrives to yes

4. Then we click **Drive Paths**, select **Add Path**, and click **Go**.
5. On the **Add Drive Path** sub-panel, we type in the device name, select drive, select library, and click **OK** (Figure 11-10).

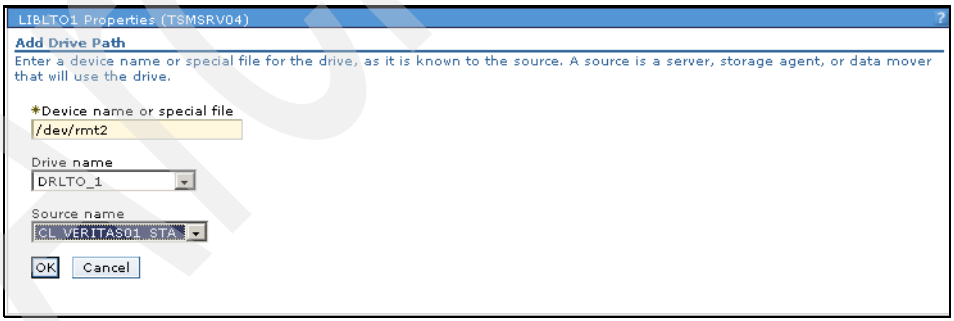

Figure 11-10 Define drive path panel

6. We repeat the add path steps for all the drives for each Storage Agent.

From the administrator command line, the above tasks can be accomplished with the server command shown in Example 11-4.

Archived

Example 11-4 Define paths using the command line

```
TSMSRV03> upd library liblto1 shared=yes resetdrives=yes
TSMSRV03> define path cl_hacmp03_sta drlto_1 srctype=server destype=drive
library=liblto1 device=/dev/rmt2
TSMSRV03> define path cl_hacmp03_sta drlto_2 srctype=server destype=drive
library=liblto1 device=/dev/rmt3
```

Storage Agent instances configuration

Here we configure three different Storage Agent instances:

1. We set up the three dsmsta.opt configuration files, in the three different instance directories, with planned TCP/IP ports and devconfig file path as for Table 11-2 on page 559; a local dsmsta.opt is shown in Example 11-5.

Example 11-5 Local instance dsmsta.opt

```
COMMmethod TCPIP
TCPPort 1502
DEVCONFIG /usr/tivoli/tsm/StorageAgent/bin/devconfig.txt
```

2. Next, we run the **/usr/tivoli/tsm/StorageAgent/bin/dsmsta setstorageserver command** to populate the devconfig.txt and dsmsta.opt files for local instances, using information from Table 11-3 on page 560, as shown in Example 11-6.

Example 11-6 The dsmsta setstorageserver command

```
# cd /usr/tivoli/tsm/StorageAgent/bin
# dsmsta setstorageserver myname=kanaga_sta mypassword=password
myhladdress=kanaga servername=tsmsrv04 serverpassword=password
hladdress=atlantic lladdress=1500
```

3. Now we do the clustered instance setup, using appropriate parameters and running environment, as shown in Example 11-7.

Example 11-7 The dsmsta setstorageserver command for clustered Storage Agent

```
# export DSMSERV_CONFIG=/opt/IBM/ISC/tsm/StorageAgent/bin/dsmsta.opt
# export DSMSERV_DIR=/usr/tivoli/tsm/StorageAgent/bin
# cd /opt/IBM/ISC/tsm/StorageAgent/bin
# dsmsta setstorageserver myname=cl_hacmp03_sta mypassword=password
myhladdress=admcnt01 servername=tsmsrv04 serverpassword=password
hladdress=atlantic lladdress=1500
```

4. We then review the results of running this command, which populates the devconfig.txt file, as shown in Example 11-8.

Archived

Example 11-8 The devconfig.txt file

```
SET STANAME KANAGA_STA
SET STAPASSWORD 2153327d37e22d1a357e47fcdf82bcfaf0
SET STAHLADDRESS KANAGA
DEFINE SERVER TSMSRV01 HLADDRESS=ATLANTIC LLADDRESS=1500
SERVERPA=21911a57cfe832900b9c6f258aa0926124
```

5. Next, we review the results of this update on the dsmsta.opt file. We see that the last line was updated with the servername, as seen in Example 11-9.

Example 11-9 Clustered Storage Agent devconfig.txt

```
COMMmethod TCPIP
TCPPort 1504
DEVCONFIG /opt/IBM/ISC/tsm/StorageAgent/bin/devconfig.txt
SERVERNAME TSMSRV04
```

Note: If **`dsmsta setstorageserver`** is run more than once, devconfig.txt and dsmsta.opt files have to be cleared up from duplicate entries.

Modifying client configuration

We then convert a LAN only Tivoli Storage Manager Client in a LAN-free enabled one and make it using the LAN-free backup method, by adding an appropriate stanza to our /usr/tivoli/tsm/client/ba/bin/dsm.sys file for the LAN-free connection for the clustered client, as shown in Example 11-10 .

Example 11-10 The /usr/tivoli/tsm/client/ba/bin/dsm.sys file

```
* Server stanza for the HACMP highly available client CL_HACMP03_CLIENT (AIX)
* this will be a client which uses the lan-free StorageAgent

SErvername                      tsmsrv04_san
   nodename                     cl_hacmp03_client
   COMMMethod                   TCPip
   TCPPort                      1500
   TCPServeraddr                atlantic
   TCPClientaddress             admcnt01

   TXNBytelimit                 256000
   resourceutilization          5
   enablelanfree                yes
   lanfreecommmethod            tcpip
   lanfreetcpport               1504
   lanfreetcpserveraddress      admcnt01

   passwordaccess               generate
   passworddir                  /opt/IBM/ISC/tsm/client/ba/bin
```

```
managedservices              schedule webclient
schedmode                    prompt
schedlogname                 /opt/IBM/ISC/tsm/client/ba/bin/dsmsched.log
errorlogname                 /opt/IBM/ISC/tsm/client/ba/bin/dsmerror.log
ERRORLOGRETENTION            7

clusternode                  yes
domain                       /opt/IBM/ISC
include                      /opt/IBM/ISC/.../* MC_SAN
```

The clients have to be restarted after dsm.sys has been modified, to have them using LAN-free operation.

Note: We also set a wider TXNBytelimit and a resourceutilization set at 5 to obtain two LAN-free backup sessions, and an include statement pointing to a management class whose B/A copy group uses a tape storage pool.

DATAREADPATH and DATAWRITEPATH node attributes

The node attributes DATAREADPATH and DATAWRITEPATH determine the restriction placed on the node. You can restrict a node to use only the LAN-free path on backup and archive (DATAWRITEPATH), and the LAN path on restore and retrieve (DATAREADPATH). Note that such a restriction can fail a backup or archive operation if the LAN-free path is unavailable. Consult the *Administrator's Reference* for more information regarding these attributes.

Start scripts with an AIX Tivoli Storage Manager server.

Locale Storage Agent instances are started at boot time by an *inittab* entry, added automatically at Storage Agent code installation time, which executes the default *rc.tsmstgagnt* placed in the default directory.

For the clustered instance we set up a start script merging the Tivoli Storage Manager server supplied sample start script with *rc.tsmstgagnt*.

We chose to use the standard HACMP application scripts directory for start and stop scripts:

1. We create the **/usr/es/sbin/cluster/local/tsmsta** directory on both nodes.
2. Then from **/usr/tivoli/tsm/server/bin/** we copy the two sample scripts to our scripts directory on the first node (Example 11-11).

Example 11-11 Example scripts copied to /usr/es/sbin/cluster/local/tsmsrv, first node

```
cd /usr/tivoli/tsm/server/bin/
cp startserver /usr/es/sbin/cluster/local/tsmsta/startcl_hacmp03_sta.sh
cp stopserver /usr/es/sbin/cluster/local/tsmsta/stopcl_hacmp03_sta.sh
```

3. Now we adapt the start script to set the correct running environment for a Storage Agent running in a directory different from the default and launch it as for the original rc.tsmstgagnt.

   Here is our script in Example 11-12.

Example 11-12 Our Storage Agent with AIX server startup script

```
#!/bin/ksh
##########################################################################
#                                                                        #
# Shell script to start a StorageAgent.                                  #
#                                                                        #
# Originated from the sample TSM server start script                     #
#                                                                        #
##########################################################################
echo “Starting Storage Agent now...”
# Start up TSM storage agent
##########################################################################
# Set the correct configuration
# dsmsta honors same variables as dsmserv does
export DSMSERV_CONFIG=/opt/IBM/ISC/tsm/StorageAgent/bin/dsmsta.opt
export DSMSERV_DIR=/usr/tivoli/tsm/StorageAgent/bin

# Get the language correct....
export LANG=en_US

# max out size of data area
ulimit -d unlimited

#OK, now fire-up the storage agent in quiet mode.
print “$(date ‘+%D %T’) Starting Tivoli Storage Manager storage agent”

cd /opt/IBM/ISC/tsm/StorageAgent/bin
$DSMSERV_DIR/dsmsta quiet &
```

4. We include the Storage Agent start script in the application server start script, after the ISC launch and before the Tivoli Storage Manager client scheduler start (Example 11-13).

Example 11-13 Application server start script

```
#!/bin/ksh

# Startup the ISC_Portal tu make the TSM Admin Center available
/opt/IBM/ISC/PortalServer/bin/startISC.sh ISC_Portal iscadmin iscadmin

# Startup the TSM Storage Agent
/usr/es/sbin/cluster/local/tsmsta/startcl_hacmp03_sta.sh
```

```
# Startup the TSM Client Acceptor Daemon
/usr/es/sbin/cluster/local/tsmcli/StartClusterTsmClient.sh
```

Then we continue with “Stop script” on page 577.

Start script with NON AIX Tivoli Storage Manager server

Locale Storage Agent instances are started at boot time by an *inittab* entry, added automatically at Storage Agent code installation time, which execute the default *rc.tsmstgagnt* placed in the default directory.

For the clustered instance, we set up a start script merging the Tivoli Storage Manager Server supplied sample scripts, the *rc.tsmstgagnt*, and inserted a query to the Tivoli Storage Manager Server database to find any tape resources that might have been left allocated to the clustered Storage Agent after a takeover. This is done, not for the allocation issue that is resolved automatically by the server at Storage Agent restart time, but for solving SCSI reserve issues that are still present when working with non-AIX servers. If it finds that condition, it issues a SCSI reset against the involved devices.

We chose to use the standard HACMP application scripts directory for start and stop scripts:

1. At first we create the **/usr/es/sbin/cluster/local/tsmsta** directory on both nodes.
2. Then from **/usr/tivoli/tsm/server/bin/** we copy the two sample scripts and their referenced executables to our script directory to the first node (Example 11-14).

Example 11-14 Copy from /usr/tivoli/tsm/server/bin to /usr/es/sbin/cluster/local/tsmsrv

```
cd /usr/tivoli/tsm/server/bin/
cp startserver /usr/es/sbin/cluster/local/tsmsta/startcl_hacmp03_sta.sh
cp stopserver /usr/es/sbin/cluster/local/tsmsta/stopcl_hacmp03_sta.sh
cp checkdev /usr/es/sbin/cluster/local/tsmsta/
cp opendev /usr/es/sbin/cluster/local/tsmsta/
cp fcreset /usr/es/sbin/cluster/local/tsmsta/
cp fctest /usr/es/sbin/cluster/local/tsmsta/
cp scsireset /usr/es/sbin/cluster/local/tsmsta/
cp scsitest /usr/es/sbin/cluster/local/tsmsta/
cp verdev /usr/es/sbin/cluster/local/tsmsta/
cp verfcdev /usr/es/sbin/cluster/local/tsmsta/
```

3. Now we adapt the start script to our environment, and use the script operator we defined for server automated operation:

a. At first we insert an SQL query to the Tivoli Storage Manager Server database that resolves the AIX device name for any drive allocated to the instance that we are starting here.

b. Then we use the discovered device names with the original provided functions.

c. We left commented out the test for all devices available.

d. At the end we set the correct running environment for a Storage Agent running in a directory different from the default and launch it as for the original rc.tsmstgagnt.

Here our script in Example 11-15.

Example 11-15 Our Storage Agent with non-AIX server startup script

```
#!/bin/ksh
##########################################################################
#                                                                        #
# Shell script to start a StorageAgent, making sure required offline storage #
# devices are available.                                                 #
#                                                                        #
# Please note commentary below indicating the places where this shell script #
# may need to be modified in order to tailor it for your environment.    #
#                                                                        #
# Originated from the TSM server sample start script                     #
#                                                                        #
##########################################################################
# Get file name of shell script
scrname=${0##*/}
# Get path to directory where shell script was found
bindir=${0%/$scrname}
#
# Define function to verify that offline storage device is available (SCSI)
VerifyDevice ()
{
        $bindir/verdev $1 &
        device[i]=$1
        process[i]=$!
        i=i+1
}
#
#
# Define function to verify that offline storage device is available (FC)
VerifyFCDevice ()
{
        $bindir/verfcdev $1 &
        device[i]=$1
        process[i]=$!
        i=i+1
```

```
}
#
# Turn on ksh job monitor mode
set -m
#
echo “Verifying that offline storage devices are available...”
integer i=0
##############################################################################
#                                                                            #
# - Setup an appropriate administrator for use instead of admin.             #
#                                                                            #
# - Insert your Storage Agent server_name as searching value for             #
#   ALLOCATED_TO and SOURCE_NAME in the SQL query.                           #
#                                                                            #
# - Use VerifyDevice or VerifyDevice in the loop below depending of the      #
#   type of connection your tape storage subsystems is using.                #
#                                                                            #
#   VerifyDevice is for SCSI-attached devices                                #
#   VerifyFCDevice is for FC-attached devices                                #
##############################################################################
# Find out if this Storage Agent instance has left any tape drive reserved in
# its previous life.
WORKDIR=/tmp
TSM_ADMIN_CMD=”dsmadmc -quiet -se=tsmsrv04_admin -id=script_operator
-pass=password”

$TSM_ADMIN_CMD -outfile=$WORKDIR/DeviceQuery.out “select DEVICE from PATHS
where DESTINATION_NAME in ( select DRIVE_NAME from DRIVES where
ALLOCATED_TO=’CL_HACMP03_STA’ and SOURCE_NAME=’CL_HACMP03_STA’)” > /dev/null
 if [ “$?” = “0” ]
  then
   echo “Tape drives have been left allocated to this instance, most likely on
a server that has died so now we need to reset them.”
   RMTS_TO_RESET=‘cat $WORKDIR/DeviceQuery.out|egrep “/dev/rmt”|sed -e
“s/\/dev\///g”‘
   echo $RMT_TO_RESET
   for RMT in $RMTS_TO_RESET
    do
     # Change verify function type below to “VerifyDevice” or “ VerifyFCDevice”
     # depending of your devtype
     VerifyFCDevice $RMT
    done
  else
   echo “No tape drives have been left allocated to this instance”
 fi
# Remove tmp work file
 if [ -f $WORKDIR/DeviceQuery.out ]
  then
   rm $WORKDIR/DeviceQuery.out
```

```
 fi

#
# Wait for all VerifyDevice processes to complete
#
wait
# Check return codes from all VerifyDevice (verdev/verfcdev) processes
integer allrc=0
tty=$(tty)
if [ $? != 0 ]
then tty=/dev/null
fi
jobs -ln | tee $tty | awk -v "encl=Done()" '{print $3,
substr($4,length(encl),length($4)-length(encl))}' | while read jobproc rc
do
  if [ -z "$rc" ]
  then rc=0
  fi
  i=0
  while (( i < ${#process[*]} ))
  do
    if [ ${process[i]} = $jobproc ] ; then break ; fi
    i=i+1
  done
  if (( i >= ${#process[*]} ))
  then
    echo "Process $jobproc not found in array!"
    exit 99
  fi
  if [ $rc != 0 ]
  then
    echo "Attempt to make offline storage device ${device[i]} available ended
with return code $rc!"
    allrc=$rc
  fi
done
##############################################################################
#                                                                            #
# Comment the following three lines if you do not want the start-up of the STA#
# server to fail if all of the devices do not become available.              #
#                                                                            #
##############################################################################
#if (( allrc ))
#then exit $allrc
#fi

echo "Starting Storage Agent now..."
# Start up TSM storage agent
##############################################################################
```

```
# Set the correct configuration
# dsmsta honors same variables as dsmserv does
export DSMSERV_CONFIG=/opt/IBM/ISC/tsm/StorageAgent/bin/dsmsta.opt
export DSMSERV_DIR=/usr/tivoli/tsm/StorageAgent/bin

# Get the language correct....
export LANG=en_US

# max out size of data area
ulimit -d unlimited

#OK, now fire-up the storage agent in quiet mode.
print "$(date '+%D %T') Starting Tivoli Storage Manager storage agent"

cd /opt/IBM/ISC/tsm/StorageAgent/bin
$DSMSERV_DIR/dsmsta quiet &
```

4. We include the Storage Agent start scripts in the application server start script, after the ISC launch and before the Tivoli Storage Manager Client scheduler start (Example 11-16).

Example 11-16 Application server start script

```
#!/bin/ksh

# Startup the ISC_Portal tu make the TSM Admin Center available
/opt/IBM/ISC/PortalServer/bin/startISC.sh ISC_Portal iscadmin iscadmin

# Startup the TSM Storage Agent
/usr/es/sbin/cluster/local/tsmsta/startcl_hacmp03_sta.sh

# Startup the TSM Client Acceptor Daemon
/usr/es/sbin/cluster/local/tsmcli/StartClusterTsmClient.sh
```

Stop script

We chose to use the standard HACMP application scripts directory for start and stop scripts.

1. We use the Tivoli Storage Manager Server code provided sample stop script as for "Start and stop scripts setup" on page 490, having it pointing to a server stanza in dsm.sys which provides connection to our storage server instance, as shown in Example 11-17.

Example 11-17 Storage agent stanza in dsm.sys

```
* Server stanza for local storagent admin connection purpose
SErvername              cl_hacmp03_sta
   COMMMethod           TCPip
```

```
TCPPort              1504
TCPServeraddress     admcnt01
ERRORLOGRETENTION    7
ERRORLOGname         /usr/tivoli/tsm/client/ba/bin/dsmerror.log
```

2. Then the Storage Agent stop script is included in the application server stop script, which shows an inverted order of execution (Example 11-18).

Example 11-18 Application server stop script

```
#!/bin/ksh

# Stop the TSM Client Acceptor Daemon
/usr/es/sbin/cluster/local/tsmcli/StopClusterTsmClient.sh

# Stop the TSM Storage Agent
/usr/es/sbin/cluster/local/tsmsta/stopcl_hacmp03_sta.sh

# Stop The Portal
/opt/IBM/ISC/PortalServer/bin/stopISC.sh ISC_Portal iscadmin iscadmin

# killing all AppServer related java processes left running

JAVAASPIDS=‘ps -ef | egrep “java|AppServer” | awk ‘{ print $2 }’‘

for PID in $JAVAASPIDS
 do
  kill  $PID
 done

exit 0
```

11.5 Testing the cluster

Here we start testing our LAN-free environment failure and recovery.

11.5.1 LAN-free client system failover while the client is backing up

Now we test recovery of a scheduled backup operation after a node crash, while two tapes are in use by the Storage Agent:

1. We verify that the cluster services are running with the **`lssrc -g cluster command`** on both nodes.
2. On the resource group secondary node, we use **`tail -f /tmp/hacmp.out`** to monitor cluster operation.

3. Then we schedule a client selective backup having the whole shared filesystems as object and wait for it to be started (Example 11-19).

Example 11-19 Client sessions starting

```
tsm: TSMSRV04>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name     Status
-------------------- -------------------- ------------- ------------- ---------
02/08/05   09:30:25  02/08/05   09:31:41  TEST_1        CL_HACMP03_C- Started
                                                        LIENT
```

4. We wait for volume opened messages on the server console (Example 11-20).

Example 11-20 Output volumes open messages

```
[...]
02/08/05   09:31:41      ANR0511I Session 183 opened output volume ABA927.
(SESSION: 183)
[...]
02/08/05   09:32:31      ANR0511I Session 189 opened output volume ABA928.
(SESSION: 189)
```

5. Then we check for data being written by the Storage Agent, querying it via command routing functionality using the **cl_hacmp03_sta:q se** command (Example 11-21).

Example 11-21 Client sessions transferring data to Storage Agent

```
ANR1687I Output for command ‘Q SE ‘ issued against server CL_HACMP03_STA
follows:

  Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- --------------------
     1 Tcp/Ip IdleW    1 S    1.3 K   1.8 K Serv- AIX-RS/- TSMSRV04
                                              er    6000
     2 Tcp/Ip IdleW    0 S   86.7 K     257 Serv- AIX-RS/- TSMSRV04
                                              er    6000
     4 Tcp/Ip IdleW    0 S   22.2 K  26.3 K Serv- AIX-RS/- TSMSRV04
                                              er    6000
   182 Tcp/Ip Run      0 S      732 496.2 M Node  AIX      CL_HACMP03_CLIENT
   183 Tcp/Ip Run      0 S    6.2 M   5.2 M Serv- AIX-RS/- TSMSRV04
                                              er    6000
   189 Tcp/Ip Run      0 S      630 447.3 M Node  AIX      CL_HACMP03_CLIENT
   190 Tcp/Ip Run      0 S    4.6 M   3.9 M Serv- AIX-RS/- TSMSRV04
                                              er    6000
```

Archived

Failure

Now we simulate a server failure:

1. Being sure that client LAN-free backup is running, we issue **halt -q** on the AIX server on which the backup is running; the **halt -q** command stops any activity immediately and powers off the server.
2. The server remains waiting for client and Storage Agent communication until idletimeout expires (the default is 15 minutes).

Recovery

Here we see how failure is managed:

1. The secondary cluster node takes over the resources and launches the application server start script.
2. At first, the clustered application (ISC portal) is restarted by the application server start script (Example 11-22).

Example 11-22 The ISC being restarted

```
ADMU0116I: Tool information is being logged in file
/opt/IBM/ISC/AppServer/logs/ISC_Portal/startServer.log
ADMU3100I: Reading configuration for server: ISC_Portal
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server ISC_Portal open for e-business; process id is 106846
```

3. Then the Storage Agent startup script is run and the Storage Agent is started (Example 11-23).

Example 11-23 The Tivoli Storage Manager Storage Agent is restarted

```
Starting Storage Agent now...
Starting Tivoli Storage Manager storage agent
```

4. Then the Tivoli Storage Manager server, accepting new connections from the restarted CL_HACMP03_STA Storage Agent, cancels the previous ones, and the Storage Agent gets I/O errors trying to access tape drives that it left reserved on the crashed AIX (Example 11-24).

Example 11-24 CL_HACMP03_STA reconnecting

```
ANR0408I Session 228 started for server CL_HACMP03_STA (AIX-RS/6000) (Tcp/Ip)
for storage agent. (SESSION: 228)
ANR0490I Canceling session 4 for node CL_HACMP03_STA (AIX-RS/6000) . (SESSION:
228)
ANR3605E Unable to communicate with storage agent.  (SESSION: 4)
ANR0490I Canceling session 5 for node CL_HACMP03_STA (AIX-RS/6000) . (SESSION:
228)
ANR0490I Canceling session 7 for node CL_HACMP03_STA (AIX-RS/6000) . (SESSION:
228)
```

```
ANR3605E Unable to communicate with storage agent. (SESSION: 7)
ANR0483W Session 4 for node CL_HACMP03_STA (AIX-RS/6000) terminated - forced by
administrator. (SESSION: 4)
ANR0483W Session 5 for node CL_HACMP03_STA (AIX-RS/6000) terminated - forced by
administrator. (SESSION: 5)
ANR0483W Session 7 for node CL_HACMP03_STA (AIX-RS/6000) terminated - forced by
administrator. (SESSION: 7)
ANR0408I Session 229 started for server CL_HACMP03_STA (AIX-RS/6000) (Tcp/Ip)
for library sharing. (SESSION: 229)
ANR0408I Session 230 started for server CL_HACMP03_STA (AIX-RS/6000) (Tcp/Ip)
for event logging. (SESSION: 230)
ANR0409I Session 229 ended for server CL_HACMP03_STA (AIX-RS/6000). (SESSION:
229)
ANR0408I Session 231 started for server CL_HACMP03_STA (AIX-RS/6000) (Tcp/Ip)
for storage agent. (SESSION: 231)
ANR0407I Session 234 started for administrator ADMIN (AIX) (Tcp/Ip
9.1.39.89(33738)). (SESSION: 234)
ANR0408I (Session: 230, Origin: CL_HACMP03_STA)  Session 2 started for server
TSMSRV04 (AIX-RS/6000) (Tcp/Ip) for library sharing.  (SESSION: 230)
[...]
ANR8779E Unable to open drive /dev/rmt3, error number=16. (SESSION: 229)
ANR8779E Unable to open drive /dev/rmt2, error number=16. (SESSION: 229)
```

5. Now the Tivoli Storage Manager server is aware of the reserve problem and resets the reserved tape drives (it can only be seen with a trace) (Example 11-25).

Example 11-25 Trace showing pvr at work with reset

```
[42][output.c][6153]: ANR8779E Unable to open drive /dev/rmt2, error
number=16.~
[42][pspvr.c][3004]: PvrCheckReserve called for /dev/rmt2.
[42][pspvr.c][3820]: getDevParent: odm_initialize successful.
[42][pspvr.c][3898]: getDevParent with rc=0.
[42][pspvr.c][3954]: getFcIdLun: odm_initialize successful.
[42][pspvr.c][4071]: getFcIdLun with rc=0.
[42][pspvr.c][3138]: SCIOLTUR - device is reserved.
[42][pspvr.c][3441]: PvrCheckReserve with rc=79.
[42][pvrmp.c][7990]: Reservation conflict for DRLTO_1 will be reset
[42][pspvr.c][3481]: PvrResetDev called for /dev/rmt2.
[42][pspvr.c][3820]: getDevParent: odm_initialize successful.
[42][pspvr.c][3898]: getDevParent with rc=0.
[42][pspvr.c][3954]: getFcIdLun: odm_initialize successful.
[42][pspvr.c][4071]: getFcIdLun with rc=0.
[42][pspvr.c][3575]: SCIOLRESET Device with scsi id 0x50700, lun
0x2000000000000 has been RESET.
```

6. And now tape volumes are dismounted (Example 11-26).

Example 11-26 Tape dismounted after SCSI reset

```
ANR8336I Verifying label of LTO volume ABA928 in drive DRLTO_1 (/dev/rmt2).
(SESSION: 15)
ANR8336I Verifying label of LTO volume ABA927 in drive DRLTO_2 (/dev/rmt3).
(SESSION: 20)
[...]
ANR8468I LTO volume ABA928 dismounted from drive DRLTO_1 (/dev/rmt2) in library
LIBLTO1. (SESSION: 15)
ANR8468I LTO volume ABA927 dismounted from drive DRLTO_2 (/dev/rmt3) in library
LIBLTO1. (SESSION: 20)
```

7. Once the Storage Agent start script completes, the CL_HACMP03_CLIENT scheduler start script is started too.
8. It searches for sessions to cancel (Example 11-27).

Example 11-27 Extract of console log showing session cancelling work

```
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_HACMP03_CLIENT'
(SESSION: 227)
[...]
ANR2017I Administrator SCRIPT_OPERATOR issued command: CANCEL SESSION 183
(SESSION: 234)
[...]
ANR2017I Administrator SCRIPT_OPERATOR issued command: CANCEL SESSION 189
(SESSION: 238)
[...]
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_HACMP03_CLIENT'
(SESSION: 240)
[...]
ANR2017I Administrator SCRIPT_OPERATOR issued command: CANCEL SESSION 183
(SESSION: 241)
[...]
ANR0483W Session 183 for node CL_HACMP03_CLIENT (AIX) terminated - forced by
administrator. (SESSION: 183)
[...]
ANR2017I Administrator SCRIPT_OPERATOR issued command: CANCEL SESSION 189
(SESSION: 242)
[...]
ANR0483W Session 189 for node CL_HACMP03_CLIENT (AIX) terminated - forced by
administrator. (SESSION: 189)
```

Archived

Note: Sessions with *_VOL_ACCESS not null increase the node mount point used number, preventing new sessions from the same node to obtain new mount points by the MAXNUMMP parameter. This session remains until commtimeout expires; refer to 10.7.3, “Client system failover while the client is backing up to tape with higher CommTimeOut” on page 543.

9. Once the session’s cancelling work finishes, the scheduler is restarted and the scheduled backup operation is restarted too (Example 11-28).

Example 11-28 The client schedule restarts

```
ANR0406I Session 244 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.89(33748)). (SESSION: 244)

tsm: TSMSRV04>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name     Status
-------------------- -------------------- ------------- ------------- ---------
02/08/05   09:30:25  02/08/05   09:31:41  TEST_1        CL_HACMP03_C- Restarted
                                                        LIENT
```

10. We can find messages in the actlog for backup operation restarting via SAN with the same tapes mounted to the Storage Agent and completing with a successful result (Example 11-29).

Example 11-29 Server log view of restarted restore operation

```
ANR0406I Session 244 started for node CL_HACMP03_CLIENT (AIX) (Tcp/Ip
9.1.39.89(33748)). (SESSION: 244)
[...]
ANR0408I Session 247 started for server CL_HACMP03_STA (AIX-RS/6000) (Tcp/Ip)
for library sharing.  (SESSION: 247)
[...]
ANR8337I LTO volume ABA928 mounted in drive DRLTO_1 (/dev/rmt2). (SESSION: 248)
ANR8337I (Session: 230, Origin: CL_HACMP03_STA)  LTO volume ABA928 mounted in
drive DRLTO_1 (/dev/rmt2).  (SESSION: 230)
ANR0511I Session 246 opened output volume ABA928. (SESSION: 246)
ANR0511I (Session: 230, Origin: CL_HACMP03_STA)  Session 13 opened output
volume ABA928. (SESSION: 230)
[...]
ANR8337I LTO volume ABA927 mounted in drive DRLTO_2 (/dev/rmt3). (SESSION: 255)
ANR8337I (Session: 237, Origin: CL_HACMP03_STA)  LTO volume ABA927 mounted in
drive DRLTO_2 (/dev/rmt3).  (SESSION: 237)
ANR0511I Session 253 opened output volume ABA927. (SESSION: 253)
ANR0511I (Session: 237, Origin: CL_HACMP03_STA)  Session 20 opened output
volume ABA928. (SESSION: 237)
```

```
[...]
ANE4971I (Session: 244, Node: CL_HACMP03_CLIENT)  LanFree data bytes:
1.57 GB (SESSION: 244)
[...]
ANR2507I Schedule TEST_1 for domain STANDARD started at 02/08/05 09:30:25 for
node CL_HACMP03_CLIENT complete successfully at 02/08/05 09:50:39. (SESSION:
244)
```

Result summary

We are able to have the HACMP cluster restarting an application with its backup environment up and running.

Tivoli Storage Manager server 5.3 or later for AIX is able to resolve SCSI reserve issues. A scheduled operation, still in its startup window, is restarted by the scheduler and obtains back the previous resources.

There is the opportunity of having a backup restarted even if, considering a database as an example, this can lead to a backup window breakthrough, thus affecting other backup operations.

We run this test, at first using command line initiated backups with the same result; the only difference is that the operation needs to be restarted manually.

11.5.2 LAN-free client system failover while the client is restoring

Now we test the ability to restart and complete a command line LAN-free restore operation, still over SAN, after a node crashed while two tapes are in use by the Storage Agent:

1. We verify that the cluster services are running with the **`lssrc -g cluster command`** on both nodes.
2. On the resource group secondary node, we use **`tail -f /tmp/hacmp.out`** to monitor cluster operation.
3. We launch a restore operation from the LAN-free enabled clustered node (Example 11-30).

Example 11-30 Client sessions starting

```
Node Name: CL_HACMP03_CLIENT
Session established with server TSMSRV04: AIX-RS/6000
  Server Version 5, Release 3, Level 0.0
  Server date/time: 02/15/05   13:24:20  Last access: 02/15/05   13:21:02

tsm> restore -subdir=yes "/opt/IBM/ISC/backups/*"
Restore function invoked.
```

```
ANS1899I ***** Examined     1,000 files *****
[...]
```

4. We wait for volumes to mount and see open messages on the server console (Example 11-31).

Example 11-31 Tape mount and open messages

```
ANR8337I LTO volume ABA927 mounted in drive DRLTO_2 (/dev/rmt3). (SESSION: 270)
ANR8337I (Session: 257, Origin: CL_HACMP03_STA)  LTO volume ABA927 mounted in
drive DRLTO_2 (/dev/rmt3). (SESSION: 257)
ANR0510I (Session: 257, Origin: CL_HACMP03_STA)  Session 16 opened input volume
ABA927. (SESSION: 257)
ANR0514I (Session: 257, Origin: CL_HACMP03_STA)  Session 16 closed volume
ABA927. (SESSION: 257)
ANR0514I Session 267 closed volume ABA927. (SESSION: 267)
ANR8337I LTO volume ABA928 mounted in drive DRLTO_1 (/dev/rmt2). (SESSION: 278)
ANR8337I (Session: 257, Origin: CL_HACMP03_STA)  LTO volume ABA928 mounted in
drive DRLTO_1 (/dev/rmt2). (SESSION: 257)
ANR0510I (Session: 257, Origin: CL_HACMP03_STA)  Session 16 opened input volume
ABA928. (SESSION: 257)
```

5. Then we check for data being read from the Storage Agent, querying it via command routing functionality using the **cl_hacmp03_sta:q se** command (Example 11-32).

Example 11-32 Checking for data being received by the Storage Agent

```
tsm: TSMSRV04>CL_HACMP03_STA:q se
ANR1699I Resolved CL_HACMP03_STA to 1 server(s) - issuing command Q SE  against
server(s).
ANR1687I Output for command ‘Q SE ‘ issued against server CL_HACMP03_STA
follows:

  Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- --------------------
     1 Tcp/Ip IdleW    0 S    6.1 K   7.0 K Serv- AIX-RS/- TSMSRV04
                                             er    6000
     4 Tcp/Ip IdleW    0 S   30.4 M  33.6 M Serv- AIX-RS/- TSMSRV04
                                             er    6000
    13 Tcp/Ip IdleW    0 S    8.8 K     257 Serv- AIX-RS/- TSMSRV04
                                             er    6000
    16 Tcp/Ip Run      0 S  477.1 M 142.0 K Node  AIX      CL_HACMP03_CLIENT
    17 Tcp/Ip Run      0 S    5.3 M   6.9 M Serv- AIX-RS/- TSMSRV04
                                             er    6000
```

Failure

Now we simulate a server crash:

1. Being sure that client LAN-free restore is running, we issue **halt -q** on the AIX server on which the backup is running; the **halt -q** command stops any activity immediately and powers off the server.

Recovery

Here we can see how failure recovery is managed:

1. The secondary cluster node takes over the resources and launches the application server start script.
2. At first, the clustered application (ISC portal) is restarted by the application server start script (Example 11-33).

Example 11-33 ISC restarting

```
ADMU0116I: Tool information is being logged in file
/opt/IBM/ISC/AppServer/logs/ISC_Portal/startServer.log
ADMU3100I: Reading configuration for server: ISC_Portal
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server ISC_Portal open for e-business; process id is 319994
```

3. Then the Storage Agent startup script is run and the Storage Agent is started (Example 11-34).

Example 11-34 Storage agent restarting.

```
Starting Storage Agent now...
Starting Tivoli Storage Manager storage agent
```

4. Then the server accepts new connections from the CL_HACMP03_STA agent and cancels the previous ones. At the same time, it unmounts the volume that was previously allocated to CL_HACMP03_STA, being aware that it has been restarted (Example 11-35).

Example 11-35 Tivoli Storage Manager server accepts new sessions, unloads tapes

ANR0408I Session 290 started for server CL_HACMP03_STA (AIX-RS/6000) (Tcp/Ip) for storage agent. (SESSION: 290)

ANR0490I Canceling session 229 for node CL_HACMP03_STA (AIX-RS/6000) . (SESSION: 290)

ANR3605E Unable to communicate with storage agent. (SESSION: 229)

ANR0490I Canceling session 232 for node CL_HACMP03_STA (AIX-RS/6000) . (SESSION: 290)

ANR3605E Unable to communicate with storage agent. (SESSION: 232)

ANR0490I Canceling session 257 for node CL_HACMP03_STA (AIX-RS/6000) . (SESSION: 290)

ANR0483W Session 229 for node CL_HACMP03_STA (AIX-RS/6000) terminated - forced by administrator. (SESSION: 229)

[...]

ANR8920I (Session: 291, Origin: CL_HACMP03_STA) Initialization and recovery has ended for shared library LIBLTO1. (SESSION: 291)

[...]

ANR8779E Unable to open drive /dev/rmt3, error number=16. (SESSION: 292)

[...]

ANR8336I Verifying label of LTO volume ABA928 in drive DRLTO_1 (/dev/rmt2). (SESSION: 278)

[...]

ANR8468I LTO volume ABA928 dismounted from drive DRLTO_1 (/dev/rmt2) in library LIBLTO1. (SESSION: 278)

---

5. Once the Storage Agent scripts completes, the clustered scheduler start script is started too.
6. It searches for previous sessions to cancel, issues cancel session commands, and in this test, a cancel command needs to be issued twice to cancel session 267 (Example 11-36).

Example 11-36 Extract of console log showing session cancelling work

```
ANR2017I Administrator SCRIPT_OPERATOR issued command: CANCEL SESSION 265
(SESSION: 297)
ANR0490I Canceling session 267 for node CL_HACMP03_CLIENT (AIX) . (SESSION:
298)
ANR0483W Session 265 for node CL_HACMP03_CLIENT (AIX) terminated - forced by
administrator. (SESSION: 265)
[...]
ANR2017I Administrator SCRIPT_OPERATOR issued command: CANCEL SESSION 267
(SESSION: 298)
ANR0490I Canceling session 267 for node CL_HACMP03_CLIENT (AIX) . (SESSION:
298)
```

```
[...]
ANR2017I Administrator SCRIPT_OPERATOR issued command: CANCEL SESSION 267
(SESSION: 301)
ANR0490I Canceling session 267 for node CL_HACMP03_CLIENT (AIX) . (SESSION:
301)
ANR0483W Session 267 for node CL_HACMP03_CLIENT (AIX) terminated - forced by
administrator. (SESSION: 267)
```

7. Once the session's cancelling work finishes, the scheduler is restarted.
8. We re-issue the restore command with the **replace=all** option (Example 11-37).

Example 11-37 The client restore re issued.

```
tsm> restore -subdir=yes -replace=all "/opt/IBM/ISC/backups/*"
Restore function invoked.

ANS1899I ***** Examined     1,000 files *****
ANS1899I ***** Examined     2,000 files *****
ANS1899I ***** Examined     3,000 files *****
ANS1899I ***** Examined     4,000 files *****
ANS1899I ***** Examined     5,000 files *****
[...]
```

9. We can find messages in the actlog (Example 11-38), and on the client (Example 11-39) for a restore operation restarting via SAN and completing with a successful result.

Example 11-38 Server log of new restore operation

```
ANR8337I (Session: 291, Origin: CL_HACMP03_STA)  LTO volume ABA927 mounted in
drive DRLTO_2 (/dev/rmt3). (SESSION: 291)
ANR0510I (Session: 291, Origin: CL_HACMP03_STA)  Session 10 opened input volume
ABA927. (SESSION: 291)
ANR0514I (Session: 291, Origin: CL_HACMP03_STA)  Session 10 closed volume
ABA927. (SESSION: 291)
ANR0514I Session 308 closed volume ABA927. (SESSION: 308)
[...]
ANR8337I LTO volume ABA928 mounted in drive DRLTO_1 (/dev/rmt2). (SESSION: 319)
ANR8337I (Session: 291, Origin: CL_HACMP03_STA)  LTO volume ABA928 mounted in
drive DRLTO_1 (/dev/rmt2). (SESSION: 291)
ANR0510I (Session: 291, Origin: CL_HACMP03_STA)  Session 10 opened input volume
ABA928. (SESSION: 291)
ANR0514I (Session: 291, Origin: CL_HACMP03_STA)  Session 10 closed volume
ABA928. (SESSION: 291)
[...]
ANE4955I (Session: 304, Node: CL_HACMP03_CLIENT)  Total number of objects
restored:    20,338 (SESSION: 304)
```

```
ANE4959I (Session: 304, Node: CL_HACMP03_CLIENT)  Total number of objects
failed:           0 (SESSION: 304)
ANE4961I (Session: 304, Node: CL_HACMP03_CLIENT)  Total number of bytes
transferred:  1.00 GB (SESSION: 304)
ANE4971I (Session: 304, Node: CL_HACMP03_CLIENT)  LanFree data bytes:
1.00 GB (SESSION: 304)
ANE4963I (Session: 304, Node: CL_HACMP03_CLIENT)  Data transfer time:
149.27 sec (SESSION: 304)
ANE4966I (Session: 304, Node: CL_HACMP03_CLIENT)  Network data transfer rate:
7,061.28 KB/sec (SESSION: 304)
ANE4967I (Session: 304, Node: CL_HACMP03_CLIENT)  Aggregate data transfer rate:
1,689.03 KB/sec (SESSION: 304)
ANE4964I (Session: 304, Node: CL_HACMP03_CLIENT)  Elapsed processing time:
00:10:24 (SESSION: 304)
```

Example 11-39 Client restore terminating successfully

```
Restoring        344,908 /opt/IBM/ISC/backups/backups/_acjvm/jre/lib/fonts/LucidaBrightRegular.ttf [Done]
Restoring        208,628 /opt/IBM/ISC/backups/backups/_acjvm/jre/lib/fonts/LucidaSansDemiBold.ttf [Done]
Restoring         91,352 /opt/IBM/ISC/backups/backups/_acjvm/jre/lib/fonts/LucidaSansDemiOblique.ttf [Done]


Restore processing finished.


Total number of objects restored:   20,338
Total number of objects failed:          0
Total number of bytes transferred:    1.00 GB
LanFree data bytes:                   1.00 GB
Data transfer time:                 149.27 sec
Network data transfer rate:       7,061.28 KB/sec
Aggregate data transfer rate:     1,689.03 KB/sec
Elapsed processing time:          00:10:24
```

Archived

```
tsm>
```

Result summary

We are able to have the HACMP cluster restarting an application with its LAN-free backup environment up and running.

Only the tape drive that was in use by the Storage Agent is reset and unloaded, the other one was under server control at failure time.

The restore operation can be restarted immediately without any intervention.

Archived

Part 4

Clustered IBM System Automation for Multiplatforms Version 1.2 environments and IBM Tivoli Storage Manager Version 5.3

In this part of the book, we discuss highly available clustering, using the Red Hat Enterprise Linux 3 Update 2 operating system with IBM System Automation for Multiplatforms Version 1.2 and Tivoli Storage Manager Version 5.3.

© Copyright IBM Corp. 2005. All rights reserved.

Archived

12

IBM Tivoli System Automation for Multiplatforms setup

In this chapter we describe Tivoli System Automation for Multiplatforms Version 1.2 cluster concepts, planning and design issues, preparing the OS and necessary drivers, and persistent binding of disk and tape devices. We also describe the installation of Tivoli System Automation and how to set up a two-node cluster.

© Copyright IBM Corp. 2005. All rights reserved.

12.1 Linux and Tivoli System Automation overview

In this section we provide some introductory information about Linux and Tivoli System Automation.

12.1.1 Linux overview

Linux is an open source UNIX-like kernel, originally created by Linus Torvalds. The term Linux is often used to mean the whole operating system, GNU/Linux. The Linux kernel, the tools, and the software needed to run an operating system are maintained by a loosely organized community of thousands of, mostly, volunteer programmers.

There are several organizations (distributors) that bundle the Linux kernel, tools, and applications to form a distribution, a package that can be downloaded or purchased and installed on a computer. Some of these distributions are commercial, others are not.

Linux is different from the other, proprietary, operating systems in many ways:

- There is no one person or organization that can be held responsible or called for support.
- Depending on the target group, the distributions differ largely in the kind of support that is available.
- Linux is available for almost all computer architectures.
- Linux is rapidly changing.

All these factors make it difficult to promise and provide generic support for Linux. As a consequence, IBM has decided on a support strategy that limits the uncertainty and the amount of testing.

IBM only supports the major Linux distributions that are targeted at enterprise customers, like Red Hat Enterprise Linux or SuSE Linux Enterprise Server. These distributions have release cycles of about one year, are maintained for five years, and require the user to sign a support contract with the distributor. They also have a schedule for regular updates. These factors mitigate the issues listed above. The limited number of supported distributions also allows IBM to work closely with the vendors to ensure interoperability and support.

For more details on the Linux distributions, please refer to:

http://www.redhat.com/
http://www.novell.com/linux/suse/index.html

12.1.2 IBM Tivoli System Automation for Multiplatform overview

Tivoli System Automation manages the availability of applications running in Linux systems or clusters on xSeries®, zSeries, iSeries™, pSeries, and AIX systems or clusters. It consists of the following features:

- High availability and resource monitoring
- Policy based automation
- Automatic recovery
- Automatic movement of applications
- Resource grouping

You can find the IBM product overview at:

http://www.ibm.com/software/tivoli/products/sys-auto-linux/

High availability and resource monitoring

Tivoli System Automation provides a high availability environment. High availability describes a system which is continuously available and which has a self-healing infrastructure to prevent downtime caused by system problems. Such an infrastructure detects improper operation of systems, transactions, and processes, and initiates corrective action without disrupting users.

Tivoli System Automation offers mainframe-like high availability by using fast detection of outages and sophisticated knowledge about application components and their relationships. It provides quick and consistent recovery of failed resources and whole applications either in place or on another system of a Linux cluster or AIX cluster without any operator intervention. Thus it relieves operators from manual monitoring, remembering application components and relationships, and therefore eliminates operator errors.

Policy based automation

Tivoli System Automation allows us to configure high availability systems through the use of policies that define the relationships among the various components. These policies can be applied to existing applications with minor modifications.

Once the relationships are established, Tivoli System Automation will assume responsibility for managing the applications on the specified nodes as configured. This reduces implementation time and the need for complex coding of applications. In addition, systems can be added without modifying scripts, and resources can be easily added, too.

There are sample policies available for IBM Tivoli System Automation. You can download them from the following Web page:

http://www.ibm.com/software/tivoli/products/sys-auto-linux/downloads.html

Archived

Automatic recovery

Tivoli System Automation quickly and consistently performs an automatic restart of failed resources or whole applications either in place or on another system of a Linux or AIX cluster. This greatly reduces system outages.

Automatic movement of applications

Tivoli System Automation manages the cluster-wide relationships among resources for which it is responsible. If applications need to be moved among nodes, the start and stop relationships, node requirements, and any preliminary or follow-up actions are automatically handled by Tivoli System Automation. This again relieves the operator from manual command entry, reducing operator errors.

Resource grouping

Resources can be grouped together in Tivoli System Automation. Once grouped, all relationships among the members of the group can be established, such as location relationships, start and stop relationships, and so on. After all of the configuration is completed, operations can be performed against the entire group as a single entity. This once again eliminates the need for operators to remember the application components and relationships, reducing the possibility of errors.

12.1.3 Tivoli System Automation terminology

The following terms are used within this redbook and within the Tivoli System Automation manual when describing Tivoli System Automation:

- Cluster / peer domain:

  The group of host systems upon which Tivoli System Automation manages resources is known as a cluster. A cluster can consist of one or more systems or nodes. The term “peer domain” is also used when referring to a cluster. The two terms are interchangeable.

- Node:

  A single host system that is part of a Tivoli System Automation cluster. Tivoli System Automation v1.2 supports up to 32 nodes within a cluster.

- Resource:

  A resource is any piece of hardware or software that can be defined to Tivoli System Automation. Resources have characteristics, or attributes, which can be defined. For example, when considering an IP address as a resource, attributes would include the IP address itself and the net mask.

- Resource attributes:

  A resource attribute describes some characteristics of a resource. There are two types of resource attributes: persistent attributes and dynamic attributes.

  - Persistent attributes: The attributes of the IP address just mentioned (the IP address itself and the net mask) are examples of persistent attributes – they describe enduring characteristics of a resource. While you could change the IP address and net mask, these characteristics are, in general, stable and unchanging.
  - Dynamic attributes: On the other hand, dynamic attributes represent changing characteristics of the resource. Dynamic attributes of an IP address, for example, would identify such things as its operational state.

- Resource class:

  A resource class is a collection of resources of the same type.

- Resource group:

  Resource groups are logical containers for a collection of resources. This container allows you to control multiple resources as a single logical entity. Resource groups are the primary mechanism for operations within Tivoli System Automation.

- Managed resource:

  A managed resource is a resource that has been defined to Tivoli System Automation. To accomplish this, the resource is added to a resource group, at which time it becomes manageable through Tivoli System Automation.

- Nominal state:

  The nominal state of a resource group indicates to Tivoli System Automation whether the resources with the group should be Online or Offline at this point in time. So setting the nominal state to “Offline” indicates that you wish for Tivoli System Automation to stop the resources in the group, and setting the nominal state to “Online” is an indication that you wish to start the resources in the resource group. You can change the value of the NominalState resource group attribute, but you cannot set the nominal state of a resource directly.

- Equivalency:

  An equivalency is a collection of resources that provides the same functionality. For example, equivalencies are used for selecting network adapters that should host an IP address. If one network adapter goes offline, IBM Tivoli System Automation selects another network adapter to host the IP address.

Archived

- Relationships:

  Tivoli System Automation allows the definition of relationships between resources in a cluster. There are two different relationship types:

  - Start-/stop relationships are used to define start and stop dependencies between resources. You can use the StartAfter, StopAfter, DependsOn, DependsOnAny, and ForcedDownBy relationships to achieve this. For example, a resource must only be started after another resource was started. You can define this by using the policy element StartAfter relationship.
  - Location relationships are applied when resources must, or should if possible, be started on the same or a different node in the cluster. Tivoli System Automation provides the following location relationships: Collocation, AntiCollocation, Affinity, AntiAffinity, and IsStartable.

- Quorum:

  The main goal of quorum operations is to keep data consistent and to protect critical resources. Quorum can be seen as the number of nodes in a cluster that are required to modify the cluster definition or perform certain cluster operations. There are two types of quorum:

  - Configuration quorum: This quorum determines when configuration changes in the cluster will be accepted. Operations affecting the configuration of the cluster or resources are only allowed when the absolute majority of nodes is online.
  - Operational quorum: This quorum is used to decide whether resources can be safely activated without creating conflicts with other resources. In case of a cluster splitting, resources can only be started in the subcluster which has a majority of nodes or has obtained a tie breaker.

- Tie breaker:

  In case of a tie in which a cluster has been partitioned into two subcluster with an equal number of nodes, the tie breaker is used to determine which subcluster will have an operational quorum.

12.2 Planning and design

Before we start the implementation of a Tivoli System Automation cluster in our Linux environment, we must consider the software and hardware requirements of the following software components:

- Tivoli System Automation for Multiplatforms Version 1.2
- Tivoli Storage Manager Version 5.3 Server
- Tivoli Storage Manager Version 5.3 Administration Center

- Tivoli Storage Manager Version 5.3 Backup/Archive Client
- Tivoli Storage Manager Version 5.3 Storage Agent

The Tivoli System Automation release notes give detailed information about required operating system versions and hardware. You can find the release notes online at:

http://publib.boulder.ibm.com/tividd/td/IBMTivoliSystemAutomationforMultiplatforms1.2.html

12.3 Lab setup

We have the following hardware components in our lab for the implementation of the Tivoli System Automation cluster, that will host different Tivoli Storage Manager software components:

- IBM 32 Bit Intel® based servers with IBM FAStT FC2-133 FC host bus adapters (HBAs)
- IBM DS4500 disk system (Firmware v6.1 with Storage Manager v9.10) with two EXP700 storage expansion units
- IBM 3582 tape library with two FC-attached LTO2 tape drives
- IBM 2005 B32 FC switch

Note: We use the most current supported combination of software components and drivers that fulfills the requirements for our lab hardware and our software requirements as they are at the time of writing. You need to check supported distributions, device driver versions, and other requirements when you plan such an environment. The online IBM HBA search tool is useful for this. It is available at:

http://knowledge.storage.ibm.com/servers/storage/support/hbasearch/interop/hbaSearch.do

We use the following steps to find our supported cluster configuration:

1. We choose a Linux distribution that meets the requirements for the components mentioned in 12.2, “Planning and design” on page 598. In our case, we use Red Hat Enterprise Linux AS 3 (RHEL AS 3). We could also use, for example, the SuSE Linux Enterprise Server 8 (SLES 8). The main difference would be the way in which we ensure persistent binding of devices. We discuss these ways to accomplish the different distributions in Persistent binding of disk and tape devices.

2. To find the necessary kernel level, we check the available versions of the necessary drivers and their kernel dependencies. All drivers are available for the 2.4.21-15.ELsmp kernel, which is shipped with Red Hat Enterprise Linux 3 Update 2. We use the following drivers:
   - IBM supported Qlogic HBA driver version 7.01.01 for HBA BIOS level 1.43
   - IBM FAStT RDAC driver version 09.10.A5.01
   - IBMtape driver version 1.5.3

Note: If you want to use the SANDISCOVERY option of the Tivoli Storage Manager Server and Storage Agent, you must also ensure to fulfill the required driver level for the HBA. You find the supported driver levels at:

http://www.ibm.com/support/docview.wss?uid=swg21193154

12.4 Preparing the operating system and drivers

During the installation of Red Hat Enterprise Linux Advanced Server 3 (RHEL AS 3) we also make sure to install the following packages:

- compat-libstdc++ (necessary for the installation of Tivoli System Automation for Multiplatforms)
- development packages (gcc, ...)
- kernel-sources

Note: Configuring NTP (Network Time Protocol) on all cluster nodes ensures correct time information on all nodes. This is very valuable once we have to compare log files from different nodes.

12.4.1 Installation of host bus adapter drivers

Although qlogic Fibre Channel drivers are shipped with RHEL AS 3, we need to install a version of the driver supported by IBM (in our case v7.01.01). We download the *non-failover* version of the driver and the readme file from:

http://www.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-54952

We verify that the HBAs have the supported firmware BIOS level, v1.43, and follow the instructions provided in the readme file, README.i2xLNX-v7.01.01.txt to install the driver. These steps are as follows:

1. We enter the HBA BIOS during startup and load the default values. After doing this, according to the readme file, we change the following parameters:
   - Loop reset delay: 8
   - LUNs per target: 0
   - Enable Target: Yes
   - Port down retry count: 12
2. In some cases the Linux QLogic HBA Driver disables an HBA after a path failure (with failover) occurred. To avoid this problem, we set the Connection Options in the QLogic BIOS to "1 - Point to Point only". More information about this issue can be found at:

   http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD101681
3. We continue with the installation as described in Section 6.4, “Building Symmetric Multi-Processor (SMP) Version of the Driver” in the readme file, README.i2xLNX-v7.01.01.txt.
   a. We prepare source headers for a Symmetric Multi-Processor (SMP) module build by opening a terminal window and changing to the kernel source directory /usr/src/linux-2.4.
   b. We verify that the kernel version information is correct in the makefile as shown in Example 12-1.

Example 12-1 Verifying the kernel version information in the Makefile

```
[root@diomede linux-2.4]# cat /proc/version
Linux version 2.4.21-15.ELsmp (bhcompile@bugs.build.redhat.com) (gcc version
3.2.3 20030502 (Red Hat Linux 3.2.3-34)) #1 SMP Thu Apr 22 00:18:24 EDT 2004
[root@diomede linux-2.4]# head -n 6 Makefile
VERSION = 2
PATCHLEVEL = 4
SUBLEVEL = 21
EXTRAVERSION = -15.ELsmp

KERNELRELEASE=$(VERSION).$(PATCHLEVEL).$(SUBLEVEL)$(EXTRAVERSION)
[root@diomede linux-2.4]#
```

   c. We copy the config file for our kernel to /usr/src/linux-2.4 as shown in Example 12-2.

Example 12-2 Copying kernel config file

```
[root@diomede linux-2.4]# cp configs/kernel-2.4.21-i686-smp.config .config
[root@diomede linux-2.4]# ls -l .config
-rw-r--r--    1 root     root        48349 Feb 24 10:33 .config
[root@diomede linux-2.4]#
```

   d. We rebuild the dependencies for the kernel with the **make dep** command.

e. We change back to the directory containing the device driver source code. There we execute **make all SMP=1 install** to build the driver modules.

f. We add the following lines to /etc/modules.conf:

```
alias scsi_hostadapter0 qla2300_conf
alias scsi_hostadapter1 qla2300
options scsi_mod max_scsi_luns=128
```

g. We load the module with **modprobe qla2300** to verify it is working correctly.

h. We rebuild the kernel ramdisk image:

```
# cd /boot
# cp -a initrd-2.4.21-15.ELsmp.img initrd-2.4.21-15.ELsmp.img.original
# mkinitrd -f initrd-2.4.21-15.ELsmp.img 2.4.21-15.ELsmp
```

i. We reboot to use the new kernel ramdisk image at startup.

Note: If you want to use the Tivoli Storage Manager SAN Device Mapping function as described in "Persistent binding of tape devices" on page 611, you need to install the SNIA (Storage Networking Industry Association) Host Bus Adapter (HBA) API support. You can do this via the **libinstall** script that is part of the driver source code.

12.4.2 Installation of disk multipath driver (RDAC)

We download the Redundant Disk Array Controller Driver (RDAC) and the readme file, linux_rdac_readme.txt from:

http://www.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-54973

We follow the instructions in the readme file, linux_rdac_readme.txt for the installation and setup. We do the following steps:

1. We disable the Auto Logical Drive Transfer (ADT/AVT) mode as it is not supported by the RDAC driver at this time. We use the script that is in the scripts directory of this DS4000 Storage Manager version 9 support for Linux CD. The name of the script file is DisableAVT_Linux.scr. We use the following steps to disable the ADT/AVT mode in our Linux host type partition:

   a. We open the DS4000 Storage Manager Enterprise Management window and highlight our subsystem

   b. We select **Tools**.

   c. We select **Execute script**.

   d. A script editing window opens. In this window:

      i. We select **File**.

      ii. We select **Load Script**.

iii. We give the full path name for the script file (<CDROM>/scripts/DisableAVT_Linux.scr) and click **OK**.

iv. We select **Tools**.

v. We select **Verify and Execute**.

2. To ensure kernel version synchronization between the driver and running kernel, we execute the following commands:

```
cd /usr/src/linux-2.4
make dep
make modules
```

3. We change to the directory that contains the RDAC source. We compile and install RDAC with the following commands:

```
make clean
make
make install
```

4. We edit the grub configuration file /boot/grub/menu.lst to use the kernel ramdisk image generated by the RDAC installation. Example 12-3 shows the grub configuration file.

Example 12-3 The grub configuration file /boot/grub/menu.lst

```
# grub.conf generated by anaconda
#
# Note that you do not have to rerun grub after making changes to this file
# NOTICE:  You do not have a /boot partition.  This means that
#          all kernel and initrd paths are relative to /, eg.
#          root (hd0,0)
#          kernel /boot/vmlinuz-version ro root=/dev/hda1
#          initrd /boot/initrd-version.img
#boot=/dev/hda
default=1
timeout=0
splashimage=(hd0,0)/boot/grub/splash.xpm.gz
title Red Hat Enterprise Linux AS (2.4.21-15.ELsmp)
        root (hd0,0)
        kernel /boot/vmlinuz-2.4.21-15.ELsmp ro root=LABEL=/ hdc=ide-scsi
        initrd /boot/initrd-2.4.21-15.ELsmp.img
title Red Hat Linux (2.4.21-15.ELsmp) with MPP support
        root (hd0,0)
        kernel /boot/vmlinuz-2.4.21-15.ELsmp ro root=LABEL=/ hdc=ide-scsi
ramdisk_size=15000
        initrd /boot/mpp-2.4.21-15.ELsmp.img
title Red Hat Enterprise Linux AS-up (2.4.21-15.EL)
        root (hd0,0)
        kernel /boot/vmlinuz-2.4.21-15.EL ro root=LABEL=/ hdc=ide-scsi
        initrd /boot/initrd-2.4.21-15.EL.img
```

Archived

5. After a reboot, we verify the correct setup of the RDAC as shown in Example 12-4.

Example 12-4 Verification of RDAC setup

```
[root@diomede linuxrdac]# lsmod | grep mpp
mpp_Vhba                82400 -59
mpp_Upper               74464   0  [mpp_Vhba]
scsi_mod               112680   9  [IBMtape sr_mod ide-scsi st mpp_Vhba qla2300
mpp_Upper sg sd_mod]
[root@diomede linuxrdac]# ls -lR /proc/mpp
/proc/mpp:
total 0
dr-xr-xr-x    4 root     root            0 Feb 24 11:46 ITSODS4500_A
crwxrwxrwx    1 root     root     254,   0 Feb 24 11:46 mppVBusNode

/proc/mpp/ITSODS4500_A:
total 0
dr-xr-xr-x    3 root     root            0 Feb 24 11:46 controllerA
dr-xr-xr-x    3 root     root            0 Feb 24 11:46 controllerB
-rw-r--r--    1 root     root            0 Feb 24 11:46 virtualLun0
-rw-r--r--    1 root     root            0 Feb 24 11:46 virtualLun1
-rw-r--r--    1 root     root            0 Feb 24 11:46 virtualLun2
[...]
```

6. Finally we execute **mppUpdate** to update the /var/mpp/devicemapping file.

12.4.3 Installation of the IBMtape driver

We download the IBMtape driver v1.5.3 for the RHEL 2.4.21-15 kernel. You can download the driver at:

http://www.ibm.com/servers/storage/support/tape

The driver is packed as an rpm file. We install the driver by executing the **rpm** command as shown in Figure 12-5.

Example 12-5 Installation of the IBMtape driver

```
[root@diomede ibmtape]# rpm -ihv IBMtape-1.5.3-2.4.21-15.EL.i386.rpm
Preparing...                ########################################### [100%]
Installing IBMtape
   1:IBMtape                ########################################### [100%]
Warning: loading /lib/modules/2.4.21-15.ELsmp/kernel/drivers/scsi/IBMtape.o
will taint the kernel: non-GPL license - USER LICENSE AGREEMENT FOR IBM DEVICE
DRIVERS
  See http://www.tux.org/lkml/#export-tainted for information about tainted
modules
Module IBMtape loaded, with warnings
```

```
IBMtape loaded
[root@diomede ibmtape]#
```

To verify that the installation was successful and the module was loaded correctly, we take a look at the attached devices as shown in Figure 12-6.

Example 12-6 Device information in /proc/scsi/IBMtape and /proc/scsi/IBMchanger

```
[root@diomede root]# cat /proc/scsi/IBMtape
IBMtape version: 1.5.3
IBMtape major number: 252
Attached Tape Devices:
Number  Model       SN                HBA                            FO Path
0       ULT3580-TD2 1110176223        QLogic Fibre Channel 2300      NA
1       ULT3580-TD2 1110177214        QLogic Fibre Channel 2300      NA
[root@diomede root]# cat /proc/scsi/IBMchanger
IBMtape version: 1.5.3
IBMtape major number: 252
Attached Changer Devices:
Number  Model       SN                HBA                            FO Path
0       ULT3582-TL  0000013108231000  QLogic Fibre Channel 2300      NA
[root@diomede root]#
```

Note: IBM provides IBMtapeutil, a tape utility program that exercises or tests the functions of the Linux device driver, IBMtape. It performs tape and medium changer operations. You can download it with the IBMtape driver.

12.5 Persistent binding of disk and tape devices

Whenever we attach a server to a storage area network (SAN), we must ensure the correct setup of our connections to SAN devices. Depending on applications and device drivers, it is necessary to set up persistent bindings on one or more of the different driver levels. Otherwise device addresses can change if changes in the SAN happen. For example, this can be caused by an outage of a single device in the SAN, causing SCSI IDs to change at reboot of the server.

12.5.1 SCSI addresses

Linux uses the following addressing scheme for SCSI devices:

- SCSI adapter (host)
- Bus (channel)
- Target id (ID)
- Logical unit number (LUN)

Archived

The following example shows an entry of /proc/scsi/scsi. We can display all entries with the command **cat /proc/scsi/scsi**.

```
Host: scsi0 Channel: 00 Id: 01 Lun: 02
  Vendor: IBM      Model: 1742-900         Rev: 0520
  Type:   Direct-Access                    ANSI SCSI revision: 03
```

This example shows the third disk (Lun: 02) of the second device (Id: 01) that is connected to the first port (Channel: 00) of the first SCSI or Fibre Channel adapter (Host: scsi0) of the system. Many SCSI or Fibre Channel adapters have only one port. For these adapters, the channel number is always 0 for all attached devices.

Without persistent binding of the target IDs, the following problem can arise. If the first device (Id: 00) has an outage and a reboot of the server is necessary, the target ID of the second device will change from 1 to 0.

Depending on the type of SCSI device, the LUN has different meanings. For disk subsystems, the LUN refers to an individual virtual disk assigned to the server. For tape libraries, LUN 0 is often used for a tape drive itself acting as a sequential access data device, while LUN 1 on the same SCSI target ID points to the same tape drive acting as a medium changer device.

12.5.2 Persistent binding of disk devices

Linux uses special device files to access hard disks. In distributions with Linux kernel 2.4, device files for SCSI disks normally start with /dev/sd, followed by one or two letters which refer to the disk. For example, the first SCSI disk is /dev/sda, the second /dev/sdb, the third /dev/sdc and so on. During startup, Linux scans for the attached disk devices. If the second SCSI disk is unavailable for some reason, /dev/sdb refers to the former third SCSI disk after a reboot.

To circumvent this problem in their Linux kernel 2.4 based distributions, SuSE and Red Hat provide tools that enable a persistent binding for device files. SuSE uses an approach based on the SCSI address of the devices. The tool is called *scsidev*. Red Hat uses the universal unique identifier (UUID) of a disk. The tool for this purpose is *devlabel*.

Tivoli System Automation also uses the SCSI address to access the tie breaker disk, which is necessary for a quorum in a two-node cluster. We recommend to make sure the SCSI addresses for disk devices are persistent regardless of whether you use SLES or RHEL.

Note: Some disk subsystems provide multipath drivers that create persistent special device files. The IBM subsystem device driver (SDD) for ESS, DS6000, and DS8000 creates persistent vpath devices in the form /dev/vpath*. If you use this driver for your disk subsystem, you do not need scsidev or devlabel to create persistent special device files for disks containing file systems. You can use the device files directly to create partitions and file systems.

Persistent binding of SCSI addresses for disk devices

When using SLES 8 with scsidev, we must ensure persistent SCSI addresses for all disk devices. If we use RHEL with devlabel, a persistent SCSI address is only necessary for the tie breaker disk used for Tivoli System Automation for Multiplatforms quorum.

We can ensure persistent SCSI addresses in different ways, depending on the storage subsystem and the driver. In every case, we must keep the order of SCSI adapters in our server. Otherwise the host number of the SCSI address can change. The only part of the SCSI address which can alter because of changes in our SAN is the target ID. So you must configure the target IDs to be persistent.

When using a DS4xxx storage server like we do in our environment, RDAC does the persistent binding. The first time the RDAC driver sees a storage array, it will arbitrarily assign a target ID for the virtual target that represents the storage array. At this point the target ID assignment is not persistent. It could change on a reboot. The **mppUpdate** utility updates the RDAC driver configuration files so that these target ID assignments are persistent and do not change across reboots. RDAC stores the mapping in /var/mpp/devicemapping. This file has the following contents in our environment:

```
0:ITSODS4500_A
```

If you use other storage subsystems that do not provide a special driver providing persistent target IDs, you can use the persistent binding functionality for target IDs of the Fibre Channel driver. See the documentation of your Fibre Channel driver for further details.

Persistent binding of disk devices with SLES 8

The scsidev utility adds device files containing the SCSI address to the directory /dev/scsi. During boot, scsidev is executed and updates the device files if necessary. Example 12-7 shows the contents of /proc/scsi/scsi. There is a local disk connected via SCSI host 0, two disks connected to a DS4300 Turbo via SCSI host 4, and two disks connected to a second DS4300 Turbo via SCSI host 4. The SCSI host 4 is a virtual host, created by the RDAC driver. As we use the RDAC driver, the SCSI IDs are persistent.

Example 12-7 Contents of /proc/scsi/scsi

```
sles8srv:~ # cat /proc/scsi/scsi
Attached devices:
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM-ESXS Model: DTN073C3UCDY10FN Rev: S25J
  Type:   Direct-Access                    ANSI SCSI revision: 03
[...]
Host: scsi4 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: VirtualDisk      Rev: 0610
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 00 Lun: 01
  Vendor: IBM      Model: VirtualDisk      Rev: 0610
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 01 Lun: 00
  Vendor: IBM      Model: VirtualDisk      Rev: 0610
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 01 Lun: 01
  Vendor: IBM      Model: VirtualDisk      Rev: 0610
  Type:   Direct-Access                    ANSI SCSI revision: 03
sles8srv:~ #
```

To access the disks and partitions, we use the SCSI devices created by scsidev. Example 12-8 shows these device files.

Example 12-8 SCSI devices created by scsidev

```
sles8srv:~ # ls -l /dev/scsi/s*
brw-rw----  1 root   disk     8,  0 Nov  5 11:29 /dev/scsi/sdh0-0c0i010
brw-rw----  1 root   disk     8,  1 Nov  5 11:29 /dev/scsi/sdh0-0c0i010p1
brw-rw----  1 root   disk     8,  2 Nov  5 11:29 /dev/scsi/sdh0-0c0i010p2
brw-rw----  1 root   disk     8,  3 Nov  5 11:29 /dev/scsi/sdh0-0c0i010p3
brw-rw----  1 root   disk     8, 16 Feb 21 13:23 /dev/scsi/sdh4-0c0i010
brw-rw----  1 root   disk     8, 17 Feb 21 13:23 /dev/scsi/sdh4-0c0i010p1
brw-rw----  1 root   disk     8, 32 Feb 21 13:23 /dev/scsi/sdh4-0c0i011
brw-rw----  1 root   disk     8, 33 Feb 21 13:23 /dev/scsi/sdh4-0c0i011p1
brw-rw----  1 root   disk     8, 48 Feb 21 13:23 /dev/scsi/sdh4-0c0i110
brw-rw----  1 root   disk     8, 49 Feb 21 13:23 /dev/scsi/sdh4-0c0i110p1
brw-rw----  1 root   disk     8, 64 Feb 21 13:23 /dev/scsi/sdh4-0c0i111
brw-rw----  1 root   disk     8, 65 Feb 21 13:23 /dev/scsi/sdh4-0c0i111p1
crw-r-----  1 root   disk    21,  0 Nov  5 11:29 /dev/scsi/sgh0-0c0i010
crw-r-----  1 root   disk    21,  1 Nov  5 11:29 /dev/scsi/sgh0-0c0i810
crw-r-----  1 root   disk    21,  2 Feb 21 13:23 /dev/scsi/sgh4-0c0i010
crw-r-----  1 root   disk    21,  3 Feb 21 13:23 /dev/scsi/sgh4-0c0i011
crw-r-----  1 root   disk    21,  4 Feb 21 13:23 /dev/scsi/sgh4-0c0i110
crw-r-----  1 root   disk    21,  5 Feb 21 13:23 /dev/scsi/sgh4-0c0i111
sles8srv:~ #
```

We use these device files in /etc/fstab to mount our file systems. For example, we access the filesystem located at the first partition of the first disk on the second DS4300 Turbo via /dev/scsi/sdh4-0c0i1l0p1. In case that the first DS4300 Turbo cannot be accessed and the server must be rebooted, this device file will still point to the correct device.

Persistent binding of disk devices with RHEL 3

RHEL provides the devlabel utility to establish a persistent binding to a disk or partition. Devlabel creates a symbolic link for each configured device. The symbolic link refers to the virtual device file, e.g. /dev/sda. Devlabel associates the name of the symbolic link with the UUID of the hard disk or partition. During startup, **`devlabel restart`** is called from the /etc/rc.sysinit script. It reads the configuration file /etc/sysconfig/devlabel and validates the symbolic links. If a link is invalid, it searches for the virtual device file that points to correct UUID and updates the link.

First we need to create the partitions on the disks. We create primary partitions on every disk where we place file systems. We use **`fdisk`** to create partitions. After we create the partitions, we must reload the Fibre Channel driver on the other node to detect the partitions there. Then we create file systems on the partitions.

> **Attention:** The UUID of a partition changes after the creation of a file system on it. Example 12-9 shows this behavior. So we ensure to only use devlabel after we created the file systems.

Example 12-9 UUID changes after file system is created

```
[root@diomede root]# devlabel printid -d /dev/sdb1
S83.3:600a0b80001742330000000e41f14177IBMVirtualDisksector63
[root@diomede root]# mkfs.ext3 /dev/sdb1
...
[root@diomede root]# devlabel printid -d /dev/sdb1
P:35e2136a-d233-4624-96bf-7719298b766a
[root@diomede root]#
```

To create persistent symbolic links, we follow these steps for the partitions on every disk device except the tie breaker disk. We need to accomplish these steps on both nodes:

1. We verify that the partition has a UUID, for example:

   ```
   [root@diomede root]# devlabel printid -d /dev/sdb1
   P:35e2136a-d233-4624-96bf-7719298b766a
   [root@diomede root]#
   ```

2. We add a persistent symbolic link for the disk:

```
[root@diomede root]# devlabel add -d /dev/sdb1 -s /dev/tsmdb1
SYMLINK: /dev/tsmdb1 -> /dev/sdb1
Added /dev/tsmdb1 to /etc/sysconfig/devlabel
[root@diomede root]#
```

3. We verify the contents of the configuration file /etc/sysconfig/devlabel. There must be an entry for the added symbolic link. Example 12-10 shows the contents of /etc/sysconfig/devlabel in our configuration for the highly available Tivoli Storage Manager Server described in Chapter 13, “Linux and Tivoli System Automation with IBM Tivoli Storage Manager Server” on page 617.

Example 12-10 Devlabel configuration file /etc/sysconfig/devlabel

```
# devlabel configuration file
#
# This file should generally not be edited by hand.
# Instead, use the /sbin/devlabel program to make changes.
# devlabel by Gary Lerhaupt <gary_lerhaupt@dell.com>
#
# format:  <SYMLINK> <DEVICE> <UUID>
# or format:  <RAWDEVICE> <DEVICE> <UUID>

/dev/tsmdb1 /dev/sdb1 P:35e2136a-d233-4624-96bf-7719298b766a
/dev/tsmdb1mr /dev/sdc1 P:69fc6ab5-677d-426e-b662-ee9b3355f42e
/dev/tsmlg1 /dev/sdd1 P:75fafbaf-250d-4504-82b7-3deda77b63c9
/dev/tsmlg1mr /dev/sde1 P:64191c25-8928-4817-a7a2-f437da50a5d8
/dev/tsmdp /dev/sdf1 P:83664f89-4c7a-4238-9b9a-c63376dda39a
/dev/tsmfiles /dev/sdf2 P:51a4688d-7392-4cf6-933b-32a8d840c0e1
/dev/tsmisc /dev/sdg1 P:4c10f0be-1fdf-4fee-8fc9-9af27926868e
```

Important: In case that you bring a failed node back online, check the devlabel configuration file /etc/sysconfig/devlabel and the symbolic links that are created by devlabel before you are bringing resources back online on this node. If some LUNs were not available during startup, you may need to reload the SCSI drivers and execute the **`devlabel restart`** command to update the symbolic links.

Persistent binding of disk devices with Kernel 2.6 based OS

With Linux kernel 2.6 the new user space solution *udev* for handling dynamic devices while keeping persistent device names is introduced. You can use udev for persistent binding of disk devices with SLES 9 and RHEL 4. See the documentation of your kernel 2.6 based enterprise Linux distribution for more information on how to use udev for persistent binding.

12.6 Persistent binding of tape devices

Device configuration on SAN-attached devices is made simpler with the Tivoli Storage Manager SAN Device Mapping function (SDM). SDM uses the SNIA (Storage Networking Industry Association) Host Bus Adapter (HBA) API to perform SAN discovery. The device serial number, manufacturer, and worldwide name are initially recorded for each storage device. When the device configuration changes, Tivoli Storage Manager can automatically update the device path information without the need for device persistent binding.

In our lab environment we use a IBM TotalStorage 3582 Tape Library with two LTO2 tape drives, each of them with one FC port. The first tape drives also acts as medium changer device. As we depend on the path to the first tape drive anyway, we do not activate the SDM function.

You can find a list of supported HBA driver versions for SDM at:

http://www.ibm.com/support/docview.wss?uid=swg21193154

12.7 Installation of Tivoli System Automation

Before we start the installation and configuration of Tivoli System Automation for Multiplatforms, we must set the management scope for RSCT for all users of Tivoli System Automation for Multiplatforms on all nodes. We set the variable permanently by setting it in the profile.

```
export CT_MANAGEMENT_SCOPE=2
```

We downloaded the Tivoli System Automation for Multiplatforms tar file from the Internet, so we extract the file, using the following command:

```
tar -xvf <tar file>
```

Now we change to the appropriate directory for our platform:

```
cd SAM12/i386
```

We install the product with the **installSAM** script as shown in Example 12-11.

Example 12-11 Installation of Tivoli System Automation for Multiplatforms

```
[root@diomede i386]# ./installSAM

installSAM: A general License Agreement and License Information specifically
for System Automation will be shown. Scroll down using the Enter key (line by
line) or Space bar (page by page). At the end you will be asked to accept the
terms to be allowed to install the product. Select Enter to continue.
```

```
[...]

installSAM: Installing System Automation on platform: i686

[...]

installSAM: The following license is installed:
Product ID: 5588
Creation date: Tue 11 May 2004 05:00:00 PM PDT
Expiration date: Thu 31 Dec 2037 03:59:59 PM PST

installSAM: Status of System Automation after installation:
 ctrmc            rsct             11754   active
 IBM.ERRM         rsct_rm          11770   active
 IBM.AuditRM      rsct_rm          11794   active
 ctcas            rsct                     inoperative
 IBM.SensorRM     rsct_rm                  inoperative
[root@diomede i386]#
```

We update to the latest fixpack level of Tivoli System Automation for Multiplatforms. The fixpacks are published in the form of tar files. We run the same steps as explained above for the normal installation. Fixpacks are available at:

http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliSystemAutomationforLinux.html

At the time of writing this book, the latest fixpack level is 1.2.0.3. We extract the tar file. Now we change to the appropriate directory for our platform:

```
cd SAM1203/<arch>
```

We update to this fixpack level by executing the **installSAM** script.

12.8 Creating a two-node cluster

Before proceeding, we make sure that all entries for the nodes of the cluster in our local /etc/hosts files on all nodes and the name server entries are identical.

As we run a two-node cluster, we need some additional configuration to detect network interface failures. The cluster software periodically tries to reach each network interface of the cluster. If there is a two-node cluster and one interface fails on one node, the other interface on the other node is not able to get a response from the peer and will also be flagged offline. To avoid this behavior, the cluster software must be told to contact a network instance outside the cluster.

Archived

The best practice is to use the default gateway of the subnet the interface is in. On each node we create the file /usr/sbin/cluster/netmon.cf. Each line of this file should contain the machine name or IP address of the external instance. An IP address should be specified in dotted decimal format. We add the IP address of our default gateway to /usr/sbin/cluster/netmon.cf.

To create this cluster, we need to:

1. Access a console on each node in the cluster and log in as root.
2. Execute **echo $CT_MANAGEMENT_SCOPE** to verify that this environment variable is set to 2.
3. Issue the **preprpnode** command on all nodes to allow communication between the cluster nodes. In our example, we issue **preprpnode diomede lochness** on both nodes.
4. Create a cluster with the name cl_itsamp running on both nodes. The following command can be issued from any node.

   ```
   mkrpdomain cl_itsamp diomede lochness
   ```

5. To look up the status of cl_itsamp, we issue the **lsrpdomain** command. The output looks like this:

   ```
   Name      OpState RSCTActiveVersion MixedVersions TSPort GSPort
   cl_itsamp Offline 2.3.4.5           No            12347  12348
   ```

   The cluster is defined but offline.

6. We issue the **startrpdomain cl_itsamp** command to bring the cluster online. When we run the **lsrpdomain** command again, we see that the cluster is still in the process of starting up, the OpState is Pending Online.

   ```
   Name      OpState        RSCTActiveVersion MixedVersions TSPort GSPort
   cl_itsamp Pending online 2.3.4.5           No            12347  12348
   ```

   After a short time the cluster is started, so when executing **lsrpdomain** again, we see that the cluster is now online:

   ```
   Name      OpState RSCTActiveVersion MixedVersions TSPort GSPort
   cl_itsamp Online  2.3.4.5           No            12347  12348
   ```

7. We set up the disk tie breaker and validate the configuration. The tie breaker disk in our example has the SCSI address 1:0:0:0 (host, channel, id, lun). We need to create the tie breaker resource, and change the quorum type afterwards. Example 12-12 shows the necessary steps.

Example 12-12 Configuration of the disk tie breaker

```
[root@diomede root]# mkrsrc IBM.TieBreaker Name="tb1" Type="SCSI" \
> DeviceInfo="Host=1 Channel=0 Id=0 Lun=0" HeartbeatPeriod=5
[root@diomede root]# chrsrc -c IBM.PeerNode OpQuorumTieBreaker="tb1"
[root@diomede root]# lsrsrc -c IBM.PeerNode
```

```
Resource Class Persistent Attributes for IBM.PeerNode
resource 1:
        CommittedRSCTVersion  = ""
        ActiveVersionChanging = 0
        OpQuorumOverride      = 0
        CritRsrcProtMethod    = 1
        OpQuorumTieBreaker    = "tb1"
        QuorumType            = 0
        QuorumGroupName       = ""
[root@diomede root]#
```

IBM provides many resource policies for Tivoli System Automation. You can download the latest version of the sam.policies rpm from:

http://www.ibm.com/software/tivoli/products/sys-auto-linux/downloads.html

We install the rpm (in our case sam.policies-1.2.1.0-0.i386.rpm) on both nodes. The policies are placed within different directories below /usr/sbin/rsct/sapolicies.

We use additional policies for the Tivoli Storage Manager server, client, and Storage Agent. If these policies are not included in the rpm you can download them on the Web page of this redbook.

Note: The policy scripts must be present on all nodes in the cluster.

12.9 Troubleshooting and tips

Here are some tips that may help you if you have any problems in a cluster.

System log file

Tivoli System Automation and the provided resource policies write logging information to the system log file /var/log/messages. When you do the initial cluster testing before using the cluster in production, you can use **`tail -f /var/log/messages`** to follow the logging information.

Excluded list of nodes

You can temporary exclude nodes from the cluster with the **`samctrl`** command. If the node that you to put on the list of excluded nodes hosts resources, the resources are moved to another node in the cluster.

You can use the command with the following parameters:

- **samctrl -u a [Node [Node [...]]]** adds one or more specified nodes to the excluded list of nodes.
- **samctrl -u d [Node [Node [...]]]** deletes one or more specified nodes to the excluded list of nodes.

Recovery resource manager

The recovery resource manager (RecoveryRM) serves as the decision engine for Tivoli System Automation. Once a policy for defining resource availabilities and relationships is defined, this information is supplied to the Recovery RM. This RM runs on every node in the cluster, with exactly one Recovery RM designated as the master. The master evaluates the monitoring information from the various resource managers. Once a situation develops that requires intervention, the Recovery RM drives the decisions that result in start or stop operations on the resources as needed.

We can display the status of the RecoveryRM Subsystem with the **lssrc** command as shown in Example 12-13.

Example 12-13 Displaying the status of the RecoveryRM with the lssrc command

```
[root@diomede root]# lssrc -ls IBM.RecoveryRM
Subsystem         : IBM.RecoveryRM
PID               : 32552
Cluster Name      : cl_itsamp
Node Number       : 1
Daemon start time : Thu 24 Feb 2005 03:49:39 PM PST

Daemon State:
   My Node Name         : diomede
   Master Node Name     : lochness (node number = 2)
   Our IVN              : 1.2.0.3
   Our AVN              : 1.2.0.3
   Our CVN              : 1109201832444 (0x1bc421d13a8)
   Total Node Count     : 2
   Joined Member Count  : 2
   Config Quorum Count  : 2
   Startup Quorum Count : 1
   Operational Quorum State: HAS_QUORUM
   In Config Quorum     : TRUE
   In Config State      : TRUE
   Replace Config State : FALSE


Information from malloc about memory use:
   Total Space    : 0x000e6000 (942080)
   Allocated Space: 0x000ca9d0 (829904)
```

```
   Unused Space   : 0x0001b630 (112176)
   Freeable Space : 0x00017d70 (97648)

Total Address Space Used : 0x0198c000 (26787840)
   Unknown               : 0x00000000 (0)
   Text                  : 0x009b3000 (10170368)
   Global Data           : 0x00146000 (1335296)
   Dynamic Data          : 0x00a88000 (11042816)
   Stack                 : 0x000f0000 (983040)
   Mapped Files          : 0x0031b000 (3256320)
   Shared Memory         : 0x00000000 (0)

[root@diomede root]#
```

Archived

13

Linux and Tivoli System Automation with IBM Tivoli Storage Manager Server

In this chapter we describe the necessary configuration steps to make the Tivoli Storage Manager server highly available with Tivoli System Automation V1.2 on Linux.

© Copyright IBM Corp. 2005. All rights reserved.

13.1 Overview

In a Tivoli System Automation environment, independent servers are configured to work together in order to enhance applications availability using shared disk subsystems.

We configure Tivoli Storage Manager server as a highly available application in this Tivoli System Automation environment. Clients can connect to the Tivoli Storage Manager server using a virtual server name.

To run properly, the Tivoli Storage Manager server needs to be installed and configured in a special way, as a resource in a resource group in Tivoli System Automation. This chapter covers all the tasks we follow in our lab environment to achieve this goal.

13.2 Planning storage

In the following sections we provide some information about our storage configuration and RAID protection. For detailed information on how to protect your Tivoli Storage Manager server, refer to: “Protecting and Recovering Your Server” in the *IBM Tivoli Storage Manager for Linux Administrator's Guide.*

Tivoli Storage Manager server

We use the following configuration for the setup of the Tivoli Storage Manager server:

- Tivoli Storage Manager mirroring for database and log volumes
- RAID0 shared disks volumes configured on separate storage subsystem arrays for database and log volumes copies:
  - /tsm/db1
  - /tsm/db1mr
  - /tsm/lg1
  - /tsm/lg1mr
- Database and log writes set to sequential (which disables DBPAGESHADOW)
- Log mode set to RollForward
- RAID1 shared disk volumes for configuration files and disk storage pools.
  - /tsm/files
  - /tsm/dp

Archived

Tivoli Storage Manager Administration Center

The Administration Center can be a critical application for environments where the administrator and operators are not confident with the IBM Tivoli Storage Manager Command Line Administrative Interface. So we decided to experiment with a clustered installation even if it is currently not supported.

We use a RAID1 protected shared disk volume for both code and data (servers connections and ISC users definitions) under a shared file system that we create and activate before to ISC code installation. The mountpoint of this file system is /tsm/isc.

13.3 Lab setup

The Tivoli Storage Manager virtual server configuration we use for the purpose of this chapter is shown in Table 13-1.

Table 13-1 Lab Tivoli Storage Manager server cluster resources

System Automation resource group:	SA-tsmserver-rg
TSM server name:	TSMSRV05
TSM server IP address:	9.1.39.54
TSM database disks[a]:	/tsm/db1, /tsm/db1mr
TSM recovery log disks:	/tsm/lg1, /tsm/lg1mr
TSM storage pool disk:	/tsm/dp
TSM configuration and log file disk:	/tsm/files

a. We choose two disk drives for the database and recovery log volumes so that we can use the Tivoli Storage Manager mirroring feature.

13.4 Installation

In this section we describe the installation of all necessary software for the Tivoli Storage Manager Server cluster.

Archived

13.4.1 Installation of Tivoli Storage Manager Server

Tivoli Storage Manager Server can be installed by either the install_server script or by directly installing the necessary rpms. We use the **rpm** command as shown in Example 13-1 to install the following installation packages:

- TIVsm-server-5.3.0-0.i386.rpm
- TIVsm-license-5.3.0-0.i386.rpm

Example 13-1 Installation of Tivoli Storage Manager Server

```
[root@diomede i686]# rpm -ihv TIVsm-server-5.3.0-0.i386.rpm
Preparing...                ########################################### [100%]
   1:TIVsm-server           ########################################### [100%]
Allocated space for db.dsm: 17825792 bytes
Allocated space for log.dsm: 9437184 bytes

Tivoli Storage Manager for Linux/i386
Version 5, Release 3, Level 0.0
[...]
************************************************************
IMPORTANT: Read the contents of file /README
           for extensions and corrections to printed
           product documentation.
************************************************************
[root@diomede i686]# rpm -ihv TIVsm-license-5.3.0-0.i386.rpm
Preparing...                ########################################### [100%]
   1:TIVsm-license          ########################################### [100%]
[root@diomede i686]#
```

We add /opt/tivoli/tsm/server/bin to our $PATH variable in our .bash_profile file. We close our shell and log in again to activate this new setting.

13.4.2 Installation of Tivoli Storage Manager Client

The Tivoli Storage Manager client does not support the default locale for Linux, en_US.UTF-8. There may be some files that cannot be backed up with the default locale, causing error messages in the dsmerror log, and the backup operation to stop. To avoid this problem, we set the locale LC_ALL to en_US or another supported locale.

> **Note:** The X Windows System X11R6 is a requirement to install the client. If it is not installed and you do not plan to use the end user GUI, you have to add the --nodeps option of rpm to disable the check for requirements.

Archived

To install Tivoli Storage Manager Client, we follow these steps:

1. We access a console and log in as root.
2. We change the directory to cdrom directory. We can find the latest information about the client in the file README.1ST. We change to the directory for our platform with **`cd tsmcli/linux86`**.
3. We enter the following commands to install the API and the Tivoli Storage Manager B/A client. This installs the command line, the GUI, and the administrative client:

   ```
   rpm -ihv TIVsm-API.i386.rpm
   rpm -ihv TIVsm-BA.i386.rpm
   ```

   We make sure to install these packages in the recommended order. This is required because the Tivoli Storage Manager API package is a prerequisite of the B/A client package.
4. The Tivoli Storage Manager installation default language is English. If you want to install an additional language, you need to install the appropriate rpm provided in the installation folder.

We add /opt/tivoli/tsm/client/ba/bin to our $PATH variable in our .bash_profile file. We close our shell and log in again to activate this new setting.

13.4.3 Installation of Integrated Solutions Console

The Tivoli System Automation cluster requires entries for all managed file systems in /etc/fstab. The following entry is necessary for the Integrated Solutions Console (ISC). We create the mount point and insert this entry to /etc/fstab on both nodes.

```
/dev/tsmisc     /tsm/isc        ext3    noauto  0 0
```

We mount the file system /tsm/isc on our first node, diomede. There we install the ISC.

> **Attention:** Never mount file systems of a shared disk concurrently on both nodes unless you use a shared disk file system. Doing so destroys the file system and probably all data of the file system will be lost. If you need a file system concurrently on multiple nodes, use a shared disk file system like the IBM General Parallel File System (GPFS).

The installation of Tivoli Storage Manager Administration Center is a two step install. First, we install the Integrated Solutions Console (ISC). Then we deploy the Tivoli Storage Manager Administration Center into the Integrated Solutions Console. Once both pieces are installed, we are able to administer Tivoli Storage Manager from a browser anywhere in our network.

Archived

Note: The installation process of the Integrated Solutions Console can take anywhere from 30 minutes to two hours to complete. The time to install depends on the speed of your processor and memory.

To install Integrated Solutions Console, we follow these steps:

1. We access a console and log in as root.
2. We change the directory to cdrom directory. We are installing with TSM_ISC_5300_<PLATFORM>.tar, so we issue the following command:

   ```
   tar -xf TSM_ISC_5300_<PLATFORM>.tar
   ```

3. We can run one of the following commands to install the ISC:
   - For InstallShield wizard install:

     ```
     ./setupISC
     ```

   - For console wizard install:

     ```
     ./setupISC -console
     ```

   - For silent install, we run the following command on a single line:

     ```
     ./setupISC -silent -W ConfigInput.adminName="<user name>"
     -W ConfigInput.adminPass="<user password>"
     -W ConfigInput.verifyPass="<user password>"
     -W PortInput.webAdminPort="<web administration port>"
     -W PortInput.secureAdminPort="<secure administration port>"
     -W MediaLocationInput.installMediaLocation="<media location>"
     -P ISCProduct.installLocation="<install location>"
     ```

     If we do not provide all parameters, default values will be used.

   We install ISC with the following command:

   ```
   [root@diomede tsm-isc]# ./setupISC -silent \
   > -W ConfigInput.adminName="iscadmin" \
   > -W ConfigInput.adminPass="itsosj" \
   > -W ConfigInput.verifyPass="itsosj" \
   > -P ISCProduct.installLocation="/tsm/isc/"
   [root@diomede tsm-isc]#
   ```

Important: If you use the silent install method, the ISC admin password will be visible in the history file of your shell. For security reasons, we recommend to remove the command from the history file (/root/.bash_history if you use bash). The same applies for the installation of the Administration Center (AC).

During the installation, **setupISC** adds the following entry to /etc/inittab:

```
iscn:23:boot:/tsm/isc/PortalServer/bin/startISC.sh ISC_Portal ISCUSER ISCPASS
```

We want Tivoli System Automation for Multiplatforms to control the startup and shutdown of ISC. So we simply delete this line or put a hash (#) in front of it.

> **Note:** All files of the ISC reside on the shared disk. We do not need to install it on the second node.

13.4.4 Installation of Administration Center

After we finish the installation of ISC, we continue with the installation of the Administration Center (AC) without unmounting /tsm/isc. As all files of the AC reside on the shared disk, we do not need to install it on the second node. To install AC, we follow these steps:

1. We access a console and log in as root.
2. We change the directory to cdrom directory. We are installing with TSMAdminCenter5300.tar, so we issue the following command:

   ```
   tar -xf TSMAdminCenter5300.tar
   ```

3. We can run one of the following commands to install the Administration Center:
   - For InstallShield wizard, we install:

     ```
     ./startInstall.sh
     ```

   - For console wizard, we iinstall:

     ```
     ./startInstall.sh -console
     ```

   - For silent install, we run the following command on a single line:

     ```
     ./startInstall.sh -silent -W AdminNamePanel.adminName="<user name>"
     -W PasswordInput.adminPass="<user password>"
     -W PasswordInput.verifyPass="<user password>"
     -W MediaLocationInput.installMediaLocation="<media location>"
     -W PortInput.webAdminPort="<web administration port>"
     -P AdminCenterDeploy.installLocation="<install location>"
     ```

     If we do not provide all parameters, default values will be used.

   We install Administration Center with the following command:

   ```
   [root@lochness tsm-admincenter]# ./startInstall.sh -silent \
   -W AdminNamePanel.adminName="iscadmin" \
   -W PasswordInput.adminPass="itsosj" \
   -W PasswordInput.verifyPass="itsosj" \
   -P ISCProduct.installLocation="/tsm/isc/"
   Running setupACLinux ...
   [root@lochness tsm-admincenter]#
   ```

Archived

Now that we have finished the installation of both ISC and AC, we stop ISC and unmount the shared filesystem /tsm/isc as shown in Example 13-2.

Example 13-2 Stop Integrated Solutions Console and Administration Center

```
[root@diomede root]# /tsm/isc/PortalServer/bin/stopISC.sh ISC_Portal ISCUSER
ISCPASS
ADMU0116I: Tool information is being logged in file
           /tsm/isc/AppServer/logs/ISC_Portal/stopServer.log
ADMU3100I: Reading configuration for server: ISC_Portal
ADMU3201I: Server stop request issued. Waiting for stop status.
ADMU4000I: Server ISC_Portal stop completed.
[root@diomede root]# umount /tsm/isc
[root@diomede root]#
```

Note: All files of the AC reside on the shared disk. We do not need to install it on the second node.

13.5 Configuration

In this section we describe preparation of shared storage disks, configuration of the Tivoli Storage Manager server, and the creation of necessary cluster resources.

13.5.1 Preparing shared storage

We need seven logical drives in our cluster configuration:

- **LUN 0:** Tie breaker disk for Tivoli System Automation for Multiplatforms quorum (RAID 1 protected).
- **LUN 1 and 2:** Disks for Tivoli Storage Manager database (RAID 0, because Tivoli Storage Manager mirrors the database volumes).
- **LUN 3 and 4:** Disks for Tivoli Storage Manager log (RAID 0, because Tivoli Storage Manager mirrors the log volumes).
- **LUN 5:** Disk for Tivoli Storage Manager disk storage pool (RAID 1 protected) and Tivoli Storage Manager server configuration and log files. The configuration and log files will be on a separate partition apart from the disk storage pool partition on this LUN. We could also use an additional LUN for the configuration and log files.
- **LUN 6:** Disk for Tivoli Storage Manager Administration Center.

Archived

Figure 13-1 shows the logical drive mapping of our configuration.

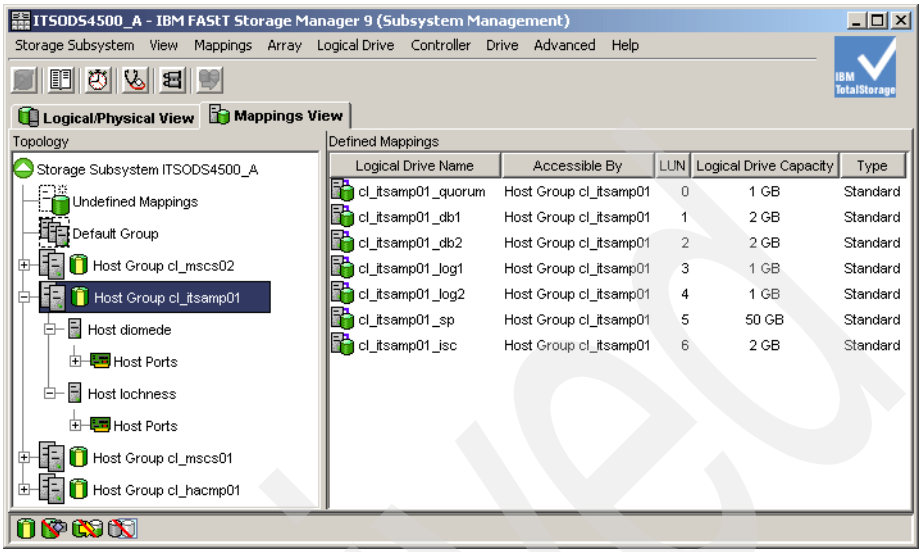

Figure 13-1 Logical drive mapping for cluster volumes

13.5.2 Tivoli Storage Manager Server configuration

In this section we describe the necessary steps to configure the Tivoli Storage Manager server.

Setting up shared disks and cleaning up default installation

Tivoli System Automation requires entries for all managed file systems in /etc/fstab. Example 13-3 shows the necessary entries for the Tivoli Storage Manager server. We create all mount points and insert these entries to /etc/fstab on both nodes.

Example 13-3 Necessary entries in /etc/fstab for the Tivoli Storage Manager server

```
/dev/tsmdb1     /tsm/db1        ext3    noauto  0 0
/dev/tsmdb1mr   /tsm/db1mr      ext3    noauto  0 0
/dev/tsmlg1     /tsm/lg1        ext3    noauto  0 0
/dev/tsmlg1mr   /tsm/lg1mr      ext3    noauto  0 0
/dev/tsmdp      /tsm/dp         ext3    noauto  0 0
/dev/tsmfiles   /tsm/files      ext3    noauto  0 0
```

To set up the database, log, and storage pool volumes, we manually mount all necessary file systems on our first node, diomede.

Archived

Attention: *Never* mount file systems of a shared disk concurrently on both nodes unless you use a shared disk file system. Doing so destroys the file system, and probably all data of the file system will be lost. If you need a file system concurrently on multiple nodes, use a shared disk file system like the IBM General Parallel File System (GPFS).

We clean up the default server installation files which are not required on both nodes as shown in Example 13-4. We remove the default created database, recovery log, space management, archive, and backup pool files.

Example 13-4 Cleaning up the default server installation

```
[root@diomede root]# cd /opt/tivoli/tsm/server/bin
[root@diomede bin]# rm db.dsm
[root@diomede bin]# rm spcmgmt.dsm
[root@diomede bin]# rm log.dsm
[root@diomede bin]# rm backup.dsm
[root@diomede bin]# rm archive.dsm
```

Server instance configuration

To configure the clustered Tivoli Storage Manager server, we follow these steps:

1. We create the dsmserv.opt configuration file and ensure that we use the TCP/IP communication method. Example 13-5 shows the appropriate content of /tsm/files/dsmserv.opt.

Example 13-5 Contents of /tsm/files/dsmserv.opt

```
*** IBM TSM Server options file
*** Refer to dsmserv.opt.smp for other options
COMMMETHOD TCPIP
TCPPORT 1500

DEVCONFIG devcnfg.out
```

2. Then we configure the local client to communicate with the server for the Tivoli Storage Manager command line administrative interface. Example 13-6 shows the stanza in /opt/tivoli/tsm/client/ba/bin/dsm.sys. We configure dsm.sys on both nodes.

Example 13-6 Server stanza in dsm.sys to enable the use of dsmadmc

```
* Server stanza for admin connection purpose
SErvername tsmsrv05_admin
COMMMethod TCPip
TCPPor 1500
TCPServeraddress 127.0.0.1
```

```
ERRORLOGRETENTION 7
ERRORLOGname /opt/tivoli/tsm/client/ba/bin/dsmerror.log
```

With this setting, we can use **dsmadmc -se=tsmsrv05_admin** to connect to the server.

3. We set up the appropriate Tivoli Storage Manager server directory environment setting for the current shell issuing the commands shown in Example 13-7.

Example 13-7 Setting up necessary environment variables

```
[root@diomede root]# cd /tsm/files
[root@diomede files]# export DSMSERV_CONFIG=./dsmserv.opt
[root@diomede files]# export DSMSERV_DIR=/opt/tivoli/tsm/server/bin
```

For more information about running the server from a directory different from the default database that was created during the server installation, see also the *IBM Tivoli Storage Manager for Linux Installation Guide*.

4. We allocate the Tivoli Storage Manager database, recovery log, and storage pools on the shared Tivoli Storage Manager volume group. To accomplish this, we will use the **dsmfmt command** to format database, log, and disk storage pools files on the shared file systems as shown in Example 13-8.

Example 13-8 Formatting database, log, and disk storage pools with dsmfmt

```
[root@diomede files]# dsmfmt -m -db /tsm/db1/vol1 500
[root@diomede files]# dsmfmt -m -db /tsm/db1mr/vol1 500
[root@diomede files]# dsmfmt -m -log /tsm/lg1/vol1 250
[root@diomede files]# dsmfmt -m -log /tsm/lg1mr/vol1 250
[root@diomede files]# dsmfmt -m -data /tsm/dp/backvol 25000
```

5. We issue the **dsmserv format command** while we are in the directory /tsm/files to initialize the server database and recovery log:

```
[root@diomede files]# dsmserv format 1 /tsm/lg1/vol1 1 /tsm/db1/vol1
```

This also creates /tsm/files/dsmserv.dsk.

6. Now we start the Tivoli Storage Manager server in the foreground as shown in Example 13-9.

Example 13-9 Starting the server in the foreground

```
[root@diomede files]# pwd
/tsm/files
[root@diomede files]# dsmserv

Tivoli Storage Manager for Linux/i386
Version 5, Release 3, Level 0.0
```

```
Licensed Materials - Property of IBM

(C) Copyright IBM Corporation 1990, 2004.
All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

ANR7800I DSMSERV generated at 05:35:17 on Dec  6 2004.
[...]
TSM:SERVER1>
```

7. We set the servername, mirror database, mirror log, and set the logmode to rollforward as shown in Example 13-10.

Example 13-10 Set up servername, mirror db and log, and set logmode to rollforward

```
TSM:SERVER1> set servername tsmsrv05
TSM:TSMSRV05> define dbcopy /tsm/db1/vol1 /tsm/db1mr/vol1
TSM:TSMSRV05> define logcopy /tsm/lg1/vol1 /tsm/lg1mr/vol1
TSM:TSMSRV05> set logmode rollforward
```

8. We define a DISK storage pool with a volume on the shared filesystem /tsm/dp (RAID1 protected) as shown in Example 13-11.

Example 13-11 Definition of the disk storage pool

```
TSM:TSMSRV05> define stgpool spd_bck disk
TSM:TSMSRV05> define volume spd_bck /tsm/dp/backvol
```

9. We define the tape library and tape drive configurations using the Tivoli Storage Manager server **`define library`**, **`define drive`**, and **`define path`** commands as shown in Example 13-12.

Example 13-12 Definition of library devices

```
TSM:TSMSRV05> define library liblto libtype=scsi shared=yes
TSM:TSMSRV05> define path tsmsrv05 liblto srctype=server desttype=library
device=/dev/IBMchanger0
TSM:TSMSRV05> define drive liblto drlto_1
TSM:TSMSRV05> define drive liblto drlto_2
TSM:TSMSRV05> define path tsmsrv05 drlto_1 srctype=server desttype=drive
library=liblto device=/dev/IBMtape0
TSM:TSMSRV05> define path tsmsrv05 drlto_2 srctype=server desttype=drive
library=liblto device=/dev/IBMtape1
TSM:TSMSRV05> define devclass libltoclass library=liblto devtype=lto
format=drive
```

Archived

10. We register the administrator “admin” with the authority “system” as shown in Example 13-13.

Example 13-13 Registration of TSM administrator

```
TSM:TSMSRV05> register admin admin admin
TSM:TSMSRV05> grant authority admin classes=system
```

We do all other necessary Tivoli Storage Manager configuration steps as we would also do on a normal installation.

13.5.3 Cluster resources for Tivoli Storage Manager Server

A Tivoli Storage Manager Server V5.3 resource group for Tivoli System Automation in Linux typically consists of the following resources:

- Tivoli Storage Manager Server resource
- IP address resource
- Multiple data resources (disks)
- Tape drive and medium changer resource

Requisites for using tape and medium changer devices

Whenever Tivoli Storage Manager Server uses a tape drive or medium changer device, it issues a SCSI RESERVE to the device. Every time a volume is mounted in a tape drive, the SCSI reservation is still present, also when it is in the IDLE status. After Tivoli Storage Manager Server finishes the use of a tape drive or medium changer device, it releases the SCSI reservation.

In a failover situation, it may happen that tape drive and medium changer devices are in use. So the failing node owns SCSI reservations that potentially affect the startup of the Tivoli Storage Manager Server on another node in the cluster.

Tivoli Storage Manager Server for Windows issues a SCSI bus reset during initialization. In a failover situation, the bus reset is expected to clear any SCSI reserves held on the tape devices. Tivoli Storage Manager Server 5.3 for AIX uses the new RESETDRIVES parameter to reset drives. If the RESETDRIVES parameter is set to YES for a library, then the reset will be performed on the library manager for the library and all drives defined to it. Tivoli Storage Manager Server V5.3 for Linux does not issue SCSI resets during initialization. In a Linux Tivoli System Automation environment we use the shell script tsmserverctrl-tape to do this. It utilizes the **sginfo** and **sg_reset** commands to issue SCSI device resets. This breaks the SCSI reservations on the devices.

Archived

Note: The tsmserverctrl-tape script uses the serial number of a device to find the correct /dev/sg* device to reset.

Configuring of the resource group and its resources

To prepare the Tivoli Storage Manager server resource group, we change to the directory /usr/sbin/rsct/sapolicies/tsmserver and copy the sample configuration file:

```
cd /usr/sbin/rsct/sapolicies/tsmserver
cp sa-tsmserver.conf.sample sa-tsmserver.conf
```

We customize the configuration file. Example 13-14 shows the example in our environment. We create a TSM administrator with operator privileges and configure the user id (TSM_USER) and the password (TSM_PASS) in the configuration file. TSM_SRV is the name of the server stanza in dsm.sys.

Note: If you run multiple Tivoli Storage Manager servers in your cluster, we suggest to create an extra directory below /usr/sbin/rsct/sapolicies for every Tivoli Storage Manager server that you run. For a second server, create for example the directory /usr/sbin/rsct/sapolicies/tsmserver2. Copy the files cfgtsmserver and sa-tsmserver.conf.sample to this directory. Rename sa-tsmserver.conf.sample to sa-tsmserver2.conf. Then you can configure this second server in the same way as the first one. Be sure to use different values for the prefix variable in the Tivoli System Automation configuration file for each server.

Example 13-14 Extract of the configuration file sa-tsmserver.conf

```
###### START OF CUSTOMIZABLE AREA ##########################################
#
# set default values
TSMSERVER_EXEC_DIR="/tsm/files"
TSMSERVER_OPT="/tsm/files/dsmserv.opt"
TSM_SRV="tsmsrv05_admin"
TSM_USER="scriptoperator"
TSM_PASS="password"

# --directory for control scripts
script_dir="/usr/sbin/rsct/sapolicies/tsmserver"

# --prefix of all TSM server resources
prefix="SA-tsmserver-"

# --list of nodes in the TSM server cluster
nodes="diomede lochness"
```

```
# --IP address and netmask for TSM server
ip_1="9.1.39.54,255.255.255.0"

# --List of network interfaces ServiceIP ip_x depends on.
#   Entries are lists of the form <network-interface-name>:<node-name>,...
nieq_1="eth0:diomede,eth0:lochness"

# --common local mountpoint for shared data
#   If more instances of <data_>, add more rows, like: data_tmp, data_proj...
#   Note: the keywords need to be unique!
data_db1="/tsm/db1"
data_db1mr="/tsm/db1mr"
data_lg1="/tsm/lg1"
data_lg1mr="/tsm/lg1mr"
data_dp="/tsm/dp"
data_files="/tsm/files"

# --serial numbers of tape units and medium changer devices
#   entries are separated with a ','
tapes="1110176223,1110177214,0000013108231000"

###### END OF CUSTOMIZABLE AREA ##############################################
```

Note: To find out the serial numbers of the tape and medium changer devices, we use the device information in the /proc file system as shown in Example 12-6 on page 605.

We verify the serial numbers of tape and medium changer devices with the **sginfo** command as shown in Example 13-15.

Example 13-15 Verification of tape and medium changer serial numbers with sginfo

```
[root@diomede root]# sginfo -s /dev/sg0
Serial Number '1110176223'

[root@diomede root]# sginfo -s /dev/sg1
Serial Number '0000013108231000'

[root@diomede root]# sginfo -s /dev/sg2
Serial Number '1110177214'

[root@diomede root]#
```

We execute the command **./cfgtsmserver** to create the necessary definition files (*.def) for Tivoli System Automation. The script **SA-tsmserver-make** which adds the resource group, resources, resource group members, equivalency, and

relationships to Tivoli System Automation is also generated by **cfgtsmserver**. Example 13-16 shows the abbreviated output.

Example 13-16 Execution of cfgtsmserver to create definition files

```
[root@diomede tsmserver]# ./cfgtsmserver
[...]
Generated resource definitions in: 'SA-tsmserver-*.def'
and commands in script: 'SA-tsmserver-make'.

Use script: 'SA-tsmserver-make' to remove and create resources based on
'SA-tsmserver-*.def' files.

[root@diomede tsmserver]# ./SA-tsmserver-make
successfully performed: 'mkrg SA-tsmserver-rg'
successfully performed: 'mkrsrc -f SA-tsmserver-server.def IBM.Application'
[...]
[root@diomede tsmserver]# ls -l *def SA-tsmserver-make
-rw-r--r--    1 root   root      483 Feb  2 08:51 SA-tsmserver-data-db1.def
-rw-r--r--    1 root   root      491 Feb  2 08:51 SA-tsmserver-data-db1mr.def
-rw-r--r--    1 root   root      479 Feb  2 08:51 SA-tsmserver-data-dp.def
-rw-r--r--    1 root   root      483 Feb  2 08:51 SA-tsmserver-data-lg1.def
-rw-r--r--    1 root   root      491 Feb  2 08:51 SA-tsmserver-data-lg1mr.def
-rw-r--r--    1 root   root      164 Feb  2 08:51 SA-tsmserver-ip-1.def
-rwx------    1 root   root    12399 Feb  2 08:51 SA-tsmserver-make
-rw-r--r--    1 root   root      586 Feb  2 08:51 SA-tsmserver-server.def
-rw-r--r--    1 root   root      611 Feb  2 08:51 SA-tsmserver-tape.def
[root@diomede tsmserver]#
```

We execute **./SA-tsmserver-make** to create the resource group and all necessary resources, equivalencies, and relationships as shown in Example 13-17.

Example 13-17 Executing the SA-tsmserver-make script

```
[root@diomede tsmserver]# ./SA-tsmserver-make
successfully performed: 'mkrg SA-tsmserver-rg'
successfully performed: 'mkrsrc -f SA-tsmserver-server.def IBM.Application'
successfully performed: 'addrgmbr -m T -g SA-tsmserver-rg
IBM.Application:SA-tsmserver-server'
successfully performed: 'mkrsrc -f SA-tsmserver-tape.def IBM.Application'
successfully performed: 'addrgmbr -m T -g SA-tsmserver-rg
IBM.Application:SA-tsmserver-tape'
successfully performed: 'mkrel -S IBM.Application:SA-tsmserver-server -G
IBM.Application:SA-tsmserver-tape -p DependsOn SA-tsmserver-server-on-tape'
[...]
[root@diomede tsmserver]#
```

Important: Depending on our needs, we can edit the tsmserverctrl-tape script to change its behavior during startup. The value of the returnAlwaysStartOK variable within the tsmserverctrl-tape script is set to 1. This means the script exits with return code 0 on every start operation, even when some SCSI resets are not successful. Tivoli System Automation recognizes the SA-tsmserver-tape resource as online and then starts the Tivoli Storage Manager Server. This is often appropriate, especially when big disk storage pools are used.

In other environments that use primarily tape storage pools we can change the value of returnAlwaysStartOK to 0. If a tape drive is unavailable on the node, a SCSI reset of drive will fail, and the script exits with return code 1. Tivoli System Automation can then try to bring the resource group online on another node, which might be able to access all tape devices. When we configure returnAlwaysStartOK to 0 we must be aware that the complete outage of a tape drive makes the successful start of the tsmserverctrl-tape script impossible until the tape drive is accessible again.

13.5.4 Cluster resources for Administration Center

We show how to set up the Administration Center as a highly available resource in the Tivoli System Automation cluster.

Important: Although our tests to run the AC in the cluster were successful, the AC is currently not supported in a clustered environment.

To prepare the Tivoli Storage Manager server resource, we change to the directory /usr/sbin/rsct/sapolicies/tsmadminc and copy the sample configuration file:

```
cd /usr/sbin/rsct/sapolicies/tsmadminc
cp sa-tsmadminc.conf.sample sa-tsmadminc.conf
```

We customize the configuration file. Example 13-18 shows the example in our environment.

Example 13-18 Extract of the configuration file sa-tsmadmin.conf

```
###### START OF CUSTOMIZABLE AREA ##############################################
#
# set default values
TSM_ADMINC_DIR="/tsm/isc"

# --directory for control scripts
script_dir="/usr/sbin/rsct/sapolicies/tsmadminc"
```

```
# --prefix of all TSM server resources
prefix="SA-tsmadminc-"

# --list of nodes in the TSM server cluster
nodes="lochness diomede"

# --IP address and netmask for TSM server
ip_1="9.1.39.69,255.255.255.0"

# --List of network interfaces ServiceIP ip_x depends on.
#   Entries are lists of the form <network-interface-name>:<node-name>,...
nieq_1="eth0:lochness,eth0:diomede"

# --common local mountpoint for shared data
#   If more instances of <data_>, add more rows, like: data_tmp, data_proj...
#   Note: the keywords need to be unique!
data_isc="/tsm/isc"

###### END OF CUSTOMIZABLE AREA ##############################################
```

Note: Compared to the configuration file of the Tivoli Storage Manager Server, we change the order of the nodes in the variables, *nodes* and *nieq_1*. During the first startup of a resource group, Tivoli System Automation tries to start the resources on the first node configured in the nodes variable if no relationships to other online resource groups conflict with it.

We execute the command **./cfgtsmadminc** to create the necessary definition files for Tivoli System Automation. Afterwards we use **./SA-tsmadminc-make** to create the resources in Tivoli System Automation. Example 13-19 shows the abbreviated output.

Example 13-19 Execution of cfgtsmadminc to create definition files

```
[root@diomede tsmadminc]# ./cfgtsmadminc
...
Generated resource definitions in: 'SA-tsmadminc-*.def'
and commands in script: 'SA-tsmadminc-make'.

Use script: 'SA-tsmadminc-make' to remove and create resources based on
'SA-tsmadminc-*.def' files.

[root@diomede tsmadminc]# ./SA-tsmadminc-make
successfully performed: 'mkrg SA-tsmadminc-rg'
successfully performed: 'mkrsrc -f SA-tsmadminc-server.def IBM.Application'
...
[root@diomede tsmadminc]#
```

13.5.5 AntiAffinity relationship

We want to ensure that the ISC with the AC do not run at the same node as the Tivoli Storage Manager Server, if possible. Tivoli System Automation provides a way to configure such relationships with the AntiAffinity relationship.
Example 13-20 shows how we create the necessary relationships with the **mkrel** command.

Example 13-20 Configuration of AntiAffinity relationship

```
[root@diomede root]# mkrel -S IBM.ResourceGroup:SA-tsmserver-rg \
-G IBM.ResourceGroup:SA-tsmadminc-rg \
-p AntiAffinity SA-tsmserver-rg-AntiAffinityTo-SA-tsmadminc-rg
[root@diomede root]# mkrel -S IBM.ResourceGroup:SA-tsmadminc-rg \
-G IBM.ResourceGroup:SA-tsmserver-rg \
-p AntiAffinity SA-tsmadminc-rg-AntiAffinityTo-SA-tsmserver-rg
[root@diomede root]#
```

13.6 Bringing the resource groups online

In this section we describe how we verify the configuration and bring the resource groups online.

13.6.1 Verify configuration

Before actually starting resource groups, we verify the Tivoli System Automation configuration. Tivoli System Automation provides several commands for this purpose.

List of resource group and their members

The **lsrg** command lists already defined resource groups and their members. You can find a detailed description of all possible parameters in its manpage. To list the members of resource groups, we execute **lsrg -m**, as shown in Example 13-21.

Example 13-21 Validation of resource group members

```
[root@diomede root]# lsrg -m

Displaying Member Resource information:
Class:Resource:Node[ManagedResource]    Mandatory MemberOf         OpState
IBM.Application:SA-tsmserver-server     True      SA-tsmserver-rg  Offline
IBM.ServiceIP:SA-tsmserver-ip-1         True      SA-tsmserver-rg  Offline
IBM.Application:SA-tsmserver-data-db1   True      SA-tsmserver-rg  Offline
IBM.Application:SA-tsmserver-data-db1mr True      SA-tsmserver-rg  Offline
```


Archived

```
IBM.Application:SA-tsmserver-data-lg1   True      SA-tsmserver-rg  Offline
IBM.Application:SA-tsmserver-data-lg1mr True      SA-tsmserver-rg  Offline
IBM.Application:SA-tsmserver-data-dp    True      SA-tsmserver-rg  Offline
IBM.Application:SA-tsmserver-tape       True      SA-tsmserver-rg  Offline
IBM.Application:SA-tsmadminc-server     True      SA-tsmadminc-rg  Offline
IBM.ServiceIP:SA-tsmadminc-ip-1         True      SA-tsmadminc-rg  Offline
IBM.Application:SA-tsmadminc-data-isc   True      SA-tsmadminc-rg  Offline
[root@diomede root]#
```

Each resource group has persistent and dynamic attributes. You can use the following parameters to show these attributes of all resource groups:

- **lsrg -A p** displays only *persistent* attributes.
- **lsrg -A d** displays only *dynamic* attributes.
- **lsrg -A b** displays *both persistent and dynamic* attributes.

Example 13-22 shows the output of the **lsrg -A b** command in our environment.

Example 13-22 Persistent and dynamic attributes of all resource groups

```
[root@diomede root]# lsrg -A b

Displaying Resource Group information:
All Attributes

Resource Group 1:
        Name                                = SA-tsmserver-rg
        MemberLocation                      = Collocated
        Priority                            = 0
        AllowedNode                         = ALL
        NominalState                        = Offline
        ExcludedList                        = {}
        ActivePeerDomain                    = cl_itsamp
        OpState                             = Offline
        TopGroup                            = SA-tsmserver-rg
        MoveStatus                          = [None]
        ConfigValidity                      =
        AutomationDetails[CompoundState]    = Satisfactory

Resource Group 2:
        Name                                = SA-tsmadminc-rg
        MemberLocation                      = Collocated
        Priority                            = 0
        AllowedNode                         = ALL
        NominalState                        = Offline
        ExcludedList                        = {}
        ActivePeerDomain                    = cl_itsamp
        OpState                             = Offline
```

Archived

```
        TopGroup                          = SA-tsmadminc-rg
        MoveStatus                        = [None]
        ConfigValidity                    =
        AutomationDetails[CompoundState] = Satisfactory
[root@diomede root]#
```

List relationships

With the **lsrel command** you can list already-defined managed relationship and their attributes. Example 13-23 shows the relationships created during execution of the SA-tsmserver-make and SA-tsmadminc-make scripts.

Example 13-23 Output of the lsrel command

```
[root@diomede root]# lsrel

Displaying Managed Relations :

Name                              Class:Resource:Node[Source]         ResourceGroup[Source]
SA-tsmserver-server-on-data-db1mr IBM.Application:SA-tsmserver-server SA-tsmserver-rg
SA-tsmserver-server-on-data-db1   IBM.Application:SA-tsmserver-server SA-tsmserver-rg
SA-tsmserver-server-on-data-lg1mr IBM.Application:SA-tsmserver-server SA-tsmserver-rg
SA-tsmserver-server-on-data-lg1   IBM.Application:SA-tsmserver-server SA-tsmserver-rg
SA-tsmserver-server-on-data-dp    IBM.Application:SA-tsmserver-server SA-tsmserver-rg
SA-tsmserver-server-on-data-files IBM.Application:SA-tsmserver-server SA-tsmserver-rg
SA-tsmserver-server-on-tape       IBM.Application:SA-tsmserver-server SA-tsmserver-rg
SA-tsmserver-server-on-ip-1       IBM.Application:SA-tsmserver-server SA-tsmserver-rg
SA-tsmserver-ip-on-nieq-1         IBM.ServiceIP:SA-tsmserver-ip-1     SA-tsmserver-rg
SA-tsmadminc-server-on-data-isc   IBM.Application:SA-tsmadminc-server SA-tsmadminc-rg
SA-tsmadminc-server-on-ip-1       IBM.Application:SA-tsmadminc-server SA-tsmadminc-rg
SA-tsmadminc-ip-on-nieq-1         IBM.ServiceIP:SA-tsmadminc-ip-1     SA-tsmadminc-rg
[root@diomede root]#
```

The **lsrel** command also provides some parameters to view persistent and dynamic attributes of a relationship. You can find a detailed description in its manpage.

13.6.2 Bringing Tivoli Storage Manager Server resource group online

We use the **chrg** command to change persistent attribute values of one or more resource groups, including starting and stopping resource groups.

The -o flag specifies the nominal state of the resource group, which can be online or offline. Example 13-24 shows how we change the nominal state of the resource group SA-tsmserver-rg to online and view the result after a few seconds with the **lsrg** command.

Example 13-24 Changing the nominal state of the SA-tsmserver-rg to online

```
[root@diomede root]# chrg -o online SA-tsmserver-rg
[root@diomede root]# lsrg -m

Displaying Member Resource information:
Class:Resource:Node[ManagedResource]    Mandatory MemberOf         OpState
IBM.Application:SA-tsmserver-server     True      SA-tsmserver-rg  Online
IBM.ServiceIP:SA-tsmserver-ip-1         True      SA-tsmserver-rg  Online
IBM.Application:SA-tsmserver-data-db1   True      SA-tsmserver-rg  Online
IBM.Application:SA-tsmserver-data-db1mr True      SA-tsmserver-rg  Online
IBM.Application:SA-tsmserver-data-lg1   True      SA-tsmserver-rg  Online
IBM.Application:SA-tsmserver-data-lg1mr True      SA-tsmserver-rg  Online
IBM.Application:SA-tsmserver-data-dp    True      SA-tsmserver-rg  Online
IBM.Application:SA-tsmserver-tape       True      SA-tsmserver-rg  Online
IBM.Application:SA-tsmadminc-server     True      SA-tsmadminc-rg  Offline
IBM.ServiceIP:SA-tsmadminc-ip-1         True      SA-tsmadminc-rg  Offline
IBM.Application:SA-tsmadminc-data-isc   True      SA-tsmadminc-rg  Offline
[root@diomede root]#
```

To find out on which node a resource is actually online, we use the **getstatus** script as shown in Example 13-25.

Example 13-25 Output of the getstatus script

```
[root@diomede root]# /usr/sbin/rsct/sapolicies/bin/getstatus
[...]
-- Resources --

                Resource Name                Node Name              State
                -------------                ---------              -----
            SA-tsmserver-server                diomede             Online
            SA-tsmserver-server               lochness            Offline
                              -                      -                  -
              SA-tsmserver-tape                diomede             Online
              SA-tsmserver-tape               lochness            Offline
                              -                      -                  -
              SA-tsmserver-ip-1                diomede             Online
              SA-tsmserver-ip-1               lochness            Offline
                              -                      -                  -
[...]
[root@diomede root]#
```

Now we know that the Tivoli Storage Manager Server runs at the node diomede.

Archived

13.6.3 Bringing Administration Center resource group online

We use again the **chrg** command to bring the Administration Center resource group online. Example 13-26 shows how we change the nominal state of the resource group SA-tsmadminc-rg to online and view the result after a while with the **lsrg** command as shown in Example 13-26.

Example 13-26 Changing the nominal state of the SA-tsmadminc-rg to online

```
[root@diomede root]# chrg -o online SA-tsmadminc-rg
[root@diomede root]# lsrg -m

Displaying Member Resource information:
Class:Resource:Node[ManagedResource]    Mandatory MemberOf         OpState
IBM.Application:SA-tsmserver-server     True      SA-tsmserver-rg  Online
IBM.ServiceIP:SA-tsmserver-ip-1         True      SA-tsmserver-rg  Online
IBM.Application:SA-tsmserver-data-db1   True      SA-tsmserver-rg  Online
IBM.Application:SA-tsmserver-data-db1mr True      SA-tsmserver-rg  Online
IBM.Application:SA-tsmserver-data-lg1   True      SA-tsmserver-rg  Online
IBM.Application:SA-tsmserver-data-lg1mr True      SA-tsmserver-rg  Online
IBM.Application:SA-tsmserver-data-dp    True      SA-tsmserver-rg  Online
IBM.Application:SA-tsmserver-tape       True      SA-tsmserver-rg  Online
IBM.Application:SA-tsmadminc-server     True      SA-tsmadminc-rg  Online
IBM.ServiceIP:SA-tsmadminc-ip-1         True      SA-tsmadminc-rg  Online
IBM.Application:SA-tsmadminc-data-isc   True      SA-tsmadminc-rg  Online
[root@diomede root]#
```

13.7 Testing the cluster

In order to check the high availability of Tivoli Storage Manager server on our lab environment, we must do some testing.

Our objective with these tests is showing how Tivoli Storage Manager on a clustered environment can respond after certain kinds of failures that affect the shared resources.

We use the Windows 2000 Backup/Archive Client 5.3.0.0 for this test. The client runs on an independent Windows 2000 workstation.

13.7.1 Testing client incremental backup using the GUI

In our first test we the Tivoli Storage Manager GUI to start an incremental backup.

Archived

Objective

The objective of this test is showing what happens when a client incremental backup is started from the Tivoli Storage Manager GUI and suddenly the node which hosts the Tivoli Storage Manager server fails.

We perform these tasks:

1. We start an incremental client backup using the GUI. We select the local drives and the System Object as shown in Figure 13-2.

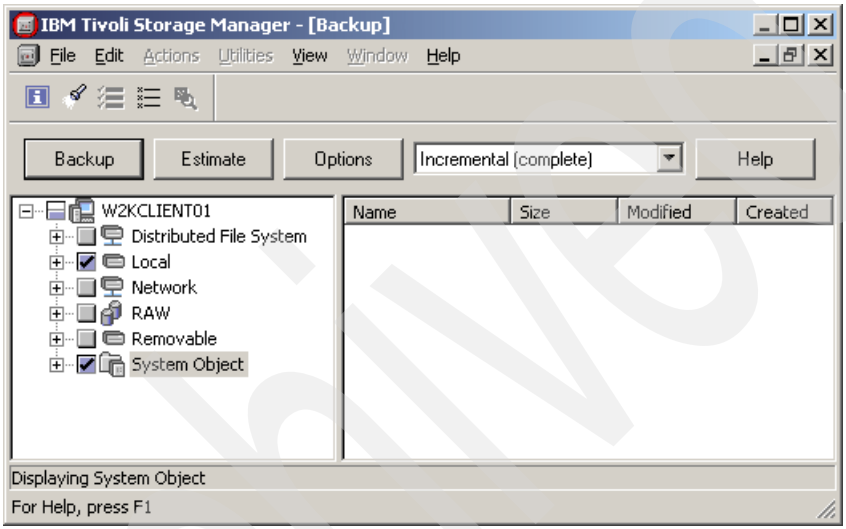

Figure 13-2 Selecting client backup using the GUI

2. Transfer of files starts as we can see in Figure 13-3.

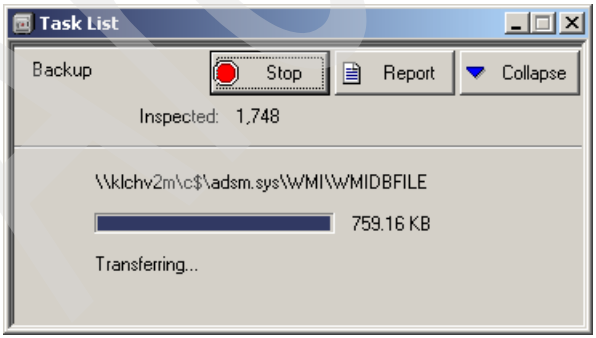

Figure 13-3 Transfer of files starts

3. While the client is transferring files to the server we unplug all power cables from the first node, diomede. On the client, backup is halted and a reopening session message is received on the GUI as shown in Figure 13-4.

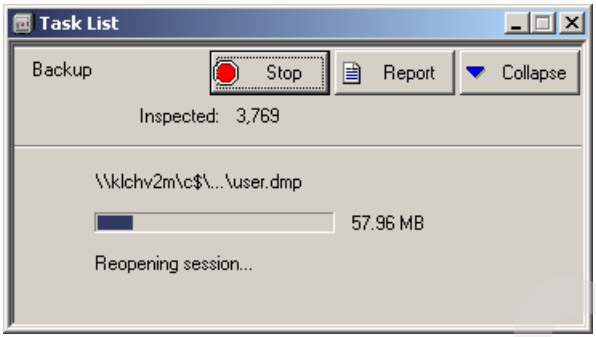

Figure 13-4 Reopening Session

4. The outage causes an automatic failover of the SA-tsmserver-rg resource group to the second node, lochness. Example 13-27 shows an extract of /var/log/messages from lochness.

Example 13-27 Log file /var/log/messages after a failover

```
Feb  2 14:36:30 lochness ConfigRM[22155]: (Recorded using libct_ffdc.a cv
2):::Error ID: :::Reference ID:  :::Template ID: 0:::Details File:
:::Location: RSCT,PeerDomain.C,1.99.7.3,15142
:::CONFIGRM_PENDINGQUORUM_ER The operational quorum state of the active peer
domain has changed to PENDING_QUORUM.  This state usually indicates that
exactly half of the nodes that are defined in the  peer domain are online.  In
this state cluster resources cannot be recovered although  none will be stopped
explicitly.
Feb  2 14:36:30 lochness RecoveryRM[22214]: (Recorded using libct_ffdc.a cv
2):::Error ID: 825....iLJ.O/pAO/72k7bO....................:::Reference ID:
:::Template ID: 0:::Details File: :::Location: RSCT,Protocol.C,1.55,2171
:::RECOVERYRM_INFO_4_ST A  member has left. Node number =  1
Feb  2 14:36:32 lochness ConfigRM[22153]: (Recorded using libct_ffdc.a cv
2):::Error ID: :::Reference ID:  :::Template ID: 0:::Details File:
:::Location: RSCT,PeerDomain.C,1.99.7.3,15138
:::CONFIGRM_HASQUORUM_ST The operational quorum state of the active peer domain
has changed to HAS_QUORUM.  In this state, cluster resources may be recovered
and controlled as needed by  management applications.
[...]
Feb  2 14:36:45 lochness
/usr/sbin/rsct/sapolicies/tsmserver/tsmserverctrl-server:[2149]: ITSAMP: TSM
server started
```

5. Now that the Tivoli Storage Manager server is restarted on lochness, the client backup goes on transferring the data as shown in Figure 13-5.

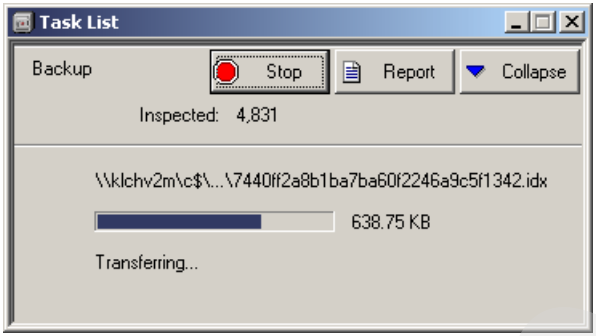

Figure 13-5 Transferring of files continues to the second node

6. Client backup ends successfully.

The result of the test shows that when you start a backup from a client and there is a failure that forces Tivoli Storage Manager server to fail, backup is halted, and when the server is up again, the client reopens a session with the server and continues transferring data.

Note: In the test we have just described, we used the disk storage pool as the destination storage pool. We also tested using a tape storage pool as the destination and we got the same results. The only difference is that when the Tivoli Storage Manager server is up again, the tape volume it used on the other node is unloaded and loaded again into the drive. The client receives a message, `Waiting for media...` while this process takes place. After the tape volume is mounted again, the backup continues and ends successfully.

13.7.2 Testing a scheduled client backup

The second test consists of a scheduled backup.

Objective

The objective of this test is to show what happens when a scheduled client backup is running and suddenly the node which hosts the Tivoli Storage Manager server fails.

Activities

We perform these tasks:

1. We schedule a client incremental backup operation using the Tivoli Storage Manager server scheduler and we associate the schedule to W2KCLIENT01 nodename.
2. At the scheduled time, a client session starts from W2KCLIENT01 as shown in Example 13-28.

Example 13-28 Activity log when the client starts a scheduled backup

```
02/09/2005 16:10:01  ANR2561I Schedule prompter contacting W2KCLIENT01 (session
                      17) to start a scheduled operation. (SESSION: 17)
02/09/2005 16:10:03  ANR8214E Session terminated when no data was read on
                      socket 14. (SESSION: 17)
02/09/2005 16:10:03  ANR0403I Session 17 ended for node W2KCLIENT01 ().
                      (SESSION: 17)
02/09/2005 16:10:03  ANR0406I Session 18 started for node W2KCLIENT01 (WinNT)
                      (Tcp/Ip dhcp38057.almaden.ibm.com(1565)).
```

3. The client starts sending files to the server as shown in Example 13-29.

Example 13-29 Schedule log file showing the start of the backup on the client

```
Executing scheduled command now.
02/09/2005 16:10:01 --- SCHEDULEREC OBJECT BEGIN SCHEDULE_1 02/09/2005 16:10:00
02/09/2005 16:10:01 Incremental backup of volume ‘\\klchv2m\c$’
02/09/2005 16:10:01 Incremental backup of volume ‘SYSTEMOBJECT’
[...]
02/09/2005 16:10:03 Directory-->  0 \\klchv2m\c$\ [Sent]
02/09/2005 16:10:03 Directory-->  0 \\klchv2m\c$\Downloads [Sent]
```

4. While the client continues sending files to the server, we force diomede to fail through a short power outage. The following sequence occurs:
   a. In the client, backup is halted and an error is received as shown in Example 13-30.

Example 13-30 Error log file when the client looses the session

```
02/09/2005 16:11:36 sessSendVerb: Error sending Verb, rc: -50
02/09/2005 16:11:36 ANS1809W Session is lost; initializing session reopen
                     procedure.
02/09/2005 16:11:37 ANS1809W Session is lost; initializing session reopen
                     procedure.
```

   b. As soon as the Tivoli Storage Manager server resource group is online on the other node, client backup restarts against the disk storage pool as shown on the schedule log file in Example 13-31.

Archived

Example 13-31 Schedule log file when backup restarts on the client

```
[...]
02/09/2005 16:11:37 Normal File-->       649,392,128
\\klchv2m\c$\Downloads\RHEL3-U2\rhel-3-U2-i386-as-disc2.iso  ** Unsuccessful **
02/09/2005 16:11:37 ANS1809W Session is lost; initializing session reopen
procedure.
02/09/2005 16:11:52 ... successful
02/09/2005 16:12:49 Retry # 1  Normal File-->       649,392,128
\\klchv2m\c$\Downloads\RHEL3-U2\rhel-3-U2-i386-as-disc2.iso [Sent]
02/09/2005 16:13:50 Normal File-->       664,571,904
\\klchv2m\c$\Downloads\RHEL3-U2\rhel-3-U2-i386-as-disc3.iso [Sent]
02/09/2005 16:14:06 Normal File-->       176,574,464
\\klchv2m\c$\Downloads\RHEL3-U2\rhel-3-U2-i386-as-disc4.iso [Sent]
[...]
```

c. The messages shown in Example 13-32 are received on the Tivoli Storage Manager server activity log after restarting.

Example 13-32 Activity log after the server is restarted

```
02/09/2005 16:11:52  ANR0406I Session 1 started for node W2KCLIENT01 (WinNT)
                      (Tcp/Ip dhcp38057.almaden.ibm.com(1585)).
[...]
02/09/2005 16:16:07  ANE4961I (Session: 1, Node: W2KCLIENT01)  Total number of
                      bytes transferred:  3.06 GB
[...]
02/09/2005 16:16:07  ANR2507I Schedule SCHEDULE_1 for domain STANDARD started
                      at 02/09/2005 04:10:00 PM for node W2KCLIENT01 completed
                      successfully at 02/09/2005 04:16:07 PM.
02/09/2005 16:16:07  ANR0403I Session 1 ended for node W2KCLIENT01 (WinNT).
```

5. Example 13-33 shows the final status of the schedule in the schedule log.

Example 13-33 Schedule log file showing backup statistics on the client

```
02/09/2005 16:16:06 --- SCHEDULEREC STATUS BEGIN
02/09/2005 16:16:06 Total number of objects inspected:    1,940
02/09/2005 16:16:06 Total number of objects backed up:    1,861
02/09/2005 16:16:06 Total number of objects updated:          0
02/09/2005 16:16:06 Total number of objects rebound:          0
02/09/2005 16:16:06 Total number of objects deleted:          0
02/09/2005 16:16:06 Total number of objects expired:          0
02/09/2005 16:16:06 Total number of objects failed:           0
02/09/2005 16:16:06 Total number of bytes transferred:     3.06 GB
02/09/2005 16:16:06 Data transfer time:                  280.23 sec
02/09/2005 16:16:06 Network data transfer rate:        11,478.49 KB/sec
02/09/2005 16:16:06 Aggregate data transfer rate:      8,803.01 KB/sec
02/09/2005 16:16:06 Objects compressed by:                    0%
02/09/2005 16:16:06 Elapsed processing time:           00:06:05
```

```
02/09/2005 16:16:06 --- SCHEDULEREC STATUS END
02/09/2005 16:16:06 --- SCHEDULEREC OBJECT END SCHEDULE_1 02/09/2005 16:10:00
02/09/2005 16:16:06 Scheduled event 'SCHEDULE_1' completed successfully.
02/09/2005 16:16:06 Sending results for scheduled event 'SCHEDULE_1'.
02/09/2005 16:16:06 Results sent to server for scheduled event 'SCHEDULE_1'.
```

Note: Depending on how long the failover process takes, we may get these error messages in dsmerror.log: `ANS5216E Could not establish a TCP/IP connection` and `ANS4039E Could not establish a session with a Tivoli Storage Manager server or client agent`). If this happens, although Tivoli Storage Manager reports in the schedule log file that the scheduled event failed with return code 12, in fact, the backup ended successfully in our tests.

Results summary

The test results show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a scheduled backup started from one client is restarted after the failover.

Note: In the test we have just described, we used the disk storage pool as the destination storage pool. We also tested using a tape storage pool as the destination and we got the same results. The only difference is that when the Tivoli Storage Manager server is up again, the tape volume it used on the other node is unloaded and loaded again into the drive. The client logs the message, `ANS1114I Waiting for mount of offline media`. in its dsmsched.log while this process takes place. After the tape volume is mounted again, the backup continues and ends successfully.

13.7.3 Testing migration from disk storage pool to tape storage pool

Our third test is a server process: migration from disk storage pool to tape storage pool.

Objective

The objective of this test is showing what happens when a disk storage pool migration process is started on the Tivoli Storage Manager server and the node that hosts the server instance fails.

Activities

For this test, we perform these tasks:

1. We use the /usr/sbin/rsct/sapolicies/bin/getstatus script to find out that the SA-tsmserver-rg is running on our first node, diomede.

2. We update the disk storage pool (SPD_BCK) high threshold migration to 0. This forces migration of data to its next storage pool, a tape storage pool (SPT_BCK). Example 13-34 shows the activity log during update of the disk storage pool and the mounting of a tape volume.

Example 13-34 Disk storage pool migration starting on the first node

```
02/09/2005 12:07:06  ANR2017I Administrator ADMIN issued command: UPDATE
                      STGPOOL SPD_BCK HIGHMIG=0 LOWMIG=0
02/09/2005 12:07:06  ANR2202I Storage pool SPD_BCK updated.
02/09/2005 12:07:06  ANR0984I Process 4 for MIGRATION started in the BACKGROUND
                      at 12:07:06 PM. (PROCESS: 4)
02/09/2005 12:07:06  ANR1000I Migration process 4 started for storage pool
                      SPD_BCK automatically, highMig=0, lowMig=0, duration=No.
                      (PROCESS: 4)
02/09/2005 12:07:41  ANR8337I LTO volume 039AKKL2 mounted in drive DRLTO_2
                      (/dev/IBMtape1). (PROCESS: 4)
02/09/2005 12:07:41  ANR0513I Process 4 opened output volume 039AKKL2.
                      (PROCESS: 4)
```

3. While migration is running, we force diomede to fail through a short power outage. The SA-tsmserver-rg resource group is brought online on the second node, lochness. The tape volume is unloaded from the drive. Since the high threshold is still 0, a new migration process is started as shown in Example 13-35.

Example 13-35 Disk storage pool migration starting on the second node

```
02/09/2005 12:09:03  ANR0984I Process 2 for MIGRATION started in the BACKGROUND
                      at 12:09:03 PM. (PROCESS: 2)
02/09/2005 12:09:03  ANR1000I Migration process 2 started for storage pool
                      SPD_BCK automatically, highMig=0, lowMig=0, duration=No.
                      (PROCESS: 2)
02/09/2005 12:09:55  ANR8439I SCSI library LIBLTO is ready for operations.
02/09/2005 12:10:24  ANR8337I LTO volume 039AKKL2 mounted in drive DRLTO_1
                      (/dev/IBMtape0). (PROCESS: 2)
02/09/2005 12:10:24  ANR0513I Process 2 opened output volume 039AKKL2.
                      (PROCESS: 2)
```

Attention: The migration process is not really restarted when the server failover occurs, as you can see comparing the process numbers for migration between Example 13-34 and Example 13-35. But the tape volume is unloaded correctly after the failover and loaded again when the new migration process starts on the server.

4. The migration ends successfully as shown in Example 13-36.

Example 13-36 Disk storage pool migration ends successfully

```
02/09/2005 12:12:30  ANR1001I Migration process 2 ended for storage pool
                      SPD_BCK. (PROCESS: 2)
02/09/2005 12:12:30  ANR0986I Process 2 for MIGRATION running in the BACKGROUND
                      processed 53 items for a total of 2,763,993,088 bytes
                      with a completion state of SUCCESS at 12:12:30 PM.
                      (PROCESS: 2)
```

Results summary

The results of our test show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a migration process that is started on the server before the failure, starts again using a new process number when the second node brings the Tivoli Storage Manager server resource group online.

13.7.4 Testing backup from tape storage pool to copy storage pool

In this section we test another internal server process, backup from a tape storage pool to a copy storage pool.

Objective

The objective of this test is to show what happens when a backup storage pool process (from tape to tape) is started on the Tivoli Storage Manager server and the node that hosts the resource fails.

Activities

For this test, we perform these tasks:

1. We use the /usr/sbin/rsct/sapolicies/bin/getstatus script to find out that the SA-tsmserver-rg is running on our first node, diomede.
2. We run the following command to start an storage pool backup from tape storage pool SPT_BCK to copy storage pool SPCPT_BCK:

   ```
   ba stg spt_bck spcpt_bck
   ```

3. A process starts for the storage pool backup, and Tivoli Storage Manager prompts to mount two tape volumes as shown in the activity log in Example 13-37.

Example 13-37 Starting a backup storage pool process

```
02/10/2005 10:40:13  ANR2017I Administrator ADMIN issued command: BACKUP
                      STGPOOL spt_bck spcpt_bck
02/10/2005 10:40:13  ANR0984I Process 2 for BACKUP STORAGE POOL started in the
                      BACKGROUND at 10:40:13 AM. (PROCESS: 2)
02/10/2005 10:40:13  ANR2110I BACKUP STGPOOL started as process 2. (PROCESS: 2)
```

```
02/10/2005 10:40:13  ANR1210I Backup of primary storage pool SPT_BCK to copy
                      storage pool SPCPT_BCK started as process 2. (PROCESS: 2)
02/10/2005 10:40:13  ANR1228I Removable volume 036AKKL2 is required for storage
                      pool backup. (PROCESS: 2)
[...]
02/10/2005 10:40:43  ANR8337I LTO volume 038AKKL2 mounted in drive DRLTO_1
                      (/dev/IBMtape0). (PROCESS: 2)
02/10/2005 10:40:43  ANR1340I Scratch volume 038AKKL2 is now defined in storage
                      pool SPCPT_BCK. (PROCESS: 2)
02/10/2005 10:40:43  ANR0513I Process 2 opened output volume 038AKKL2.
                      (PROCESS: 2)
02/10/2005 10:41:15  ANR8337I LTO volume 036AKKL2 mounted in drive DRLTO_2
                      (/dev/IBMtape1). (PROCESS: 2)
02/10/2005 10:41:15  ANR0512I Process 2 opened input volume 036AKKL2. (PROCESS:
                      2)
```

4. While the process is started and the two tape volumes are mounted on both drives, we force a short power outage on diomede. The SA-tsmserver-rg resource group is brought online on the second node, lochness. Both tape volumes are unloaded from the drives. The storage pool backup process is not restarted as we can see in Example 13-38.

Example 13-38 After restarting the server the storage pool backup doesn't restart

```
02/10/2005 10:51:21  ANR2100I Activity log process has started.
02/10/2005 10:51:21  ANR4726I The NAS-NDMP support module has been loaded.
[...]
02/10/2005 10:51:21  ANR0993I Server initialization complete.
[...]
02/10/2005 10:52:19  ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                      (SESSION: 2)
02/10/2005 10:52:19  ANR0944E QUERY PROCESS: No active processes found.
                      (SESSION: 2)
[...]
02/10/2005 10:54:10  ANR8439I SCSI library LIBLTO is ready for operations.
```

5. The backup storage pool process does not restart again unless we start it manually. If we do this, Tivoli Storage Manager does not copy again those versions already copied while the process was running before the failover.

   To be sure that the server copied something before the failover, and that starting a new backup for the same primary tape storage pool will copy the rest of the files on the copy storage pool, we use the following tips:

   – We run the following Tivoli Storage Manager command:

     ```
     q content 038AKKL2
     ```

     We do this to check that there is something copied onto the volume that was used by Tivoli Storage Manager for the copy storage pool.

– We run the backup storage pool command again:

```
ba stg spt_bck spcpt_bck
```

– When the backup ends, we use the following commands:

```
q occu stg=spt_bck
q occu stg=spcpt_bck
```

If backup versions were migrated from disk storage pool to tape storage pool both commands should report the same information.

Results summary

The results of our test show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a backup storage pool process (from tape to tape) started on the server before the failure, does not restart when the second node brings the Tivoli Storage Manager server instance online.

Both tapes are correctly unloaded from tape drives when the Tivoli Storage Manager server is again online, but the process is not restarted unless we run the command again.

There is no difference between a scheduled process or a manual process using the administrative interface.

13.7.5 Testing server database backup

The following test consists of backing up the server database.

Objective

The objective of this test is to show what happens when a Tivoli Storage Manager server database backup process is started on the Tivoli Storage Manager server and the node that hosts the resource fails.

Activities

For this test, we perform these tasks:

1. We use the /usr/sbin/rsct/sapolicies/bin/getstatus script to find out that the SA-tsmserver-rg is running on our first node, diomede.
2. We run the following command to start a full database backup:

```
backup db t=full devc=LIBLTOCLASS
```

3. A process starts for database backup and Tivoli Storage Manager prompts to mount a scratch tape volume as shown in the activity log in Example 13-39.

Example 13-39 Starting a database backup on the server

```
02/10/2005 14:16:43  ANR2017I Administrator ADMIN issued command: BACKUP DB
                      t=full devc=LIBLTOCLASS  (SESSION: 5)
02/10/2005 14:16:43  ANR0984I Process 3 for DATABASE BACKUP started in the
                        BACKGROUND at 02:16:43 PM. (SESSION: 5, PROCESS: 3)
02/10/2005 14:16:43  ANR2280I Full database backup started as process 3.
                        (SESSION: 5, PROCESS: 3)
02/10/2005 14:17:14  ANR8337I LTO volume 037AKKL2 mounted in drive DRLTO_2
                      (/dev/IBMtape1). (SESSION: 5, PROCESS: 3)
02/10/2005 14:17:14  ANR0513I Process 3 opened output volume 037AKKL2.
                      (SESSION: 5, PROCESS: 3)
02/10/2005 14:17:17  ANR1360I Output volume 037AKKL2 opened (sequence number
                      1). (SESSION: 5, PROCESS: 3)
02/10/2005 14:17:18  ANR4554I Backed up 10496 of 20996 database pages.
                      (SESSION: 5, PROCESS: 3)
```

4. While the process is started and the two tape volumes are mounted on both drives, we force a failure on diomede. The SA-tsmserver-rg resource group is brought online on the second node, lochness. The tape volumes is unloaded from the drive. The database backup process is not restarted, as we can see in the activity log in Example 13-40.

Example 13-40 After the server is restarted database backup does not restart

```
02/10/2005 14:21:04  ANR2100I Activity log process has started.
02/10/2005 14:21:04  ANR4726I The NAS-NDMP support module has been loaded.
[...]
02/10/2005 14:21:04  ANR0993I Server initialization complete.
[...]
02/10/2005 14:22:03  ANR8439I SCSI library LIBLTO is ready for operations.
[...]
02/10/2005 14:23:19  ANR2017I Administrator ADMIN issued command: QUERY PROCESS

02/10/2005 14:23:19  ANR0944E QUERY PROCESS: No active processes found.
                      (SESSION: 3)
```

5. If we want to do a database backup, we can start it now with the same command we used before.

Results summary

The results of our test show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a database backup process started on the server before the failure, does not restart when the second node brings the Tivoli Storage Manager server instance online.

The tape volume is correctly unloaded from the tape drive where it was mounted when the Tivoli Storage Manager server is again online, but the process is not restarted unless you run the command again.

There is no difference between a scheduled process or a manual process using the administrative interface.

13.7.6 Testing inventory expiration

In this section we test another server task: an inventory expiration process.

Objective

The objective of this test is to show what happens when Tivoli Storage Manager server is running the inventory expiration process and the node that hosts the server instance fails.

Activities

For this test, we perform these tasks:

1. We use the /usr/sbin/rsct/sapolicies/bin/getstatus script to find out that the SA-tsmserver-rg is running on our first node, diomede.
2. We run the following command to start an inventory expiration process:

   ```
   expire inventory
   ```

3. A process starts for inventory expiration as shown in Example 13-41.

Example 13-41 Starting inventory expiration

```
02/10/2005 15:34:53  ANR0984I Process 13 for EXPIRATION started in the
                      BACKGROUND at 03:34:53 PM. (PROCESS: 13)
02/10/2005 15:34:53  ANR0811I Inventory client file expiration started as
                      process 1. (PROCESS: 13)
02/10/2005 15:34:53  ANR4391I Expiration processing node W2KCLIENT01, filespace
                      SYSTEM OBJECT, fsId 18, domain STANDARD, and management
                       class DEFAULT - for BACKUP type files. (PROCESS: 13)
02/10/2005 15:34:53  ANR4391I Expiration processing node RH9CLIENT01, filespace
                      /home, fsId 5, domain STANDARD, and management class
                      DEFAULT - for BACKUP type files. (PROCESS: 13)
```

4. While Tivoli Storage Manager server is expiring objects, we force a failure on the node that hosts the server instance. The SA-tsmserver-rg resource group is brought online on the second node, lochness. The inventory expiration process is not started any more. There are no errors in the activity log.
5. If we want to start the process again, we just have to run the same command again.

Archived

Results summary

The results of our test show that after a failure on the node that hosts the Tivoli Storage Manager server instance, an inventory expiration process started on the server before the failure does not restart when the second node brings the Tivoli Storage Manager server instance online.

There is no error inside the Tivoli Storage Manager server database, and we can restart the process again when the server is online.

Archived

14

Linux and Tivoli System Automation with IBM Tivoli Storage Manager Client

In this chapter we discuss the details related to the installation and configuration of the Tivoli Storage Manager client V5.3, installed on RHEL V3 U2 and running as a highly available application under the control of Tivoli System Automation V1.2. The installation on another Linux distribution supported by both Tivoli System Automation V1.2 and Tivoli Storage Manager client V5.3 should work in the same way as the installation described in this chapter for RHEL V3.

© Copyright IBM Corp. 2005. All rights reserved.

14.1 Overview

An application made *highly available* needs a backup program product that has been made *highly available* too.

Tivoli System Automation allows scheduled Tivoli Storage Manager client operations to continue processing during a failover situation.

Tivoli Storage Manager in a Tivoli System Automation environment can back up anything that Tivoli Storage Manager can normally back up. However, we must be careful when backing up non-clustered resources due to the after-failover effects.

Local resources should never be backed up or archived from clustered Tivoli Storage Manager nodes. Local Tivoli Storage Manager nodes should be used for local resources.

The Tivoli Storage Manager client code will be installed on all cluster nodes, and three client nodes will be defined, one clustered and two local nodes. The dsm.sys file will be located in the default directory /opt/tivoli/tsm/client/ba/bin on each node. It contains a stanza unique for each local client, and a stanza for the clustered client which will be the same on all nodes. All cluster resource groups which are highly available will have its own Tivoli Storage Manager client. In our lab environment, a NFS server will be an application in a resource group, and will have the Tivoli Storage Manager client included.

For the clustered client node, the dsm.opt file and inclexcl.lst files will be highly available, and located on the application shared disk. The Tivoli Storage Manager client environment variables which reference these option files will be used by the StartCommand configured in Tivoli System Automation.

Archived

14.2 Planning and design

There must be a requirement to configure an Tivoli System Automation for Multiplatforms Tivoli Storage Manager client. The most common requirement would be an application, such as highly available file server that has been configured and running under Tivoli System Automation for Multiplatforms control. In such cases, the Tivoli Storage Manager client will be configured within the same resource group as this application. This ensures the Tivoli Storage Manager client is tightly coupled with the application which requires backup and recovery services.

We are testing the configuration and clustering for one or more Tivoli Storage Manager client node instance and demonstrating the possibility of restarting a client operation just after the takeover of a crashed node.

Our design considers a two node cluster, with two local Tivoli Storage Manager client nodes to be used with local storage resources and a clustered client node to manage shared storage resources backup and archive.

To distinguish the three client nodes, we use different paths for configuration files and running directory, different TCP/IP addresses, and different TCP/IP ports as shown in Table 14-1.

Table 14-1 Tivoli Storage Manager client distinguished configuration

Node name	Node directory	TCP/IP address	TCP/IP port
diomede	/opt/tivoli/tsm/client/ba/bin	9.1.39.165	1501
lochness	/opt/tivoli/tsm/client/ba/bin	9.1.39.167	1501
cl_itsamp02_client	/mnt/nfsfiles/tsm/client/ba/bin	9.1.39.54	1503

- We use default local paths for the local client nodes instances and a path on a shared filesystem for the clustered one.
- Default port 1501 is used for the local client nodes agent instances while 1503 is used for the clustered one.
- Persistent addresses are used for local Tivoli Storage Manager resources.
- After reviewing the *Backup-Archive Clients Installation and User's Guide*, we then proceed to complete our environment configuration as shown in Table 14-2.

Archived

Table 14-2 Client nodes configuration of our lab

Node 1	
TSM nodename	DIOMEDE
dsm.opt location	/opt/tivoli/tsm/client/ba/bin
Backup domain	/, /usr, /var, /home, /opt
Client Node high level address	9.1.39.165
Client Node low level address	1501
Node 2	
TSM nodename	LOCHNESS
dsm.opt location	/opt/tivoli/tsm/client/ba/bin
Backup domain	/, /usr, /var, /home, /opt
Client Node high level address	9.1.39.167
Client Node low level address	1501
Virtual node	
TSM nodename	CL_ITSAMP02_CLIENT
dsm.opt location	/mnt/nfsfiles/tsm/client/ba/bin
Backup domain	/mnt/nfsfiles
Client Node high level address	9.1.39.54
Client Node low level address	1503

14.3 Lab setup

In our test environment, we configure a highly available NFS file service as an example application. A detailed description how to manage a highly available NFS server with Tivoli System Automation can be found in the paper, “Highly available NFS server with Tivoli System Automation for Linux“, available at:

http://www.ibm.com/software/tivoli/products/sys-auto-linux/downloads.html

The Tivoli System Automation configuration files for the NFS server are located in /usr/sbin/rsct/sapolicies/nfsserver.

Archived

14.4 Installation

We need to install Tivoli System Automation V1.2 and the Tivoli Storage Manager client V5.3 on the nodes in the cluster. We use the Tivoli Storage Manager server V5.3 running on the Windows 2000 cluster to back up and restore data. For the installation and configuration of the Tivoli Storage Manager server in this test, refer to Chapter 5, “Microsoft Cluster Server and the IBM Tivoli Storage Manager Server” on page 77.

14.4.1 Tivoli System Automation V1.2 installation

We have installed, configured, and tested Tivoli System Automation prior to this point, and will utilize this infrastructure to hold our highly available application, and our highly available Tivoli Storage Manager client. To reference the Tivoli System Automation installation, see “Installation of Tivoli System Automation” on page 611.

14.4.2 Tivoli Storage Manager Client Version 5.3 installation

We have installed the Tivoli Storage Manager client V5.3 prior to this point, and will focus our efforts on the configuration in this chapter. To reference the client installation, refer to “Installation of Tivoli Storage Manager Client” on page 620.

14.5 Configuration

Before we can actually use the clustered Tivoli Storage Manager client, we must configure the clustered Tivoli Storage Manager client and the Tivoli System Automation resource group that should use the clustered Tivoli Storage Manager client.

14.5.1 Tivoli Storage Manager Client configuration

To configure the Tivoli Storage Manager Client, we follow these steps:

1. We execute the following Tivoli Storage Manager command on the Tivoli Storage Manager server:

   ```
   register node cl_itsamp02_client itsosj passexp=0
   ```

   Important: We set the passexp to 0, so the password will not expire, because we have to store the password file for the clustered client on both nodes locally. If we enable the password expiry, we must ensure to update the password file on all nodes after a password change manually.

2. Then we mount the intended application resource shared disk on one node, diomede. There we create a directory to hold the Tivoli Storage Manager configuration and log files. The path is /mnt/nfsfiles/tsm/client/ba/bin, in our case, with the mount point for the file system being /mnt/nfsfiles.

> **Note:** Depending on your needs, it may be desirable to use a dedicated file system for the Tivoli Storage Manager client configuration and log files. In certain situations, log files may grow very fast. This can lead to filling up a file system completely. Placing log files on a dedicated file system can limit the impact of such a situation.

3. We copy the default dsm.opt.smp to /mnt/nfsfiles/tsm/client/ba/bin/dsm.opt (on the shared disk) and edit the file with the servername to be used by this client instance as shown in Example 14-1.

Example 14-1 dsm.opt file contents located in the application shared disk

```
************************************************************************
* IBM Tivoli Storage Manager                                           *
************************************************************************
*  This servername is the reference for the highly available TSM       *
*  client.                                                             *
************************************************************************

SErvername       tsmsrv01_ha
```

4. We add the necessary stanza into dsm.sys on each node. This stanza for the clustered Tivoli Storage Manager client has the same contents on all nodes, as shown in Example 14-2. Each node has its own copy of the dsm.sys file on its local file system, containing also stanzas for the local Tivoli Storage Manager client nodes. The file is located at the default location /opt/tivoli/tsm/client/ba/bin/dsm.sys. We use the following options:

   a. The *passworddir* parameter points to a shared directory. Tivoli Storage Manager for Linux Client encrypts the password file with the host name. So it is necessary to create the password file locally on each node. We set the passworddir parameter in dsm.sys to the local directory /usr/sbin/rsct/sapolicies/nfsserver.

   b. The *managedservices* parameter is set to "*schedule webclient*", to have the *dsmc sched* waked up by the client acceptor daemon at schedule start time, as suggested in the *UNIX and Linux Backup-Archive Clients Installation and User's Guide*.

Archived

c. Last, but most important, we add a domain statement for our shared file system. Domain statements are required to tie each file system to the corresponding Tivoli Storage Manager client node. Without that, each node will save all of the local mounted file systems during incremental backups. See Example 14-2.

> **Important:** When domain statements, one or more, are used in a client configuration, **only** those domains (file systems) will be backed up during incremental backup.

Example 14-2 Stanza for the clustered client in dsm.sys

```
* Server stanza for the ITSAMP highly available client connection purpose
SErvername              tsmsrv01_ha
nodename                cl_itsamp02_client
   COMMMethod           TCPip
   TCPPort              1500
   TCPServeraddress     9.1.39.73
   HTTPPORT             1582
   ERRORLOGRETENTION    7
   ERRORLOGname         /mnt/nfsfiles/tsm/client/ba/bin/dsm_error.log
   passwordaccess       generate
   passworddir          /usr/sbin/rsct/sapolicies/nfsserver
   managedservices      schedule webclient
   domain               /mnt/nfsfiles
```

5. We connect to the Tivoli Storage Manager server using **dsmc -server=tsmsrv01_ha** from the Linux command line. This will generate the TSM.PWD file as shown in Example 14-3. We issue this step on each node to create the password file on every node.

Example 14-3 Creation of the password file TSM.PWD

```
[root@diomede nfsserver]# pwd
/usr/sbin/rsct/sapolicies/nfsserver
[root@diomede nfsserver]# dsmc -se=tsmsrv01_ha
IBM Tivoli Storage Manager
Command Line Backup/Archive Client Interface
  Client Version 5, Release 3, Level 0.0
  Client date/time: 02/14/2005 17:56:08
(c) Copyright by IBM Corporation and other(s) 1990, 2004. All Rights Reserved.

Node Name: CL_ITSAMP02_CLIENT
Please enter your user id <CL_ITSAMP02_CLIENT>:

Please enter password for user id "CL_ITSAMP02_CLIENT":

Session established with server TSMSRV01: Windows
```

```
  Server Version 5, Release 3, Level 0.0
  Server date/time: 02/14/2005 17:59:55  Last access: 02/14/2005 17:59:46

tsm> quit
[root@diomede nfsserver]# ls -l TSM.PWD
-rw-------    1 root     root          151 Feb 14 17:56 TSM.PWD
[root@diomede nfsserver]#
```

14.5.2 Tivoli Storage Manager client resource configuration

A Tivoli Storage Manager client resource is controlled by the tsmclientctrl-cad script. This script is used to start, stop, and monitor the Tivoli Storage Manager Client Acceptor Daemon (CAD). It is able to cancel old client sessions that may be present on the Tivoli Storage Manager server when executing the failover. This happens especially when using a higher value for the Tivoli Storage Manager CommTimeOut parameter. It is necessary to cancel these old sessions, as they still count for maximum number of points. “Troubleshooting” on page 545 describes this behavior in detail for the AIX Tivoli Storage Manager client.

The script is used in the following way:

```
tsmclientctrl-cad { start | stop | status } <TSM_CLIENT_HA_DIR> <prefix>
<TSM_NODE> <TSM_SRV> <TSM_USER> <TSM_PASS>
```

The parameters have the following meanings:

- TSM_CLIENT_HA_DIR: The directory, where the Tivoli Storage Manager client configuration and log files for the clustered client are located
- prefix: The prefix of the Tivoli System Automation resource group - this is necessary to create a unique pid file for this clustered Tivoli Storage Manager client
- TSM_NODE: The Tivoli Storage Manager client nodename, necessary to cancel old client sessions
- TSM_SRV: The Tivoli Storage Manager server name, necessary to cancel old client sessions
- TSM_USER: The Tivoli Storage Manager user with operator privileges, necessary to cancel old client sessions
- TSM_PASS: The password for the specified Tivoli Storage Manager user, necessary to cancel old client sessions

Archived

To configure the Tivoli System Automation resource, we follow these steps:

1. We change to the directory where the control scripts for the clustered application we want to back up are stored. In our example this is /usr/sbin/rsct/sapolicies/nfsserver/. Within this directory, we create a symbolic link to the script which controls the Tivoli Storage Manager client CAD in the Tivoli System Automation for Multiplatforms environment. We accomplish these steps on both nodes as shown in Example 14-4.

Example 14-4 Creation of the symbolic link that point to the Client CAD script

```
[root@diomede root]# cd /usr/sbin/rsct/sapolicies/nfsserver
[root@diomede nfsserver]# ln -s \
> /usr/sbin/rsct/sapolicies/tsmclient/tsmclientctrl-cad nfsserverctrl-tsmclient
[root@diomede nfsserver]#
```

2. We configure the cluster application for Tivoli System Automation for Multiplatforms, in our case the NFS server. The necessary steps to configure a NFS server for Tivoli System Automation for Multiplatforms are described in detail in the paper "Highly available NFS server with Tivoli System Automation for Linux", available at:

   http://www.ibm.com/software/tivoli/products/sys-auto-linux/downloads.html

3. We ensure that the resources of the cluster application resource group are offline. We use the Tivoli System Automation for Multiplatforms **lsrg -m** command on any node for this purpose. The output of the command is shown in Example 14-5.

Example 14-5 Output of the lsrg -m command before configuring the client

```
Displaying Member Resource information:
Class:Resource:Node[ManagedResource]       Mandatory MemberOf         OpState
IBM.Application:SA-nfsserver-server        True      SA-nfsserver-rg  Offline
IBM.ServiceIP:SA-nfsserver-ip-1            True      SA-nfsserver-rg  Offline
IBM.Application:SA-nfsserver-data-nfsfiles True      SA-nfsserver-rg  Offline
```

4. The necessary resource for the Tivoli Storage Manager client CAD should depend on the NFS server resource of the clustered NFS server. In that way it is guaranteed that all necessary file systems are mounted before the Tivoli Storage Manager client CAD is started by Tivoli System Automation for Multiplatforms. To configure that behavior we do the following steps. We execute these steps only on the first node, diomede.

   a. We prepare the configuration file for the SA-nfsserver-tsmclient resource.

      All parameters for the StartCommand, StopCommand, and MonitorCommand must be on a single line in this file. Example 14-6 shows the contents of the file with line breaks between the parameters.

Note: We enter the nodename parameter for the StartCommand, StopCommand, and MonitorCommand in uppercase letters. This is necessary, as the nodename will be used for an SQL query in Tivoli Storage Manager. We also use an extra Tivoli Storage Manager user, called “scriptoperator”, which is necessary to query and reset Tivoli Storage Manager sessions. Be sure that this user can access the Tivoli Storage Manager server.

Example 14-6 Definition file SA-nfsserver-tsmclient.def

```
PersistentResourceAttributes::
Name=SA-nfsserver-tsmclient
ResourceType=1
StartCommand=/usr/sbin/rsct/sapolicies/nfsserver/nfsserverctrl-tsmclient start
/mnt/nfsfiles/tsm/client/ba/bin SA-nfsserver- CL_ITSAMP02_CLIENT tsmsrv01_ha
scriptoperator password
StopCommand=/usr/sbin/rsct/sapolicies/nfsserver/nfsserverctrl-tsmclient stop
/mnt/nfsfiles/tsm/client/ba/bin SA-nfsserver- CL_ITSAMP02_CLIENT tsmsrv01_ha
scriptoperator password
MonitorCommand=/usr/sbin/rsct/sapolicies/nfsserver/nfsserverctrl-tsmclient
status /mnt/nfsfiles/tsm/client/ba/bin SA-nfsserver- CL_ITSAMP02_CLIENT
tsmsrv01_ha scriptoperator password
StartCommandTimeout=180
StopCommandTimeout=60
MonitorCommandTimeout=9
MonitorCommandPeriod=10
ProtectionMode=0
NodeNameList={'diomede','lochness'}
UserName=root
```

Note: We use a StartCommandTimouout of 180 seconds, as it may take some time to cancel all old Tivoli Storage Manager client sessions.

b. We manually add the SA-nfsserver-tsmclient resource to Tivoli System Automation for Multiplatforms with the command **mkrsrc -f SA-nfsserver-tsmclient.def IBM.Application**.

c. Now that the resource is known by Tivoli System Automation for Multiplatforms, we add it to the resource group SA-nfsserver-rg with the command **addrgmbr -m T -g SA-nfsserver-rg IBM.Application:SA-nfsserver-tsmclient**.

d. Finally we configure the dependency with the command: **mkrel -S IBM.Application:SA-nfsserver-tsmclient -G IBM.Application:SA-nfsserver-server -p DependsOn SA-nfsserver-tsmclient-on-server**. We verify the relationships with the **lsrel command**. The output of the command is shown in Example 14-7.

Example 14-7 Output of the lsrel command

```
Displaying Managed Relations :

Name                                Class:Resource:Node[Source]          ResourceGroup[Source]
SA-nfsserver-server-on-ip-1         IBM.Application:SA-nfsserver-server    SA-nfsserver-rg
SA-nfsserver-server-on-data-nfsfiles IBM.Application:SA-nfsserver-server   SA-nfsserver-rg
SA-nfsserver-ip-on-nieq-1           IBM.ServiceIP:SA-nfsserver-ip-1        SA-nfsserver-rg
SA-nfsserver-tsmclient-on-server    IBM.Application:SA-nfsserver-tsmclient SA-nfsserver-rg
```

5. Now we start the resource group with the **chrg -o online SA-nfsserver-rg** command.
6. To verify that all necessary resources are online, we use again the **lsrg -m** command. Example 14-8 shows the output of this command.

Example 14-8 Output of the lsrg -m command while resource group is online

```
Displaying Member Resource information:
Class:Resource:Node[ManagedResource]       Mandatory MemberOf         OpState
IBM.Application:SA-nfsserver-server        True      SA-nfsserver-rg  Online
IBM.ServiceIP:SA-nfsserver-ip-1            True      SA-nfsserver-rg  Online
IBM.Application:SA-nfsserver-data-nfsfiles True      SA-nfsserver-rg  Online
IBM.Application:SA-nfsserver-tsmclient     True      SA-nfsserver-rg  Online
```

14.6 Testing the cluster

In order to check the high availability of Tivoli Storage Manager client on our lab environment, we must do some testing.

Our objective with these tests is to know how Tivoli Storage Manager can respond, on a clustered environment, after certain kinds of failures that affect the shared resources.

For the purpose of this section, we use a Tivoli Storage Manager server installed on an Windows 2000 machine: TSMSRV01. We use a tape storage pool for incremental backup and restore. Incremental backup of small files to tape storage pools is not a best practice. The following tests also work with disk storage pools in our test environment.

Archived

14.6.1 Testing client incremental backup

In this section we discuss how to test the client incremental backup.

Objective

The objective of this test is to show what happens when a client incremental backup is started for a virtual node on the cluster, and the cluster node that hosts the resources at that moment suddenly fails.

Activities

To do this test, we perform these tasks:

1. We use the /usr/sbin/rsct/sapolicies/bin/getstatus script to find out that the SA-nfsserver-rg resource group is online on our first node, diomede.
2. We schedule a client incremental backup operation using the Tivoli Storage Manager server scheduler and we associate the schedule to CL_ITSAMP02_CLIENT nodename.
3. At the scheduled time, a client session for CL_ITSAMP02_CLIENT nodename starts on the server as shown in Example 14-9.

Example 14-9 Session for CL_ITSAMP02_CLIENT starts

```
02/15/2005 11:51:10  ANR0406I Session 35 started for node CL_ITSAMP02_CLIENT
                      (Linux86) (Tcp/Ip 9.1.39.165(32800)). (SESSION: 35)
02/15/2005 11:51:20  ANR0406I Session 36 started for node CL_ITSAMP02_CLIENT
                      (Linux86) (Tcp/Ip 9.1.39.165(32801)). (SESSION: 36)
```

4. The client starts sending files to the server as we can see on the schedule log file /mnt/nfsfiles/tsm/client/ba/bin/dsmsched.log shown in Example 14-10.

Example 14-10 Schedule log file during starting of the scheduled backup

```
02/15/2005 11:49:14 --- SCHEDULEREC QUERY BEGIN
02/15/2005 11:49:14 --- SCHEDULEREC QUERY END
02/15/2005 11:49:14 Next operation scheduled:
02/15/2005 11:49:14
------------------------------------------------------------
02/15/2005 11:49:14 Schedule Name:         SCHEDULE_1
02/15/2005 11:49:14 Action:                Incremental
02/15/2005 11:49:14 Objects:
02/15/2005 11:49:14 Options:
02/15/2005 11:49:14 Server Window Start:   11:50:00 on 02/15/2005
02/15/2005 11:49:14
------------------------------------------------------------
02/15/2005 11:49:14
Executing scheduled command now.
02/15/2005 11:49:14 --- SCHEDULEREC OBJECT BEGIN SCHEDULE_1 02/15/2005 11:50:00
```

```
02/15/2005 11:49:14 Incremental backup of volume '/mnt/nfsfiles'
02/15/2005 11:49:16 ANS1898I ***** Processed       500 files *****
02/15/2005 11:49:17 ANS1898I ***** Processed     1,000 files *****
02/15/2005 11:49:18 ANS1898I ***** Processed     1,500 files *****
```

5. While the client continues sending files to the server, we force a failover by unplugging the eth0 network connection of diomede. The client loses its connection with the server, and the session terminates, as we can see on the Tivoli Storage Manager server activity log shown in Example 14-11.

Example 14-11 Activity log entries while diomede fails

```
02/15/2005 11:54:22  ANR0514I Session 36 closed volume 021AKKL2. (SESSION: 36)

02/15/2005 11:54:22  ANR0480W Session 36 for node CL_ITSAMP02_CLIENT
(Linux86)

                terminated - connection with client severed. (SESSION:

                36)
```

6. The other node, lochness, brings the resources online. When the Tivoli Storage Manager Scheduler starts, the client restarts the backup as we show on the schedule log file in Example 14-12. The backup restarts, since the schedule is still within the startup window.

Example 14-12 Schedule log file dsmsched.log after restarting the backup

```
/favorites_PA_1_0_38.ear/favorites.war/resources/com/ibm/psw/wcl/renderers/menu
[Sent]
02/15/2005 11:52:04 Directory-->              4,096
/mnt/nfsfiles/root/isc-backup-2005-02-03-11-15/PortalServer/installedApps
/favorites_PA_1_0_38.ear/favorites.war/resources/com/ibm/psw/wcl/renderers/scri
pts [Sent]
02/15/2005 11:54:03 Scheduler has been started by Dsmcad.
02/15/2005 11:54:03 Querying server for next scheduled event.
02/15/2005 11:54:03 Node Name: CL_ITSAMP02_CLIENT
02/15/2005 11:54:28 Session established with server TSMSRV01: Windows
02/15/2005 11:54:28   Server Version 5, Release 3, Level 0.0
02/15/2005 11:54:28   Server date/time: 02/15/2005 11:56:23  Last access:
02/15/2005 11:55:07

02/15/2005 11:54:28 --- SCHEDULEREC QUERY BEGIN
02/15/2005 11:54:28 --- SCHEDULEREC QUERY END
02/15/2005 11:54:28 Next operation scheduled:
02/15/2005 11:54:28
------------------------------------------------------------
02/15/2005 11:54:28 Schedule Name:         SCHEDULE_1
```

```
02/15/2005 11:54:28 Action:                Incremental
02/15/2005 11:54:28 Objects:
02/15/2005 11:54:28 Options:
02/15/2005 11:54:28 Server Window Start:   11:50:00 on 02/15/2005
02/15/2005 11:54:28
--------------------------------------------------------------
02/15/2005 11:54:28 Scheduler has been stopped.

02/15/2005 11:56:29 Scheduler has been started by Dsmcad.
02/15/2005 11:56:29 Querying server for next scheduled event.
02/15/2005 11:56:29 Node Name: CL_ITSAMP02_CLIENT
02/15/2005 11:56:54 Session established with server TSMSRV01: Windows
02/15/2005 11:56:54   Server Version 5, Release 3, Level 0.0
02/15/2005 11:56:54   Server date/time: 02/15/2005 11:58:49  Last access:
02/15/2005 11:56:23

02/15/2005 11:56:54 --- SCHEDULEREC QUERY BEGIN
02/15/2005 11:56:54 --- SCHEDULEREC QUERY END
02/15/2005 11:56:54 Next operation scheduled:
02/15/2005 11:56:54
--------------------------------------------------------------
02/15/2005 11:56:54 Schedule Name:         SCHEDULE_1
02/15/2005 11:56:54 Action:                Incremental
02/15/2005 11:56:54 Objects:
02/15/2005 11:56:54 Options:
02/15/2005 11:56:54 Server Window Start:   11:50:00 on 02/15/2005
02/15/2005 11:56:54
--------------------------------------------------------------
02/15/2005 11:56:54
Executing scheduled command now.
02/15/2005 11:56:54 --- SCHEDULEREC OBJECT BEGIN SCHEDULE_1 02/15/2005 11:50:00
02/15/2005 11:56:54 Incremental backup of volume '/mnt/nfsfiles'
02/15/2005 11:56:55 ANS1898I ***** Processed     5,000 files *****
02/15/2005 11:56:56 ANS1898I ***** Processed    11,000 files *****
02/15/2005 11:57:05 Normal File-->                 0
/mnt/nfsfiles/.sa-ctrl-data-DO_NOT_DELETE [Sent]
02/15/2005 11:57:05 Directory-->               4,096
/mnt/nfsfiles/root/isc-backup-2005-02-03-11-15/PortalServer/installedApps
/favorites_PA_1_0_38.ear/favorites.war/resources/com/ibm/psw/wcl/renderers/menu
/html [Sent]
02/15/2005 11:57:05 Normal File-->            37,764
/mnt/nfsfiles/root/isc-backup-2005-02-03-11-15/PortalServer/installedApps
/favorites_PA_1_0_38.ear/favorites.war/resources/com/ibm/psw/wcl/renderers/menu
/html/context_ie.js [Sent]
```

In the Tivoli Storage Manager server activity log we can see how the connection was lost and a new session starts again for CL_ITSAMP02_CLIENT as shown in Example 14-13.

Example 14-13 Activity log entries while the new session for the backup starts

```
02/15/2005 11:55:07  ANR0406I Session 39 started for node CL_ITSAMP02_CLIENT
                      (Linux86) (Tcp/Ip 9.1.39.167(32830)). (SESSION: 39)
02/15/2005 11:55:07  ANR1639I Attributes changed for node CL_ITSAMP02_CLIENT:
                      TCP Name from diomede to lochness, TCP Address from
                      9.1.39.165 to 9.1.39.167, GUID from
                      b4.cc.54.42.fb.6b.d9.11.ab.61.00.0d.60.49.4c.39 to
                      22.77.12.20.fc.6b.d9.11.84.80.00.0d.60.49.6a.62.
                      (SESSION: 39)
02/15/2005 11:55:07  ANR0403I Session 39 ended for node CL_ITSAMP02_CLIENT
                      (Linux86). (SESSION: 39)
02/15/2005 11:55:12  ANR8468I LTO volume 021AKKL2 dismounted from drive DRLTO_1
                      (mt0.0.0.4) in library LIBLTO. (SESSION: 36)
...
02/15/2005 11:58:49  ANR0406I Session 41 started for node CL_ITSAMP02_CLIENT
                      (Linux86) (Tcp/Ip 9.1.39.167(32833)). (SESSION: 41)
02/15/2005 11:59:00  ANR0406I Session 42 started for node CL_ITSAMP02_CLIENT
                      (Linux86) (Tcp/Ip 9.1.39.167(32834)). (SESSION: 42)
02/15/2005 11:59:28  ANR8337I LTO volume 021AKKL2 mounted in drive DRLTO_2
                      (mt1.0.0.4). (SESSION: 42)
02/15/2005 11:59:28  ANR0511I Session 42 opened output volume 021AKKL2.
                      (SESSION: 42)
...
02/15/2005 12:06:29  ANR0514I Session 42 closed volume 021AKKL2. (SESSION: 42)
02/15/2005 12:06:29  ANR2507I Schedule SCHEDULE_1 for domain STANDARD started
                      at 02/15/2005 11:50:00 for node CL_ITSAMP02_CLIENT
                      completed successfully at 02/15/2005 12:06:29. (SESSION:
                      41)
```

7. The incremental backup ends without errors as we see on the schedule log file in Example 14-14.

Example 14-14 Schedule log file reports the successfully completed event

```
02/15/2005 12:04:34 --- SCHEDULEREC OBJECT END SCHEDULE_1 02/15/2005 11:50:00
02/15/2005 12:04:34 Scheduled event ‘SCHEDULE_1’ completed successfully.
02/15/2005 12:04:34 Sending results for scheduled event ‘SCHEDULE_1’.
02/15/2005 12:04:34 Results sent to server for scheduled event ‘SCHEDULE_1’.
```

Results summary

The test results show that after a failure on the node that hosts the Tivoli Storage Manager scheduler service resource, a scheduled incremental backup started on one node is restarted and successfully completed on the other node that takes the failover.

Archived

This is true if the startup window used to define the schedule is not elapsed when the scheduler services restarts on the second node.

14.6.2 Testing client restore

In this section we discuss how to test the client restore operation.

Objective

The objective of this test is to show what happens when a client restore is started for a virtual node on the cluster, and the cluster node that hosts the resources at that moment suddenly fails.

Activities

To do this test, we perform these tasks:

1. We use the /usr/sbin/rsct/sapolicies/bin/getstatus script to find out that the SA-nfsserver-rg resource group is online on our first node, diomede.
2. We schedule a client restore operation using the Tivoli Storage Manager server scheduler and we associate the schedule to CL_ITSAMP02_CLIENT nodename.
3. At the scheduled time a client session for CL_ITSAMP02_CLIENT nodename starts on the server as shown in Example 14-15.

Example 14-15 Activity log entries during start of the client restore

```
02/16/2005 12:08:05  ANR0406I Session 36 started for node CL_ITSAMP02_CLIENT
                      (Linux86) (Tcp/Ip 9.1.39.165(32779)). (SESSION: 36)
...
02/16/2005 12:08:41  ANR8337I LTO volume 021AKKL2 mounted in drive DRLTO_2
                      (mt1.0.0.4). (SESSION: 36)
02/16/2005 12:08:41  ANR0510I Session 36 opened input volume 021AKKL2.
                      (SESSION: 36)
```

4. The client starts restoring files as we can see on the schedule log file in Example 14-16.

Example 14-16 Schedule log entries during start of the client restore

```
02/16/2005 12:08:03 --- SCHEDULEREC OBJECT BEGIN SCHEDULE_2 02/16/2005 12:05:00
02/16/2005 12:08:03 Restore function invoked.

02/16/2005 12:08:04 ANS1247I Waiting for files from the server...Restoring
4,096 /mnt/nfsfiles/root [Done]
02/16/2005 12:08:04 Restoring           4,096 /mnt/nfsfiles/root/.gconf [Done]
...
```

```
02/16/2005 12:08:08 Restoring           4,096
/mnt/nfsfiles/root/tsmi686/cdrom/license/i386/jre/lib/images/ftp [Done]
02/16/2005 12:08:40  ** Interrupted **
02/16/2005 12:08:40 ANS1114I Waiting for mount of offline media.
02/16/2005 12:08:40 Restoring             161 /mnt/nfsfiles/root/.ICEauthority
[Done]
02/16/2005 12:08:40 Restoring             526 /mnt/nfsfiles/root/.Xauthority
[Done]
...
```

5. While the client is restoring the files, we force diomede to fail (unplugging network cable for eth0). The client loses its connection with the server, and the session is terminated as we can see on the Tivoli Storage Manager server activity log shown in Example 14-17.

Example 14-17 Activity log entries during the failover

```
02/16/2005 12:10:30  ANR0514I Session 36 closed volume 021AKKL2. (SESSION: 36)
02/16/2005 12:10:30  ANR8336I Verifying label of LTO volume 021AKKL2 in drive
                       DRLTO_2 (mt1.0.0.4). (SESSION: 36)
02/16/2005 12:10:30  ANR0480W Session 36 for node CL_ITSAMP02_CLIENT (Linux86)
                      terminated - connection with client severed. (SESSION:
                      36)
```

6. Lochness brings the resources online. When the Tivoli Storage Manager scheduler service resource is again online on lochness and queries the server, if the startup window for the scheduled operation is not elapsed, the restore process restarts from the beginning, as we can see on the schedule log file in Example 14-18.

Example 14-18 Schedule log entries during restart of the client restore

```
02/16/2005 12:10:01 Restoring      77,475,840
/mnt/nfsfiles/root/itsamp/1.2.0-ITSAMP-FP03linux.tar [Done]
02/16/2005 12:12:04 Scheduler has been started by Dsmcad.
02/16/2005 12:12:04 Querying server for next scheduled event.
02/16/2005 12:12:04 Node Name: CL_ITSAMP02_CLIENT
02/16/2005 12:12:29 Session established with server TSMSRV01: Windows
02/16/2005 12:12:29   Server Version 5, Release 3, Level 0.0
02/16/2005 12:12:29   Server date/time: 02/16/2005 12:12:30  Last access:
02/16/2005 12:11:13

02/16/2005 12:12:29 --- SCHEDULEREC QUERY BEGIN
02/16/2005 12:12:29 --- SCHEDULEREC QUERY END
02/16/2005 12:12:29 Next operation scheduled:
02/16/2005 12:12:29
------------------------------------------------------------
02/16/2005 12:12:29 Schedule Name:         SCHEDULE_2
02/16/2005 12:12:29 Action:                Restore
```

```
02/16/2005 12:12:29 Objects:               /mnt/nfsfiles/root/
02/16/2005 12:12:29 Options:               -subdir=yes
02/16/2005 12:12:29 Server Window Start:   12:05:00 on 02/16/2005
02/16/2005 12:12:29
--------------------------------------------------------------
02/16/2005 12:12:29 Scheduler has been stopped.

02/16/2005 12:14:30 Scheduler has been started by Dsmcad.
02/16/2005 12:14:30 Querying server for next scheduled event.
02/16/2005 12:14:30 Node Name: CL_ITSAMP02_CLIENT
02/16/2005 12:14:55 Session established with server TSMSRV01: Windows
02/16/2005 12:14:55   Server Version 5, Release 3, Level 0.0
02/16/2005 12:14:55   Server date/time: 02/16/2005 12:14:56  Last access:
02/16/2005 12:12:30

02/16/2005 12:14:55 --- SCHEDULEREC QUERY BEGIN
02/16/2005 12:14:55 --- SCHEDULEREC QUERY END
02/16/2005 12:14:55 Next operation scheduled:
02/16/2005 12:14:55
--------------------------------------------------------------
02/16/2005 12:14:55 Schedule Name:         SCHEDULE_2
02/16/2005 12:14:55 Action:                Restore
02/16/2005 12:14:55 Objects:               /mnt/nfsfiles/root/
02/16/2005 12:14:55 Options:               -subdir=yes
02/16/2005 12:14:55 Server Window Start:   12:05:00 on 02/16/2005
02/16/2005 12:14:55
--------------------------------------------------------------
02/16/2005 12:14:55
Executing scheduled command now.
02/16/2005 12:14:55 --- SCHEDULEREC OBJECT BEGIN SCHEDULE_2 02/16/2005 12:05:00
02/16/2005 12:14:55 Restore function invoked.

02/16/2005 12:14:56 ANS1247I Waiting for files from the server...Restoring
4,096 /mnt/nfsfiles/root/.gconf [Done]

02/16/2005 12:14:56 Restoring           4,096 /mnt/nfsfiles/root/.gconfd [Done]
...
02/16/2005 12:15:13 ANS1946W File /mnt/nfsfiles/root/itsamp/C57NWML.tar exists,
skipping
02/16/2005 12:16:09  ** Interrupted **
02/16/2005 12:16:09 ANS1114I Waiting for mount of offline media.
02/16/2005 12:16:09 Restoring          55,265
/mnt/nfsfiles/root/itsamp/sam.policies-1.2.1.0-0.i386.rpm [Done]
```

7. In the activity log of Tivoli Storage Manager server, we see that a new session is started for CL_MSCS01_SA as shown in Example 14-19.

Example 14-19 Activity log entries during restart of the client restore

```
02/16/2005 12:11:13  ANR0406I Session 38 started for node CL_ITSAMP02_CLIENT
                      (Linux86) (Tcp/Ip 9.1.39.167(32789)). (SESSION: 38)
02/16/2005 12:11:13  ANR1639I Attributes changed for node CL_ITSAMP02_CLIENT:
                      TCP Name from diomede to lochness, TCP Address from
                      9.1.39.165 to 9.1.39.167, GUID from
                      b4.cc.54.42.fb.6b.d9.11.ab.61.00.0d.60.49.4c.39 to
                      22.77.12.20.fc.6b.d9.11.84.80.00.0d.60.49.6a.62.
                      (SESSION: 38)
02/16/2005 12:11:13  ANR0403I Session 38 ended for node CL_ITSAMP02_CLIENT
                      (Linux86). (SESSION: 38)
...
02/16/2005 12:14:56  ANR0406I Session 40 started for node CL_ITSAMP02_CLIENT
                      (Linux86) (Tcp/Ip 9.1.39.167(32791)). (SESSION: 40)
02/16/2005 12:15:39  ANR8337I LTO volume 021AKKL2 mounted in drive DRLTO_1
                      (mt0.0.0.4). (SESSION: 40)
02/16/2005 12:15:39  ANR0510I Session 40 opened input volume 021AKKL2.
                      (SESSION: 40)
```

8. When the restore completes, we can see the final statistics in the schedule log file of the client for a successful operation as shown in Example 14-20.

Example 14-20 Schedule log entries after client restore finished

```
02/16/2005 12:19:23
Restore processing finished.
02/16/2005 12:19:25 --- SCHEDULEREC STATUS BEGIN
02/16/2005 12:19:25 Total number of objects restored:     7,052
02/16/2005 12:19:25 Total number of objects failed:           0
02/16/2005 12:19:25 Total number of bytes transferred:     1.79 GB
02/16/2005 12:19:25 Data transfer time:                  156.90 sec
02/16/2005 12:19:25 Network data transfer rate:        11,979.74 KB/sec
02/16/2005 12:19:25 Aggregate data transfer rate:      6,964.13 KB/sec
02/16/2005 12:19:25 Elapsed processing time:           00:04:29
02/16/2005 12:19:25 --- SCHEDULEREC STATUS END
02/16/2005 12:19:25 --- SCHEDULEREC OBJECT END SCHEDULE_2 02/16/2005 12:05:00
02/16/2005 12:19:25 --- SCHEDULEREC STATUS BEGIN
02/16/2005 12:19:25 --- SCHEDULEREC STATUS END
02/16/2005 12:19:25 Scheduled event ‘SCHEDULE_2’ completed successfully.
02/16/2005 12:19:25 Sending results for scheduled event ‘SCHEDULE_2’.
02/16/2005 12:19:25 Results sent to server for scheduled event ‘SCHEDULE_2’.
```

Results summary

The test results show that after a failure on the node that hosts the Tivoli Storage Manager client scheduler instance, a scheduled restore operation started on this node is started again on the second node of the cluster when the service is online.

This is true if the startup window for the scheduled restore operation is not elapsed when the scheduler client is online again on the second node.

Note: The restore is not restarted from the point of failure, but started from the beginning. The scheduler queries the Tivoli Storage Manager server for a scheduled operation, and a new session is opened for the client after the failover.

Archived

15

Linux and Tivoli System Automation with IBM Tivoli Storage Manager Storage Agent

This chapter describes the use of Tivoli Storage Manager for Storage Area Network (also known as *Storage Agent*) to back up shared data of in a Linux Tivoli System Automation cluster using the LAN-free path.

© Copyright IBM Corp. 2005. All rights reserved.

15.1 Overview

We can configure the Tivoli Storage Manager client and server so that the client, through a Storage Agent, can move its data directly to storage on a SAN. This function, called LAN-free data movement, is provided by IBM Tivoli Storage Manager for Storage Area Networks.

> **Note:** For clustering of the Storage Agent, the Tivoli Storage Manager server needs to support the new resetdrives parameter. For Tivoli Storage Manager V5.3, the AIX Tivoli Storage Manager server supports this new parameter.

For more information about the tape drive SCSI reserve and reasons why clustering a Storage Agent, see “Overview” on page 556.

15.2 Planning and design

On our servers local Storage Agents running with default environment setting are configured too. We can have more than one dsmsta running on a single machine as for servers and clients. Port 1502 is used for the local instances while 1504 is used for the clustered one as shown in Table 15-1.

Table 15-1 Storage Agents configuration

STA instance	Instance path	TCP/IP address	TCP/IP port
diomede_sta	/opt/tivoli/tsm/StorageAgent/bin	9.1.39.165	1502
lochness_sta	/usr/tivoli/tsm/StorageAgent/bin	9.1.39.167	1502
cl_itsamp02_sta	/mnt/nfsfiles/tsm/StorageAgent/bin	9.1.39.54	1504

Here we are using TCP/IP as communication method, but shared memory also applies.

15.3 Installation

We install the Storage Agent via the **`rpm -ihv`** command on both nodes. We also create a symbolic link to the **`dsmsta`** executable. Example 15-1 shows the necessary steps.

Example 15-1 Installation of the TIVsm-stagent rpm on both nodes

```
[root@diomede i686]# rpm -ihv TIVsm-stagent-5.3.0-0.i386.rpm
Preparing...                ########################################### [100%]
   1:TIVsm-stagent          ########################################### [100%]
[root@diomede i686]# ln -s /opt/tivoli/tsm/StorageAgent/bin/dsmsta \
> /usr/bin/dsmsta
[root@diomede i686]#
```

15.4 Configuration

We need to configure the Storage Agent, the backup/archive client, and the necessary Tivoli System Automation resources. We explain the necessary steps in this section.

15.4.1 Storage agents

To enable the use of the Storage Agents, we must configure them to the Tivoli Storage Manager server and do some local configuration of the Storage Agents themselves.

Configure paths on the server

We need to configure the paths for the Storage Agent on the server. We do this with the following commands entered within the Tivoli Storage Manager administration console:

```
DEFINE PATH cl_itsamp02_sta drlto_1 libr=liblto destt=drive srct=server
devi=/dev/IBMtape0
DEFINE PATH cl_itsamp02_sta drlto_2 libr=liblto destt=drive srct=server
devi=/dev/IBMtape1
```

Configure Storage Agents as servers to the server

First we need to make the Tivoli Storage Manager server aware of the three Storage Agents. This can be done on the command line or via the Administration Center (AC). We show the configuration via the AC for the local Storage Agent of the first node, diomede. We configure the local Storage Agent for the second node, lochness, and the clustered Storage Agent in the same way with the appropriate parameters.

Within the AC, we choose the “Enterprise Management” to view the list of managed servers. We click the server TSMSRV03 (where the clustered Storage Agent will connect) as shown in Figure 15-1.

Archived

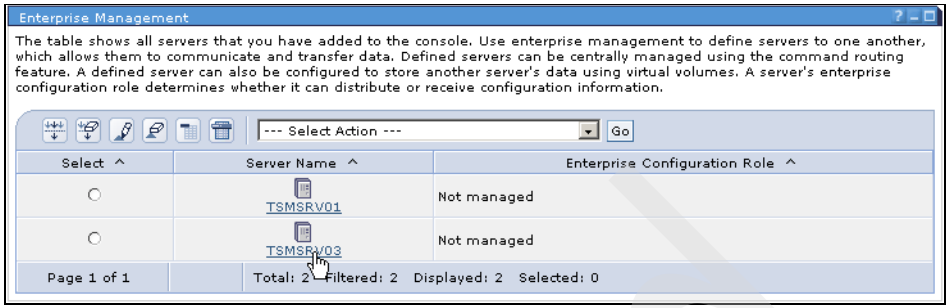

Figure 15-1 Selecting the server in the Enterprise Management panel

We can now open the list of servers defined to TSMSRV03. We choose “Define Server...” and click **Go** as shown in Figure 15-2.

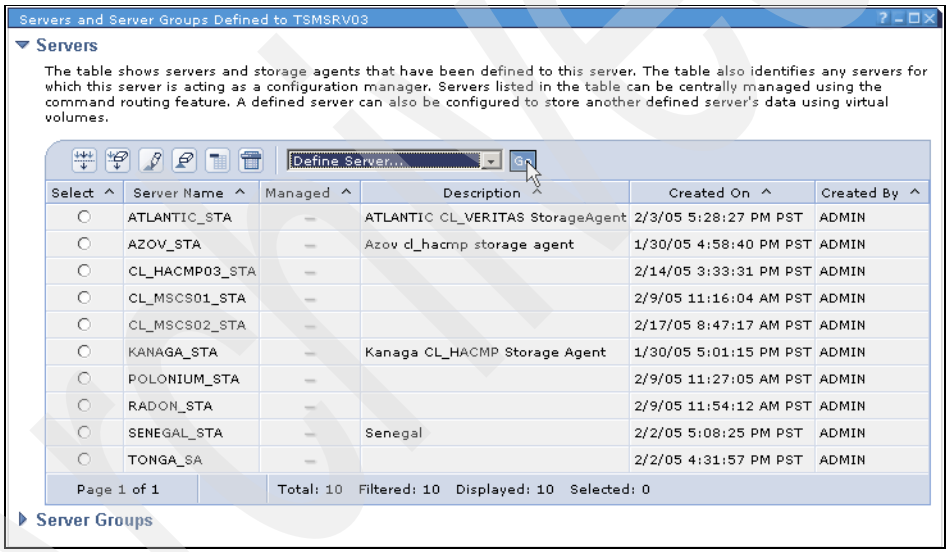

Figure 15-2 Servers and Server Groups defined to TSMSRV03

A wizard that will lead us through the configuration process is started as shown in Figure 15-3. We click **Next** to continue.

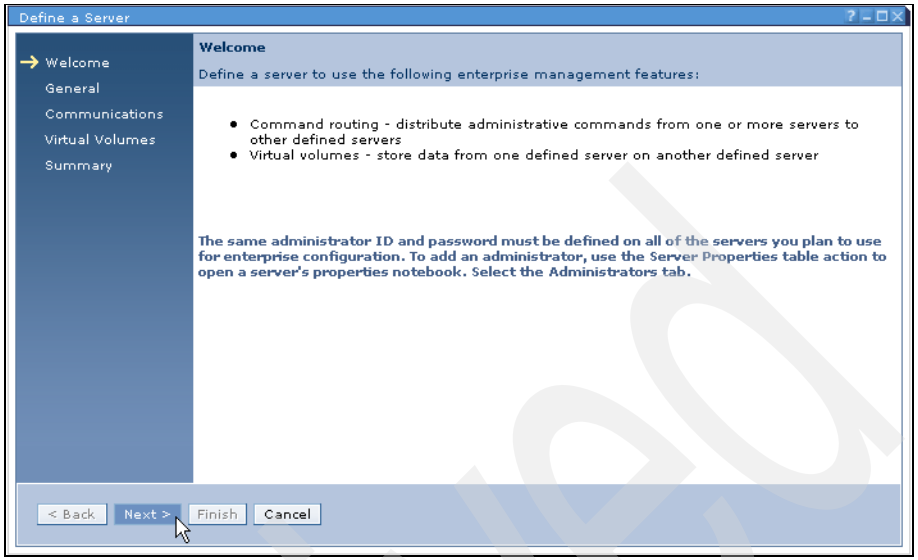

Figure 15-3 Define a Server - step one

We enter the server name of the Storage Agent, its password, and a description in the second step of the wizard as shown in Figure 15-4.

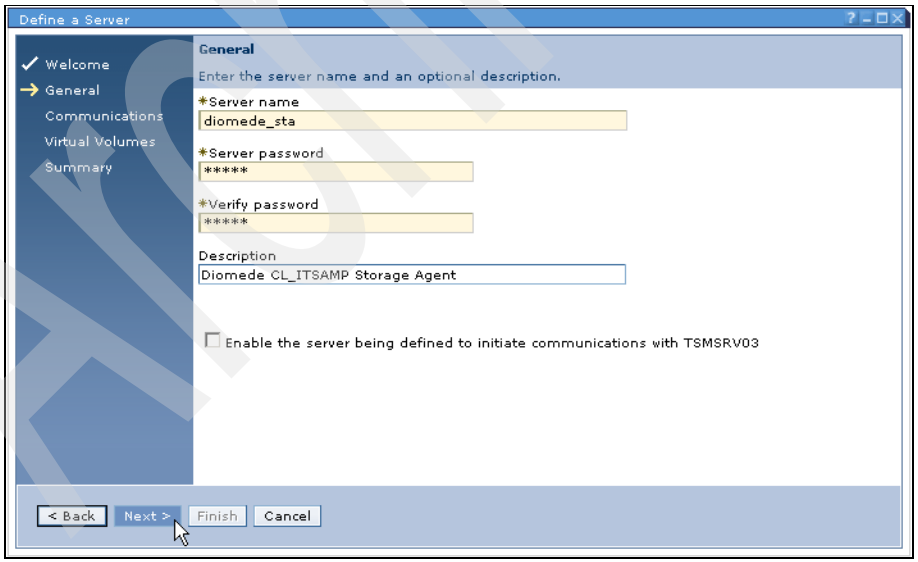

Figure 15-4 Define a Server - step two

Archived

In the next step we configure the TCP/IP address and port number and click **Next** as shown in Figure 15-5.

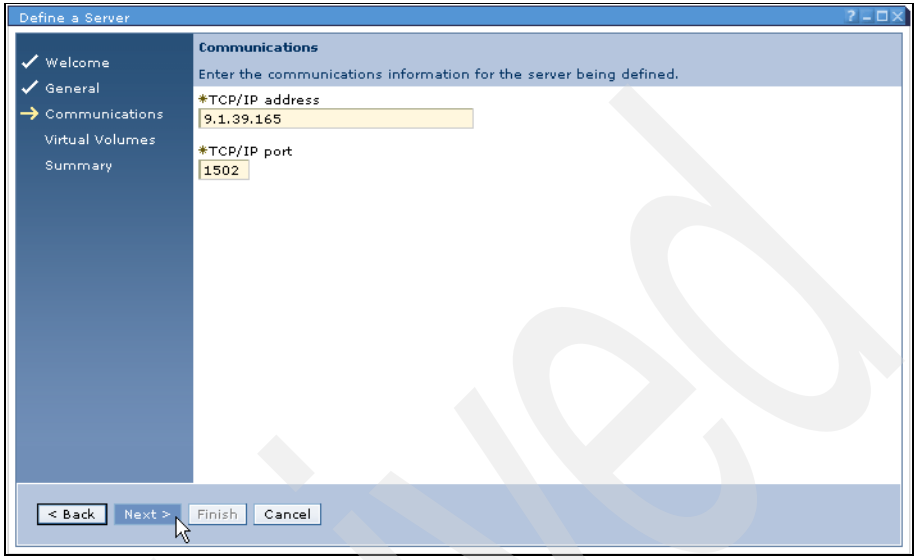

Figure 15-5 Define a Server - step three

We do not configure the use of virtual volumes, so we simply click **Next** as shown in Figure 15-6.

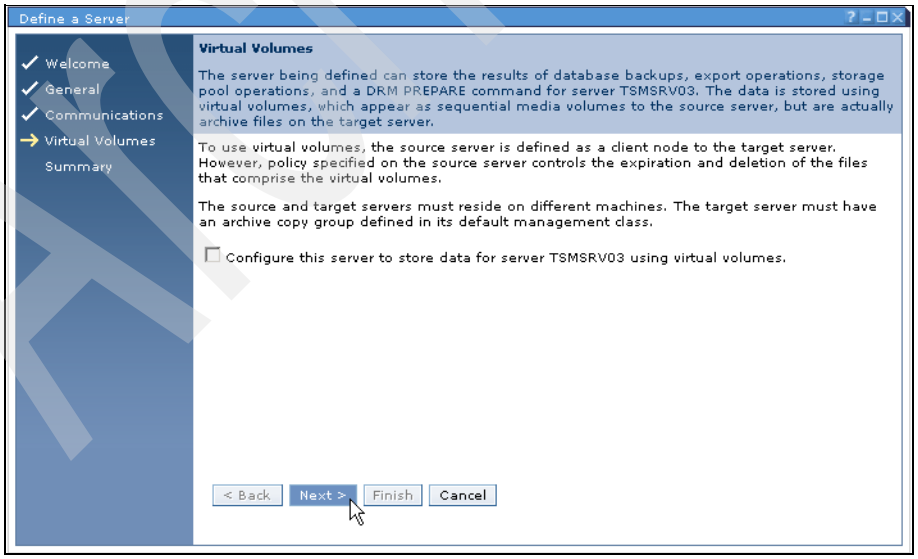

Figure 15-6 Define a Server - step four

Archived

We get a summary of the configured parameters to verify them. We click **Finish** as shown in Figure 15-7.

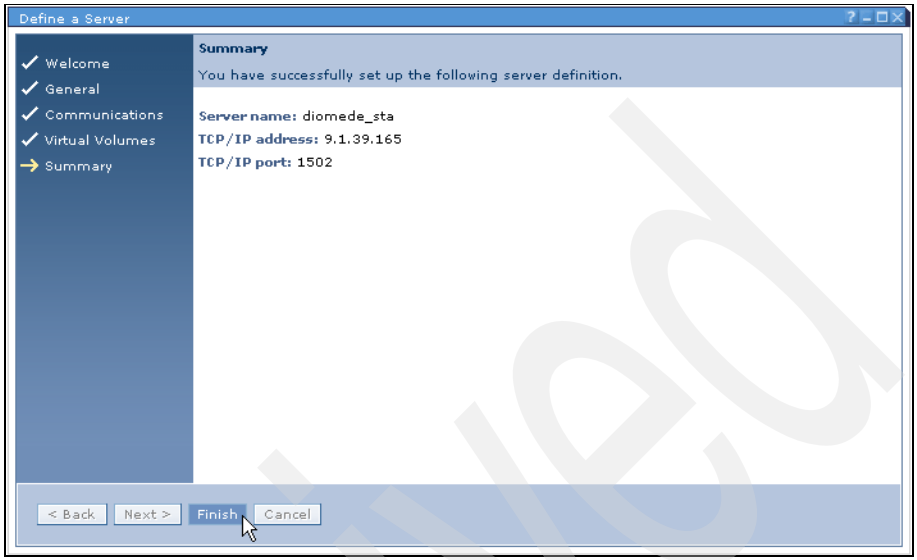

Figure 15-7 Define a Server - step five

Storage agent instances configuration

1. We set up three dsmsta.opt configuration files, located in the three different instance directories. We configure TCP/IP ports and devconfig file path according to our planning information in Table 15-1 on page 674. To create dsmsta.opt for the clustered instance, we mount the intended application resource shared disk on one node, diomede. There we create a directory to hold the Tivoli Storage Manager Storage Agent configuration files. In our case, the path is /mnt/nfsfiles/tsm/StorageAgent/bin, with the mount point for the filesystem being /mnt/nfsfiles. Example 15-2 shows the dsmsta.opt file for the clustered instance.

Example 15-2 Clustered instance /mnt/nfsfiles/tsm/StorageAgent/bin/dsmsta.opt

```
COMMmethod TCPIP
TCPPort 1504
DEVCONFIG /mnt/nfsfiles/tsm/StorageAgent/bin/devconfig.txt
```

2. We run the **`dsmsta setstorageserver command`** to populate the devconfig.txt and dsmsta.opt files for local instances. We run the it on both nodes with the appropriate values for the parameters. Example 15-3 shows the execution of the command on our first node, diomede. To verify the setup, we optionally issue the **`dsmsta`** command without any parameters. This starts the Storage Agent in foreground. We stop the Storage Agent with the **`halt`** command.

Archived

Example 15-3 The dsmsta setstorageserver command

```
[root@diomede root]# cd /opt/tivoli/tsm/StorageAgent/bin
[root@diomede bin]# dsmsta setstorageserver myname=diomede_sta \
mypassword=admin myhladdress=9.1.39.165 servername=tsmsrv03 \
serverpassword=password hladdress=9.1.39.74 lladdress=1500

Tivoli Storage Manager for Linux/i386
Version 5, Release 3, Level 0.0

Licensed Materials - Property of IBM

(C) Copyright IBM Corporation 1990, 2004.
All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

ANR7800I DSMSERV generated at 05:54:26 on Dec  6 2004.
ANR7801I Subsystem process ID is 18615.
ANR0900I Processing options file dsmsta.opt.
ANR4726I The ICC support module has been loaded.
ANR1432I Updating device configuration information to defined files.
ANR1433I Device configuration information successfully written to
/opt/tivoli/tsm/StorageAgent/bin/devconfig.txt.
ANR2119I The SERVERNAME option has been changed in the options file.
ANR0467I The SETSTORAGESERVER command completed successfully.
[root@diomede bin]#
```

3. For the clustered instance setup, we need to configure some environment variables. Example 15-4 shows the necessary steps to run the **dsmsta setstorageserver** command for the clustered instance. We can again use the **dsmsta** command without any parameters to verify the setup.

Example 15-4 The dsmsta setstorageserver command for clustered STA

```
[root@diomede root]# export \
> DSMSERV_CONFIG=/mnt/nfsfiles/tsm/StorageAgent/bin/dsmsta.opt
[root@diomede root]# export DSMSERV_DIR=/opt/tivoli/tsm/StorageAgent/bin
[root@diomede root]# cd /mnt/nfsfiles/tsm/StorageAgent/bin
[root@diomede bin]# dsmsta setstorageserver myname=cl_itsamp02_sta \
> mypassword=admin myhladdress=9.1.39.54 servername=tsmsrv03 \
> serverpassword=password hladdress=9.1.39.74 lladdress=1500
...
ANR0467I The SETSTORAGESERVER command completed successfully.
[root@diomede bin]#
```

4. We then review the results of running this command, which populates the devconfig.txt file as shown in Example 15-5.

Example 15-5 The devconfig.txt file

```
[root@diomede bin]# cat devconfig.txt
SET STANAME CL_ITSAMP02_STA
SET STAPASSWORD 21ff10f62b9caf883de8aa5ce017f536a1
SET STAHLADDRESS 9.1.39.54
DEFINE SERVER TSMSRV03 HLADDRESS=9.1.39.74 LLADDRESS=1500
SERVERPA=21911a57cfe832900b9c6f258aa0926124
[root@diomede bin]#
```

5. Next, we review the results of this update on the dsmsta.opt file. We see that the last line was updated with the servername, as seen in Example 15-6.

Example 15-6 Clustered Storage Agent dsmsta.opt

```
[root@diomede bin]# cat dsmsta.opt
COMMmethod TCPIP
TCPPort 1504
DEVCONFIG /mnt/nfsfiles/tsm/StorageAgent/bin/devconfig.txt

SERVERNAME TSMSRV03
[root@diomede bin]#
```

15.4.2 Client

1. We execute the following Tivoli Storage Manager commands on the Tivoli Storage Manager server tsmsrv03 to create three client nodes:

   ```
   register node diomede itsosj passexp=0
   register node lochness itsosj passexp=0
   register node cl_itsamp02_client itsosj passexp=0
   ```

2. We ensure that /mnt/nfsfiles is still mounted on diomede. We create a directory to hold the Tivoli Storage Manager client configuration files. In our case, the path is /mnt/nfsfiles/tsm/client/ba/bin.

3. We copy the default dsm.opt.smp to the shared disk directory as dsm.opt and edit the file with the servername to be used by this client instance. The contents of the file is shown in Example 15-7.

Example 15-7 dsm.opt file contents located in the application shared disk

```
************************************************************************
* IBM Tivoli Storage Manager                                           *
************************************************************************
*  This servername is the reference for the highly available TSM       *
*  client.                                                             *
************************************************************************
SErvername       tsmsrv03_san
```

4. We edit /opt/tivoli/tsm/client/ba/bin/dsm.sys on both nodes to configure server stanzas using the Storage Agent. Example 15-8 shows the server stanza for the clustered Tivoli Storage Manager client. This server stanza must be present in dsm.sys on both nodes. The stanzas for the local clients are only present in dsm.sys on the appropriate client. From now on we concentrate only on the clustered client. The setup of the local clients is the same as in a non-clustered environment.

Example 15-8 Server stanza in dsm.sys for the clustered client

```
* Server stanza for the ITSAMP highly available client to the atlantic (AIX)
* this will be a client which uses the LAN-free StorageAgent
SErvername                  tsmsrv03_san
   nodename                 cl_itsamp02_client
   COMMMethod               TCPip
   TCPPort                  1500
   TCPServeraddress         9.1.39.74
   HTTPPORT                 1582
   TCPClientaddress         9.1.39.54

   TXNBytelimit             256000
   resourceutilization      5
   enablelanfree            yes
   lanfreecommmethod        tcpip
   lanfreetcpport           1504
   lanfreetcpserveraddress  9.1.39.54

   passwordaccess           generate
   passworddir              /usr/sbin/rsct/sapolicies/nfsserver
   managedservices          schedule webclient
   schedmode                prompt
   schedlogname             /mnt/nfsfiles/tsm/client/ba/bin/dsmsched.log
   errorlogname             /mnt/nfsfiles/tsm/client/ba/bin/dsmerror.log
   ERRORLOGRETENTION        7

   domain                   /mnt/nfsfiles
   include                  /mnt/nfsfiles/.../*
```

Important: When domain statements, one or more, are used in a client configuration **only** that domains (file systems) will be backed up during incremental backup.

5. We issue this step again on both nodes. We connect to the Tivoli Storage Manager server using **`dsmc -server=tsmsrv03_san`** from the Linux command line. This will generate the TSM.PWD file as shown in Example 15-9.

Note: Tivoli Storage Manager for Linux Client encrypts the password file with the hostname. So it is necessary to create the password file locally on all nodes.

Example 15-9 Creation of the password file TSM.PWD

```
[root@diomede nfsserver]# pwd
/usr/sbin/rsct/sapolicies/nfsserver
[root@diomede nfsserver]# dsmc -se=tsmsrv03_san
IBM Tivoli Storage Manager
Command Line Backup/Archive Client Interface
  Client Version 5, Release 3, Level 0.0
  Client date/time: 02/18/2005 10:54:06
(c) Copyright by IBM Corporation and other(s) 1990, 2004. All Rights Reserved.

Node Name: CL_ITSAMP02_CLIENT
ANS9201W LAN-free path failed.
Node Name: CL_ITSAMP02_CLIENT
Please enter your user id <CL_ITSAMP02_CLIENT>:

Please enter password for user id "CL_ITSAMP02_CLIENT":

Session established with server TSMSRV03: AIX-RS/6000
  Server Version 5, Release 3, Level 0.0
  Server date/time: 02/18/2005 10:46:31  Last access: 02/18/2005 10:46:31

tsm> quit
[root@diomede nfsserver]# ls -l TSM.PWD
-rw-------    1 root     root          152 Feb 18 10:54 TSM.PWD
[root@diomede nfsserver]#
```

15.4.3 Resource configuration for the Storage Agent

The highly available Storage Agent instance will be used by the highly available Tivoli Storage Manager client instance.

Note: The approach we show here needs the library to be configured with the resetdrives option on the Tivoli Storage Manager server. For Tivoli Storage Manager V5.3, the AIX Tivoli Storage Manager server supports this new parameter. If you use a Tivoli Storage Manager server that does not support the resetdrives option you need also to configure the SCSI reset for the drives. You can use the same script that is used for clustering of the Tivoli Storage Manager server in Linux. Refer to “Requisites for using tape and medium changer devices” on page 629.

Archived

We configure the Tivoli System Automation for Multiplatforms resources for the Tivoli Storage Manager client and the Storage Agent by following these steps:

1. We change to the directory where the control scripts for the clustered application we want to back up are stored. In our example this is /usr/sbin/rsct/sapolicies/nfsserver/. Within this directory, we create symbolic links to the script which controls the Tivoli Storage Manager Client CAD and the Storage Agent in the Tivoli System Automation for Multiplatforms environment. We accomplish these steps on both nodes as shown in Example 15-10.

Example 15-10 Creation of the symbolic link that points to the Storage Agent script

```
[root@diomede root]# cd /usr/sbin/rsct/sapolicies/nfsserver
[root@diomede nfsserver]# ln -s \
> /usr/sbin/rsct/sapolicies/tsmclient/tsmclientctrl-cad nfsserverctrl-tsmclient
[root@diomede nfsserver]# ln -s \
> /usr/sbin/rsct/sapolicies/tsmclient/tsmstactrl-sta nfsserverctrl-tsmsta
[root@diomede nfsserver]#
```

2. We ensure that the resources of the cluster application resource group are offline. We use the Tivoli System Automation for Multiplatforms **lsrg -m** command on any node for this purpose. The output of the command is shown in Example 15-11.

Example 15-11 Output of the lsrg -m command before configuring the Storage Agent

```
Displaying Member Resource information:
Class:Resource:Node[ManagedResource]       Mandatory MemberOf         OpState
IBM.Application:SA-nfsserver-server        True      SA-nfsserver-rg  Offline
IBM.ServiceIP:SA-nfsserver-ip-1            True      SA-nfsserver-rg  Offline
IBM.Application:SA-nfsserver-data-nfsfiles True      SA-nfsserver-rg  Offline
```

3. The necessary resource for the Tivoli Storage Manager client CAD should depend on the Storage Agent resource. And the Storage Agent resource itself should depend on the NFS server resource of the clustered NFS server. In that way it is guaranteed that all necessary file systems are mounted before the Storage Agent or the Tivoli Storage Manager client CAD are started by Tivoli System Automation for Multiplatforms. To configure that behavior we do the following steps. We execute these steps only on the first node, diomede.

   a. We prepare the configuration file for the SA-nfsserver-tsmsta resource. All parameters for the StartCommand, StopCommand, and MonitorCommand must be on a single line in this file. Example 15-12 shows the contents of the file with line breaks between the parameters.

Example 15-12 Definition file SA-nfsserver-tsmsta.def

```
PersistentResourceAttributes::
```

```
Name=SA-nfsserver-tsmsta
ResourceType=1
StartCommand=/usr/sbin/rsct/sapolicies/nfsserver/nfsserverctrl-tsmsta start
/mnt/nfsfiles/tsm/StorageAgent/bin
/mnt/nfsfiles/tsm/StorageAgent/bin/dsmsta.opt SA-nfsserver-
StopCommand=/usr/sbin/rsct/sapolicies/nfsserver/nfsserverctrl-tsmsta stop
/mnt/nfsfiles/tsm/StorageAgent/bin
/mnt/nfsfiles/tsm/StorageAgent/bin/dsmsta.opt SA-nfsserver-
MonitorCommand=/usr/sbin/rsct/sapolicies/nfsserver/nfsserverctrl-tsmsta status
/mnt/nfsfiles/tsm/StorageAgent/bin
/mnt/nfsfiles/tsm/StorageAgent/bin/dsmsta.opt SA-nfsserver-
StartCommandTimeout=60
StopCommandTimeout=60
MonitorCommandTimeout=9
MonitorCommandPeriod=10
ProtectionMode=0
NodeNameList={'diomede','lochness'}
UserName=root
```

b. We prepare the configuration file for the SA-nfsserver-tsmclient resource. All parameters for the StartCommand, StopCommand, and MonitorCommand must be on a single line in this file. Example 15-13 shows the contents of the file with line breaks between the parameters.

> **Note:** We enter the nodename parameter for the StartCommand, StopCommand, and MonitorCommand in uppercase letters. This is necessary, as the nodename will be used for an SQL query in Tivoli Storage Manager. We also use an extra Tivoli Storage Manager user, called “scriptoperator”, which is necessary to query and reset Tivoli Storage Manager sessions. Be sure that this user can access the Tivoli Storage Manager server.

Example 15-13 Definition file SA-nfsserver-tsmclient.def

```
PersistentResourceAttributes::
Name=SA-nfsserver-tsmclient
ResourceType=1
StartCommand=/usr/sbin/rsct/sapolicies/nfsserver/nfsserverctrl-tsmclient start
/mnt/nfsfiles/tsm/client/ba/bin SA-nfsserver- CL_ITSAMP02_CLIENT tsmsrv03_san
scriptoperator password
StopCommand=/usr/sbin/rsct/sapolicies/nfsserver/nfsserverctrl-tsmclient stop
/mnt/nfsfiles/tsm/client/ba/bin SA-nfsserver- CL_ITSAMP02_CLIENT tsmsrv03_san
scriptoperator password
MonitorCommand=/usr/sbin/rsct/sapolicies/nfsserver/nfsserverctrl-tsmclient
status /mnt/nfsfiles/tsm/client/ba/bin SA-nfsserver- CL_ITSAMP02_CLIENT
tsmsrv03_san scriptoperator password
StartCommandTimeout=180
```

```
StopCommandTimeout=60
MonitorCommandTimeout=9
MonitorCommandPeriod=10
ProtectionMode=0
NodeNameList={'diomede','lochness'}
UserName=root
```

c. We manually add the SA-nfsserver-tsmsta and SA-nfsserver-tsmclient resources to Tivoli System Automation for Multiplatforms with the following commands:

```
mkrsrc -f SA-nfsserver-tsmsta.def IBM.Application
mkrsrc -f SA-nfsserver-tsmclient.def IBM.Application
```

d. Now that the resources are known by Tivoli System Automation for Multiplatforms, we add them to the resource group SA-nfsserver-rg with the commands:

```
addrgmbr -m T -g SA-nfsserver-rg IBM.Application:SA-nfsserver-tsmsta
addrgmbr -m T -g SA-nfsserver-rg IBM.Application:SA-nfsserver-tsmclient
```

e. We configure the dependency of the Storage Agent:

```
mkrel -S IBM.Application:SA-nfsserver-tsmsta -G
IBM.Application:SA-nfsserver-server -p DependsOn
SA-nfsserver-tsmsta-on-server
```

f. Finally we configure the dependency of the Tivoli Storage Manager Client:

```
mkrel -S IBM.Application:SA-nfsserver-tsmclient -G
IBM.Application:SA-nfsserver-tsmsta -p DependsOn
SA-nfsserver-tsmclient-on-tsmsta
```

We verify the relationships with the **lsrel command**. The output of the command is shown in Example 15-14.

Example 15-14 Output of the lsrel command

```
Displaying Managed Relations :

Name                                Class:Resource:Node[Source]
ResourceGroup[Source]
SA-nfsserver-server-on-data-nfsfiles IBM.Application:SA-nfsserver-server    SA-nfsserver-rg
SA-nfsserver-server-on-ip-1          IBM.Application:SA-nfsserver-server    SA-nfsserver-rg
SA-nfsserver-ip-on-nieq-1            IBM.ServiceIP:SA-nfsserver-ip-1        SA-nfsserver-rg
SA-nfsserver-tsmclient-on-tsmsta     IBM.Application:SA-nfsserver-tsmclient SA-nfsserver-rg
SA-nfsserver-tsmsta-on-server        IBM.Application:SA-nfsserver-tsmsta    SA-nfsserver-rg
```

4. Now we start the resource group with the **chrg -o online SA-nfsserver-rg** command.

5. To verify that all necessary resources are online, we use again the **lsrg -m** command. Example 15-15 shows the output of this command.

Example 15-15 Output of the lsrg -m command while resource group is online

```
Displaying Member Resource information:
Class:Resource:Node[ManagedResource]       Mandatory MemberOf         OpState
IBM.Application:SA-nfsserver-server        True      SA-nfsserver-rg  Online
IBM.ServiceIP:SA-nfsserver-ip-1            True      SA-nfsserver-rg  Online
IBM.Application:SA-nfsserver-data-nfsfiles True      SA-nfsserver-rg  Online
IBM.Application:SA-nfsserver-tsmsta        True      SA-nfsserver-rg  Online
IBM.Application:SA-nfsserver-tsmclient     True      SA-nfsserver-rg  Online
```

15.5 Testing the cluster

Here we show how we test the clustered Storage Agent environment.

15.5.1 Backup

For this first test, we do a failover during a LAN-free backup process.

Objective

The objective of this test is to show what happens when a LAN-free client incremental backup is started for a virtual node on the cluster using the Storage Agent created for this group, and the node that hosts the resources at that moment suddenly fails.

Activities

To do this test, we perform these tasks:

1. We use the /usr/sbin/rsct/sapolicies/bin/getstatus script to find out that the SA-nfsserver-rg resource group is online on our second node, diomede.
2. We schedule a client incremental backup operation using the Tivoli Storage Manager server scheduler and we associate the schedule to CL_ITSAMP02_CLIENT nodename.
3. At the scheduled time, the client starts to back up files as we can see in the schedule log file in Example 15-16 on page 687.

Example 15-16 Scheduled backup starts

```
02/25/2005 10:05:03 Scheduler has been started by Dsmcad.
02/25/2005 10:05:03 Querying server for next scheduled event.
02/25/2005 10:05:03 Node Name: CL_ITSAMP02_CLIENT
02/25/2005 10:05:03 Session established with server TSMSRV03: AIX-RS/6000
02/25/2005 10:05:03   Server Version 5, Release 3, Level 0.0
02/25/2005 10:05:03   Server date/time: 02/25/2005 10:05:03  Last access:
02/25/2005 10:01:02
```


Archived

```
02/25/2005 10:05:03 --- SCHEDULEREC QUERY BEGIN
02/25/2005 10:05:03 --- SCHEDULEREC QUERY END
02/25/2005 10:05:03 Next operation scheduled:
02/25/2005 10:05:03
------------------------------------------------------------
02/25/2005 10:05:03 Schedule Name:         INCR_BACKUP
02/25/2005 10:05:03 Action:                Incremental
02/25/2005 10:05:03 Objects:
02/25/2005 10:05:03 Options:               -subdir=yes
02/25/2005 10:05:03 Server Window Start:   10:05:00 on 02/25/2005
02/25/2005 10:05:03
------------------------------------------------------------
02/25/2005 10:05:03
Executing scheduled command now.
02/25/2005 10:05:03 --- SCHEDULEREC OBJECT BEGIN INCR_BACKUP 02/25/2005
10:05:00
02/25/2005 10:05:03 Incremental backup of volume '/mnt/nfsfiles'
02/25/2005 10:05:04 Directory-->              4,096 /mnt/nfsfiles/ [Sent]
02/25/2005 10:05:04 Directory-->             16,384 /mnt/nfsfiles/lost+found
[Sent]
02/25/2005 10:05:05 ANS1898I ***** Processed       500 files *****
02/25/2005 10:05:05 Directory-->              4,096 /mnt/nfsfiles/root [Sent]
02/25/2005 10:05:05 Directory-->              4,096 /mnt/nfsfiles/tsm [Sent]
[...]
02/25/2005 10:05:07 Normal File-->          341,631
/mnt/nfsfiles/root/ibmtape/IBMtape-1.5.3-2.4.21-15.EL.i386.rpm [Sent]
[...]
02/25/2005 10:05:07 ANS1114I Waiting for mount of offline media.
02/25/2005 10:05:08 ANS1898I ***** Processed     1,500 files *****
02/25/2005 10:05:08 Retry # 1  Directory-->               4,096 /mnt/nfsfiles/
[Sent]
02/25/2005 10:05:08 Retry # 1  Directory-->              16,384
/mnt/nfsfiles/lost+found [Sent]
02/25/2005 10:05:08 Retry # 1  Directory-->               4,096
/mnt/nfsfiles/root [Sent]
02/25/2005 10:05:08 Retry # 1  Directory-->               4,096
/mnt/nfsfiles/tsm [Sent]
[...]
02/25/2005 10:06:11 Retry # 1  Normal File-->           341,631
/mnt/nfsfiles/root/ibmtape/IBMtape-1.5.3-2.4.21-15.EL.i386.rpm
 [Sent]
```

4. The client session for CL_ITSAMP02_CLIENT nodename starts on the server. At the same time, several sessions are also started for CL_ITSAMP02_STA for Tape Library Sharing and the Storage Agent prompts the Tivoli Storage Manager server to mount a tape volume, as we can see in Example 15-17.

Example 15-17 Activity log when scheduled backup starts

```
02/25/05   10:05:03  ANR0406I Session 1319 started for node CL_ITSAMP02_CLIENT
                      (Linux86) (Tcp/Ip 9.1.39.165(33850)). (SESSION: 1319)
02/25/05   10:05:04  ANR0406I Session 1320 started for node CL_ITSAMP02_CLIENT
                      (Linux86) (Tcp/Ip 9.1.39.165(33852)). (SESSION: 1320)
02/25/05   10:05:04  ANR0406I (Session: 1312, Origin: CL_ITSAMP02_STA)  Session
                      8 started for node CL_ITSAMP02_CLIENT (Linux86) (Tcp/Ip
                      dhcp39054.almaden.ibm.com(33853)). (SESSION: 1312)
02/25/05   10:05:04  ANR0408I Session 1321 started for server CL_ITSAMP02_STA
                      (Linux/i386) (Tcp/Ip) for storage agent.  (SESSION: 1321)
02/25/05   10:05:04  ANR0408I (Session: 1312, Origin: CL_ITSAMP02_STA)  Session
                      9 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      storage agent.  (SESSION: 1312)
02/25/05   10:05:04  ANR0415I Session 1321 proxied by CL_ITSAMP02_STA started
                      for node CL_ITSAMP02_CLIENT. (SESSION: 1321)
02/25/05   10:05:04  ANR0408I Session 1322 started for server CL_ITSAMP02_STA
                      (Linux/i386) (Tcp/Ip) for library sharing.  (SESSION:
                      1322)
02/25/05   10:05:04  ANR0408I (Session: 1312, Origin: CL_ITSAMP02_STA)  Session
                      10 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 1312)
02/25/05   10:05:04  ANR0409I (Session: 1312, Origin: CL_ITSAMP02_STA)  Session
                      10 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                      1312)
02/25/05   10:05:04  ANR0409I Session 1322 ended for server CL_ITSAMP02_STA
                      (Linux/i386). (SESSION: 1322)
02/25/05   10:05:07  ANR0408I Session 1323 started for server CL_ITSAMP02_STA
                      (Linux/i386) (Tcp/Ip) for library sharing.  (SESSION:
                      1323)
02/25/05   10:05:07  ANR0408I (Session: 1312, Origin: CL_ITSAMP02_STA)  Session
                      11 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 1312)
02/25/05   10:05:15  ANR0406I Session 1324 started for node CL_ITSAMP02_CLIENT
                      (Linux86) (Tcp/Ip 9.1.39.165(33858)). (SESSION: 1324)
02/25/05   10:05:15  ANR0406I (Session: 1312, Origin: CL_ITSAMP02_STA)  Session
                      13 started for node CL_ITSAMP02_CLIENT (Linux86) (Tcp/Ip
                      dhcp39054.almaden.ibm.com(33859)). (SESSION: 1312)
02/25/05   10:05:15  ANR0408I Session 1325 started for server CL_ITSAMP02_STA
                      (Linux/i386) (Tcp/Ip) for storage agent.  (SESSION: 1325)
02/25/05   10:05:15  ANR0408I (Session: 1312, Origin: CL_ITSAMP02_STA)  Session
                      14 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      storage agent.  (SESSION: 1312)
02/25/05   10:05:15  ANR0415I Session 1325 proxied by CL_ITSAMP02_STA started
                      for node CL_ITSAMP02_CLIENT. (SESSION: 1325)
02/25/05   10:05:16  ANR0409I (Session: 1312, Origin: CL_ITSAMP02_STA)  Session
                      14 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                      1312)
02/25/05   10:05:16  ANR0403I (Session: 1312, Origin: CL_ITSAMP02_STA)  Session
```

Archived

```
                      13 ended for node CL_ITSAMP02_CLIENT (Linux86). (SESSION:
                      1312)
02/25/05   10:05:17  ANR0403I Session 1324 ended for node CL_ITSAMP02_CLIENT
                      (Linux86). (SESSION: 1324)
02/25/05   10:05:17  ANR0403I Session 1325 ended for node CL_ITSAMP02_CLIENT
                      (Linux86). (SESSION: 1325)
```

5. After a few seconds the Tivoli Storage Manager server mounts the tape volume 030AKK in drive DRLTO_2, and it informs the Storage Agent about the drive where the volume is mounted. The Storage Agent CL_ITSAMP02_STA opens then the tape volume as an output volume and starts sending data to the DRLTO_2 as shown in Example 15-18.

Example 15-18 Activity log when tape is mounted

```
02/25/05   10:05:34  ANR8337I LTO volume 030AKK mounted in drive DRLTO_2
                      (/dev/rmt1). (SESSION: 1323)
02/25/05   10:05:34  ANR0409I (Session: 1312, Origin: CL_ITSAMP02_STA)  Session
                      11 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                      1312)
02/25/05   10:05:34  ANR0409I Session 1323 ended for server CL_ITSAMP02_STA
                      (Linux/i386). (SESSION: 1323)
02/25/05   10:05:34  ANR2997W The server log is 85 percent full. The server
                      will delay transactions by 3 milliseconds.
02/25/05   10:05:34  ANR8337I (Session: 1312, Origin: CL_ITSAMP02_STA)  LTO
                      volume 030AKK mounted in drive DRLTO_2 (/dev/IBMtape1).
                      (SESSION: 1312)
02/25/05   10:05:34  ANR0511I Session 1321 opened output volume 030AKK.
                      (SESSION: 1321)
02/25/05   10:05:34  ANR0511I (Session: 1312, Origin: CL_ITSAMP02_STA)  Session
                      9 opened output volume 030AKK. (SESSION: 1312)
```

6. While the client is backing up the files, we execute a manual failover to lochness by executing the command **samctrl -u a diomede**. This command adds diomede to the list of excluded nodes, which leads to a failover. The Storage Agent and the client are stopped on diomede. We get a message in the activity log of the server, indicating that the session was severed, as shown in Example 15-19.

Example 15-19 Activity log when failover takes place

```
02/25/05   10:06:57  ANR3605E Unable to communicate with storage agent.
                      (SESSION: 1314)
02/25/05   10:06:57  ANR3605E Unable to communicate with storage agent.
                      (SESSION: 1311)
02/25/05   10:06:59  ANR0480W Session 1321 for node CL_ITSAMP02_CLIENT
                      (Linux86) terminated - connection with client severed.
                      (SESSION: 1321)
```

The tape volume is still mounted in tape drive DRLTO_2.

7. Resources are brought online on our second node, lochness. During startup of the SA-nfsserver-tsmclient resource, the tsmclientctrl-cad script searches for old sessions to cancel as shown in the activity log in Example 15-20. Refer to “Tivoli Storage Manager client resource configuration” on page 660 for detailed information about why we need to cancel old sessions.

Example 15-20 Activity log when tsmclientctrl-cad script searches for old sessions

```
02/25/05   10:07:18  ANR0407I Session 1332 started for administrator
                      SCRIPTOPERATOR (Linux86) (Tcp/Ip 9.1.39.167(33081)).
                      (SESSION: 1332)
02/25/05   10:07:18  ANR2017I Administrator SCRIPTOPERATOR issued command:
                      select SESSION_ID,CLIENT_NAME from SESSIONS where
                      CLIENT_NAME='CL_ITSAMP02_CLIENT' (SESSION: 1332)
02/25/05   10:07:18  ANR2034E SELECT: No match found using this criteria.
                      (SESSION: 1332)
02/25/05   10:07:18  ANR2017I Administrator SCRIPTOPERATOR issued command:
                      ROLLBACK  (SESSION: 1332)
02/25/05   10:07:18  ANR0405I Session 1332 ended for administrator
                      SCRIPTOPERATOR (Linux86). (SESSION: 1332)
```

8. The CAD is started on lochness as shown in dsmwebcl.log in Example 15-21.

Example 15-21 dsmwebcl.log when the CAD starts

```
02/25/2005 10:07:18 (dsmcad) IBM Tivoli Storage Manager
02/25/2005 10:07:18 (dsmcad) Client Acceptor - Built Dec  7 2004 10:24:17
02/25/2005 10:07:18 (dsmcad) Version 5, Release 3, Level 0.0
02/25/2005 10:07:18 (dsmcad) Dsmcad is working in   Webclient   Schedule
mode.
02/25/2005 10:07:18 (dsmcad) ANS3000I HTTP communications available on port
1582.
02/25/2005 10:07:18 (dsmcad) Command will be executed in 1 minute.
```

9. The CAD connects to the Tivoli Storage Manager server. This is logged in the actlog as shown in Example 15-22.

Example 15-22 Actlog when CAD connects to the server

```
02/25/05   10:07:19  ANR0406I Session 1333 started for node CL_ITSAMP02_CLIENT
                      (Linux86) (Tcp/Ip 9.1.39.167(33083)). (SESSION: 1333)
02/25/05   10:07:19  ANR1639I Attributes changed for node CL_ITSAMP02_CLIENT:
                      TCP Name from diomede to lochness, TCP Address from  to
                      9.1.39.167, GUID from b4.cc.54.42.fb.6b.d9.11.ab.61.00.0-
                      d.60.49.4c.39 to 22.77.12.20.fc.6b.d9.11.84.80.00.0d.60.-
                      49.6a.62. (SESSION: 1333)
02/25/05   10:07:19  ANR0403I Session 1333 ended for node CL_ITSAMP02_CLIENT
                      (Linux86). (SESSION: 1333)
```

10. Now that the Storage Agent is also up it connects to the Tivoli Storage Manager server, too. The tape volume is now unmounted as shown in Example 15-23.

Example 15-23 Actlog when Storage Agent connects to the server

```
02/25/05   10:07:35  ANR0408I (Session: 1328, Origin: CL_ITSAMP02_STA)  Session
                      7 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 1328)
02/25/05   10:07:35  ANR0408I Session 1334 started for server CL_ITSAMP02_STA
                      (Linux/i386) (Tcp/Ip) for library sharing.  (SESSION:
                      1334)
02/25/05   10:07:35  ANR0409I Session 1334 ended for server CL_ITSAMP02_STA
                      (Linux/i386). (SESSION: 1334)
02/25/05   10:07:35  ANR0409I (Session: 1328, Origin: CL_ITSAMP02_STA)  Session
                      7 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                      1328)
02/25/05   10:07:35  ANR8336I Verifying label of LTO volume 030AKK in drive
                      DRLTO_2 (/dev/rmt1). (SESSION: 1323)
02/25/05   10:08:11  ANR8468I LTO volume 030AKK dismounted from drive DRLTO_2
                      (/dev/rmt1) in library LIBLTO. (SESSION: 1323)
```

11. The backup schedule is restarted as shown in the schedule log in Example 15-24.

Example 15-24 Schedule log when schedule is restarted

```
02/25/2005 10:08:19 --- SCHEDULEREC QUERY BEGIN
02/25/2005 10:08:19 --- SCHEDULEREC QUERY END
02/25/2005 10:08:19 Next operation scheduled:
02/25/2005 10:08:19
------------------------------------------------------------
02/25/2005 10:08:19 Schedule Name:         INCR_BACKUP
02/25/2005 10:08:19 Action:                Incremental
02/25/2005 10:08:19 Objects:
02/25/2005 10:08:19 Options:               -subdir=yes
02/25/2005 10:08:19 Server Window Start:   10:05:00 on 02/25/2005
02/25/2005 10:08:19
------------------------------------------------------------
02/25/2005 10:08:19
Executing scheduled command now.
02/25/2005 10:08:19 --- SCHEDULEREC OBJECT BEGIN INCR_BACKUP 02/25/2005
10:05:00
02/25/2005 10:08:19 Incremental backup of volume '/mnt/nfsfiles'
02/25/2005 10:08:21 ANS1898I ***** Processed       500 files *****
02/25/2005 10:08:22 ANS1898I ***** Processed     1,500 files *****
02/25/2005 10:08:22 ANS1898I ***** Processed     3,500 files *****
[...]
```

The tape volume is mounted again as shown in the activity log in Example 15-25.

Example 15-25 Activity log when the tape volume is mounted again

```
02/25/05   10:08:19  ANR0406I Session 1335 started for node CL_ITSAMP02_CLIENT
                      (Linux86) (Tcp/Ip 9.1.39.167(33091)). (SESSION: 1335)
02/25/05   10:08:19  ANR1639I Attributes changed for node CL_ITSAMP02_CLIENT:
                      TCP Address from 9.1.39.167 to . (SESSION: 1335)
02/25/05   10:08:22  ANR0406I Session 1336 started for node CL_ITSAMP02_CLIENT
                      (Linux86) (Tcp/Ip 9.1.39.167(33093)). (SESSION: 1336)
02/25/05   10:08:22  ANR0406I (Session: 1328, Origin: CL_ITSAMP02_STA)  Session
                      10 started for node CL_ITSAMP02_CLIENT (Linux86) (Tcp/Ip
                      dhcp39054.almaden.ibm.com(33094)). (SESSION: 1328)
02/25/05   10:08:22  ANR2997W The server log is 85 percent full. The server
                      will delay transactions by 3 milliseconds.
02/25/05   10:08:22  ANR0408I Session 1337 started for server CL_ITSAMP02_STA
                      (Linux/i386) (Tcp/Ip) for storage agent.  (SESSION: 1337)
02/25/05   10:08:23  ANR0408I (Session: 1328, Origin: CL_ITSAMP02_STA)  Session
                      11 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      storage agent.  (SESSION: 1328)
02/25/05   10:08:23  ANR0415I Session 1337 proxied by CL_ITSAMP02_STA started
                      for node CL_ITSAMP02_CLIENT. (SESSION: 1337)
02/25/05   10:08:23  ANR0408I Session 1338 started for server CL_ITSAMP02_STA
                      (Linux/i386) (Tcp/Ip) for library sharing.  (SESSION:
                      1338)
02/25/05   10:08:23  ANR0408I (Session: 1328, Origin: CL_ITSAMP02_STA)  Session
                      12 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 1328)
02/25/05   10:08:23  ANR0409I (Session: 1328, Origin: CL_ITSAMP02_STA)  Session
                      12 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                      1328)
02/25/05   10:08:23  ANR0409I Session 1338 ended for server CL_ITSAMP02_STA
                      (Linux/i386). (SESSION: 1338)
02/25/05   10:08:23  ANR0408I Session 1339 started for server CL_ITSAMP02_STA
                      (Linux/i386) (Tcp/Ip) for library sharing.  (SESSION:
                      1339)
02/25/05   10:08:23  ANR0408I (Session: 1328, Origin: CL_ITSAMP02_STA)  Session
                      13 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 1328)
02/25/05   10:08:31  ANR0406I Session 1340 started for node CL_ITSAMP02_CLIENT
                      (Linux86) (Tcp/Ip 9.1.39.167(33099)). (SESSION: 1340)
02/25/05   10:08:31  ANR0406I (Session: 1328, Origin: CL_ITSAMP02_STA)  Session
                      15 started for node CL_ITSAMP02_CLIENT (Linux86) (Tcp/Ip
                      dhcp39054.almaden.ibm.com(33100)). (SESSION: 1328)
02/25/05   10:08:31  ANR0408I Session 1341 started for server CL_ITSAMP02_STA
                      (Linux/i386) (Tcp/Ip) for storage agent.  (SESSION: 1341)
02/25/05   10:08:31  ANR0408I (Session: 1328, Origin: CL_ITSAMP02_STA)  Session
                      16 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
```

```
                    storage agent.  (SESSION: 1328)
02/25/05   10:08:31  ANR0415I Session 1341 proxied by CL_ITSAMP02_STA started
                     for node CL_ITSAMP02_CLIENT. (SESSION: 1341)
02/25/05   10:08:33  ANR0409I (Session: 1328, Origin: CL_ITSAMP02_STA)  Session
                     16 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                     1328)
02/25/05   10:08:33  ANR0403I (Session: 1328, Origin: CL_ITSAMP02_STA)  Session
                     15 ended for node CL_ITSAMP02_CLIENT (Linux86). (SESSION:
                     1328)
02/25/05   10:08:33  ANR0403I Session 1340 ended for node CL_ITSAMP02_CLIENT
                     (Linux86). (SESSION: 1340)
02/25/05   10:08:33  ANR0403I Session 1341 ended for node CL_ITSAMP02_CLIENT
                     (Linux86). (SESSION: 1341)
02/25/05   10:08:49  ANR8337I LTO volume 030AKK mounted in drive DRLTO_1
                     (/dev/rmt0). (SESSION: 1339)
02/25/05   10:08:49  ANR0409I (Session: 1328, Origin: CL_ITSAMP02_STA)  Session
                     13 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                     1328)
02/25/05   10:08:49  ANR0409I Session 1339 ended for server CL_ITSAMP02_STA
                     (Linux/i386). (SESSION: 1339)
02/25/05   10:08:49  ANR8337I (Session: 1328, Origin: CL_ITSAMP02_STA)  LTO
                     volume 030AKK mounted in drive DRLTO_1 (/dev/IBMtape0).
                     (SESSION: 1328)
02/25/05   10:08:49  ANR0511I Session 1337 opened output volume 030AKK.
                     (SESSION: 1337)
02/25/05   10:08:49  ANR0511I (Session: 1328, Origin: CL_ITSAMP02_STA)  Session
                     11 opened output volume 030AKK. (SESSION: 1328)
```

12. The backup finishes successfully as shown in the schedule log in Example 15-26. We remove diomede from the list of excluded nodes with the **samctrl -u d diomede** command.

Example 15-26 Schedule log shows that the schedule completed successfully

```
02/25/2005 10:17:41 --- SCHEDULEREC OBJECT END INCR_BACKUP 02/25/2005 10:05:00
02/25/2005 10:17:41 Scheduled event 'INCR_BACKUP' completed successfully.
02/25/2005 10:17:41 Sending results for scheduled event 'INCR_BACKUP'.
02/25/2005 10:17:42 Results sent to server for scheduled event 'INCR_BACKUP'.
```

Results summary

The test results show that after a failure on the node that hosts both the Tivoli Storage Manager client scheduler as well as the Storage Agent shared resources, a scheduled incremental backup started on one node for LAN-free is restarted and successfully completed on the other node, also using the SAN path.

This is true if the startup window used to define the schedule is not elapsed when the scheduler service restarts on the second node.

The Tivoli Storage Manager server on AIX resets the SCSI bus when the Storage Agent is restarted on the second node. This permits us to dismount the tape volume from the drive where it was mounted before the failure. When the client restarts the LAN-free operation the same Storage Agent commands the server to mount again the tape volume to continue the backup.

15.5.2 Restore

Our second test is a scheduled restore using the SAN path while a failover takes place.

Objective

The objective of this test is to show what happens when a LAN-free restore is started for a virtual node on the cluster, and the node that hosts the resources at that moment suddenly fails.

Activities

To do this test, we perform these tasks:

1. We use the /usr/sbin/rsct/sapolicies/bin/getstatus script to find out that the SA-nfsserver-rg resource group is online on our first node, diomede.
2. We schedule a client restore operation using the Tivoli Storage Manager server scheduler and we associate the schedule to CL_ITSAMP02_CLIENT nodename.
3. At the scheduled time, the client starts the restore as shown in the schedule log in Example 15-27.

Example 15-27 Scheduled restore starts

```
02/25/2005 11:50:42 Scheduler has been started by Dsmcad.
02/25/2005 11:50:42 Querying server for next scheduled event.
02/25/2005 11:50:42 Node Name: CL_ITSAMP02_CLIENT
02/25/2005 11:50:42 Session established with server TSMSRV03: AIX-RS/6000
02/25/2005 11:50:42   Server Version 5, Release 3, Level 0.0
02/25/2005 11:50:42   Server date/time: 02/25/2005 11:50:42  Last access:
02/25/2005 11:48:41

02/25/2005 11:50:42 --- SCHEDULEREC QUERY BEGIN
02/25/2005 11:50:42 --- SCHEDULEREC QUERY END
02/25/2005 11:50:42 Next operation scheduled:
02/25/2005 11:50:42
------------------------------------------------------------
```

```
02/25/2005 11:50:42 Schedule Name:         RESTORE_ITSAMP
02/25/2005 11:50:42 Action:                Restore
02/25/2005 11:50:42 Objects:               /mnt/nfsfiles/root/*.*
02/25/2005 11:50:42 Options:               -subdir=yes
02/25/2005 11:50:42 Server Window Start:   11:50:00 on 02/25/2005
02/25/2005 11:50:42
--------------------------------------------------------------
02/25/2005 11:50:42
Executing scheduled command now.
02/25/2005 11:50:42 --- SCHEDULEREC OBJECT BEGIN RESTORE_ITSAMP 02/25/2005
11:50:00
02/25/2005 11:50:42 Restore function invoked.

02/25/2005 11:50:43 ANS1899I ***** Examined     1,000 files *****
02/25/2005 11:50:43 ANS1899I ***** Examined     2,000 files *****
[...]
02/25/2005 11:51:21 Restoring           4,096
/mnt/nfsfiles/root/tsmi686/cdrom/noarch [Done]
02/25/2005 11:51:21  ** Interrupted **
02/25/2005 11:51:21 ANS1114I Waiting for mount of offline media.
02/25/2005 11:52:25 Restoring             161 /mnt/nfsfiles/root/.ICEauthority
[Done]
[...]
```

4. A session for CL_ITSAMP02_CLIENT nodename starts on the server. At the same time several sessions are also started for CL_ITSAMP02_STA for Tape Library Sharing and the Storage Agent prompts the Tivoli Storage Manager server to mount a tape volume. The tape volume is mounted in drive DRLTO_2. All of these messages in the actlog are shown in Example 15-28.

Example 15-28 Actlog when the schedule restore starts

```
02/25/05   11:50:42  ANR0406I Session 1391 started for node CL_ITSAMP02_CLIENT
                      (Linux86) (Tcp/Ip 9.1.39.165(33913)). (SESSION: 1391)
02/25/05   11:50:45  ANR0406I (Session: 1367, Origin: CL_ITSAMP02_STA)  Session
                      15 started for node CL_ITSAMP02_CLIENT (Linux86) (Tcp/Ip
                      dhcp39054.almaden.ibm.com(33914)). (SESSION: 1367)
02/25/05   11:50:45  ANR0408I Session 1392 started for server CL_ITSAMP02_STA
                      (Linux/i386) (Tcp/Ip) for storage agent.  (SESSION: 1392)
02/25/05   11:50:45  ANR0408I (Session: 1367, Origin: CL_ITSAMP02_STA)  Session
                      16 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      storage agent.  (SESSION: 1367)
02/25/05   11:50:45  ANR0415I Session 1392 proxied by CL_ITSAMP02_STA started
                      for node CL_ITSAMP02_CLIENT. (SESSION: 1392)
02/25/05   11:51:17  ANR0408I Session 1393 started for server CL_ITSAMP02_STA
                      (Linux/i386) (Tcp/Ip) for library sharing.  (SESSION:
                      1393)
02/25/05   11:51:17  ANR0408I (Session: 1367, Origin: CL_ITSAMP02_STA)  Session
                      17 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
```

```
                    library sharing.  (SESSION: 1367)
02/25/05   11:51:17  ANR0409I (Session: 1367, Origin: CL_ITSAMP02_STA)  Session
                     17 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                     1367)
02/25/05   11:51:17  ANR0409I Session 1393 ended for server CL_ITSAMP02_STA
                     (Linux/i386). (SESSION: 1393)
02/25/05   11:51:17  ANR0408I Session 1394 started for server CL_ITSAMP02_STA
                     (Linux/i386) (Tcp/Ip) for library sharing.  (SESSION:
                     1394)
02/25/05   11:51:17  ANR0408I (Session: 1367, Origin: CL_ITSAMP02_STA)  Session
                     18 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                     library sharing.  (SESSION: 1367)
02/25/05   11:51:17  ANR0409I Session 1394 ended for server CL_ITSAMP02_STA
                     (Linux/i386). (SESSION: 1394)
02/25/05   11:51:17  ANR0409I (Session: 1367, Origin: CL_ITSAMP02_STA)  Session
                     18 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                     1367)
02/25/05   11:51:21  ANR0408I Session 1395 started for server CL_ITSAMP02_STA
                     (Linux/i386) (Tcp/Ip) for library sharing.  (SESSION:
                     1395)
02/25/05   11:51:21  ANR0408I (Session: 1367, Origin: CL_ITSAMP02_STA)  Session
                     19 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                     library sharing.  (SESSION: 1367)
02/25/05   11:51:47  ANR8337I LTO volume 030AKK mounted in drive DRLTO_2
                     (/dev/rmt1). (SESSION: 1395)
02/25/05   11:51:48  ANR0409I (Session: 1367, Origin: CL_ITSAMP02_STA)  Session
                     19 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                     1367)
02/25/05   11:51:48  ANR0409I Session 1395 ended for server CL_ITSAMP02_STA
                     (Linux/i386). (SESSION: 1395)
02/25/05   11:51:48  ANR8337I (Session: 1367, Origin: CL_ITSAMP02_STA)  LTO
                     volume 030AKK mounted in drive DRLTO_2 (/dev/IBMtape1).
                     (SESSION: 1367)
02/25/05   11:51:48  ANR0510I (Session: 1367, Origin: CL_ITSAMP02_STA)  Session
                     15 opened input volume 030AKK. (SESSION: 1367)
```

5. While the client is restoring the files, we execute a manual failover to lochness by executing the command **samctrl -u a diomede**. This command adds diomede to the list of excluded nodes, which leads to a failover. The Storage Agent and the client are stopped on diomede. We get a message in the activity log of the server, indicating that the session was severed, as shown in Example 15-29.

Example 15-29 Actlog when resources are stopped at diomede

```
02/25/05   11:53:14  ANR0403I Session 1391 ended for node CL_ITSAMP02_CLIENT
                     (Linux86). (SESSION: 1391)
02/25/05   11:53:14  ANR0514I (Session: 1367, Origin: CL_ITSAMP02_STA)  Session
                     15 closed volume 030AKK. (SESSION: 1367)
```

```
02/25/05   11:53:14  ANR0408I Session 1397 started for server CL_ITSAMP02_STA
                      (Linux/i386) (Tcp/Ip) for library sharing.  (SESSION:
                      1397)
02/25/05   11:53:14  ANR0408I (Session: 1367, Origin: CL_ITSAMP02_STA)  Session
                      20 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 1367)
02/25/05   11:53:14  ANR0409I (Session: 1367, Origin: CL_ITSAMP02_STA)  Session
                      20 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                      1367)
02/25/05   11:53:14  ANR0408I Session 1398 started for server CL_ITSAMP02_STA
                      (Linux/i386) (Tcp/Ip) for library sharing.  (SESSION:
                      1398)
02/25/05   11:53:14  ANR0409I Session 1397 ended for server CL_ITSAMP02_STA
                      (Linux/i386). (SESSION: 1397)
02/25/05   11:53:14  ANR0408I (Session: 1367, Origin: CL_ITSAMP02_STA)  Session
                      21 started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 1367)
02/25/05   11:53:14  ANR0409I (Session: 1367, Origin: CL_ITSAMP02_STA)  Session
                      21 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                      1367)
02/25/05   11:53:14  ANR0409I (Session: 1367, Origin: CL_ITSAMP02_STA)  Session
                      16 ended for server TSMSRV03 (AIX-RS/6000). (SESSION:
                      1367)
02/25/05   11:53:14  ANR0403I Session 1392 ended for node CL_ITSAMP02_CLIENT
                      (Linux86). (SESSION: 1392)
02/25/05   11:53:14  ANR0480W (Session: 1367, Origin: CL_ITSAMP02_STA)  Session
                      15 for node CL_ITSAMP02_CLIENT (Linux86) terminated -
                      connection with client severed. (SESSION: 1367)
02/25/05   11:53:14  ANR0409I Session 1398 ended for server CL_ITSAMP02_STA
                      (Linux/i386). (SESSION: 1398)
02/25/05   11:53:14  ANR2997W The server log is 89 percent full. The server
                      will delay transactions by 3 milliseconds.
02/25/05   11:53:16  ANR0991I (Session: 1367, Origin: CL_ITSAMP02_STA)  Storage
                      agent shutdown complete. (SESSION: 1367)
02/25/05   11:53:16  ANR3605E Unable to communicate with storage agent.
                      (SESSION: 1366)
02/25/05   11:53:16  ANR3605E Unable to communicate with storage agent.
                      (SESSION: 1369)
```

The tape volume is still mounted in tape drive DRLTO_2.

6. Resources are brought online on our second node, lochness. The restore schedule is restarted as shown in the schedule log in Example 15-30.

Example 15-30 Schedule restarts at lochness

```
02/25/2005 11:54:38 Scheduler has been started by Dsmcad.
02/25/2005 11:54:38 Querying server for next scheduled event.
02/25/2005 11:54:38 Node Name: CL_ITSAMP02_CLIENT
02/25/2005 11:54:38 Session established with server TSMSRV03: AIX-RS/6000
```

```
[...]
Executing scheduled command now.
02/25/2005 11:54:38 --- SCHEDULEREC OBJECT BEGIN RESTORE_ITSAMP 02/25/2005
11:50:00
02/25/2005 11:54:38 Restore function invoked.

02/25/2005 11:54:39 ANS1898I ***** Processed     3,000 files *****
02/25/2005 11:54:39 ANS1946W File /mnt/nfsfiles/root/.ICEauthority exists,
skipping
[...]
02/25/2005 11:54:47  ** Interrupted **
02/25/2005 11:54:47 ANS1114I Waiting for mount of offline media.
02/25/2005 11:55:56 Restoring          30,619
/mnt/nfsfiles/root/isc-backup-2005-02-03-11-15/AppServer/temp/DefaultNode/ISC_P
o
rtal/AdminCenter_PA_1_0_69/AdminCenter.war/jsp/5.3.0.0/common/_server_5F_prop_5
F_nbcommun.class [Done]
```

The tape volume is unmounted and then mounted again.

7. The backup finishes successfully as shown in the schedule log in Example 15-31. We remove diomede from the list of excluded nodes with the **samctrl -u d diomede** command.

Example 15-31 Restore finishes successfully

```
02/25/2005 12:00:02 --- SCHEDULEREC STATUS END
02/25/2005 12:00:02 Scheduled event 'RESTORE_ITSAMP' completed successfully.
02/25/2005 12:00:02 Sending results for scheduled event 'RESTORE_ITSAMP'.
02/25/2005 12:00:02 Results sent to server for scheduled event
'RESTORE_ITSAMP'.
```

Attention: notice that the restore process is started from the beginning. It is not restarted.

Results summary

The test results show that after a failure on the node that hosts the Tivoli Storage Manager client scheduler instance, a scheduled restore operation started on this node using the LAN-free path is started again from the beginning on the second node of the cluster when the service is online.

This is true if the startup window for the scheduled restore operation is not elapsed when the scheduler client is online again on the second node.

Also notice that the restore is not restarted from the point of failure, but started from the beginning. The scheduler queries the Tivoli Storage Manager server for

a scheduled operation and a new session is opened for the client after the failover.

Archived

Part 5

Establishing a VERITAS Cluster Server Version 4.0 infrastructure on AIX with IBM Tivoli Storage Manager Version 5.3

In this part of the book, we provide details on the planning, installation, configuration, testing, and troubleshooting of a VERITAS Cluster Server Version 4.0 running on AIX V5.2 and hosting the Tivoli Storage Manager Version 5.3 as a highly available application.

© Copyright IBM Corp. 2005. All rights reserved.

Archived

16

The VERITAS Cluster Server for AIX

This chapter[1] introduces VERITAS Cluster Server for AIX, which is a high availability software package that is designed to reduce both planned and unplanned downtime in a business critical environment.

Topics discussed include:

- Executive overview
- Components of a VERITAS cluster
- Cluster resources
- Cluster configurations
- Cluster communications
- Cluster installation and setup
- Cluster administration facilities
- HACMP and VERITAS Cluster Server compared

[1] This chapter was originally written in the IBM Redbook SG24-6619, then updated with version changes.

© Copyright IBM Corp. 2005. All rights reserved.

16.1 Executive overview

VERITAS Cluster Server is a leading open systems clustering solution on Sun Solaris and is also available on HP/UX, AIX, Linux, and Windows 2003. It is scalable up to 32 nodes in an AIX cluster, and supports the management of multiple VCS clusters (Windows or UNIX) from a single Web or Java based Graphical User Interface (GUI). However, individual clusters must be comprised of systems running the same operating system.

VERITAS Cluster Server has similar function to IBM High Availability Cluster Multi Processing (HACMP) product, eliminating single points of failure through the provision of redundant components, automatic detection of application, adapter, network, and node failures, and managing failover to a remote server with no apparent outage to the end user.

The VCS GUI based cluster management console provides a common administrative interface in a cross platform environment. There is also integration with other VERITAS products, such the VERITAS Volume Replicator and VERITAS Global Cluster Server.

16.2 Components of a VERITAS cluster

A VERITAS cluster is comprised nodes, external shared disk, networks, applications, and clients. Specifically, a cluster is defined as all servers with the same cluster ID connected via a set of redundant heartbeat paths:

- **Nodes:** Nodes in a VERITAS cluster are called cluster servers. There can be up to 32 cluster servers in an AIX VERITAS cluster, and up to 32 nodes on other platforms. A node will run an application or multiple applications, and can be added to or removed from a cluster dynamically.
- **Shared external disk devices:** VERITAS Cluster Server supports a number of third-party storage vendors, and works in small computer system interface (SCSI), network storage attached (NAS), and storage area networks (SAN) environments. In addition, VERITAS offer a Cluster Server Storage Certification Suite (SCS) for OEM disk vendors to certify their disks for use with VCS. Contact VERITAS directly for more information about SCS.
- **Networks and disk channels:** These channels, in VCS cluster networks, are required for both heartbeat communication, to determine the status of resources in the cluster, and also for client traffic. VCS uses its own protocol, Low Latency Transport (LLT), for cluster heartbeat communication. A second protocol, Group Membership Services/Atomic Broadcast (GAB), is used for communicating cluster configuration and state information between servers in the cluster. The LLT and GAB protocols are used instead of a TCP/IP based

communication mechanism. VCS requires a minimum of two dedicated private heartbeat connections, or high-priority network links, for cluster communication. To enable active takeover of resources, should one of these heartbeat paths fail, a third dedicated heartbeat connection is required.

Client traffic is sent and received over public networks. This public network can also be defined as a low-priority network, so should there be a failure of the dedicated high-priority networks, heartbeats can be sent at a slower rate over this secondary network.

A further means of supporting heartbeat traffic is via disk, using what is called a GABdisk. Heartbeats are written to and read from a specific area of a disk by cluster servers. Disk channels can only be used for cluster membership communication, not for passing information about a cluster's state. Note that the use of a GABdisk limits the number of servers in a cluster to eight, and not all vendors disk arrays support GABdisks.

Ethernet is the only supported network type for VCS.

16.3 Cluster resources

Resources to be made highly available include network adapters, shared storage, IP addresses, applications, and processes. Resources have a type associated with them and you can have multiple instances of a resource type. Control of each resource type involves bringing the resource online, taking it offline, and monitoring its health:

- **Agents:** For each resource type, VCS has a cluster agent that controls the resource. Types of VCS agents include:
  - **Bundled agents:** These are standard agents that come bundled with the VCS software for basic resource types, such as disk, IP, and mount. Examples of actual agents are *Application*, *IP*, *DiskGroup*, and *Mount*. For additional information, see the *VERITAS Bundled Agents Reference Guide*.
  - **Enterprise agents:** These are for applications, and are purchased separately from VCS. Enterprise agents exist for products such as DB2®, Oracle, and VERITAS Netbackup.
  - **Storage agents:** These also exist to provide access and control over storage components, such as the VERITAS ServPoint (NAS) appliance.
  - **Custom agents:** These can be created using the VERITAS developer agent for additional resource types, including applications for which there is no enterprise agent. See the *VERITAS Cluster Server Agents Developers Guide* for information about creating new cluster agents.

VERITAS cluster agents are multithreaded, so they support the monitoring of multiple instances of a resource type.

- **Resource categories:** A resource also has a category associated with it that determines how VCS handles the resource. Resources categories include:
  - **On-Off:** VCS starts and stops the resource as required (most resources are On-Off).
  - **On-Only:** Brought online by VCS, but is not stopped when the related service group is taken offline. An example of this kind of resource would be starting a daemon.
  - **Persistent:** VCS cannot take the resource online or offline, but needs to use it, so it monitors its availability. An example would be the network card that an IP address is configured upon.
- **Service group:** A set of resources that are logically grouped to provide a service. Individual resource dependencies must be explicitly defined when the service group is created to determine the order resources are brought online and taken offline. When VERITAS cluster server is started, the cluster server engine examines resource dependencies and starts all the required agents. A cluster server can support multiple service groups.

  Operations are performed on resources and also on service groups. All resources that comprise a service group will move if any resource in the service group needs to move in response to a failure. However, where there are multiple service groups running on a cluster server, only the affected service group is moved.

  The service group type defines takeover relationships, which are termed either failover or parallel, as follows:

  - **Failover:** This type of service group runs only one cluster server at a time and supports failover of resources between cluster server nodes. Failover can be both unplanned (unexpected resource outage) and planned, for example, for maintenance purposes. Although the nodes, which can take over a service group, will be defined, there are three methods by which the destination failover node is decided:
    - **Priority:** The *SystemList* attribute is used to set the priority for a cluster server. The server with the lowest defined priority that is in the running state becomes the target system. Priority is determined by the order the servers are defined in the SystemList with the first server in the list being the lowest priority server. This is the default method of determining the target node at failover, although priority can also be set explicitly.
    - **Round:** The system running the smallest number of service groups becomes the target.

- **Load:** The cluster server with the most available capacity becomes the target node. To determine available capacity, each service group is assigned a capacity. This value is used in the calculation to determine the fail-over node, based on the service groups active on the node.

– **Parallel:** These service groups are active on all cluster nodes that run resources simultaneously. Applications must be able to run on multiple servers simultaneously with no data corruption. This type of service group is sometimes also described as *concurrent*. A parallel resource group is used for things like Web hosting.

The Web VCS interface is typically defined as a service group and kept highly available. It should be noted, however, that although actions can be initiated from the browser, it is not possible to add or remove elements from the configuration via the browser. The Java VCS console should be used for making configuration changes.

In addition, service group dependencies can be defined. Service group dependencies apply when a resource is brought online, when a resource faults, and when the service group is taken offline. Service group dependencies are defined in terms of a parent and child, and a service group can be both a child and parent. Service group dependencies are defined by three parameters:

– Category
– Location
– Type

Values for these parameters are:

– Online/offline
– Local/global/remote
– Soft/hard

As an example, take two service groups with a dependency of online, remote, and soft. The category online means that the parent service group must wait for the child service group to be brought on online before it is started. Use of the remote location parameter requires that the parent and child must necessarily be on different servers. Finally, the type soft has implications for service group behavior should a resource fault. See the *VERITAS Cluster Server User Guide* for detailed descriptions of each option. Configuring service group dependencies adds complexity, so must be carefully planned.

► **Attributes:** All VCS components have attributes associated with them that are used to define their configuration. Each attribute has a *data type* and *dimension*. Definitions for data types and dimensions are detailed in the *VERITAS Cluster Server User Guide*. An example of a resource attribute is the IP address associated with a network interface card.

- **System zones:** VCS supports system zones, which are a subset of systems for a service group to use at initial failover. The service group will choose a host within its system zone before choosing any other host.

16.4 Cluster configurations

Here is the VERITAS terminology describing supported cluster configurations:

- **Asymmetric:** There is a defined primary and a dedicated backup server. Only the primary server is running a production workload.
- **Symmetric:** There is a two-node cluster where each cluster server is configured to provide a highly available service and acts as a backup to the other.
- **N-to-1:** There are N production cluster servers and a single backup server. This setup relies on the concept that failure of multiple servers at any one time is relatively unlikely. In addition, the number of slots in a server limits the total number of nodes capable of being connected in this cluster configuration.
- **N+1:** An extra cluster server is included as a spare. Should any of the N production servers fail, its service groups will move to the spare cluster server. When the failed server is recovered, it simply joins as a spare so there is no further interruption to service to failback the service group.
- **N-to-N:** There are multiple service groups running on multiple servers, which can be failed to potentially different servers.

16.5 Cluster communication

Cross cluster communication is required to achieve automated failure detection and recovery in a high availability environment. Essentially, all cluster servers in a VERITAS cluster must run:

- **High availability daemon (HAD):** This is the primary process and is sometimes referred to as the cluster server engine. A further process, *hashadow*, monitors HAD and can restart it if required. VCS agents monitor the state of resources and pass information to their local HAD. The HAD then communicates information about cluster status to the other HAD processes using the GAB and LLT protocols.
- **Group membership services/atomic broadcast (GAB):** This operates in the kernel space, monitors cluster membership, tracks cluster status (resources and service groups), and distributes this information among cluster nodes using the low latency transport layer.

- **Low latency transport (LLT):** Low latency transport operates in kernel space, supporting communication between servers in a cluster, and handles heartbeat communication. LLT runs directly on top of the DLPI layer in UNIX. LLT load balances cluster communication over the private network links.

A critical question related to cluster communication is, “What happens when communication is lost between cluster servers?” VCS uses heartbeats to determine the health of its peers and requires a minimum of two heartbeat paths, either private, public, or disk based. With only a single heartbeat path, VCS is unable to determine the difference between a network failure and a system failure. The process of handling loss of communication on a single network as opposed to a multiple network is called *jeopardy*. So, if there is a failure on all communication channels, the action taken depends on what channels have been lost and the state of the channels prior to the failure. Essentially, VCS will take action such that only one node has a service group at any one time; in some instances, disabling failover to avoid possible corruption of data. A full discussion is included in “Network partitions and split-brain” in Chapter 13, “Troubleshooting and Recovery”, in the *VERITAS Cluster Server User Guide*.

16.6 Cluster installation and setup

Installation of VCS on AIX is via **`installp`** or SMIT. It should be noted, however, that if **`installp`** is used, then LLT, GAB, and the main.cf file must be configured manually. Alternatively, the installvcs script can be used to handle the installation of the required software and initial cluster configuration.

After the VCS software has been installed, configuration is typically done via the VCS GUI interface. The first step is to carry out careful planning of the desired high availability environment. There are no specific tools in VCS to help with this process. Once this has been done, service groups are created and resources are added to them, including resource dependencies. Resources are chosen from the bundled agents and enterprise agents or, if there are no existing agents for a particular resource, a custom agent can be built. After the service groups have been defined, the cluster definition is automatically synchronized to all cluster servers.

Under VCS, the cluster configuration is stored in ASCII files. The two main files are the *main.cf* and *types.cf*:

- **main.cf:** Defines the entire cluster
- **types.cf:** Defines the resources

These files are user readable and can be edited in a text editor. A new cluster can be created based on these files as templates.

16.7 Cluster administration facilities

Administration in a VERITAS cluster is generally carried out via the cluster manager GUI interface. The cluster manager provides a graphical view of cluster status for resources, service groups, heartbeat communication, etc.

- **Administration security:** A VCS administrator can have one of five user categories. These include *Cluster Administrator*, *Cluster Operator*, *Group Administrator*, *Group Operator*, and *Cluster Guest*. Functions within these categories overlaps. The Cluster Administrator has full privileges and the ClusterGuest has read only function. User categories are set implicitly for the cluster by default, but can be also be set explicitly for individual service groups.
- **Logging:** VCS generates both error messages and log entries for activity in the cluster from both the cluster engine and each of the agents. Log files related to the cluster engine can be found in the /var/VRTSsvc/log directory, and agent log files in the $VCS_HOME/log directory. Each VCS message has a tag, which is used to indicate the type of the message. Tags are of the form TAG_A-E, where TAG_A is an error message and TAG_D indicates that an action has occurred in the VCS cluster. Log files are ASCII text and user readable. However, the cluster management interface is typically used to view logs.
- **Monitoring and diagnostic tools:** VCS can monitor both system events and applications. Event triggers allow the system administrator to define actions to be performed when a service group or resource hits a particular trigger. Triggers can also be used to carry out an action before the service group comes online or goes offline. The action is typically a script, which can be edited by the user. The event triggers themselves are pre-defined. Some can be enabled by administrators, where others are enabled by default. In addition, VCS provides simple network management protocol (SNMP), management interface base (MIB), and simple mail transfer protocol (SMTP) notification. The severity level of a notification is configurable. Event notification is implemented in VCS using *triggers*.
- **Emulation tools:** There are no emulation tools in the current release of VERITAS Cluster Server for AIX Version 2.0.

16.8 HACMP and VERITAS Cluster Server compared

The following section describes HACMP and highlights where terminology and operation differ between HACMP and VERITAS Cluster Server (VCS). HACMP and VCS have fairly comparable function, but differ in some areas. VCS has support for cross-platform management, is integrated with other VERITAS products, and uses a GUI interface as its primary management interface.

HACMP is optimized for AIX and pSeries servers, and is tightly integrated with the AIX operating system. HACMP can readily utilize availability functions in the operating system to extend its capabilities to monitoring and managing of non-cluster events.

16.8.1 Components of an HACMP cluster

An HACMP cluster is similarly comprised nodes, external shared disk, networks, applications, and clients:

- **Nodes:** The nodes in an HACMP cluster are called cluster nodes, compared with VCS cluster server. There can be up to 32 nodes in an HACMP/ES cluster, including in a concurrent access configuration. A node will run an application or multiple applications, and can be added to or removed from a cluster dynamically.
- **Shared external disk devices:** HACMP has built-in support for a wide variety of disk attachments, including Fibre Channel and several varieties of SCSI. HACMP provides an interface for OEM disk vendors to provide additional attachments for NAS, SAN, and other disks.
- **Networks:** IP networks in an HACMP cluster are used for both heartbeat/message communication, to determine the status of the resources in the cluster, and also for client traffic. HACMP uses an optimized heartbeat protocol over IP. Supported IP networks include Ethernet, FDDI, token-ring, SP-Switch, and ATM. Non-IP networks are also supported to prevent the TCP/IP network from becoming a single point of failure in a cluster. Supported non-IP networks include serial (rs232), target mode SSA (TMSSA), and target mode SCSI (TMSCSI) via the shared disk cabling. Public networks in HACMP carry both heartbeat/message and client traffic. Networks based on X.25 and SNA are also supported as cluster resources. Cluster configuration information is propagated over the public TCP/IP networks in an HACMP cluster. However, heartbeats and messages, including cluster status information, is communicated over all HACMP networks.

16.8.2 Cluster resources

Resources to be made highly available include network adapters, shared storage, IP addresses, applications, and processes. Resources have a type, and you can have multiple instances of a resource type.

- **HACMP event scripts:** Both HACMP and VCS support built-in processing of common cluster events. HACMP provides a set of predefined event scripts that handle bringing resources online, taking them offline, and moving them if required. VCS uses bundled agents. HACMP provides an event customization process and VCS provides a means to develop agents:

- **Application server:** This is the HACMP term used to describe how applications are controlled in an HACMP environment. Each application server is comprised of a start and stop script, which can be customized on a per node basis. Sample start and stop scripts are available for download for common applications at no cost.
- **Application monitor:** Both HACMP and VCS have support for application monitoring, providing for retry/restart recovery, relocation of the application, and for different processing requirements, based on the node where the application is being run.

The function of an application server coupled with an application monitor is similar to a VCS enterprise agent.

- **Resource group:** This is equivalent to a VCS service group, and is the term used to define a set of resources that comprise a service. The type of a resource group defines takeover relationships, which includes:
  - **Cascading:** A list of participating nodes is defined for a resource group, with the order of nodes indicating the node priority for the resource group. Resources are owned by the highest priority node available. If there is a failure, then the next active node with the highest priority will take over. Upon reintegration of a previously failed node, the resource group will move back to the preferred highest priority node.
    - **Cascading without fall back (CWOF):** This is a feature of cascading resource groups which allows a previously stopped cluster node to be reintegrated into a running HACMP cluster without initiating a take back of resources. The environment once more becomes fully highly available and the system administrator can choose when to move the resource group(s) back to the server where they usually run.
    - **Dynamic node priority (DNP) policy:** It is also possible to set a dynamic node priority (DNP) policy, which can be used at failover time to determine the best takeover node. Each potential takeover node is queried regarding the DNP policy, which might be something like least loaded. DNP uses the Event Management component of RSCT and is therefore available with HACMP/ES only. Obviously, it only makes sense to have a DNP policy where there are more than two nodes in a cluster. Similarly, the use of Load to determine the takeover node in a VCS cluster is only relevant where there are more than two cluster servers. There is an extensive range of possible values that can be used to define a DNP policy; run the **`haemqvar -h`** ***`cluster_name`*** command to get a full list.
  - **Rotating:** A list of participating nodes is defined for a resource group, with the order indicating the node priority for a resource group. When a cluster node is started, it will try to bring online the resource group for which it has the highest priority. Once all rotating resource groups have been brought

online, any additional cluster nodes that participate in the resource group join as standby. Should there be a failure, a resource group will move to an available standby (with the highest priority) and remain there. At reintegration of a previously failed node, there is no take back, and the server simply joins as standby.

– **Concurrent:** Active on multiple nodes at the same time. Applications in a concurrent resource group are active on all cluster nodes, and access the same shared data. Concurrent resource groups are typically used for applications that handle access to the data, although the cluster lock daemon *cllockd* is also provided with HACMP to support locking in this environment. Raw logical volumes must be used with concurrent resources groups. An example of an application that uses concurrent resource groups is Oracle 9i Real Application Cluster.

In HACMP Version 4.5 or later, resource groups are brought online in parallel by default to minimize the total time required to bring resources online. It is possible, however, to define a temporal order if resource groups need to be brought online sequentially. Other resource group dependencies can be scripted and executed via pre- and post-events to the main cluster events.

HACMP does not have an equivalent to VCS system zones.

16.8.3 Cluster configurations

HACMP and VCS are reasonably comparable in terms of supported cluster configurations, although the terminology differs. HACMP cluster configurations include:

- **Standby configurations:** Support a traditional hardware configuration where there is redundant equipment available as a hot standby. Can have both cascading and rotating resources in a hot standby configuration.
- **Takeover configurations:** All cluster nodes do useful work and act as a backup to each other. Takeover configurations include cascading mutual takeover, concurrent, one-to-one, one-to-any, any-to-one, and any-to-any.
- **Concurrent:** All cluster nodes are active and have simultaneous access to the same shared resources.

16.8.4 Cluster communications

Cross cluster communication is a part of all high availability software, and in HACMP this task is carried out by the following components:

- **Cluster manager daemon (clstrmgr):** This can be considered similar to the VCS cluster engine and must be running on all active nodes in an HACMP

cluster. In the classic feature of HACMP, the clstrmgr is responsible for monitoring nodes and networks for possible failure, and keeping track of the cluster peers. In the enhanced scalability feature of HACMP (HACMP/ES), some of the clstrmgr function is carried out by other components, specifically, the group services and topology services components of RSCT.16.8.3, “Cluster configurations” on page 713 The clstrmgr executes scripts in response to changes in the cluster (events) to maintain availability in the clustered environment.

- **Cluster SMUX peer daemon (clsmuxpd):** This provides cluster based simple network management protocol (SNMP) support to client applications and is integrated with Tivoli Netview via HATivoli in a bundled HACMP plug-in. VCS has support for SNMP. There are two additional HACMP daemons: the cluster lock daemon (cllockd) and cluster information daemon (clinfo). Only clstrmgr and clsmuxpd need to be running in the cluster.
- **Reliable scalable cluster technology (RSCT):** This is used extensively in HACMP/ES for heartbeat and messaging, monitoring cluster status, and event monitoring. RSCT is part of the AIX 5L base operating system and is comprised of:
  - **Group services:** Co-ordinates distributed messaging and synchronization tasks.
  - **Topology services:** Provides heartbeat function, enables reliable messaging, and co-ordinates membership of nodes and adapters in the cluster.
  - **Event management:** Monitors system resources and generates events when resource status changes.

HACMP and VCS both have a defined method to determine whether a remote system is alive, and a defined response to the situation where communication has been lost between all cluster nodes. These methods essentially achieve the same result, which is to avoid multiple nodes trying to grab the same resources.

16.8.5 Cluster installation and setup

Installation of HACMP for AIX software is via the standard AIX install process using **`installp`**, from the command line or via SMIT. Installation of HACMP will automatically update a number of AIX files, such as /etc/services and /etc/inittab. No further system related configuration is required following the installation of the HACMP software.

The main **`smit`** HACMP configuration menu (fast path **`smitty hacmp`**) outlines the steps required to configure a cluster. The cluster topology is defined first and synchronized via the network to all nodes in the cluster and then the resource groups are set up. Resource groups can be created on a single HACMP node

and the definitions propagated to all other nodes in the cluster. The resources, which comprise the resource group, have implicit dependencies that are captured in the HACMP software logic.

HACMP configuration information is held in the object data manager (ODM) database, providing a secure but easily shareable means of managing the configuration. A *cluster snapshot* function is also available, which captures the current cluster configuration in two ASCII user readable files. The output from the snapshot can then be used to clone an existing HACMP cluster or to re-apply an earlier configuration. In addition, the snapshot can be easily modified to capture additional user-defined configuration information as part of the HACMP snapshot. VCS does not have a snapshot function *per se*, but allows for the current configuration to be dumped to file. The resulting VCS configuration files can be used to clone cluster configurations. There is no VCS equivalent to applying a cluster snapshot.

16.8.6 Cluster administration facilities

Cluster management is typically via the System Management Interface Tool (SMIT). The HACMP menus are tightly integrated with SMIT and are easy to use. There is also close integration with the AIX operating system.

- **Administration security:** HACMP employs AIX user management to control access to cluster management function. By default, the user must have root privilege to make any changes. AIX *roles* can be defined if desired to provide a more granular level of user control. Achieving high availability requires good change management, and this includes restricting access to users who can modify the configuration.
- **Logging:** HACMP log files are simple ASCII text files. There are separate logs for messages from the cluster daemons and for cluster events. The primary log file for cluster events is the hacmp.out file, which is by default in /tmp. The system administrator can define a non default directory for individual HACMP log files. The contents of the log files can be viewed via SMIT or a Web browser. In addition, RSCT logs are also maintained for HACMP/ES.
- **Monitoring and diagnostic tools:** HACMP has extensive event monitoring capability based on the RSCT technology, and it is possible to define a custom HACMP event to run in response to the outcome of event monitoring. In addition, multiple pre- and post-events can be scripted for all cluster events to tailor them for local conditions. HACMP and VCS both support flexible notification methods, SNMP, SMTP, and e-mail notification. HACMP uses the AIX error notification facility and can be configured to react to any error reported to AIX. VCS is based on event triggers and reacts to information from agents. HACMP also supports pager notification.

- **Emulation tools:** Actions in an HACMP cluster can be emulated. There is no emulation function in VCS.

Both HACMP and VCS provide tools to enable maintenance and change in a cluster without downtime. HACMP has the cluster single point of control (CSPOC) and dynamic reconfiguration capability (DARE). CSPOC allows a cluster change to be made on a single node in the cluster and for the change to be applied to all nodes. Dynamic reconfiguration uses the **`cldare`** command to change configuration, status, and location of resource groups dynamically. It is possible to add nodes, remove nodes, and support rolling operating system or other software upgrades. VCS has the same capabilities and cluster changes are automatically propagated to other cluster servers. However, HACMP has the unique ability to emulate migrations for testing purposes.

16.8.7 HACMP and VERITAS Cluster Server high level feature comparison summary

Table 16-1 provides a high level feature comparison of HACMP and VERITAS Cluster Server, followed by Table 16-2, which compares support hardware and software environments. It should be understood that both HACMP and VERITAS Cluster Server have extensive functions that can be used to build highly available environments and the online documentation for each product must be consulted.

Table 16-1 HACMP/VERITAS Cluster Server feature comparison

Feature	HACMP	VCS for AIX
Resource/service group failover.	Yes, only affected resource group moved in response to a failure. Resource group moved as an entity.	Yes, only affected service group moved in response to a failure. Service group moved as an entity.
IP address takeover.	Yes.	Yes.
Local swap of IP address.	Yes.	Yes.
Management interfaces.	CLI and SMIT menus.	CLI, Java-based GUI, and Web console.
Cross-platform cluster management.	No.	Yes, but with the requirement that nodes in a cluster be homogenous.
Predefined resource agents.	N/A. Management of resources integrated in the logic of HACMP.	Yes.

Feature	HACMP	VCS for AIX
Predefined application agents.	No. Sample application server start/stop scripts available for download.	Yes, Oracle, DVB2, and VVR.
Automatic cluster synchronization of volume group changes.	Yes.	N/A.
Ability to define resource relationships.	Yes, majority of resource relationships integral in HACMP logic. Others can be scripted.	Yes, via CLI and GUI.
Ability to define resource/service group relationships.	Yes, to some extent via scripting.	Yes, via CLI and GUI.
Ability to decide fail-over node at time of failure based on load.	Yes, dynamic node priority with cascading resource group. Number of ways to define load via RSCT.	Yes, load option of fallover service group. Single definition of load.
Add/remove nodes without bringing the cluster down.	Yes.	Yes.
Ability to start/shutdown cluster without bringing applications down.	Yes.	Yes.
Ability to stop individual components of the resource/service group.	No.	Yes.
User level security for administration.	Based on the operating system with support for roles.	Five security levels of user management.
Integration with backup/recovery software.	Yes, with Tivoli Storage Manager.	Yes, with VERITAS NetBackup.
Integration with disaster recovery software.	Yes, with HAGEO.	Yes, with VERITAS Volume Replicator and VERITAS Global Cluster Server.
Emulation of cluster events.	Yes.	Yes.

Archived

Table 16-2 HACMP/VERITAS Cluster Server environment support

Environment	HACMP	VCS for AIX
Operating system	AIX 4.X/5L 5.3.	AIX 4.3.3/5L 5.2. VCS on AIX 4.3.3 uses AIX LVM, JFS/JFS2 only.
Network connectivity	Ethernet (10/100 Mbs), Gigabit Ethernet, ATM, FDDI, Token-ring, and SP switch.	Ethernet (10/100 Mbs) and Gigabit Ethernet.
Disk connectivity	SCSI, Fibre Channel, and SSA.	SCSI, Fibre Channel, and SSA.
Maximum servers in a cluster	32 with HACMP Enhanced Scalability (ES) feature, eight with HACMP feature.	32.
Maximum servers - Concurrent disk access	32 - Raw logical volumes only.	N/A.
LPAR support	Yes.	Yes.
SNA	Yes.	No.
Storage subsystems	See *HACMP Version 4.5 for AIX Release Notes* available for download at http://www.ibm.com/wwoi. Search for 5765-E54.	See *VERITAS Cluster Server 4.0 for AIX Release Notes*, available for download from http://support.veritas.com.

Archived

17

Preparing VERITAS Cluster Server environment

This chapter describes how our team planned, installed, configured, and tested the Veritas Cluster Server v4.0 on AIX V5.2.

This chapter provides the steps to do the following tasks:

- Review the infrastructure plan for the VCS cluster and AIX
- Do the infrastructure preparations for the Tivoli Storage Manager applications
- Install VCS v4.0 on AIX V5.2

© Copyright IBM Corp. 2005. All rights reserved.

17.1 Overview

In this chapter we discuss (and demonstrate) the installation of our Veritas cluster on AIX. It is critical that all the related Veritas documentation be reviewed and understood.

17.2 AIX overview

We will be using AIX V5.2 ML4, with the AIX JFS2 file systems, and the AIX Logical Volume Manager.

17.3 VERITAS Cluster Server

We begin with the assumption that the reader already understands the high availability concepts, and specifically, concepts related directly to the Veritas product suite. We do not discuss Veritas concepts for architecture or design in this chapter. Instead we focus entirely on implementation (installation and configuration) and testing.

Our VCS cluster running on AIX V5.2, will consist of a two-node cluster, with two Service Groups, one group per node:

- sg_tsmsrv
  - Tivoli Storage Manager server and its associated resources
  - IP and NIC assigned
  - Volume Group and mounted file systems
- sg_isc_sta_tsmcli
  - Tivoli Storage Manager client
  - Tivoli Storage Manager Storage Agent:
    - Integrated Solutions Console
    - Tivoli Storage Manager Administration Center
  - IP and NIC assigned
  - Volume Group and mounted file systems

Archived

For specific updates and changes to the Veritas Cluster Server we highly recommend referencing the following Veritas documents, which can be found at:

http://support.veritas.com

These are the documents you may find helpful:

1. Release Notes
2. Getting Started Guide
3. Installation Guide
4. User Guide
5. Latest breaking news for Storage Solutions and Clustered File Solutions 4.0 for AIX:

   http://support.veritas.com/docs/269928

17.4 Lab environment

Our lab configuration is shown in Figure 17-1, which illustrates the logical layout of the **cl_veritas01** cluster. One factor which determined our disk requirements and planning for this cluster was the decision to use Tivoli Storage Manager mirroring, which requires four disks: two for the database and two for the recovery log).

These logical disks are configured in five (5) separate arrays on the DS4500 storage subsystem. There is one array for each LUN.

Archived

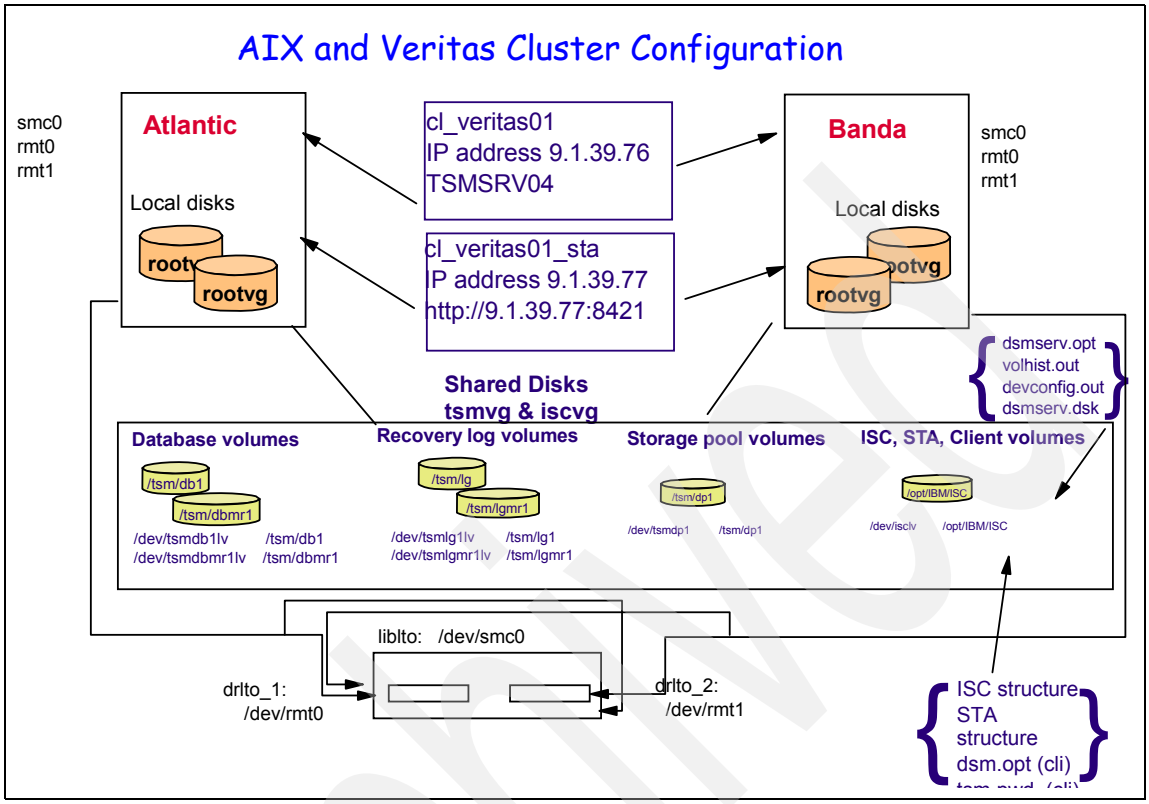

Figure 17-1 cl_veritas01 cluster physical resource layout

We are using a dual fabric SAN, with the paths shown for the disk access in Figure 17-2. This diagram also shows the heartbeat and IP connections.

Archived

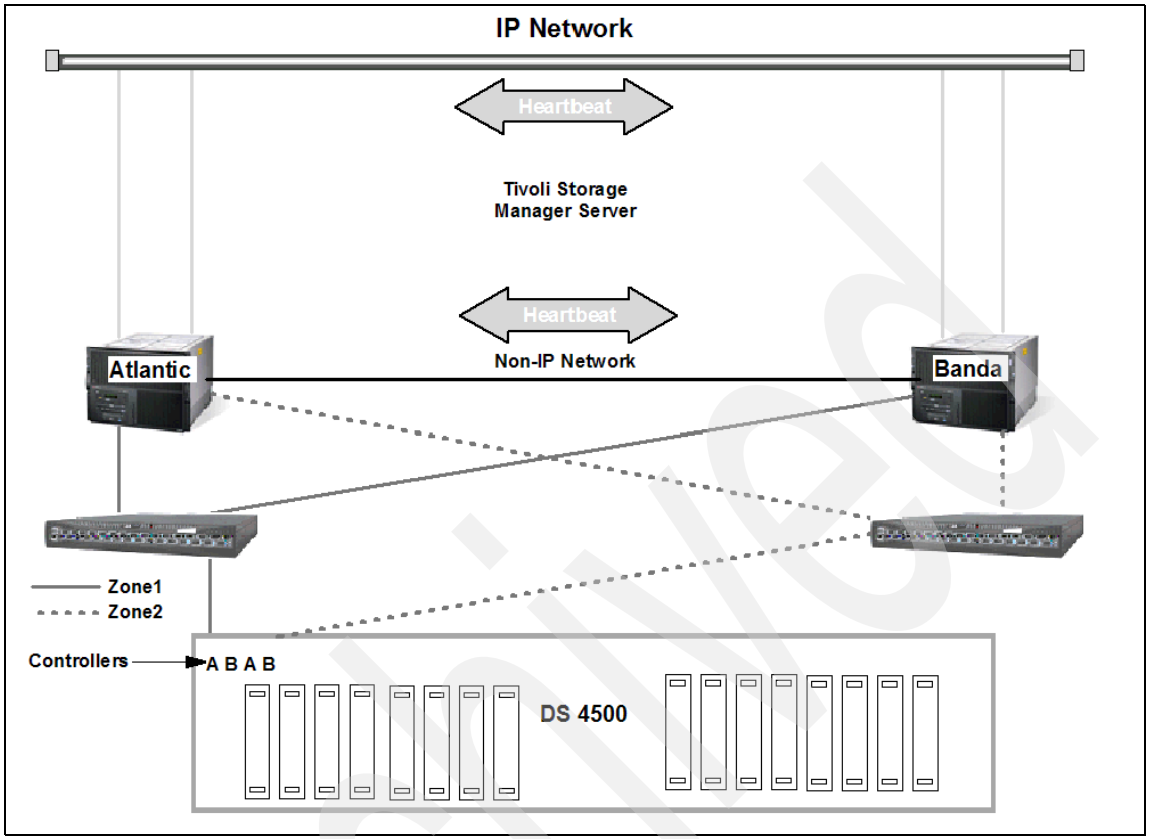

Figure 17-2 Network, SAN (dual fabric), and Heartbeat logical layout

17.5 VCS pre-installation

In this section we describe VCS pre-installation.

17.5.1 Preparing network connectivity

For this cluster, we will be implementing one private ethernet network, one disk heartbeat network, and two public NIC interfaces.

Archived

Private Ethernet network preparation

Here are the steps to follow:

1. We wire one adapter per machine using an ethernet cross-over cable. We use exactly the same adapter location and type of adapter for this connection between the two nodes. We use a cross-over cable for connecting two 10/100 integrated adapters.
2. Then, we connect the second and third adapters to the public (production) ethernet switch for each node.
3. We then configure the private network for IP communication, and validate (test) the connection. Once we determine the connection works, we remove the IP configuration using the **rmdev -ld en0** AIX command.
4. We also create a .rhosts file in the root directory for each node as shown in Example 17-1 and Example 17-2.

Example 17-1 Atlantic .rhosts

```
banda root
```

Example 17-2 Banda .rhosts

```
atlantic root
```

5. Then we configure a basic /etc/hosts file with the two nodes' IP addresses and a loopback address as shown in Example 17-3 and Example 17-4.

Example 17-3 atlantic /etc/hosts file

```
127.0.0.1 loopback localhost      # loopback (lo0) name/address
9.1.39.92 atlantic
9.1.39.94 banda
```

Example 17-4 banda /etc/hosts file

```
127.0.0.1 loopback localhost      # loopback (lo0) name/address
9.1.39.92 atlantic
9.1.39.94 banda
```

17.5.2 Installing the Atape drivers

Here are the steps to follow:

1. We then install the Atape driver using the **smitty installp** AIX command. This is required, as our library is an IBM 3582 LTO library.
2. We verify that the tape library and drives are visible to AIX using the **lsdev -Cc tape** command.

17.5.3 Preparing the storage

Here are the steps to follow:

1. Initially, we determine what the WWPNs are for the FC HBAs on the hosts systems to be configured. These systems are running AIX V5.2, so the command to determine this is shown in Example 17-5.

Example 17-5 The AIX command lscfg to view FC disk details

```
banda:/usr/tivoli/tsm/client/ba/bin# lscfg -vl fcs0 |grep Z8
        Device Specific.(Z8)........20000000C932A75D
banda:/usr/tivoli/tsm/client/ba/bin# lscfg -vl fcs1 |grep Z8
        Device Specific.(Z8)........20000000C932A865

Atlantic:/opt/local/tsmsrv# lscfg -vl fcs0 |grep Z8
        Device Specific.(Z8)........20000000C932A80A
Atlantic:/opt/local/tsmsrv# lscfg -vl fcs1 |grep Z8
        Device Specific.(Z8)........20000000C9329B6F
```

2. Next, we ensure we have fiber connectivity to the switch (visually checking the light status of both the adapter and the corresponding switch ports).
3. Then, we log into the SAN switch and assign alias and zones for the SAN disk and tape devices, and the FC HBAs listed in Example 17-5. The summary of the switch configuration is shown in Figure 17-3 and Figure 17-4.

```
Zone:          z_ATLANTIC_0_3583        10:00:00:00:c9:32:a8:0a
                                        20:01:00:60:45:16:1f:f5

Zone:          z_ATLANTIC_0_DS4500A_a        10:00:00:00:c9:32:a8:0a
                                             20:04:00:a0:b8:17:44:32
                                             20:05:00:a0:b8:17:44:32

Zone:          z_ATLANTIC_1_3582_dr1        10:00:00:00:c9:32:a8:0a
                                            50:03:08:c1:40:46:70:03

Zone:          z_ATLANTIC_1_3582_dr2        50:03:08:c1:40:46:70:06
                                            10:00:00:00:c9:32:a8:0a
```

Figure 17-3 Atlantic zoning

Archived

```
Zone:                z_BANDA_1_3582_dr1        50:03:08:c1:40:46:70:03
                                               10:00:00:00:c9:32:a8:65

Zone:                z_BANDA_1_3582_dr2        10:00:00:00:c9:32:a8:65
                                               50:03:08:c1:40:46:70:06

Zone:                z_BANDA_1_3583          10:00:00:00:c9:32:a8:65
                                             20:01:00:60:45:16:1f:f5

Zone:                z_BANDA_1_DS4500A_b        10:00:00:00:c9:32:a8:65
                                                20:05:00:a0:b8:17:44:32
                                                20:04:00:a0:b8:17:44:32
```

Figure 17-4 Banda zoning

4. Then, we go to the DS4500 storage subsystem assign LUNs to the adapter WWPNs for Banda and Atlantic. The summary of this is shown in Figure 17-5.

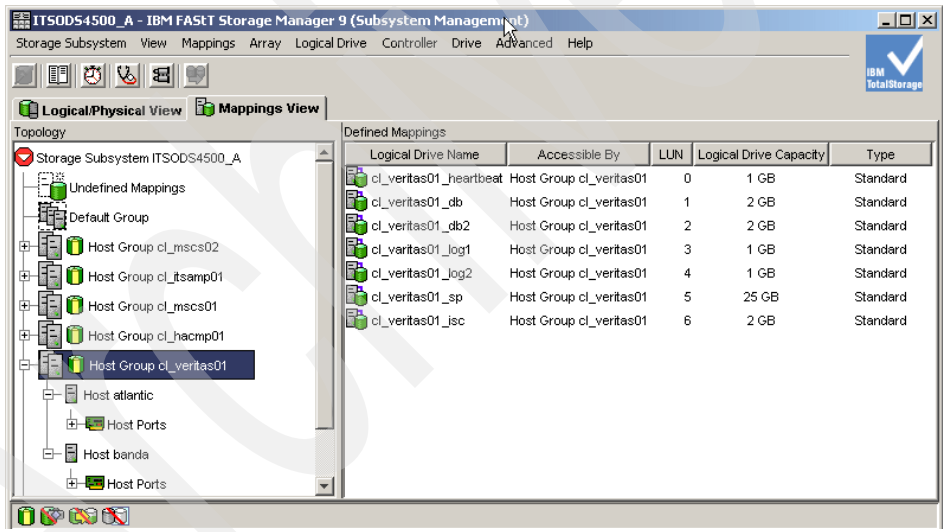

Figure 17-5 DS4500 LUN configuration for cl_veritas01

5. We then run **`cfgmgr -S`** on Atlantic, then Banda.
6. We verify the availability of volumes with **`lspv`** as shown in Example 17-6.

Example 17-6 The lspv command output

```
hdisk0          0009cdcaeb48d3a3                    rootvg          active
hdisk1          0009cdcac26dbb7c                    rootvg          active
hdisk2          0009cdcab5657239                    None
hdisk3          none                                None
```

```
hdisk4          0009cdaad089888c                    None
hdisk5          0009cdcad0b400e5                    None
hdisk6          0009cdaad089898d                    None
hdisk7          0009cdcad0b4020c                    None
hdisk8          0009cdaad0898a9c                    None
hdisk9          0009cdcad0b40349                    None
```

7. We validate that the storage subsystems configured LUNs map the same to both operating systems physical volumes, using **lscfg -vpl hdiskx** command for all disks; however, only the first one is shown in Example 17-7.

Example 17-7 The lscfg command

```
atlantic:/# lscfg -vpl hdisk4
  hdisk4           U0.1-P2-I4/Q1-W200400A0B8174432-L1000000000000 1742-900
(900) Disk Array Device

banda:/# lscfg -vpl hdisk4
  hdisk4           U0.1-P2-I5/Q1-W200400A0B8174432-L1000000000000 1742-900
(900) Disk Array Device
```

Create a non-concurrent shared volume group - Server

We now create a shared volume and the file systems required for the *Tivoli Storage Manager server*. This same procedure will also be used for setting up the storage resources for the *Integrated Solutions Console* and *Administration Center*.

1. We create the non-concurrent shared volume group on a node, using the **mkvg command,** as shown in Example 17-8.

Example 17-8 The mkvg command to create the volume group

```
mkvg -n -y tsmvg -V 47 hdisk4 hdisk5 hdisk6 hdisk7 hdisk8
```

Important:

- Do not activate the volume group AUTOMATICALLY at system restart. Set to **no** (-n flag) so that the volume group can be activated as appropriate by the cluster event scripts.
- Use the **lvlstmajor command** on each node to determine a free major number common to all nodes.
- If using SMIT, use the default fields that are already populated whereever possible, unless the site has specific requirements.

Archived

2. Then we create the logical volumes using the **mklv command** (Example 17-9). This will create the logical volumes for the jfs2log, Tivoli Storage Manager disk storage pools and configuration files on the RAID1 volume.

Example 17-9 The mklv commands to create the logical volumes

```
/usr/sbin/mklv -y tsmvglg -t jfs2log tsmvg 1 hdisk8
/usr/sbin/mklv -y tsmlv -t jfs2 tsmvg 1 hdisk8
/usr/sbin/mklv -y tsmdp1lv -t jfs2 tsmvg 790 hdisk8
```

3. Next, we create the logical volumes for Tivoli Storage Manager database and log files on the RAID-0 volumes, using the **mklv** command as shown in Example 17-10.

Example 17-10 The mklv commands used to create the logical volumes

```
/usr/sbin/mklv -y tsmdb1lv -t jfs2 tsmvg 63 hdisk4
/usr/sbin/mklv -y tsmdbmr1lv -t jfs2 tsmvg 63 hdisk5
/usr/sbin/mklv -y tsmlg1lv -t jfs2 tsmvg 32 hdisk6
/usr/sbin/mklv -y tsmlgmr1lv -t jfs2 tsmvg 32 hdisk7
```

4. We then format the jfs2log device, which will then be used when we create the file systems, as seen in Example 17-11.

Example 17-11 The logform command

```
      logform /dev/tsmvglg
logform: destroy /dev/rtsmvglg (y)?y
```

5. Then, we create the file systems on the previously defined logical volumes using the **crfs** command. All these commands are shown in Example 17-12.

Example 17-12 The crfs commands used to create the file systems

```
/usr/sbin/crfs -v jfs2 -d tsmlv -m /tsm/files -A no -p rw -a agblksize=4096
/usr/sbin/crfs -v jfs2 -d tsmdb1lv -m /tsm/db1 -A no -p rw -a agblksize=4096
/usr/sbin/crfs -v jfs2 -d tsmdbmr1lv -m /tsm/dbmr1 -A no -p rw -a
agblksize=4096
/usr/sbin/crfs -v jfs2 -d tsmlg1lv -m /tsm/lg1 -A no -p rw -a agblksize=4096
/usr/sbin/crfs -v jfs2 -d tsmlgmr1lv -m /tsm/lgmr1 -A no -p rw -a
agblksize=4096
/usr/sbin/crfs -v jfs2 -d tsmdp1lv -m /tsm/dp1 -A no -p rw -a agblksize=4096
```

6. We then vary offline the shared volume group, seen in Example 17-13.

Example 17-13 The varyoffvg command

```
varyoffvg tsmvg
```

7. We then run **cfgmgr -S** on the second node, and check for tsmvg's PVIDs presence on the second node.

> **Important:** If PVIDs are not present, issue the **chdev -l** *hdiskname* **-a pv=yes** for the required physical volumes:
>
> ```
> chdev -l hdisk4 -a pv=yes
> ```

8. We then import the volume group tsmvg on the second node, as demonstrated in Example 17-14.

Example 17-14 The importvg command

```
importvg -y tsmvg -V 47 hdisk4
```

9. Then, we change the *tsmvg* volume group, so it will not varyon (activate) at boot time, as shown in Example 17-15.

Example 17-15 The chvg command

```
chvg -a n tsmvg
```

10. We then varyoff the *tsmvg* volume group on the second node, as shown in Example 17-16.

Example 17-16 The varyoffvg command

```
varyoffvg tsmvg
```

Create a shared volume group - ISC and Administration Centre

We now create a non-concurrent shared volume and the file systems required for the *Integrated Solutions Console and Administration Center*. This same procedure will be used for creating the Tivoli Storage Manager server disk environment.

1. We create the non-concurrent shared volume group on a node, using the **mkvg** command as seen in Example 17-17.

Example 17-17 The mkvg command to create the volume group

```
mkvg -n -y iscvg -V 48 hdisk9
```

Archived

Important:

- Do not activate the volume group AUTOMATICALLY at system restart. Set to **no** (-n flag) so that the volume group can be activated as appropriate by the cluster event scripts.
- Use the **lvlstmajor** command on each node to determine a free major number common to all nodes
- If using SMIT, use the default fields that are already populated wherever possible, unless the site has specific requirements.

2. Then we create the logical volumes using the **mklv command**, as shown in Example 17-18. This will create the logical volumes for the jfs2log, Tivoli Storage Manager disk storage pools, and configuration files on the RAID1 volume.

Example 17-18 The mklv commands to create the logical volumes

```
/usr/sbin/mklv -y iscvglg -t jfs2log iscvg 1 hdisk9
/usr/sbin/mklv -y isclv -t jfs2 iscvg 100 hdisk9
```

3. We then format the jfs2log device, which will then be used when we create the file systems which is shown in Example 17-19.

Example 17-19 The logform command

```
      logform /dev/iscvglg
logform: destroy /dev/rtsmvglg (y)?y
```

4. Then, we create the file systems on the previously defined logical volumes using the **crfs** command as seen in Example 17-20.

Example 17-20 The crfs commands used to create the file systems

```
/usr/sbin/crfs -v jfs2 -d isclv -m /opt/IBM/ISC -A no -p rw -a agblksize=4096
```

5. Then, we set the volume group not to varyon automatically by using the **chvg command** as seen in Example 17-21.

Example 17-21 The chvg command

```
chvg -a n iscvg
```

6. We then vary offline the shared volume group, seen in Example 17-22.

Example 17-22 The varyoffvg command

```
      varyoffvg iscvg
```

17.5.4 Installing the VCS cluster software

Here are the steps to follow:

1. We only execute the VCS software on one node, and VCS will install the software on the second node. To facilitate this operation, we create a .rhosts file in both systems root directory, as shown in Example 17-23.

Example 17-23 .rhosts file

```
atlantic        root
banda           root
```

2. Next, we start the VCS installation script from an AIX command line, as shown in Example 17-24, which then spawns the installation screen sequence.

Example 17-24 VCS installation script

```
Atlantic:/opt/VRTSvcs/install# ./installvcs
```

3. We then reply to the first screen with the two node names for our cluster, as shown in Figure 17-6.

```
# ./installvcs

                    VERITAS CLUSTER SERVER 4.0 INSTALLATION PROGRAM

Copyright (c) 2004 VERITAS Software Corporation. All rights reserved.

VERITAS, the VERITAS Logo and all other VERITAS product names and slogans
are trademarks or registered trademarks of VERITAS Software Corporation.
VERITAS and the VERITAS Logo Reg. U.S. Pat. & Tm. Off. Other product names
and/or slogans mentioned herein may be trademarks or registered trademarks
of their respective companies.

Enter the system names separated by spaces on which to install VCS: atlantic ban
da
```

Figure 17-6 Veritas Cluster Server 4.0 Installation Program

Archived

4. This results in a cross system check verifying connectivity and environment as seen in Figure 17-7. We press Return to continue.

```
Checking system communication:

    Checking OS version on atlantic ............................ AIX 5.2.0
    Checking VRTSvcs.rte fileset ............................ not installed
    Verifying communication with banda ..................... ping successful
    Attempting rsh with banda ............................... rsh successful
    Attempting rcp with banda ............................... rcp successful
    Checking OS version on banda ............................... AIX 5.2.0
    Checking VRTSvcs.rte fileset ............................ not installed
    Creating log directory on banda ................................ Done

Logs for installvcs are being created in /var/tmp/installvcs208093610.

Using /usr/bin/rsh and /usr/bin/rcp to communicate with remote systems.

Initial system check completed successfully.

Press [Return] to continue: ▮
```

Figure 17-7 VCS system check results

5. The VCS filesets are now installed. Then we review the summary, as shown in Figure 17-8, then press Return to continue.

```
VERITAS Infrastructure fileset installation:

Installing VERITAS Infrastructure filesets on atlantic:

    Checking VRTScpi.rte fileset ............. version 4.0.4001.0 installed
    Checking VRTSvlic fileset ................... version 3.2.6.1 installed

Installing VERITAS Infrastructure filesets on banda:

    Checking VRTScpi.rte fileset ............. version 4.0.4001.0 installed
    Checking VRTSvlic fileset ................... version 3.2.6.1 installed

VERITAS Infrastructure filesets installed successfully.

Press [Return] to continue:
```

Figure 17-8 Summary of the VCS Infrastructure fileset installation

Archived

6. We then enter the VCS license key and press Enter, as seen in Figure 17-9.

```
Each system requires a VCS product license before installation.  License
keys for additional product features should also be added at this time.

Some license keys are node locked and are unique per system.  Other license
keys, such as demo keys and site license keys, are registered on all systems
and must be entered on the first system.

VCS Licensing Verification:

    Checking VCS license key on atlantic ..................... not licensed
Enter a VCS license key for atlantic: [?]
```

Figure 17-9 License key entry screen

7. Next, we are prompted with a choice of optional VCS filesets to install, we accept the default option of all filesets, and press Enter to continue as shown in Figure 17-10.

```
installvcs can install the following optional VCS filesets:

VRTSvxfen.rte     VERITAS I/O Fencing
VRTSvcs.man       VERITAS Cluster Server Man Pages
VRTSvcs.doc       VERITAS Cluster Server Documentation
VRTScssim.rte     VERITAS Cluster Server Simulator
VRTScscm.rte      VERITAS Cluster Server Cluster Manager

     1)  Install all of the optional filesets
     2)  Install none of the optional filesets
     3)  View fileset descriptions and select optional filesets

Select the optional filesets to be installed on all systems? [1-3,q,?] (1) ▌
```

Figure 17-10 Choice of which filesets to install

8. After selecting the default option to install all of the filesets by pressing Enter, a summary screen appears listing all the filesets which will be installed as shown in Figure 17-11. We then press Return to continue.

Archived

```
installvcs will install the following VCS filesets:
VRTSperl.rte     VERITAS Perl 5.8.0 Redistribution
VRTSveki         VERITAS Kernel Interface
VRTSllt.rte      VERITAS Low Latency Transport
VRTSgab.rte      VERITAS Group Membership and Atomic Broadcast
VRTSvxfen.rte    VERITAS I/O Fencing
VRTSvcs.rte      VERITAS Cluster Server
VRTSvcsag.rte    VERITAS Cluster Server Bundled Agents
VRTSvcs.msg.en_US VERITAS Cluster Server Message Catalogs
VRTSvcs.man      VERITAS Cluster Server Man Pages
VRTSvcs.doc      VERITAS Cluster Server Documentation
VRTSjre.rte      VERITAS Java Runtime Environment Redistribution
VRTScutil.rte    VERITAS Cluster Utilities
VRTScssim.rte    VERITAS Cluster Server Simulator
VRTScscw.rte     VERITAS Cluster Server Configuration Wizards
VRTSweb.rte      VERITAS Java Web Server
VRTSvcsw.rte     VERITAS Cluster Manager (Web Console)
VRTScscm.rte     VERITAS Cluster Server Cluster Manager

Press [Return] to continue: ▌
```

Figure 17-11 Summary of filesets chosen to install

9. Next, after pressing Enter, we see the VCS installation program validating its prerequisites prior to installing the filesets. The output is shown in Example 17-25. We then press Return to continue.

Example 17-25 The VCS checking of installation requirements

```
VERITAS CLUSTER SERVER 4.0 INSTALLATION PROGRAM

Checking system installation requirements:

Checking VCS installation requirements on atlantic:

    Checking VRTSperl.rte fileset ........................... not installed
    Checking VRTSveki fileset ............................... not installed
    Checking VRTSllt.rte fileset............................ not installed
    Checking VRTSgab.rte fileset............................ not installed
    Checking VRTSvxfen.rte fileset......................... not installed
    Checking VRTSvcs.rte fileset............................ not installed
    Checking VRTSvcsag.rte fileset......................... not installed
    Checking VRTSvcs.msg.en_US fileset...................... not installed
    Checking VRTSvcs.man fileset............................ not installed
    Checking VRTSvcs.doc fileset............................ not installed
    Checking VRTSjre.rte fileset............................ not installed
    Checking VRTScutil.rte fileset......................... not installed
    Checking VRTScssim.rte fileset......................... not installed
    Checking VRTScscw.rte fileset........................... not installed
    Checking VRTSweb.rte fileset............................ not installed
    Checking VRTSvcsw.rte fileset........................... not installed
    Checking VRTScscm.rte fileset........................... not installed
```

```
    Checking required AIX patch bos.rte.tty-5.2.0.14 on atlantic...
bos.rte.tty-5.2.0.50 installed
    Checking file system space............... required space is available
    Checking had process..................................... not running
    Checking hashadow process................................ not running
    Checking CmdServer process............................... not running
    Checking notifier process................................ not running
    Checking vxfen driver............... vxfen check command not installed
    Checking gab driver................... gab check command not installed
    Checking llt driver...................................... not running
    Checking veki driver..................................... not running

Checking VCS installation requirements on banda:

    Checking VRTSperl.rte fileset........................... not installed
    Checking VRTSveki fileset............................... not installed
    Checking VRTSllt.rte fileset............................ not installed
    Checking VRTSgab.rte fileset............................ not installed
    Checking VRTSvxfen.rte fileset.......................... not installed
    Checking VRTSvcs.rte fileset............................ not installed
    Checking VRTSvcsag.rte fileset.......................... not installed
    Checking VRTSvcs.msg.en_US fileset...................... not installed
    Checking VRTSvcs.man fileset............................ not installed
    Checking VRTSvcs.doc fileset............................ not installed
    Checking VRTSjre.rte fileset............................ not installed
    Checking VRTScutil.rte fileset.......................... not installed
    Checking VRTScssim.rte fileset.......................... not installed
    Checking VRTScscw.rte fileset........................... not installed
    Checking VRTSweb.rte fileset............................ not installed
    Checking VRTSvcsw.rte fileset........................... not installed
    Checking VRTScscm.rte fileset........................... not installed
    Checking required AIX patch bos.rte.tty-5.2.0.14 on banda...
bos.rte.tty-5.2.0.50 installed
    Checking file system space............... required space is available
    Checking had process..................................... not running
    Checking hashadow process................................ not running
    Checking CmdServer process............................... not running
    Checking notifier process................................ not running
    Checking vxfen driver............... vxfen check command not installed
    Checking gab driver................... gab check command not installed
    Checking llt driver...................................... not running
    Checking veki driver..................................... not running

Installation requirement checks completed successfully.
Press [Return] to continue:
```

Archived

10. The panel which offers the option to configure VCS now appears. We then choose the default option by pressing Enter, as shown in Figure 17-12.

```
It is possible to install VCS filesets without performing configuration.

VCS cannot be started without proper configuration.

It is optional to configure VCS now. If you choose to configure VCS later,
you can either do so manually or run the installvcs -configure command.
Are you ready to configure VCS? [y,n,q] (y)
```

Figure 17-12 VCS configuration prompt screen

11. We then press Enter at the prompt for the screen as shown in Figure 17-13.

```
installvcs will now ask sets of VCS configuration-related questions.

When a [b] is presented after a question, 'b' may be entered to go back to
the first question of the configuration set.

When a [?] is presented after a question, '?' may be entered for help or
additional information about the question.

Following each set of questions, the information you have entered will be
presented for confirmation.  To repeat the set of questions and correct any
previous errors, enter 'n' at the confirmation prompt.

No configuration changes are made to the systems until all configuration
questions are completed and VCS is installed successfully.

Press [Return] to continue:
```

Figure 17-13 VCS installation screen instructions

12. Next, we enter the cluster_name, cluster_id, and the heartbeat NICs for the cluster, as shown in Figure 17-14.

Archived

```
To configure VCS the following is required:

        A unique Cluster name
        A unique Cluster ID number between 0-255
        Two or more NIC cards per system used for heartbeat links

        One or more heartbeat links are configured as private links
        One heartbeat link may be configured as a low priority link

All systems are being configured to create one cluster

Enter the unique cluster name: [?] cl_veritas01
Enter the unique Cluster ID number between 0-255: [b,?] 01

    Discovering NICs on atlantic ................... discovered en0 en1 en2

Enter the NIC for the first private heartbeat link on atlantic: [b,?] en0
Would you like to configure a second private heartbeat link? [y,n,q,b,?] (y) y
Enter the NIC for the second private heartbeat link on atlantic: [b,?] en2
```

Figure 17-14 VCS cluster configuration screen

13. Next, the VCS summary screen is presented, which we review and then accept the values by pressing Enter, as shown in Figure 17-15.

```
Cluster information verification:

        Cluster Name: cl_veritas01
        Cluster ID Number: 1
        Private Heartbeat NICs for atlantic: link1=en0 link2=en2
        Private Heartbeat NICs for banda: link1=en0 link2=en2

Is this information correct? [y,n,q] (y)
```

Figure 17-15 VCS screen reviewing the cluster information to be set

14. We are then presented with an option to set the password for the admin user, which we decline by accepting the default and pressing Enter, which is shown in Figure 17-16.

```
The following information is required to add VCS users:

        A user name
        A password for the user
        User privileges (Administrator, Operator, or Guest)

Do you want to set the password for the Admin user
(default password='password')? [y,n,q] (n)
```

Figure 17-16 VCS setup screen to set a non-default password for the admin user

Archived

15. We accept the default password for the administrative user, and decline on the option to add additional users, which is shown in Figure 17-17.

```
The following information is required to add VCS users:

        A user name
        A password for the user
        User privileges (Administrator, Operator, or Guest)

Do you want to set the password for the Admin user
(default password='password')? [y,n,q] (n)
Do you want to add another user to the cluster? [y,n,q] (y) n
```

Figure 17-17 VCS adding additional users screen

16. Next, the summary screen is presented, which we review. We then accept the default by pressing Enter, as shown in Figure 17-18.

```
VCS User verification:

        User: admin     Privilege: Administrators

        Passwords are not displayed

Is this information correct? [y,n,q] (y)
```

Figure 17-18 VCS summary for the privileged user and password configuration

17. Then, we respond to the Cluster Manager Web Console configuration prompt by pressing Enter (accepting the default), as shown in Figure 17-19.

```
The following information is required to configure Cluster Manager:

        A public NIC used by each system in the cluster
        A Virtual IP address and netmask for Cluster Manager

Do you want to configure Cluster Manager (Web Console) [y,n,q] (y)
```

Figure 17-19 VCS prompt screen to configure the Cluster Manager Web console

18. We answer the prompts for configuring the Cluster Manager Web Console and then press Enter, which then results in the summary screen displaying as seen in Figure 17-20.

Archived

```
Cluster Manager (Web Console) verification:

        NIC: en1
        IP: 9.1.39.77
        Netmask: 255.255.255.0

Is this information correct? [y,n,q] (y)
```

Figure 17-20 VCS screen summarizing Cluster Manager Web Console settings

19. The following screen prompts us to configure SMTP notification, which we decline, as shown in Figure 17-21. Then we press Return to continue.

```
The following information is required to configure SMTP notification:

        The domain-based hostname of the SMTP server
        The email address of each SMTP recipient
        A minimum severity level of messages to send to each recipient

Do you want to configure SMTP notification? [y,n,q] (y) n
```

Figure 17-21 VCS screen prompt to configure SNTP notification

20. On the following panel, we decline the opportunity to configure SNMP notification for our lab environment, as shown in Figure 17-22.

```
The following information is required to configure SNMP notification:

        System names of SNMP consoles to receive VCS trap messages
        SNMP trap daemon port numbers for each console
        A minimum severity level of messages to send to each console

Do you want to configure SNMP notification? [y,n,q] (y) n
```

Figure 17-22 VCS screen prompt to configure SNMP notification

21. The option to install VCS simultaneously or consecutively is given, and we choose consecutively (answer no to the prompt), which allows for better error handling, as shown in Figure 17-23.

Archived

```
VCS can be installed on systems consecutively or simultaneously. Installing
on systems consecutively takes more time but allows for better error
handling.

Would you like to install Cluster Server on all systems simultaneously?
[y,n,q,?] (y) n
```

Figure 17-23 VCS prompt for a simultaneous installation of both nodes

22. The install summary follows, and is shown in Example 17-26.

Example 17-26 The VCS install method prompt and install summary

```
VERITAS CLUSTER SERVER 4.0 INSTALLATION PROGRAM

Installing Cluster Server 4.0.0.0 on atlantic:

    Installing VRTSperl 4.0.2.0 on atlantic............ Done 1 of 51 steps
    Installing VRTSveki 1.0.0.0 on atlantic............ Done 2 of 51 steps
    Installing VRTSllt 4.0.0.0 on atlantic............ Done 3 of 51 steps
    Installing VRTSgab 4.0.0.0 on atlantic............ Done 4 of 51 steps
    Installing VRTSvxfen 4.0.0.0 on atlantic........... Done 5 of 51 steps
    Installing VRTSvcs 4.0.0.0 on atlantic............ Done 6 of 51 steps
    Installing VRTSvcsag 4.0.0.0 on atlantic........... Done 7 of 51 steps
    Installing VRTSvcsmg 4.0.0.0 on atlantic........... Done 8 of 51 steps
    Installing VRTSvcsmn 4.0.0.0 on atlantic........... Done 9 of 51 steps
    Installing VRTSvcsdc 4.0.0.0 on atlantic.......... Done 10 of 51 steps
    Installing VRTSjre 1.4.0.0 on atlantic............ Done 11 of 51 steps
    Installing VRTScutil 4.0.0.0 on atlantic.......... Done 12 of 51 steps
    Installing VRTScssim 4.0.0.0 on atlantic.......... Done 13 of 51 steps
    Installing VRTScscw 4.0.0.0 on atlantic........... Done 14 of 51 steps
    Installing VRTSweb 4.1.0.0 on atlantic............ Done 15 of 51 steps
    Installing VRTSvcsw 4.1.0.0 on atlantic........... Done 16 of 51 steps
    Installing VRTScscm 4.1.0.0 on atlantic........... Done 17 of 51 steps

Installing Cluster Server 4.0.0.0 on banda:

    Copying VRTSperl.rte.bff.gz to banda.............. Done 18 of 51 steps
    Installing VRTSperl 4.0.2.0 on banda.............. Done 19 of 51 steps
    Copying VRTSveki.bff.gz to banda................... Done 20 of 51 steps
    Installing VRTSveki 1.0.0.0 on banda.............. Done 21 of 51 steps
    Copying VRTSllt.rte.bff.gz to banda............... Done 22 of 51 steps
    Installing VRTSllt 4.0.0.0 on banda............... Done 23 of 51 steps
    Copying VRTSgab.rte.bff.gz to banda............... Done 24 of 51 steps
    Installing VRTSgab 4.0.0.0 on banda............... Done 25 of 51 steps
    Copying VRTSvxfen.rte.bff.gz to banda............. Done 26 of 51 steps
    Installing VRTSvxfen 4.0.0.0 on banda............. Done 27 of 51 steps
```

```
    Copying VRTSvcs.rte.bff.gz to banda............... Done 28 of 51 steps
    Installing VRTSvcs 4.0.0.0 on banda............... Done 29 of 51 steps
    Copying VRTSvcsag.rte.bff.gz to banda............. Done 30 of 51 steps
    Installing VRTSvcsag 4.0.0.0 on banda............. Done 31 of 51 steps
    Copying VRTSvcs.msg.en_US.bff.gz to banda......... Done 32 of 51 steps
    Installing VRTSvcsmg 4.0.0.0 on banda............. Done 33 of 51 steps
    Copying VRTSvcs.man.bff.gz to banda............... Done 34 of 51 steps
    Installing VRTSvcsmn 4.0.0.0 on banda............. Done 35 of 51 steps
    Copying VRTSvcs.doc.bff.gz to banda............... Done 36 of 51 steps
    Installing VRTSvcsdc 4.0.0.0 on banda............. Done 37 of 51 steps
    Copying VRTSjre.rte.bff.gz to banda............... Done 38 of 51 steps
    Installing VRTSjre 1.4.0.0 on banda............... Done 39 of 51 steps
    Copying VRTScutil.rte.bff.gz to banda............. Done 40 of 51 steps
    Installing VRTScutil 4.0.0.0 on banda............. Done 41 of 51 steps
    Copying VRTScssim.rte.bff.gz to banda............. Done 42 of 51 steps
    Installing VRTScssim 4.0.0.0 on banda............. Done 43 of 51 steps
    Copying VRTScscw.rte.bff.gz to banda.............. Done 44 of 51 steps
    Installing VRTScscw 4.0.0.0 on banda.............. Done 45 of 51 steps
    Copying VRTSweb.rte.bff.gz to banda............... Done 46 of 51 steps
    Installing VRTSweb 4.1.0.0 on banda............... Done 47 of 51 steps
    Copying VRTSvcsw.rte.bff.gz to banda.............. Done 48 of 51 steps
    Installing VRTSvcsw 4.1.0.0 on banda.............. Done 49 of 51 steps
    Copying VRTScscm.rte.bff.gz to banda.............. Done 50 of 51 steps
    Installing VRTScscm 4.1.0.0 on banda.............. Done 51 of 51 steps

Cluster Server installation completed successfully.

Press [Return] to continue:
```

23. We then review the installation results and press Enter to continue, which then produces the screen as shown in Figure 17-24.

```
Configuring Cluster Server:

    Creating Cluster Server configuration files ...................... Done
    Copying configuration files to atlantic .......................... Done
    Updating pse.conf on atlantic ....................... pse.conf updated
    Copying configuration files to banda ............................. Done
    Updating pse.conf on banda .......................... pse.conf updated

Cluster Server configured successfully.

Press [Return] to continue:
```

Figure 17-24 VCS completes the server configuration successfully

24. Then, we press Enter and accept the prompt default to start the cluster server processes as seen in Figure 17-25.

```
Do you want to start Cluster Server processes now? [y,n,q] (y)

Starting Cluster Server:

    Starting LLT on atlantic ..................................... Started
    Starting LLT on banda ........................................ Started
    Starting GAB on atlantic ..................................... Started
    Starting GAB on banda ........................................ Started
    Starting Cluster Server on atlantic .......................... Started
    Starting Cluster Server on banda ............................. Started
    Confirming Cluster Server startup ................... 2 systems RUNNING

Cluster Server was started successfully.

Press [Return] to continue: ▌
```

Figure 17-25 Results screen for starting the cluster server processes

25. We then press Enter and the process is completed successfully as shown in Figure 17-26.

```
Installation of Cluster Server 4.0.0.0 has completed successfully.

The installation summary is saved at:

        /opt/VRTS/install/logs/installvcs208093610.summary

The installvcs log is saved at:

        /opt/VRTS/install/logs/installvcs208093610.log

The installation response file is saved at:

        /opt/VRTS/install/logs/installvcs208093610.response

When installvcs installs software, some software may be applied rather than
committed. It is the responsibility of the system administrator to commit
the software, which can be performed later with the -c option of the
installp command.

# ▌
```

Figure 17-26 Final VCS installation screen

Archived

18

VERITAS Cluster Server on AIX and IBM Tivoli Storage Manager Server

In this chapter we provide details regarding the installation of the Tivoli Storage Manager V5.3 server software, and configuring it as an application within a VCS Service Group. We then do some testing of VCS and the Tivoli Storage Manager server functions within the VCS cluster.

© Copyright IBM Corp. 2005. All rights reserved.

18.1 Overview

In the following topics, we discuss (and demonstrate) the physical installation of the application software (Tivoli Storage Manager server and the Tivoli Storage Manager Backup Archive client).

18.2 Installation of Tivoli Storage Manager Server

We will begin with the installation of the Tivoli Storage Manager server component, after reviewing all the installation and readme documents.

18.2.1 Tivoli Storage Manager Server AIX filesets

For up-to-date information, always refer to the readme file that comes with the latest maintenance or patches you are going to install.

Server code

Use normal AIX install procedures (installp) to install server code filesets according to your environment at the latest level on both cluster nodes:

32-bit hardware, 32-bit AIX kernel

tivoli.tsm.server.com
tivoli.tsm.server.rte
tivoli.tsm.msg.en_US.server
tivoli.tsm.license.cert
tivoli.tsm.license.rte
tivoli.tsm.webcon
tivoli.tsm.msg.en-US.devices
tivoli.tsm.devices.aix5.rte

64-bit hardware, 64-bit AIX kernel

tivoli.tsm.server.com
tivoli.tsm.server.aix5.rte64
tivoli.tsm.msg.en_US.server
tivoli.tsm.license.cert
tivoli.tsm.license.rte
tivoli.tsm.webcon
tivoli.tsm.msg.en-US.devices
tivoli.tsm.devices.aix5.rte

Archived

64-bit hardware, 32-bit AIX kernel

tivoli.tsm.server.com
tivoli.tsm.server.rte
tivoli.tsm.msg.en_US.server
tivoli.tsm.license.cert
tivoli.tsm.license.rte
tivoli.tsm.webcon
tivoli.tsm.msg.en-US.devices
tivoli.tsm.devices.aix5.rte

18.2.2 Tivoli Storage Manager Client AIX filesets

Important: The *Command Line Administrative Interface* is necessary to be installed during this process (**dsmadmc** command).

Even if we were not planning to utilize the Tivoli Storage Manager client, we would still need these components installed on both servers, as the scripts for VCS starting and stopping the client require the **dsmadmc** command. In addition, we will be performing some initial Tivoli Storage Manager server configuration using the **dsmadmc** command line.

tivoli.tsm.client.api.32bit
tivoli.tsm.client.ba.32bit.base
tivoli.tsm.client.ba.32bit.common
tivoli.tsm.client.ba.32bit.web

18.2.3 Tivoli Storage Manager Client Installation

We will install the Tivoli Storage Manager client into the default location of /usr/tivoli/tsm/client/ba/bin and the API into /usr/tivoli/tsm/client/api/bin on all systems in the cluster.

1. First we change into the directory which holds our installation images, and issue the **smitty installp** AIX command as shown in Figure 18-1.

Archived

```
                         Install and Update Software

Move cursor to desired item and press Enter.

  Install Software
  Update Installed Software to Latest Level (Update All)
  Install Software Bundle
  Update Software by Fix (APAR)
  Install and Update from ALL Available Software

F1=Help             F2=Refresh          F3=Cancel           F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 18-1 The smit install and update panel

2. Then, for the input device we used a **dot**, implying the current directory as shown in Figure 18-2.

```
                 Install and Update from ALL Available Software

Type or select a value for the entry field.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* INPUT device / directory for software                 [.]                +

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 18-2 Launching SMIT from the source directory, only dot (.) is required

Archived

3. For the next **smit** panel, we select a LIST using the F4 key.
4. We then select the required filesets to install using the F7 key, as seen in Figure 18-3.

```
                         SOFTWARE to install

Move cursor to desired item and press F7. Use arrow keys to scroll.
    ONE OR MORE items can be selected.
Press Enter AFTER making all selections.


  tivoli.tsm.books.en_US.client                                      ALL
   + 5.3.0.0  TSM BOOKS - Using UNIX Clients - HTML Format
   + 5.3.0.0  TSM BOOKS - Using UNIX Clients - PDF Format

  tivoli.tsm.client.api.32bit                                        ALL
>  + 5.3.0.0  TSM Client - Application Programming Interface 32bit

  tivoli.tsm.client.api.64bit                                        ALL
   + 5.3.0.0  TSM Client - 64bit Application Programming Interface

  tivoli.tsm.client.ba                                               ALL
>  + 5.3.0.0  TSM Client - Backup/Archive Base Files
>  + 5.3.0.0  TSM Client - Backup/Archive Common Files
>  + 5.3.0.0  TSM Client - Backup/Archive WEB Client
   + 5.3.0.0  TSM Client - IMAGE Backup Client
   + 5.3.0.0  TSM Client - NAS Backup Client

F1=Help             F2=Refresh            F3=Cancel
F7=Select           F8=Image              F10=Exit
Enter=Do            /=Find                n=Find Next
```

Figure 18-3 AIX installp filesets chosen for client installation

Archived

5. After making the selection and pressing Enter, we change the default **smit** panel options to allow for a detailed preview first, as shown in Figure 18-4.

```
                    Install and Update from ALL Available Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* INPUT device / directory for software                  .
* SOFTWARE to install                                   [tivoli.tsm.client.api> +
  PREVIEW only? (install operation will NOT occur)       yes                    +
  COMMIT software updates?                               no                     +
  SAVE replaced files?                                   no                     +
  AUTOMATICALLY install requisite software?              yes                    +
  EXTEND file systems if space needed?                   yes                    +
  OVERWRITE same or newer versions?                      no                     +
  VERIFY install and check file sizes?                   no                     +
  DETAILED output?                                       yes                    +
  Process multiple volumes?                              yes                    +
  ACCEPT new license agreements?                         no                     +
  Preview new LICENSE agreements?                        no                     +

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 18-4 Changing the defaults to preview with detail first prior to installing

6. Following a successful preview, we change the **smit** panel configuration to reflect a detailed and committed installation as shown in Figure 18-5.

```
                    Install and Update from ALL Available Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* INPUT device / directory for software                  .
* SOFTWARE to install                                   [tivoli.tsm.client.api> +
  PREVIEW only? (install operation will NOT occur)       no                     +
  COMMIT software updates?                               yes                    +
  SAVE replaced files?                                   no                     +
  AUTOMATICALLY install requisite software?              yes                    +
  EXTEND file systems if space needed?                   yes                    +
  OVERWRITE same or newer versions?                      no                     +
  VERIFY install and check file sizes?                   no                     +
  DETAILED output?                                       yes                    +
  Process multiple volumes?                              yes                    +
  ACCEPT new license agreements?                         no                     +
  Preview new LICENSE agreements?                        no                     +

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 18-5 The smit panel demonstrating a detailed and committed installation

Archived

7. Finally, we review the installed filesets using the **AIX command lslpp** as shown in Figure 18-6.

```
kanaga/images/TSM/client: lslpp -h |grep -p tsm.client
  tivoli.tsm.client.api.32bit
                  5.3.0.0   COMMIT       COMPLETE     01/24/05     17:16:11

  tivoli.tsm.client.ba.32bit.base
                  5.3.0.0   COMMIT       COMPLETE     01/24/05     17:16:18

  tivoli.tsm.client.ba.32bit.common
                  5.3.0.0   COMMIT       COMPLETE     01/24/05     17:16:18

  tivoli.tsm.client.ba.32bit.web
                  5.3.0.0   COMMIT       COMPLETE     01/24/05     17:16:18
```

Figure 18-6 AIX lslpp command to review the installed filesets

8. Finally, we repeat this same process on the other node in this cluster.

18.2.4 Installing the Tivoli Storage Manager server software

We will install the Tivoli Storage Manager server into the default location of /usr/tivoli/tsm/server/bin on all systems in the cluster which could host the Tivoli Storage Manager server if a failover were to occur.

1. First we change into the directory which holds our installation images, and issue the **smitty installp** AIX command, which presents the first install panel as shown in Figure 18-7.

```
                         Install and Update Software

Move cursor to desired item and press Enter.

  Install Software
  Update Installed Software to Latest Level (Update All)
  Install Software Bundle
  Update Software by Fix (APAR)
  Install and Update from ALL Available Software

F1=Help             F2=Refresh          F3=Cancel           F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 18-7 The smit software installation panel

Archived

2. Then, for the input device we used a **dot**, implying the current directory as shown in Figure 18-8.

```
                    Install and Update from ALL Available Software

Type or select a value for the entry field.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* INPUT device / directory for software                 [.]                +

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 18-8 The smit input device panel

3. Next, we select the filesets which will be required for our clustered environment, using the F7 key. Our selection is shown in Figure 18-9.

Archived

```
                 Install and Update from ALL Available Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* INPUT device / directory for software                 .
*
                           SOFTWARE to install

    Move cursor to desired item and press F7. Use arrow keys to scroll.
        ONE OR MORE items can be selected.
    Press Enter AFTER making all selections.


    > tivoli.tsm.devices.aix5.rte                                      ALL
       + 5.3.0.0  IBM Tivoli Storage Manager Device Support runtime

    > tivoli.tsm.license.aix5.rte64                                    ALL
       + 5.3.0.0  IBM Tivoli Storage Manager 64 bit  License Registration

    > tivoli.tsm.license.cert                                          ALL
       + 5.3.0.0  IBM Tivoli Storage Manager License Certificates

    > tivoli.tsm.loc.server.ela                                        ALL
       + 5.3.0.0  IBM Tivoli Storage Manager Server E-Lic

    > tivoli.tsm.msg.EN_US.devices                                     ALL
       + 5.3.0.0  IBM Tivoli Storage Manager Devices SMIT Menus, US English

    > tivoli.tsm.msg.EN_US.server                                      ALL
       + 5.3.0.0  IBM Tivoli Storage Manager Server Msgs, US English utf8

    > tivoli.tsm.msg.en_US.devices                                     ALL
       + 5.3.0.0  IBM Tivoli Storage Manager Devices SMIT Menus, US English

    > tivoli.tsm.msg.en_US.server                                      ALL
       + 5.3.0.0  IBM Tivoli Storage Manager Server Msgs, US English

    > tivoli.tsm.server.aix5.rte64                                     ALL
       + 5.3.0.0  IBM Tivoli Storage Manager 64 bit Server Runtime

    > tivoli.tsm.server.com                                            ALL
       + 5.3.0.0  IBM Tivoli Storage Manager Server common services

    > tivoli.tsm.server.webcon                                         ALL
       + 5.3.0.0  IBM Tivoli Storage Manager Web Console intfc

    F1=Help             F2=Refresh          F3=Cancel
F1  F7=Select           F8=Image            F10=Exit
F5  Enter=Do            /=Find              n=Find Next
```

Figure 18-9 The smit selection screen for filesets

4. We then press Enter after the selection has been made.
5. On this next panel presented, we change the default values for preview, commit, detailed, accept. This allows us to verify that we have all the prerequisites installed prior to running a commit installation. The changes to these default options are shown in Figure 18-10.

Archived

```
                    Install and Update from ALL Available Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* INPUT device / directory for software                  .
* SOFTWARE to install                                   []                    +
  PREVIEW only? (install operation will NOT occur)      yes                   +
  COMMIT software updates?                              no                    +
  SAVE replaced files?                                  no                    +
  AUTOMATICALLY install requisite software?             yes                   +
  EXTEND file systems if space needed?                  yes                   +
  OVERWRITE same or newer versions?                     no                    +
  VERIFY install and check file sizes?                  no                    +
  DETAILED output?                                      yes                   +
  Process multiple volumes?                             yes                   +
  ACCEPT new license agreements?                        yes                   +
  Preview new LICENSE agreements?                       no                    +


F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 18-10 The smit screen showing non-default values for a detailed preview

6. After we successfully complete the preview, we change the installation panel to reflect a detailed, committed installation and accepting new license agreements. This is shown in Figure 18-11.

```
                    Install and Update from ALL Available Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* INPUT device / directory for software                  .
* SOFTWARE to install                                   [tivoli.tsm.devices.ai> +
  PREVIEW only? (install operation will NOT occur)      no                    +
  COMMIT software updates?                              yes                   +
  SAVE replaced files?                                  no                    +
  AUTOMATICALLY install requisite software?             yes                   +
  EXTEND file systems if space needed?                  yes                   +
  OVERWRITE same or newer versions?                     no                    +
  VERIFY install and check file sizes?                  no                    +
  DETAILED output?                                      yes                   +
  Process multiple volumes?                             yes                   +
  ACCEPT new license agreements?                        yes                   +
  Preview new LICENSE agreements?                       no                    +


F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 18-11 The final smit install screen with selections and a commit installation

Archived

7. After the installation has been successfully completed, we review the installed filesets from the AIX command line with the **lslpp** command, as shown in Figure 18-12.

```
  tivoli.tsm.loc.server.ela
                  5.3.0.0   COMMIT       COMPLETE     01/24/05     15:22:45

  tivoli.tsm.msg.EN_US.server
                  5.3.0.0   COMMIT       COMPLETE     01/24/05     15:22:42

  tivoli.tsm.msg.en_US.server
                  5.3.0.0   COMMIT       COMPLETE     01/24/05     15:22:44

  tivoli.tsm.server.aix5.rte64
                  5.3.0.0   COMMIT       COMPLETE     01/24/05     15:22:34

  tivoli.tsm.server.com
                  5.3.0.0   COMMIT       COMPLETE     01/24/05     15:22:30

  tivoli.tsm.server.webcon
                  5.3.0.0   COMMIT       COMPLETE     01/24/05     15:22:54

  tivoli.tsm.server.aix5.rte64
                  5.3.0.0   COMMIT       COMPLETE     01/24/05     15:22:42

  tivoli.tsm.server.com
                  5.3.0.0   COMMIT       COMPLETE     01/24/05     15:22:33

kanaga/images/TSM/server:
```

Figure 18-12 AIX lslpp command listing of the server installp images

8. Lastly, we repeat all of these processes on the other cluster node.

18.3 Configuration for clustering

Now we provide details about the configuration of the Veritas Cluster Server, including the configuration of the Tivoli Storage Manager server as a highly available application.

We will prepare the environments prior to configuring this application in the VCS cluster, and ensure that the Tivoli Storage Manager server and client communicate properly prior to HA configuration.

VCS will require start, stop, monitor, and clean scripts for most of the applications. Creating and testing these prior to implementing the Service Group configuration is a good approach.

Archived

18.3.1 Tivoli Storage Manager server configuration

In 17.5, “VCS pre-installation” on page 723, we prepared the needed storage, network, and volume resources. We now utilize these resources during the Tivoli Storage Manager server configuration, and develop the start and stop scripts to be used by the VCS cluster:

1. First, we remove the entry from /etc/inittab on both nodes, which auto starts the IBM Tivoli Storage Manager server, StorageAgent, and ISC, using the **rmitab autosrvr** command, as shown in Example 18-1.

Example 18-1 The AIX rmitab command

```
banda:/# rmitab autosrvr
banda:/# rmitab autostgagnt
banda:/# rmitab iscn

Atlantic:/# rmitab autosrvr
Atlantic:/# rmitab autostgagnt
Atlantic:/# rmitab iscn
```

2. We stop the default server installation instance, if running, as shown in Example 18-2. Using the **kill** command (without the **-9** option) will shut down the Tivoli Storage Manager server process and the associated threads.

Example 18-2 Stop the initial server installation instance

```
# ps -ef|grep dsmserv
root 41304 176212 0 09:52:48 pts/3 0:00 grep dsmserv
root 229768     1 0 07:39:36 - 0:56 /usr/tivoli/tsm/server/bin/dsmserv quiet
# kill 229768
```

3. Next, we set up the appropriate IBM Tivoli Storage Manager server directory environment setting for the current shell issuing the following commands, as shown in Example 18-3.

Example 18-3 The variables which must be exported in our environment

```
# export DSMSERV_CONFIG=/tsm/files/dsmserv.opt
# export DSMSERV_DIR=/usr/tivoli/tsm/server/bin
```

4. Then, we clean up the default server installation files which are not required, and must be completed on both nodes. We will remove the default created database, recovery log, space management, archive, and backup files created. We also remove the dsmserv.opt and dsmserv.dsk files which will be located on the shared disk. These commands are shown in Example 18-4.

Archived

Example 18-4 Files to remove after the initial server installation

```
# cd /usr/tivoli/tsm/server/bin
# mv dsmserv.opt /tsm/files
# mv dsmserv.dsk /tsm/files
# rm db.dsm
# rm spcmgmt.dsm
# rm log.dsm
# rm backup.dsm
# rm archive.dsm
```

5. Next, we configure IBM Tivoli Storage Manager to use the TCP/IP communication method. See the *Installation Guide* for more information on specifying server and client communications. We verify that the /tsm/files/dsmserv.opt file reflects our requirements.
6. Then we configure the local client to communicate with the server, (only basic communication parameters in dsm.sys found in the /usr/tivoli/tsm/client/ba/bin directory). We will use this initially for the Command Line Administrative Interface. This configuration stanza is shown in Example 18-5.

Example 18-5 The server stanza for the client dsm.sys file

```
* Server stanza for admin connection purpose
SErvername tsmsrv04_admin
COMMMethod TCPip
TCPPor 1500
TCPServeraddress 127.0.0.1
ERRORLOGRETENTION 7
ERRORLOGname /usr/tivoli/tsm/client/ba/bin/dsmerror.log
```

Tip: For information about running the server from a directory different from the default database that was created during the server installation, see the Installation Guide, which can be found at:

http://publib.boulder.ibm.com/infocenter/tivihelp/index.jsp?topic=/com.ibm.i

7. Allocate the IBM Tivoli Storage Manager database, recovery log, and storage pools on the shared IBM Tivoli Storage Manager volume group. To accomplish this, we will use the **dsmfmt** command to format database, log, and disk storage pool files on the shared file systems. This is shown in Example 18-6.

Archived

Example 18-6 dsmfmt command to create database, recovery log, storage pool files

```
# dsmfmt -m -db /tsm/db1/vol1 2000
# dsmfmt -m -db /tsm/dbmr1/vol1 2000
# dsmfmt -m -log /tsm/lg1/vol1 1000
# dsmfmt -m -log /tsm/lgmr1/vol1 1000
# dsmfmt -m -data /tsm/dp1/bckvol1 25000
```

8. We change the current directory to the new server directory and we then issue the **dsmserv format command** to install the database which will create the dsmserv.dsk, as shown in Example 18-7.

Example 18-7 The dsmserv format command to prepare the recovery log

```
# cd /tsm/files
# dsmserv format 1 /tsm/lg1/vol1 1 /tsm/db1/vol1
```

9. Next, we start the Tivoli Storage Manager server in the foreground by issuing the command dsmserv from the installation directory and with the environment variables set within the running shell, as shown in Example 18-8.

Example 18-8 An example of starting the server in the foreground

```
dsmserv
```

10. Once the Tivoli Storage Manager server has completed the started, we run the Tivoli Storage Manager server commands; **set servername**, and then mirror database and log, as shown in Example 18-9.

Example 18-9 The server setup for use with our shared disk files

```
TSM:SERVER1> set servername tsmsrv04
TSM:TSMSRV04> define dbcopy /tsm/db1/vol1 /tsm/dbmr1/vol1
TSM:TSMSRV04> define logcopy /tsm/lg1/vol1 /tsm/lgmr1/vol1
```

11. We then define a DISK storage pool with a volume on the shared filesystem /tsm/dp1 which is configured as a RAID1 protected storage device, shown here in Example 18-10.

Example 18-10 The define commands for the diskpool

```
TSM:TSMSRV04> define stgpool spd_bck disk
TSM:TSMSRV04> define volume spd_bck /tsm/dp1/bckvol1
```

12. We now define the tape library and tape drive configurations using the define library, define drive and define path commands, demonstrated in Example 18-11.

Example 18-11 An example of define library, define drive and define path commands

```
TSM:TSMSRV04> define library liblto libtype=scsi
TSM:TSMSRV04> define path tsmsrv04 liblto srctype=server desttype=libr
device=/dev/smc0
TSM:TSMSRV04> define drive liblto drlto_1
TSM:TSMSRV04> define drive liblto drlto_2
TSM:TSMSRV04> define path tsmsrv04 drlto_1 srctype=server desttype=drive
libr=liblto device=/dev/rmt0
TSM:TSMSRV04> define path tsmsrv04 drlto_2 srctype=server desttype=drive
libr=liblto device=/dev/rmt1
```

13. We will now register the *admin* administrator with the *system* authority with the register admin and grant authority commands. Also, we will need another ID for our scripts, and we will call this one script_operator, as shown in Example 18-12.

Example 18-12 The register admin and grant authority commands

```
TSM:TSMSRV04> reg admin admin admin
   TSM:TSMSRV04> grant authority admin classes=system
TSM:TSMSRV04> reg admin script_operator password
   TSM:TSMSRV04> grant authority script_operator classes=system
```

18.4 Veritas Cluster Manager configuration

The installation process configured the cluster and core services for us, now we need to configure the Service Groups and their associated Resources for the Tivoli Storage Manager server, client, Storage Agent, and the ISC.

18.4.1 Preparing and placing application startup scripts

We will develop and test our start, stop, clean and monitor scripts for all of our applications, then place them in the /opt/local directory on each node, which is a local filesystem within the rootvg.

Scripts for the Tivoli Storage Manager server

We placed the scripts for the server in the rootvg, /opt filesystem, in the directory /opt/local/tsmsrv.

1. The start script, which is supplied with Tivoli Storage Manager as a sample for HACMP, works fine for this VCS environment. We placed the script in our /opt/local/tsmsrv directory as /opt/local/tsmsrv/startTSMsrv.sh is shown in Example 18-13.

Example 18-13 /opt/local/tsmsrv/startTSMsrv.sh

```
#!/bin/ksh
##############################################################################
#                                                                            #
# Shell script to start a TSM server.                                        #
#                                                                            #
# Please note commentary below indicating the places where this shell script  #
# may need to be modified in order to tailor it for your environment.        #
#                                                                            #
##############################################################################
#                                                                            #
# Update the cd command below to change to the directory that contains the   #
# dsmserv.dsk file and change the export commands to point to the dsmserv.opt #
# file and /usr/tivoli/tsm/server/bin directory for the TSM server being     #
# started.  The export commands are currently set to the defaults.           #
#                                                                            #
##############################################################################
echo "Starting TSM now..."

cd /tsm/files

export DSMSERV_CONFIG=/tsm/files/dsmserv.opt
export DSMSERV_DIR=/usr/tivoli/tsm/server/bin

# Allow the server to pack shared memory segments
export EXTSHM=ON

# max out size of data area
ulimit -d unlimited

# Make sure we run in the correct threading environment
export AIXTHREAD_MNRATIO=1:1
export AIXTHREAD_SCOPE=S

##############################################################################
#                                                                            #
# set the server language.  These two statements need to be modified by the  #
# user to set the appropriate language.                                      #
#                                                                            #
##############################################################################
export LC_ALL=en_US
export LANG=en_US

#OK, now fire-up the server in quiet mode.
$DSMSERV_DIR/dsmserv quiet &
```

2. We then placed the stop script in as /opt/local/tsmsrv/stopTSMsrv.sh, shown in Example 18-14.

Example 18-14 /opt/local/tsmsrv/stopTSMsrv.sh

```
#!/bin/ksh
##############################################################################
# Shell script to stop a TSM AIX server.
# Please note that changes must be made to the dsmadmc command below in order
# to tailor it for your environment:
#
#    1. Set -servername= to the TSM server name on the SErvername option
#       in the /usr/tivoli/tsm/client/ba/bin/dsm.sys file.
#    2. Set -id= and -password= to a TSM userid that has been granted
#       operator authority, as described in the section:
#       "Chapter 3. Customizing Your Tivoli Storage Manager System -
#       Adding Administrators", in the Quick Start manual.
#
#    3. Edit the path in the LOCKFILE= statement to the directory where your
#       dsmserv.dsk file exists for this server.
#
#
# Author: Steve Pittman
#
# Date: 12/6/94
#
# Modifications:
#
#  4/20/2004  Bohm.  IC39681, fix incorrect indentation.
#
# 10/21/2002  David Bohm.  IC34520, don't exit from the script if there are
#             kernel threads running.
#
#  7/03/2001  David Bohm.  Made changes for support of the TSM server.
#             General clean-up.
#
#
#
##############################################################################
#
# Set seconds to sleep.
secs=2
# TSM lock file
LOCKFILE="/tsm/files/adsmserv.lock"

echo "Stopping the TSM server now..."

# Check to see if the adsmserv.lock file exists.  If not then the server is not
running
if [[ -f $LOCKFILE ]]; then

  read J1 J2 J3 PID REST < $LOCKFILE
```

```
  /usr/tivoli/tsm/client/ba/bin/dsmadmc -servername=tsmsrv04_admin -id=admin
-password=admin -noconfirm << EOF
halt
EOF
  echo "Waiting for TSM server running on pid $PID to stop..."

# Make sure all of the threads have ended
  while [[ `ps -m -o THREAD -p $PID | grep -c $PID` > 0 ]]; do
    sleep $secs
  done
fi
exit 0
```

3. Next, we placed the clean script in as /opt/local/tsmsrv/cleanTSMsrv.sh, shown in Example 18-15.

Example 18-15 /opt/local/tsmsrv/cleanTSMsrv.sh

```
#!/bin/ksh
##############################################################################
# Shell script to stop a TSM AIX server.
#
# Please note that changes must be made to the dsmadmc command below in order
# to tailor it for your environment: #
#    1. Set -servername= to the TSM server name on the SErvername option
#       in the /usr/tivoli/tsm/client/ba/bin/dsm.sys file. #
#    2. Set -id= and -password= to a TSM userid that has been granted
#       operator authority, as described in the section:
#       "Chapter 3. Customizing Your Tivoli Storage Manager System -
#       Adding Administrators", in the Quick Start manual.
#
#    3. Edit the path in the LOCKFILE= statement to the directory where your
#       dsmserv.dsk file exists for this server. #
# Author: Steve Pittman
#

# Date: 12/6/94
#
# Modifications:
#
#  4/20/2004  Bohm.  IC39681, fix incorrect indentation.
#
# 10/21/2002  David Bohm.  IC34520, don't exit from the script if there are
#             kernel threads running.
#
#  7/03/2001  David Bohm.  Made changes for support of the TSM server.
#             General clean-up.
```

```
##############################################################################
#
# Set seconds to sleep.
secs=2
# TSM lock file
LOCKFILE="/tsm/files/adsmserv.lock"

echo "Stopping the TSM server now..."

# Check to see if the adsmserv.lock file exists.  If not then the server is not
running
if [[ -f $LOCKFILE ]]; then

  read J1 J2 J3 PID REST < $LOCKFILE
  /usr/tivoli/tsm/client/ba/bin/dsmadmc -servername=tsmsrv04_admin -id=admin
-password=admin -noconfirm << EOF
halt
EOF
  echo "Waiting for TSM server running on pid $PID to stop..."

# Make sure all of the threads have ended
  while [[ `ps -m -o THREAD -p $PID | grep -c $PID` > 0 ]]; do
    sleep $secs
  done
fi
exit 0
atlantic:/opt/local/tsmsrv#
atlantic:/opt/local/tsmsrv#
atlantic:/opt/local/tsmsrv# cleanTSMsrv.sh
/usr/bin/ksh: cleanTSMsrv.sh:  not found.
atlantic:/opt/local/tsmsrv# ls
cleanTSMsrv.sh  monTSMsrv.sh    startTSMsrv.sh  stopTSMsrv.sh
atlantic:/opt/local/tsmsrv# cat cleanTSMsrv.sh
#!/bin/ksh
# killing TSM server process if the stop fails

TSMSRVPID=`ps -ef | egrep "dsmserv" | awk '{ print $2 }'`

for PID in $TSMSRVPID
 do
  kill  $PID
 done

exit 0
```

4. Lastly, we placed the monitor script in as /opt/local/tsmsrv/monitorTSMsrv.sh, shown in Example 18-16.

Example 18-16 /opt/local/tsmsrv/monTSMsrv.sh

```
#!/bin/ksh
#######################################################
#
# Module:       monitortsmsrv04.sh
#
# Function:     Simple query to ensure TSM is running and responsive
#
# Author:       Dan Edwards (IBM Canada Ltd.)
#
# Date:         February 09, 2005
#
#######################################################

# Define some variables for use throughout the script
export ID=admin                 # TSM admin ID
export PASS=admin               # TSM admin password

#
#Query tsmsrv looking for a response
#
/usr/tivoli/tsm/client/ba/bin/dsmadmc -id=${ID} -pa=${PASS} "q session"
>/dev/console 2>&1
#
if [ $? -gt 0 ]
        then exit 100
fi
#
exit 110
```

Tip: The return codes for the monitor are important, RC=100 means the application is OFFLINE, and the RC=110 means the application is ONLINE with the highest level of confidence.

5. We then test the scripts to ensure that everything works as expected, prior to configuring VCS.

Hint: It is possible to configure just a process monitoring, instead of using a script, which in most cases will work very well. In the case of a Tivoli Storage Manager server, the process could be listed in the process tree, yet not responding to connection requests. For this reason, using the **dsmadmc** command will allow confirmation that connections are possible. Using a more complex query could also improve state determination if required.

Archived

18.4.2 Service Group and Application configuration

Now we configure Tivoli Storage Manager as Service Group and application.

1. First, we will use the command line options to configure the sg_tsmsrv Service Group, as shown in Example 18-17.

Example 18-17 Adding a Service Group sg_tsmsrv

```
hagrp -add sg_tsmsrv
hagrp -modify sg_tsmsrv SystemList banda 0 atlantic 1
hagrp -modify sg_tsmsrv AutoStartList banda atlantic
hagrp -modify sg_tsmsrv Parallel 0
hagrp -modify sg_tsmsrv_tsmcli AutoStartList banda atlantic
hagrp -modify sg_tsmsrv Parallel 0
```

2. Next, we add the NIC Resource for this Service Group. This monitors the NIC layer to determine if there is connectivity to the network, as shown in Example 18-18.

Example 18-18 Adding a NIC Resource

```
hares -add NIC_en1 NIC sg_tsmsrv
hares -modify NIC_en1 Critical 1
hares -modify NIC_en1 PingOptimize 1
hares -modify NIC_en1 Device en1
hares -modify NIC_en1 NetworkType ether
hares -modify NIC_en1 NetworkHosts -delete -keys
hares -probe NIC_en1 -sys banda
hares -probe NIC_en1 -sys atlantic
hares -modify NIC_en1 Enabled 1
```

3. Next, we add the IP Resource for this Service Group. This will be the IP Address that the Tivoli Storage Manager server will be contacted at, no matter on which node it resides, as shown in Example 18-19.

Example 18-19 Configuring an IP Resource in the sg_tsmsrv Service Group

```
hares -add ip_tsmsrv IP sg_tsmsrv
hares -modify ip_tsmsrv Critical 1
hares -modify ip_tsmsrv Device en1
hares -modify ip_tsmsrv Address 9.1.39.76
hares -modify ip_tsmsrv NetMask 255.255.255.0
hares -modify ip_tsmsrv Options ""
hares -probe ip_tsmsrv -sys banda
hares -probe ip_tsmsrv -sys atlantic
hares -link ip_tsmsrv NIC_en1
hares -modify ip_tsmsrv Enabled 1
```

4. Then, we add the LVMVG Resource to the Service Group sg_tsmsrv, as shown in Example 18-20.

Archived

Example 18-20 Adding the LVMVG Resource to the sg_tsmsrv Service Group

```
hares -add vg_tsmsrv LVMVG sg_tsmsrv
hares -modify vg_tsmsrv Critical 1
hares -modify vg_tsmsrv MajorNumber 47
hares -modify vg_tsmsrv ImportvgOpt n
hares -modify vg_tsmsrv SyncODM 1
hares -modify vg_tsmsrv VolumeGroup iscvg
hares -modify vg_tsmsrv OwnerName ""
hares -modify vg_tsmsrv GroupName ""
hares -modify vg_tsmsrv Mode ""
hares -modify vg_tsmsrv VaryonvgOpt ""
hares -probe vg_tsmsrv -sys banda
hares -probe vg_tsmsrv -sys atlantic
```

5. Then, we add the Mount Resources to the sg_tsmsrv Service Group, as shown in Example 18-21.

Example 18-21 Configuring the Mount Resource in the sg_tsmsrv Service Group

```
hares -add m_tsmsrv_db1 Mount sg_tsmsrv
hares -modify m_tsmsrv_db1 Critical 1
hares -modify m_tsmsrv_db1 SnapUmount 0
hares -modify m_tsmsrv_db1 MountPoint /tsm/db1
hares -modify m_tsmsrv_db1 BlockDevice /dev/tsmdb1lv
hares -modify m_tsmsrv_db1 FSType jfs2
hares -modify m_tsmsrv_db1 MountOpt ""
hares -modify m_tsmsrv_db1 FsckOpt -y
hares -probe m_tsmsrv_db1 -sys banda
hares -probe m_tsmsrv_db1 -sys atlantic
hares -link m_tsmsrv_db1 vg_tsmsrv
hares -modify m_tsmsrv_db1 Enabled 1

hares -add m_tsmsrv_dbmr1 Mount sg_tsmsrv
hares -modify m_tsmsrv_dbmr1 Critical 1
hares -modify m_tsmsrv_dbmr1 SnapUmount 0
hares -modify m_tsmsrv_dbmr1 MountPoint /tsm/dbmr1
hares -modify m_tsmsrv_dbmr1 BlockDevice /dev/tsmdbmr1lv
hares -modify m_tsmsrv_dbmr1 FSType jfs2
hares -modify m_tsmsrv_dbmr1 MountOpt ""
hares -modify m_tsmsrv_dbmr1 FsckOpt -y
hares -probe m_tsmsrv_dbmr1 -sys banda
hares -probe m_tsmsrv_dbmr1 -sys atlantic
hares -link m_tsmsrv_dbmr1 vg_tsmsrv
hares -modify m_tsmsrv_dbmr1 Enabled 1

hares -add m_tsmsrv_lg1 Mount sg_tsmsrv
hares -modify m_tsmsrv_lg1 Critical 1
hares -modify m_tsmsrv_lg1 SnapUmount 0
hares -modify m_tsmsrv_lg1 MountPoint /tsm/lg1
```

```
hares -modify m_tsmsrv_lg1 BlockDevice /dev/tsmlg1lv
hares -modify m_tsmsrv_lg1 FSType jfs2
hares -modify m_tsmsrv_lg1 MountOpt ""
hares -modify m_tsmsrv_lg1 FsckOpt -y
hares -probe m_tsmsrv_lg1 -sys banda
hares -probe m_tsmsrv_lg1 -sys atlantic
hares -link m_tsmsrv_lg1 vg_tsmsrv
hares -modify m_tsmsrv_lg1 Enabled 1

hares -add m_tsmsrv_lgmr1 Mount sg_tsmsrv
hares -modify m_tsmsrv_lgmr1 Critical 1
hares -modify m_tsmsrv_lgmr1 SnapUmount 0
hares -modify m_tsmsrv_lgmr1 MountPoint /tsm/lgmr1
hares -modify m_tsmsrv_lgmr1 BlockDevice /dev/tsmlgmr1lv
hares -modify m_tsmsrv_lgmr1 FSType jfs2
hares -modify m_tsmsrv_lgmr1 MountOpt ""
hares -modify m_tsmsrv_lgmr1 FsckOpt -y
hares -probe m_tsmsrv_lgmr1 -sys banda
hares -probe m_tsmsrv_lgmr1 -sys atlantic
hares -link m_tsmsrv_lgmr1 vg_tsmsrv
hares -modify m_tsmsrv_lgmr1 Enabled 1

hares -add m_tsmsrv_dp1 Mount sg_tsmsrv
hares -modify m_tsmsrv_dp1 Critical 1
hares -modify m_tsmsrv_dp1 SnapUmount 0
hares -modify m_tsmsrv_dp1 MountPoint /tsm/dp1
hares -modify m_tsmsrv_dp1 BlockDevice /dev/tsmdp1lv
hares -modify m_tsmsrv_dp1 FSType jfs2
hares -modify m_tsmsrv_dp1 MountOpt ""
hares -modify m_tsmsrv_dp1 FsckOpt -y
hares -probe m_tsmsrv_dp1 -sys banda
hares -probe m_tsmsrv_dp1 -sys atlantic
hares -link m_tsmsrv_dp1 vg_tsmsrv
hares -modify m_tsmsrv_dp1 Enabled 1

hares -add m_tsmsrv_files Mount sg_tsmsrv
hares -modify m_tsmsrv_files Critical 1
hares -modify m_tsmsrv_files SnapUmount 0
hares -modify m_tsmsrv_files MountPoint /tsm/files
hares -modify m_tsmsrv_files BlockDevice /dev/tsmlv
hares -modify m_tsmsrv_files FSType jfs2
hares -modify m_tsmsrv_files MountOpt ""
hares -modify m_tsmsrv_files FsckOpt -y
hares -probe m_tsmsrv_files -sys banda
hares -probe m_tsmsrv_files -sys atlantic
hares -link m_tsmsrv_files vg_tsmsrv
hares -modify m_tsmsrv_files Enabled 1
```

Archived

6. Then, we configure the Application Resource for the sg_tsmsrv Service Group as shown in Example 18-22.

Example 18-22 Adding and configuring the app_tsmsrv Application

```
hares -add app_tsmsrv Application sg_tsmsrv
hares -modify app_tsmsrv User ""
hares -modify app_tsmsrv StartProgram /opt/local/tsmsrv/startTSMsrv.sh
hares -modify app_tsmsrv StopProgram /opt/local/tsmsrv/stopTSMsrv.sh
hares -modify app_tsmsrv CleanProgram /opt/local/tsmsrv/cleanTSMsrv.sh
hares -modify app_tsmsrv MonitorProgram /opt/local/tsmsrv/monTSMsrv.sh
hares -modify app_tsmsrv PidFiles -delete -keys
hares -modify app_tsmsrv MonitorProcesses -delete -keys
hares -probe app_tsmsrv -sys banda
hares -probe app_tsmsrv -sys atlantic
hares -link app_tsmsrv m_tsmsrv_files
hares -link app_tsmsrv m_tsmsrv_dp1
hares -link app_tsmsrv m_tsmsrv_lgmr1
hares -link app_tsmsrv m_tsmsrv_lg1
hares -link app_tsmsrv m_tsmsrv_db1mr1
hares -link app_tsmsrv m_tsmsrv_db1
hares -link app_tsmsrv ip_tsmsrv
hares -modify app_tsmsrv Enabled 1
```

7. Then, from within the Veritas Cluster Manager GUI, we review the setup and links, which demonstrate the resources in a child-parent relationship, as shown in Figure 18-13.

Archived

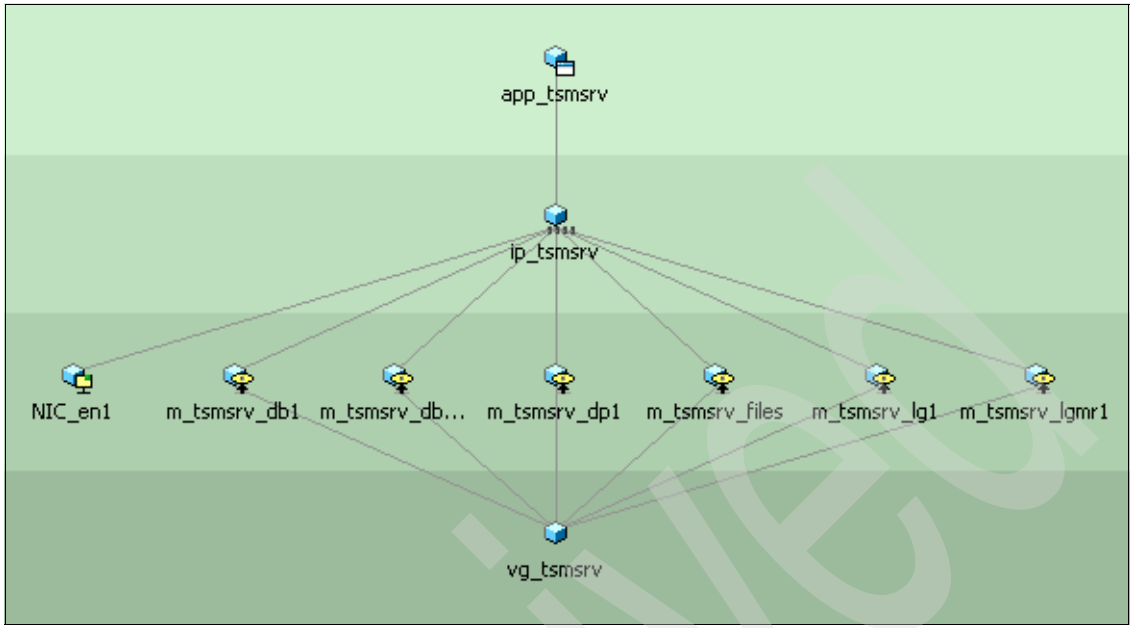

Figure 18-13 Child-parent relationships within the sg_tsmsrv Service Group.

8. Next, we review the main.cf file, which is shown in Example 18-23.

Example 18-23 The sg_tsmsrv Service Group: /etc/VRTSvcs/conf/config/main.cf file

```
group sg_tsmsrv (
        SystemList = {banda = 0, atlantic = 1}
        AutoStartList = {banda, atlantic}
        )

        Application app_tsmsrv (
                StartProgram = "/opt/local/tsmsrv/startTSMsrv.sh"
                StopProgram = "/opt/local/tsmsrv/stopTSMsrv.sh"
                CleanProgram = "/opt/local/tsmsrv/cleanTSMsrv.sh"
                MonitorProcesses = {"/usr/tivoli/tsm/server/bin/dsmserv quiet"
}
                )

        IP ip_tsmsrv (
                ComputeStats = 1
                Device = en1
                Address = "9.1.39.76"
                NetMask = "255.255.255.0"
                )

        LVMVG vg_tsmsrv (
```

Archived

```
        VolumeGroup = tsmvg
        MajorNumber = 47
        )

Mount m_tsmsrv_db1 (
        MountPoint = "/tsm/db1"
        BlockDevice = "/dev/tsmdb1lv"
        FSType = jfs2
        FsckOpt = "-y"
        )

Mount m_tsmsrv_dbmr1 (
        MountPoint = "/tsm/dbmr1"
        BlockDevice = "/dev/tsmdbmr1lv"
        FSType = jfs2
        FsckOpt = "-y"
        )

Mount m_tsmsrv_dp1 (
        MountPoint = "/tsm/dp1"
        BlockDevice = "/dev/tsmdp1lv"
        FSType = jfs2
        FsckOpt = "-y"
        )

Mount m_tsmsrv_files (
        MountPoint = "/tsm/files"
        BlockDevice = "/dev/tsmlv"
        FSType = jfs2
        FsckOpt = "-y"
        )

Mount m_tsmsrv_lg1 (
        MountPoint = "/tsm/lg1"
        BlockDevice = "/dev/tsmlg1lv"
        FSType = jfs2
        FsckOpt = "-y"
        )

Mount m_tsmsrv_lgmr1 (
        MountPoint = "/tsm/lgmr1"
        BlockDevice = "/dev/tsmlgmr1lv"
        FSType = jfs2
        FsckOpt = "-y"
        )

NIC NIC_en1 (
        Device = en1
        NetworkType = ether
```

Archived

```
        )

app_tsmsrv requires ip_tsmsrv
ip_tsmsrv requires NIC_en1
ip_tsmsrv requires m_tsmsrv_db1
ip_tsmsrv requires m_tsmsrv_db1mr1
ip_tsmsrv requires m_tsmsrv_dp1
ip_tsmsrv requires m_tsmsrv_files
ip_tsmsrv requires m_tsmsrv_lg1
ip_tsmsrv requires m_tsmsrv_lgmr1
m_tsmsrv_db1 requires vg_tsmsrv
m_tsmsrv_db1mr1 requires vg_tsmsrv
m_tsmsrv_dp1 requires vg_tsmsrv
m_tsmsrv_files requires vg_tsmsrv
m_tsmsrv_lg1 requires vg_tsmsrv
m_tsmsrv_lgmr1 requires vg_tsmsrv


// resource dependency tree
//
//      group sg_tsmsrv
//      {
//      Application app_tsmsrv
//          {
//          IP ip_tsmsrv
//              {
//              NIC NIC_en1
//              Mount m_tsmsrv_db1
//                  {
//                  LVMVG vg_tsmsrv
//                  }
//              Mount m_tsmsrv_db1mr1
//                  {
//                  LVMVG vg_tsmsrv
//                  }
//              Mount m_tsmsrv_dp1
//                  {
//                  LVMVG vg_tsmsrv
//                  }
//              Mount m_tsmsrv_files
//                  {
//                  LVMVG vg_tsmsrv
//                  }
//              Mount m_tsmsrv_lg1
//                  {
//                  LVMVG vg_tsmsrv
//                  }
//              Mount m_tsmsrv_lgmr1
//                  {
```

Archived

```
//                  LVMVG vg_tsmsrv
//                  }
//              }
//          }
//      }
```

Note: Observe the relationship tree for this configuration, which is critical, ensuring that the correct resource becomes available or stopped in the appropriate order.

9. Next, we are now ready to place the resources online and test.

18.5 Testing the cluster

We have installed and configured the Veritas Cluster Manager, and the sg_tsmsrv Service Group. Now, it is time to test the solution to ensure that it behaves as we expect.

18.5.1 Core VCS cluster testing

Here we are testing basic cluster functions.This can help in problem determination if something goes wrong later on during setup and further testing.

We determine state of the cluster services by doing the **hastatus** command from the AIX command line, and tail the main cluster log, on both systems in the cluster.

18.5.2 Node Power Failure

Initially, this test is run with the applications OFFLINE.

1. First, we verify that the Service Groups are OFFLINE using the Veritas **hastatus** command, as shown in Example 18-24.

Example 18-24 The results return from hastatus

```
banda:/# hastatus
attempting to connect....connected

group           resource             system               message
--------------- -------------------- -------------------- --------------------
                                     atlantic             RUNNING
                                     banda                RUNNING
sg_tsmsrv                            banda                OFFLINE
sg_tsmsrv                            atlantic             OFFLINE
```

2. Next, we clear the VCS log by doing the **command cp /dev/null /var/VRTSvcs/log/engine_A.log.** For testing purposes, clearing the log prior, then copying the contents of the complete log after the test to an appropriately named file, is a good methodology to reduce the log data you must sort through for a test, yet preserving the historical integrity of the test results.
3. Then, we do the AIX command, **tail -f /var/VRTSvcs/log/engine_A.log**. This allows us to monitor the transition real-time.
4. Next we fail Banda by **pulling the power plug**. The results of the **hastatus** log on the surviving node (Atlantic) is shown in Example 18-25, and the result tail of the engine_A.log on Atlantic is shown in Example 18-26.

Example 18-25 hastatus log from the surviving node, Atlantic

```
Atlantic:/var/VRTSvcs/log# hastatus
attempting to connect....connected

group           resource             system               message
--------------- -------------------- -------------------- --------------------
                                     atlantic             RUNNING
                                     banda                *FAULTED*
```

Example 18-26 tail -f /var/VRTSvcs/log/engine_A.log from surviving node, Atlantic

```
VCS INFO V-16-1-10077 Received new cluster membership
VCS NOTICE V-16-1-10080 System (atlantic) - Membership: 0x1, Jeopardy: 0x0
VCS ERROR V-16-1-10079 System banda (Node '1') is in Down State - Membership:
0x1
VCS ERROR V-16-1-10322 System banda (Node '1') changed state from RUNNING to
FAULTED
```

5. Then, we restart Banda and wait for the cluster to recover, then review the **hastatus**, which has returned to full cluster membership. This is shown in Example 18-27.

Example 18-27 The recovered cluster using hastatus

```
banda:/# hastatus
attempting to connect....connected

group           resource             system               message
--------------- -------------------- -------------------- --------------------
                                     atlantic             RUNNING
                                     banda                RUNNING
sg_tsmsrv                            banda                OFFLINE
sg_tsmsrv                            atlantic             OFFLINE
```

Archived

6. We then repeat this process for the other node, Atlantic.

Results

Once the cluster recovers, we repeat the process for the other node, ensuring that full cluster recovery occurs. Once the test has occurred on both nodes, and recovery details have been confirmed as functioning correctly, this test is complete.

18.5.3 Start Service Group (bring online)

1. To begin, we review the current cluster status, confirming that all resources are offline, as shown from the **hastatus** command output, detailed in Example 18-28.

Example 18-28 Current cluster status from the hastatus output

```
banda:/# hastatus
attempting to connect....connected

group           resource            system              message
--------------- ------------------- ------------------- -------------------
                                    atlantic            RUNNING
                                    banda               RUNNING
sg_tsmsrv                           banda               OFFLINE
sg_tsmsrv                           atlantic            OFFLINE
```

2. We then clear the log using **cp /dev/null /var/VRTSvcs/logengine_A.log** and then start a **tail -f /var/VRTSvcs/logengine_A.log**.
3. Next, from Atlantic (it can be done on any node) we bring the sg_tsmsrv Service Group online on Banda using the **hagrp** command from the AIX command line, as shown in Example 18-29.

Example 18-29 hagrp -online command

```
Atlantic:/opt/local/tsmcli# hagrp -online sg_tsmsrv -sys banda -localclus
```

4. We then view the **hastatus | grep banda** and verify the results as shown in Example 18-30.

Example 18-30 hastatus of the online transition for the sg_tsmsrv

```
banda:/# hastatus | grep ONLINE
attempting to connect....connected
sg_tsmsrv                           banda               ONLINE
sg_tsmsrv                           banda               ONLINE
                vg_tsmsrv           banda               ONLINE
                ip_tsmsrv           banda               ONLINE
```

```
                m_tsmsrv_db1         banda                ONLINE
                m_tsmsrv_db1mr1      banda                ONLINE
                m_tsmsrv_lg1         banda                ONLINE
                m_tsmsrv_lgmr1       banda                ONLINE
                m_tsmsrv_dp1         banda                ONLINE
                m_tsmsrv_files       banda                ONLINE
                app_tsmsrv           banda                ONLINE
                NIC_en1              banda                ONLINE
                NIC_en1              atlantic             ONLINE
```

5. Then we review the engine_A.log shown in Example 18-31.

Example 18-31 tail -f /var/VRTSvcs/log/engine_A.log

```
VCS INFO V-16-1-50135 User root fired command: hagrp -online sg_tsmsrv banda
localclus from localhost
.
.
.
VCS NOTICE V-16-1-10447 Group sg_tsmsrv is online on system banda
```

18.5.4 Stop Service Group (bring offline)

1. Before every test, we check the status for cluster services, resource groups and resources on both nodes; In Example 18-32 we are verifying using **hastatus**. For this test, we expect that all applications are offline, as we are just testing the cluster's core functionality.

Example 18-32 Verify available cluster resources using the hastatus command

```
banda:/var/VRTSvcs/log# hastatus
attempting to connect....connected

group           resource             system               message
--------------- -------------------- -------------------- --------------------
                                     atlantic             RUNNING
                                     banda                RUNNING
sg_tsmsrv                            banda                ONLINE
sg_tsmsrv                            atlantic             OFFLINE
---------------------------------------------------------------------------
sg_tsmsrv                            banda                ONLINE
sg_tsmsrv                            atlantic             OFFLINE
---------------------------------------------------------------------------
                vg_tsmsrv            banda                ONLINE
                vg_tsmsrv            atlantic             OFFLINE
                ip_tsmsrv            banda                ONLINE
                ip_tsmsrv            atlantic             OFFLINE
---------------------------------------------------------------------------
```

Archived

```
                m_tsmsrv_db1         banda                ONLINE
                m_tsmsrv_db1         atlantic             OFFLINE
                m_tsmsrv_db1mr1      banda                ONLINE
                m_tsmsrv_db1mr1      atlantic             OFFLINE
                m_tsmsrv_lg1         banda                ONLINE
-----------------------------------------------------------------------
                m_tsmsrv_lg1         atlantic             OFFLINE
                m_tsmsrv_lgmr1       banda                ONLINE
                m_tsmsrv_lgmr1       atlantic             OFFLINE
                m_tsmsrv_dp1         banda                ONLINE
                m_tsmsrv_dp1         atlantic             OFFLINE
-----------------------------------------------------------------------
                m_tsmsrv_files       banda                ONLINE
                m_tsmsrv_files       atlantic             OFFLINE
                app_tsmsrv           banda                ONLINE
                app_tsmsrv           atlantic             OFFLINE
                NIC_en1              banda                ONLINE
-----------------------------------------------------------------------
                NIC_en1              atlantic             ONLINE
                vg_tsmsrv            banda                ONLINE
                vg_tsmsrv            atlantic             OFFLINE
                ip_tsmsrv            banda                ONLINE
                ip_tsmsrv            atlantic             OFFLINE
                m_tsmsrv_db1         banda                ONLINE
-----------------------------------------------------------------------
                m_tsmsrv_db1         atlantic             OFFLINE
                m_tsmsrv_db1mr1      banda                ONLINE
                m_tsmsrv_db1mr1      atlantic             OFFLINE
                m_tsmsrv_lg1         banda                ONLINE
                m_tsmsrv_lg1         atlantic             OFFLINE
-----------------------------------------------------------------------
                m_tsmsrv_lgmr1       banda                ONLINE
                m_tsmsrv_lgmr1       atlantic             OFFLINE
                m_tsmsrv_dp1         banda                ONLINE
                m_tsmsrv_dp1         atlantic             OFFLINE
                m_tsmsrv_files       banda                ONLINE
-----------------------------------------------------------------------
                m_tsmsrv_files       atlantic             OFFLINE
                app_tsmsrv           banda                ONLINE
                app_tsmsrv           atlantic             OFFLINE
                NIC_en1              banda                ONLINE
                NIC_en1              atlantic             ONLINE
```

2. Now, we bring the applications OFFLINE using the **hagrp -offline** command, as shown in Example 18-33.

Example 18-33 hagrp -offline command

```
Atlantic:/opt/local/tsmcli# hagrp -offline sg_tsmsrv -sys banda -localclus
```

3. Now, we review the **hastatus** output as shown in Example 18-34.

Example 18-34 hastatus output for the Service Group OFFLINE

```
banda:/var/VRTSvcs/log# hastatus
attempting to connect....connected

group           resource             system               message
--------------- -------------------- -------------------- --------------------
                                     atlantic             RUNNING
                                     banda                RUNNING
sg_tsmsrv                            banda                OFFLINE
sg_tsmsrv                            atlantic             OFFLINE
```

4. Then, we review the /var/VRTSvcs/log/engine_A.log, as shown in Example 18-35.

Example 18-35 tail -f /var/VRTSvcs/log/engine_A.log

```
2005/02/17 12:12:38 VCS NOTICE V-16-1-10446 Group sg_tsmsrv is offline on
system banda
```

18.5.5 Manual Service Group switch

Here are the steps to follow for this test:

1. For this test, all Service Groups are on one node (Banda), and will be switched to Atlantic, using the Cluster Manager GUI. As with all tests, we clear the engine_A.log using **cp /dev/null /var/VRTSvcs/log/engine_A.log**. The **hastatus | grep ONLINE** output prior to starting the transition is shown in Example 18-36.

Example 18-36 hastatus output prior to the Service Groups switching nodes

```
^banda:/var/VRTSvcs/log# Nhastatus | grep ONILIn
banda:/var/VRTSvcs/log# hastatus |grep ONLINE
attempting to connect....connected
sg_tsmsrv                            banda                ONLINE
sg_tsmsrv                            banda                ONLINE
                vg_tsmsrv            banda                ONLINE
                ip_tsmsrv            banda                ONLINE
                m_tsmsrv_db1         banda                ONLINE
                m_tsmsrv_db1mr1      banda                ONLINE
                m_tsmsrv_lg1         banda                ONLINE
                m_tsmsrv_lgmr1       banda                ONLINE
```

```
m_tsmsrv_dp1          banda                ONLINE
m_tsmsrv_files        banda                ONLINE
app_tsmsrv            banda                ONLINE
NIC_en1               banda                ONLINE
NIC_en1               atlantic             ONLINE
```

2. Now, we switch the Service Groups using the Cluster Manager GUI, as shown in Figure 18-14.

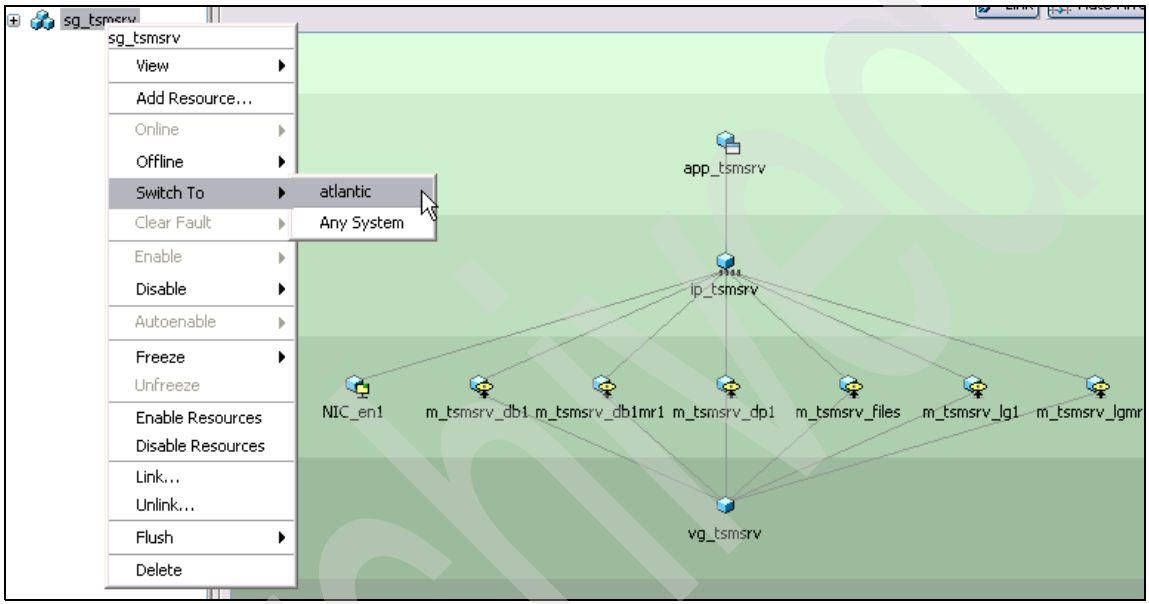

Figure 18-14 VCS Cluster Manager GUI switching Service Group to another node

3. Then, we click **Yes** to start the process as shown in Figure 18-15.

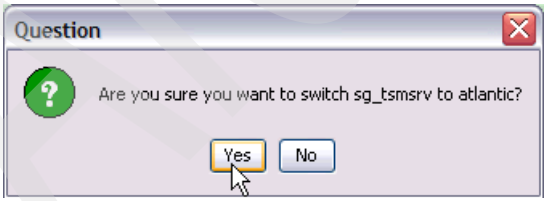

Figure 18-15 Prompt to confirm the switch

Tip: This process can be completed using the command line as well:

```
banda:/var/VRTSvcs/log# hagrp -switch sg_tsmsrv -to atlantic -localclus
```

4. Now, we monitor the transition which can be seen using the Cluster Manager GUI, and review the results in **hastatus** and the engine_A.log. The two logs are shown in Example 18-37 and Example 18-38.

Example 18-37 hastatus output of the Service Group switch

```
^banda:/var/VRTSvcs/log# hastatus |grep ONLINE
attempting to connect....connected
sg_tsmsrv                           atlantic            ONLINE
sg_tsmsrv                           atlantic            ONLINE
                vg_tsmsrv           atlantic            ONLINE
                ip_tsmsrv           atlantic            ONLINE
                m_tsmsrv_db1        atlantic            ONLINE
                m_tsmsrv_db1mr1     atlantic            ONLINE
                m_tsmsrv_lg1        atlantic            ONLINE
                m_tsmsrv_lgmr1      atlantic            ONLINE
                m_tsmsrv_dp1        atlantic            ONLINE
                m_tsmsrv_files      atlantic            ONLINE
                app_tsmsrv          atlantic            ONLINE
                NIC_en1             banda               ONLINE
                NIC_en1             atlantic            ONLINE
```

Example 18-38 tail -f /var/VRTSvcs/log/engine_A.log from surviving node, Atlantic

```
VCS INFO V-16-1-50135 User root fired command: hagrp -switch sg_tsmsrv atlantic
localclus from localhost
VCS NOTICE V-16-1-10208 Initiating switch of group sg_tsmsrv from system banda
to system atlantic
VCS NOTICE V-16-1-10300 Initiating Offline of Resource app_tsmsrv (Owner:
unknown, Group: sg_tsmsrv) on System banda
.
.
.
VCS NOTICE V-16-1-10447 Group sg_tsmsrv is online on system banda
VCS NOTICE V-16-1-10448 Group sg_tsmsrv failed over to system atlantic
```

Results

In this test, our Service Group has completed the switch and are now online on Atlantic. This completes the test successfully.

18.5.6 Manual fallback (switch back)

Here are the steps to follow for this test:

1. Before every test, we check the status for cluster services, resource groups, and resources on both nodes. In Example 18-39 we are verifying using **hastatus**.

Example 18-39 hastatus output of the current cluster state

```
banda:/# hastatus |grep ONLINE
attempting to connect....connected
sg_tsmsrv                              atlantic             ONLINE
sg_tsmsrv                              atlantic             ONLINE
                vg_tsmsrv              atlantic             ONLINE
                ip_tsmsrv              atlantic             ONLINE
                m_tsmsrv_db1           atlantic             ONLINE
                m_tsmsrv_db1mr1        atlantic             ONLINE
                m_tsmsrv_lg1           atlantic             ONLINE
                m_tsmsrv_lgmr1         atlantic             ONLINE
                m_tsmsrv_dp1           atlantic             ONLINE
                m_tsmsrv_files         atlantic             ONLINE
                app_tsmsrv             atlantic             ONLINE
                NIC_en1                banda                ONLINE
                NIC_en1                atlantic             ONLINE
```

2. For this test, we will use the AIX command line to switch the Service Group back to Banda, as shown in Example 18-40.

Example 18-40 hargrp -switch command to switch the Service Group back to Banda

```
banda:/# hagrp -switch sg_tsmsrv -to banda -localclus
```

3. We then review the results in the engine_A.log, as shown in Example 18-41.

Example 18-41 /var/VRTSvcs/log/engine_A.log segment for the switch back to Banda

```
VCS NOTICE V-16-1-10208 Initiating switch of group sg_tsmsrv from system
atlantic to system banda
VCS NOTICE V-16-1-10300 Initiating Offline of Resource app_tsmsrv (Owner:
unknown, Group: sg_tsmsrv) on System atlantic
.
.
.
VCS NOTICE V-16-1-10447 Group sg_tsmsrv is online on system banda
VCS NOTICE V-16-1-10448 Group sg_tsmsrv failed over to system banda
```

Results

Once we have the Service Group is back on Banda, this test is now complete.

18.5.7 Public NIC failure

Here, we are testing a failure situation on the public NIC.

Objective

We will now test the failure of a critical resource within the Service Group, the public NIC. First, we will test the reaction of the cluster when the NIC fails (physically disconnected), then document the cluster's recovery behavior once the NIC is plugged back in. We anticipate that the Service Group sg_tsmsrv will fault the NIC_en1 on Atlantic, then failover to Banda. Once sg_tsmsrv resources come online on Banda, we will replace the ethernet cable, which should produce a recovery of the resource, then we will manually switch sg_tsmsrv back to Atlantic.

Test sequence

Here are the steps to follow for this test:

1. For this test, one Service Group will be on each node, As with all tests, we clear the engine_A.log using **cp /dev/null /var/VRTSvcs/log/engine_A.log**.
2. Next, we physically disconnect the ethernet cable from the EN1 device on Atlantic. This is defined as a critical resource for the Service Group in which the Tivoli Storage Manager server is the Application. We will then observe the results in both logs being monitored.
3. Then we will review the engine_A.log file to understand the transition actions, which is shown in Example 18-42.

Example 18-42 /var/VRTSvcs/log/engine_A.log output for the failure activity

```
VCS INFO V-16-1-10077 Received new cluster membership
VCS NOTICE V-16-1-10080 System (banda) - Membership: 0x3, Jeopardy: 0x2
VCS ERROR V-16-1-10087 System banda (Node '1') is in Regardy Membership -
Membership: 0x3, Jeopardy: 0x2
.
.
.
VCS WARNING V-16-10011-5607 (atlantic) NIC:NIC_en1:monitor:Packet count test
failed: Resource is offline
VCS WARNING V-16-10011-5607 (atlantic) NIC:NIC_en1:monitor:Packet count test
failed: Resource is offline
VCS INFO V-16-1-10307 Resource NIC_en1 (Owner: unknown, Group: sg_tsmsrv) is
offline on atlantic (Not initiated by VCS)
VCS NOTICE V-16-1-10300 Initiating Offline of Resource app_t
tsmsrv (Owner: unknown, Group: sg_tsmsrv) on System atlantic
.
.
.
VCS INFO V-16-1-10298 Resource app_tsmsrv (Owner: unknown, Group: sg_tsmsrv) is
online on banda (VCS initiated)
VCS NOTICE V-16-1-10447 Group sg_tsmsrv is online on system banda
VCS NOTICE V-16-1-10448 Group sg_tsmsrv failed over to system banda
```

```
VCS WARNING V-16-10011-5607 (atlantic) NIC:NIC_en1:monitor:Packet count test
failed: Resource is offline
VCS WARNING V-16-10011-5607 (atlantic) NIC:NIC_en1:monitor:p
Packet count test failed: Resource is offline
.
.
.
VCS WARNING V-16-10011-5607 (atlantic) NIC:NIC_en1:monitor:packet count test
failed: Resource is offline
VCS WARNING V-16-10011-5607 (atlantic) NIC:NIC_en1:monitor:packet count test
failed: Resource is offline
```

4. As a result of the failed NIC, which is a critical resource for sg_tsmsrv the Service Group fails over to Banda (from Atlantic).
5. Next, we plug the ethernet cable back into the NIC and monitor for a state change, and now the cluster ONLINE resources show that EN1 on Atlantic is back ONLINE, however there is no failback (resources are stable on Banda) and the cluster knows it is now capable of failing over to Atlantic for both NICs if required. The **hastatus** of the NIC1 transition is shown in Example 18-43.

Example 18-43 hastatus of the ONLINE resources

```
# hastatus |grep ONLINE
attempting to connect....connected
sg_tsmsrv                          banda               ONLINE
sg_tsmsrv                          banda               ONLINE
               vg_tsmsrv           banda               ONLINE
               ip_tsmsrv           banda               ONLINE
               m_tsmsrv_db1        banda               ONLINE
               m_tsmsrv_db1mr1     banda               ONLINE
               m_tsmsrv_lg1        banda               ONLINE
               m_tsmsrv_lgmr1      banda               ONLINE
               m_tsmsrv_dp1        banda               ONLINE
               m_tsmsrv_files      banda               ONLINE
               app_tsmsrv          banda               ONLINE
               NIC_en1             banda               ONLINE
               NIC_en1             atlantic            ONLINE
```

6. Then, we review the contents of the engine_A.log, which is shown in Example 18-44.

Example 18-44 /var/VRTSvcs/log/engine_A.log output for the recovery activity

```
VCS INFO V-16-1-10077 Received new cluster membership
VCS NOTICE V-16-1-10080 System (banda) - Membership: 0x3, Jeopardy: 0x0
VCS NOTICE V-16-1-10086 System banda (Node '1') is in Regular Membership -
Membership: 0x3
```

```
VCS INFO V-16-1-10299 Resource NIC_en1 (Owner: unknown, Group: sg_tsmsrv) is
online on atlantic (Not initiated by VCS)
```

7. At this point we manually switch the sg_tsmsrv back over to Atlantic, with the ONLINE resources shown in **hastatus** in Example 18-45, which then concludes this test.

Example 18-45 hastatus of the online resources fully recovered from the failure test

```
hastatus |grep ONLINE
attempting to connect....connected
sg_tsmsrv                           atlantic            ONLINE
sg_tsmsrv                           atlantic            ONLINE
               vg_tsmsrv            atlantic            ONLINE
               ip_tsmsrv            atlantic            ONLINE
               m_tsmsrv_db1         atlantic            ONLINE
               m_tsmsrv_db1mr1      atlantic            ONLINE
               m_tsmsrv_lg1         atlantic            ONLINE
               m_tsmsrv_lgmr1       atlantic            ONLINE
               m_tsmsrv_dp1         atlantic            ONLINE
               m_tsmsrv_files       atlantic            ONLINE
               app_tsmsrv           atlantic            ONLINE
               NIC_en1              banda               ONLINE
               NIC_en1              atlantic            ONLINE
```

18.5.8 Failure of the server during a client backup

We will be testing the Tivoli Storage Manager server during a client backup.

Objective

In this test we are verifying client operation which originates from Azov, survives a server failure on Atlantic, and the subsequent takeover by the node Banda.

Preparation

Here are the steps to follow:

1. We verify that the cluster services are running with the **hastatus | grep ONLINE** command. We see that the sg_tsmsrv Service Group is currently on Atlantic, shown in Example 18-46.

Example 18-46 hastatus | grep ONLINE output

```
hastatus |grep ONLINE
attempting to connect....connected
sg_tsmsrv                           atlantic            ONLINE
sg_tsmsrv                           atlantic            ONLINE
               vg_tsmsrv            atlantic            ONLINE
```


Archived

```
ip_tsmsrv            atlantic             ONLINE
m_tsmsrv_db1         atlantic             ONLINE
m_tsmsrv_db1mr1      atlantic             ONLINE
m_tsmsrv_lg1         atlantic             ONLINE
m_tsmsrv_lgmr1       atlantic             ONLINE
m_tsmsrv_dp1         atlantic             ONLINE
m_tsmsrv_files       atlantic             ONLINE
app_tsmsrv           atlantic             ONLINE
NIC_en1              banda                ONLINE
NIC_en1              atlantic             ONLINE
```

2. On Banda, we us the **AIX command tail -f /var/VRTSvcs/log/engine_A.log** to monitor cluster operation.
3. Then we start a client incremental backup with the command line and see metadata and data sessions starting on Atlantic (Tivoli Storage Manager server), sessions 37 and 38, as shown in Example 18-47.

Example 18-47 Client sessions starting

```
Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- --------------------

    36 Tcp/Ip Run      0 S    3.0 K     201 Admin AIX      ADMIN
    37 Tcp/Ip IdleW    0 S    1.2 K     670 Node  AIX      AZOV
    38 Tcp/Ip Run      0 S      393  17.0 M Node  AIX      AZOV
```

4. On the server, we verify that data is being transferred via the **query session** command, noticing session 38, which is now sending data, as shown in Example 18-47.

Failure

Here are the steps to follow for this test:

1. To ensure that the client backup is running, we issue **halt -q** on the AIX server running the Tivoli Storage Manager server Atlantic, then issue the **halt -q** command, which stops the AIX system immediately and powers off the system.
2. The client stops sending data to server and keeps retrying (Example 18-48).

Example 18-48 client stops sending data

```
ANS1809W Session is lost; initializing session reopen procedure.
A Reconnection attempt will be made in 00:00:12
```

3. From the cluster point of view, we view the contents of the engine_A.log, as shown in Example 18-49.

Example 18-49 Cluster log demonstrating the change of cluster membership status

```
VCS INFO V-16-1-10077 Received new cluster membership
VCS NOTICE V-16-1-10080 System (banda) - Membership: 0x2, Jeopardy: 0x0
VCS ERROR V-16-1-10079 System atlantic (Node '0') is in Down State -
Membership: 0x2
VCS ERROR V-16-1-10322 System atlantic (Node '0') changed state from RUNNING to
FAULTED
VCS NOTICE V-16-1-10446 Group sg_tsmsrv is offline on system atlantic
VCS INFO V-16-1-10493 Evaluating banda as potential target node for group
sg_tsmsrv
VCS INFO V-16-1-10493 Evaluating atlantic as potential target node for group
sg_tsmsrv
VCS INFO V-16-1-10494 System atlantic not in RUNNING state
VCS NOTICE V-16-1-10301 Initiating Online of Resource vg_tsmsrv (Owner:
unknown, Group: sg_tsmsrv) on System banda
```

Recovery

The failover from Atlantic to Banda happens in approximately 5 minutes, of which most of the failover time is managing volumes that are marked DIRTY, and must be **fcsk**'d by VCS. We show the details of the engine_A.log for the ONLINE process and the completion in Example 18-50.

Example 18-50 engine_A.log online process and completion summary

```
VCS INFO V-16-2-13001 (banda) Resource(m_tsmsrv_files): Output of the completed
operation (online)
Replaying log for /dev/tsmlv.
mount: /dev/tsmlv on /tsm/files: Unformatted or incompatible media
The superblock on /dev/tsmlv is dirty.  Run a full fsck to fix.
/dev/tsmlv:   438500
mount: /dev/tsmlv on /tsm/files: Device busy

****************
The current volume is: /dev/tsmlv
locklog: failed on open, tmpfd=-1, errno:26
**Phase 1 - Check Blocks, Files/Directories, and Directory Entries
**Phase 2 - Count links
**Phase 3 - Duplicate Block Rescan and Directory Connectedness
**Phase 4 - Report Problems
**Phase 5 - Check Connectivity
**Phase 7 - Verify File/Directory Allocation Maps
**Phase 8 - Verify Disk Allocation Maps
    32768 kilobytes total disk space.
        1 kilobytes in 2 directories.
```

```
       36 kilobytes in 8 user files.
    32396 kilobytes are available for use.
File system is clean.
.
.
.
VCS INFO V-16-1-10298 Resource app_tsmsrv (Owner: unknown, Group: sg_tsmsrv) is
online on banda (VCS initiated)
VCS NOTICE V-16-1-10447 Group sg_tsmsrv is online on system
banda
```

Once the server is restarted, and the Tivoli Storage Manager server and client re-establish the sessions, the data flow begins again, as seen in Example 18-51 and Example 18-52.

Example 18-51 The restarted Tivoli Storage Manager accept client rejoin.

```
ANR8441E Initialization failed for SCSI library LIBLTO.
ANR2803I License manager started.
ANR8200I TCP/IP driver ready for connection with clients
                      on port 1500.
ANR2560I Schedule manager started.
ANR0993I Server initialization complete.
ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli is
                      now ready for use.
ANR2828I Server is licensed to support Tivoli Storage
                      Manager Basic Edition.
ANR2828I Server is licensed to support Tivoli Storage
                      Manager Extended Edition.
ANR1305I Disk volume /tsm/dp1/bckvol1 varied online.
ANR0406I Session 1 started for node AZOV (AIX) (Tcp/Ip
                      9.1.39.74(33513)). (SESSION: 1)
ANR0406I Session 2 started for node AZOV (AIX) (Tcp/Ip
                      9.1.39.74(33515)). (SESSION: 2)
```

Example 18-52 The client reconnect and continue operations

```
Directory-->              4,096 /usr/lpp/X11/Xamples/programs/xmag [Sent]
Directory-->              4,096 /usr/lpp/X11/Xamples/programs/xman [Sent]
Directory-->              4,096 /usr/lpp/X11/Xamples/programs/xmh [Sent]
Directory-->                256 /usr/lpp/X11/Xamples/programs/xprop [Sent]
Directory-->              256 /usr/lpp/X11/Xamples/programs/xrefresh [Sent]
Directory-->              4,096 /usr/lpp/X11/Xamples/programs/xsm [Sent]
Directory-->              256 /usr/lpp/X11/Xamples/programs/xstdcmap [Sent]
Directory-->                256 /usr/lpp/X11/Xamples/programs/xterm [Sent]
Directory-->                  256 /usr/lpp/X11/Xamples/programs/xwininfo [Sent]
```

Results

Due to the nature of this failure methodology (crashing the server during writes), this recovery example would be considered a real test. This test was successful.

> **Attention:** It is important to emphasize that these tests are only appropriate using test data, and should only be performed after the completion of a FULL Tivoli Storage Manager database backup.

18.5.9 Failure of the server during a client scheduled backup

We repeat the same test using a scheduled backup operation. The results are essentially the same (no **fcsk** was required) and the event for the schedule showed an exception of RC=12, however the backup completed entirely. We verified in both the server and client logs that the backup completed successfully.

In both cases the VCS cluster is able to manage the server failure and make the sg_tsmsrv Service Group available to client in about 1 minute (unless disk **fsck**'s are required) and the client is able to continue its operations successfully to the end.

18.5.10 Failure during disk to tape migration operation

We will be testing the Tivoli Storage Manager server while it is performing disk to tape migration.

Objectives

Here we test the recovery of a failure during a disk to tape migration operation and we will verify that the operation continues.

Preparation

Here are the steps to follow for this test:

1. We verify that the cluster services are running with the **hastatus** command.
2. On Banda, we clean the engine log with the command **cp /dev/null /var/VRTSvcs/log/engine_A.log**
3. On Banda we use **tail -f /var/VRTSvcs/log/engine_A.log** to monitor cluster operation.
4. We have a disk storage pool, having a tape storage pool as next. The disk storage pool is currently at 34% utilized.
5. Lowering the **highMig** threshold to zero, we start the migration to tape.

6. We wait for a tape cartridge mount, monitor using the Tivoli Storage Manager command **q mount** and **q proc** commands. These commands, and the output for them, are shown in Example 18-53.

Example 18-53 Command query mount and process

```
tsm: TSMSRV04>q mount
ANR8330I LTO volume ABA990 is mounted R/W in drive DRLTO_1 (/dev/rmt2), status:
IN USE.

tsm: TSMSRV04>q proc

 Process Process Description  Status
 Number
-------- -------------------- ------------------------------------------------
       1 Migration            Disk Storage Pool SPD_BCK, Moved Files: 6676,
                               Moved Bytes: 203,939,840, Unreadable Files: 0,
                               Unreadable Bytes: 0. Current Physical File
                               (bytes): 25,788,416 Current output volume:
                               ABA990.
```

7. Next the Tivoli Storage Manager actlog shows the following entry for this mount (Example 18-54).

Example 18-54 Actlog output showing the mount of volume ABA990

```
ANR1340I Scratch volume ABA990 is now defined in storage
                      pool SPT_BCK. (PROCESS: 1)
ANR0513I Process 1 opened output volume ABA990. (PROCESS: 1)
```

8. Then after a few minutes of data transfer we crash the Tivoli Storage Manager server.

Failure

We use the **halt -q** command to stop AIX immediately and power off the server.

Recovery

Banda now takes over the resources. As we have seen before in this testing chapter, the superblock is marked DIRTY on the shared drives, and VCS does an **fsck** to reset the bit and mount all the required disk resources.

The Service Group which contains the Tivoli Storage Manager server Applications is then restarted.

Once the server is restarted, the migration restarts because of the used percentage still above the **highMig** percentage (which is still currently zero).

As we have experienced with the testing on our other cluster platforms, this process completes successfully. The Tivoli Storage Manager actlog summary shows the completed lines for this operation in Example 18-55.

Example 18-55 Actlog output demonstrating the completion of the migration

```
ANR0515I Process 1 closed volume ABA990. (PROCESS: 1)
ANR0513I Process 1 opened output volume ABA990. (PROCESS: 1)
ANR1001I Migration process 1 ended for storage pool SPD_BCK. (PROCESS: 1)
ANR0986I Process 1 for MIGRATION running in the BACKGROUND
                      processed 11201 items for a total of 561,721,344 bytes
                     with a completion state of SUCCESS at 16:39:17(PROCESS:1)
```

Finally, we return the cluster configuration back to where we started, with the sg_tsmsrv hosted on Atlantic, and this test has completed.

Result summary

The actual recovery time from the halt to the process continuing was approximately 10 minutes. Again, this time will vary depending on the activity on the Tivoli Storage Manager server at the time of failure, as devices must be cleaned (**fsck** of disks), reset (tapes), and potentially media unmounted and then mounted again as the process starts up.

In the case of Tivoli Storage Manager migration, this was restarted due to the **highMig** value still being set lower than the current utilization of the storage pool.

The tape volume which was in use for the migration remained in a read/write state after the recovery, and was the volume re-mounted and reused to complete the process.

18.5.11 Failure during backup storage pool operation

Here we describe how to handle failure during backup storage pool operation.

Objectives

Here we test the recovery of a failure situation, in which the Tivoli Storage Manager server is currently performing a tape storage pool backup operation. We will confirm that we are able to restart the process without special intervention, after the Tivoli Storage Manager server recovers. We do not expect the operation to restart, as this is a command initiated process (unlike the migration or expiration processes).

Preparation

Here are the steps to follow for this test:

1. We verify that the cluster services are running with the **hastatus** command.
2. On the secondary node (the node which the sg_tsmsrv will failover to), we use **tail -f /var/VRTSvcs/log/engine_A.log** to monitor cluster operation.
3. We have a primary sequential storage pool called SPT_BCK containing an amount of backup data and a copy storage pool called SPC_BCK.
4. **Backup stg SPT_BCK SPC_BCK** command is issued.
5. We wait for a tape cartridges mount using the Tivoli Storage Manager commands **q mount**, as shown in Example 18-56.

Example 18-56 q mount output

```
tsm: TSMSRV04>q mount
ANR8379I Mount point in device class CLLTO1 is waiting for the volume mount to
complete, status: WAITING FOR VOLUME.
ANR8330I LTO volume ABA990 is mounted R/W in drive DRLTO_2 (/dev/rmt3), status:
IN USE.
ANR8334I         2 matches found.
```

6. Then we check for data being transferred from disk to tape using the **query process** command, as shown in Example 18-57.

Example 18-57 q process output

```
tsm: TSMSRV04>q proc

 Process Process Description  Status
 Number
-------- ------------------- ---------------------------------------------
       3 Backup Storage Pool  Primary Pool SPT_BCK, Copy Pool SPC_BCK, Files
                               Backed Up: 3565, Bytes Backed Up: 143,973,320,
                               Unreadable Files: 0, Unreadable Bytes: 0.
                               Current Physical File (bytes): 7,808,841 Current
                               input volume: ABA927. Current output volume:
                               ABA990.
```

7. Once data transfer is confirmed we fail the server banda.

Failure

We use the **halt -q** command to stop immediately AIX and power off the server.

Archived

Recovery

The cluster node atlantic takes over the Service Group, which we can see using **hastatus**, as shown in Example 18-58.

Example 18-58 VCS hastatus command output after the failover

```
atlantic:/var/VRTSvcs/log# hastatus
attempting to connect....connected

group           resource             system               message
--------------- -------------------- -------------------- --------------------
                                     atlantic             RUNNING
                                     banda                *FAULTED*
sg_tsmsrv                            atlantic             ONLINE
                vg_tsmsrv            banda                OFFLINE
                vg_tsmsrv            atlantic             ONLINE
                ip_tsmsrv            banda                OFFLINE
                ip_tsmsrv            atlantic             ONLINE
                m_tsmsrv_db1         banda                OFFLINE
-----------------------------------------------------------------------------
                m_tsmsrv_db1         atlantic             ONLINE
                m_tsmsrv_db1mr1      banda                OFFLINE
                m_tsmsrv_db1mr1      atlantic             ONLINE
                m_tsmsrv_lg1         banda                OFFLINE
                m_tsmsrv_lg1         atlantic             ONLINE
-----------------------------------------------------------------------------
                m_tsmsrv_lgmr1       banda                OFFLINE
                m_tsmsrv_lgmr1       atlantic             ONLINE
                m_tsmsrv_dp1         banda                OFFLINE
                m_tsmsrv_dp1         atlantic             ONLINE
                m_tsmsrv_files       banda                OFFLINE
-----------------------------------------------------------------------------
                m_tsmsrv_files       atlantic             ONLINE
                app_tsmsrv           banda                OFFLINE
                app_tsmsrv           atlantic             ONLINE
                NIC_en1              banda                ONLINE
                NIC_en1              atlantic             ONLINE
```

The Tivoli Storage Manager server is restarted on Atlantic, and after monitoring and reviewing the process status, there are no storage pool backups which restart.

At this point, we then restart the backup storage pool by re-issuing the command **Backup stg SPT_BCK SPC_BCK**.

Archived

Example 18-59 q process after the backup storage pool command has restarted

```
tsm: TSMSRV04>q proc

 Process Process Description  Status
 Number
-------- -------------------- ------------------------------------------------
       1 Backup Storage Pool  Primary Pool SPT_BCK, Copy Pool SPC_BCK, Files
                               Backed Up: 81812, Bytes Backed Up:
                               4,236,390,075, Unreadable Files: 0, Unreadable
                               Bytes: 0. Current Physical File (bytes):
                               26,287,875 Current input volume: ABA927. Current
                               output volume: ABA990.
```

8. Then, we review the process with data flow, as shown in Example 18-59. In addition, we also observe that the same tape volume is mounted and used as before, using q mount, as shown in Example 18-60.

Example 18-60 q mount after the takeover and restart of Tivoli Storage Manager

```
tsm: TSMSRV04>q mount
ANR8330I LTO volume ABA927 is mounted R/W in drive DRLTO_2 (/dev/rmt3), status:
IN USE.
ANR8330I LTO volume ABA990 is mounted R/W in drive DRLTO_1 (/dev/rmt2), status:
IN USE.
ANR8334I         2 matches found.
```

This process continues until completion, and terminates successfully. We then return the cluster to the starting position by during a manual switch of the Service Group.“Manual fallback (switch back)” on page 777

Results

In this case the cluster is failed over, and Tivoli Storage Manager is back operating in 4 minutes (approximately). This slightly extended time was due to having two tapes in use which had to be unmounted during the reset operation, then remounted once the command was re-issued.

Backup storage pool process has to be restarted, and completed with a consistent state.

The Tivoli Storage Manager database survives the failure with all volumes synchronized (even when **fsck** filesystem checks are required).

The tape volumes involved in failure have remained in a read/write state and reused.

If administration scripts are used for scheduling and rescheduling activities, it is possible that this process will restart after the failover has completed.

18.5.12 Failure during database backup operation

This section describes how to handle a failure situation during database backup.

Objectives

Now we test the recovery of a Tivoli Storage Manager server node failure, while performing a full database backup. Regardless of the outcome, we would not consider the volume credible for disaster recovery (limit your risk by re-doing the operation if there is a failure during a full Tivoli Storage Manager database backup).

Preparation

Here are the steps to follow for this test:

1. We verify that the cluster services are running with the **hastatus** command on Atlantic.
2. Then, on the node Banda (which the sg_tsmsrv will failover to), we use **tail -f /var/VRTSvcs/log/engine_A.log** to monitor cluster operation.
3. We issue a **backup db type=full devc=lto1**.
4. Then we wait for a tape mount and for the first *ANR4554I* message.

Failure

We use the **halt -q** command to stop immediately AIX and power off the server.

Recovery

The sequence of events for the recovery of this failure is as follows:

1. The node Banda takes over the resources.
2. The tape is unloaded by reset issued during cluster takeover operations.
3. The Tivoli Storage Manager server is restarted.
4. Then we check the state of database backup in execution at halt time with **q vol** and **q libv** commands.
5. We see that volume state has been reserved for database backup, but the operation is not finished.
6. We used **BACKUP DB t=f devc=lto1** to start a new database backup process.
7. The new process skips the previous volume, takes a new one, and completes.

8. Then we have to return the failed DBB volume to the scratch pool, using the command **upd libv LIBLTO <volid> status=scr**.
9. At the end of testing, we return the cluster operation back to Atlantic.

Result summary

In this situation the cluster is able to manage the server failure and make Tivoli Storage Manager available in a short period of time.

The database backup has to be restarted.

The tape volume used in the database backup process running at failure time has remained in a non-scratch status, to which has to be returned using an **update libv** command.

Anytime there is a failover of a Tivoli Storage Manager server environment, it is essential to understand what processes were in progress, and validate the successful completion. In the case of a full database backup being interrupted, the task is to clean up by removing the backup which was started prior to the failover, and ensuring that another backup completes after the failover.

Archived

19

VERITAS Cluster Server on AIX with the IBM Tivoli Storage Manager StorageAgent

This chapter describes our installation, configuration, and testing related to the Tivoli Storage Manager Storage Agent, and its configuration as a highly available Veritas Cluster Server application.

© Copyright IBM Corp. 2005. All rights reserved.

19.1 Overview

We will configure the Tivoli Storage Manager client and server so that the client, through a Storage Agent, can move its data directly to storage on a SAN. This function, called LAN-free data movement, is provided by IBM Tivoli Storage Manager for Storage Area Networks.

As part of the configuration, a Storage Agent is installed on the client system. Tivoli Storage Manager supports both tape libraries and FILE libraries. This feature supports SCSI, 349X, and ACSLS tape libraries.

For more information on configuring Tivoli Storage Manager for LAN-free data movement, see the *IBM Tivoli Storage Manager Storage Agent User's Guide*.

The configuration procedure we follow will depend on the type of environment we want to implement, which in this testing environment will be a highly available Storage Agent only. We will not configure the local Storage Agents. There is rarely a need for a locally configured Storage Agent within a cluster, as the application data will reside as part of the clustered shared disks, which our Tivoli Storage Manager client and Storage Agent must move with. This is the same reason that the application, Tivoli Storage Manager client, and Storage Agents are configured within the same VCS Service Group, as separate applications.

Tape drives SCSI reserve concern

When a server running Tivoli Storage Manager server or Storage Agent crashes while using a tape drive, its SCSI reserve remains, preventing other servers accessing the tape resources.

A new library parameter called *resetdrives*, which specifies whether the server performs a target reset when the server is restarted or when a library client or Storage Agent re-connection is established, has been made available in AIX Tivoli Storage Manager server for AIX from release 5.3. This parameter only applies to SCSI, 3494, Manual, and ACSLS type libraries.

An external SCSI reset is still needed to free up tape resources in case the server fails if the library server is other than V5.3 or later running on AIX.

For setting up Tivoli Storage Manager Storage Agents with a library server running on platforms different from AIX, we adapted a sample script, provided for starting the server in previous versions, and also for the startup, the Storage Agent within a cluster.

We can't have the cluster software doing this, using tape resource management, because it will reset all of the drives, even if in use from the server or other Storage Agents.

Why cluster a Storage Agent?

In a clustered client environment, Storage Agents can be local or a cluster resource, for both backup/archive and API clients. They can be accessed, using shared memory communication, with a specific port number or TCP/IP communication with loopback address and specific port number, or accessed using highly available TCP/IP addresses.

The advantage of clustering a Storage Agent, in a client failover scenario, is to have Tivoli Storage Manager reacting immediately when Storage Agent restarts.

When a Storage Agent restarts, Tivoli Storage Manager server checks for the resources previously allocated to that Storage Agent, then issues SCSI resets if needed. Otherwise Tivoli Storage Manager reacts on a time-out only basis to Storage Agent failures.

19.2 Planning and design

For this implementation, we will be testing the configuration and clustering for one Tivoli Storage Manager Storage Agent instance and demonstrating the possibility of restarting a LAN-free backup just after the takeover of a failed cluster node.

Our design considers a two-node cluster, with one virtual (clustered) Storage Agent to be used by a clustered application which relies on a clustered client for backup and restore, as described in Table 19-1.

Table 19-1 Storage Agent configuration for our design

STA instance	Instance path	TCP/IP address	TCP/IP port
cl_veritas01_sta	/opt/IBM/ISC/tsm/Storageagent/bin	9.1.39.77	1502

We install the Storage Agent on both nodes in the local filesystem to ensure it is referenced locally in each node, within AIX ODM. Then we copy the configuration files into the shared disk structure.

Here we are using TCP/IP as communication method, but shared memory also applies only if the Storage Agent and the Tivoli Storage Manager server remain on the same physical node.

Archived

A complete environment configuration is shown in Table 19-2, Table 19-3, and Table 19-4.

Table 19-2 .LAN-free configuration of our lab

Node 1	
TSM nodename	
dsm.opt location	
Storage Agent name	
dsmsta.opt and devconfig.txt location	
Storage Agent high level address	
Storage Agent low level address	
LAN-free communication method	
Node 2	
TSM nodename	
dsm.opt location	
Storage Agent name	
dsmsta.opt and devconfig.txt location	
Storage Agent high level address	
Storage Agent low level address	
LAN-free communication method	
Virtual node	
TSM nodename	cl_veritas01_client
dsm.opt location	/opt/IBM/ISC/tsm/client/ba/bin
Storage Agent name	cl_veritas01_sta
dsmsta.opt and devconfig.txt location	/opt/IBM/ISC/tsm/Storageagent/bin
Storage Agent high level address	9.1.39.77
Storage Agent low level address	1502
LAN-free communication method	Tcpip

Archived

Table 19-3 Server information

Server information	
Servername	TSMSRV03
High level address	9.1.39.74
Low level address	1500
Server password for server-to-server communication	password

Our Storage Area Network devices are listed inTable 19-4.

Table 19-4 Storage Area Network devices

SAN devices	
Disk	IBM DS4500 Disk Storage Subsystem
Library	IBM LTO 3583 Tape Library
Tape drives	3580 Ultrium 1
Tape drive device name	drlto_1: /dev/rmt2 drlto_2: /dev/rmt3

19.3 Lab setup

We use the lab already set up for clustered client testing in 17.4, “Lab environment” on page 721.

Once the installation and configuration of Tivoli Storage Manager Storage Agent has finished, we need to modify the existing client configuration to make it use the LAN-free backup.

19.4 Tivoli Storage Manager Storage Agent installation

We will install the AIX Storage Agent V5.3 for LAN-free backup services on both nodes of the VCS cluster. This installation will be a standard installation, following the Storage Agent User’s Guide, which can be located online at:

http://publib.boulder.ibm.com/infocenter/tivihelp/index.jsp?topic=/com.ibm.i

Archived

At this point, our team has already installed the Tivoli Storage Manager server and Tivoli Storage Manager client, which will have been configured for high availability. We have also configured and verified the communication paths between the client and server.

After reviewing the readme file and the User's Guide, we then proceed to fill out the Configuration Information Worksheet provided in Table 19-2 on page 796.

Using the AIX command **smitty installp**, we install the filesets for the Tivoli Storage Manager Storage Agent. This installation is standard, with the agent being installed on both clients in the default locations.

19.5 Storage agent configuration

We begin our configuration of the LAN-free client by registering our Storage Agent on TSMSRV03, then set up our definitions locally, and lastly, we configure our drive paths on the remote Tivoli Storage Manager server.

Locally, we have already defined the VCS Service Group sg_isc_sta_tsmcli, which hosts the shared disk resource. We will activate the shared disk to facilitate our setup of the Storage Agent configuration files as follows:

1. First we register our Storage Agent server with the Tivoli Storage Manager server we will be connecting to, in this case TSMSRV03.
2. Next, we run the **/usr/tivoli/tsm/StorageAgent/bin/dsmsta setstorageserver command** to populate the devconfig.txt and dsmsta.opt files, as shown in Example 19-1.

Example 19-1 The dsmsta setstorageserver command

```
dsmsta setstorageserver myname=cl_veritas01_sta mypassword=password
myhladdress=9.1.39.77 servername=tsmsrv03 serverpassword=password
hladdress=9.1.39.74 lladdress=1500
```

3. We then review the results of running this command, which populates the devconfig.txt file as shown in Example 19-2.

Example 19-2 The devconfig.txt file

```
SET STANAME CL_VERITAS01_STA
SET STAPASSWORD 2128bafb1915d7ee7cc49f9e116493280c
SET STAHLADDRESS 9.1.39.77
DEFINE SERVER TSMSRV03 HLADDRESS=9.1.39.74 LLADDRESS=1500
SERVERPA=21911a57cfe832900b9c6f258aa0926124
```

4. Next, we review the results of this update on the dsmsta.opt file. We also see the configurable parameters we have included, as well as the last line added by the update just completed, which adds the servername, as shown in Example 19-3.

Example 19-3 dsmsta.opt file change results

```
SANDISCOVERY ON
COMMmethod TCPIP
TCPPort 1502
DEVCONFIG /opt/IBM/ISC/tsm/StorageAgent/bin/devconfig.txt
SERVERNAME TSMSRV03
```

5. Then, we add a two stanzas to our /usr/tivoli/tsm/client/ba/bin/dsm.sys file for the LAN-free connection and a direct connection to the Storage Agent (for use with the **dsmadmc** command), as shown in Example 19-4.

Example 19-4 dsm.sys stanzas for Storage Agent configured as highly available

```
* StorageAgent Server stanza for admin connection purpose
SErvername              cl_veritas01_sta
   COMMMethod           TCPip
   TCPPort              1502
   TCPServeraddress     9.1.39.77
   ERRORLOGRETENTION    7
   ERRORLOGname         /usr/tivoli/tsm/client/ba/bin/dsmerror.log

*******************************************************************
*    Clustered Storage Agents Labs Stanzas                        *
*******************************************************************
* Server stanza for the LAN-free atlantic client to the tsmsrv03 (AIX)
* this will be a client which uses the LAN-free StorageAgent
SErvername                     tsmsrv03_san
   nodename                    cl_veritas01_client
   COMMMethod                  TCPip
   TCPPort                     1500
   TCPClientaddress            9.1.39.77
   TCPServeraddr               9.1.39.74

   TXNBytelimit                256000
   resourceutilization         5
   enablelanfree               yes
   lanfreecommmethod           tcpip
   lanfreetcpport              1502
   lanfreetcpserveraddress     9.1.39.77
```

Archived

```
schedmode                    prompt
passwordaccess               generate
passworddir                  /opt/IBM/ISC/tsm/client/ba/bin/atlantic
schedlogname                 /opt/IBM/ISC/tsm/client/ba/bin/dsmsched.log
errorlogname                 /opt/IBM/ISC/tsm/client/ba/bin/dsmerror.log
ERRORLOGRETENTION            7
```

6. Now we configure our LAN-free tape paths by using the ISC administration interface, connecting to TSMSRV03. We start the ISC, then select Tivoli Storage Manager, then Storage Devices, then the library associated to the server TSMSRV03.
7. We choose **Drive Paths**, as seen in Figure 19-1.

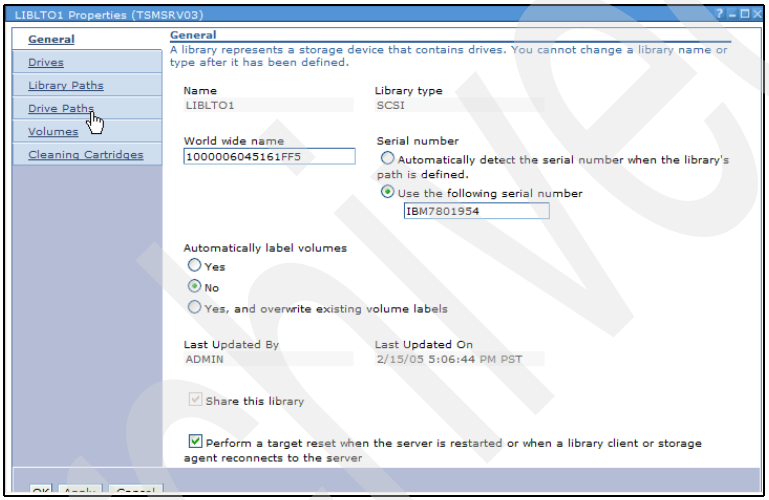

Figure 19-1 Administration Center screen to select drive paths

Archived

8. We select **Add Path** and click **Go**, as seen in Figure 19-2.

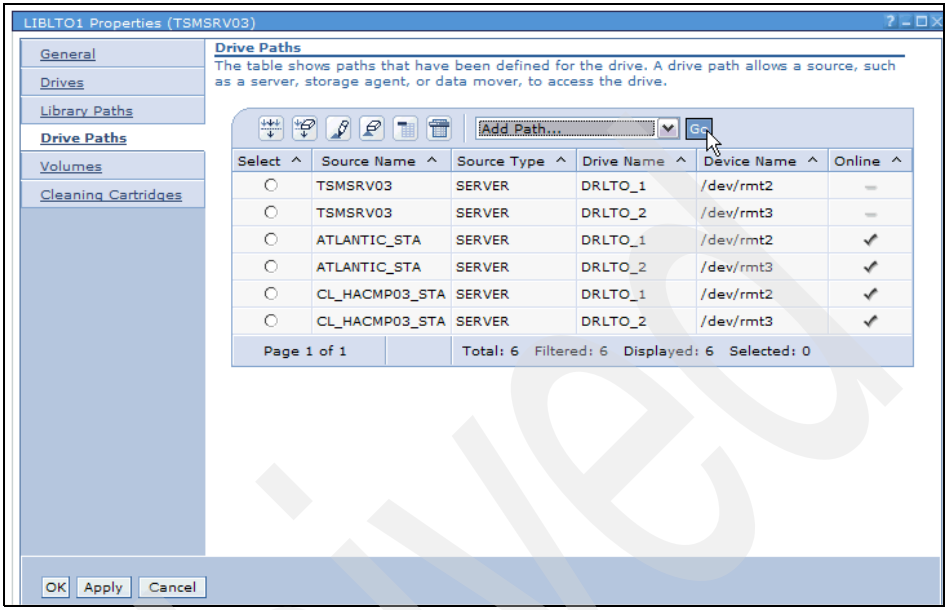

Figure 19-2 Administration Center screen to add a drive path

9. Then, we fill out the next panel with the local special device name, and select the corresponding device which has been defined on TSMSRV03, as seen in Figure 19-3.

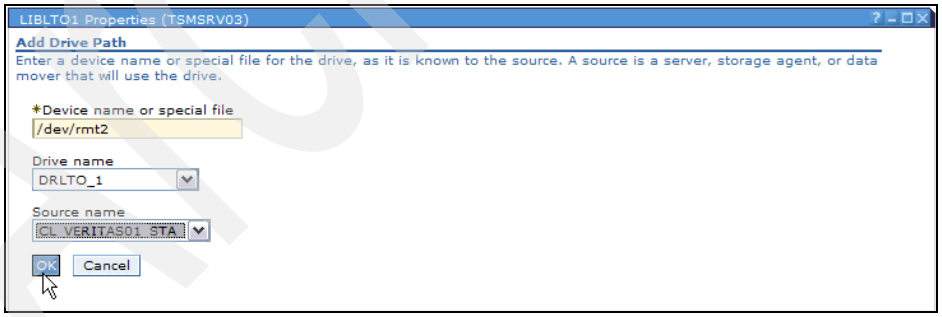

Figure 19-3 Administration Center screen to define DRLTO_1

Archived

10. For the next panel, we click **Close Message**, as seen in Figure 19-4.

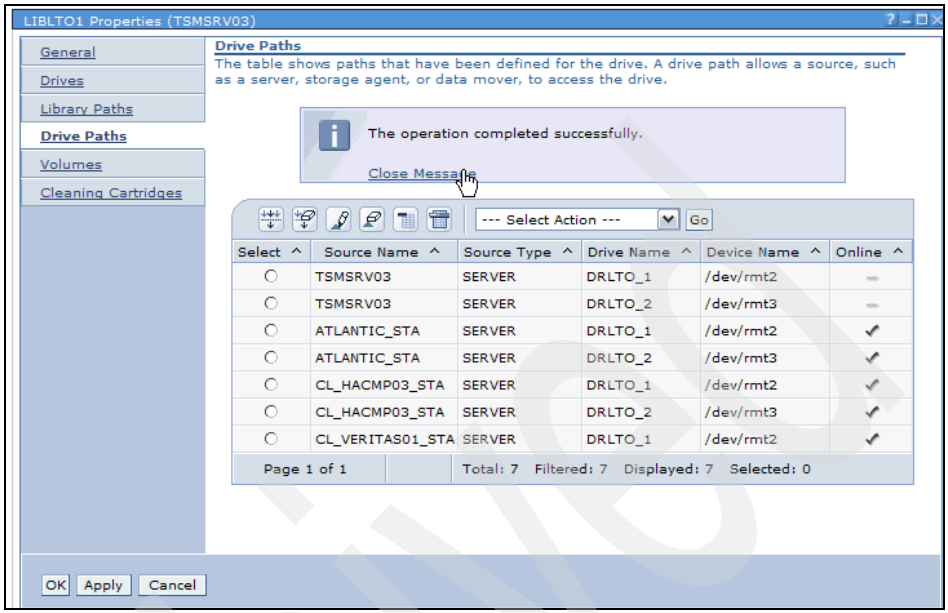

Figure 19-4 Administration Center screen to review completed adding drive path

11. We then select add drive path to add the second drive, as shown in Figure 19-5.

Archived

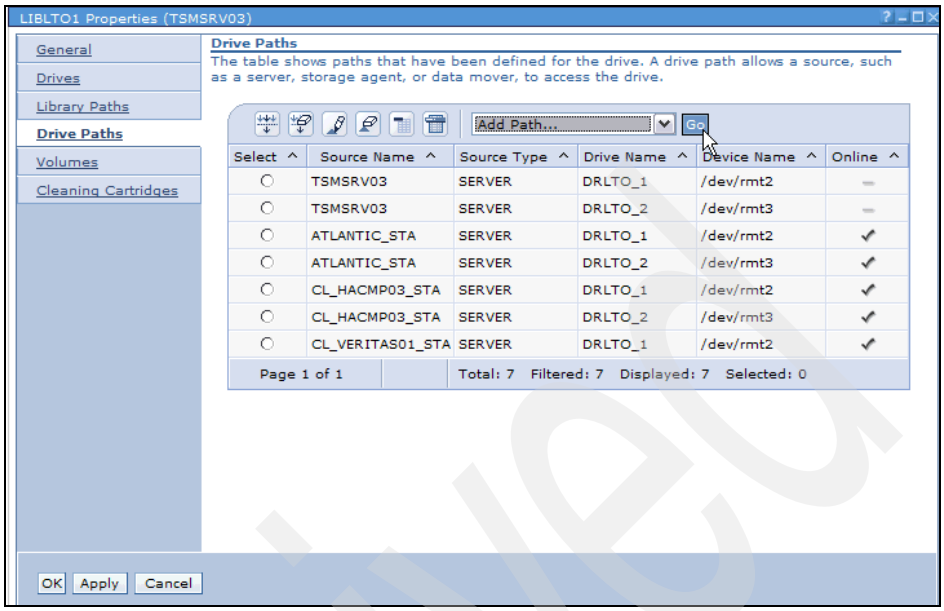

Figure 19-5 Administration Center screen to define a second drive path

12. We then fill out the panel to configure the second drive path to our local special device file and the TSMSRV03 drive equivalent, as seen in Figure 19-6.

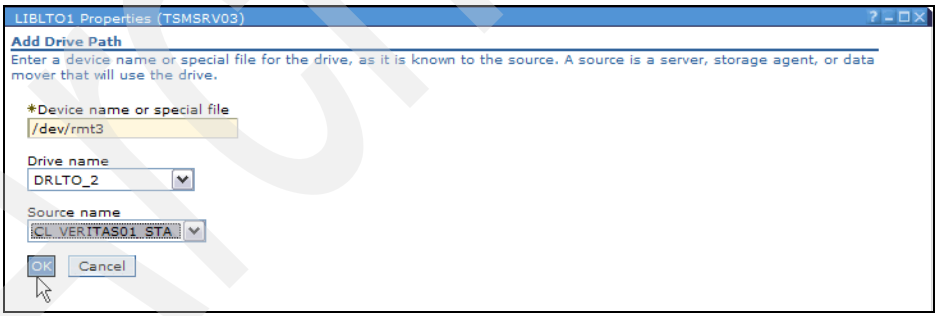

Figure 19-6 Administration Center screen to define a second drive path mapping

13. Finally, we click **OK**, and now we have our drives configured for the cl_veritas01_sta Storage Agent.

Archived

19.6 Configuring a cluster application

In the following sections we describe how to configure the cluster application.

Scripts for the Tivoli Storage Manager Storage Agent

We place the scripts for the server in the rootvg, /opt filesystem, in the directory /opt/local/tsmsta:

1. First, we place the start script in the directory as /opt/local/tsmsrv/startSTA.sh as shown in Example 19-5.

Example 19-5 /opt/local/tsmsta/startSTA.sh

```
#!/bin/ksh
##############################################################################
#                                                                            #
# Shell script to start a StorageAgent.                                      #
#                                                                            #
# Originated from the sample TSM server start script                         #
#                                                                            #
##############################################################################
echo "Starting Storage Agent now..."
# Start up TSM storage agent
##############################################################################
#
# Set the correct configuration
# dsmsta honors same variables as dsmserv does
export DSMSERV_CONFIG=/opt/IBM/ISC/tsm/StorageAgent/bin/dsmsta.opt
export DSMSERV_DIR=/opt/IBM/ISC/tsm/StorageAgent/bin
#export DSMSERV_DIR=/usr/tivoli/tsm/StorageAgent/bin

# Get the language correct....
export LANG=en_US

# max out size of data area
ulimit -d unlimited

#OK, now fire-up the storage agent in quiet mode.
print "$(date '+%D %T') Starting Tivoli Storage Manager storage agent"

cd /opt/IBM/ISC/tsm/StorageAgent/bin
$DSMSERV_DIR/dsmsta quiet &
exit 0
```

Archived

2. We then place the stop script in the directory as /opt/local/tsmsrv/stopSTA.sh, as shown in Example 19-6.

Example 19-6 /opt/local/tsmsta/stopSTA.sh

```
#!/bin/ksh
# killing the StorageAgent server process

##############################################################################
#
# Shell script to stop a TSM AIX Storage Agent.
# Please note that changes must be made to the dsmadmc command below in order
# to tailor it for your environment:
#
#    1. Set -servername= to the TSM server name on the SErvername option
##       in the /usr/tivoli/tsm/client/ba/bin/dsm.sys file.
#    2. Set -id= and -password= to a TSM userid that has been granted
##       operator authority, as described in the section:
##       "Chapter 3. Customizing Your Tivoli Storage Manager System -
##       Adding Administrators", in the Quick Start manual.
#    3. Edit the path in the LOCKFILE= statement to the directory where your
##       Storage Agent runs.
#
##############################################################################
#
# Set seconds to sleep.
secs=5
# TSM lock file
LOCKFILE="/opt/IBM/ISC/tsm/StorageAgent/bin/adsmserv.lock"

echo "Stopping the TSM Storage Agent now..."

# Check to see if the adsmserv.lock file exists.  If not then the server is not
running
if [[ -f $LOCKFILE ]]; then

  read J1 J2 J3 PID REST < $LOCKFILE
  /usr/tivoli/tsm/client/ba/bin/dsmadmc -servername=cl_veritas01_sta -id=admin
-password=admin -noconfirm << EOF
halt
EOF
  echo "Waiting for TSM server Storage Agent on pid $PID to stop..."

# Make sure all of the threads have ended
  while [[ `ps -m -o THREAD -p $PID | grep -c $PID` > 0 ]]; do
    sleep $secs
  done
fi
```

```
# Just in case the above doesn't stop the STA, then we'll hit it with a hammer
STAPID=`ps -af | egrep "dsmsta" | awk '{ print $2 }'`

for PID in $STAPID
 do
  kill -9 $PID
 done

exit 0
```

3. Next, we place the clean script in the directory /opt/local/tsmsta/cleanSTA.sh, as shown in Example 19-7.

Example 19-7 /opt/local/tsmsta/cleanSTA.sh

```
#!/bin/ksh
# killing StorageAgent server process if the stop fails

STAPID=`ps -af | egrep "dsmsta" | awk '{ print $2 }'`

for PID in $STAPID
 do
  kill $PID
 done

LINES=`ps -af | grep "/opt/IBM/ISC/tsm/StorageAgent/bin/dsmsta quiet" | awk
'{print $2}' | wc | awk '{print $1}'` >/dev/console 2>&1

STAPID=`ps -af | egrep "dsmsta" | awk '{ print $2 }'`
if [ $LINES -gt 1 ]
   then
        for PID in $STAPID
         do
                kill -9 $PID
         done
fi
exit 0
```

4. Lastly, we monitor the storageagent using the script monSTA.sh, as shown in Example 19-8.

Example 19-8 monSTA.sh script

```
#!/bin/ksh
# Monitoring for the existance of the ISC
# killing all AppServer related java processes left running

LINES=`ps -ef | egrep dsmsta | awk '{print $2}' | wc | awk '{print $1}'`
>/dev/console 2>&1
```

```
if [ $LINES -gt 1 ]
        then exit 110
fi
sleep 10
exit 100
```

5. We now add the Clustered Storage Agent into the VCS configuration, by adding an additional application within the same Service Group (sg_isc_sta_tsmcli). As this new application, we will use the same shared disk as the ISC (iscvg). Observe the **unlink** and **link** commands as we establish the parent-child relationship with the tsmcli application. This is all accomplished using the commands shown in Example 19-9.

Example 19-9 VCS commands to add app_sta application into sg_isc_sta_tsmcli

```
haconf -makerw
hares -add app_sta Application sg_isc_sta_tsmcli
hares -modify app_sta Critical 1
hares -modify app_sta User ""
hares -modify app_sta StartProgram /opt/local/tsmsta/startSTA.sh
hares -modify app_sta StopProgram /opt/local/tsmsta/stopSTA.sh
hares -modify app_sta CleanProgram /opt/local/tsmsta/cleanSTA.sh
hares -modify app_sta MonitorProgram /opt/local/tsmsta/monSTA.sh
hares -modify app_sta PidFiles -delete -keys
hares -modify app_sta MonitorProcesses ""
hares -probe app_sta -sys banda
hares -probe app_sta -sys atlantic
hares -unlink app_tsmcad app_pers_ip
hares -link app_sta app_pers_ip
hares -link app_tsmcad app_sta
hares -modify app_sta Enabled 1
haconf -dump -makero
```

Archived

6. Next we review the Veritas Cluster Manager GUI to ensure that everything is linked as expected, which is shown in Figure 19-7.

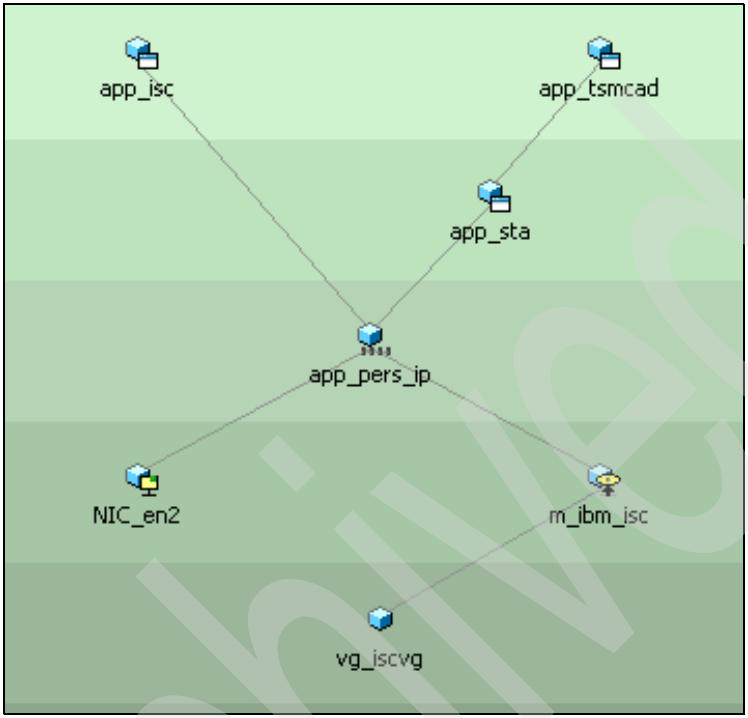

Figure 19-7 Veritas Cluster Manager GUI, sg_isc_sta_tsmcli resource relationship

7. Next, we review the /etc/VRTSvcs/conf/config/main.cf file, as shown in Example 19-10.

Example 19-10 The completed /etc/VRTSvcs/conf/config/main.cf file

```
group sg_isc_sta_tsmcli (
        SystemList = { banda = 0, atlantic = 1 }
        AutoStartList = { banda, atlantic }
        )

        Application app_isc (
                Critical = 0
                StartProgram = "/opt/local/isc/startISC.sh"
                StopProgram = "/opt/local/isc/stopISC.sh"
                CleanProgram = "/opt/local/isc/cleanISC.sh"
                MonitorProgram = "/opt/local/isc/monISC.sh"
                )

        Application app_sta (
```

```
                StartProgram = "/opt/local/tsmsta/startSTA.sh"
                StopProgram = "/opt/local/tsmsta/stopSTA.sh"
                CleanProgram = "/opt/local/tsmsta/cleanSTA.sh"
                MonitorProgram = "/opt/local/tsmsta/monSTA.sh"
                MonitorProcesses = { "" }
                )

        Application app_tsmcad (
                Critical = 0
                StartProgram = "/opt/local/tsmcli/startTSMcli.sh"
                StopProgram = "/opt/local/tsmcli/stopTSMcli.sh"
                CleanProgram = "/opt/local/tsmcli/stopTSMcli.sh"
               MonitorProcesses = { "/usr/tivoli/tsm/client/ba/bin/dsmc sched"
}
                )

        IP app_pers_ip (
                Device = en2
                Address = "9.1.39.77"
                NetMask = "255.255.255.0"
                )

        LVMVG vg_iscvg (
                VolumeGroup = iscvg
                MajorNumber = 48
                )

        Mount m_ibm_isc (
                MountPoint = "/opt/IBM/ISC"
                BlockDevice = "/dev/isclv"
                FSType = jfs2
                FsckOpt = "-y"
                )

        NIC NIC_en2 (
                Device = en2
                NetworkType = ether
                )

        app_isc requires app_pers_ip
        app_pers_ip requires NIC_en2
        app_pers_ip requires m_ibm_isc
        app_sta requires app_pers_ip
        app_tsmcad requires app_sta
        m_ibm_isc requires vg_iscvg


        // resource dependency tree
        //
```

Archived

```
//      group sg_isc_sta_tsmcli
//      {
//      Application app_isc
//          {
//          IP app_pers_ip
//              {
//              NIC NIC_en2
//              Mount m_ibm_isc
//                  {
//                  LVMVG vg_iscvg
//                  }
//              }
//          }
//      Application app_tsmcad
//          {
//          Application app_sta
//              {
//              IP app_pers_ip
//                  {
//                  NIC NIC_en2
//                  Mount m_ibm_isc
//                      {
//                      LVMVG vg_iscvg
//                      }
//                  }
//              }
//          }
//      }
```

8. We are now ready to put this resource online and test it.

19.7 Testing

We will now begin to test the cluster environment.

19.7.1 Veritas Cluster Server testing

Here we are testing basic cluster functions.This can help in problem determination if something goes wrong later on during setup and further testing.

We determine state of the cluster services by using the **hastatus** command from the AIX command line, and run a **tail** on the main cluster log, on both systems in the cluster.

Archived

19.7.2 Node power failure

Initially, this test is run with the applications OFFLINE:

1. First, we verify that the Service Groups are OFFLINE using the Veritas **hastatus** command, as shown in Example 19-11.

Example 19-11 The results return from hastatus

```
banda:/# hastatus
attempting to connect....connected

group           resource             system               message
--------------- -------------------- -------------------- --------------------
                                     atlantic             RUNNING
                                     banda                RUNNING
sg_tsmsrv                            banda                OFFLINE
sg_tsmsrv                            atlantic             OFFLINE
-------------------------------------------------------------------------
sg_isc_sta_tsmcli                      banda                OFFLINE
sg_isc_sta_tsmcli                       atlantic             OFFLINE
```

2. Next, we clear the VCS log by doing the **command cp /dev/null /var/VRTSvcs/log/engine_A.log.** For testing purposes, clearing the log prior, then copping the contents of the complete log after the test to an appropriately named file is a good methodology to reduce the log data you must sort through for a test, yet preserving the historical integrity of the test results.
3. Then, we do the AIX command tail -f **/var/VRTSvcs/log/engine_A.log.** This allows us to monitor the transition real-time.
4. Next we fail Banda by **pulling the power plug**. The results of the **hastatus** log on the surviving node (Atlantic) is shown in Example 19-12, and the result **tail** of the engine_A.log on Atlantic is shown in Example 19-13.

Example 19-12 hastatus log from the surviving node, Atlantic

```
Atlantic:/var/VRTSvcs/log# hastatus
attempting to connect....connected

group           resource             system               message
--------------- -------------------- -------------------- --------------------
                                     atlantic             RUNNING
                                     banda                *FAULTED*
```

Archived

Example 19-13 tail -f /var/VRTSvcs/log/engine_A.log from surviving node, Atlantic

```
VCS INFO V-16-1-10077 Received new cluster membership
VCS NOTICE V-16-1-10080 System (atlantic) - Membership: 0x1, Jeopardy: 0x0
VCS ERROR V-16-1-10079 System banda (Node '1') is in Down State - Membership:
0x1
VCS ERROR V-16-1-10322 System banda (Node '1') changed state from RUNNING to
FAULTED
```

5. Then, we restart Banda and wait for the cluster to recover, then review the **hastatus**, which has returned to full cluster membership. This is shown in Example 19-14.

Example 19-14 The recovered cluster using hastatus

```
banda:/# hastatus
attempting to connect....connected

group           resource            system              message
--------------- -------------------- -------------------- --------------------
                                    atlantic            RUNNING
                                    banda               RUNNING
sg_tsmsrv                           banda               OFFLINE
sg_tsmsrv                           atlantic            OFFLINE
---------------------------------------------------------------------------
sg_isc_sta_tsmcli                    banda               OFFLINE
sg_isc_sta_tsmcli                    atlantic            OFFLINE
```

6. We then repeat this process for the other node, Atlantic.

Results

Once the cluster recovers, we repeat the process for the other node, ensuring that full cluster recovery occurs. Once the test has occurred on both nodes, and recovery details have been confirmed as functioning correctly, this test is complete.

19.7.3 Start Service Group (bring online)

Here are the steps we follow for this test:

1. To begin, we review the current cluster status, confirming that all resources are offline, as shown from the **hastatus** command output, detailed in Example 19-15.

Archived

Example 19-15 Current cluster status from the hastatus output

```
banda:/# hastatus
attempting to connect....connected

group           resource             system               message
--------------- -------------------- -------------------- --------------------
                                     atlantic             RUNNING
                                     banda                RUNNING
sg_tsmsrv                            banda                OFFLINE
sg_tsmsrv                            atlantic             OFFLINE
-----------------------------------------------------------------------------
sg_isc_sta_tsmcli                      banda                OFFLINE
sg_isc_sta_tsmcli                      atlantic             OFFLINE
```

2. We then clear the log using **cp /dev/null /var/VRTSvcs/logengine_A.log** and then start a **tail -f /var/VRTSvcs/logengine_A.log**.
3. Next, from Atlantic (this can be done on any node), we bring the sg_isc_sta_tsmcli and the sg_tsmsrv Service Groups online on Banda using the **hagrp** command from the AIX command line, as shown in Example 19-16.

Example 19-16 hagrp -online command

```
Atlantic:/opt/local/tsmcli# hagrp -online sg_isc_sta_tsmcli -sys banda
-localclus
Atlantic:/opt/local/tsmcli# hagrp -online sg_tsmsrv -sys banda -localclus
```

4. We then view the **hastatus | grep banda** and verify the results as shown in Example 19-17.

Example 19-17 hastatus of online transition for sg_isc_sta_tsmcli Service Group

```
banda:/# hastatus | grep ONLINE
attempting to connect....connected
sg_tsmsrv                            banda                ONLINE
sg_isc_sta_tsmcli                      banda                ONLINE
sg_tsmsrv                            banda                ONLINE
sg_isc_sta_tsmcli                      banda                ONLINE
                 vg_tsmsrv           banda                ONLINE
                 ip_tsmsrv           banda                ONLINE
                 m_tsmsrv_db1        banda                ONLINE
                 m_tsmsrv_db1mr1     banda                ONLINE
                 m_tsmsrv_lg1        banda                ONLINE
                 m_tsmsrv_lgmr1      banda                ONLINE
                 m_tsmsrv_dp1        banda                ONLINE
                 m_tsmsrv_files      banda                ONLINE
                 app_tsmsrv          banda                ONLINE
                 NIC_en1             banda                ONLINE
                 NIC_en1             atlantic             ONLINE
```

Archived

```
app_isc              banda                ONLINE
app_pers_ip          banda                ONLINE
vg_iscvg             banda                ONLINE
m_ibm_isc            banda                ONLINE
app_sta              banda                ONLINE
app_tsmcad           banda                ONLINE
```

5. Then we review the engine_A.log shown in Example 19-18.

Example 19-18 tail -f /var/VRTSvcs/log/engine_A.log

```
VCS INFO V-16-1-50135 User root fired command: hagrp -online sg_isc_sta_tsmcli
banda  localclus  from localhost
VCS NOTICE V-16-1-10166 Initiating manual online of group sg_isc_sta_tsmcli on
system banda
VCS INFO V-16-1-50135 User root fired command: hagrp -online sg_tsmsrv banda
localclus  from localhost
VCS NOTICE V-16-1-10301 Initiating Online of Resource vg_iscvg (Owner: unknown,
Group: sg_isc_sta_tsmcli) on System banda

.
.
.
VCS NOTICE V-16-1-10447 Group sg_tsmsrv is online on system banda
VCS NOTICE V-16-1-10447 Group sg_isc_sta_tsmcli is online on system banda
```

19.7.4 Stop Service Group (bring offline)

Here are the steps we follow for this test:

1. Before every test, we check the status for cluster services, resource groups, and resources on both nodes; In Example 19-19 we are verifying using **hastatus**. For this test, we expect that all applications are offline, as we are just testing the cluster's core functionality.

Example 19-19 Verify available cluster resources using the hastatus command

```
banda:/var/VRTSvcs/log# hastatus
attempting to connect....connected

group           resource             system               message
--------------- -------------------- -------------------- --------------------
                                     atlantic             RUNNING
                                     banda                RUNNING
sg_tsmsrv                            banda                ONLINE
sg_tsmsrv                            atlantic             OFFLINE
---------------------------------------------------------------------------
sg_isc_sta_tsmcli                      banda                ONLINE
sg_isc_sta_tsmcli                      atlantic             OFFLINE
```

```
sg_tsmsrv                          banda               ONLINE
sg_tsmsrv                          atlantic            OFFLINE
sg_isc_sta_tsmcli                   banda               ONLINE
--------------------------------------------------------------------------
sg_isc_sta_tsmcli                   atlantic            OFFLINE
               vg_tsmsrv           banda               ONLINE
               vg_tsmsrv           atlantic            OFFLINE
               ip_tsmsrv           banda               ONLINE
               ip_tsmsrv           atlantic            OFFLINE
--------------------------------------------------------------------------
               m_tsmsrv_db1        banda               ONLINE
               m_tsmsrv_db1        atlantic            OFFLINE
               m_tsmsrv_db1mr1     banda               ONLINE
               m_tsmsrv_db1mr1     atlantic            OFFLINE
               m_tsmsrv_lg1        banda               ONLINE
--------------------------------------------------------------------------
               m_tsmsrv_lg1        atlantic            OFFLINE
               m_tsmsrv_lgmr1      banda               ONLINE
               m_tsmsrv_lgmr1      atlantic            OFFLINE
               m_tsmsrv_dp1        banda               ONLINE
               m_tsmsrv_dp1        atlantic            OFFLINE
--------------------------------------------------------------------------
               m_tsmsrv_files      banda               ONLINE
               m_tsmsrv_files      atlantic            OFFLINE
               app_tsmsrv          banda               ONLINE
               app_tsmsrv          atlantic            OFFLINE
               NIC_en1             banda               ONLINE
--------------------------------------------------------------------------
               NIC_en1             atlantic            ONLINE
               app_isc             banda               ONLINE
               app_isc             atlantic            OFFLINE
               app_pers_ip         banda               ONLINE
               app_pers_ip         atlantic            OFFLINE
--------------------------------------------------------------------------
               vg_iscvg            banda               ONLINE
               vg_iscvg            atlantic            OFFLINE
               m_ibm_isc           banda               ONLINE
               m_ibm_isc           atlantic            OFFLINE
               app_sta             banda               ONLINE
--------------------------------------------------------------------------
               app_sta             atlantic            OFFLINE
               app_tsmcad          banda               ONLINE
               app_tsmcad          atlantic            OFFLINE
               NIC_en2             banda               ONLINE
               NIC_en2             atlantic            ONLINE
--------------------------------------------------------------------------
               vg_tsmsrv           banda               ONLINE
               vg_tsmsrv           atlantic            OFFLINE
               ip_tsmsrv           banda               ONLINE
```

Archived

```
                ip_tsmsrv            atlantic             OFFLINE
                m_tsmsrv_db1         banda                ONLINE
-----------------------------------------------------------------------
                m_tsmsrv_db1         atlantic             OFFLINE
                m_tsmsrv_db1mr1      banda                ONLINE
                m_tsmsrv_db1mr1      atlantic             OFFLINE
                m_tsmsrv_lg1         banda                ONLINE
                m_tsmsrv_lg1         atlantic             OFFLINE
-----------------------------------------------------------------------
                m_tsmsrv_lgmr1       banda                ONLINE
                m_tsmsrv_lgmr1       atlantic             OFFLINE
                m_tsmsrv_dp1         banda                ONLINE
                m_tsmsrv_dp1         atlantic             OFFLINE
                m_tsmsrv_files       banda                ONLINE
-----------------------------------------------------------------------
                m_tsmsrv_files       atlantic             OFFLINE

group           resource             system               message
--------------- -------------------- -------------------- --------------------
                app_tsmsrv           banda                ONLINE
                app_tsmsrv           atlantic             OFFLINE
                NIC_en1              banda                ONLINE
                NIC_en1              atlantic             ONLINE
-----------------------------------------------------------------------
                app_isc              banda                ONLINE
                app_isc              atlantic             OFFLINE
                app_pers_ip          banda                ONLINE
                app_pers_ip          atlantic             OFFLINE
                vg_iscvg             banda                ONLINE
-----------------------------------------------------------------------
                vg_iscvg             atlantic             OFFLINE
                m_ibm_isc            banda                ONLINE
                m_ibm_isc            atlantic             OFFLINE
                app_sta              banda                ONLINE
                app_sta              atlantic             OFFLINE
-----------------------------------------------------------------------
                app_tsmcad           banda                ONLINE
                app_tsmcad           atlantic             OFFLINE
                NIC_en2              banda                ONLINE
                NIC_en2              atlantic             ONLINE
```

2. Now, we bring the applications OFFLINE using the **hagrp -offline** command, as shown in Example 19-20.

Example 19-20 hagrp -offline command

```
Atlantic:/opt/local/tsmcli# hagrp -offline sg_isc_sta_tsmcli -sys banda
-localclus
Atlantic:/opt/local/tsmcli# hagrp -offline sg_tsmsrv -sys banda -localclus
```

3. Now, we review the **hastatus** output as shown in Example 19-21.

Example 19-21 hastatus output for the Service Group OFFLINE

```
banda:/var/VRTSvcs/log# hastatus
attempting to connect....connected

group           resource             system               message
--------------- -------------------- -------------------- --------------------
                                     atlantic             RUNNING
                                     banda                RUNNING
sg_tsmsrv                            banda                OFFLINE
sg_tsmsrv                            atlantic             OFFLINE
-----------------------------------------------------------------------------
sg_isc_sta_tsmcli                      banda                OFFLINE
sg_isc_sta_tsmcli                      atlantic             OFFLINE
```

4. Then, we review the /var/VRTSvcs/log/engine_A.log, as shown in Example 19-22.

Example 19-22 tail -f /var/VRTSvcs/log/engine_A.log

```
VCS NOTICE V-16-1-10446 Group sg_tsmsrv is offline on system banda
VCS NOTICE V-16-1-10446 Group sg_isc_sta_tsmcli is offline on system banda
```

19.7.5 Manual Service Group switch

Here are the steps we follow for this test:

1. For this test, all Service Groups are on one node (Banda), and will be switched to Atlantic, using the Cluster Manager GUI. As with all tests, we clear the engine_A.log using **cp /dev/null /var/VRTSvcs/log/engine_A.log**. The **hastatus | grep ONLINE** output prior to starting the transition is shown in Example 19-23.

Example 19-23 hastatus output prior to the Service Groups switching nodes

```
banda:/var/VRTSvcs/log# hastatus |grep ONLINE
attempting to connect....connected
sg_tsmsrv                            banda                ONLINE
sg_isc_sta_tsmcli                      banda                ONLINE
sg_tsmsrv                            banda                ONLINE
```

```
sg_isc_sta_tsmcli                    banda              ONLINE
                vg_tsmsrv           banda              ONLINE
                ip_tsmsrv           banda              ONLINE
                m_tsmsrv_db1        banda              ONLINE
                m_tsmsrv_db1mr1     banda              ONLINE
                m_tsmsrv_lg1        banda              ONLINE
                m_tsmsrv_lgmr1      banda              ONLINE
                m_tsmsrv_dp1        banda              ONLINE
                m_tsmsrv_files      banda              ONLINE
                app_tsmsrv          banda              ONLINE
                NIC_en1             banda              ONLINE
                NIC_en1             atlantic           ONLINE
                app_isc             banda              ONLINE
                app_pers_ip         banda              ONLINE
                vg_iscvg            banda              ONLINE
                m_ibm_isc           banda              ONLINE
                app_sta             banda              ONLINE
                app_tsmcad          banda              ONLINE
```

2. Now, we switch the Service Groups using the Cluster Manager GUI, as shown in Figure 19-8.

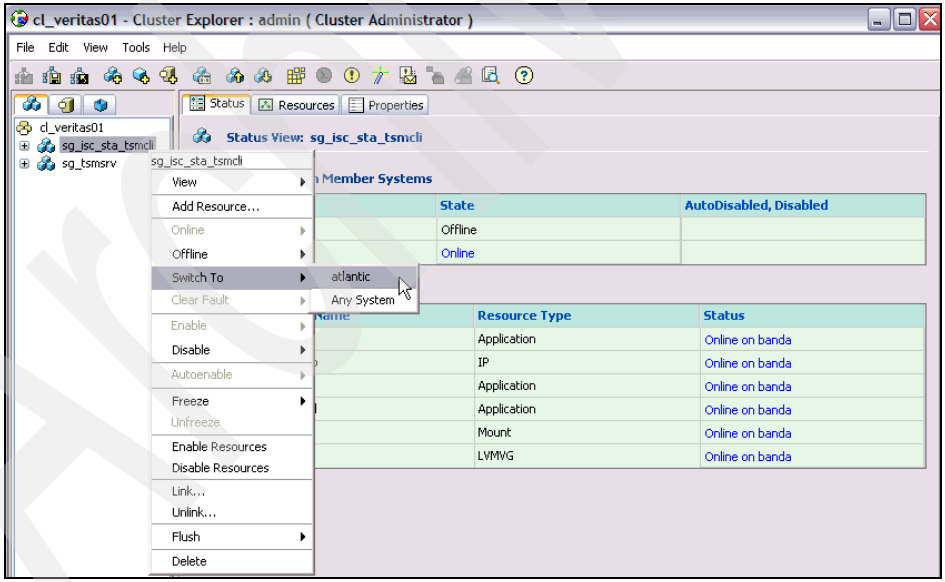

Figure 19-8 VCS Cluster Manager GUI switching Service Group to another node

3. Then, we click **Yes** to start the process as shown in Figure 19-9.

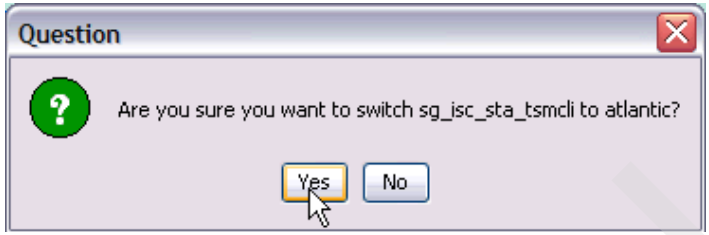

Figure 19-9 Prompt to confirm the switch

Tip: This process can be completed using the command line as well:

```
banda:/var/VRTSvcs/log# hagrp -switch sg_isc_sta_tsmcli -to atlantic
-localclus
banda:/var/VRTSvcs/log# hagrp -switch sg_tsmsrv -to atlantic -localclus
```

4. Now, we monitor the transition which can be seen using the Cluster Manager GUI, and review the results in **hastatus** and the engine_A.log. The two logs are shown in Example 19-24 and Example 19-25.

Example 19-24 hastatus output of the Service Group switch

```
^Cbanda:/var/VRTSvcs/log# hastatus |grep ONLINE
attempting to connect....connected
sg_tsmsrv                          atlantic            ONLINE
sg_isc_sta_tsmcli                   atlantic            ONLINE
sg_tsmsrv                          atlantic            ONLINE
sg_isc_sta_tsmcli                   atlantic            ONLINE
                vg_tsmsrv          atlantic            ONLINE
                ip_tsmsrv          atlantic            ONLINE
                m_tsmsrv_db1       atlantic            ONLINE
                m_tsmsrv_db1mr1    atlantic            ONLINE
                m_tsmsrv_lg1       atlantic            ONLINE
                m_tsmsrv_lgmr1     atlantic            ONLINE
                m_tsmsrv_dp1       atlantic            ONLINE
                m_tsmsrv_files     atlantic            ONLINE
                app_tsmsrv         atlantic            ONLINE
                NIC_en1            banda               ONLINE
                NIC_en1            atlantic            ONLINE
                app_isc            atlantic            ONLINE
                app_pers_ip        atlantic            ONLINE
                vg_iscvg           atlantic            ONLINE
                m_ibm_isc          atlantic            ONLINE
                app_sta            atlantic            ONLINE
                app_tsmcad         atlantic            ONLINE
```

Archived

Example 19-25 tail -f /var/VRTSvcs/log/engine_A.log from surviving node, Atlantic

```
VCS NOTICE V-16-1-10208 Initiating switch of group sg_isc_sta_tsmcli from
system banda to system atlantic
VCS NOTICE V-16-1-10300 Initiating Offline of Resource app_isc (Owner: unknown,
Group: sg_isc_sta_tsmcli) on System banda
VCS INFO V-16-1-50135 User root fired command: hagrp -switch sg_tsmsrv
atlantic  localclus  from localhost
VCS NOTICE V-16-1-10208 Initiating switch of group sg_tsmsrv from system banda
to system atlantic
VCS NOTICE V-16-1-10300 Initiating Offline of Resource app_tsmsrv (Owner:
unknown, Group: sg_tsmsrv) on System banda
.
.
.
VCS NOTICE V-16-1-10447 Group sg_tsmsrv is online on system banda
VCS NOTICE V-16-1-10447 Group sg_isc_sta_tsmcli is online on system atlantic
VCS NOTICE V-16-1-10448 Group sg_tsmsrv failed over to system atlantic
VCS NOTICE V-16-1-10448 Group sg_isc_sta_tsmcli failed over to system atlantic
```

Results

In this test, our Service Groups have completed the switch and are now online on Atlantic. This completes the test successfully.

19.7.6 Manual fallback (switch back)

Here are the steps we follow for this test:

1. Before every test we check the status for cluster services, resource groups and resources on both nodes; In Example 19-26 we are verifying using **hastatus**.

Example 19-26 hastatus output of the current cluster state

```
banda:/# hastatus |grep ONLINE
attempting to connect....connected
sg_tsmsrv                          atlantic             ONLINE
sg_isc_sta_tsmcli                    atlantic             ONLINE
sg_tsmsrv                          atlantic             ONLINE
sg_isc_sta_tsmcli                    atlantic             ONLINE
               vg_tsmsrv           atlantic             ONLINE
               ip_tsmsrv           atlantic             ONLINE
               m_tsmsrv_db1        atlantic             ONLINE
               m_tsmsrv_db1mr1     atlantic             ONLINE
               m_tsmsrv_lg1        atlantic             ONLINE
               m_tsmsrv_lgmr1      atlantic             ONLINE
               m_tsmsrv_dp1        atlantic             ONLINE
               m_tsmsrv_files      atlantic             ONLINE
```

```
app_tsmsrv          atlantic            ONLINE
NIC_en1             banda               ONLINE
NIC_en1             atlantic            ONLINE
app_isc             atlantic            ONLINE
app_pers_ip         atlantic            ONLINE
vg_iscvg            atlantic            ONLINE
m_ibm_isc           atlantic            ONLINE
app_sta             atlantic            ONLINE
app_tsmcad          atlantic            ONLINE
```

2. For this test, we will use the AIX command line to switch the Service Group back to Banda, as shown in Example 19-27.

Example 19-27 hargrp -switch command to switch the Service Group back to Banda

```
banda:/# hagrp -switch sg_tsmsrv -to banda -localclus
banda:/# hagrp -switch sg_isc_sta_tsmcli -to banda -localclus
```

3. We then review the results in the engine_A.log, as shown in Example 19-28.

Example 19-28 /var/VRTSvcs/log/engine_A.log segment for the switch back to Banda

```
VCS NOTICE V-16-1-10208 Initiating switch of group sg_tsmsrv from system
atlantic to system banda
VCS NOTICE V-16-1-10208 Initiating switch of group sg_isc_sta_tsmcli from
system atlantic to system banda
VCS NOTICE V-16-1-10300 Initiating Offline of Resource app_tsmsrv (Owner:
unknown, Group: sg_tsmsrv) on System atlantic
.
.
.
VCS NOTICE V-16-1-10447 Group sg_tsmsrv is online on system banda
VCS NOTICE V-16-1-10448 Group sg_tsmsrv failed over to system banda
VCS NOTICE V-16-1-10447 Group sg_isc_sta_tsmcli is online on system banda
VCS NOTICE V-16-1-10448 Group sg_isc_sta_tsmcli failed over to system banda
```

Results

Once we have the Service Group back on Banda, this test is now complete.

Archived

19.7.7 Public NIC failure

Testing the public NIC to ensure the behavior is what is expected.

Objective

Now we test the failure of a critical resource within the Service Group, the public NIC. First, we test the reaction of the cluster when the NIC fails (is physically disconnected), then we document the cluster's recovery behavior once the NIC is plugged back in. We anticipate that the Service Group sg_tsmsrv will fault the NIC_en1 on Atlantic, then failover to Banda. Once sg_tsmsrv resources come online on Banda, we replace the ethernet cable, which should produce a recovery of the resource, then we manually switch sg_tsmsrv back to Atlantic.

Test sequence

Here are the steps we follow for this test:

1. For this test, one Service Group will be on each node. As with all tests, we clear the engine_A.log using **cp /dev/null /var/VRTSvcs/log/engine_A.log**.
2. Next, we physically disconnect the ethernet cable from the EN1 device on Atlantic. This is defined as a critical resource for the Service Group in which the TSM server is the application. We will then observe the results in both logs being monitored.
3. Then we review the engine_A.log file to understand the transition actions, which is shown in Example 19-29.

Example 19-29 /var/VRTSvcs/log/engine_A.log output for the failure activity

```
VCS INFO V-16-1-10077 Received new cluster membership
VCS NOTICE V-16-1-10080 System (banda) - Membership: 0x3, Jeopardy: 0x2
VCS ERROR V-16-1-10087 System banda (Node '1') is in Regardy Membership -
Membership: 0x3, Jeopardy: 0x2
.
.
.
VCS WARNING V-16-10011-5607 (atlantic) NIC:NIC_en1:monitor:packet count test
failed: Resource is offline
VCS WARNING V-16-10011-5607 (atlantic) NIC:NIC_en1:monitor:packet count test
failed: Resource is offline
VCS INFO V-16-1-10307 Resource NIC_en1 (Owner: unknown, Group: sg_tsmsrv) is
offline on atlantic (Not initiated by VCS)
VCS NOTICE V-16-1-10300 Initiating Offline of Resource app_tsmsrv (Owner:
unknown, Group: sg_tsmsrv) on System atlantic
.
.
.
VCS INFO V-16-1-10298 Resource app_tsmsrv (Owner: unknown, Group: sg_tsmsrv) is
online on banda (VCS initiated)
```

```
VCS NOTICE V-16-1-10447 Group sg_tsmsrv is online on system banda
VCS NOTICE V-16-1-10448 Group sg_tsmsrv failed over to system banda
VCS WARNING V-16-10011-5607 (atlantic) NIC:NIC_en1:monitor:packet count test
failed: Resource is offline
VCS WARNING V-16-10011-5607 (atlantic) NIC:NIC_en1:monitor:packet count test
failed: Resource is offline
.
.
.
VCS WARNING V-16-10011-5607 (atlantic) NIC:NIC_en1:monitor:packet count test
failed: Resource is offline
VCS WARNING V-16-10011-5607 (atlantic) NIC:NIC_en1:monitor:packet count test
failed: Resource is offline
```

4. As a result of the failed NIC, which is a critical resource for sg_tsmsrv the Service Group fails over to Banda (from Atlantic).
5. Next, we plug the ethernet cable back into the NIC and monitor for a state change, and now the cluster ONLINE resources show that EN1 on Atlantic is back ONLINE, however there is no failback (resources are stable on Banda) and the cluster knows it is now capable of failing over to Atlantic for both NICs if required. The **hastatus** of the NIC1 transition is shown in Example 19-30.

Example 19-30 hastatus of the ONLINE resources

```
# hastatus |grep ONLINE
attempting to connect....connected
sg_tsmsrv                          banda                ONLINE
sg_isc_sta_tsmcli                   banda                ONLINE
sg_tsmsrv                          banda                ONLINE
sg_isc_sta_tsmcli                   banda                ONLINE
                vg_tsmsrv          banda                ONLINE
                ip_tsmsrv          banda                ONLINE
                m_tsmsrv_db1       banda                ONLINE
                m_tsmsrv_db1mr1    banda                ONLINE
                m_tsmsrv_lg1       banda                ONLINE
                m_tsmsrv_lgmr1     banda                ONLINE
                m_tsmsrv_dp1       banda                ONLINE
                m_tsmsrv_files     banda                ONLINE
                app_tsmsrv         banda                ONLINE
                NIC_en1            banda                ONLINE
                NIC_en1            atlantic             ONLINE
                app_isc            banda                ONLINE
                app_pers_ip        banda                ONLINE
                vg_iscvg           banda                ONLINE
                m_ibm_isc          banda                ONLINE
                app_sta            banda                ONLINE
                app_tsmcad         banda                ONLINE
```

Archived

6. Then, we review the contents of the engine_A.log, which is shown in Example 19-31.

Example 19-31 /var/VRTSvcs/log/engine_A.log output for the recovery activity

```
VCS INFO V-16-1-10077 Received new cluster membership
VCS NOTICE V-16-1-10080 System (banda) - Membership: 0x3, Jeopardy: 0x0
VCS NOTICE V-16-1-10086 System banda (Node '1') is in Regular Membership -
Membership: 0x3
VCS INFO V-16-1-10299 Resource NIC_en1 (Owner: unknown, Group: sg_tsmsrv) is
online on atlantic (Not initiated by VCS)
```

7. At this point we manually switch the sg_tsmsrv back over to Atlantic, with the ONLINE resources shown in **hastatus** in Example 19-32, which then concludes this test.

Example 19-32 hastatus of the online resources fully recovered from the failure test

```
hastatus |grep ONLINE
attempting to connect....connected
sg_tsmsrv                           atlantic            ONLINE
sg_isc_sta_tsmcli                    banda               ONLINE
sg_tsmsrv                           atlantic            ONLINE
sg_isc_sta_tsmcli                    banda               ONLINE
               vg_tsmsrv            atlantic            ONLINE
               ip_tsmsrv            atlantic            ONLINE
               m_tsmsrv_db1         atlantic            ONLINE
               m_tsmsrv_db1mr1      atlantic            ONLINE
               m_tsmsrv_lg1         atlantic            ONLINE
               m_tsmsrv_lgmr1       atlantic            ONLINE
               m_tsmsrv_dp1         atlantic            ONLINE
               m_tsmsrv_files       atlantic            ONLINE
               app_tsmsrv           atlantic            ONLINE
               NIC_en1              banda               ONLINE
               NIC_en1              atlantic            ONLINE
               app_isc              banda               ONLINE
               app_pers_ip          banda               ONLINE
               vg_iscvg             banda               ONLINE
               m_ibm_isc            banda               ONLINE
               app_sta              banda               ONLINE
               app_tsmcad           banda               ONLINE
```

19.7.8 LAN-free client system failover while the client is backing up

Now we test the ability of a scheduled backup operation over SAN to restart and complete, still over SAN, after the node Banda fails while a tape is in use by the Storage Agent cl_veritas01_sta:

1. We verify that the cluster services are running with the **hastatus** command.
2. On Atlantic (which is the surviving node), we use **tail -f /var/VRTSvcs/log/engine_A.log** to monitor cluster operation.
3. Then we schedule a client selective backup having the whole shared file systems as an object, as shown in Example 19-33.

Example 19-33 Client selective backup schedule configured on TSMSRV03

```
Policy Domain Name: STANDARD
                 Schedule Name: RESTORE
                   Description:
                        Action: Restore
                       Options: -subdir=yes -replace=yes
                       Objects: /mnt/nfsfiles/root/*
                      Priority: 5
               Start Date/Time: 02/22/05   10:44:27
                      Duration: 15 Minute(s)
                 Schedule Style: Classic
                        Period: One Time
                   Day of Week: Any
                         Month:
                  Day of Month:
                 Week of Month:
                    Expiration:
Last Update by (administrator): ADMIN
         Last Update Date/Time: 02/22/05   10:44:27
              Managing profile:
```

4. Then wait for the session to start, monitoring this using **query session** on the Tivoli Storage Manager server TSMSRV03, as shown in Example 19-34.

Example 19-34 Client sessions starting

```
6,585    Tcp/Ip    IdleW    12 S     1.9 K     1.2 K    Serv-    AIX-RS/-    CL_VERITAS01_STA
6,588    Tcp/Ip    IdleW    12 S     3.5 K     1.6 K    Serv-    AIX-RS/-    CL_VERITAS01_STA
6,706    Tcp/Ip    IdleW     3 S     1,002       642    Node     AIX         CL_VERITAS01_CLIENT
6,707    Tcp/Ip    RecvW    13 S       349     8.1 M    Node     AIX         CL_VERITAS01_CLIENT
6,708    Tcp/Ip    Run       0 S       474   119.5 M    Node     AIX         CL_VERITAS01_CLIENT
```

5. We wait for volume to be mounted either by monitoring the server console or doing a **query mount** as shown in Example 19-35.

Example 19-35 Tivoli Storage Manager server volume mounts

```
tsm: TSMSRV03>q mount
ANR8330I LTO volume 030AKK is mounted R/W in drive DRLTO_2 (/dev/rmt1), status: IN USE.
ANR8330I LTO volume 031AKK is mounted R/W in drive DRLTO_1 (/dev/rmt0), status: IN USE.
ANR8334I         2 matches found.
```

Archived

Failure

Being sure that client LAN-free backup is running, we issue **halt -q** on the AIX server on Atlantic, on which backup is running; the **halt -q** command stops any activity immediately and powers off the server.

The server remains waiting for client and Storage Agent communication until idletimeout expires (the default is 15 minutes). The Tivoli Storage Manager server reports the failure on the server console as shown in Example 19-36.

Example 19-36 The sessions being cancelled at the time of failure

```
ANR0490I Canceling session 6585 for node CL_VERITAS01_STA (AIX-RS/6000) .
ANR3605E Unable to communicate with storage agent.
ANR0490I Canceling session 6588 for node CL_VERITAS01_STA (AIX-RS/6000) .
ANR3605E Unable to communicate with storage agent.
```

Recovery

Here are the steps we follow:

1. The second node, Atlantic takes over the resources and launches the application server start script. Once this happens, the Tivoli Storage Manager server logs the difference in physical node names, reserved devices are reset, and the Storage Agent is started, as seen in the server actlog, shown in Example 19-37.

Example 19-37 TSMSRV03 actlog of the cl_veritas01_sta recovery process

```
ANR0408I Session 6721 started for server CL_VERITAS01_STA (AIX-RS/6000)
(Tcp/Ip) for event logging.
ANR0409I Session 6720 ended for server CL_VERITAS01_STA (AIX-RS/6000).
ANR0408I Session 6722 started for server CL_VERITAS01_STA (AIX-RS/6000)
(Tcp/Ip) for storage agent.
ANR0407I Session 6723 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
9.1.39.92(33332)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 6723 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 6724 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
9.1.39.42(33333)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 6724 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 6725 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
9.1.39.92(33334)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 6725 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 6726 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
```

```
9.1.39.42(33335)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 6726 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 6727 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
9.1.39.92(33336)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 6727 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 6728 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
9.1.39.42(33337)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 6728 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 6729 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
9.1.39.92(33338)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 6729 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 6730 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
9.1.39.42(33339)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 6730 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 6731 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
9.1.39.92(33340)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 6731 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 6732 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
9.1.39.42(33341)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 6732 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 6733 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
9.1.39.92(33342)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 6733 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 6734 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
9.1.39.42(33343)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 6734 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 6735 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
9.1.39.92(33344)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 6735 ended for administrator SCRIPT_OPERATOR (AIX).
```

```
ANR0407I Session 6736 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
9.1.39.42(33345)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 6736 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 6737 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
9.1.39.92(33346)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 6737 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 6738 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
9.1.39.42(33347)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 6738 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0406I Session 6739 started for node CL_VERITAS01_CLIENT (AIX) (Tcp/Ip
9.1.39.92(33349)).
ANR1639I Attributes changed for node CL_VERITAS01_CLIENT: TCP Name from banda
to atlantic, GUID from 00.00.00.00.75.8e.11.d9.ac.29.08.63.09.01.27.5e to
00.00.00.01.75.8f.11.d9.b4.d1.08.63.09.01.27.5c.
ANR0406I Session 6740 started for node CL_VERITAS01_CLIENT (AIX) (Tcp/Ip
9.1.39.42(33351)).
```

2. Now, we review the current process situation, as seen in Example 19-38. We see that there are currently 6 CL_VERITAS01_CLIENT sessions. The three older sessions (6706, 6707, 6708) will be cancelled by the logic imbedded within our startTSMcli.sh script. Once this happens, there will only be three client sessions remaining.

Example 19-38 Server process view during LAN-free backup recovery

```
6,706    Tcp/Ip    IdleW    8.3 M      1.0 K      682     Node    AIX         CL_VERITAS01_CLIENT
6,707    Tcp/Ip    RecvW    8.2 M        424     16.9 M    Node    AIX         CL_VERITAS01_CLIENT
6,708    Tcp/Ip    IdleW    8.2 M        610    132.0 M    Node    AIX         CL_VERITAS01_CLIENT
6,719    Tcp/Ip    IdleW      7 S      1.4 K      722     Serv-   AIX-RS/-    CL_VERITAS01_STA
6,721    Tcp/Ip    IdleW    3.4 M        257      1.4 K    Serv-   AIX-RS/-    CL_VERITAS01_STA
6,722    Tcp/Ip    IdleW      7 S        674      639     Serv-   AIX-RS/-    CL_VERITAS01_STA
6,739    Tcp/Ip    IdleW    3.1 M        978      621     Node    AIX         CL_VERITAS01_CLIENT
6,740    Tcp/Ip    MediaW   3.4 M        349      8.1 M   Node    AIX         CL_VERITAS01_CLIENT
6,742    Tcp/Ip    MediaW   3.1 M        349      7.5 M   Node    AIX         CL_VERITAS01_CLIENT
```

3. Once the Storage Agent scripts completes, the clustered scheduler start script begins. The startup of the client and Storage Agent will first search for previous tape using sessions to cancel. First, we observe the older Storage Agent sessions being terminated, as shown in Example 19-29.

Example 19-39 Extract of console log showing session cancelling work

```
ANR0483W Session 6159 for node CL_VERITAS01_STA
                          (AIX-RS/6000) terminated - forced by administrator.
                          (SESSION: 6159)
ANR0483W Session 6161 for node CL_VERITAS01_STA
                          (AIX-RS/6000) terminated - forced by administrator.
                          (SESSION: 6161)
ANR0483W Session 6162 for node CL_VERITAS01_STA
                          (AIX-RS/6000) terminated - forced by administrator.
                          (SESSION: 6162)
```

Note: Sessions with *_VOL_ACCESS not null increase the node mount point used number, preventing new sessions from the same node to obtain new mount points by the MAXNUMMP parameter. To assist in managing this, the node point points were increased from the default of 1 to 3.

4. Once the session's cancelling work finishes, the scheduler is restarted and the scheduled backup operation is restarted, as seen from the client log, shown in Example 19-40.

Example 19-40 dsmsched.log output showing failover transition, schedule restarting

```
02/22/05   17:16:59 Normal File-->              117
/opt/IBM/ISC/AppServer/installedApps/DefaultNo
de/wps.ear/wps.war/themes/html/ps/com/ibm/ps/uil/nls/TB_help_pushed_24.gif [Sent]
02/22/05   17:16:59 Normal File-->              111
/opt/IBM/ISC/AppServer/installedApps/DefaultNo
de/wps.ear/wps.war/themes/html/ps/com/ibm/ps/uil/nls/TB_help_unavail_24.gif [Sent]
02/22/05   17:18:48 Querying server for next scheduled event.
02/22/05   17:18:48 Node Name: CL_VERITAS01_CLIENT
02/22/05   17:18:48 Session established with server TSMSRV03: AIX-RS/6000
02/22/05   17:18:48   Server Version 5, Release 3, Level 0.0
02/22/05   17:18:48   Server date/time: 02/22/05   17:18:30  Last access: 02/22/05   17:15:45

02/22/05   17:18:48 --- SCHEDULEREC QUERY BEGIN
02/22/05   17:18:48 --- SCHEDULEREC QUERY END
02/22/05   17:18:48 Next operation scheduled:
02/22/05   17:18:48 ------------------------------------------------------------
02/22/05   17:18:48 Schedule Name:         TEST_SCHED
02/22/05   17:18:48 Action:                Selective
02/22/05   17:18:48 Objects:               /opt/IBM/ISC/*
02/22/05   17:18:48 Options:               -subdir=yes
02/22/05   17:18:48 Server Window Start:   17:10:08 on 02/22/05
02/22/05   17:18:48 ------------------------------------------------------------
02/22/05   17:18:48
Executing scheduled command now.
02/22/05   17:18:48 --- SCHEDULEREC OBJECT BEGIN TEST_SCHED 02/22/05   17:10:08
```

```
02/22/05   17:18:48 Selective Backup function invoked.

02/22/05   17:18:49 ANS1898I ***** Processed     1,500 files *****
02/22/05   17:18:49 Directory-->               4,096 /opt/IBM/ISC/ [Sent]
02/22/05   17:18:49 Directory-->               4,096 /opt/IBM/ISC/AppServer [Sent]
```

5. Backup completion then occurs, with the summary as shown in Example 19-41.

Example 19-41 Backup during a failover shows a completed successful summary

```
02/22/05   17:31:34 ANS1804E Selective Backup processing of '/opt/IBM/ISC/*' finished with
failures.

02/22/05   17:31:34 --- SCHEDULEREC STATUS BEGIN
02/22/05   17:31:34 Total number of objects inspected:   24,466
02/22/05   17:31:34 Total number of objects backed up:   24,465
02/22/05   17:31:34 Total number of objects updated:         0
02/22/05   17:31:34 Total number of objects rebound:         0
02/22/05   17:31:34 Total number of objects deleted:         0
02/22/05   17:31:34 Total number of objects expired:         0
02/22/05   17:31:34 Total number of objects failed:          1
02/22/05   17:31:34 Total number of bytes transferred:    696.29 MB
02/22/05   17:31:34 LanFree data bytes:                      0  B
02/22/05   17:31:34 Data transfer time:                  691.72 sec
02/22/05   17:31:34 Network data transfer rate:        1,030.76 KB/sec
02/22/05   17:31:34 Aggregate data transfer rate:        931.36 KB/sec
02/22/05   17:31:34 Objects compressed by:                   0%
02/22/05   17:31:34 Elapsed processing time:           00:12:45
02/22/05   17:31:34 --- SCHEDULEREC STATUS END
02/22/05   17:31:34 --- SCHEDULEREC OBJECT END TEST_SCHED 02/22/05   17:10:08
02/22/05   17:31:34 Scheduled event 'TEST_SCHED' completed successfully.
02/22/05   17:31:34 Sending results for scheduled event 'TEST_SCHED'.
```

Result summary

We are able to have the VCS cluster restarting an application with its backup environment up and running.

Locked resources are discovered and freed up.

Scheduled operation is restarted via by the scheduler and obtain back the previous resources.

There is the opportunity of having a backup restarted even if, considering a database as an example, this can lead to a backup window breakthrough, thus affecting other backup operations.

Archived

We run this test, at first using command line initiated backups with the same result; the only difference is that the operation needs to be restarted manually.

19.7.9 LAN-free client failover while the client is restoring

We will now do a client restore test, which is using LAN-free communications.

Objective

In this test we are verifying how a restore operation scenario is managed in a client takeover scenario.

For this test we will use a scheduled restore, which after the failover recovery, will re-start the restore operation which was interrupted. We use a scheduled operation with parameter **`replace=all`**, so the restore operation is restarted from beginning on restart, with no prompting.

If we were to use a manual restore with a command line (and wildcard), this would be restarted from the point of failure with the Tivoli Storage Manager client command **`restart restore`**.

Preparation

Here are the steps we follow for this test:

1. We verify that the cluster services are running with the **`hastatus`** command.
2. Then we schedule a restore with client node CL_VERITAS01_CLIENT association.

Example 19-42 Restore schedule

```
Day of Month:
                Week of Month:
                   Expiration:
Last Update by (administrator): ADMIN
         Last Update Date/Time: 02/21/05   10:26:04
              Managing profile:

            Policy Domain Name: STANDARD
               Schedule Name: RESTORE_TEST
                 Description:
                      Action: Restore
                     Options: -subdir=yes -replace=all
                     Objects: /opt/IBM/ISC/backup/*.*
                    Priority: 5
             Start Date/Time: 02/21/05   18:30:44
                    Duration: Indefinite
              Schedule Style: Classic
```

```
                          Period: One Time
                     Day of Week: Any
                           Month:
                    Day of Month:
                   Week of Month:
                      Expiration:
Last Update by (administrator): ADMIN
         Last Update Date/Time: 02/21/05   18:52:26
              Managing profile:
```

3. We wait for the client session to start and data beginning to be transferred to Banda, and finally session 8.645 shows data being sent to CL_VERITAS01_CLIENT, as seen in Example 19-43.

Example 19-43 Client restore sessions starting

```
8,644   Tcp/Ip   IdleW    1.9 M      1.6 K       722   Node    AIX        CL_VERITAS01_CLIENT
8,645   Tcp/Ip   SendW      0 S    152.9 M     1.0 K   Node    AIX        CL_VERITAS01_CLIENT
8,584   Tcp/Ip   IdleW     24 S      1.9 K     1.2 K   Serv-   AIX-RS/-   CL_VERITAS01_STA
8,587   Tcp/Ip   IdleW     24 S      7.4 K     4.5 K   Serv-   AIX-RS/-   CL_VERITAS01_STA
8,644   Tcp/Ip   IdleW    2.3 M      1.6 K       722   Node    AIX        CL_VERITAS01_CLIENT
8,645   Tcp/Ip   SendW     16 S    238.2 M     1.0 K   Node    AIX        CL_VERITAS01_CLIENT
8,648   Tcp/Ip   IdleW     19 S        257     1.0 K   Serv-   AIX-RS/-   CL_VERITAS01_STA
```

4. Also, we look for the input volume being mounted and opened for the restore, as seen in Example 19-44.

Example 19-44 Query the mounts looking for the restore data flow starting

```
tsm: TSMSRV03>q mount
ANR8330I LTO volume 030AKK is mounted R/W in drive DRLTO_1 (/dev/rmt0), status: IN USE.
ANR8334I         1 matches found.
```

Failure

Here are the steps we follow for this test:

1. Once satisfied that the client restore is running, we issue **halt -q** on the AIX server running the Tivoli Storage Manager client (Banda). The **halt -q** command stops AIX immediately and powers off the server.
2. Atlantic (the surviving node) is not yet receiving data after the failover, and we see from the Tivoli Storage Manager server that the current sessions remain in idlew and recvw states, as shown in Example 19-45.

Example 19-45 Query session command during the transition after failover of banda

```
8,644   Tcp/Ip   IdleW    1.9 M      1.6 K       722   Node    AIX        CL_VERITAS01_CLIENT
8,645   Tcp/Ip   SendW      0 S    152.9 M     1.0 K   Node    AIX        CL_VERITAS01_CLIENT
8,584   Tcp/Ip   IdleW     24 S      1.9 K     1.2 K   Serv-   AIX-RS/-   CL_VERITAS01_STA
8,587   Tcp/Ip   IdleW     24 S      7.4 K     4.5 K   Serv-   AIX-RS/-   CL_VERITAS01_STA
8,644   Tcp/Ip   IdleW    2.3 M      1.6 K       722   Node    AIX        CL_VERITAS01_CLIENT
8,645   Tcp/Ip   SendW     16 S    238.2 M     1.0 K   Node    AIX        CL_VERITAS01_CLIENT
8,648   Tcp/Ip   IdleW     19 S        257     1.0 K   Serv-   AIX-RS/-   CL_VERITAS01_STA
```

Recovery

Here are the steps we follow for this test:

1. Atlantic takes over the resources and launches the Tivoli Storage Manager start script.
2. We can see from the server console log in Example 19-46 which is showing the same events occurred in the backup test previously completed.
   a. The select searching for a tape holding session.
   b. The **`cancel`** command for the session found above.
   c. A new select with no result because the first **`cancel session`** command is successful.
   d. The restarted client scheduler querying for schedules.
   e. The schedule is still in the window, so a new restore operation is started and it obtains its input volume.

Example 19-46 The server log during restore restart

```
ANR0408I Session 8648 started for server CL_VERITAS01_STA (AIX-RS/6000) (Tcp/Ip) for event
logging.
ANR2017I Administrator ADMIN issued command: QUERY SESSION
ANR3605E Unable to communicate with storage agent.
ANR0482W Session 8621 for node RADON_STA (Windows) terminated - idle for more than 15 minutes.
ANR0408I Session 8649 started for server RADON_STA (Windows) (Tcp/Ip) for storage agent.
ANR0408I Session 8650 started for server CL_VERITAS01_STA (AIX-RS/6000) (Tcp/Ip) for storage
agent.
ANR0490I Canceling session 8584 for node CL_VERITAS01_STA (AIX-RS/6000) .
ANR3605E Unable to communicate with storage agent.
ANR0490I Canceling session 8587 for node CL_VERITAS01_STA (AIX-RS/6000) .
ANR3605E Unable to communicate with storage agent.
ANR0483W Session 8584 for node CL_VERITAS01_STA (AIX-RS/6000) terminated - forced by
administrator.
ANR0483W Session 8587 for node CL_VERITAS01_STA (AIX-RS/6000) terminated - forced by
administrator.
ANR0408I Session 8651 started for server CL_VERITAS01_STA (AIX-RS/6000) (Tcp/Ip) for library
sharing.
```

Archived

```
ANR0408I Session 8652 started for server CL_VERITAS01_STA (AIX-RS/6000) (Tcp/Ip) for event
logging.
ANR0409I Session 8651 ended for server CL_VERITAS01_STA (AIX-RS/6000).
ANR0408I Session 8653 started for server CL_VERITAS01_STA (AIX-RS/6000) (Tcp/Ip) for storage
agent.
ANR3605E Unable to communicate with storage agent.
ANR0407I Session 8655 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip9.1.39.42(33530)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select SESSION_ID,CLIENT_NAME from
SESSIONS where
CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 8655 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 8656 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip9.1.39.92(33531)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select SESSION_ID,CLIENT_NAME from
SESSIONS where
CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 8656 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 8657 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip9.1.39.42(33532)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select SESSION_ID,CLIENT_NAME from
SESSIONS where
CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 8657 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 8658 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip9.1.39.92(33533)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select SESSION_ID,CLIENT_NAME from
SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 8658 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 8659 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip9.1.39.42(33534)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select SESSION_ID,CLIENT_NAME from
SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 8659 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 8660 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip9.1.39.92(33535)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select SESSION_ID,CLIENT_NAME from
SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 8660 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 8661 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip9.1.39.42(33536)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select SESSION_ID,CLIENT_NAME from
SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 8661 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 8662 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip9.1.39.92(33537)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select SESSION_ID,CLIENT_NAME from
SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 8662 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 8663 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip9.1.39.42(33538)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select SESSION_ID,CLIENT_NAME from
SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 8663 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 8664 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip9.1.39.92(33539)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select SESSION_ID,CLIENT_NAME from
SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 8664 ended for administrator SCRIPT_OPERATOR (AIX).
```

```
ANR0407I Session 8665 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip9.1.39.42(33540)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select SESSION_ID,CLIENT_NAME from
SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 8665 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 8666 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip9.1.39.92(33541)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select SESSION_ID,CLIENT_NAME from
SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 8666 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 8667 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip9.1.39.42(33542)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select SESSION_ID,CLIENT_NAME from
SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 8667 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 8668 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip9.1.39.92(33543)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select SESSION_ID,CLIENT_NAME from
SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 8668 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 8669 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip9.1.39.42(33544)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select SESSION_ID,CLIENT_NAME from
SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 8669 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0407I Session 8670 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip9.1.39.92(33545)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select SESSION_ID,CLIENT_NAME from
SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 8670 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0406I Session 8671 started for node CL_VERITAS01_CLIENT (AIX) (Tcp/Ip 9.1.39.42(33547)).
ANR1639I Attributes changed for node CL_VERITAS01_CLIENT: TCP Name from banda to atlantic, GUID
from 00.00.00.00.75.8e.11.d9.ac.29.08.63.09.01.27.5e to
00.00.00.01.75.8f.11.d9.b4.d1.08.63.09.01.27.5c.
ANR0408I Session 8672 started for server CL_VERITAS01_STA (AIX-RS/6000) (Tcp/Ip) for storage
agent.
ANR0415I Session 8672 proxied by CL_VERITAS01_STA started for node CL_VERITAS01_CLIENT.
```

3. We then see a new session appear in MediaW (8,672), which will take over the restore data send from the original session 8.645, which is still in SendW status, as seen in Example 19-47.

Example 19-47 Addition restore session begins, completes restore after the failover

```
8,644   Tcp/Ip   IdleW    4.5 M      1.6 K       722   Node    AIX        CL_VERITAS01_CLIENT
8,645   Tcp/Ip   SendW    2.5 M    238.2 M     1.0 K   Node    AIX        CL_VERITAS01_CLIENT
8,648   Tcp/Ip   IdleW    2.5 M        257     1.0 K   Serv-   AIX-RS/-   CL_VERITAS01_STA
8,650   Tcp/Ip   IdleW      4 S      1.3 K       678   Serv-   AIX-RS/-   CL_VERITAS01_STA
8,652   Tcp/Ip   IdleW     34 S        257     1.8 K   Serv-   AIX-RS/-   CL_VERITAS01_STA
8,653   Tcp/Ip   IdleW      4 S      4.3 K     3.4 K   Serv-   AIX-RS/-   CL_VERITAS01_STA
8,671   Tcp/Ip   IdleW     34 S      1.6 K       725   Node    AIX        CL_VERITAS01_CLIENT
8,672   Tcp/Ip   MediaW    34 S      1.5 K     1.0 K   Node    AIX        CL_VERITAS01_CLIENT
```

4. We then view the transition point for the end and then restart in the dsmsched.log on the client, as seen in Example 19-48.

Example 19-48 dsmsched.log output demonstrating the failure and restart transition

```
------------------------------------------------------------
Schedule Name:         RESTORE
Action:                Restore
Objects:               /opt/IBM/ISC/backup/*.*
Options:               -subdir=yes -replace=all
Server Window Start:   11:30:00 on 02/23/05
------------------------------------------------------------
Executing scheduled command now.
--- SCHEDULEREC OBJECT BEGIN RESTORE 02/23/05   11:30:00
Restore function invoked.

** Interrupted **
ANS1114I Waiting for mount of offline media.
Restoring   1,034,141,696 /opt/IBM/ISC/backup/520005.tar [Done]
Restoring   1,034,141,696 /opt/IBM/ISC/backup/520005.tar [Done]
** Interrupted **
ANS1114I Waiting for mount of offline media.
Restoring     403,398,656 /opt/IBM/ISC/backup/VCS_TSM_package.tar [Done]
Restoring     403,398,656 /opt/IBM/ISC/backup/VCS_TSM_package.tar [Done]
```

5. Next, we review the Tivoli Storage Manager server sessions, as seen in Example 19-49.

Example 19-49 Server sessions after the restart of the restore operation.

```
8,644   Tcp/Ip   IdleW    12.8 M     1.6 K      722   Node    AIX        CL_VERITAS01_CLIENT
8,648   Tcp/Ip   IdleW    10.8 M      257     1.0 K   Serv-   AIX-RS/-   CL_VERITAS01_STA
8,650   Tcp/Ip   IdleW     2 S       1.5 K      810   Serv-   AIX-RS/-   CL_VERITAS01_STA
8,652   Tcp/Ip   IdleW    8.8 M      257     1.8 K   Serv-   AIX-RS/-   CL_VERITAS01_STA
8,653   Tcp/Ip   IdleW     2 S       5.0 K     3.6 K   Serv-   AIX-RS/-   CL_VERITAS01_STA
8,671   Tcp/Ip   IdleW    8.8 M      1.6 K      725   Node    AIX        CL_VERITAS01_CLIENT
8,672   Tcp/Ip   SendW     0 S     777.0 M     1.0 K   Node    AIX        CL_VERITAS01_CLIENT
```

6. The new restore operation completes successfully, as we confirm in the Client log, as shown in Example 19-50.

Archived

Example 19-50 dsmsched.log output of completed summary of failover restore test

```
--- SCHEDULEREC STATUS BEGIN
Total number of objects restored:         4
Total number of objects failed:           0
Total number of bytes transferred:     1.33 GB
LanFree data bytes:                    1.33 GB
Data transfer time:                  114.55 sec
Network data transfer rate:        12,256.55 KB/sec
Aggregate data transfer rate:      2,219.52 KB/sec
Elapsed processing time:           00:10:32
--- SCHEDULEREC STATUS END
--- SCHEDULEREC OBJECT END RESTORE 02/23/05   11:30:00
```

Result summary

The cluster is able to manage the client failure and make Tivoli Storage Manager client scheduler available in about 1 minute. The client is able to restart its operations successfully to the end (although the actual session numbers will change, there is no user intervention required).

Since this is a scheduled restore with **`replace=all`**, it is restarted from the beginning and completes successfully, overwriting the previously restored data.

Archived

Archived

20

VERITAS Cluster Server on AIX with IBM Tivoli Storage Manager Client and ISC applications

This chapter provides details about the configuration of the Veritas Cluster Server, including the configuration of the Tivoli Storage Manager client as a highly available application. We also include the Integrated Support Console as a highly available application.

© Copyright IBM Corp. 2005. All rights reserved.

20.1 Overview

We will prepare the environments prior to configuring these applications in the VCS cluster. All Tivoli Storage Manager components must communicate properly prior to HA configuration, including the products installed on the cluster shared disks.

VCS will require start, stop, monitor and clean scripts for most of the applications. Creating and testing these prior to implementing the Service Group configuration is a good approach.

20.2 Planning

There must be a requirement to configure a highly available Tivoli Storage Manager client. The most common requirement would be an application, such as a database product that has been configured and running under VCS control. In such cases, the Tivoli Storage Manager client will be configured within the same Service Group as an application. This ensures that the Tivoli Storage Manager client is tightly coupled with the application which requires backup and recovery services.

Table 20-1 Tivoli Storage Manager client configuration

Node name	Node directory	TCP/IP address	TCP/IP port
atlantic	/usr/tivoli/tsm/client/ba/bin	9.1.39.92	1501
banda	/usr/tivoli/tsm/client/ba/bin	9.1.39.94	1501
cl_veritas01_client	/opt/IBM/ISC/tsm/client/ba/bin	9.1.39.77	1502

For the purposes of this setup exercise, we will install the Integrated Solutions Console (ISC) and the Tivoli Storage Manager Administration Center onto the shared disk (simulating a client application). This feature, which is used for Tivoli Storage Manager administration, will become a highly available application, along with the Tivoli Storage Manager client.

The ISC was not designed with high availability in mind, and installation of this product on a shared disk, as a highly available application, is not officially supported, but is certainly possible. Another important note about the ISC is that its database must be backed up with the product offline to ensure database consistency. Refer to the ISC documentation for specific backup and recovery instructions.

Archived

20.3 Tivoli Storage Manager client installation

We installed the client software locally on both nodes in the previous chapter, shown in 18.2.3, “Tivoli Storage Manager Client Installation” on page 745.

20.3.1 Preparing the client for high availability

Now we will configure the Tivoli Storage Manager client for to be a virtual node.

1. First, we copy the dsm.opt file onto the shared disk location in which we will store our Tivoli Storage Manager client, storage agent, /opt/IBM/ISC/tsm/client/ba/bin.
2. Edit the /opt/IBM/ISC/tsm/client/ba/bin/dsm.opt file to reflect the servername which it will contact. For this purpose, tsmsrv06 will be the server we will connect to, as shown in Example 20-1.

Example 20-1 /opt/IBM/ISC/tsm/client/ba/bin/dsm.opt file content

```
banda:/opt/IBM/ISC/tsm/client/ba/bin# more dsm.opt
servername tsmsrv06
```

3. Next, we edit the /usr/tivoli/tsm/client/ba/bin/dsm.sys file and create the stanza which links the dsm.opt file shown in Example 20-1 and the dsm.sys file stanza shown in Example 20-2.

Example 20-2 /usr/tivoli/tsm/client/ba/bin/dsm.sys stanza, links clustered dsm.opt file

```
banda:/opt/IBM/ISC/tsm/client/ba/bin# grep -p tsmsrv06
/usr/tivoli/tsm/client/ba/bin/dsm.sys
* Server stanza for Win2003 server connection purpose
SErvername                      tsmsrv06
nodename                         cl_veritas01_client
   COMMMethod                    TCPip
   TCPPort                       1500
   TCPServeraddress              9.1.39.47
   ERRORLOGRETENTION             7
   ERRORLOGname                  /opt/IBM/ISC/tsm/client/ba/bin/dsmerror.log
   passworddir                   /opt/IBM/ISC/tsm/client/ba/bin/banda
   passwordaccess                generate
   managedservices              schedule webclient
   inclexcl                      /opt/IBM/ISC/tsm/client/ba/bin/inclexcl.lst
```

4. Then we ensure the changed (dsm.sys) file is copied (or ftp’d) over the other node (Atlantic in this case).same on both nodes on their local disks, with the exception of the passworddir for the highly available client, which will point to its own directory on the shared disk as shown in Example 20-3.

Archived

Example 20-3 The path and file difference for the passworddir option

```
banda:/opt/local/isc# grep passworddir /usr/tivoli/tsm/client/ba/bin/dsm.sys
   passworddir                  /opt/IBM/ISC/tsm/client/ba/bin/banda

atlantic:/# grep passworddir /usr/tivoli/tsm/client/ba/bin/dsm.sys
   passworddir          /opt/IBM/ISC/tsm/client/ba/bin/atlantic
```

5. Next, we set the password with the server, on each node one at a time, and verify the connection and authentication.

Tip: We have the TSM.PWD file written on the shared disk, in a separate directory for each physical node. Essentially there will be four Tivoli Storage Manager client passwords in use, one for each node's local backups (TSM.PWD is written to the default location), and one for each node's high availability backup. The reason for this is that the option clusternode=yes does not support VCS, only MSCS and HACMP.

20.4 Installing the ISC and the Administration Center

The installation of Tivoli Storage Manager Administration Center is a two-step install. First install the Integrated Solutions Console. Then deploy the Tivoli Storage Manager Administration Center into the Integrated Solutions Console. Once both pieces are installed, you will be able to administer Tivoli Storage Manager from a browser any where in your network.

In addition, these two software components will be contained within a Service Group, Application Resource within our VCS cluster. To achieve this, these software packages will be installed onto shared disk, and on the second node in the Tivoli Storage Manager cluster. This will make this cluster configuration an active/active configuration.

Integrated Solutions Console installation

We will install the Integrated Solutions Console (ISC) onto our shared disk resource. This is not a supported configuration, however, based on the design of the ISC, it is currently the only way to make this software product highly available.

Why make the ISC highly available?

This console has been positioned as a central point of access for the management of Tivoli Storage Manager servers. Prior to the ISC/AC introduction, the access of the Tivoli Storage Manager server was through the server itself. Now, with the exception of the administration command line, the ISC/AC is the only control method.

Given this, there may be many Tivoli Storage Manager servers (10's or 100's) accessed using this single console. All Tivoli Storage Manager server tasks, including adding, updating, and health checking (monitoring) is performed using this facility.

This single point of failure (access failure), leads our team to include the ISC and AC into our HA application configurations. Now, we will install and configure the ISC, as shown in the following steps:

1. First we extract the contents of the file TSM_ISC_5300_AIX.tar as shown in Example 20-4.

Example 20-4 The tar command extraction

```
tar xvf TSM_ISC_5300_AIX.tar
```

2. Then we change directory into iscinstall and run the **setupISC** InstallShield command, as shown in Example 20-5.

Example 20-5 Integrated Solutions Console installation script

```
banda:/install/ISC/# setupISC
```

Note: Depending on what the screen and graphics requirements would be, the following options exist for this installation.

Run one of the following commands to install the runtime:

- For InstallShield wizard install, run:

```
setupISC
```

- For console wizard install, run:

```
setupISC -console
```

- For silent install, run the following command on a single line:

```
setupISC -silent -W ConfigInput.adminName="<user name>"
```

Flags:

```
-W ConfigInput.adminPass="<user password>"
-W ConfigInput.verifyPass="<user password>"
-W PortInput.webAdminPort="<web administration port>"
-W PortInput.secureAdminPort="<secure administration port>"
-W MediaLocationInput.installMediaLocation="<media location>"
-P ISCProduct.installLocation="<install location>"
```

3. Then, we follow the Java based installation process, as shown in Figure 20-1. This is the introduction screen, in which we click **Next**.

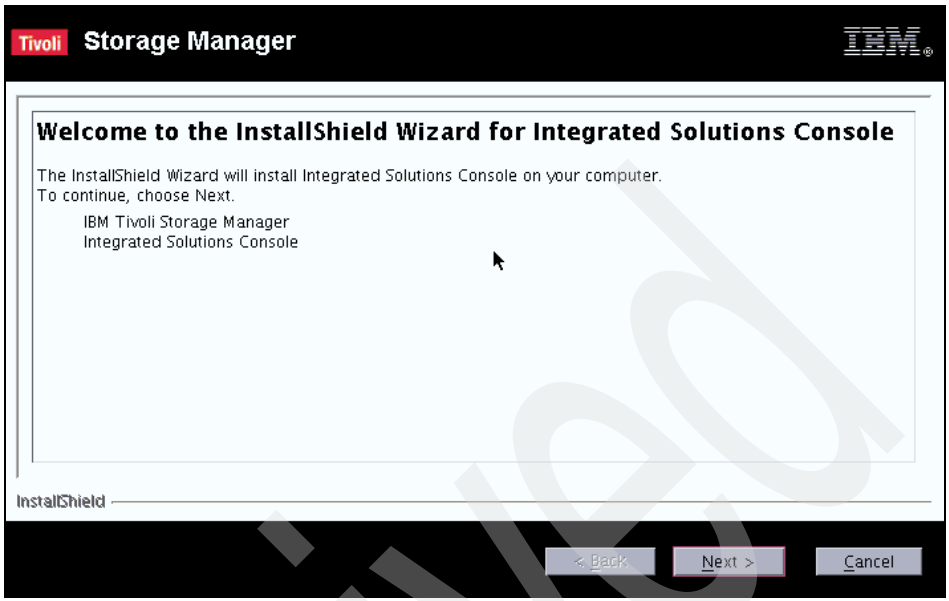

Figure 20-1 ISC installation screen

4. We review the licensing details, then click **Next**, as shown in Figure 20-2.

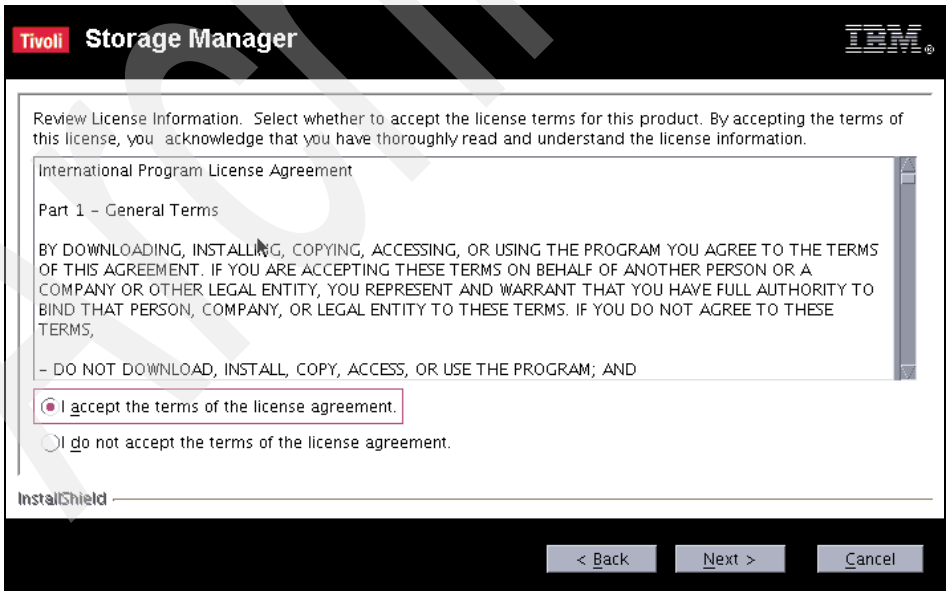

Figure 20-2 ISC installation screen, license agreement

Archived

5. This is followed by the location of the source files, which we verify and click **Next** as shown in Figure 20-3.

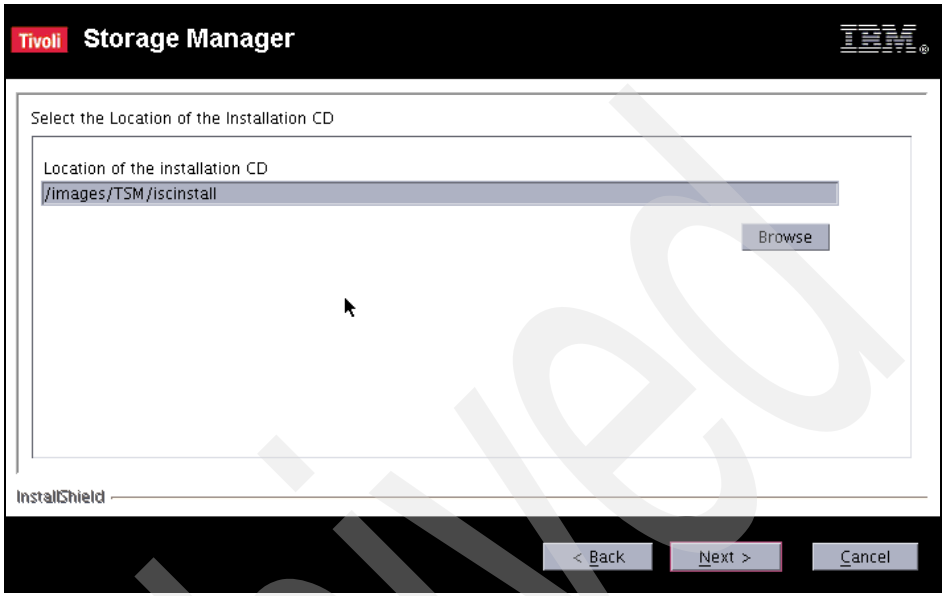

Figure 20-3 ISC installation screen, source path

6. Then, at this point, we ensure that the VG iscvg is online and the /opt/IBM/ISC is mounted. Then, we type in our target path and click **Next**, as shown in Figure 20-4.

Archived

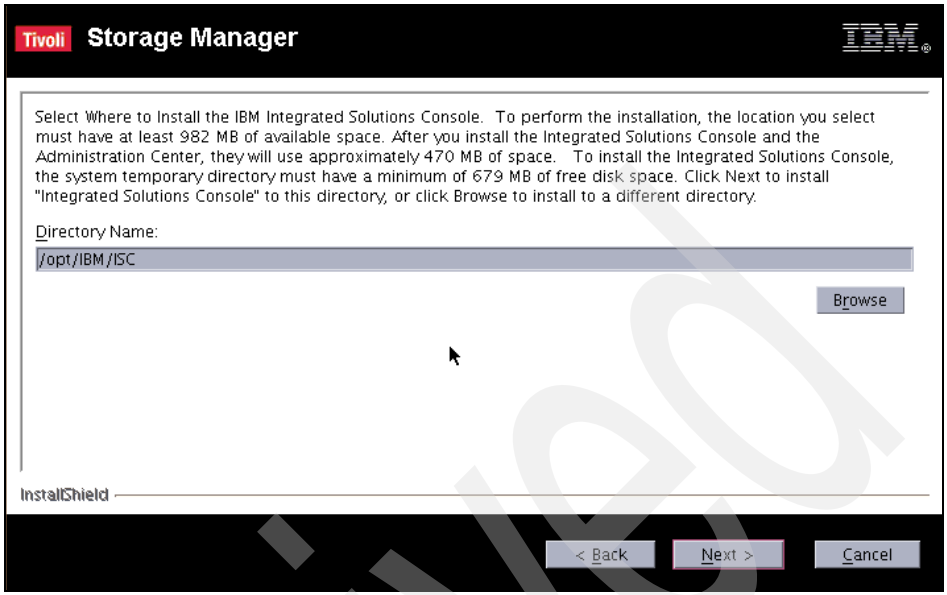

Figure 20-4 ISC installation screen, target path - our shared disk for this node

7. Next, we establish our userID and password to log into the ISC once the installation is complete. We fill in the details and click **Next**, as shown in Figure 20-5.

Archived

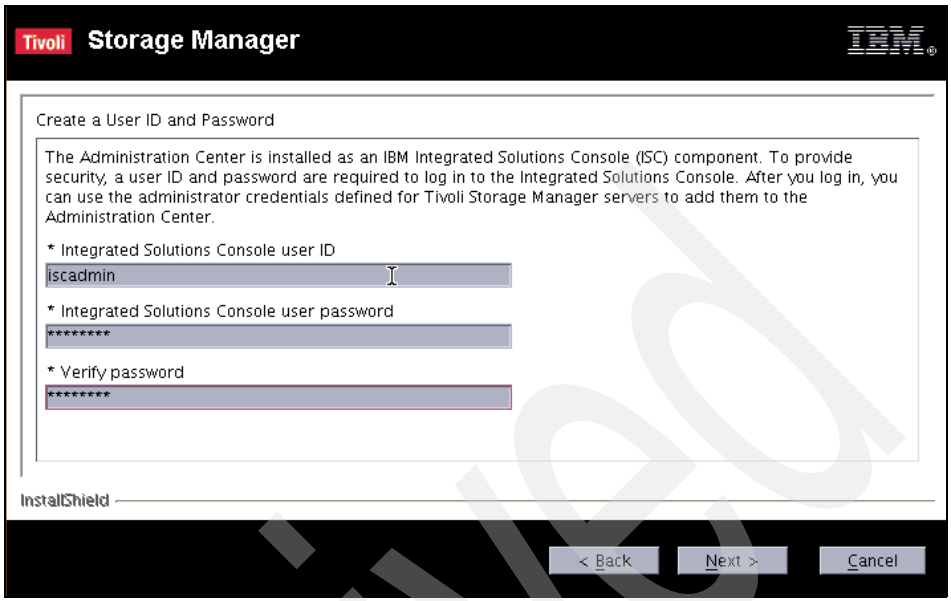

Figure 20-5 ISC installation screen, establishing a login and password

8. Next, we then select the HTTP ports, which we leave as the default and click **Next**, as shown in Figure 20-6.

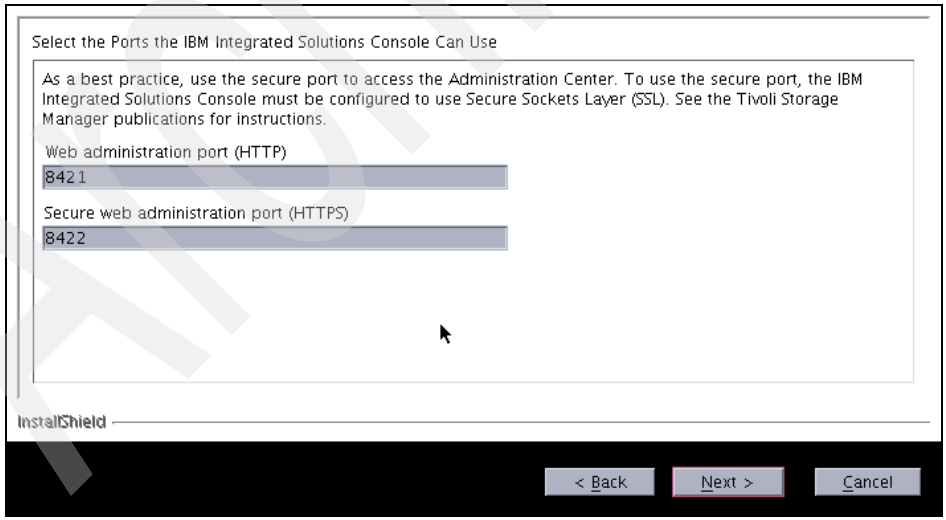

Figure 20-6 ISC installation screen establishing the ports which will be used

Archived

9. We now review the installation selections and the space requirements, then click **Next** as shown in Figure 20-7.

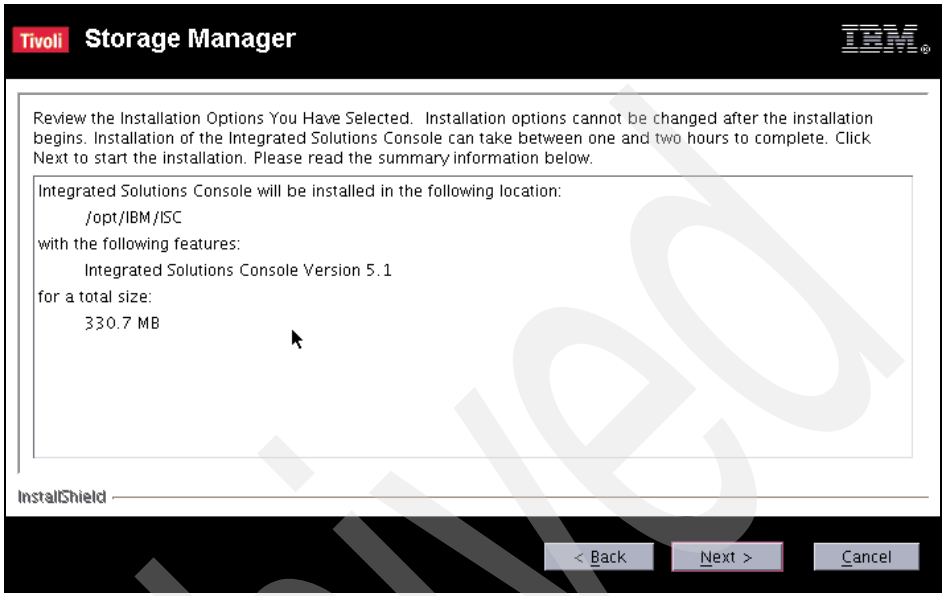

Figure 20-7 ISC installation screen, reviewing selections and disk space required

10. We then review the summary of the successful completion of the installation, and click **Next** to continue, as shown in Figure 20-7.

Archived

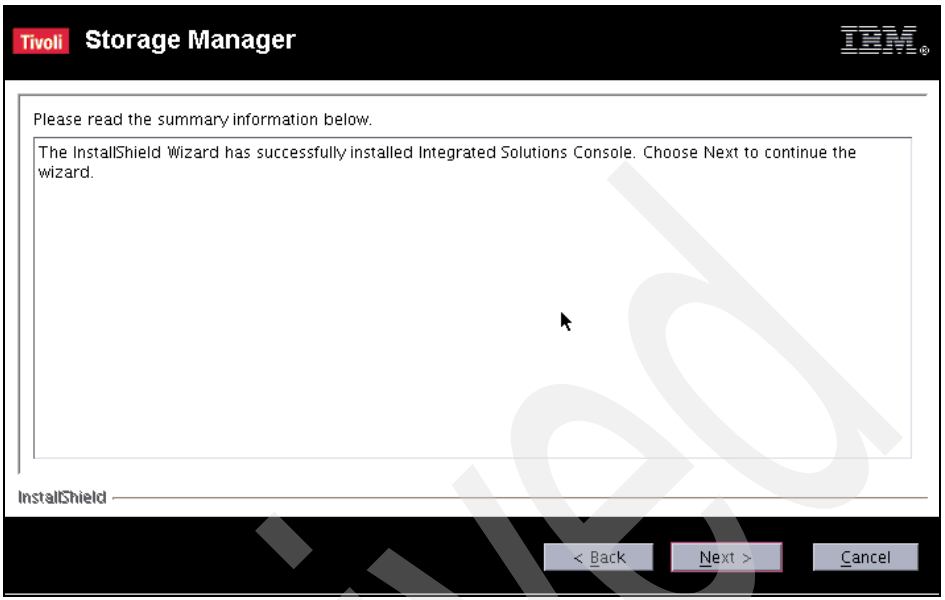

Figure 20-8 ISC installation screen showing completion

11. The final screen appears now, and we select **Done**, as shown in Figure 20-9.

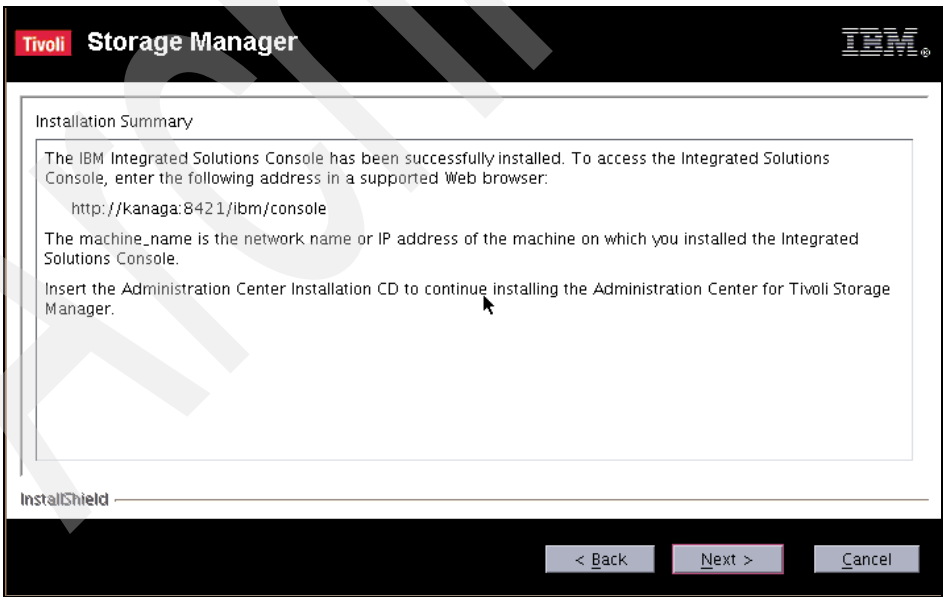

Figure 20-9 ISC installation screen, final summary providing URL for connection

Archived

Tivoli Storage Manager Administration Center

1. First, we start by reviewing the contents of the Administration Center installation directory, as seen in Example 20-6.

Example 20-6 Administration Center install directory

```
Atlantic:/install/TSM/AdminCenter/acinstall# ls -l
total 139744
-rw-r-----   1 501      300         7513480 Nov 29 17:30 AdminCenter.war
-rw-r--r--   1 501      300         6481802 Nov 11 17:09 ISCAction.jar
drwxr-xr-x   2 501      300             256 Nov 02 09:06 META-INF
-rw-r--r--   1 501      300            6795 Nov 29 17:30 README
-rw-r-----   1 501      300           15978 Nov 23 08:26 README.INSTALL
drwxr-xr-x   3 501      300             256 Nov 29 17:30 Tivoli
-rw-r--r--   1 501      300           18266 Nov 29 17:30 dsminstall.jar
-rw-r--r--   1 501      300        22052682 Nov 29 17:30 help.jar
drwxr-xr-x   2 501      300             256 Nov 29 17:30 jacl
-rw-r-----   1 501      300           79853 Nov 11 17:18 license.txt
-rw-r--r--   1 501      300              13 Oct 21 18:01 media.inf
-rwxr-xr-x   1 501      300        35355831 Nov 11 17:18 setupAC
drwxr-xr-x   2 501      300             256 Nov 29 17:30 shared
-rw-r-----   1 501      300             152 Nov 01 14:17 startInstall.bat
-rwxr-xr-x   1 501      300             647 Nov 23 07:56 startInstall.sh
```

2. We then review the readme files prior to running the install script.
3. Then, we issue the **startInstall.sh** command, which spawns the following Java screens.
4. The first screen is an introduction, and we click **Next**, as seen in Figure 20-10.

Archived

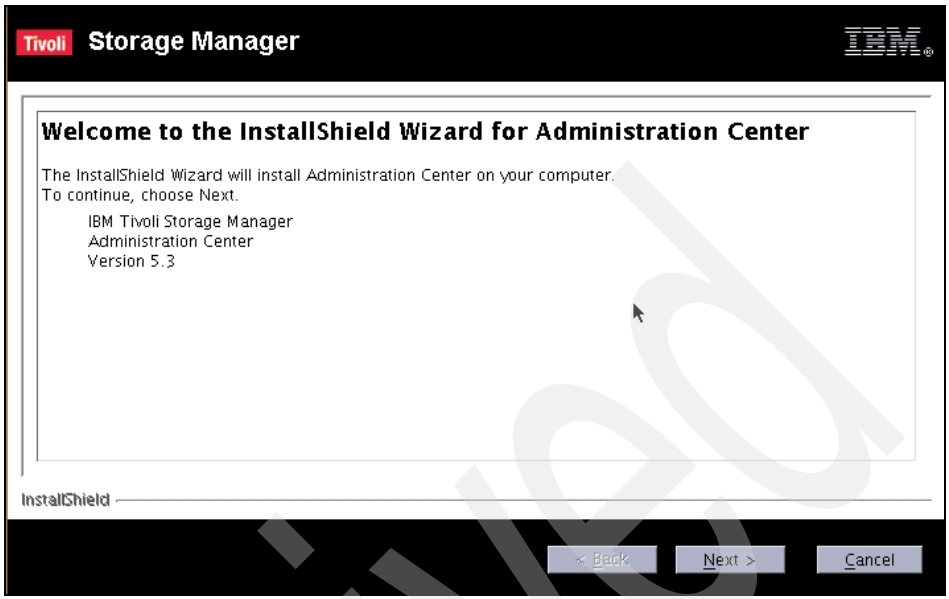

Figure 20-10 Welcome wizard screen

5. Next, we get a panel giving the space requirements, and we click **Next**, as shown in Figure 20-11.

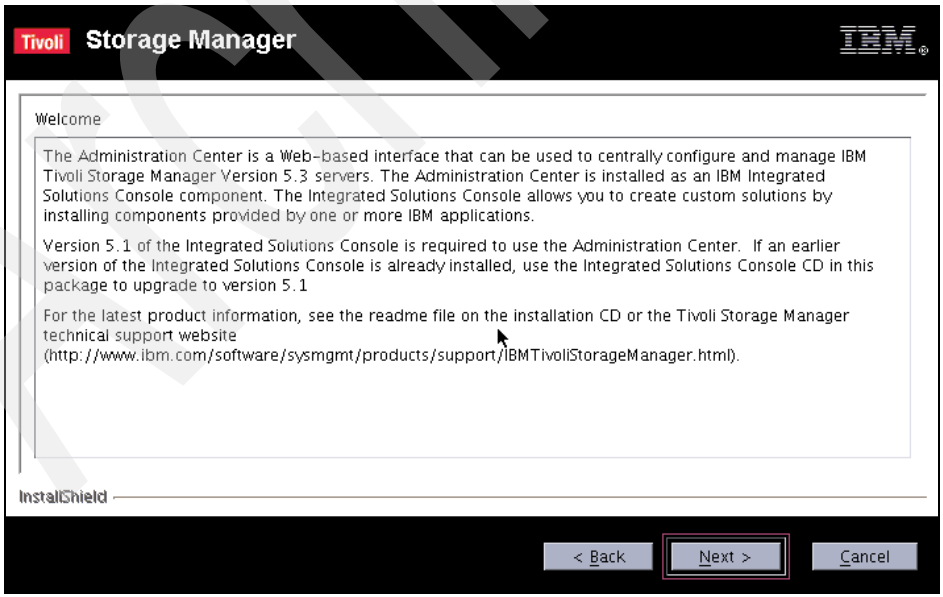

Figure 20-11 Review of AC purpose and requirements

Archived

6. We then accept the terms of the license and click **Next**, as shown in Figure 20-12.

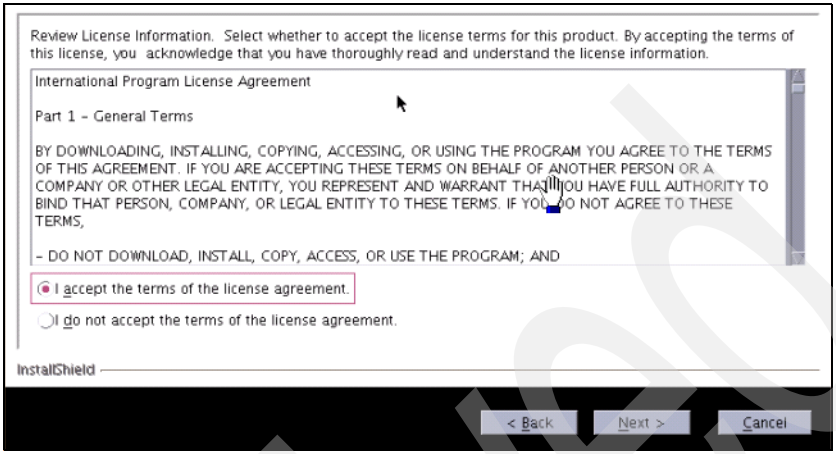

Figure 20-12 AC Licensing panel

7. Next, we validate the ISC installation environment, check that the information is correct, then click **Next**, as seen in Figure 20-13.

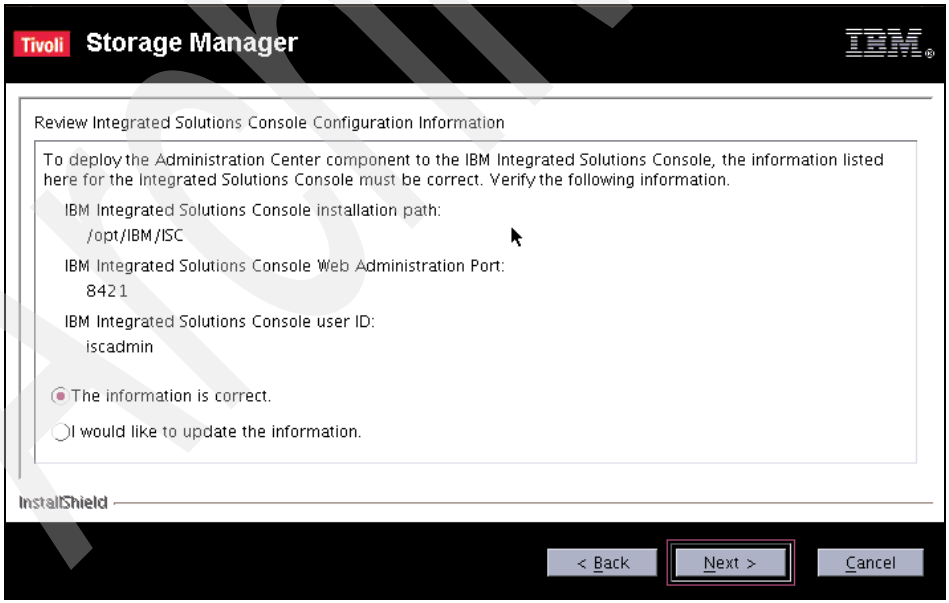

Figure 20-13 Validation of the ISC installation environment

8. Next, we are prompted for the ISC userid and password and then click **Next**, as shown in Figure 20-14.

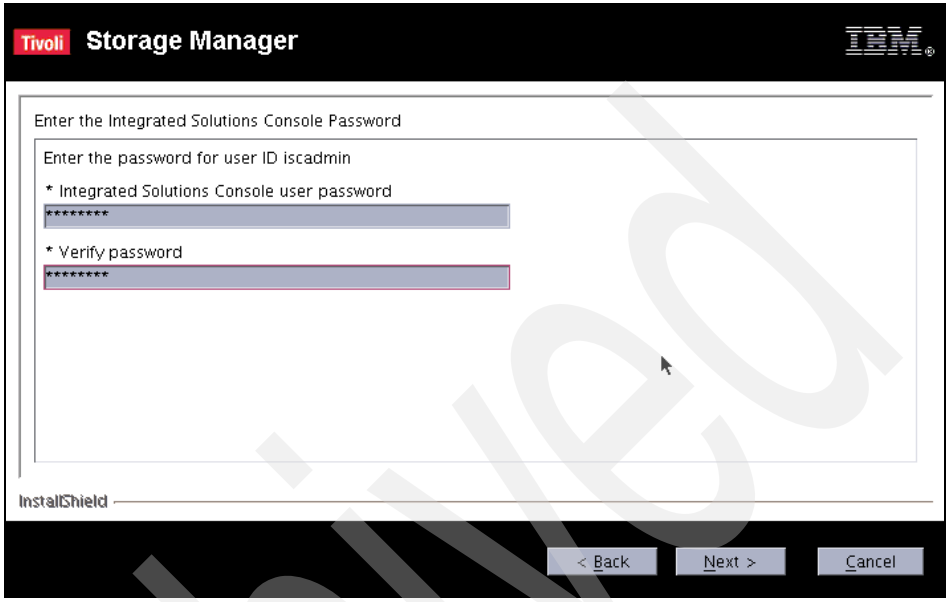

Figure 20-14 Prompting for the ISC userid and password

9. The we confirm the AC installation directory, as shown in Figure 20-15.

Archived

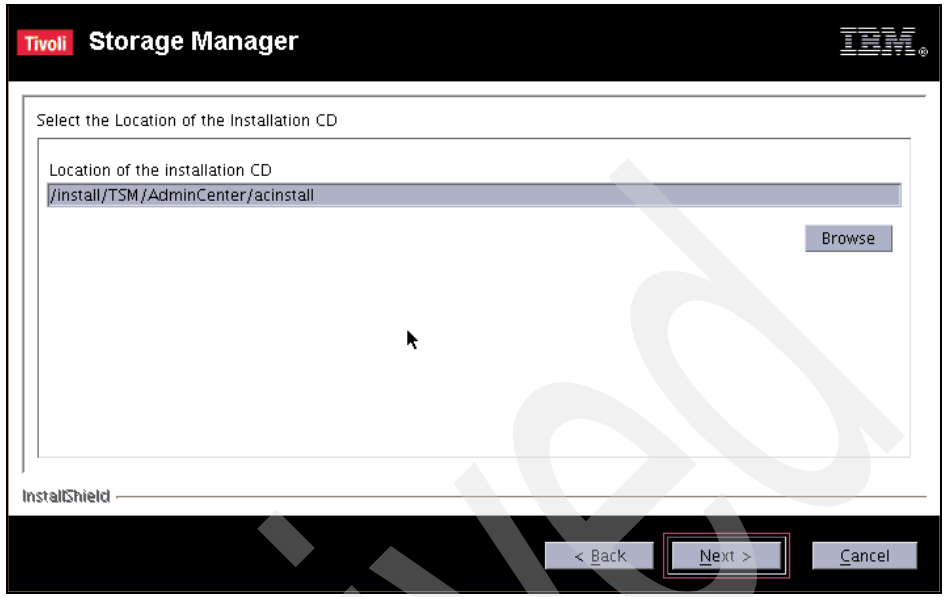

Figure 20-15 AC installation source directory

10. We then confirm the installation directory and required space, and click **Next** as shown in Figure 20-16.

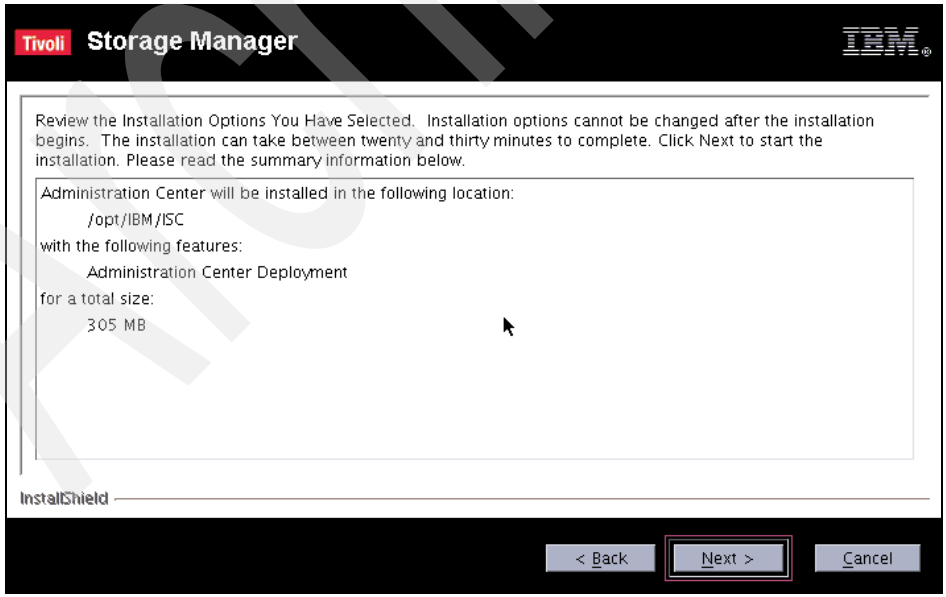

Figure 20-16 AC target source directory

Archived

11. We see the installation progression screen, shown in Figure 20-17.

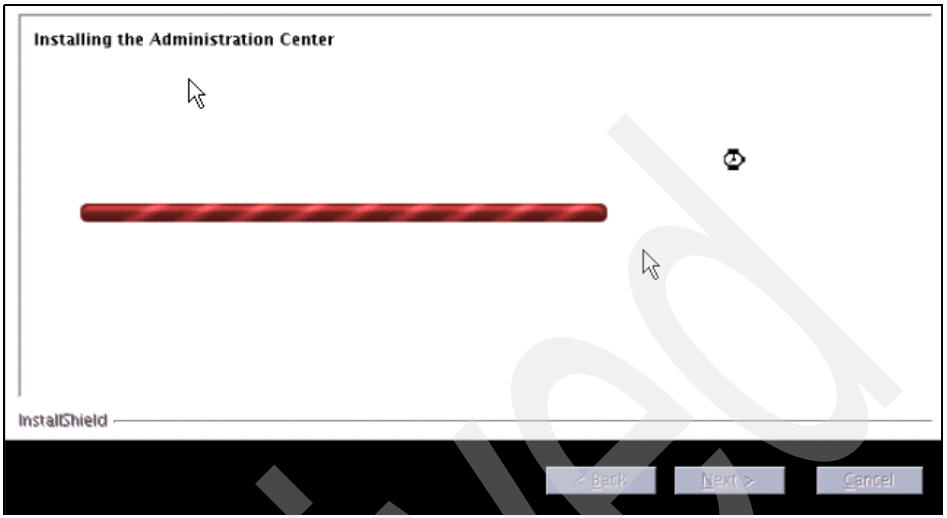

Figure 20-17 AC progress screen

12. Next, a successful completion screen appears, as shown in Figure 20-18.

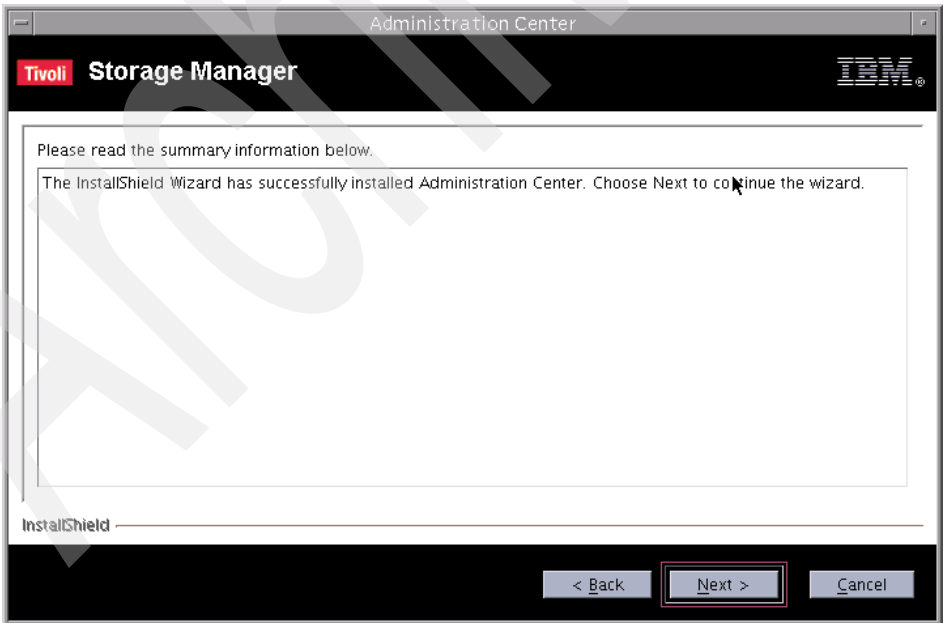

Figure 20-18 AC successful completion

Archived

13. We get a summary of the installation, which includes the URL with port, shown Figure 20-19.

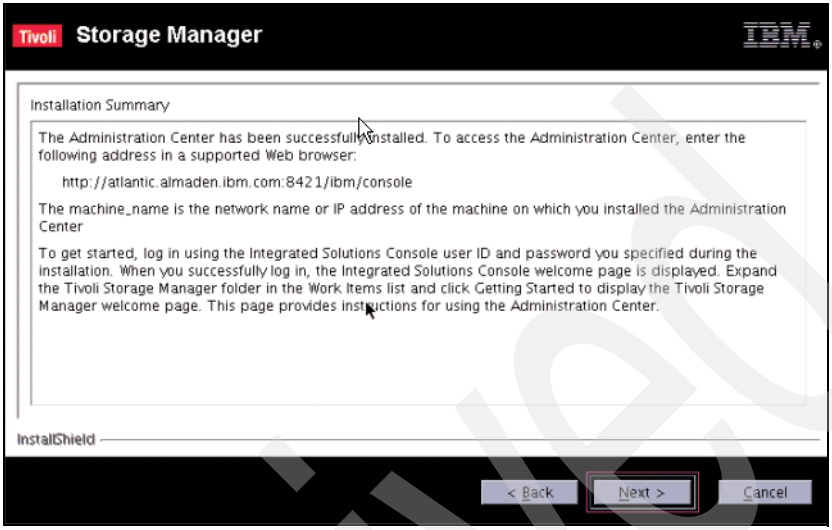

Figure 20-19 Summary and review of the port and URL to access the AC

14. Finally, we click **Finish** to complete the installation as shown in Figure 20-20.

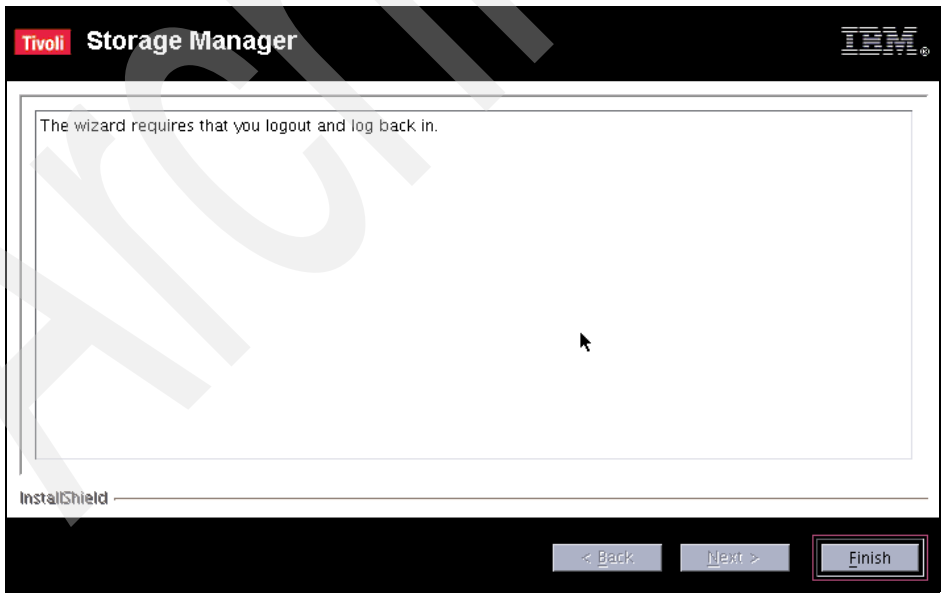

Figure 20-20 Final AC screen

Archived

20.5 Veritas Cluster Manager configuration

The installation process configured the core cluster services for us, now we need to configure the Service Groups and their associated resources for the Tivoli Storage Manager client and the ISC.

20.5.1 Preparing and placing application startup scripts

We will develop and test our start, stop, clean, and monitor scripts for all of our applications, then place them in the /opt/local directory on each node, which is a local filesystem within the rootvg.

Scripts for the client CAD

We placed the scripts for the server in the rootvg, /opt filesystem, in the directory /opt/local/tsmcli.

1. The start script /opt/local/tsmcli/startTSMcli.sh is shown in Example 20-7.

Example 20-7 /opt/local/tsmcli/startTSMcli.sh

```
#!/bin/ksh
set -x
##############################################################################
# Tivoli Storage Manager                                                     *
#                                                                            *
##############################################################################
#
#   The start script is used in the following cases:
#   1. when HACMP is started and resource groups are "activated"
#   2. when a failover occurs and the resource group is started on another node
#   3. when fallback occurs (a failed node re-enters the cluster) and the
#      resource group is transferred back to the node re-entering the cluster.
#
#   Name: StartClusterTsmclient.sh
#
#   Function: A sample shell script to start the client acceptor daemon (CAD)
#   for the TSM Backup-Archive Client. The client system options file must be
#   configured (using the MANAGEDSERVICES option) to allow the CAD to manage
#   the client scheduler. HACMPDIR can be specified as an environment variable.
#   The default HACMPDIR is /ha_mnt1/tsmshr
#
##############################################################################

if [[ $VERBOSE_LOGGING = "high" ]]
then
   set -x
fi
```

```
#Set the name of this script.
myname=${0##*/}

#Set the hostname for the HADIR
hostname=`hostname`

# Set default HACMP DIRECTORY if environment variable not present
if [[ $HADIR = "" ]]
then
     HADIR=/opt/IBM/ISC/tsm/client/ba/bin/$hostname
fi

PIDFILE=$HADIR/hacad.pids
#export DSM variables
export DSM_DIR=/usr/tivoli/tsm/client/ba/bin
export DSM_CONFIG=$HADIR/dsm.opt

##############################################
# Function definitions.
##############################################

function CLEAN_EXIT
{

#There should be only one process id in this file
#if more than one cad, then display error message

wc $HADIR/hacad.pids |awk '{print $2}' >$INP |
if [[ $INP > 1 ]]
then
        msg_p1="WARNING: Unable to determine HACMP CAD"

else
        msg_p1="HACMP CAD process successfully logged in the pidfile"

fi
    print "$myname: Start script completed.  $msg_p1"
    exit 0
}

#Create a function to first start the cad and then capture the cad pid in a
file

START_CAD()
{

#Capture the process ids of all CAD processes on the system
        ps -ae |grep dsmcad | awk '{print $1}' >$HADIR/hacad.pids1
```

```
#Start the client accepter daemon in the background
        nohup $DSM_DIR/dsmcad &

#wait for 3 seconds for true cad daemon to start
        sleep 3

#Capture the process ids of all CAD processes on the system
        ps -ae |grep dsmcad | awk '{print $1}' >$HADIR/hacad.pids2

#Get the HACMP cad from the list of cads on the system
        diff $HADIR/hacad.pids1 $HADIR/hacad.pids2 |grep ">" |awk '{print$2}'
>$PIDFILE

}

# Now invoke the above function to start the Client Accepter Daemon (CAD)
# to allow connections from the web client interface

START_CAD

#Display exit status
CLEAN_EXIT

exit
```

2. We then place the stop script in the directory as /opt/local/tsmcli/stopTSMcli.sh, shown in Example 20-8.

Example 20-8 /opt/local/tsmcli/stopTSMcli.sh

```
#!/bin/ksh
##############################################################################
# Tivoli Storage Manager                                                     *
#                                                                            *
##############################################################################
#
#   The stop script is used in the following situations
#   1. When HACMP is stopped
#   2. When a failover occurs due to a failure of one component of the resource
#      groups, the other members are stopped so that the entire group can be
#      restarted on the target node in the failover
#   3. When a fallback occurs and the resource group is stopped on the node
#      currently hosting it to allow transfer back to the node re-entering the
#      cluster.
#
#   Name: StopClusterTsmclient.sh
#
#   Function: A sample shell script to stop the client acceptor daemon (CAD)
#   and all other processes started by CAD for the TSM Backup-Archive Client.
```

```
#   The client system options file must be configured (using the
#   MANAGEDSERVICES option) to allow the CAD to manage the client scheduler.
#   HADIR can be specified as an environment variable. The default HADIR is
#   /ha_mnt1/tsmshr This variable must be customized.
#
##############################################################################
#!/bin/ksh

if [[ $VERBOSE_LOGGING = "high" ]]
then
   set -x
fi

#Set the name of this script.
myname=${0##*/}

#Set the hostname for the HADIR
hostname=`hostname`

# Set default HACMP DIRECTORY if environment variable not present
if [[ $HADIR = "" ]]
then
     HADIR=/opt/IBM/ISC/tsm/client/ba/bin/$hostname
fi

PIDFILE=$HADIR/hacad.pids
CPIDFILE=$HADIR/hacmp.cpids

#export DSM variables
export DSM_DIR=/usr/tivoli/tsm/client/ba/bin
export DSM_CONFIG=$HADIR/dsm.opt

#define some local variables
final_rc=0;

##############################################
# Function definitions.
##############################################

# Exit function
function CLEAN_EXIT
{

    # Display final message
    if (( $final_rc==0 ))
    then
       # remove pid file.
       if [[ -a $PIDFILE ]]
       then
```

```
            rm $PIDFILE
         fi
         # remove cpid file.
         if [[ -a $CPIDFILE ]]
         then
            rm $CPIDFILE
         fi

         msg_p1="$pid successfully deleted"
      else
          msg_p1="HACMP stop script failed "
      fi
      print "$myname: Processing completed.  $msg_p1"
      exit $final_rc
}

function bad_pidfile
{
    print "$myname: pid file not found or not readable $PIDFILE"
    final_rc=1
    CLEAN_EXIT
}

function bad_cpidfile
{
    print "$myname: cpid file not readable $CPIDFILE"
    final_rc=2
    CLEAN_EXIT
}

function validate_pid
{
   #There should be only one process id in this file
   #if more than one cad, then exit

   wc $HADIR/hacad.pids |awk '{print $2}' >$INP |
   if [[ $INP > 1 ]]
   then
        print "$myname: Unable to determine HACMP CAD"
        final_rc=3
        clean_exit
   fi
}

# Function to read/kill child processes
function kill_child
{
   # If cpid file exists, is not empty, and is not readable then
   # display error message
```

```
   if [[ -s $PIDFILE ]] && [[ ! -r $PIDFILE ]]
   then
      bad_cpidfile
   fi

   # delete child processes
   while read -r cpid;
   do
      kill -9 $cpid
   done <$CPIDFILE
}

# Function to read/kill CAD process and get child processes
function read_pid
{
   while read -r pid;
   do
      # Get all child processes of HACMP CAD
      ps -ef |grep $pid | awk '{print $2}' >/$CPIDFILE

      # Kill any child processes
      kill_child

      # Kill HACMP CAD
      kill -9 $pid

   done <$PIDFILE
   final_rc=0
}

# Main function
function CAD_STOP
{
   # Check if pid file exists, is not empty, and is readable
   if [[ ! -s $PIDFILE ]] && [[ ! -r $PIDFILE ]]
   then
      bad_pidfile
   fi

   #Make sure there is only one CAD in PID file
   validate_pid

   # read and stop hacmp CAD
   read_pid

   # Call exit function to display final message and exit
   CLEAN_EXIT
}
```

```
# Now invoke the above function to stop the Client Accepter Daemon (CAD)
# and all child processes

CAD_STOP
```

3. Next, we placed the clean script in as /opt/local/tsmsrv/cleanTSMsrv.sh, as shown in Example 20-9.

Example 20-9 /opt/local/tsmcli/cleanTSMcli.sh

```
#!/bin/ksh
# CLean script for VCS
# TSM client and scheduler process if the stop fails
# Only used by VCS if there is no other option

TSMCLIPID=`ps -ef | egrep "sched|dsmcad" | awk '{ print $2 }'`

echo $TSMCLIPID
for PID in $TSMCLIPID
 do
  kill -9 $PID
 done

exit 0
```

4. Lastly, we use the process monitoring for the client CAD and do not use a script. The process we will monitor is /usr/tivoli/tsm/client/ba/bin/dsmcad. This will be configured within VCS in 20.5.2, “Configuring Service Groups and applications” on page 865.

Scripts for the ISC

We placed the scripts for the server in the rootvg, /opt filesystem, in the directory /opt/local/isc.

1. First, we place the start script in the directory as /opt/local/isc/startISC.sh, shown in Example 20-10.

Example 20-10 /opt/local/isc/startISC.sh

```
#!/bin/ksh
# Startup the ISC_Portal
# This startup will also make the TSM Admin Center available
# There is aproaximately a 60-70 second start delay, prior to the script
# returning RC=0
#
/opt/IBM/ISC/PortalServer/bin/startISC.sh ISC_Portal iscadmin iscadmin
>/dev/console 2>&1
```

```
if [ $? -ne 0 ]
        then exit 1
fi
exit 0
```

2. Next, we place the stop script in the directory as /opt/local/isc/stopISC.sh, shown in Example 20-11.

Example 20-11 /opt/local/isc/stopISC.sh

```
#!/bin/ksh
# Stop The ISC_Portal and the TSM Administration Centre
/opt/IBM/ISC/PortalServer/bin/stopISC.sh ISC_Portal iscadmin iscadmin
if [ $? -ne 0 ]
        then exit 1
fi
exit 0
```

3. Then, we place the clean script in the directory as /opt/local/isc/cleanSTA.sh, as shown in Example 20-12.

Example 20-12 /opt/local/isc/cleanISC.sh

```
#!/bin/ksh
# killing ISC server process if the stop fails

ISCPID=`ps -af | egrep "AppServer|ISC_Portal" | awk '{ print $2 }'`

for PID in $ISCPID
 do
  kill -9 $PID
 done

exit 0
```

4. Lastly, we place the monitor script in the directory as /opt/local/isc/monISC.sh, shown in Example 20-13.

Example 20-13 /opt/local/isc/monISC.sh

```
#!/bin/ksh
# Monitoring for the existance of the ISC

LINES=`ps -ef | egrep "AppServer|ISC_Portal" | awk '{print $2}' | wc | awk
'{print $1}'` >/dev/console 2>&1
if [ $LINES -gt 1 ]
        then exit 110
fi
exit 100
```

20.5.2 Configuring Service Groups and applications

For this Service Group configuration, we use the command line approach:

1. First, to start, we change the default value for 'Type=Application, OnlineTimeout = 300', which is not enough for the ISC startup time, as shown in Example 20-14.

Example 20-14 Changing the OnlineTimeout for the ISC

```
hatype -modify Application OnlineTimeout 600
```

2. Then, we add the Service Group in VCS, first making the configuration readwrite, then adding the Service Group, then doing a series of modify commands, which define which nodes will participate, and their order, and the autostart list, as shown in Example 20-15.

Example 20-15 Adding a Service Group

```
haconf -makerw
hagrp -add sg_isc_sta_tsmcli
hagrp -modify sg_isc_sta_tsmcli SystemList banda 0 atlantic 1
hagrp -modify sg_isc_sta_tsmcli AutoStartList banda atlantic
hagrp -modify sg_isc_sta_tsmcli Parallel 0
```

3. Then, we add the LVMVG Resource to the Service Group sg_isc_sta_tsmcli, as depicted in Example 20-16. We set only the values that are relevant to starting Volume Groups (Logical Volume Manager).

Example 20-16 Adding an LVMVG Resource

```
hares -add vg_iscvg LVMVG sg_isc_sta_tsmcli
hares -modify vg_iscvg Critical 1
hares -modify vg_iscvg MajorNumber 48
hares -modify vg_iscvg ImportvgOpt n
hares -modify vg_iscvg SyncODM 1
hares -modify vg_iscvg VolumeGroup iscvg
hares -modify vg_iscvg OwnerName ""
hares -modify vg_iscvg GroupName ""
hares -modify vg_iscvg Mode ""
hares -modify vg_iscvg VaryonvgOpt ""
hares -probe vg_iscvg -sys banda
hares -probe vg_iscvg -sys atlantic
```

4. Next, we add the Mount Resource (mount point), which is also a resource configured within the Service Group sg_isc_sta_tsmcli as shown in Example 20-17. Note the link command at the bottom, which is the first parent-child resource relationship we establish.

Archived

Example 20-17 Adding the Mount Resource to the Service Group sg_isc_sta_tsmcli

```
hares -add m_ibm_isc Mount sg_isc_sta_tsmcli
hares -modify m_ibm_isc Critical 1
hares -modify m_ibm_isc SnapUmount 0
hares -modify m_ibm_isc MountPoint /opt/IBM/ISC
hares -modify m_ibm_isc BlockDevice /dev/isclv
hares -modify m_ibm_isc FSType jfs2
hares -modify m_ibm_isc MountOpt ""
hares -modify m_ibm_isc FsckOpt "-y"
hares -probe m_ibm_isc -sys banda
hares -probe m_ibm_isc -sys atlantic
hares -link m_ibm_isc vg_iscvg
```

5. Next, we add the NIC Resource for this Service Group. This monitors the NIC layer to determine if there is connectivity to the network. This is shown in Example 20-18.

Example 20-18 Adding a NIC Resource

```
hares -add NIC_en2 NIC sg_isc_sta_tsmcli
hares -modify NIC_en2 Critical 1
hares -modify NIC_en2 PingOptimize 1
hares -modify NIC_en2 Device en2
hares -modify NIC_en2 NetworkType ether
hares -modify NIC_en2 NetworkHosts -delete -keys
hares -modify NIC_en2 Enabled 1
hares -probe NIC_en2 -sys banda
hares -probe NIC_en2 -sys atlantic
```

6. Now, we add an IP Resource to the Service Group sg_isc_sta_tsmcli, as shown in Example 20-19. This resource will be linked to the NIC resource, implying that the NIC must be available prior to bringing the IP online.

Example 20-19 Adding an IP Resource

```
hares -add app_pers_ip IP sg_isc_sta_tsmcli
VCS NOTICE V-16-1-10242 Resource added. Enabled attribute must be set before
agent monitors
hares -modify app_pers_ip Critical 1
hares -modify app_pers_ip Device en2
hares -modify app_pers_ip Address 9.1.39.77
hares -modify app_pers_ip NetMask 255.255.255.0
hares -modify app_pers_ip Options ""
hares -probe app_pers_ip -sys banda
hares -probe app_pers_ip -sys atlantic
hares -link app_pers_ip NIC_en2
```

Archived

7. Then, to add the clustered Tivoli Storage Manager client, we add the additional Application Resource app_tsmcad within the Service Group sg_isc_sta_tsmcli, as shown in Example 20-20.

Example 20-20 VCS commands to add tsmcad application to the sg_isc_sta_tsmcli

```
hares -add app_tsmcad Application sg_isc_sta_tsmcli
hares -modify app_tsmcad User ""
hares -modify app_tsmcad StartProgram /opt/local/tsmcli/startTSMcli.sh
hares -modify app_tsmcad StopProgram /opt/local/tsmcli/stopTSMcli.sh
hares -modify app_tsmcad CleanProgram /opt/local/tsmcli/stopTSMcli.sh
hares -modify app_tsmcad MonitorProgram /opt/local/tsmcli/monTSMcli.sh
hares -modify app_tsmcad PidFiles -delete -keys
hares -modify app_tsmcad MonitorProcesses /usr/tivoli/tsm/client/ba/bin/dsmcad
hares -probe app_tsmcad -sys banda
hares -probe app_tsmcad -sys atlantic
hares -link app_tsmcad app_pers_ip
```

8. Next, we add an Application Resource app_isc to the Service Group sg_isc_sta_tsmcli, as shown in Example 20-21.

Example 20-21 Adding app_isc Application to the sg_isc_sta_tsmcli Service Group

```
hares -add app_isc Application sg_isc_sta_tsmcli
hares -modify app_isc User ""
hares -modify app_isc StartProgram /opt/local/isc/startISC.sh
hares -modify app_isc StopProgram /opt/local/isc/stopISC.sh
hares -modify app_isc CleanProgram /opt/local/isc/cleanISC.sh
hares -modify app_isc MonitorProgram /opt/local/isc/monISC.sh
hares -modify app_isc PidFiles -delete -keys
hares -modify app_isc MonitorProcesses -delete -keys
hares -probe app_isc -sys banda
hares -probe app_isc -sys atlantic
hares -link app_isc app_pers_ip
haconf -dump -makero
```

9. Next, we review the main.cf file which reflects the sg_isc_sta_tsmcli Service Group, as shown in Example 20-22.

Example 20-22 Example of the main.cf entries for the sg_isc_sta_tsmcli

```
group sg_isc_sta_tsmcli (
        SystemList = { banda = 0, atlantic = 1 }
        AutoStartList = { banda, atlantic }
        )

        Application app_isc (
                Critical = 0
                StartProgram = "/opt/local/isc/startISC.sh"
                StopProgram = "/opt/local/isc/stopISC.sh"
```

```
                CleanProgram = "/opt/local/isc/cleanISC.sh"
                MonitorProgram = "/opt/local/isc/monISC.sh"
                )

        Application app_tsmcad (
                Critical = 0
                StartProgram = "/opt/local/tsmcli/startTSMcli.sh"
                StopProgram = "/opt/local/tsmcli/stopTSMcli.sh"
                CleanProgram = "/opt/local/tsmcli/stopTSMcli.sh"
                MonitorProcesses = { "/usr/tivoli/tsm/client/ba/bin/dsmc sched"
}
                )

        IP app_pers_ip (
                Device = en2
                Address = "9.1.39.77"
                NetMask = "255.255.255.0"
                )

        LVMVG vg_iscvg (
                VolumeGroup = iscvg
                MajorNumber = 48
                )

        Mount m_ibm_isc (
                MountPoint = "/opt/IBM/ISC"
                BlockDevice = "/dev/isclv"
FSType = jfs2
                FsckOpt = "-y"
                )

        NIC NIC_en2 (
                Device = en2
                NetworkType = ether
                )

        app_isc requires app_pers_ip
        app_pers_ip requires NIC_en2
        app_pers_ip requires m_ibm_isc
        app_tsmcad requires app_pers_ip
        m_ibm_isc requires vg_iscvg

        // resource dependency tree
        //
        //      group sg_isc_sta_tsmcli
        //      {
        //      Application app_isc
        //          {
        //          IP app_pers_ip
```

```
        //              {
        //              NIC NIC_en2
        //              Mount m_ibm_isc
        //                  {
        //                  LVMVG vg_iscvg
        //                  }
        //              }
        //          }
        //      Application app_tsmcad
        //          {
        //          IP app_pers_ip
        //              {

//              NIC NIC_en2
        //              Mount m_ibm_isc
        //                  {
        //                  LVMVG vg_iscvg
        //                  }
        //              }
        //          }
        //      }
```

10. Now, we review the configuration for the sg_isc_sta_tsmcli Service Group using the Veritas Cluster Manager GUI, as shown in Figure 20-21.

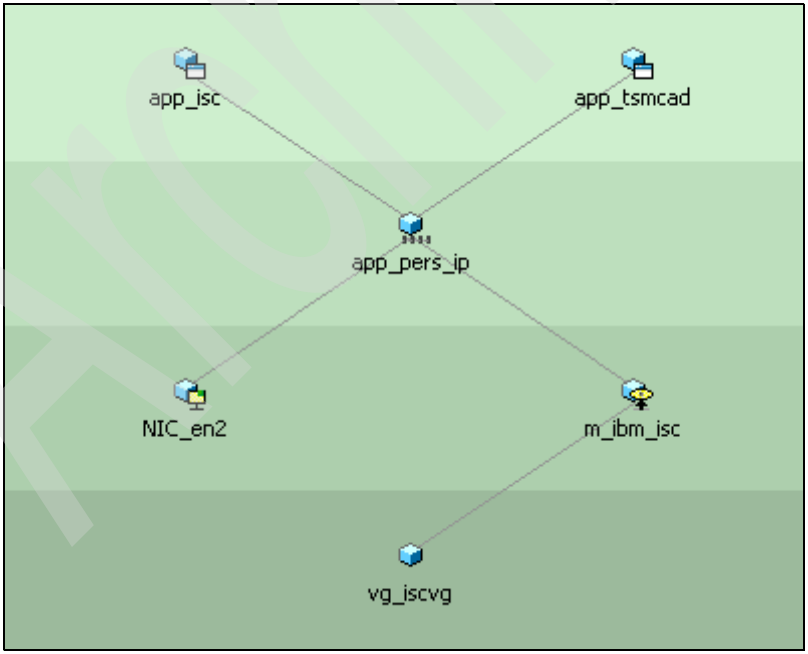

Figure 20-21 GUI diagram, child-parent relation, sg_isc_sta_tsmcli Service Group

20.6 Testing the highly available client and ISC

We now move into testing the highly available Tivoli Storage Manager client. After these tests begin, we will crash the server which has begun the test (the AIX command used is **halt -q**). We also explain the sequence of events as we progress through the various stages of testing.

20.6.1 Cluster failure during a client back up

Now we test the ability of a scheduled backup operation to restart and complete, after a node crash (backup going direct to tape):

1. We verify that the cluster services are running with the **lssrc -g cluster command** on both nodes.
2. On the resource group secondary node, we use **tail -f /tmp/VCS.out** to monitor cluster operation.
3. Then we schedule a client selective backup having the whole shared filesystem as an object and wait for it being started using query session on the Tivoli Storage Manager server (Example 20-23).

Example 20-23 Client sessions starting

```
Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- --------------------
tsm: TSMSRV03>q se
    58 Tcp/Ip SendW    0 S      701     139 Admin AIX      ADMIN
    59 Tcp/Ip IdleW   38 S      857     501 Node  AIX      CL_VERITAS01_CLIENT
    60 Tcp/Ip Run      0 S      349   8.1 M Node  AIX      CL_VERITAS01_CLIENT
```

4. We wait for volume opened messages on server console (Example 20-24).

Example 20-24 Volume opened messages on server console

```
ANR0406I Session 59 started for node CL_VERITAS01_CLIENT (AIX) (Tcp/Ip
9.1.39.42(32869)).
ANR0406I Session 60 started for node CL_VERITAS01_CLIENT (AIX) (Tcp/Ip
9.1.39.92(32870)).
ANR1639I Attributes changed for node CL_VERITAS01_CLIENT: TCP Address from
9.1.39.42 to 9.1.39.92.
ANR8337I LTO volume 030AKK mounted in drive DRLTO_2 (/dev/rmt1).
ANR0511I Session 60 opened output volume 030AKK.
ANR0407I Session 61 started for administrator ADMIN (AIX) (Tcp/Ip
```

Archived

Failure

This is the only step needed for this test:

1. Being sure that client LAN-free backup is running, we issue **halt -q** on the AIX server on Atlantic, for which the backup is running; the **halt -q** command stops any activity immediately and powers off the server.

Recovery

These are the steps we follow for this test:

1. The second node, Banda takes over the resources and starts up the Service Group and Application start script.
2. Next, the clustered scheduler start script is started. Once this happens, the Tivoli Storage Manager server logs the difference in physical node names on the server console, as shown in Example 20-25.

Example 20-25 Server console log output for the failover reconnection

```
ANR0406I Session 221 started for node CL_VERITAS01_CLIENT (AIX) (Tcp/Ip
9.1.39.94(33515)).
ANR1639I Attributes changed for node CL_VERITAS01_CLIENT: TCP Name from
atlantic to banda,
GUID from 00.00.00.01.75.8f.11.d9.b4.d1.08.63.09.01.27.5c to
00.00.00.00.75.8e.11.d9.ac.29.08.63.09.01.27.5e.
ANR0403I Session 221 ended for node CL_VERITAS01_CLIENT (AIX).
```

3. Once the session's cancelling work finishes, the scheduler is restarted and the scheduled backup operation is restarted, as shown in Example 20-26.

Example 20-26 The client schedule restarts.

```
ANR0403I Session 221 ended for node CL_VERITAS01_CLIENT (AIX).
ANR0406I Session 222 started for node CL_VERITAS01_CLIENT (AIX) (Tcp/Ip
9.1.39.43(33517)).
ANR0406I Session 223 started for node CL_VERITAS01_CLIENT (AIX) (Tcp/Ip
9.1.39.94(33519)).
ANR0403I Session 223 ended for node CL_VERITAS01_CLIENT (AIX).
ANR0403I Session 222 ended for node CL_VERITAS01_CLIENT (AIX).
ANR0406I Session 224 started for node CL_VERITAS01_CLIENT (AIX) (Tcp/Ip
9.1.39.43(33521)).
```

4. The Tivoli Storage Manager command **q session** still shows the backup in progress, as shown in Example 20-27.

Archived

Example 20-27 q session shows the backup and dataflow continuing

```
tsm: TSMSRV03>q se

  Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- --------------------

    58 Tcp/Ip SendW    0 S    3.1 K     139 Admin AIX      ADMIN
    59 Tcp/Ip IdleW  9.9 M      905     549 Node  AIX      CL_VERITAS01_CLIENT
    60 Tcp/Ip RecvW  9.9 M      574 139.6 M Node  AIX      CL_VERITAS01_CLIENT
```

5. Next, we see from the server actlog that the session is closed and the tape unmounted, as shown in Example 20-28.

Example 20-28 Unmounting the tape once the session is complete

```
ANR8336I Verifying label of LTO volume 030AKK in drive DRLTO_2 (/dev/rmt1).
ANR8468I LTO volume 030AKK dismounted from drive DRLTO_2 (/dev/rmt1) in library
LIBLTO.
```

6. We can find messages in the actlog for backup operation restarting in a completed successful message, as shown in Example 20-29.

Example 20-29 Server actlog output of the session completing successfully

```
ANR2507I Schedule TEST_SCHED for domain STANDARD started at 02/19/05 19:52:08
for node CL_VERITAS01_CLIENT completed successfully at 02/19/05 19:52:08.
```

Result summary

We are able to have the VCS cluster restarting an application with its backup environment up and running.

Locked resources are discovered and freed up.

Scheduled operation is restarted via by the scheduler and obtains back the previous resources.

There is the opportunity of having a backup restarted even if, considering a database as an example, this can lead to a backup window breakthrough, thus affecting other backup operations.

We run this test, at first using command line initiated backups with the same result; the only difference is that the operation needs to be restarted manually.

Archived

20.6.2 Cluster failure during a client restore

In this test we are verifying how a restore operation scenario is managed in a client takeover scenario.

Objective

For this test we will use a scheduled restore, which, after the failover recovery, will restart the restore operation that was interrupted. We will use a scheduled operation with the parameter **replace=all**, so the restore operation is restarted from the beginning on restart, with no prompting.

If we were to use a manual restore with a command line (and wildcard), this would be restarted from the point of failure with the Tivoli Storage Manager client command **restart restore**.

Preparation

These are the steps we follow for this test:

1. We verify that the cluster services are running with the **hastatus** command.
2. Then we schedule a restore with client node CL_VERITAS01_CLIENT association (Example 20-30).

Example 20-30 Schedule a restore with client node CL_VERITAS01_CLIENT

```
Day of Month:
                Week of Month:
                   Expiration:
Last Update by (administrator): ADMIN
        Last Update Date/Time: 02/21/05   10:26:04
            Managing profile:

           Policy Domain Name: STANDARD
               Schedule Name: RESTORE_TEST
                 Description:
                      Action: Restore
                     Options: -subdir=yes -replace=all
                     Objects: /install/*.*
                    Priority: 5
             Start Date/Time: 02/21/05   18:30:44
                    Duration: Indefinite
              Schedule Style: Classic
                      Period: One Time
                 Day of Week: Any
                       Month:
                Day of Month:
                Week of Month:
                   Expiration:
```

```
Last Update by (administrator): ADMIN
         Last Update Date/Time: 02/21/05   18:52:26
              Managing profile:
```

3. We wait for the client session to start and data beginning to be transferred to Banda, as seen in Example 20-31.

Example 20-31 Client sessions starting

```
tsm: TSMSRV06>q se

  Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- --------------------
   290 Tcp/Ip Run      0 S   32.5 K     139 Admin AIX      ADMIN
   364 Tcp/Ip Run      0 S    1.9 K     211 Admin AIX      ADMIN
   366 Tcp/Ip IdleW  7.6 M  241.0 K   1.9 K Admin DSMAPI   ADMIN
   407 Tcp/Ip SendW    1 S   33.6 M   1.2 K Node  AIX      CL_VERITAS01_CLIENT
```

4. Also, we look for the input volume being mounted and opened for the restore, as seen in Example 20-32.

Example 20-32 Mount of the restore tape as seen from the server actlog

```
ANR8337I LTO volume 030AKK mounted in drive DRLTO_2 (/dev/rmt1).
ANR0511I Session 60 opened output volume 020AKK.
```

Failure

These are the steps we follow for this test:

1. Once satisfied that the client restore is running, we issue **halt -q** on the AIX server running the Tivoli Storage Manager client (Banda). The **halt -q** command stops AIX immediately and powers off the server.
2. The server is not receiving data to server, and sessions remain in idlew and recvw state.

Recovery

These are the steps we follow for this test:

1. Atlantic takes over the resources and launches the Tivoli Storage Manager cad start script.
2. In Example 20-33 we can see the server console showing that the same events occurred in the backup test previously completed:
   a. The select searching for a tape holding session.
   b. The **cancel** command for the session found above.

c. A new select with no result because the first **cancel session** command is successful.

d. The restarted client scheduler querying for schedules.

e. The schedule is still in the window, so a new restore operation is started, and it obtains its input volume.

Example 20-33 The server log during restore restart

```
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR0405I Session 415 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0514I Session 407 closed volume 020AKKL2.
ANR0480W Session 407 for node CL_VERITAS01_CLIENT (AIX) terminated - connection
with client severed.
ANR8336I Verifying label of LTO volume 020AKKL2 in drive DRLTO_1 (mt0.0.0.2).
ANR0407I Session 416 started for administrator SCRIPT_OPERATOR (AIX) (Tcp/Ip
9.1.39.92(32911)).
ANR2017I Administrator SCRIPT_OPERATOR issued command: select
SESSION_ID,CLIENT_NAME from SESSIONS where CLIENT_NAME='CL_VERITAS01_CLIENT'
ANR2034E SELECT: No match found using this criteria.
ANR2017I Administrator SCRIPT_OPERATOR issued command: ROLLBACK
ANR0405I Session 416 ended for administrator SCRIPT_OPERATOR (AIX).
ANR0406I Session 417 started for node CL_VERITAS01_CLIENT (AIX) (Tcp/Ip
9.1.39.92(32916)).
ANR1639I Attributes changed for node CL_VERITAS01_CLIENT: TCP Name from banda
to atlantic, TCP Address from 9.1.39.43 to 9.1.39.92, GUID from
00.00.00.00.75.8e.11.d9.ac.29.08.63.09.01.27.5e to
00.00.00.01.75.8f.11.d9.b4.d1.08.63.09.01.27.5c.
ANR0403I Session 417 ended for node CL_VERITAS01_CLIENT (AIX).
ANR0406I Session 430 started for node CL_VERITAS01_CLIENT (AIX) (Tcp/Ip
9.1.39.42(32928)).
ANR1639I Attributes changed for node CL_VERITAS01_CLIENT: TCP Address from
9.1.39.92 to 9.1.39.42.
```

3. The new restore operation completes successfully.

4. In the client log we can see the restore start, interruption and restart.

Example 20-34 The Tivoli Storage Manager client log

```
SCHEDULEREC QUERY BEGIN
SCHEDULEREC QUERY END
Next operation scheduled:
------------------------------------------------------------
Schedule Name:         RESTORE_TEST
Action:                Restore
Objects:               /install/*.*
Options:               -subdir=yes -replace=all
Server Window Start:   18:30:44 on 02/21/05
```

```
------------------------------------------------------------
Executing scheduled command now.
--- SCHEDULEREC OBJECT BEGIN RESTORE_TEST
Restore function invoked.
.
.
.
Restoring          71,680 /install/AIX_ML05/U800869.bff [Done]
Restoring         223,232 /install/AIX_ML05/U800870.bff [Done]
Restore processing finished.
--- SCHEDULEREC STATUS BEGIN
Total number of objects restored:     1,774
Total number of objects failed:           0
Total number of bytes transferred:     1.03 GB
Data transfer time:                1,560.33 sec
Network data transfer rate:          693.54 KB/sec
Aggregate data transfer rate:        623.72 KB/sec
Elapsed processing time:           00:28:55
SCHEDULEREC STATUS END
SCHEDULEREC OBJECT END RESTORE_TEST 02/21/05   18:30:44
SCHEDULEREC STATUS BEGIN
SCHEDULEREC STATUS END
Scheduled event 'RESTORE_TEST' completed successfully.
```

Result summary

The cluster is able to manage client failure and make Tivoli Storage Manager client scheduler available in about 1 minute and the client is able to restart its operations successfully to the end.

Since this is a scheduled restore with **`replace=all`**, it is restarts from the beginning and completes successfully, overwriting the previously restored data.

Important: In every failure test done, we have traced and documented from the client perspective. We will not mention the ISC at all, however, this application fails every time the client does, and totally recovers on the surviving node every time during these tests. After every failure, we log into the ISC to make server schedule changes, or others for other reasons, so the application is constantly accessed, and during multiple server failure tests, the ISC has always recovered.

Archived

Part 6

Establishing a VERITAS Cluster Server Version 4.0 infrastructure on Windows with IBM Tivoli Storage Manager Version 5.3

In this part of the book, we describe how we set up Tivoli Storage Manager Version 5.3 products to be used with Veritas Cluster Server Version 4.0 in Microsoft Windows 2003 environments.

© Copyright IBM Corp. 2005. All rights reserved.

Archived

21

Installing the VERITAS Storage Foundation HA for Windows environment

This chapter describes how our team planned, installed, configured, and tested the Storage Foundation HA for Windows on Windows 2003.

We explain how to do the following tasks:

- Plan, install, and configure the Storage Foundation HA for Windows for the Tivoli Storage Manager application
- Test the clustered environment prior to deployment of the Tivoli Storage Manager application.

© Copyright IBM Corp. 2005. All rights reserved.

21.1 Overview

VERITAS Storage Foundation HA for Windows is a package that comprises two high availability technologies:

- VERITAS Storage Foundation for Windows
- VERITAS Cluster Server

VERITAS Storage Foundation for Windows allows storage management. VERITAS Cluster Server is the clustering solution itself.

21.2 Planning and design

For our VCS environment running on Windows 2003, we will implement a two-node cluster, with two resource groups, one for Tivoli Storage Manager Server and Client, and another one for the Integrated Solutions Console, Tivoli Storage Manager administration tool. We will be using two private networks for the heartbeat.

We install the basic package of VERITAS Storage Foundation HA for Windows. For specific configurations and more information on the product, we highly recommend referencing the following VERITAS documents, available at:

http://support.veritas.com

These are the documents:

- *Release Notes*
- *Getting Started Guide*
- *Installation Guide*
- *Administrator's Guide*

21.3 Lab environment

Table 21-1 shows the lab we use to set up our Windows 2003 in two servers, SALVADOR and OTTAWA.

Archived

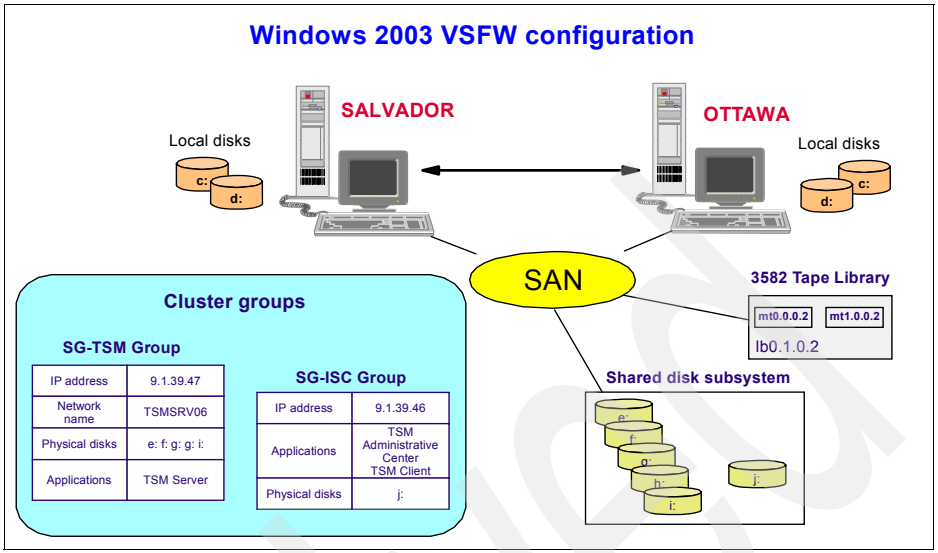

Figure 21-1 Windows 2003 VSFW configuration

The details of this configuration for the servers SALVADOR and OTTAWA are shown in Table 21-1, Table 21-2 and Table 21-3 below. One factor which determines our disk requirements and planning for this cluster is the decision of using Tivoli Storage Manager database and recovery log mirroring. This requires four disks, two for the database and two for the recovery log.

Table 21-1 Cluster server configuration

VSFW Cluster	
Cluster name	CL_VCS02
Node 1	
Name	SALVADOR
Private network IP addresses	10.0.0.1 and 10.0.1.1
Public network IP address	9.1.39.44
Node 2	
Name	OTTAWA
Private network IP addresses	10.0.0.2 10.0.1.2
Public network IP address	9.1.39.45

Archived

Table 21-2 Service Groups in VSFW

Service Group 1	
Name	SG-ISC
IP address	9.1.39.46
Network name	ADMCNT06
Physical disks	j:
Applications	IBM WebSphere Application Center ISC Help Service
Service Group 2	
Name	SG-TSM
IP address	9.1.39.47
Network name	TSMSRV06
Physical disks	e: f: g: h: i:
Applications	TSM Server

Table 21-3 DNS configuration

Domain	
Name	TSMVERITAS.COM
Node 1	
DNS name	salvador.tsmveritas.com
Node 2	
DNS name	ottawa.tsmveritas.com

21.4 Before VSFW installation

Before we install VSFW, we need to prepare Windows 2003 with the necessary configuration.

21.4.1 Installing Windows 2003

For our lab we choose to install Windows 2003 Advanced Server. Since we do not have other servers to be domain controllers, we install Active Directory and DNS Servers in both nodes.

Archived

21.4.2 Preparing network connectivity

For this cluster, we will be implementing two private ethernet networks and one production LAN interface. For ease of use, we rename the network connections icons to Private1, Private2, and Public, as shown in Figure 21-2.

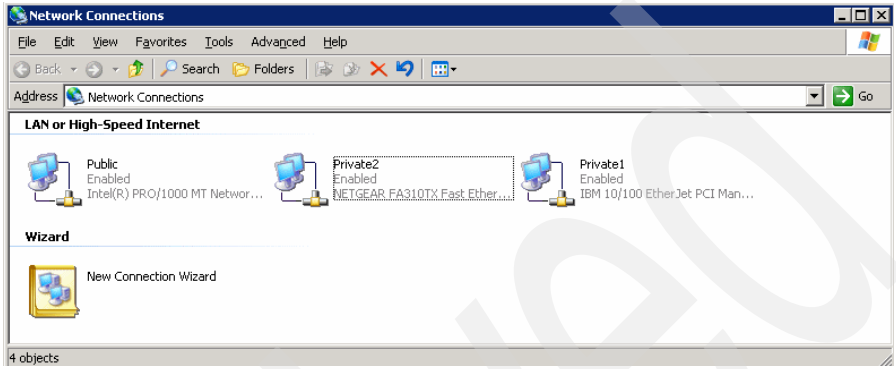

Figure 21-2 Network connections

The two network cards have some special settings shown below:

1. We wire two adapters per machine using an ethernet cross-over cable. We use the exact same adapter location and type of adapter for this connection between the two nodes.
2. We then configure the two private networks for IP communication. We set the link speed of the nic cards to 10 Mbps/Half Duplex and disable Netbios over TCP/IP
3. We run **ping** to test the connections.

21.4.3 Domain membership

All nodes must be members of the same domain and have access to a DNS server. In this lab we set up the servers both as domain controllers as well as DNS Servers. If this is your scenario, use **dcpromo.exe** to promote the servers to domain controllers.

Promoting the first server

These are the steps we followed:

1. We set up our network cards so that the servers point to each other for primary DNS resolution, and to themselves for secondary resolution.
2. We run **dcpromo** and create a new domain, a new tree, and a new forest.
3. We take note of the password used for the administrator account.

Archived

4. We let the setup install DNS server.
5. We wait until the setup finishes and boot the server.
6. We configure the DNS server and create a *Reverse Lookup Zones* for all our network addresses. We make them active directory integrated zones.
7. We define new hosts for each of the nodes with the option of creating the associated pointer (PTR) record.
8. We test DNS using **nslookup** from a command prompt.
9. We look for any error messages in the event viewer.

Promoting the other servers

These are the steps we followed:

1. We run **dcpromo** and join the domain created above, selecting **Additional domain controller for an existing domain**.
2. We use the password set up in step 3 on page 883 above.
3. When the server boots, we install **DNS server**.
4. We check if DNS is replicated correctly using **nslookup.**
5. We look for any error messages in the event viewer.

21.4.4 Setting up external shared disks

On the DS4500 side we prepare the LUNs that will be designated to our servers. A summary of the configuration is shown in Figure 21-3.

Attention: While configuring shared disks, we always have only one server up at a time, to avoid corruption. To proceed, we shut down all servers, turn on the storage device, and turn on only one of the nodes.

Archived

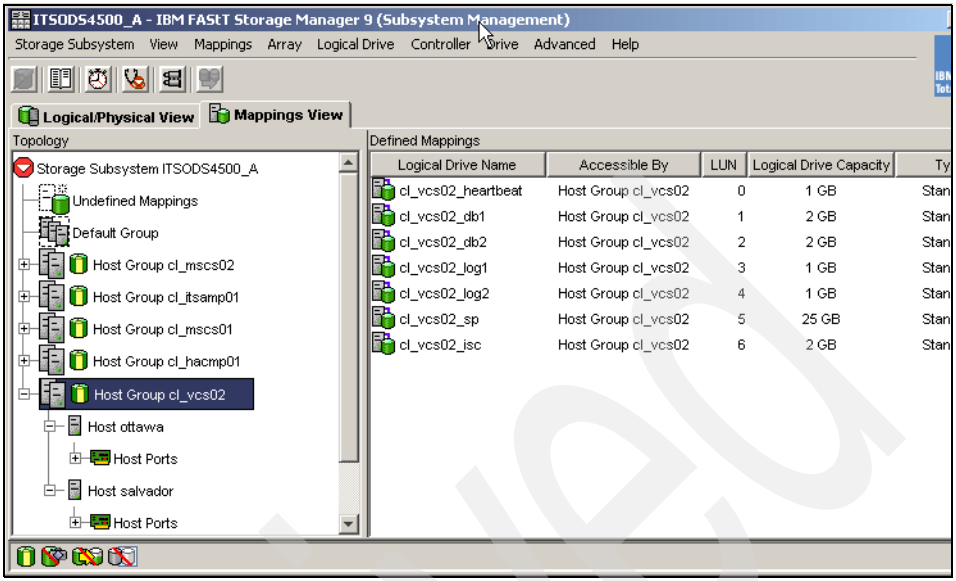

Figure 21-3 LUN configuration

For Windows 2003 and DS4500, we upgrade de QLOGIC drives and install the Redundant Disk Array Controller (RDAC) according to the manufacturer's manual, so that Windows recognizes the storage disks. Since we have dual path to the storage, if we do not install the RDAC, Windows will see duplicate drives. The device manager should look similar to Figure 21-4 on the items, Disk drivers and SCSI and RAID controllers.

Archived

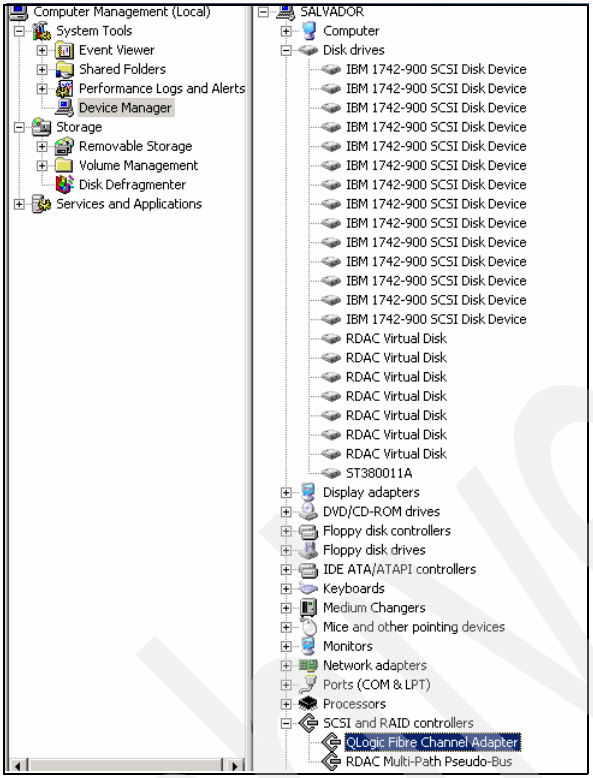

Figure 21-4 Device manager with disks and SCSI adapters

Configuring shared disks

Prior to installing the VSFW, we create the shared disks and partitions in Windows. VSFW can be set up either with or without disk partitioning in Windows:

1. We double-click Disk Management and follow the Write Signature and Upgrade Disk Wizard. We select all disks for the Write Signature part, but we choose not to upgrade any of the disks to dynamic.
2. When finished, the disk manager will now show that all disks are online but with unallocated partitions.
3. We create new partitions on each unallocated disk, assigning the maximum size. We also assign a letter to each partition, following our plan in Table 21-2 on page 882, and format them with NTFS.
4. We check disk access in Windows Explorer. We create any file on the drives and also try to delete them.

Archived

5. When we turn the second node, on we check the partitions. If the letters are not set correctly, we change them to match the ones you set up on the first node. We also test write/delete file access from the other node.

Note: VERITAS Cluster Server can also work with dynamic disks, provided that they are created with the VERITAS Storage Foundation for Windows, using the VERITAS Enterprise Administration GUI (VEA). For more information, refer to the *VERITAS Storage Foundation 4.2 for Windows Administrator's Guide*.

21.5 Installing the VSFW software

We will only execute the VSFW installation software on one node, and VSFW will simultaneously install the software on the second node. In order for this operation to be successful, we set the Windows driver signing options to ignore on both nodes. This is done in **Control Panel** →**System** →**Hardware tab** →**Driver Signing** and selecting **Ignore - Install all files, regardless of file signature**. We will reverse this at the end of the installation.

Important: Failure to change this setting will cause the installation of the remote node to be rejected when it validates the environment (Figure 21-14 on page 892). For the local node, we only have to be sure it is *not* set to **block**.

These are the steps we followed:

1. We run the setup.exe on the CD and choose **Storage Foundation HA 4.2 for Windows**, as shown in Figure 21-5.

Archived

Figure 21-5 Choosing the product to install

2. We choose the complete installation and click **Next** (Figure 21-6).

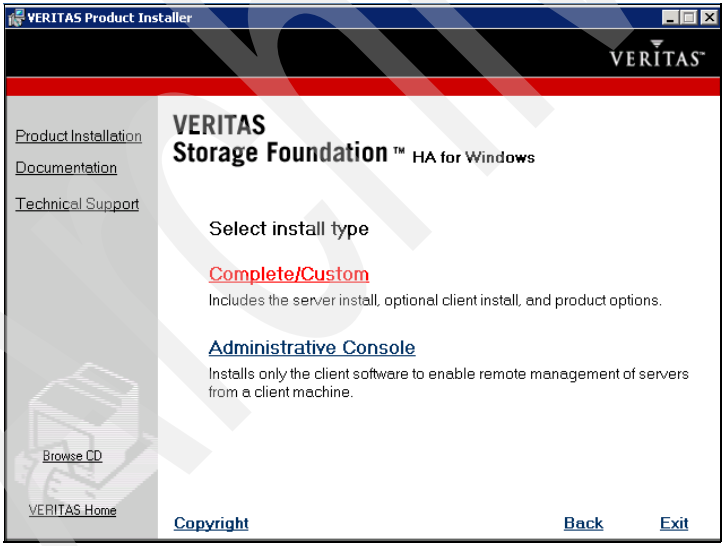

Figure 21-6 Choose complete installation

3. The files are unpacked, and the welcome page appears, as shown as in Figure 21-7. We read the prerequisites, confirming that we have disabled the driver signing option, and click **Next**.

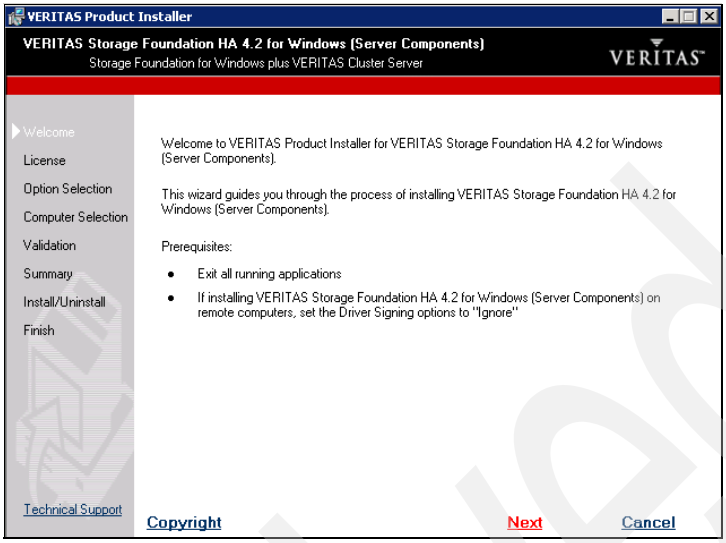

Figure 21-7 Pre-requisites - attention to the driver signing option

4. We read and accept the license agreement shown in Figure 21-8 and click **Next**.

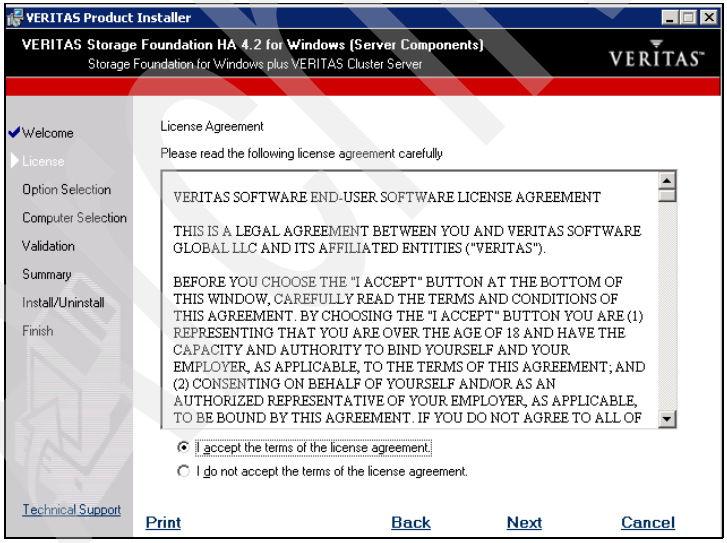

Figure 21-8 License agreement

5. We enter the license key (Figure 21-9), click **Add** so it is moved to the list below, and then click **Next**.

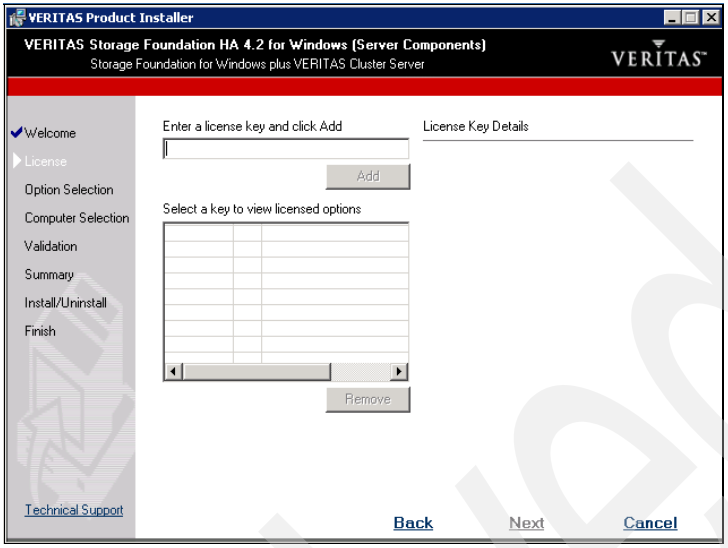

Figure 21-9 License key

6. Since we are installing only the basic software, we leave all boxes clear in Figure 21-10.

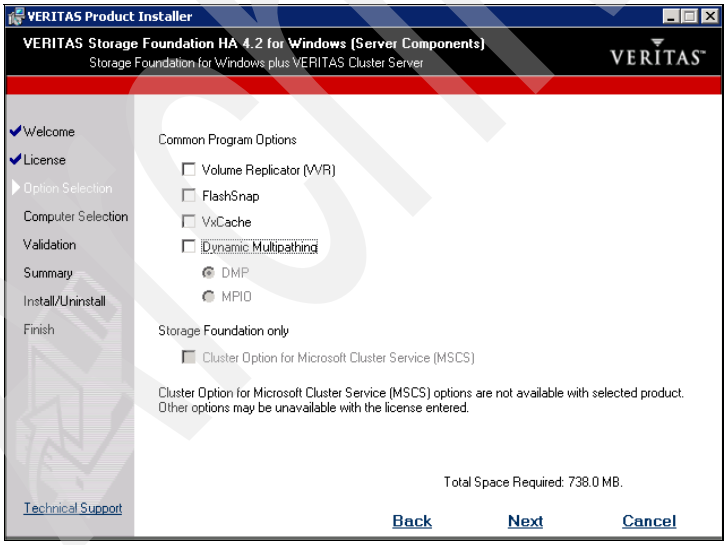

Figure 21-10 Common program options

Archived

7. We will not install the Global Campus Option (for clusters in geographically different locations) or any of the other applications, so we leave all boxes clear in Figure 21-11.

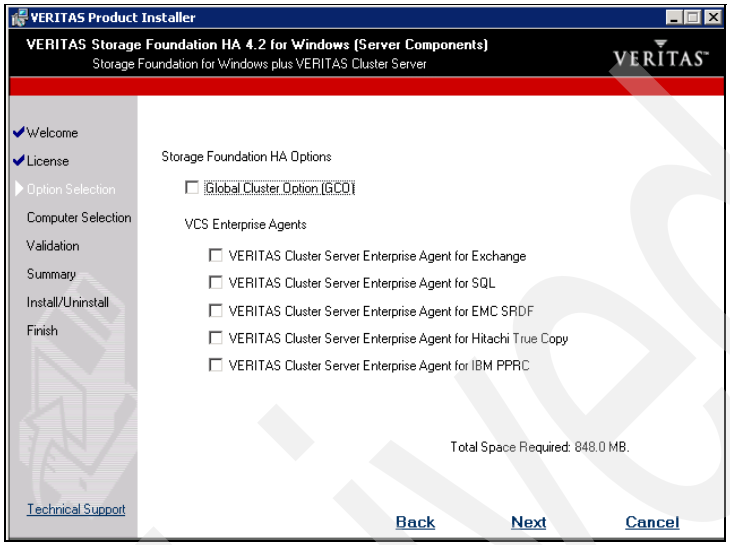

Figure 21-11 Global cluster option and agents

8. We choose to install the client components and click **Next** in Figure 21-12.

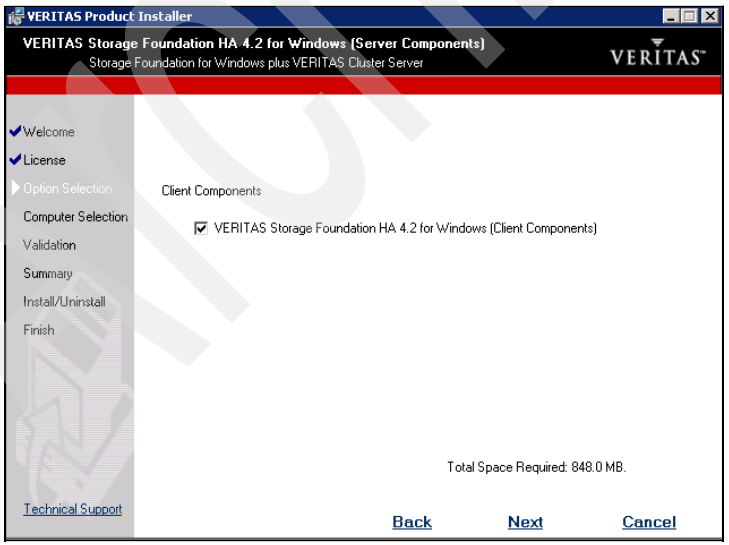

Figure 21-12 Install the client components

Archived

9. Using the arrow boxes, we choose to install the software on both machines. After highlighting each server, we click **Add** as shown in Figure 21-13. We leave the default install path. We confirm the information and click **Next**.

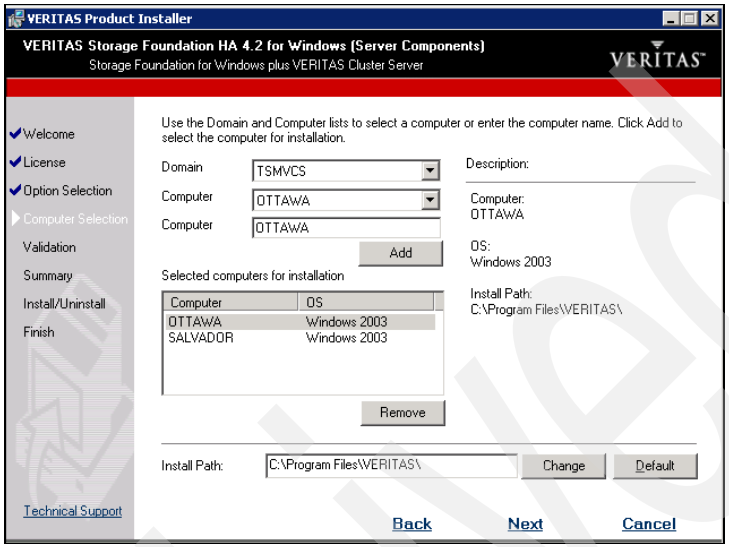

Figure 21-13 Choosing the servers and path

10. The installer will validate the environment and inform us if the setup is possible, as shown in Figure 21-14.

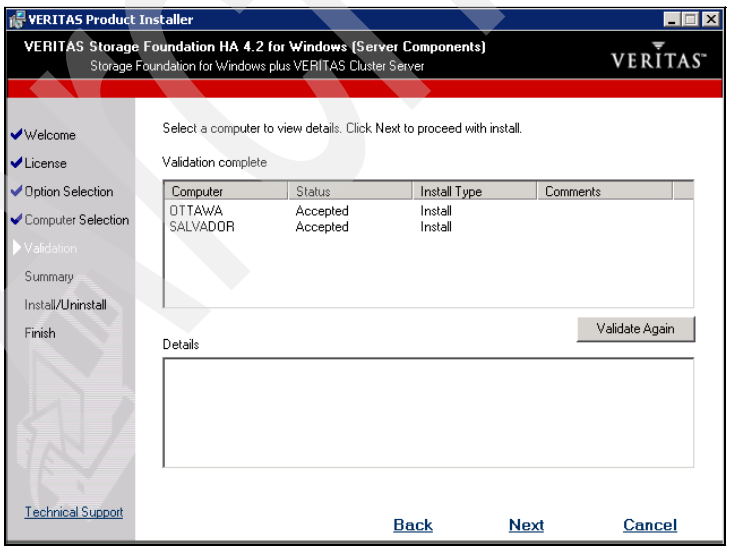

Figure 21-14 Testing the installation

11. We review the summary shown in Figure 21-15 and click **Install**.

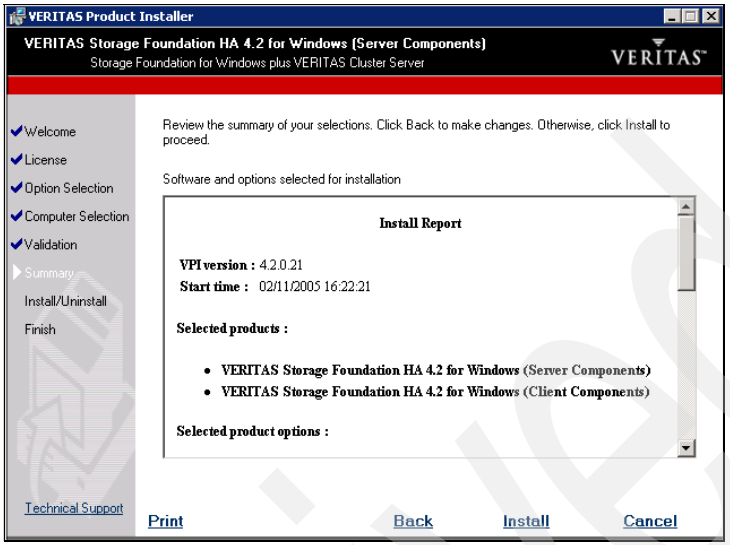

Figure 21-15 Summary of the installation

12. The installation process begins as shown in Figure 21-16.

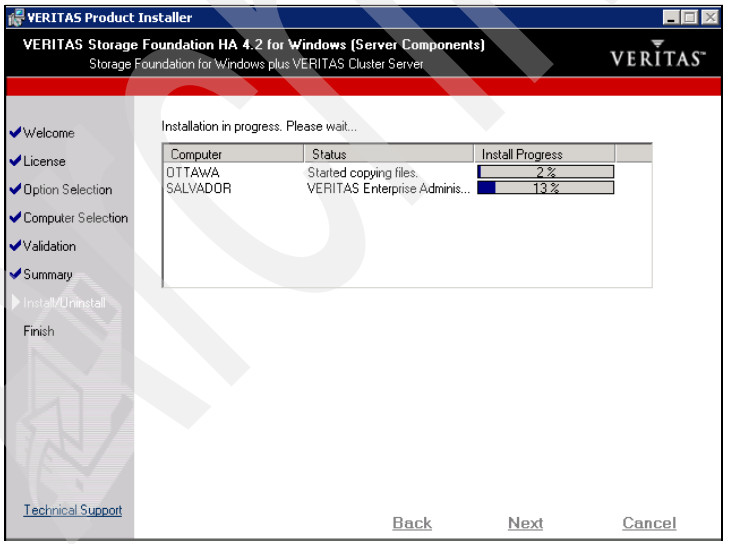

Figure 21-16 Installation progress on both nodes

13. When the installation finishes, we review the installation report summary as shown in Figure 21-17 and click **Next**.

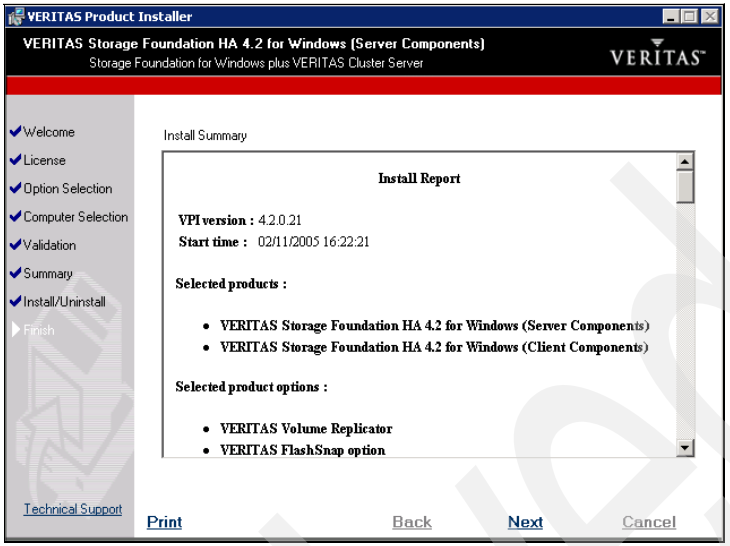

Figure 21-17 Install report

14. As shown in Figure 21-18, the installation now asks for the reboot of the remote server (OTTAWA). We click **Reboot** and wait until the remote server is back.

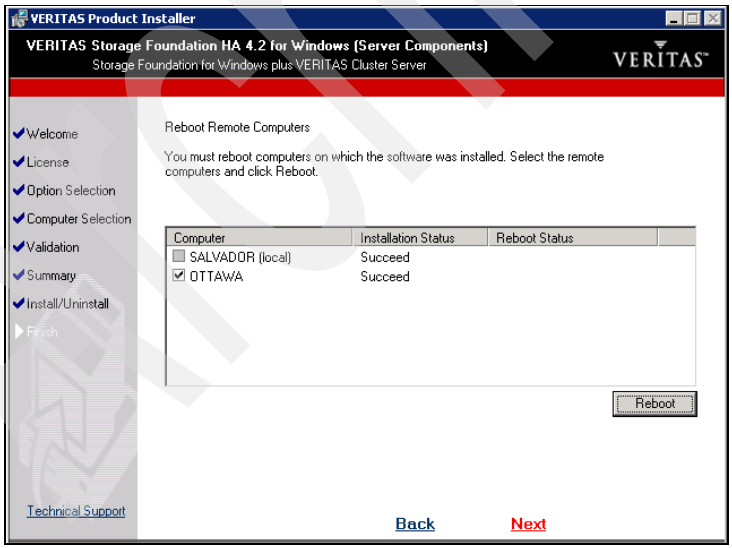

Figure 21-18 Reboot remote server

15. The installer shows the server is online again (Figure 21-19) so we click **Next**.

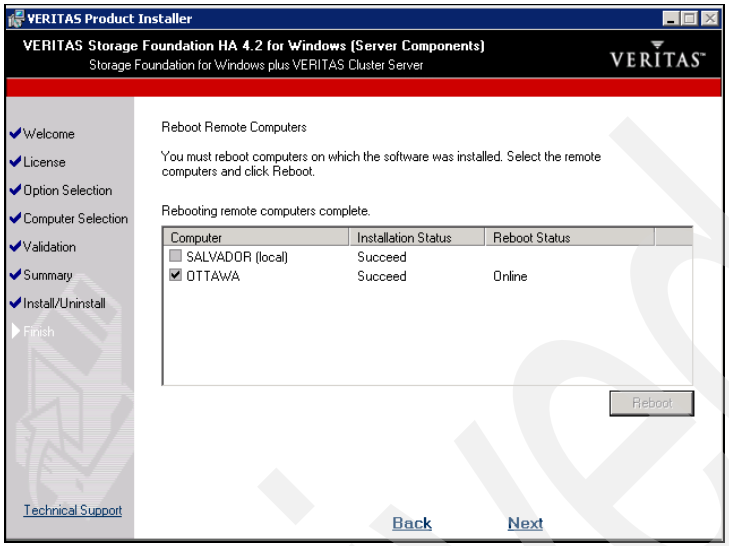

Figure 21-19 Remote server online

16. The installation is now complete. We have to reboot SALVADOR as shown in Figure 21-20. We click **Finish** and we are prompted to reboot the server.

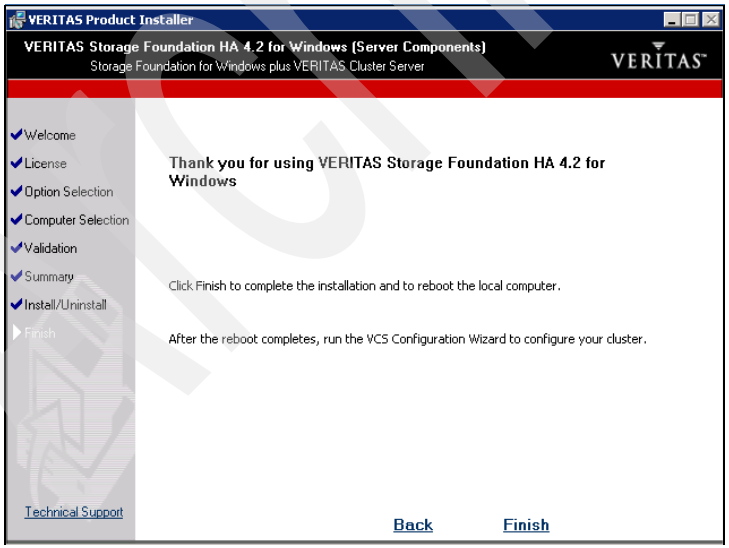

Figure 21-20 Installation complete

Archived

17. When the servers are back and installation is complete, we reset the driver signing option to Warn: **Control Panel →System →Hardware tab →Driver Signing** and then select **Warn - DIsplay message before installing an unsigned file**.

21.6 Configuring VERITAS Cluster Server

Now that the product is installed, we need to configure the environment. This can be done on any of the nodes with the VCS Configuration Wizard.

1. We open the wizard selecting **Start →All Programs →VERITAS →VERITAS Cluster Server →VCS Configuration Wizard**. When the welcome page appears we click **Next**.
2. On the Configuration Options page, in Figure 21-21, we choose Cluster Operations and click **Next**.

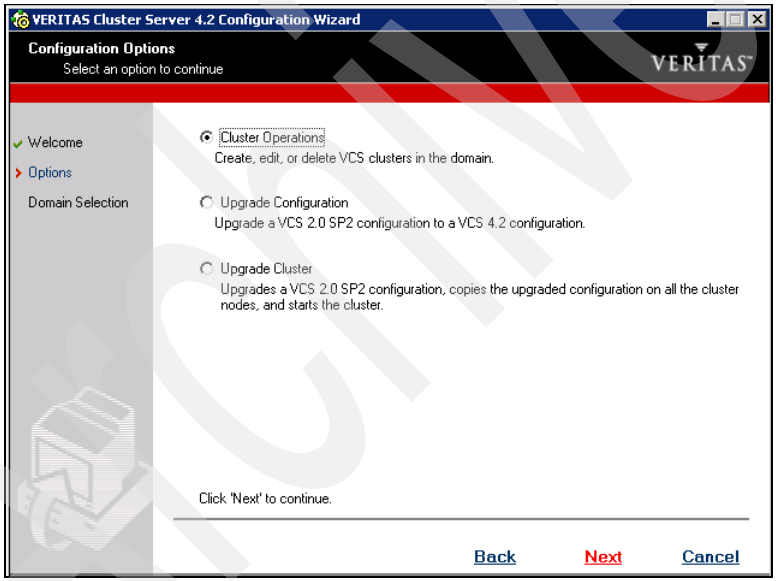

Figure 21-21 Start cluster configuration

3. On the Domain Selection page in Figure 21-22, we confirm the domain name and clear the check box **Specify systems and users manually**.

Archived

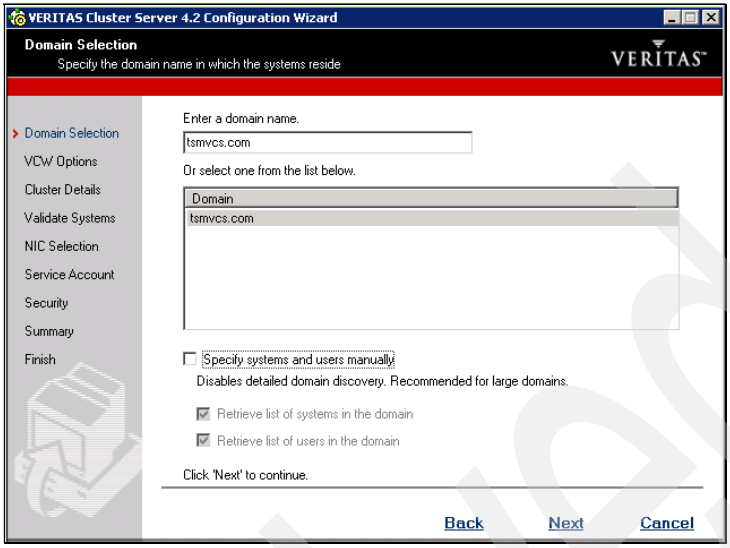

Figure 21-22 Domain and user selection

4. On the Cluster Configuration Options in Figure 21-23, we choose **Create New Cluster** and click **Next**.

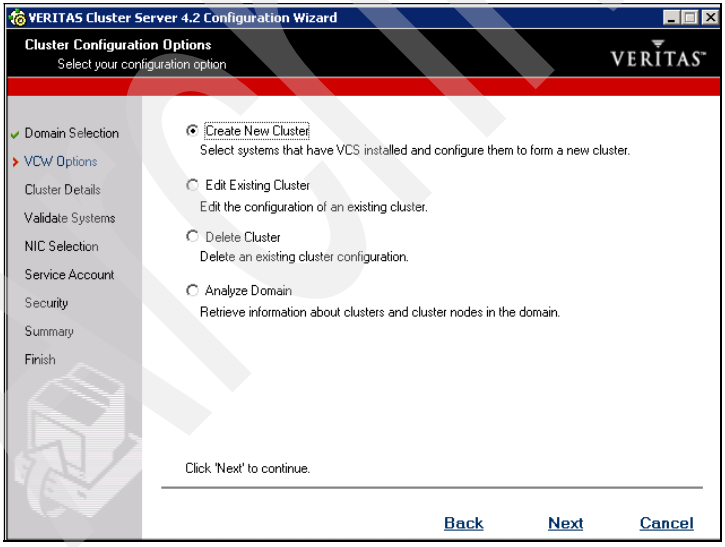

Figure 21-23 Create new cluster

Archived

5. We input the Cluster Name, the Cluster ID (accept the suggested one), the Operating System, and select the nodes that form the cluster, as shown in Figure 21-24.

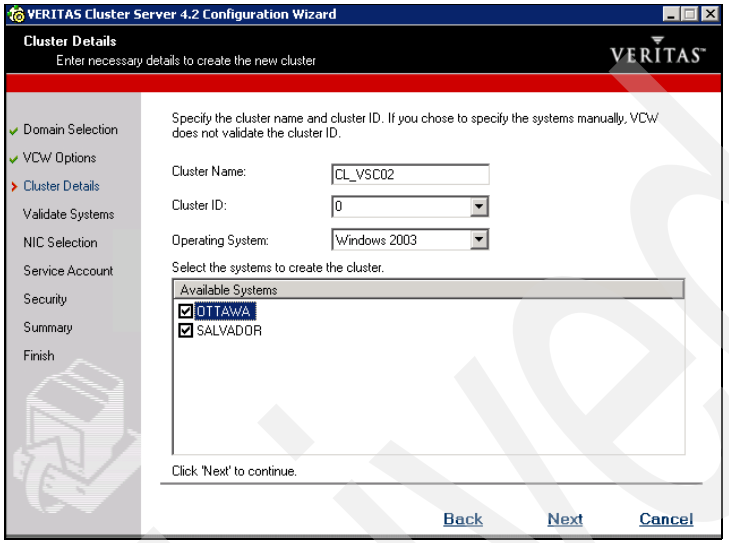

Figure 21-24 Cluster information

6. The wizard validates both nodes and when it finishes, it shows the status as in Figure 21-25. We can click **Next**.

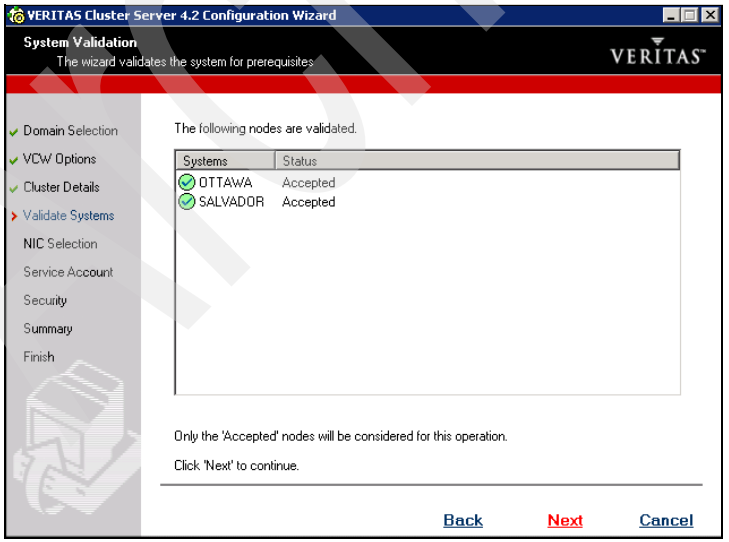

Figure 21-25 Node validation

7. We select the two private networks on each system as shown in Figure 21-26 and click **Next**.

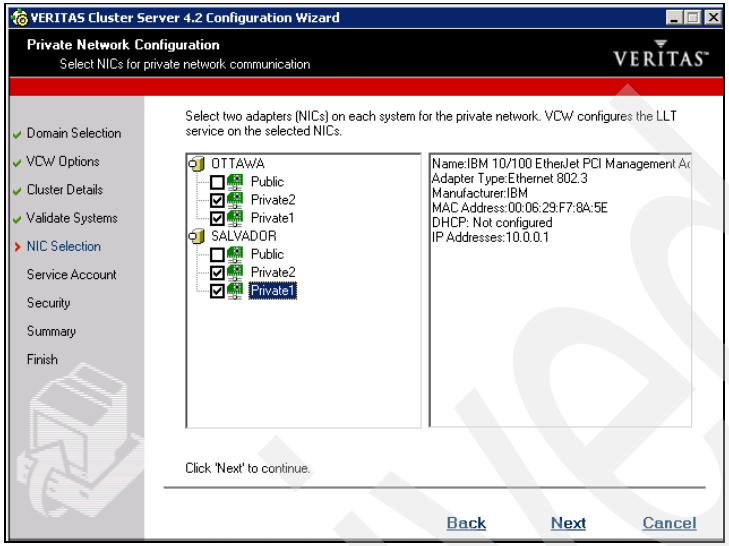

Figure 21-26 NIC selection for private communication

8. In Figure 21-27, we choose to use the Administrator account to start the VERITAS Cluster Helper Service. (However, in a production environment, we recommend to create another user.)

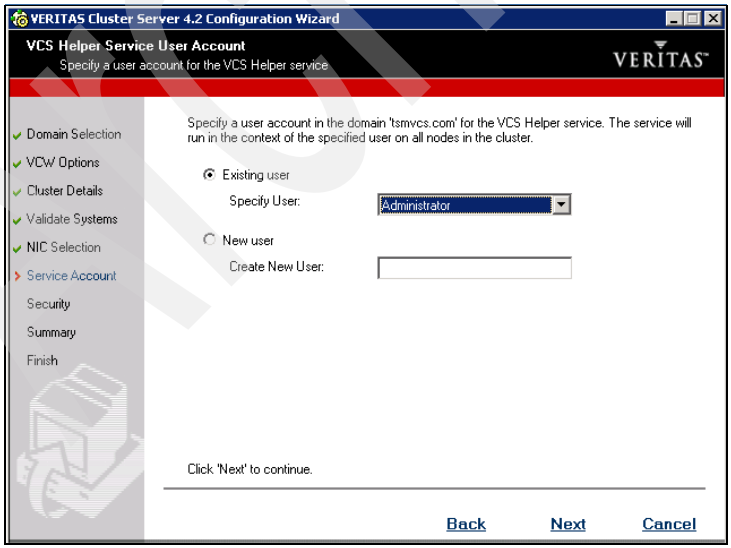

Figure 21-27 Selection of user account

9. We input the password (Figure 21-28) and click **OK**.

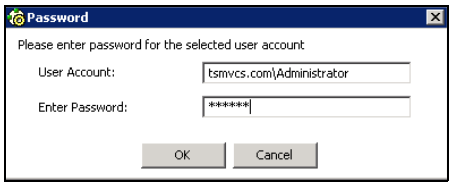

Figure 21-28 Password information

10. In Figure 21-29, we have the choice of using a secure cluster or a non-secure cluster. For our environment, we choose a non-secure environment and accept the user name and password for the VCS administrator account. The default password is *password*.

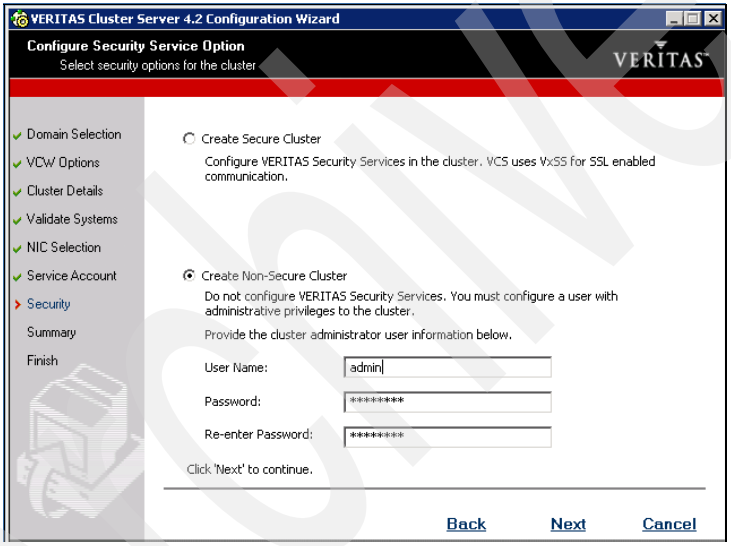

Figure 21-29 Setting up secure or non secure cluster

Archived

11. We read the summary in Figure 21-30 and click **Configure**.

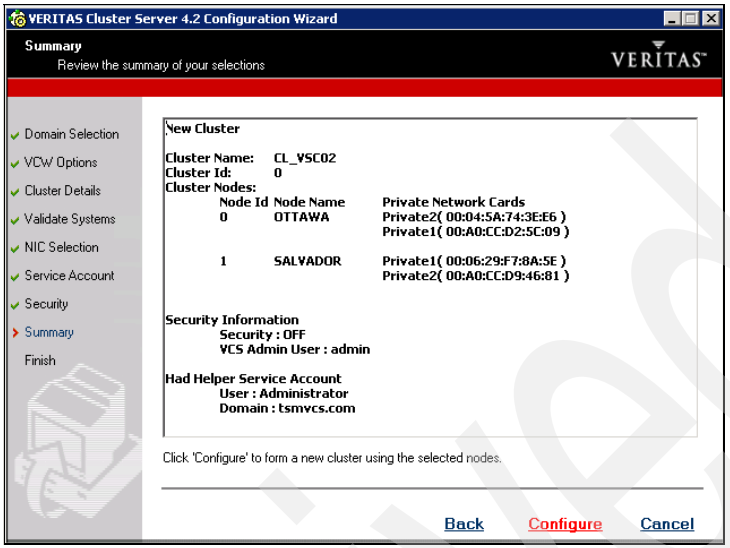

Figure 21-30 Summary prior to actual configuration

12. When the basic configuration finishes as shown in Figure 21-31, we could continue with the wizard and configure the Web console and notification. Since we are not going to use these features, we click **Finish**.

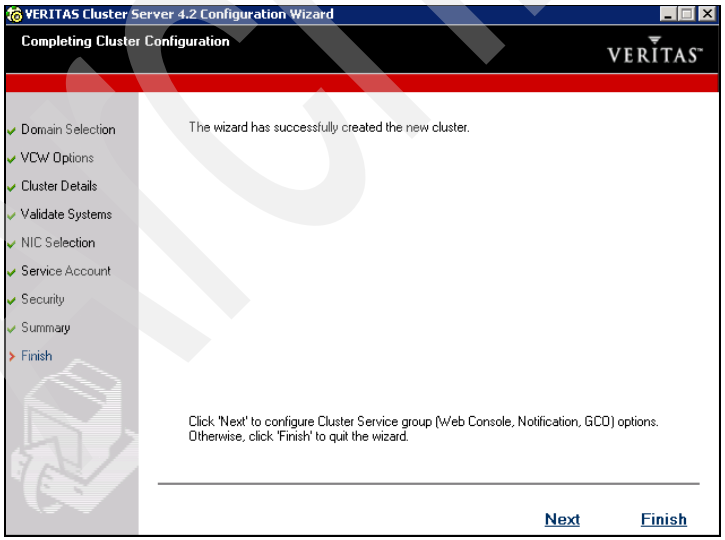

Figure 21-31 End of configuration

Archived

VERITAS Cluster Server is now created but with no resources defined. We will be creating the resources for each of our test environments in the next chapters.

21.7 Troubleshooting

VERITAS has some command line tools that can help in troubleshooting. One of them is **havol**, which queries the drives and inform, among other things, the signature and partition of the disks.

We run **havol** with the **-scsitest -l** parameters to discover the disk signatures as shown in Figure 21-32. To obtain more detailed information, we can use **havol -getdrive**, which will create a file driveinfo.txt in the path in which the command was executed.

```
C:\Documents and Settings\Administrator.TSMVCS>havol -scsitest -l

-----------------------------------------------------------------
Disk#   Signature  Port# Path# Target# Lun# Type      Status
-----------------------------------------------------------------
   0   2520421946     0     0       0    0  Internal  Active
   1    743349378     2     0       3    0  External  Active
   2    743349379     2     0       3    1  External  Active
   3    743349376     2     0       3    2  External  Active
   4    743349377     2     0       3    3  External  Active
   5    743349382     2     0       3    4  External  Active
   6    743349383     2     0       3    5  External  Active
   7    743349380     2     0       3    6  External  Active
-----------------------------------------------------------------
* 8 disks are visible from this system
```

Figure 21-32 The Havol utility - Disk signatures

To verify cluster operations, there is the **hasys** command. If we issue **hasys -display**, we will receive a detailed report of our cluster present state.

For logging, we can always refer to the Windows event viewer and to the engine logs located at %VCS_HOME%\log\engine*.txt.

For further information on other administrative tools, please refer to the *VERITAS Cluster Server 4.2 Administrator's Guide.*

Archived

22

VERITAS Cluster Server and the IBM Tivoli Storage Manager Server

This chapter discusses how we set up Tivoli Storage Manager server to work in a Windows 2003 Enterprise Edition with *Veritas Cluster Server 4.2* (VCS) for high availability.

© Copyright IBM Corp. 2005. All rights reserved.

22.1 Overview

Tivoli Storage Manager server is a cluster aware application and is supported in VCS environments.

Tivoli Storage Manager server needs to be installed and configured in a special way, as a shared application in the VCS.

This chapter covers all the tasks we follow in our lab environment to achieve this goal.

22.2 Planning and design

When planning our Tivoli Storage Manager server cluster environment, we should:

- Identify disk resources to be used by Tivoli Storage Manager. We should not partition a disk and use it with other applications that might reside in the same server, so that a problem in any of the applications will not affect the others.
- Have a TCP/IP address for the Tivoli Storage Manager server.
- Create one separate cluster resource for each Tivoli Storage Manager instance, with the corresponding disk resources.
- Check disk space on each node for the installation of Tivoli Storage Manager server. We highly recommend that the same drive letter and path be used on each machine.
- Use an additional shared SCSI bus so that Tivoli Storage Manager can provide tape drive failover support.

Note: Refer to “Appendix A” *of the IBM Tivoli Storage Manager for Windows: Administrator’s Guide* for instructions on how to manage SCSI tape failover.

For additional planning and design information, refer to *Tivoli Storage Manager for Windows Installation Guide*, *Tivoli Storage Manager Administrator’s Guide,* and *Tivoli Storage Manager for SAN for Windows Storage Agent User’s Guide*.

22.3 Lab setup

Our clustered lab environment consists of two Windows 2003 Enterprise Edition servers as described in Chapter 21, “Installing the VERITAS Storage Foundation HA for Windows environment” on page 879.

Figure 22-1 shows our Tivoli Storage Manager clustered server environment:

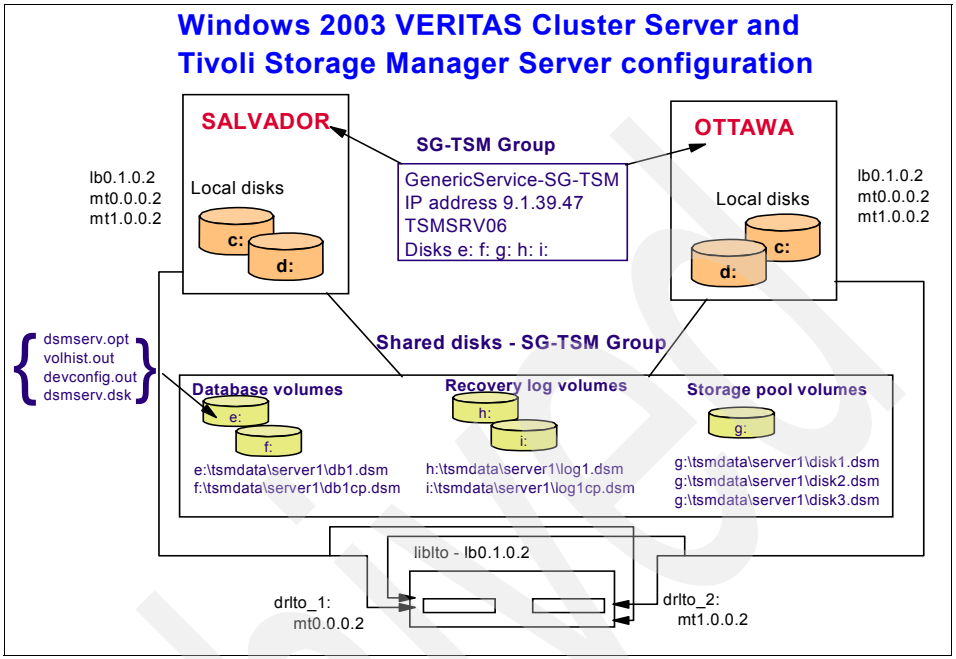

Figure 22-1 Tivoli Storage Manager clustering server configuration

Archived

Table 22-1, Table 22-2, and Table 22-3 show the specifics of our Windows VCS environment and Tivoli Storage Manager virtual server configuration that we use for the purpose of this chapter.

Table 22-1 Lab Tivoli Storage Manager server service group

Resource group SG-TSM	
TSM server name	TSMSRV06
TSM server IP address	9.1.39.47
TSM database disks[a]	e: h:
TSM recovery log disks	f: i:
TSM storage pool disk	g:
TSM service	TSM Server1

a. We choose two disk drives for the database and recovery log volumes so that we can use the Tivoli Storage Manager mirroring feature.

Table 22-2 ISC service group

Resource group SG-ISC	
ISC name	ADMCNT06
ISC IP address	9.1.39.46
ISC disk	j:
ISC services	ISC Help Service IBM WebSphere Application Server V5 - ISC Runtime Service

Archived

Table 22-3 Tivoli Storage Manager virtual server configuration in our lab

Server parameters	
Server name	TSMSRV06
High level address	9.1.39.47
Low level address	1500
Server password	itsosj
Recovery log mode	roll-forward
Libraries and drives	
Library name	LIBLTO
Drive 1	DRLTO_1
Drive 2	DRLTO_2
Device names	
Library device name	lb0.1.0.2
Drive 1 device name	mt0.0.0.2
Drive 2 device name	mt1.0.0.2
Primary Storage Pools	
Disk Storage Pool	SPD_BCK (nextstg=SPT_BCK)
Tape Storage Pool	SPT_BCK
Copy Storage Pool	
Tape Storage Pool	SPCPT_BCK
Policy	
Domain name	STANDARD
Policy set name	STANDARD
Management class name	STANDARD
Backup copy group	STANDARD (default, DEST=SPD_BCK)
Archive copy group	STANDARD (default)

Archived

22.3.1 Installation of IBM tape device drivers

The two servers are attached to the *Storage Area Network*, so that both can see the IBM 3582 Tape Library as well as the two IBM 3580 tape drives.

Since IBM Tape Libraries use their own device drivers to work with Tivoli Storage Manager, we have to download and install the latest available version of the IBM LTO drivers for 3582 Tape Library and 3580 Ultrium 2 tape drives.

We use the Windows device manager menu to update the device drives, specifying the path to where we made the download.

We do not show the whole installation process in this book. Refer to the *IBM Ultrium Device Drivers Installation and User's Guide* for a detailed description of this task.

After the successful installation of the drivers, both nodes recognize the 3582 medium changer and the 3580 tape drives, as shown in Table 22-2.

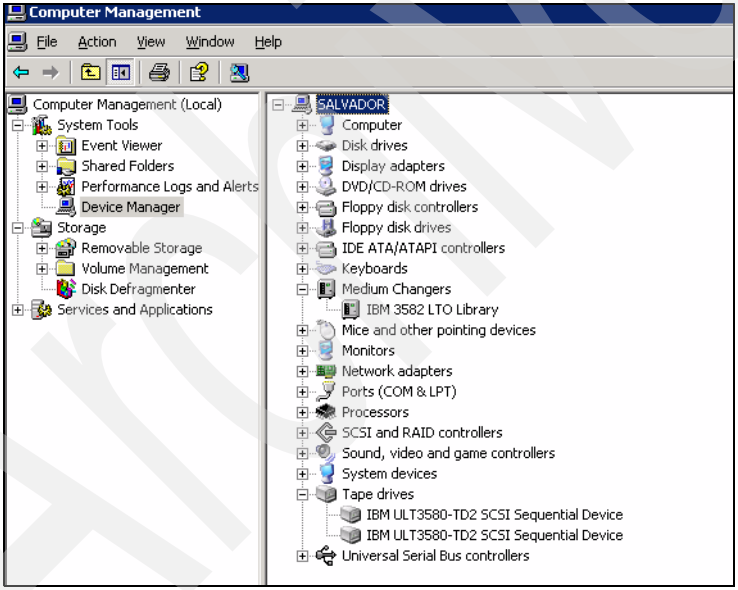

Figure 22-2 IBM 3582 and IBM 3580 device drivers on Windows Device Manager

Archived

22.4 Tivoli Storage Manager installation

We install Tivoli Storage Manager on the local disk of each node, one at a time, since there will be a reboot at the end. We use the same drive letter for each node.

After Tivoli Storage Manager server is installed on both nodes, we configure the VCS for the failover.

To install Tivoli Storage Manager, we follow the same process described in 5.3.1, “Installation of Tivoli Storage Manager server” on page 80

22.5 Configuration of Tivoli Storage Manager for VCS

When the installation of Tivoli Storage Manager packages on both nodes of the cluster is completed, we can proceed with the configuration.

The Tivoli Storage Manager configuration wizard does not recognize the VCS as it does with MSCS. The configuration is done the same way we would do it for single servers with no cluster installed. The important factor here is to inform the system of the correct location of the common files.

When we start the configuration procedure on the first node, a Tivoli Storage Manager server instance is created and started. For the second node, we need to create a server instance and the service, using the same files on the shared folders for the database, log, and storage pool. This can be done using the configuration wizard in the management console again or manually. We discuss both methods here.

22.5.1 Configuring Tivoli Storage Manager on the first node

As for now, our cluster environment has no resources. The disks are still being seen by both servers simultaneously. To avoid disk corruption, we shut down one of the servers during the configuration of the first node. In production environments, VCS may already be configured with disk drives. In this case, make sure the disks that are going to be used by Tivoli Storage Manager are all hosted by one of the servers.

1. We open the Tivoli Storage Manager Management Console (**Start** →**Programs** →**Tivoli Storage Manager** →**Management Console**) to start the initialization.

2. The Initial Configuration Task List for the Tivoli Storage Manager menu, Figure 22-3, shows a list of the tasks needed to configure a server with all of the basic information. To let the wizard guide us throughout the process, we select **Standard Configuration**. We then click **Start**.

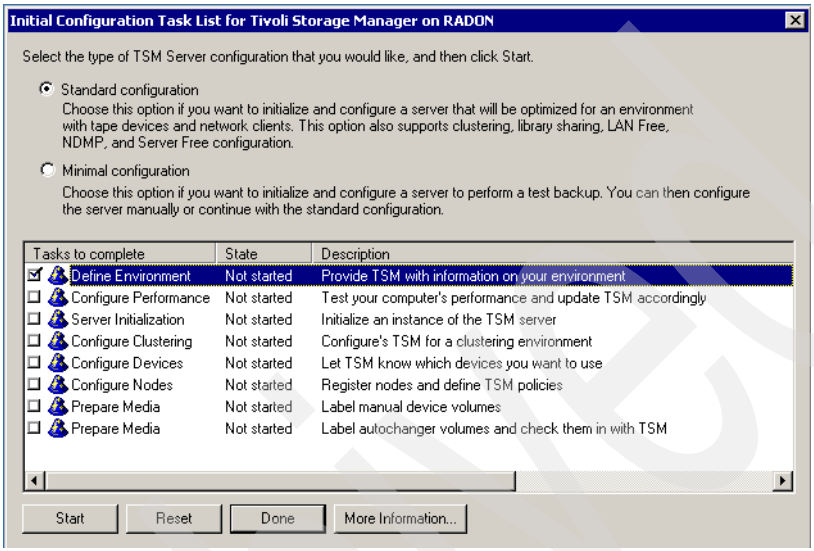

Figure 22-3 Initial Configuration Task List

3. The Welcome menu for the first task, **Define Environment**, displays (Figure 22-4). We click **Next**.

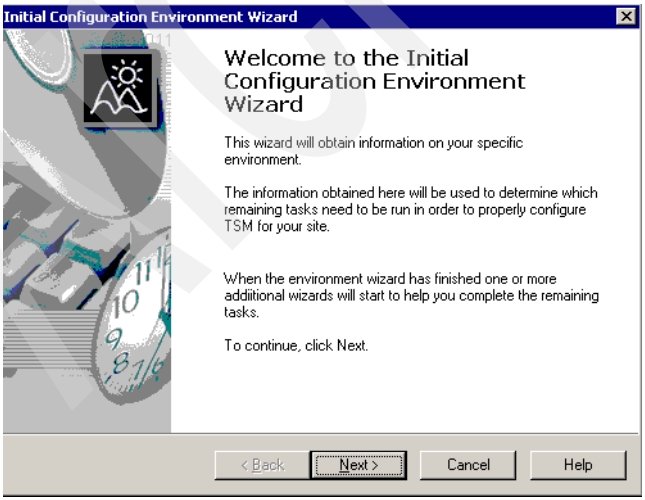

Figure 22-4 Welcome Configuration wizard

4. To have additional information displayed during the configuration, we select **Yes** and click **Next** as shown in Figure 22-5.

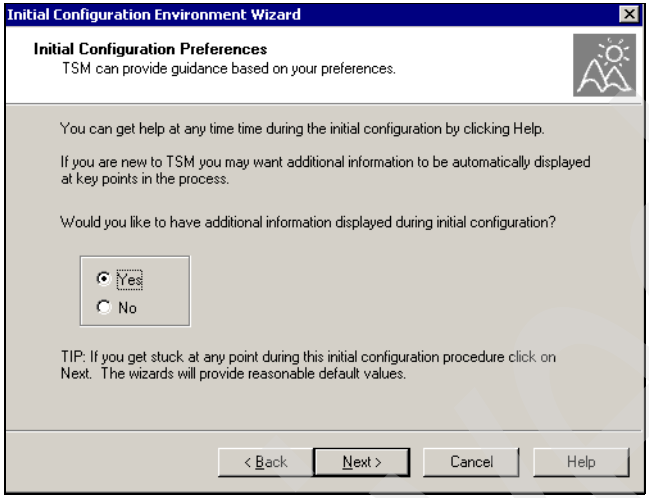

Figure 22-5 Initial configuration preferences

5. Tivoli Storage Manager can be installed Standalone (for only one client), or Network (when there are more clients). In most cases we have more than one client. We select **Network** and then click **Next** as shown in Figure 22-6.

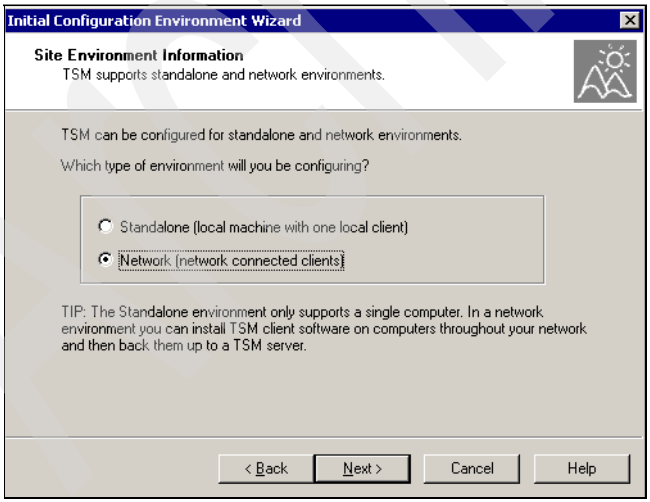

Figure 22-6 Site environment information

6. The Initial Configuration Environment is done. We click **Finish** in Figure 22-7.

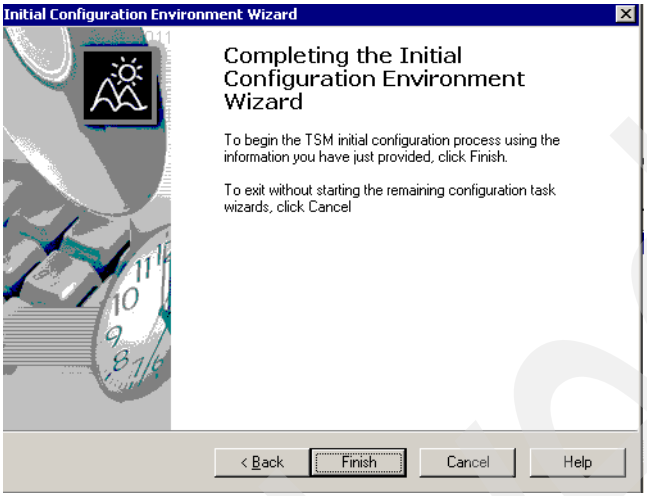

Figure 22-7 Initial configuration

7. The next task is to run the Performance Configuration Wizard. In Figure 22-8 we click **Next**.

Figure 22-8 Welcome Performance Environment wizard

Archived

8. In Figure 22-9 we provide information about our own environment. Tivoli Storage Manager will use this information for tuning. For our lab we used the defaults. In a production server, we would select the values that best fit the environment. We click **Next**.

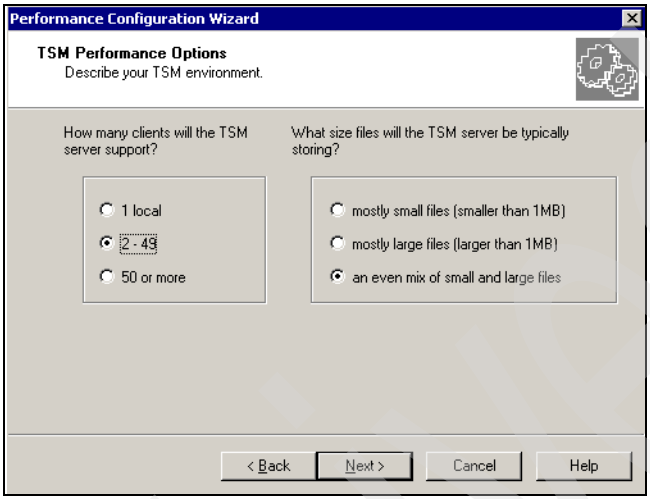

Figure 22-9 Performance options

9. The wizard starts to analyze the hard drives as shown in Figure 22-10. When the process ends, we click **Finish**.

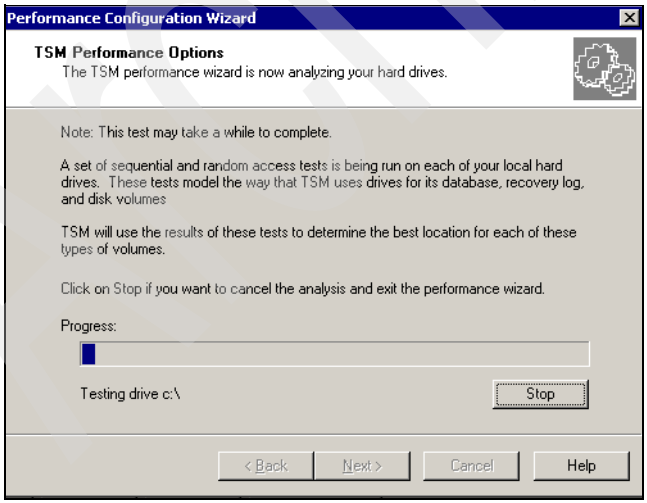

Figure 22-10 Drive analysis

10. The Performance Configuration task completes as shown in Figure 22-11.

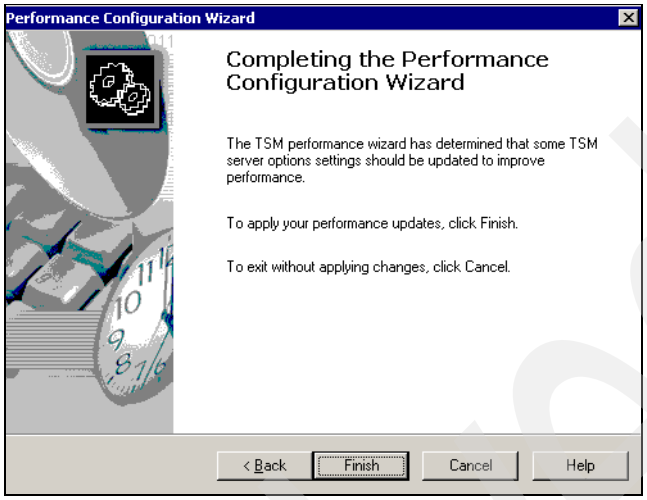

Figure 22-11 Performance wizard

11. The next step is the initialization of the Tivoli Storage Manager server instance. In Figure 22-12 we click **Next**.

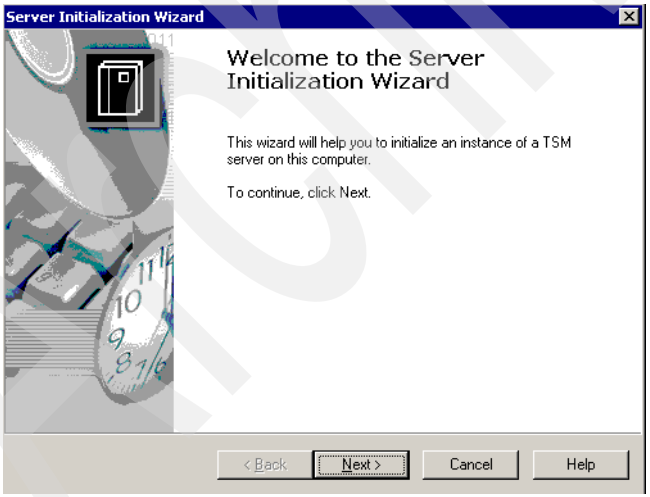

Figure 22-12 Server instance initialization wizard

Archived

12. In Figure 22-13 we select the directory where the files used by Tivoli Storage Manager server will be placed. It is possible to choose any disk on the Tivoli Storage Manager Service Group. We change the drive letter to use e: and click **Next**.

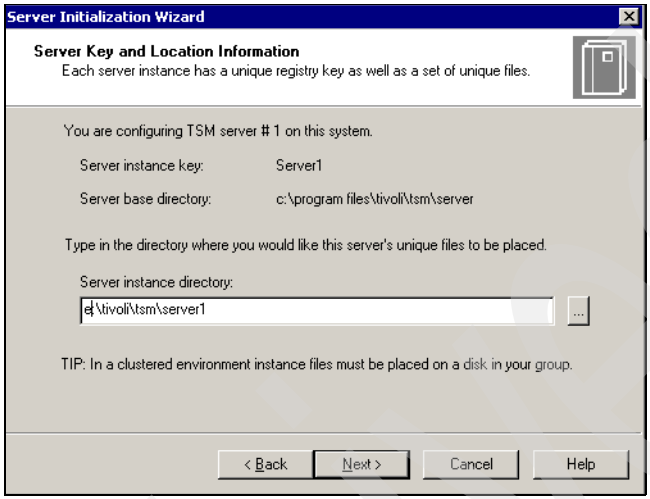

Figure 22-13 Server initialization wizard

13. In Figure 22-14 we type the complete path and sizes of the initial volumes to be used for database, recovery log, and disk storage pools. We base our values on Table 22-1 on page 906, where we describe our cluster configuration for Tivoli Storage Manager server.

    We also check the two boxes on the two bottom lines to let Tivoli Storage Manager create additional volumes as needed.

    With the selected values we will initially have a 1000 MB size database volume with name db1.dsm, a 500 MB size recovery log volume called log1.dsm, and a 5 GB size storage pool volume of name disk1.dsm. If we need, we can create additional volumes later.

    We input our values and click **Next**.

Archived

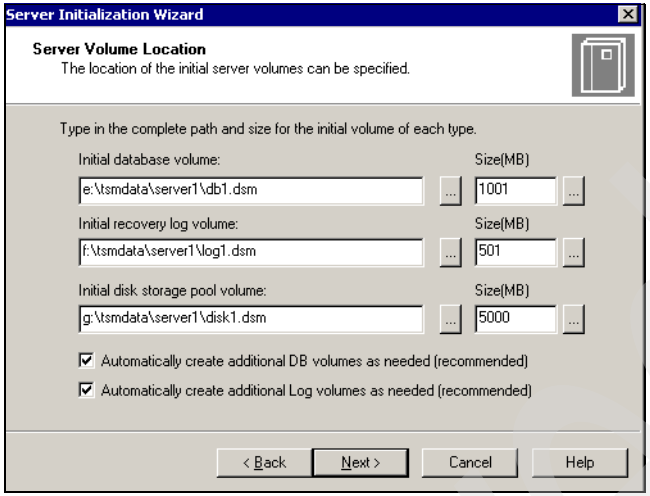

Figure 22-14 Server volume location

14. On the server service logon parameters shown in Figure 22-15, we select the Windows account and user ID that Tivoli Storage Manager server instance will use when logging onto Windows. We recommend to leave the defaults and click **Next**.

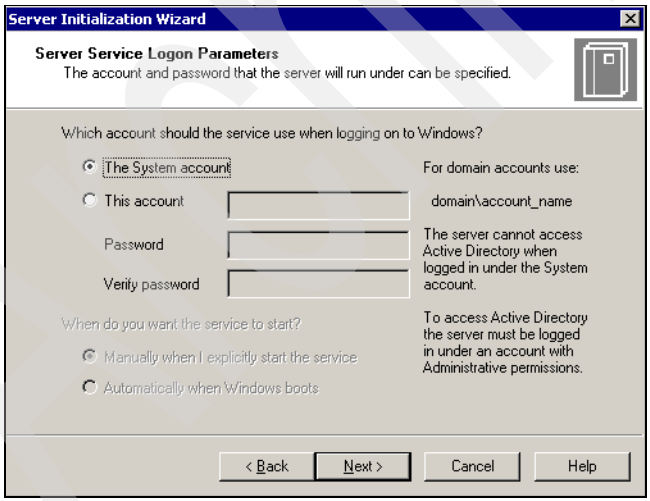

Figure 22-15 Server service logon parameters

15. In Figure 22-16, we provide the server name and password. The server password is used for server-to-server communications. We will need it later on with Storage Agent.This password can also be set later using the administrator interface. We click **Next**.

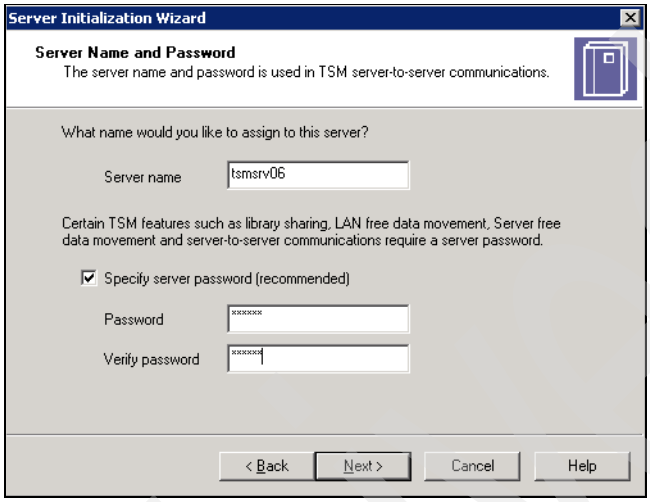

Figure 22-16 Server name and password

16. We click **Finish** in Figure 22-17 to start the process of creating the server instance.

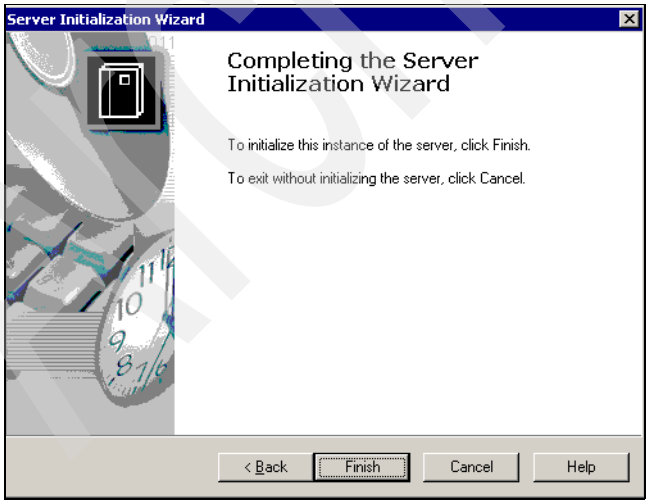

Figure 22-17 Completing the Server Initialization Wizard

17. The wizard starts the process of the server initialization and shows a progress bar as in Figure 22-18.

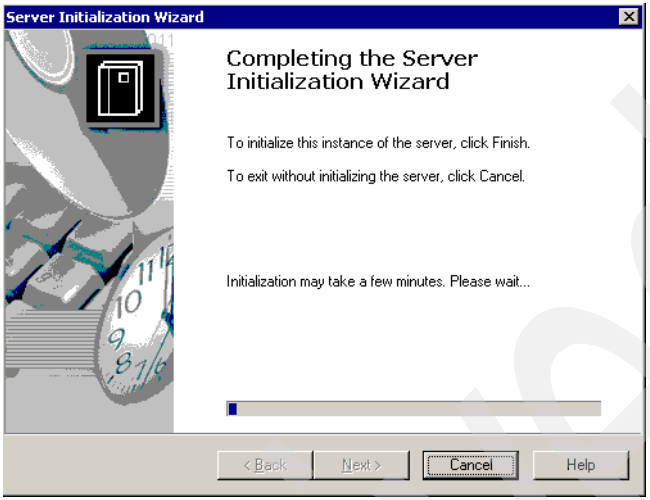

Figure 22-18 Completing the server installation wizard

18. If the initialization ends without any errors, we receive the following informational message (Figure 22-19). We click **OK**.

Figure 22-19 TSM server has been initialized

At this time, we could continue with the initial configuration wizard, to set up devices, nodes, and label media. However, for the purpose of this book, we will stop here. We click **Cancel** when the Device Configuration welcome menu displays.

So far Tivoli Storage Manager server instance is installed and started on SALVADOR. If we open the Tivoli Storage Manager console, we can check that the service is running as shown in Figure 22-20.

Archived

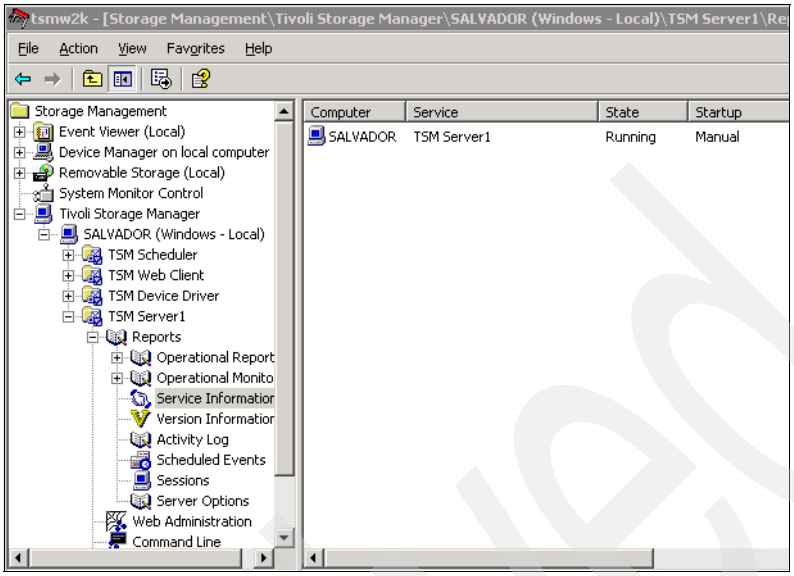

Figure 22-20 Tivoli Storage Manager console

Important: Before starting the initial configuration for Tivoli Storage Manager on the second node, you must stop the instance on the first node.

19. We stop the Tivoli Storage Manager server instance on SALVADOR before going on with the configuration on OTTAWA.

22.5.2 Configuring Tivoli Storage Manager on the second node

In this section we describe two ways to configure Tivoli Storage Manager on the second node of the VCS: by using the wizard, and by manual configuration.

Using the wizard to configure the second node

We can use again the wizard. We will need to delete the files created for the database, logs, and storage pools in drives E: F: and G:. To use the wizard, we then do the following tasks:

1. Delete the files under e:\tsmdata, f:\tsmdata and g:\tsmdata.
2. Run the wizard, repeating steps 1 through 18 of “Configuring Tivoli Storage Manager on the first node” on page 909. The wizard will detect files under e:\program files\tivoli\tsm\server1 and ask to overwrite them. We choose to overwrite.
3. In the end, we confirm that the service is set to manual and is running.

Manually configuring the second node

The server initialization process creates two keys in the registry besides creating the database, recovery log, and storage pool files. After the configuration of the first node, the only thing we would need is to recreate the keys in the registry.

> **Attention:** *Using the registry incorrectly can cause serious damage to the system.* We advise that only experienced administrators should run the following steps, at their own risk, and taking all the necessary precautions.

To copy the keys from one server to the other, we would do the following tasks:

1. Run regedit.exe on the first node.
2. Export the following keys to files in a shared disk. These files have a **reg** extension:
   - For the Tivoli Storage Manager Server instance, **HKLM\SOFTWARE\IBM\ADSM\CurrentVersion\Server\Server1**
   - For the Tivoli Storage Manager Server1 service, **HKLM\SYSTEM\CurrentControlSet\Services\TSM Server1**
3. Double-click the files, on the second node.
4. Boot the second node.
5. Start Tivoli Storage Manager Server1 instance and test. If there are disks already configured in the VCS, move the resources to the second node first.

22.6 Creating service group in VCS

Now that Tivoli Storage Manager Server is installed and configured on both nodes, we will create a service group in VCS using the Application Configuration Wizard.

1. We click **Start** →**Programs** →**VERITAS** →**VERITAS Cluster Service** → **Application Configuration Wizard**.
2. We review the welcome page in Figure 22-21 and click **Next**.

Archived

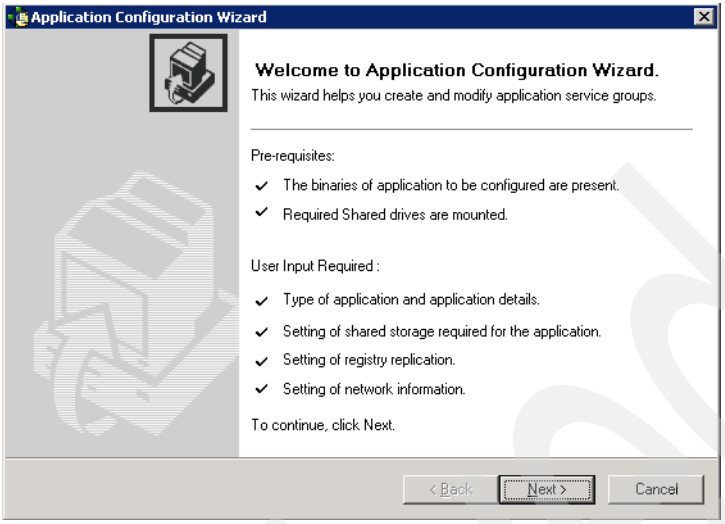

Figure 22-21 Starting the Application Configuration Wizard

3. Since we do not have any group created, we are able only to check the Create service group option as shown in Figure 22-22. We click **Next**.

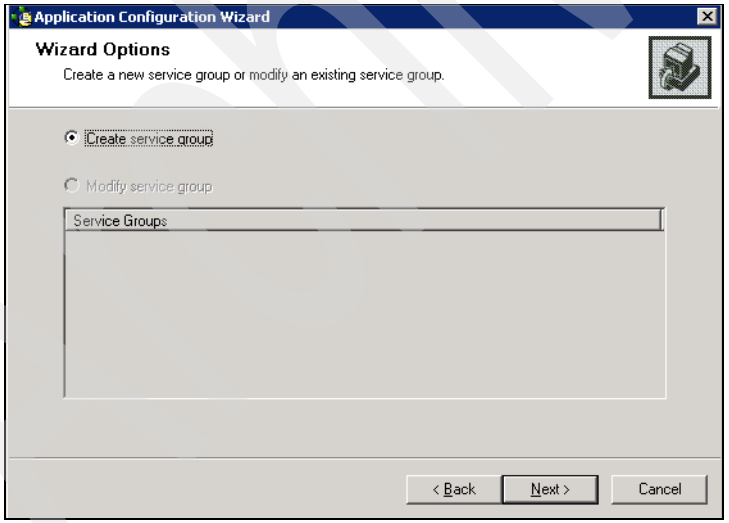

Figure 22-22 Create service group option

4. We specify the group name and choose the servers that will hold them, as in Figure 22-23. We can set the priority between the servers, moving them with the down and up arrows. We click **Next**.

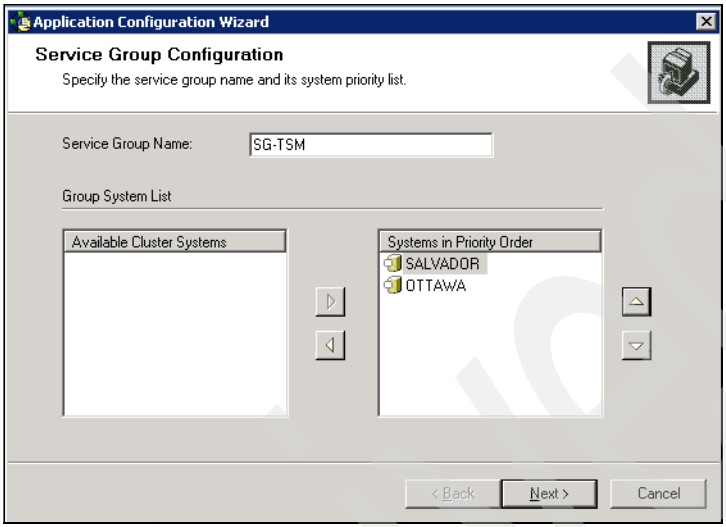

Figure 22-23 Service group configuration

5. Since it is the first time we are using the cluster after it was set up, we receive a warning saying that the configuration is in read-only mode and needs to be changed, as shown in Figure 22-24. We click **Yes**.

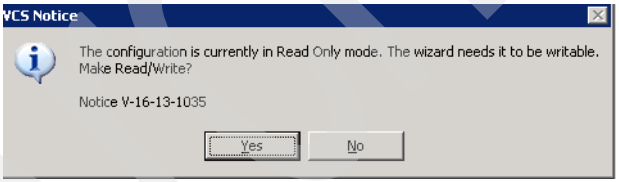

Figure 22-24 Change configuration to read-write

6. The wizard will start a process of discovering all necessary objects to create the service group, as shown in Figure 22-25. We wait until this process ends.

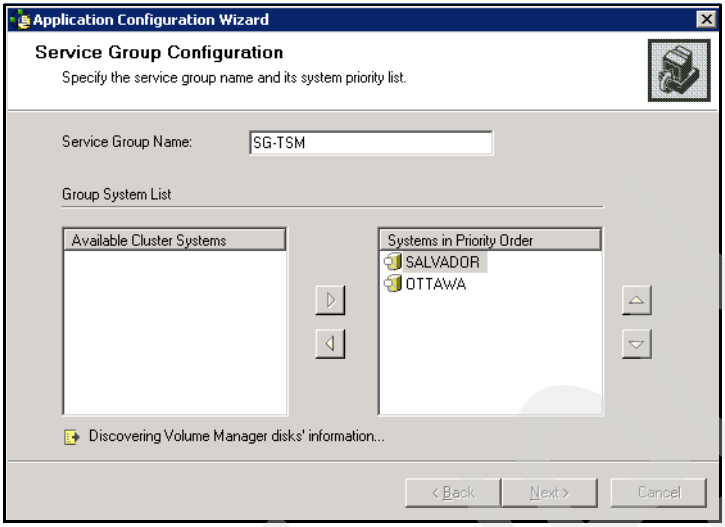

Figure 22-25 Discovering process

7. We then define what kind of application group this is. In our case, it is a generic service application, since it is the Tivoli Storage Manager Server 1 service in Windows that need to be brought online/offline by the cluster during a failover. We choose **Generic Service** from the drop-down list in Figure 22-26 and click **Next**.

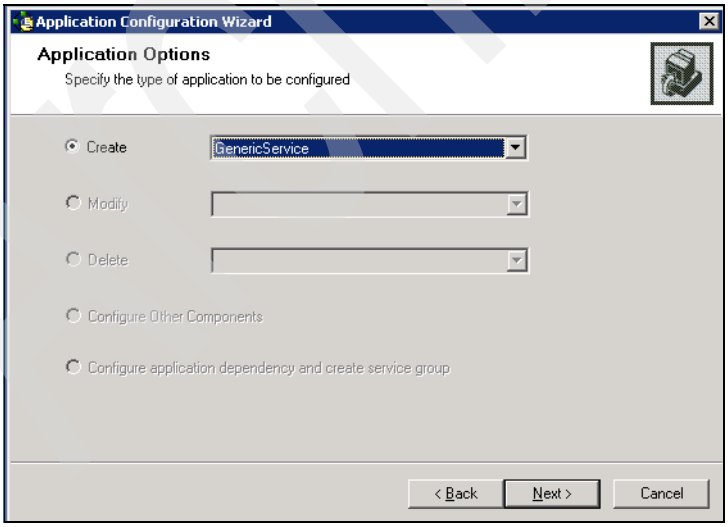

Figure 22-26 Choosing the kind of application

8. We click the button next to the **Service Name** line and choose the TSM Server1 service from the drop-down list as shown in Figure 22-27.

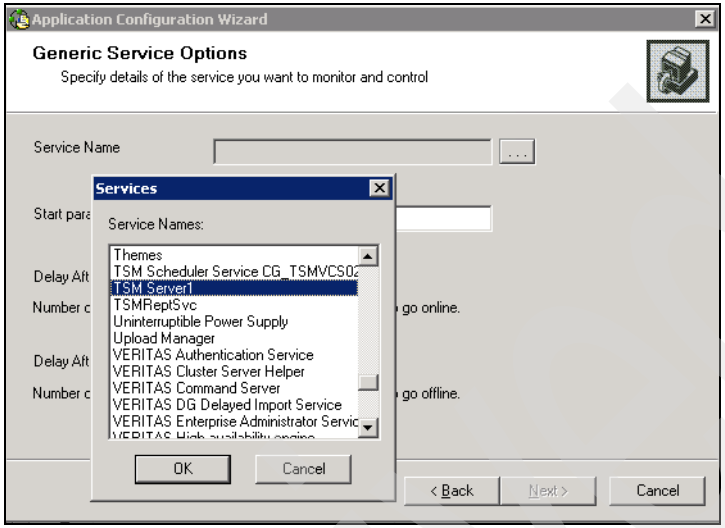

Figure 22-27 Choosing TSM Server1 service

9. We confirm the name of the service chosen and click **Next** in Figure 22-28.

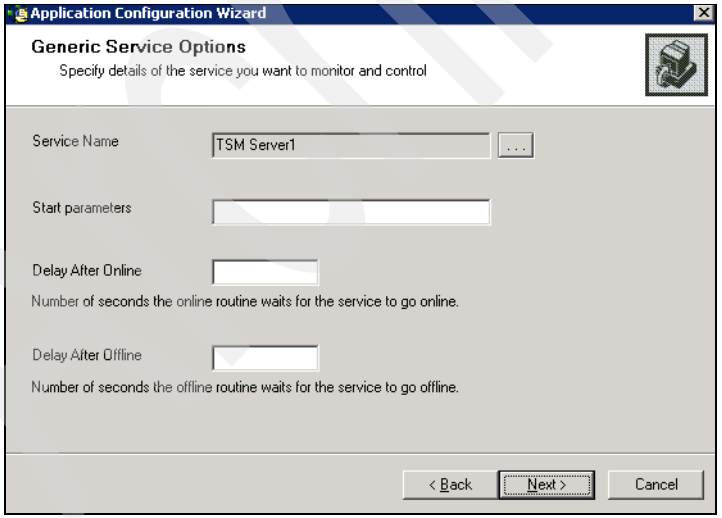

Figure 22-28 Confirming the service

10. In Figure 22-29 we choose to start the service with the LocalSystem account.

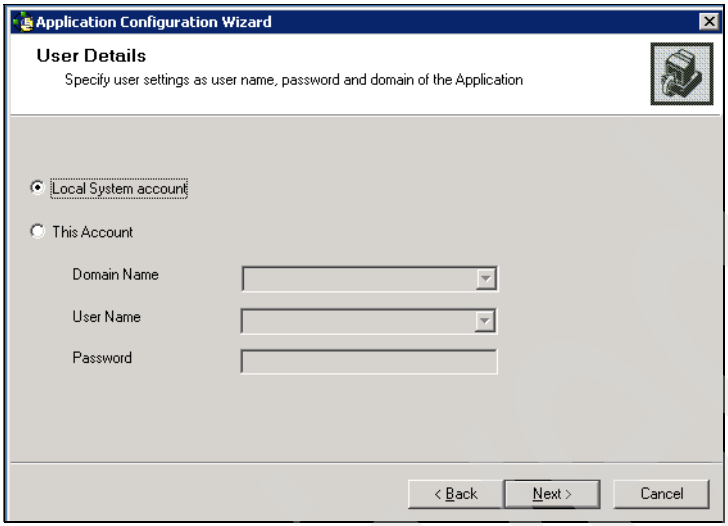

Figure 22-29 Choosing the service account

11. We select the drives that will be used by our Tivoli Storage Manager server. We refer to Table 22-1 on page 906 to confirm the drive letters. We select the letters as in Figure 22-30 and click **Next**.

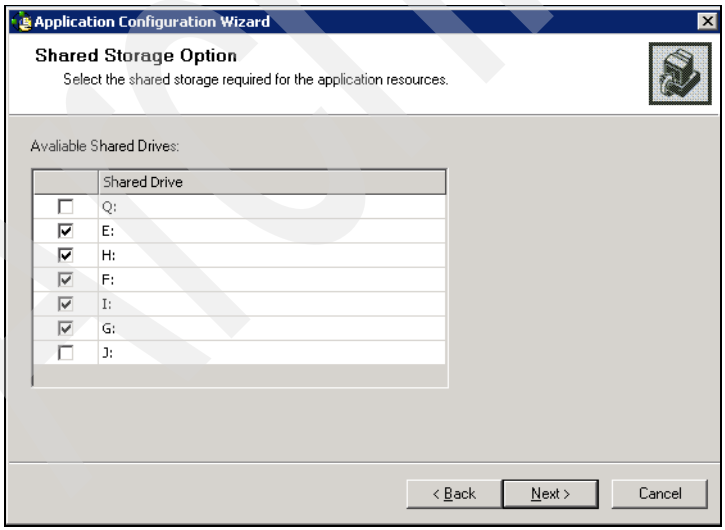

Figure 22-30 Selecting the drives to be used

12. We receive a summary of the application resource with the name and user account as in Figure 22-31. We confirm and click **Next**.

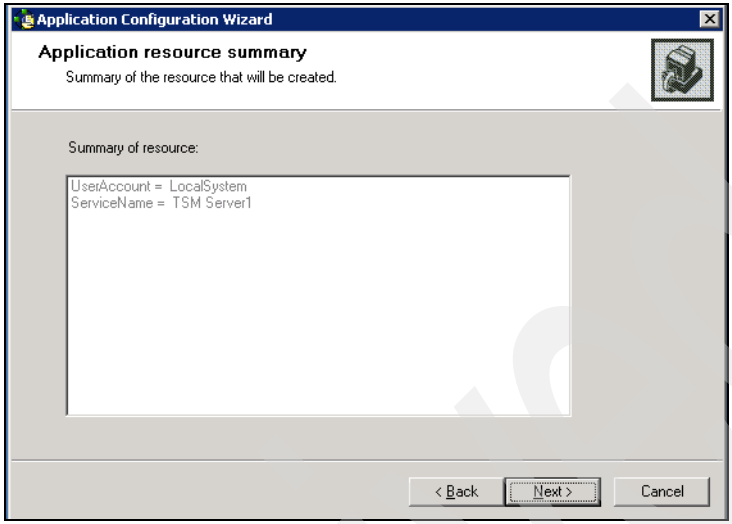

Figure 22-31 Summary with name and account for the service

13. We need two more resources for the TSM Group: IP and a Name. So in Figure 22-32 we will choose **Configure Other Components** and then click **Next**.

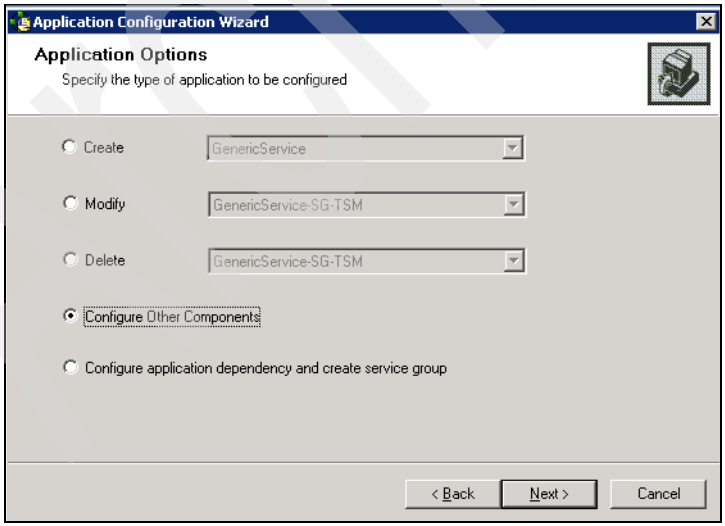

Figure 22-32 Choosing additional components

14. In Figure 22-33 we choose to create **Network Component** (IP address) and **Lanman Component** (Name) and click **Next**.

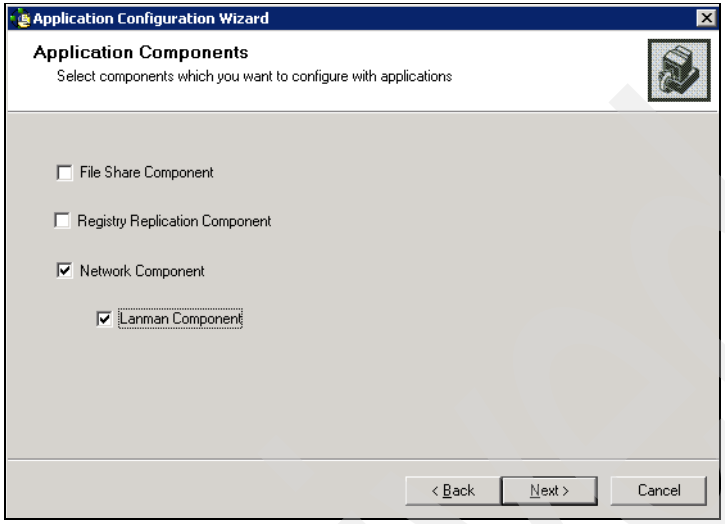

Figure 22-33 Choosing other components for IP address and Name

15. In Figure 22-34 we specify the name of the Tivoli Storage Manager server and the IP address we will use to connect our clients and click **Next**. We refer to Table 22-1 on page 906 for the necessary information.

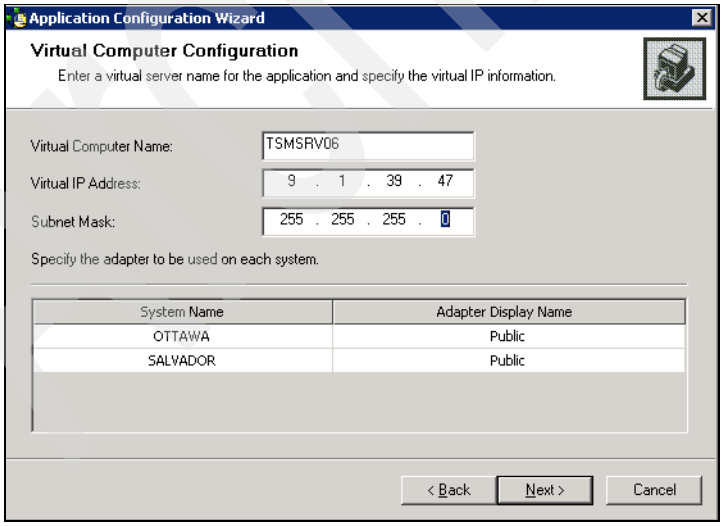

Figure 22-34 Specifying name and IP address

16. We now do not need any other resources to be configured. We choose **Configure application dependency and create service group** in Figure 22-35 and click **Next**.

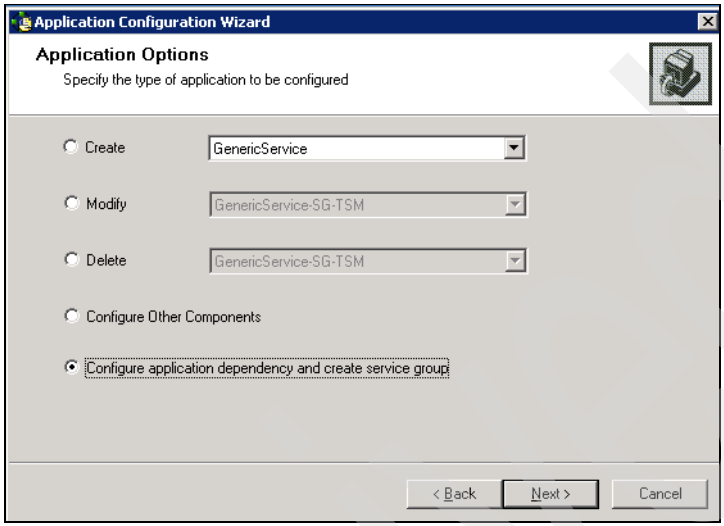

Figure 22-35 Completing the application options

17. The wizard brings up the summary of all resources to be created, as shown in Figure 22-36.

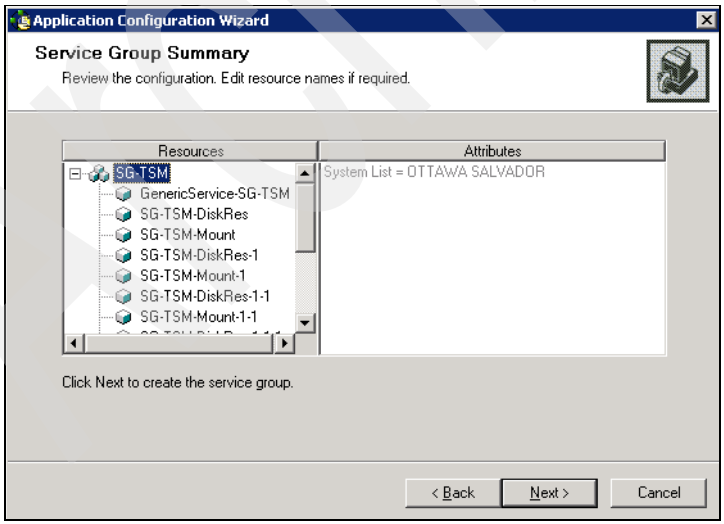

Figure 22-36 Service Group Summary

18. The default names of the resources are not very clear, so with the F2 key we change the resources, naming the drives and disk resources with the corresponding letter as shown in Figure 22-37. We have to be careful and match the right disk with the right letter. We refer to the **hasys** output in Figure 21-32 on page 902 and look in the attributes list to match them.

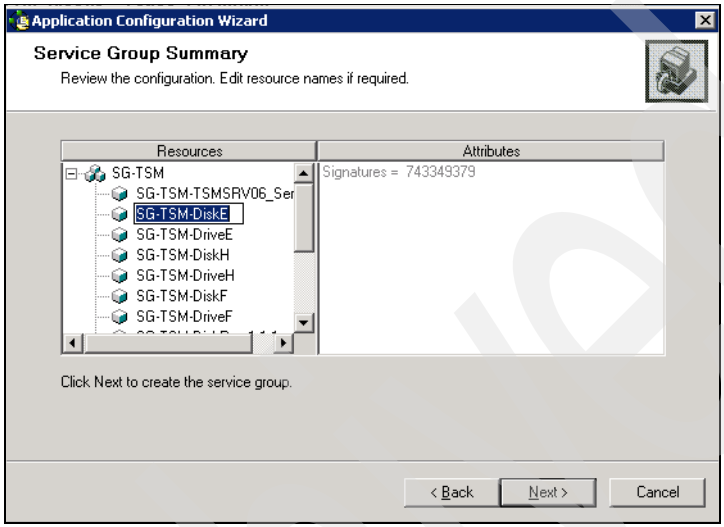

Figure 22-37 Changing resource names

19. We confirm we want to create the service group by clicking **Yes** in Figure 22-38.

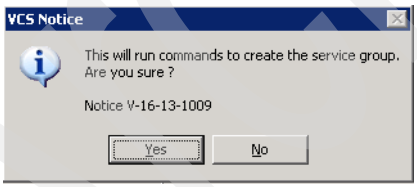

Figure 22-38 Confirming the creation of the service group

Archived

20. The process begins as shown in Figure 22-39.

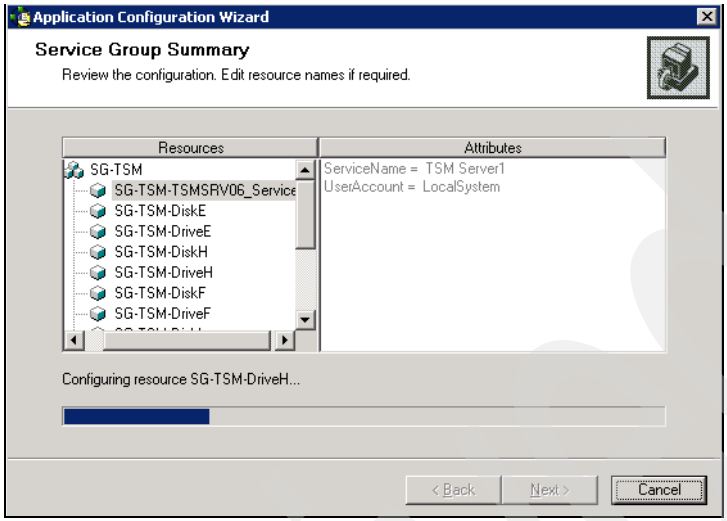

Figure 22-39 Creating the service group

21. When the process completes, we confirm that we want to bring the resources online and **click Finish** as shown in Figure 22-40. We could also uncheck the **Bring the service group online** option and do it in the Java Console.

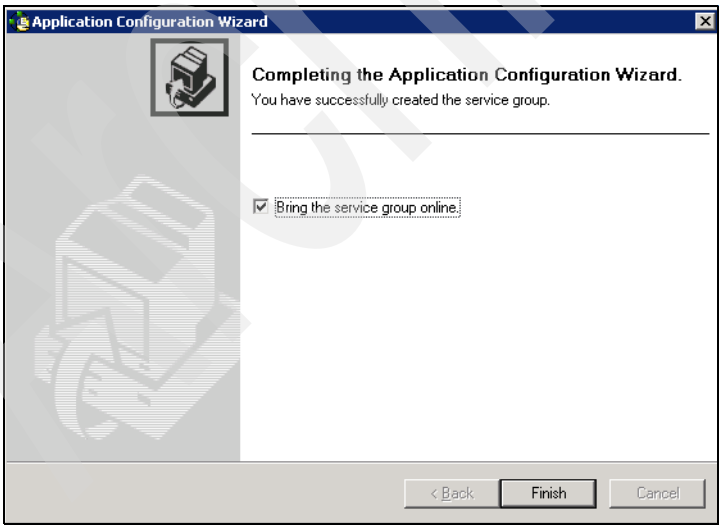

Figure 22-40 Completing the wizard

22. We now open the Java Console to administer the cluster and check configurations. To open the Java Console, either click the desktop icon or select **Start** →**Programs** →**VERITAS** →**VERITAS Cluster Manager (Java Console)**. The cluster monitor opens as shown in Figure 22-41.

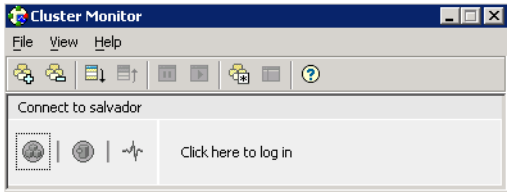

Figure 22-41 Cluster Monitor

23. We log on the console, specifying name and password, and the Java Console (also known as the Cluster Explorer) is displayed as shown in Figure 22-42. We navigate in the console and check the resources created.

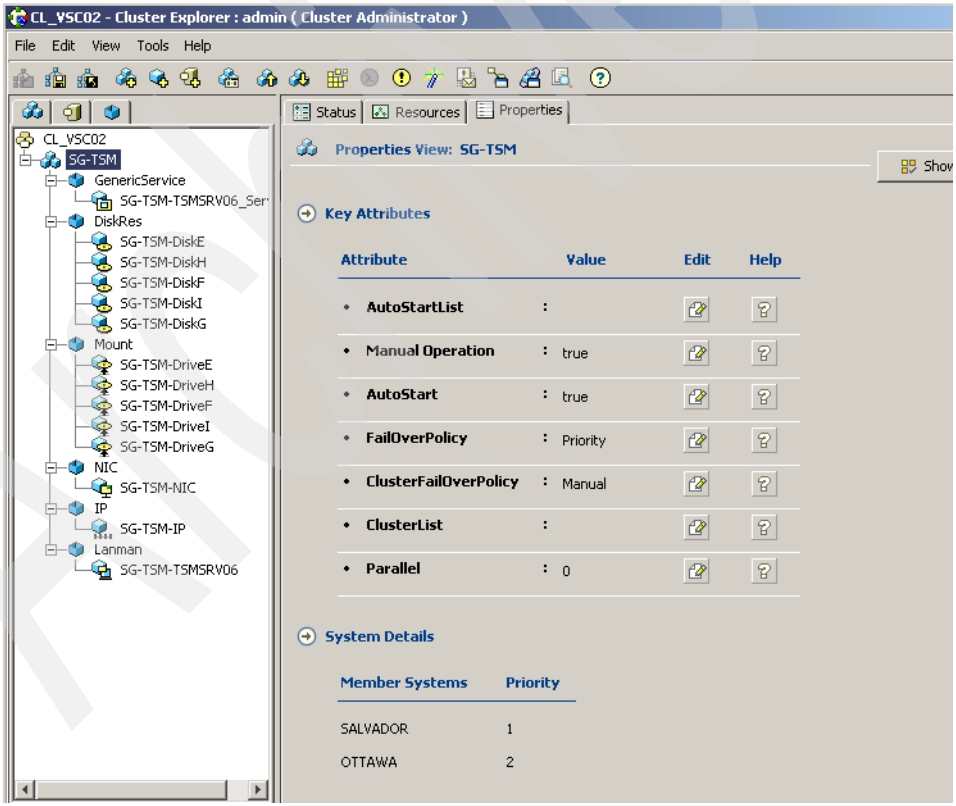

Figure 22-42 Resources online

24. If we click the **Resources** tab on the right panel we will see the dependencies created by the wizard, as shown in Figure 22-43, which illustrates the order that resources are brought online, from bottom to top.

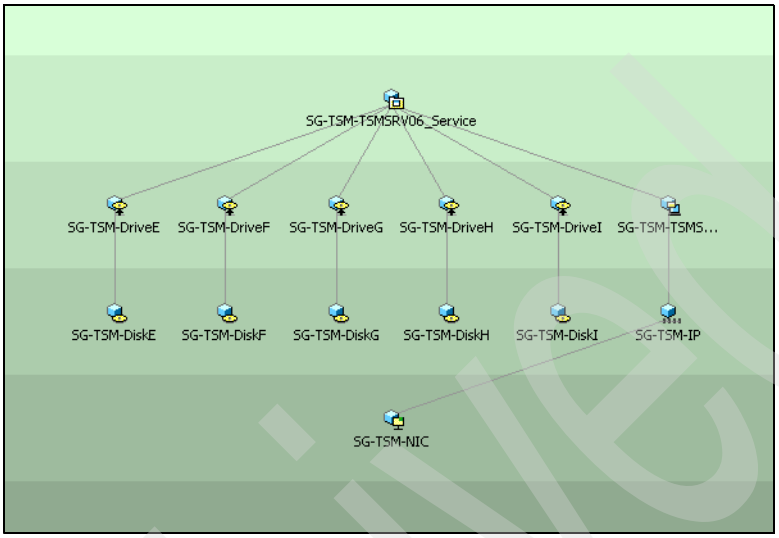

Figure 22-43 Link dependencies

22.7 Testing the Cluster

To test the cluster functionality, we use the Cluster Explorer and perform the following tasks:

- Switching the service group from one server to another. We verify that resources fail over and are brought online on the other node.
- Switching the service group to one node and stopping the Cluster service. We verify that all resources fail over and come online on the other node
- Switching the service group to one node and shutting it down. We verify that all resources fail over and come online on the other node.
- Switching the service group to one node and removing the public network cable from that node. We verify that the groups will fail over and come online on the other node.

Archived

22.8 IBM Tivoli Storage Manager Administrative Center

With IBM Tivoli Storage Manager V5.3.0, the Administrative Web interface has been replaced with the Administrative Center. This is a Web-based interface to centrally configure and manage any Tivoli Storage Manager V5.3.0 server.

IBM Tivoli Storage Manager Administrative Center consists of two components:

- The Integrated Solutions Console (ISC)
- The Administration Center

ISC allows us to install components provided by multiple IBM applications, and access them from a single interface. It is a requirement to install the Administrative Center.

22.8.1 Installing the Administrative Center in a clustered environment

To install the Administrative Center in the cluster, we follow the same procedures outlined in “Installing the ISC and Administration Center for clustering” on page 92.

For this installation, we will be using drive J, which is still not being seen by VCS. Again, in order to avoid disk corruption, since both servers see this drive at this time, we will perform each installation with only one server up at a time, and bring them both online before configuring the service group in VCS.

22.8.2 Creating the service group for the Administrative Center

After installing ISC and the Administration Center on both nodes, we will use the Application Configuration Wizard again to create the service group with all the necessary resources:

1. We click **Start** →**Programs** →**VERITAS** →**VERITAS Cluster Service** →**Application Configuration Wizard**.
2. We review the welcome page in Figure 22-44 and click **Next**.

Archived

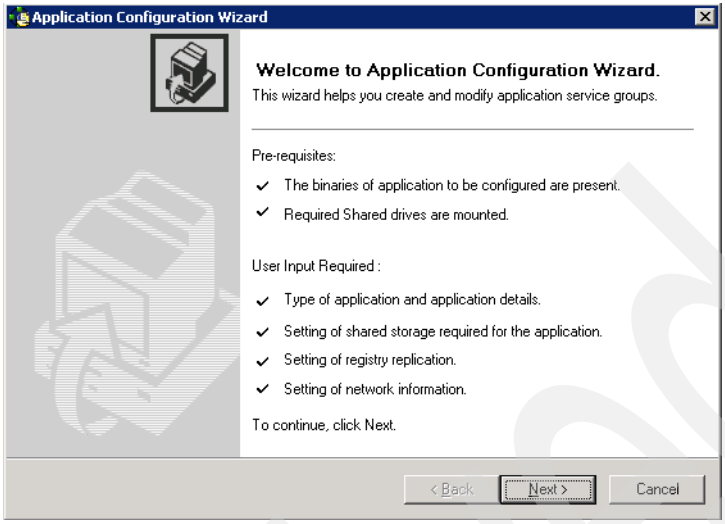

Figure 22-44 Starting the Application Configuration Wizard

3. We select the **Create service group** option as shown in Figure 22-45 and click **Next**.

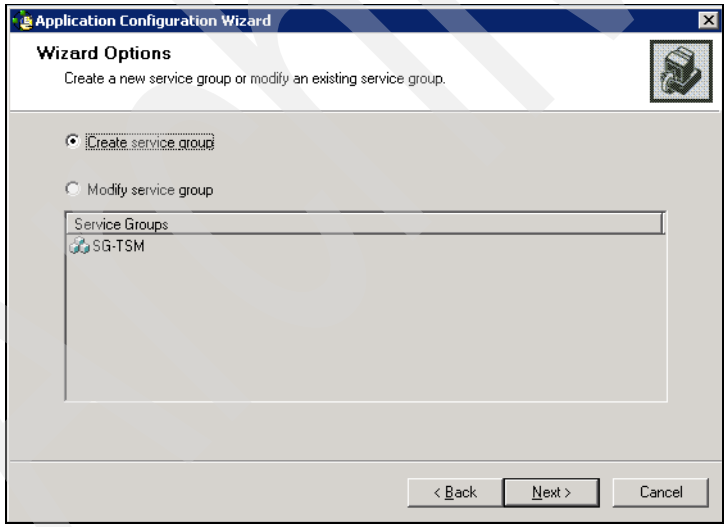

Figure 22-45 Create service group option

4. We specify the group name and choose the servers that will hold them, as in Figure 22-46. We can set the priority between the servers, moving them with the down and up arrows. We click **Next**.

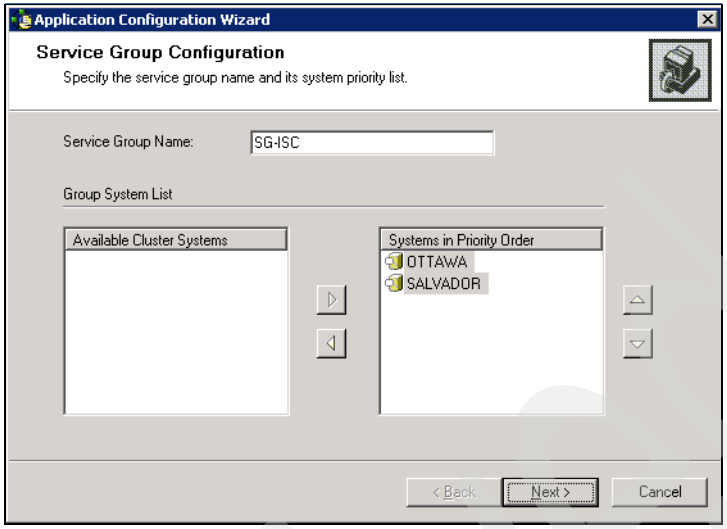

Figure 22-46 Service group configuration

5. The wizard will start a process of discovering all necessary objects to create the service group, as shown in Figure 22-47. We wait until this process ends.

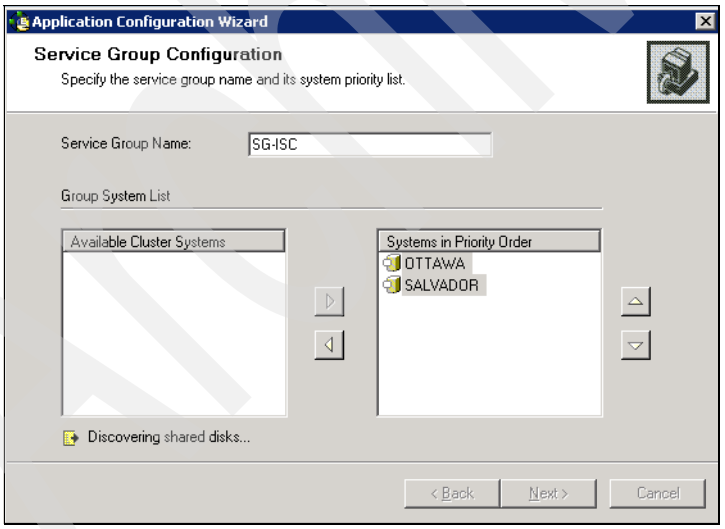

Figure 22-47 Discovering process

6. We then define what kind of application group this is. In our case there are two services: *ISC Help Service* and *IBM WebSphere Application Server V5 - ISC Runtime Service*. We choose **Generic Service** from the drop-down list in Figure 22-48 and click **Next**.

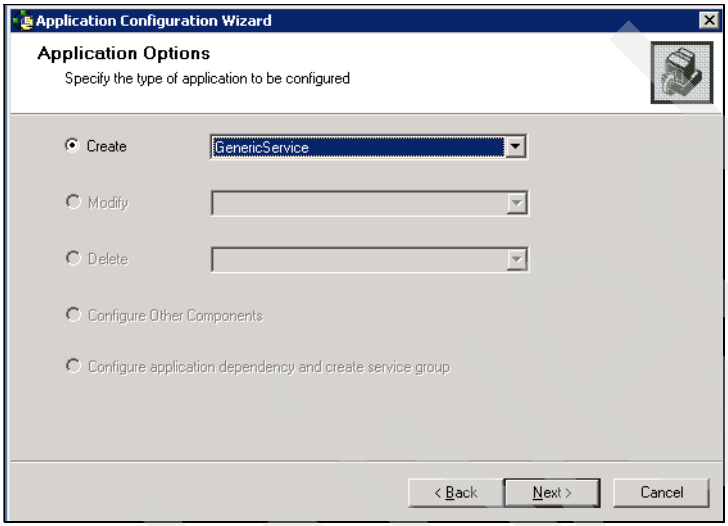

Figure 22-48 Choosing the kind of application

7. We click the button next to the **Service Name** line and choose the service *ISC Help Service* from the drop-down list as shown in Figure 22-49.

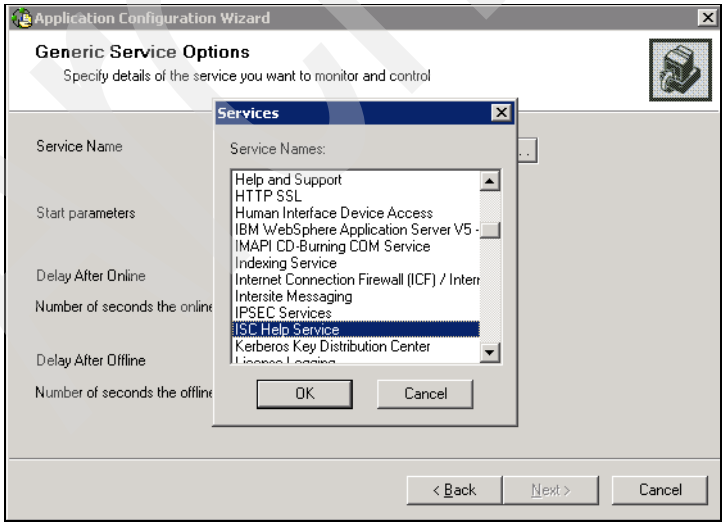

Figure 22-49 Choosing TSM Server1 service

8. We confirm the name of the service chosen and click **Next** in Figure 22-50.

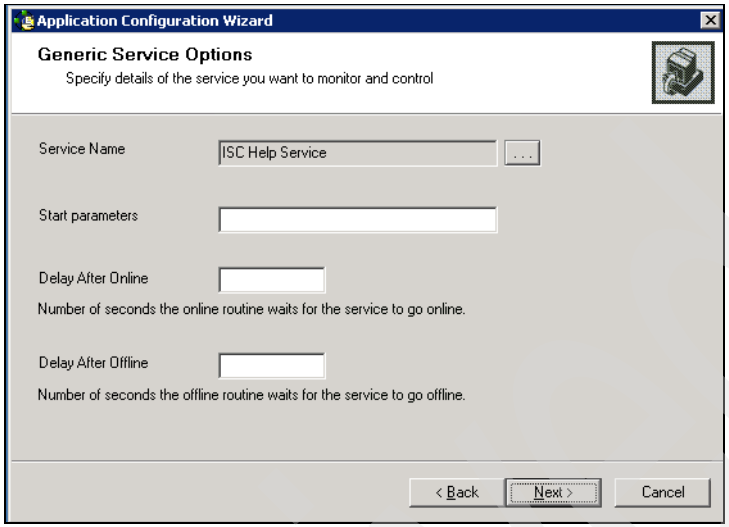

Figure 22-50 Confirming the service

9. In Figure 22-51 we choose to start the service with the LocalSystem account.

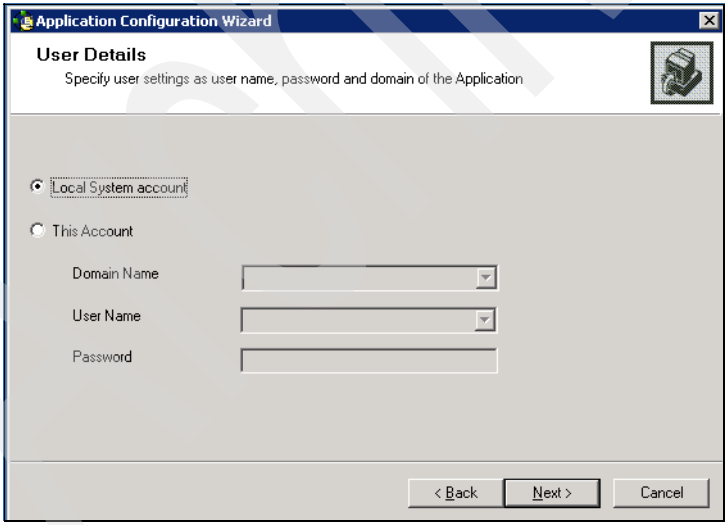

Figure 22-51 Choosing the service account

10. We select the drives that will be used by the Administration Center. We refer to Table 22-1 on page 906 to confirm the drive letters. We select the letters as in Figure 22-52 and click **Next**.

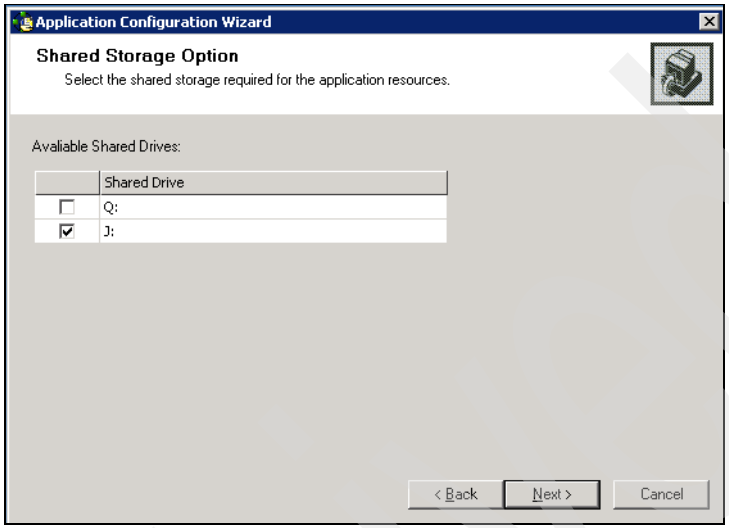

Figure 22-52 Selecting the drives to be used

11. We receive a summary of the application resource with the name and user account as in Figure 22-53. We confirm and click **Next**.

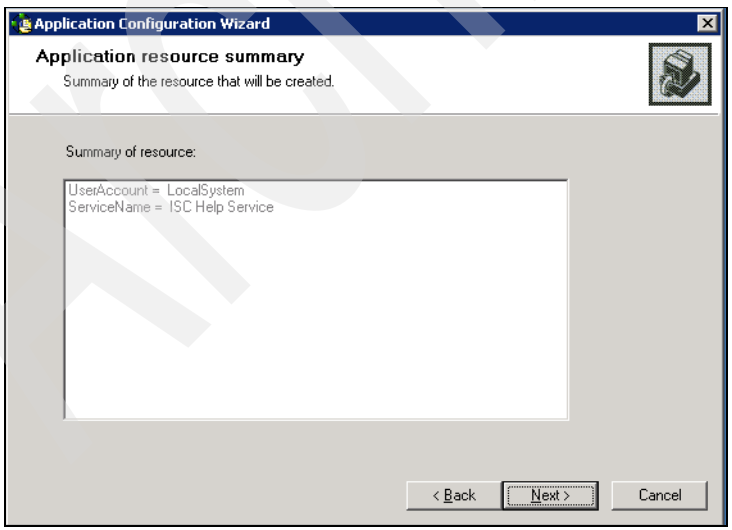

Figure 22-53 Summary with name and account for the service

12. We need to include one more service, that is *IBM WebSphere Application Server V5 - ISC Runtime Service*. We repeat steps 6 to 11 changing the service name.

13. We need two more resources for this group: IP and a Name. So in Figure 22-54 we choose **Configure Other Components** and then click **Next**.

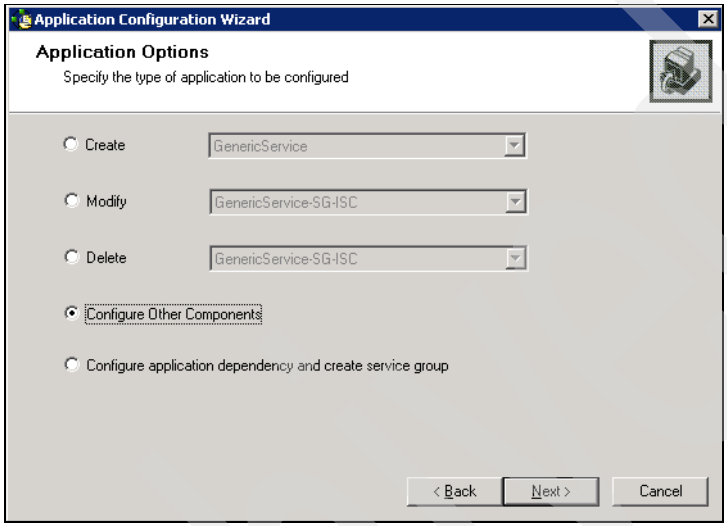

Figure 22-54 Choosing additional components

14. In Figure 22-55 we choose to create **Network Component** (IP address) and **Lanman Component** (Name) and click **Next**.

Archived

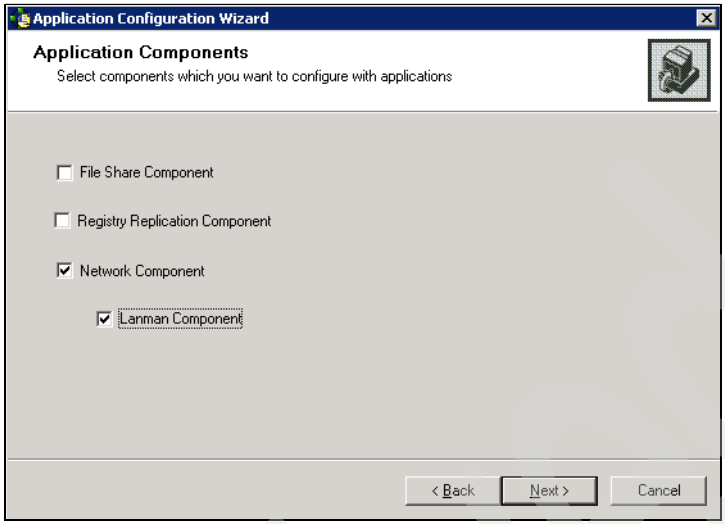

Figure 22-55 Choosing other components for IP address and Name

15. In Figure 22-56 we specify the name of the Tivoli Storage Manager server and the IP address we will use to connect our clients and click **Next**. We refer to Table 22-1 for the necessary information.

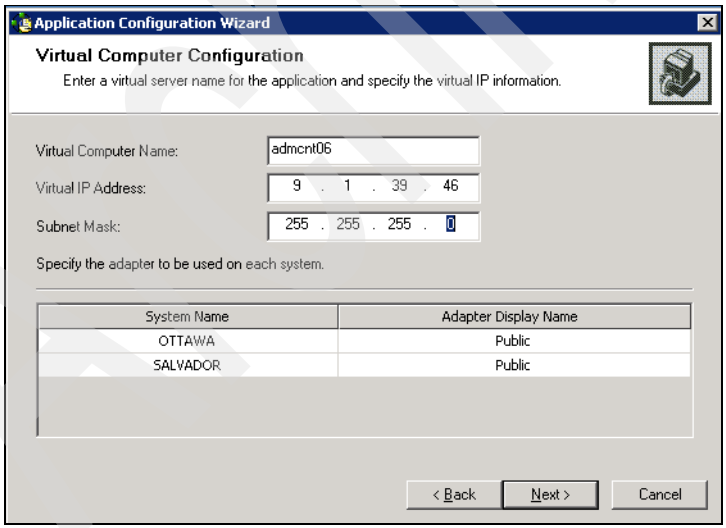

Figure 22-56 Informing name and ip address

16. We do not need any other resources to be configured. We choose **Configure application dependency and create service group** in Figure 22-57 and click **Next**.

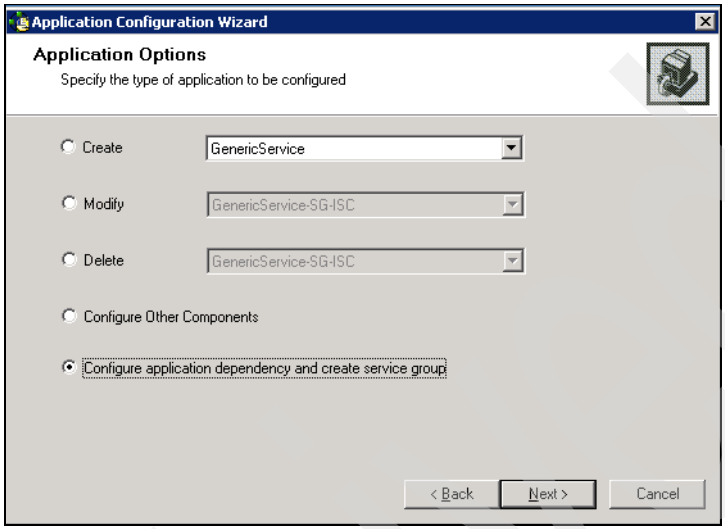

Figure 22-57 Completing the application options

17. We review the information presented in the summary in Figure 22-58.

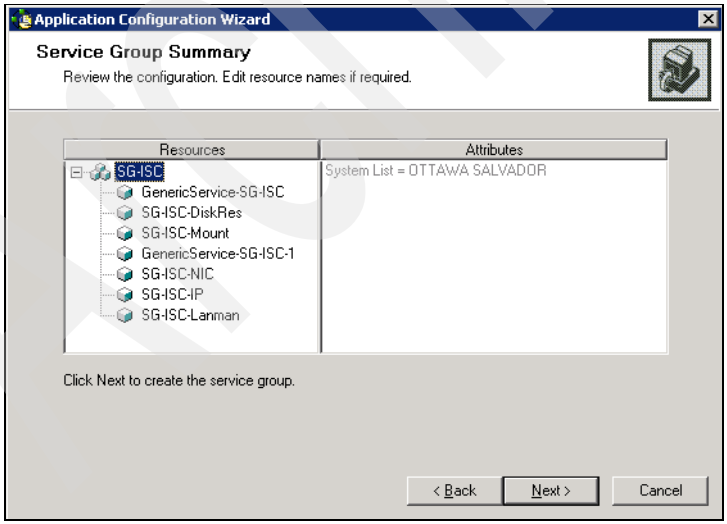

Figure 22-58 Service Group Summary

18. For a clearer information of the resources, we use the F2 key and change the name of the services, disk and mount resources so that they reflect their actual names, as shown in Figure 22-59.

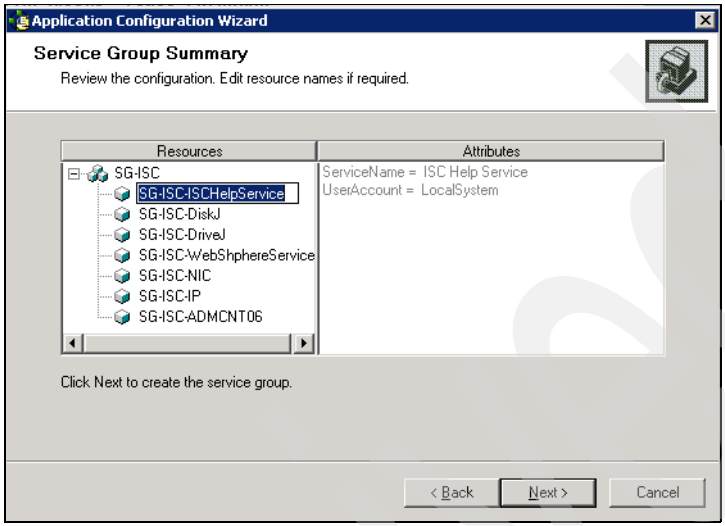

Figure 22-59 Changing the names of the resources

19. We confirm we want to create the service group clicking **Yes** in Figure 22-60.

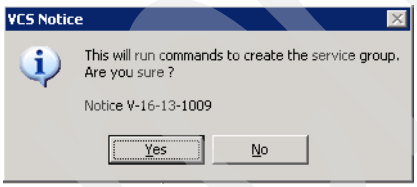

Figure 22-60 Confirming the creation of the service group

20. The process begins now as shown in Figure 22-61.

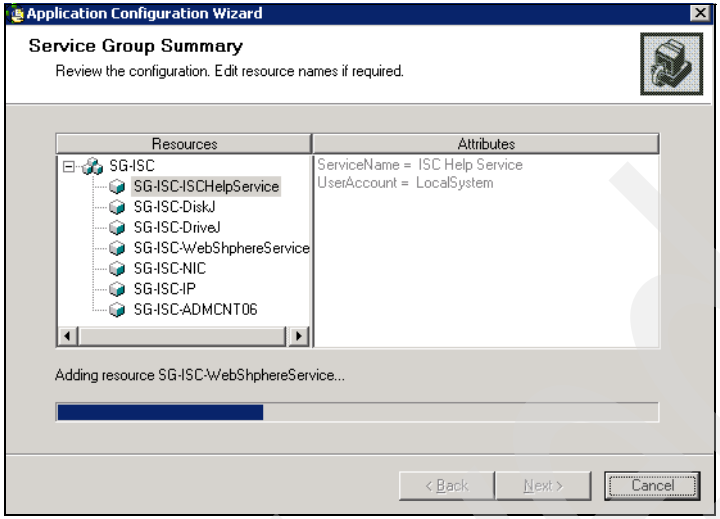

Figure 22-61 Creating the service group

21. When the process completes, uncheck the **Bring the service group online** option as shown in Figure 22-62. Because of the two services, we need to confirm the dependencies first

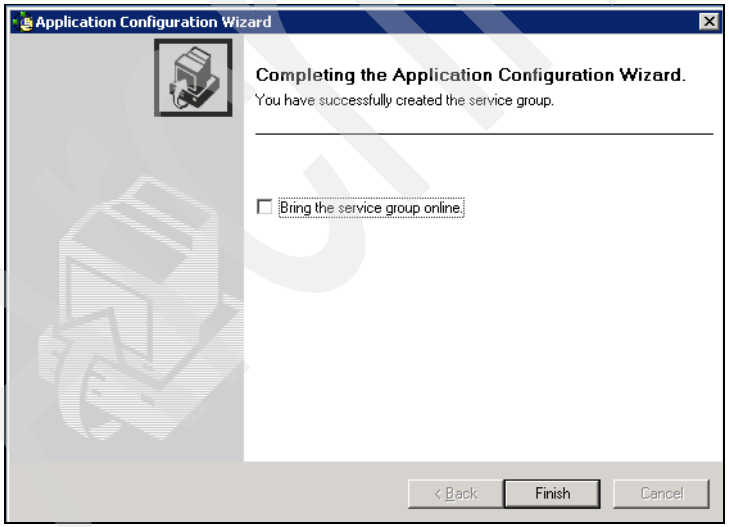

Figure 22-62 Completing the wizard

22. We now open the Java Console to administer the cluster and check configurations. We need to change the links, so we open the Resource tag in the right panel. *IBM WebSphere Application Server V5 - ISC Runtime Service*

needs to be started prior to the *ISC Help Service*. The link should be changed to match Figure 22-63. After changing, we bring the group online.

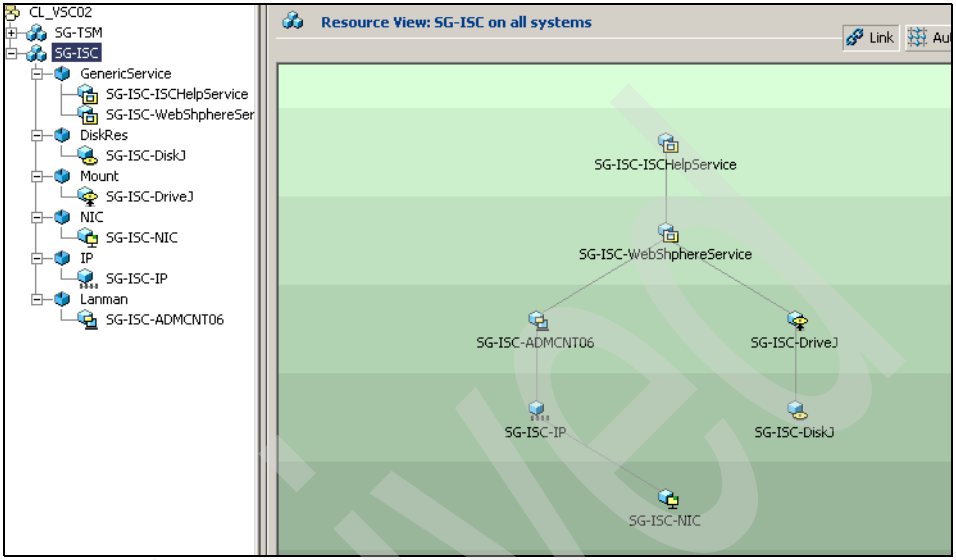

Figure 22-63 Correct link for the ISC Service Group

23. To validate the group, we switch it to the other node and access the ISC using a browser and pointing to either the name: admcnt06 or the ip 9.1.39.46 as shown in Figure 22-64. We can also include the name and IP in the DNS server.

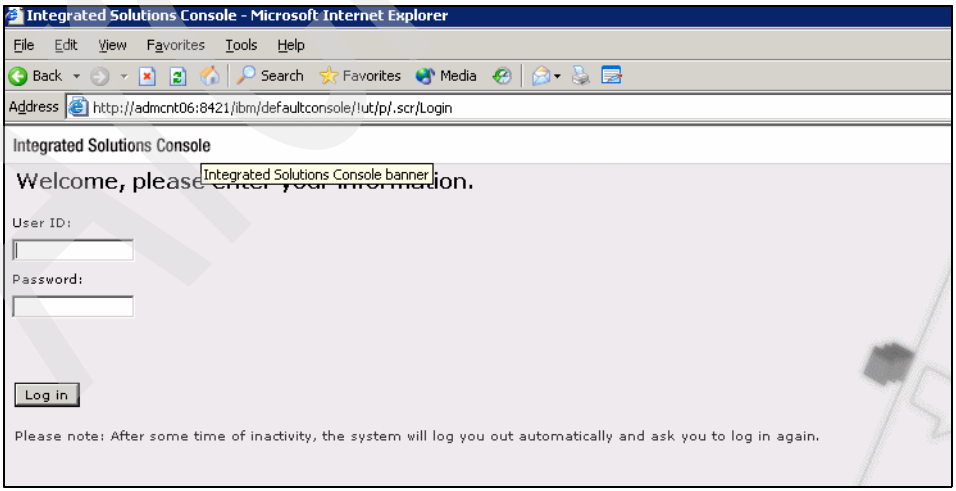

Figure 22-64 Accessing the administration center

22.9 Configuring Tivoli Storage Manager devices

Before starting the tests of the Tivoli Storage Manager Server, we create the necessary storage devices such as library, drives, and storage groups, using the Administration Center. We created the devices based on Table 22-3 on page 907.

22.10 Testing the Tivoli Storage Manager on VCS

In order to check the high availability of Tivoli Storage Manager server on our lab VCS environment, we must do some testing.

Our objective with these tests is to show how Tivoli Storage Manager, on a Veritas clustered environment, manages its own resources to achieve high availability and how it responds after certain kinds of failures that affect the shared resources.

22.10.1 Testing incremental backup using the GUI client

Our first test uses the Tivoli Storage Manager GUI to start an incremental backup.

Objective

The objective of this test is to show what happens when a client incremental backup is started from the Tivoli Storage Manager GUI and suddenly the node which hosts the Tivoli Storage Manager server in the VCS fails.

Activities

To do this test, we perform these tasks:

1. We open the Veritas Cluster Manager console to check which node hosts the Tivoli Storage Manager Service Group as shown in Figure 22-65.

Archived

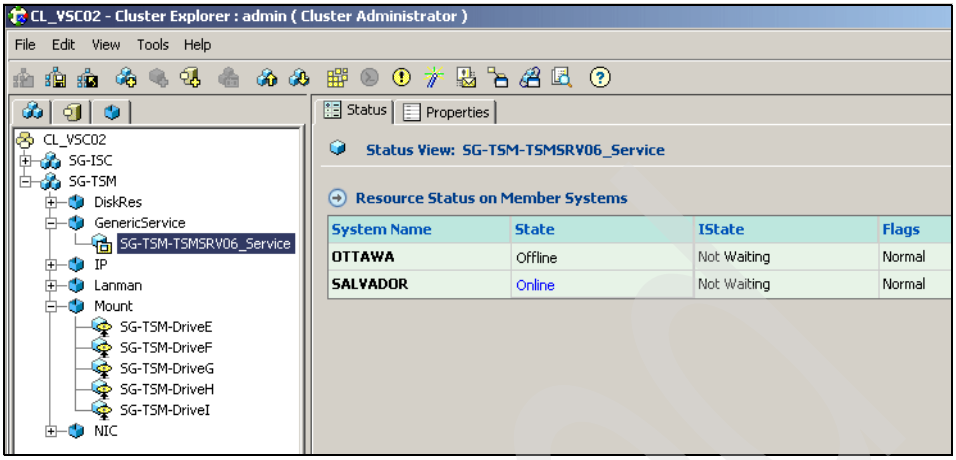

Figure 22-65 Veritas Cluster Manager console shows TSM resource in SALVADOR

2. We start an incremental backup from RADON (one of the two nodes of the Windows 2000 MSCS), using the Tivoli Storage Manager backup/archive GUI client. We select the local drives, the System State, and the System Services as shown in Figure 22-66.

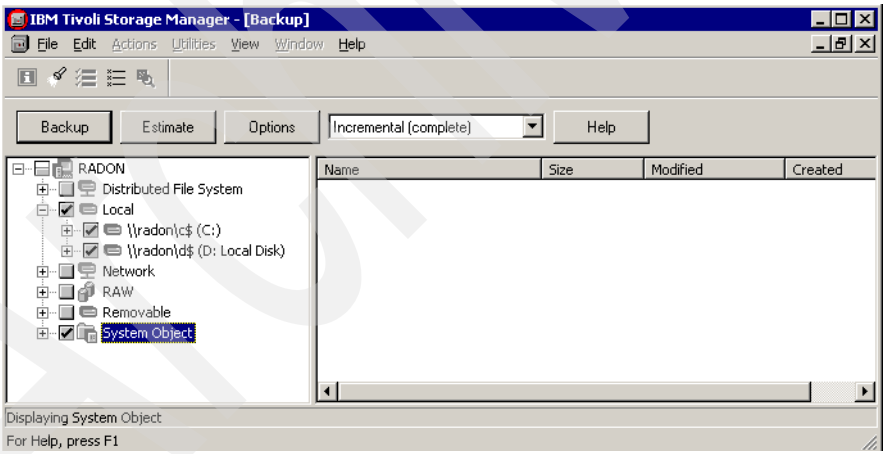

Figure 22-66 Starting a manual backup using the GUI from RADON

3. The transfer of files starts as we can see in Figure 22-67.

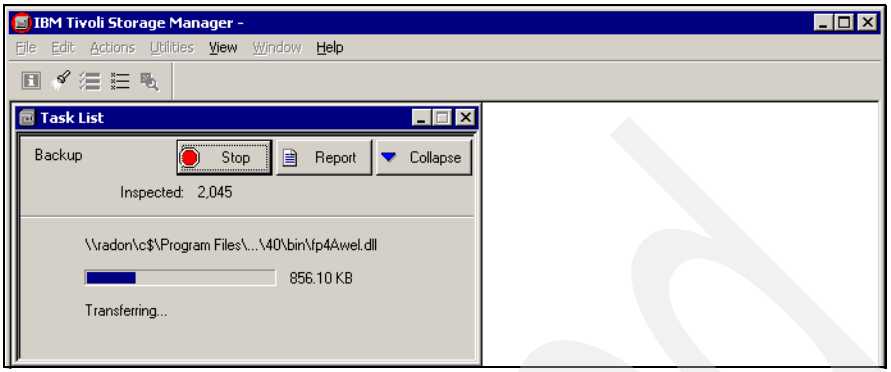

Figure 22-67 RADON starts transferring files to the TSMSRV06 server

4. While the client is transferring files to the server we force a failure on SALVADOR, the node that hosts the Tivoli Storage Manager server. When Tivoli Storage Manager restarts on the second node, we can see in the GUI client that backup is held and a reopening session message is received, as shown in Figure 22-68.

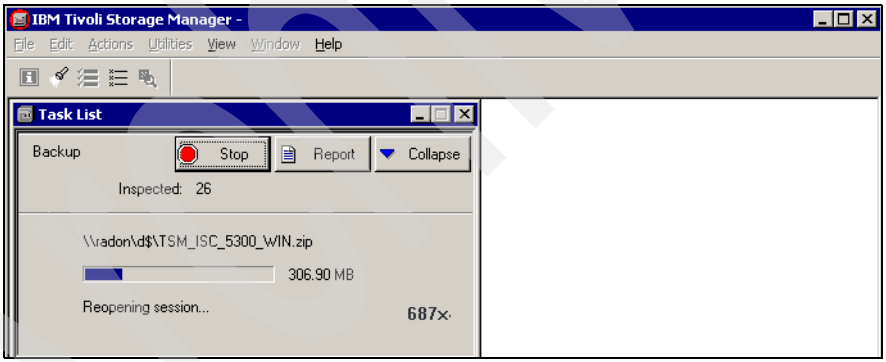

Figure 22-68 RADON loses its session, tries to reopen new connection to server

5. When the connection is re-established, the client continues sending files to the server, as shown in Figure 22-69.

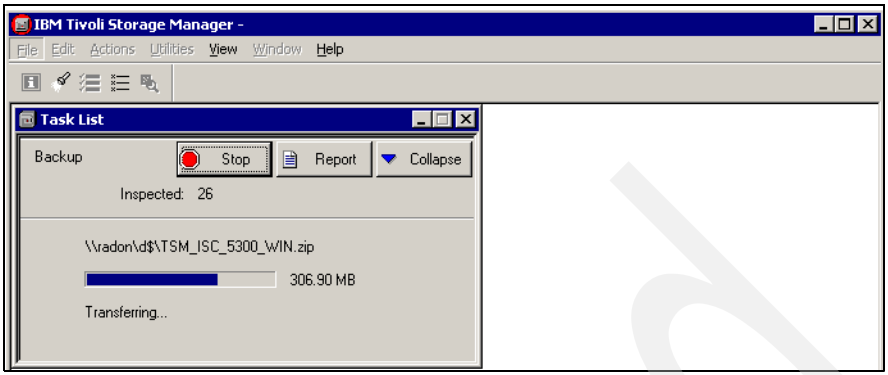

Figure 22-69 RADON continues transferring the files again to the server

6. RADON ends its backup successfully.

Results summary

The result of the test shows that when you start a backup from a client and there is a failure that forces Tivoli Storage Manager server to fail in a VCS, backup is held, and when the server is up again, the client reopens a session with the server and continues transferring data.

> **Note:** In the test we have just described, we used disk storage pool as the destination storage pool. We also tested using a tape storage pool as destination and we got the same results. The only difference is that when the Tivoli Storage Manager server is again up, the tape volume it was using on the first node is unloaded from the drive and loaded again into the second drive, and the client receives a “media wait” message while this process takes place. After the tape volume is mounted, the backup continues and ends successfully.

22.10.2 Testing a scheduled incremental backup

The second test consists of a scheduled backup.

Objective

The objective of this test is to show what happens when a scheduled client backup is running and suddenly the node which hosts the Tivoli Storage Manager server in the VCS fails.

Activities

We perform these tasks:

1. We open the Veritas Cluster Manager console to check which node hosts the Tivoli Storage Manager Service Group: SALVADOR.
2. We schedule a client incremental backup operation using the Tivoli Storage Manager server scheduler and we associate the schedule to the Tivoli Storage Manager client installed on RADON.
3. A client session starts from RADON as shown in Figure 22-70.

```
IBM Tivoli Storage Manager

Scheduled Start      Actual Start          Schedule Name Node Name     Status
-------------------- --------------------- ------------- ------------- ---------
02/18/2005 11:47:08                        INCR_BACKUP   RADON         Pending

tsm: TSMSRV06>q se

  Sess Comm.  Sess      Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ -------- -------- ------ -------- --------------------
     3 Tcp/Ip Run      0 S   67.5 K      519 Admin WinNT    ADMIN

tsm: TSMSRV06>q ev * *

Scheduled Start      Actual Start          Schedule Name Node Name     Status
-------------------- --------------------- ------------- ------------- ---------
02/18/2005 11:47:08                        INCR_BACKUP   RADON         Pending

tsm: TSMSRV06>q ev * *

Scheduled Start      Actual Start          Schedule Name Node Name     Status
-------------------- --------------------- ------------- ------------- ---------
02/18/2005 11:47:08  02/18/2005 11:48:24   INCR_BACKUP   RADON         Started

tsm: TSMSRV06>_
```

Figure 22-70 Scheduled backup started for RADON in the TSMSRV06 server

4. The client starts sending files to the server as shown in Figure 22-71.

Archived

```
dsmsched - Notepad
File  Edit  Format  Help
Executing scheduled command now.
02/18/2005 11:57:30 --- SCHEDULEREC OBJECT BEGIN INCR_BACKUP 02/18/2005 11:47:
02/18/2005 11:57:30 Incremental backup of volume '\\radon\c$'
02/18/2005 11:57:30 Incremental backup of volume '\\radon\d$'
02/18/2005 11:57:30 Incremental backup of volume 'SYSTEMOBJECT'
02/18/2005 11:57:30 Backup System Object: 'COM+ Database'.

02/18/2005 11:57:30 Backup System Object: 'Event Log'.

02/18/2005 11:57:30 Backup System Object: 'MSCS Database'.

02/18/2005 11:57:30 Backup System Object: 'System Volume'.

02/18/2005 11:57:30 Backup System Object: 'Windows NT Directory Services'.

02/18/2005 11:57:30 Backup System Object: 'Registry'.

02/18/2005 11:57:30 Backup System Object: 'WMI Database'.

02/18/2005 11:57:30 Backup System Object: 'System and Boot Files'.

02/18/2005 11:57:54 Directory-->                    0 \\radon\c$\ [Sent]
02/18/2005 11:57:54 Directory-->                    0 \\radon\c$\CmdFileCache [
02/18/2005 11:57:54 Directory-->                    0 \\radon\c$\Documents and
02/18/2005 11:57:54 Directory-->                    0 \\radon\c$\drivers_lto [s
02/18/2005 11:57:54 Directory-->                    0 \\radon\c$\fast [Sent]
02/18/2005 11:57:54 Directory-->                    0 \\radon\c$\Inetpub [Sent]
02/18/2005 11:57:54 Directory-->                    0 \\radon\c$\Program Files
02/18/2005 11:57:54 Directory-->                    0 \\radon\c$\RECYCLER [Sent
02/18/2005 11:57:54 Directory-->                    0 \\radon\c$\System Volume
02/18/2005 11:57:54 Directory-->                    0 \\radon\c$\TSM_Images [Se
02/18/2005 11:57:54 Directory-->                    0 \\radon\c$\WINNT [Sent]
```

Figure 22-71 Schedule log file in RADON shows the start of the scheduled backup

5. While the client continues sending files to the server, we force SALVADOR to fail. The following sequence occurs:

   a. In the client, the connection is lost, just as we can see in Figure 22-72.

```
dsmsched - Notepad
File  Edit  Format  Help
02/18/2005 11:59:14 Normal File-->            123 \\radon\c$\Documents and Settings\clustera
02/18/2005 11:59:14 Normal File-->            127 \\radon\c$\Documents and Settings\clustera
02/18/2005 11:59:14 Normal File-->            136 \\radon\c$\Documents and Settings\clustera
02/18/2005 11:59:14 Normal File-->            130 \\radon\c$\Documents and Settings\clustera
02/18/2005 11:59:14 Normal File-->            122 \\radon\c$\Documents and Settings\clustera
02/18/2005 11:59:14 Normal File-->            122 \\radon\c$\Documents and Settings\clustera
02/18/2005 11:59:14 Directory-->                0 \\radon\c$\Documents and Settings\clustera
02/18/2005 11:59:14 Directory-->                0 \\radon\c$\Documents and Settings\clustera
02/18/2005 11:59:14 Directory-->                0 \\radon\c$\Documents and Settings\clustera
02/18/2005 11:59:21 Retry # 1  Normal File-->       23,158,315 \\radon\c$\Documents and Settin
02/18/2005 11:59:21 ANS1809W Session is lost; initializing session reopen procedure.
```

Figure 22-72 RADON loses its connection with the TSMSRV06 server

   b. In the Veritas Cluster Manager console, SALVADOR goes down and OTTAWA receives the resources.

   c. When the Tivoli Storage Manager server instance resource is online (now hosted by OTTAWA), the schedule restarts as shown on the activity log in Figure 22-73.

```
IBM Tivoli Storage Manager
                          port 1580.
02/18/2005 11:52:04  ANR8200I TCP/IP driver ready for connection with clients
                          on port 1500.
02/18/2005 11:52:04  ANR4747W The web administrative interface is no longer
                          supported. Begin using the Integrated Solutions Console
                          instead.
02/18/2005 11:52:09  ANR0406I Session 1 started for node RADON (WinNT) (Tcp/Ip
                          RADON(2461)). (SESSION: 1)
02/18/2005 11:52:16  ANR8439I SCSI library LIBLTO is ready for operations.
02/18/2005 11:53:23  ANR0407I Session 2 started for administrator ADMIN (WinNT)
                          (Tcp/Ip RADON(2462)). (SESSION: 2)
02/18/2005 11:53:27  ANR2017I Administrator ADMIN issued command: QUERY SESSION
more...   (<ENTER> to continue, 'C' to cancel)

                          (SESSION: 2)
02/18/2005 11:53:29  ANR2017I Administrator ADMIN issued command: QUERY ACTLOG
                          (SESSION: 2)

tsm: TSMSRV06>q se

  Sess Comm.  Sess      Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State     Time    Sent   Recvd Type
------ ------ ------ ------- ------- ------- ----- -------- --------------------
     1 Tcp/Ip RecvW    1 S      426 184.0 M Node  WinNT    RADON
     2 Tcp/Ip Run      0 S   21.7 K     161 Admin WinNT    ADMIN

tsm: TSMSRV06>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name     Status
-------------------- -------------------- ------------- ------------- ---------
02/18/2005 11:47:08  02/18/2005 11:48:24  INCR_BACKUP   RADON         Restarted

tsm: TSMSRV06>_
```

Figure 22-73 In the event log the scheduled backup is restarted

6. The backup ends, just as we can see in the schedule log file of RADON in Figure 22-74.

```
dsmsched - Notepad
File  Edit  Format  Help
02/18/2005 13:08:07 Normal File-->          270,608 \\radon\c$\WINNT\winhlp32.exe [Sent]
02/18/2005 13:08:07 Normal File-->          193,296 \\radon\c$\WINNT\winrep.exe [Sent]
02/18/2005 13:08:07 Successful incremental backup of 'System and Boot Files'

02/18/2005 13:08:07 --- SCHEDULEREC STATUS BEGIN
02/18/2005 13:08:07 Total number of objects inspected:|  20,514
02/18/2005 13:08:07 Total number of objects backed up:   20,166
02/18/2005 13:08:07 Total number of objects updated:          0
02/18/2005 13:08:07 Total number of objects rebound:          0
02/18/2005 13:08:07 Total number of objects deleted:          0
02/18/2005 13:08:07 Total number of objects expired:          0
02/18/2005 13:08:07 Total number of objects failed:           2
02/18/2005 13:08:07 Total number of bytes transferred:     4.32 GB
02/18/2005 13:08:07 Data transfer time:                3,841.11 sec
02/18/2005 13:08:07 Network data transfer rate:        1,181.32 KB/sec
02/18/2005 13:08:07 Aggregate data transfer rate:      1,070.92 KB/sec
02/18/2005 13:08:07 Objects compressed by:                -33%
02/18/2005 13:08:07 Elapsed processing time:           01:10:37
02/18/2005 13:08:07 --- SCHEDULEREC STATUS END
02/18/2005 13:08:07 --- SCHEDULEREC OBJECT END INCR_BACKUP 02/18/2005 11:47:08
02/18/2005 13:08:07 ANS1512E Scheduled event 'INCR_BACKUP' failed.  Return code = 12.
```

Figure 22-74 Schedule log file in RADON shows the end of the scheduled backup

In Figure 22-74 the scheduled log file displays the event as failed with a “return code = 12”. However, if we look at this file in detail, each volume was backed up successfully, as we can see in Figure 22-75.

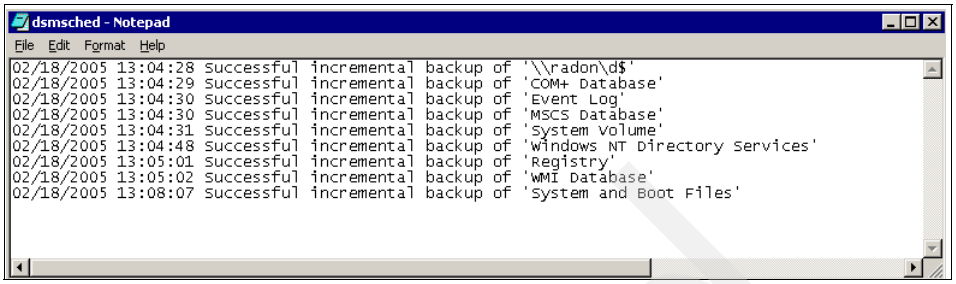

Figure 22-75 Every volume was successfully backed up by RADON

Attention: The scheduled event can end as *failed* with return code = 12 or as *completed* with return code = 8. It depends on the elapsed time until the second node of the cluster brings the resource online. In both cases, however, the backup completes successfully for each drive as we can see in Figure 22-75.

Results summary

The test results show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a scheduled backup started from one client is restarted after the failover on the other node of the VCS.

In the event log, the schedule can display *failed* instead of *completed*, with a return code = 12, if the elapsed time since the first node lost the connection, is too long. In any case, the incremental backup for each drive ends successfully.

Note: In the test we have just described, we used a disk storage pool as the destination storage pool. We also tested using a tape storage pool as destination and we got the same results. The only difference is that when the Tivoli Storage Manager server is again up, the tape volume it was using on the first node is unloaded from the drive and loaded again into the second drive, and the client receives a “media wait” message while this process takes place. After the tape volume is mounted the backup continues and ends successfully.

22.10.3 Testing migration from disk storage pool to tape storage pool

Our third test is a server process: migration from disk storage pool to tape storage pool.

Objective

The objective of this test is to show what happens when a disk storage pool migration process is started on the Tivoli Storage Manager server and the node that hosts the server instance fails.

Activities

For this test, we perform these tasks:

1. We open the Veritas Cluster Manager console to check which node hosts the Tivoli Storage Manager Service Group: OTTAWA.
2. We update the disk storage pool (SPD_BCK) high threshold migration to 0. This forces migration of backup versions to its next storage pool, a tape storage pool (SPT_BCK).
3. A process starts for the migration task and Tivoli Storage Manager prompts the tape library to mount a tape volume. After some seconds the volume is mounted as we show in Figure 22-76.

```
IBM Tivoli Storage Manager
02/18/2005 13:32:13   ANR0984I Process 2 for MIGRATION started in the BACKGROUND
                       at 13:32:13. (PROCESS: 2)
02/18/2005 13:32:13   ANR1000I Migration process 2 started for storage pool
                       SPD_BCK automatically, highMig=0, lowMig=0, duration=No.
                       (PROCESS: 2)
02/18/2005 13:32:15   ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                        (SESSION: 8)
02/18/2005 13:32:23   ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                        (SESSION: 8)
02/18/2005 13:32:25   ANR2017I Administrator ADMIN issued command: QUERY PROCESS
more...   (<ENTER> to continue, 'C' to cancel)

                        (SESSION: 8)
02/18/2005 13:32:39   ANR8337I LTO volume 022AKKL2 mounted in drive DRLTO_1
                       (mt0.0.0.2). (PROCESS: 2)
02/18/2005 13:32:39   ANR1340I Scratch volume 022AKKL2 is now defined in storage
                       pool SPT_BCK. (PROCESS: 2)
02/18/2005 13:32:39   ANR0513I Process 2 opened output volume 022AKKL2.
                       (PROCESS: 2)
```

Figure 22-76 Migration task started as process 2 in the TSMSRV06 server

4. While migration is running, we force a failure on OTTAWA. At this time the process has already migrated thousands of files, as we can see in Figure 22-77.

```
IBM Tivoli Storage Manager
tsm: TSMSRV06>q pr

 Process Process Description  Status
  Number
-------- -------------------- ---------------------------------------------
       2 Migration            Disk Storage Pool SPD_BCK, Moved Files: 4124,
                               Moved Bytes: 1,921,490,944, Unreadable Files: 0,
                               Unreadable Bytes: 0. Current Physical File
                               (bytes): 26,202,112 Current output volume:
                               022AKKL2.

tsm: TSMSRV06>_
```

Figure 22-77 Migration has already transferred 4124 files to the tape storage pool

The following sequence occurs:

a. In the Veritas Cluster Manager console, OTTAWA is out of the cluster and SALVADOR starts to bring the resources online.

b. After a short period of time the resources are online in SALVADOR.

c. When the Tivoli Storage Manager server instance resource is online (hosted by SALVADOR), the tape volume is unloaded from the drive. Since the high threshold is still 0, a new migration process is started and the server prompts to mount the same tape volume as shown in Figure 22-78.

```
IBM Tivoli Storage Manager
02/18/2005 13:36:10  ANR2560I Schedule manager started.
02/18/2005 13:36:10  ANR8200I TCP/IP driver ready for connection with clients
                      on port 1500.
02/18/2005 13:36:10  ANR8280I HTTP driver ready for connection with clients on
                      port 1580.
02/18/2005 13:36:10  ANR4747W The web administrative interface is no longer
                      supported. Begin using the Integrated Solutions Console
                      instead.
02/18/2005 13:36:11  ANR0407I Session 1 started for administrator ADMIN (WinNT)
                      (Tcp/Ip RADON(2808)). (SESSION: 1)
02/18/2005 13:36:11  ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                       (SESSION: 1)
02/18/2005 13:36:11  ANR0944E QUERY PROCESS: No active processes found.
                      (SESSION: 1)
02/18/2005 13:36:11  ANR2017I Administrator ADMIN issued command: ROLLBACK
more...   (<ENTER> to continue, 'C' to cancel)

                      (SESSION: 1)
02/18/2005 13:36:12  ANR0984I Process 2 for MIGRATION started in the BACKGROUND
                      at 13:36:12. (PROCESS: 2)
02/18/2005 13:36:12  ANR1000I Migration process 2 started for storage pool
                      SPD_BCK automatically, highMig=0, lowMig=0, duration=No.
                      (PROCESS: 2)
02/18/2005 13:36:38  ANR2017I Administrator ADMIN issued command: QUERY ACTLOG
                      (SESSION: 1)
02/18/2005 13:36:44  ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                       (SESSION: 1)
02/18/2005 13:37:55  ANR8439I SCSI library LIBLTO is ready for operations.
02/18/2005 13:38:22  ANR8337I LTO volume 022AKKL2 mounted in drive DRLTO_1
                      (mt0.0.0.2). (PROCESS: 2)
02/18/2005 13:38:22  ANR0513I Process 2 opened output volume 022AKKL2.
                      (PROCESS: 2)
```

Figure 22-78 Migration starts again in OTTAWA

5. The migration task ends successfully as we can see on the activity log in Figure 22-79.

```
Select IBM Tivoli Storage Manager
                      (SESSION: 1)
02/18/2005 13:40:18  ANR1001I Migration process 2 ended for storage pool
                      SPD_BCK. (PROCESS: 2)
02/18/2005 13:40:18  ANR0986I Process 2 for MIGRATION running in the BACKGROUND
more...   (<ENTER> to continue, 'C' to cancel)

                      processed 15223 items for a total of 2,548,236,288 bytes
                      with a completion state of SUCCESS at 13:40:18. (PROCESS:
                      2)
02/18/2005 13:42:55  ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                       (SESSION: 1)
```

Figure 22-79 Migration process ends successfully

Results summary

The results of our test show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a migration process started on the server before the failure, starts again when the second node on the VCS brings the Tivoli Storage Manager server instance online. This is true if the high threshold is still set to the value that caused the migration process to start.

The migration process starts from the last transaction committed into the database before the failure. In our test, before the failure, 4124 files were migrated to the tape storage pool, SPT_BCK. Those files are not migrated again when the process starts in OTTAWA.

22.10.4 Testing backup from tape storage pool to copy storage pool

In this section we test another internal server process, backup from a tape storage pool to a copy storage pool.

Objective

The objective of this test is to show what happens when a backup storage pool process (from tape to tape) is started on the Tivoli Storage Manager server and the node that hosts the resource fails.

Activities

For this test, we perform these tasks:

1. We open the Veritas Cluster Manager console to check which node hosts the Tivoli Storage Manager Service Group: SALVADOR.
2. We run the following command to start an storage pool backup from our primary tape storage pool SPT_BCK to our copy storage pool SPCPT_BCK:

   ```
   ba stg spt_bck spcpt_bck
   ```

3. A process starts for the storage pool backup task and Tivoli Storage Manager prompts to mount two tape volumes, one of them from the scratch pool because it is the first time we back up the primary tape storage pool against the copy storage pool. We show these events in Figure 22-80.

Archived

```
IBM Tivoli Storage Manager
02/18/2005 13:49:23  ANR1210I Backup of primary storage pool SPT_BCK to copy
                      storage pool SPCPT_BCK started as process 1. (SESSION: 1,
                      PROCESS: 1)
02/18/2005 13:49:24  ANR1228I Removable volume 022AKKL2 is required for storage
                      pool backup. (SESSION: 1, PROCESS: 1)
more...   (<ENTER> to continue, 'C' to cancel)

02/18/2005 13:49:25  ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                       (SESSION: 1)
02/18/2005 13:49:50  ANR8337I LTO volume 022AKKL2 mounted in drive DRLTO_1
                      (mt0.0.0.2). (SESSION: 1, PROCESS: 1)
02/18/2005 13:49:50  ANR0512I Process 1 opened input volume 022AKKL2. (SESSION:
                      1, PROCESS: 1)
02/18/2005 13:50:10  ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                       (SESSION: 1)
02/18/2005 13:50:20  ANR8337I LTO volume 023AKKL2 mounted in drive DRLTO_2
                      (mt1.0.0.2). (SESSION: 1, PROCESS: 1)
02/18/2005 13:50:20  ANR1340I Scratch volume 023AKKL2 is now defined in storage
                      pool SPCPT_BCK. (SESSION: 1, PROCESS: 1)
02/18/2005 13:50:20  ANR0513I Process 1 opened output volume 023AKKL2.
                      (SESSION: 1, PROCESS: 1)
```

Figure 22-80 Process 1 is started for the backup storage pool task

4. When the process is started, the two tape volumes are mounted on both drives as we show in Figure 22-81. We force a failure on SALVADOR.

```
IBM Tivoli Storage Manager
tsm: TSMSRV06>q pr

 Process Process Description  Status
  Number
-------- -------------------- ---------------------------------------------
       1 Backup Storage Pool  Primary Pool SPT_BCK, Copy Pool SPCPT_BCK, Files
                               Backed Up: 6990, Bytes Backed Up: 2,390,924,996,
                               Unreadable Files: 0, Unreadable Bytes: 0.
                               Current Physical File (bytes): 24,302,259
                               Current input volume: 022AKKL2. Current output
                               volume: 023AKKL2.

tsm: TSMSRV06>q mo
ANR8330I LTO volume 022AKKL2 is mounted R/W in drive DRLTO_1 (mt0.0.0.2),
status: IN USE.
ANR8330I LTO volume 023AKKL2 is mounted R/W in drive DRLTO_2 (mt1.0.0.2),
status: IN USE.
ANR8334I           2 matches found.

tsm: TSMSRV06>_
```

Figure 22-81 Process 1 has copied 6990 files in copy storage pool tape volume

The following sequence takes place:

a. In the Veritas Cluster Manager console, OTTAWA starts to bring the resources online while SALVADOR fails.

b. After a short period of time, the resources are online on OTTAWA.

c. When the Tivoli Storage Manager server instance resource is online (hosted by OTTAWA), the tape library dismounts both tape volumes from the drives. However, in the activity log there is no process started and there is no track of the process that was started before the failure in the server, as we see in Figure 22-82.

```
IBM Tivoli Storage Manager
02/18/2005 13:51:55  ANR8330I LTO volume 022AKKL2 is mounted R/W in drive
                      DRLTO_1 (mt0.0.0.2), status: IN USE. (SESSION: 1)
02/18/2005 13:51:55  ANR8330I LTO volume 023AKKL2 is mounted R/W in drive
                      DRLTO_2 (mt1.0.0.2), status: IN USE. (SESSION: 1)
02/18/2005 13:51:55  ANR8334I          2 matches found. (SESSION: 1)
02/18/2005 13:54:13  ANR4726I The NAS-NDMP support module has been loaded.
02/18/2005 13:54:13  ANR4726I The Centera support module has been loaded.
02/18/2005 13:54:13  ANR4726I The ServerFree support module has been loaded.
02/18/2005 13:54:14  ANR1305I Disk volume G:\ISMDATA\SERVER1\DISK1.DSM varied
                      online.
02/18/2005 13:54:14  ANR2803I License manager started.
02/18/2005 13:54:14  ANR0993I Server initialization complete.
02/18/2005 13:54:14  ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli is
                      now ready for use.
more...   (<ENTER> to continue, 'C' to cancel)

02/18/2005 13:54:14  ANR8260I Named Pipes driver ready for connection with
                      clients.
02/18/2005 13:54:14  ANR2828I Server is licensed to support Tivoli Storage
                      Manager Basic Edition.
02/18/2005 13:54:14  ANR2560I Schedule manager started.
02/18/2005 13:54:14  ANR8280I HTTP driver ready for connection with clients on
                      port 1580.
02/18/2005 13:54:14  ANR4747W The web administrative interface is no longer
                      supported. Begin using the Integrated Solutions Console
                      instead.
02/18/2005 13:54:14  ANR8200I TCP/IP driver ready for connection with clients
                      on port 1500.
02/18/2005 13:54:21  ANR0407I Session 1 started for administrator ADMIN (WinNT)
                      (Tcp/Ip RADON(2844)). (SESSION: 1)
02/18/2005 13:54:21  ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                       (SESSION: 1)
02/18/2005 13:54:21  ANR0944E QUERY PROCESS: No active processes found.
                      (SESSION: 1)
02/18/2005 13:54:21  ANR2017I Administrator ADMIN issued command: ROLLBACK
```

Figure 22-82 Backup storage pool task is not restarted when TSMSRV06 is online

5. The backup storage pool process does not restart again unless we start it manually.

6. If the backup storage pool process sent enough data before the failure so that the server was able to commit the transaction in the database, when the Tivoli Storage Manager server starts again in the second node, those files already copied in the copy storage pool tape volume and committed in the server database, are valid copied versions.

   However, there are still files not copied from the primary tape storage pool. If we want to be sure that the server copies all the files from this primary storage pool, we need to repeat the command. Those files committed as copied in the database will not be copied again.

   This happens both using roll-forward recovery log mode as well as normal recovery log mode.

   In our particular test, there was no tape volume in the copy storage pool before starting the backup storage pool process in the first node, because it was the first time we used this command.

   If you look at Figure 22-80 on page 956, there is an informational message in the activity log telling us that the scratch volume 023AKKL2 is now defined in the copy storage pool.

- When the server is again online in OTTAWA, we run the command:

  ```
  q vol
  ```

This reports the volume 023AKKL2 as a valid tape volume for the copy storage pool SPCPT_BCK, as we show in Figure 22-83.

```
IBM Tivoli Storage Manager
tsm: TSMSRV06>q vol

Volume Name                  Storage      Device       Estimated    Pct    Volume
                             Pool Name    Class Name    Capacity    Util    Status
----------------------------- ----------- ----------- ---------- ------ --------
022AKKL2                     SPT_BCK      CLLTO1       381,468.0    1.2   Filling
023AKKL2                     SPCPT_BCK    CLLTO1       381,468.0    0.5   Filling
G:\TSMDATA\SERVER1\DISK-     SPD_BCK      DISK           5,000.0    0.0   On-Line
 1.DSM

tsm: TSMSRV06>_
```

Figure 22-83 Volume 023AKKL2 defined as valid volume in the copy storage pool

- We run the command q occupancy against the copy storage pool and the Tivoli Storage Manager server reports the information in Figure 22-84.

```
IBM Tivoli Storage Manager
tsm: TSMSRV06>q occu stg=spcpt_bck

Node Name   Type Filespace    FSID Storage     Number of  Physical   Logical
                 Name              Pool Name      Files     Space     Space
                                                        Occupied  Occupied
                                                            (MB)      (MB)
---------- ---- ---------- ----- ---------- --------- --------- ---------
RADON       Bkup \\radon\d$      3 SPCPT_BCK      4,951  2,020.41  2,020.41

tsm: TSMSRV06>_
```

Figure 22-84 Occupancy for the copy storage pool after the failover

This means that the transaction was committed to the database before the failure in SALVADOR. Those files are valid copies.

- To be sure that the server copies the rest of the files, we start a new backup from the same primary storage pool, SPT_BCK to the copy storage pool, SPCPT_BCK.
- When the backup ends successfully, we use the following commands:

```
q occu stg=spt_bck
q occu stg=spcpt_bck
```

This reports the information in Figure 22-85.

Archived

```
IBM Tivoli Storage Manager

tsm: TSMSRV06>q occu stg=spt_bck

Node Name   Type Filespace    FSID Storage     Number of  Physical   Logical
                 Name               Pool Name       Files     Space     Space
                                                           Occupied  Occupied
                                                               (MB)      (MB)
---------- ---- ---------- ----- ---------- --------- -------- --------
RADON       Bkup \\radon\d$     3 SPT_BCK        4,951  2,020.41  2,020.41
RADON       Bkup \\radon\c$     4 SPT_BCK       13,238  2,159.83  2,159.83
RADON       Bkup SYSTEM         5 SPT_BCK        1,985    269.67    269.67
                  OBJECT

tsm: TSMSRV06>q occu stg=spcpt_bck

Node Name   Type Filespace    FSID Storage     Number of  Physical   Logical
                 Name               Pool Name       Files     Space     Space
                                                           Occupied  Occupied
                                                               (MB)      (MB)
---------- ---- ---------- ----- ---------- --------- -------- --------
RADON       Bkup \\radon\d$     3 SPCPT_BCK      4,951  2,020.41  2,020.41
RADON       Bkup \\radon\c$     4 SPCPT_BCK     13,238  2,159.83  2,159.83
RADON       Bkup SYSTEM         5 SPCPT_BCK      1,985    269.67    269.67
                  OBJECT

tsm: TSMSRV06>_
```

Figure 22-85 Occupancy is the same for primary and copy storage pools

If we do not have more primary storage pools, as in our case, both commands report exactly the same information.

7. If the backup storage pool task does not process enough data to commit the transaction into the database, when the Tivoli Storage Manager server starts again in the second node, those files copied in the copy storage pool tape volume before the failure are not recorded in the Tivoli Storage Manager server database. So, if we start a new backup storage pool task, they will be copied again.

   If the tape volume used for the copy storage pool before the failure was taken from the scratch pool in the tape library, (as in our case), it is given back to scratch status in the tape library.

   If the tape volume used for the copy storage pool before the failure had already data belonging to back up storage pool tasks from other days, the tape volume is kept in the copy storage pool but the new information written it is not valid.

   If we want to be sure that the server copies all the files from this primary storage pool, we need to repeat the command.

   This happens both using roll-forward recovery log mode as well as normal recovery log mode.

   In a test we made with recovery log in normal mode, also with no tape volumes in the copy storage pool, the server also mounted a scratch volume that was defined in the copy storage pool. However, when the server started on the second node after the failure, the tape volume was deleted from the copy storage pool.

Archived

Results summary

The results of our test show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a backup storage pool process (from tape to tape) started on the server before the failure, does not restart when the second node on the VCS brings the Tivoli Storage Manager server instance online.

Both tapes are correctly unloaded from the tape drives when the Tivoli Storage Manager server is again online, but the process is not restarted unless you run the command again.

Depending on the amount of data already sent when the task failed (if it was committed to the database or not), the files copied before the failure in the copy storage pool tape volume will be reflected on the database or not.

If enough information was copied to the copy storage pool tape volume so that the transaction was committed before the failure, when the server restarts in the second node, the information is recorded in the database and the files copied are valid copies.

If the transaction was not committed to the database, there is no information in the database about the process, and the files copied into the copy storage pool before the failure will need to be copied again.

This situation happens either if the recovery log is set to roll-forward mode or it is set to normal mode.

In any of these cases, to be sure that all information is copied from the primary storage pool to the copy storage pool, you should repeat the command.

There is no difference between a scheduled backup storage pool process or a manual process using the administrative interface. In our lab we tested both methods and the results were the same.

22.10.5 Testing server database backup

The following test consists of backing up the server database.

Objective

The objective of this test is to show what happens when a Tivoli Storage Manager server database backup process starts on the Tivoli Storage Manager server and the node that hosts the resource fails.

Activities

For this test, we perform these tasks:

1. We open the Veritas Cluster Manager console to check which node hosts the Tivoli Storage Manager Service Group: OTTAWA.
2. We start a full database backup.
3. Process 1 starts for database backup and Tivoli Storage Manager prompts to mount a scratch tape volume as shown in Figure 22-86.

```
IBM Tivoli Storage Manager

02/18/2005 14:28:49  ANR0984I Process 2 for DATABASE BACKUP started in the
                      BACKGROUND at 14:28:49. (SESSION: 1, PROCESS: 2)
02/18/2005 14:28:49  ANR2280I Full database backup started as process 2.
                      (SESSION: 1, PROCESS: 2)
02/18/2005 14:28:50  ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                      (SESSION: 1)
02/18/2005 14:28:54  ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                      (SESSION: 1)
02/18/2005 14:28:55  ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                      (SESSION: 1)
02/18/2005 14:28:56  ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                      (SESSION: 1)
02/18/2005 14:28:57  ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                      (SESSION: 1)
02/18/2005 14:29:00  ANR2017I Administrator ADMIN issued command: QUERY PROCESS
                      (SESSION: 1)
02/18/2005 14:29:03  ANR2017I Administrator ADMIN issued command: QUERY MO
```

Figure 22-86 Process 1 started for a database backup task

4. While the backup is running and the tape volume is mounted we force a failure on OTTAWA, just as we show in Figure 22-87.

```
IBM Tivoli Storage Manager
                      (SESSION: 1)
02/18/2005 14:29:16  ANR8379I Mount point in device class CLLTO1 is waiting for
                      the volume mount to complete, status: WAITING FOR VOLUME.
                      (SESSION: 1)
02/18/2005 14:29:16  ANR8334I          1 matches found. (SESSION: 1)
02/18/2005 14:29:17  ANR2017I Administrator ADMIN issued command: QUERY MO
                      (SESSION: 1)
02/18/2005 14:29:17  ANR8379I Mount point in device class CLLTO1 is waiting for
                      the volume mount to complete, status: WAITING FOR VOLUME.
                      (SESSION: 1)
02/18/2005 14:29:17  ANR8334I          1 matches found. (SESSION: 1)
02/18/2005 14:29:18  ANR8337I LTO volume 027AKKL2 mounted in drive DRLTO_1
                      (mt0.0.0.2). (SESSION: 1, PROCESS: 2)
02/18/2005 14:29:18  ANR0513I Process 2 opened output volume 027AKKL2.
                      (SESSION: 1, PROCESS: 2)
02/18/2005 14:29:19  ANR2017I Administrator ADMIN issued command: QUERY MO
                      (SESSION: 1)
02/18/2005 14:29:19  ANR8330I LTO volume 027AKKL2 is mounted R/W in drive
                      DRLTO_1 (mt0.0.0.2), status: IN USE. (SESSION: 1)
02/18/2005 14:29:19  ANR8334I          1 matches found. (SESSION: 1)
02/18/2005 14:29:21  ANR1360I Output volume 027AKKL2 opened (sequence number
                      1). (SESSION: 1, PROCESS: 2)
more...   <<ENTER> to continue, 'C' to cancel>

02/18/2005 14:29:21  ANR4554I Backed up 1664 of 3363 database pages. (SESSION:
                      1, PROCESS: 2)
02/18/2005 14:29:21  ANR4554I Backed up 3328 of 3363 database pages. (SESSION:
                      1, PROCESS: 2)
02/18/2005 14:29:59  ANR4726I The NAS-NDMP support module has been loaded.
02/18/2005 14:30:00  ANR4726I The Centera support module has been loaded.
02/18/2005 14:30:00  ANR4726I The ServerFree support module has been loaded.
02/18/2005 14:30:00  ANR1305I Disk volume G:\TSMDATA\SERVER1\DISK1.DSM varied
```

Figure 22-87 While the database backup process is started OTTAWA fails

The following sequence occurs:

a. In the Veritas Cluster Manager console, SALVADOR tries to bring the resources online while OTTAWA fails.

b. After a few minutes the resources are online on SALVADOR.

c. When the Tivoli Storage Manager server instance resource is online (hosted by SALVADOR), the tape volume is unloaded from the drive by the tape library automatic system. There is no process started on the server for any database backup and there is no track in the server database for that backup.

5. We query the volume history and there is no record for the tape volume 027AKKL2, which is the tape volume that was mounted by the server before the failure in OTTAWA. We can see this in Figure 22-88.

```
IBM Tivoli Storage Manager
tsm: TSMSRV06>q volh t=dbb

       Date/Time: 02/18/2005 14:27:49
     Volume Type: BACKUPFULL
   Backup Series: 2
Backup Operation: 0
      Volume Seq: 1
    Device Class: FILE
     Volume Name: G:\TSMDATA\08765669.DBB
 Volume Location:
         Command:

tsm: TSMSRV06>_
```

Figure 22-88 Volume history does not report any information about 027AKKL2

6. We query the library inventory. The tape volume status displays as private and its last use reports as dbbackup. We see this in Figure 22-89.

```
Select IBM Tivoli Storage Manager
tsm: TSMSRV06>q libvol

Library Name  Volume Name  Status     Owner      Last Use   Home     Device
                                                            Element  Type
------------  -----------  ---------  ---------  ---------  -------  ------
LIBLTO        020AKKL2     Private    TSMSRV06              4,096    LTO
LIBLTO        021AKKL2     Private    TSMSRV06              4,097    LTO
LIBLTO        022AKKL2     Private    TSMSRV06   Data       4,098    LTO
LIBLTO        023AKKL2     Scratch                           4,099    LTO
LIBLTO        026AKKL2     Scratch                           4,102    LTO
LIBLTO        027AKKL2     Private    TSMSRV06   DbBackup   4,116    LTO
LIBLTO        028AKKL2     Private    TSMSRV06              4,104    LTO
LIBLTO        029AKKL2     Private    TSMSRV06              4,105    LTO
LIBLTO        030AKKL2     Private    TSMSRV06              4,106    LTO
more...   <<ENTER> to continue, 'C' to cancel> _
```

Figure 22-89 The library volume inventory displays the tape volume as private

7. Since the database backup was not considered as valid, we must update the library inventory to change the status to scratch, using the following command:

```
upd libvol liblto 027akkl2 status=scratch
```

8. We repeat the database backup to have a valid and recent copy.

Results summary

The results of our test show that after a failure on the node that hosts the Tivoli Storage Manager server instance, a database backup process that started on the server before the failure, does not restart when the second node on the VCS brings the Tivoli Storage Manager server instance online.

The tape volume is correctly unloaded from the tape drive where it was mounted when the Tivoli Storage Manager server is again online, but the process does not end successfully. It is not restarted unless you run the command.

There is no difference between a scheduled process or a manual process using the administrative interface.

Important: The tape volume used for the database backup before the failure is not useful. It is reported as a private volume in the library inventory but it is not recorded as valid backup in the volume history file. It is necessary to update the tape volume in the library inventory to scratch and start again a new database backup process

Archived

Archived

23

VERITAS Cluster Server and the IBM Tivoli Storage Manager Client

This chapter describes the implementation of Tivoli Storage Manager backup/archive client on our Windows 2003 VCS clustered environment.

Archived

© Copyright IBM Corp. 2005. All rights reserved.

23.1 Overview

When servers are set up in a clustered environment, applications can be active on different nodes at different times.

Tivoli Storage Manager backup/archive client is designed to support its implementation on an VCS environment. However, it needs to be installed and configured following certain rules in order to run properly.

This chapter covers all the tasks we follow to achieve this goal.

23.2 Planning and design

You need to gather the following information to plan a backup strategy with Tivoli Storage Manager:

- Configuration of your cluster resource groups
- IP addresses and network names
- Shared disks that need to be backed up
- Tivoli Storage Manager nodenames used by each service group

Note: To back up the Windows 2003 system state or system services on local disks, Tivoli Storage Manager client must be connected to a Tivoli Storage Manager Version 5.2.0 or higher.

Plan the names of the various services and resources so that they reflect your environment and ease your work.

Archived

23.3 Lab setup

Our lab environment consists of a Windows 2003 Enterprise Server cluster with two node VCS, OTTAWA and SALVADOR.

The Tivoli Storage Manager backup/archive client configuration for this cluster is shown in Figure 23-1.

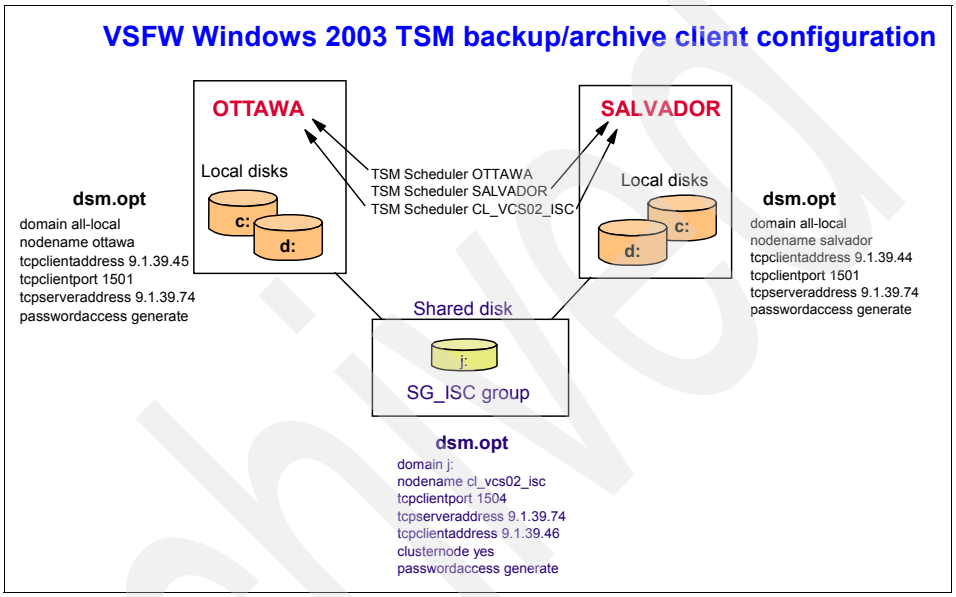

Figure 23-1 Tivoli Storage Manager backup/archive clustering client configuration

Refer to Table 21-1 on page 881, Table 21-2 on page 882, and Table 21-3 on page 882 for details of the VCS configuration used in our lab.

Archived

Table 23-1 and Table 23-2 show the specific Tivoli Storage Manager backup/archive client configuration we use for the purpose of this chapter.

Table 23-1 Tivoli Storage Manager backup/archive client for local nodes

Local node 1	
TSM nodename	OTTAWA
Backup domain	c: d: systemstate systemservices
Scheduler service name	TSM Scheduler OTTAWA
Client Acceptor service name	TSM Client Acceptor OTTAWA
Remote Client Agent service name	TSM Remote Client Agent OTTAWA
Local node 2	
TSM nodename	SALVADOR
Backup domain	c: d: systemstate systemservices
Scheduler service name	TSM Scheduler SALVADOR
Client Acceptor service name	TSM Client Acceptor SALVADOR
Remote Client Agent service name	TSM Remote Client Agent SALVADOR

Table 23-2 Tivoli Storage Manager backup/archive client for virtual node

Virtual node 1	
TSM nodename	CL_VCS02_ISC
Backup domain	j:
Scheduler service name	TSM Scheduler CL_VCS02_ISC
Client Acceptor service name	TSM Client Acceptor CL_VCS02_ISC
Remote Client Agent service name	TSM Remote Client Agent CL_VCS02_ISC
Service Group name	SG-ISC

23.4 Installation of the backup/archive client

The steps for installing the Tivoli Storage Manager backup/archive client in this environment were the same outlined in Chapter 6, “Microsoft Cluster Server and the IBM Tivoli Storage Manager Client” on page 241.

23.5 Configuration

In this section we describe how to configure the Tivoli Storage Manager backup/archive client in the cluster environment. This is a two-step procedure:

1. Configuring Tivoli Storage Manager client on local disks
2. Configuring Tivoli Storage Manager client on shared disks

23.5.1 Configuring Tivoli Storage Manager client on local disks

The configuration for the backup of the local disks is the same as for any standalone client:

1. We create a nodename for each server (OTTAWA and SALVADOR) on the Tivoli Storage Manager server.
2. We create the option file (dsm.opt) for each node on the local drive.

> **Important:** You should *only* use the `domain` option if not all local drives are going to be backed up. The default, if you do not specify anything, is backing up all local drives and system objects. Also, do not include any cluster drive in the domain parameter.

3. We generate the password locally by either opening the backup-archive GUI or issuing a query on the command prompt, such as **`dsmc q se`**.
4. We create the local Tivoli Storage Manager services as needed for each node:
   - Tivoli Storage Manager Scheduler
   - Tivoli Storage Manager Client Acceptor
   - Tivoli Storage Manager Remote Client Agent

23.5.2 Configuring Tivoli Storage Manager client on shared disks

The configuration of Tivoli Storage Manager client to back up shared disks is slightly different for virtual nodes on VCS.

For every resource group that has shared disks with backup requirements, we need to define a node name, an option file and an associated Tivoli Storage Manager scheduler service. If we want to use the Web client to access that virtual node from a browser, we also have to install the Web client services for that particular resource group.

For details of the nodenames, resources and services used for this part of the chapter, refer to Table 23-1 on page 968 and Table 23-2 on page 968.

Each resource group needs its own unique nodename. This ensures that Tivoli Storage Manager client correctly manages the disk resources in case of failure on any physical node, independently of the node who hosts the resources at that time.

As you can see in the tables mentioned above, we create one node in the Tivoli Storage Manager server database:

CL_VCS02_ISC: for the TSM_ISC Service Group

The configuration process consists, for each group, of the following tasks:

1. Creation of the option files
2. Password generation
3. Creation of the Tivoli Storage Manager Scheduler service
4. Creation of a resource for scheduler service in VCS

We describe each activity in the following sections.

Creation of the option files

For each group on the cluster we need to create an option file that will be used by the Tivoli Storage Manager nodename attached to that group.

The option file should be located on one of the shared disks hosted by this group. This ensures that both physical nodes have access to the file.

The dsm.opt file *must* contain at least the following options:

- nodename: Specifies the name that this group uses when it backs up data to the Tivoli Storage Manager server
- domain: Specifies the disk drive letters managed by this group
- clusternode yes: Specifies that it is a virtual node of a cluster. This is the main difference between the option file for a virtual node and the option file for a physical node.

If we plan to use the schedmode prompted option to schedule backups, and we plan to use the *Web client interface* for each virtual node, we also should specify the following options:

- tcpclientaddress: Specifies the unique IP address for this resource group
- tcpclientport: Specifies a different TCP port for each node
- httpport: Specifies a different http port to contact with.

There are other options we can specify, but the ones mentioned above are a requirement for a correct implementation of the client.

Archived

In our environment we create the dsm.opt file in the j:\tsm directory.

Option file for TSM_ISC Service Group

The dsm.opt file for this group contains the following options:

```
nodename cl_vcs02_isc
passwordaccess generate
tcpserveraddress 9.1.39.74
errorlogretention 7
errorlogname j:\tsm\dsmerror.log
schedlogretention 7
schedlogname j:\tsm\dsmsched.log
domain j:
clusternode yes
schedmode prompted
tcpclientaddress 9.1.39.46
tcpclientport 1504
httpport 1584
```

Password generation

Important: The steps below require that we run the following commands on both nodes while they own the resources. We recommend to move all resources to one of the nodes, complete the tasks for this node, and then move all resources to the other node and repeat the tasks.

The Windows registry of each server needs to be updated with the password that was used to create the nodename in the Tivoli Storage Manager server. Since the dsm.opt for the Service Group is in a different location as the default, we need to specify the path using the -optfile option:

1. We run the following commands from a MS-DOS prompt in the Tivoli Storage Manager client directory (c:\program files\tivoli\tsm\baclient):

   ```
   dsmc q se -optfile=j:\tsm\dsm.opt
   ```

2. Tivoli Storage Manager prompts the nodename for the client (the specified in dsm.opt). If it is correct, press Enter.
3. Tivoli Storage Manager next asks for a password. We type the password we used to register this node in the Tivoli Storage Manager server.
4. The result is shown in Example 23-1.

Example 23-1 Registering the node password

```
C:\Program Files\Tivoli\TSM\baclient>dsmc q se -optfile=j:\tsm\dsm.opt
IBM Tivoli Storage Manager
Command Line Backup/Archive Client Interface
  Client Version 5, Release 3, Level 0.0
```


Archived

```
  Client date/time: 02/21/2005 11:03:03
(c) Copyright by IBM Corporation and other(s) 1990, 2004. All Rights Reserved.

Node Name: CL_VCS02_ISC
Please enter your user id <CL_VCS02_ISC>:

Please enter password for user id "CL_VCS02_ISC": ******

Session established with server TSMSRV06: Windows
  Server Version 5, Release 3, Level 0.0
  Server date/time: 02/21/2005 11:03:03  Last access: 02/21/2005 11:03:03

TSM Server Connection Information

Server Name.............: TSMSRV06
Server Type.............: Windows
Server Version..........: Ver. 5, Rel. 3, Lev. 0.0
Last Access Date........: 02/21/2005 11:03:03
Delete Backup Files.....: "No"
Delete Archive Files....: "Yes"

Node Name...............: CL_VCS02_ISC
User Name...............:
```

5. We move the resources to the other node and repeat steps 1 to 3.

Creation of the Tivoli Storage Manager Scheduler service

For backup automation, using the Tivoli Storage Manager scheduler, we need to create and configure one scheduler service for each resource group.

> **Important:** We must create the scheduler service for each Service Group exactly with the same name, which is case sensitive, on each of the physical nodes and on the Veritas Cluster Explorer, otherwise failover will not work.

1. We need to be sure we run the commands on the node that hosts all resources.
2. We begin the installation of the scheduler service for each group in OTTAWA. This is the node that hosts the resources. We use the **dsmcutil** program. This utility is located on the Tivoli Storage Manager client installation path (c:\program files\tivoli\tsm\baclient).

   In our lab we installed one scheduler service, for our TSM_ISC Service Group.
3. We open an MS-DOS command line and, in the Tivoli Storage Manager client installation path we issue the following command:

```
dsmcutil inst sched /name:"TSM Scheduler CL_VCS02_ISC"
/clientdir:"c:\program files\tivoli\tsm\baclient" /optfile:j:\tsm\dsm.opt
/node:CL_VCS02_ISC /password:itsosj /clustername:CL_VCS02 /clusternode:yes
/autostart:no
```

4. The result is shown in Example 23-2.

Example 23-2 Creating the schedule on each node

```
C:\Program Files\Tivoli\TSM\baclient>dsmcutil inst sched /name:"TSM Scheduler
CL_VCS02_ISC" /clientdir:"c:\program files\tivoli\tsm\baclient" /optfile:j:
\tsm\dsm.opt /node:CL_VCS02_ISC /password:itsosj /clustername:CL_VCS02
/clusternode:yes /autostart:no

TSM Windows NT Client Service Configuration Utility
Command Line Interface - Version 5, Release 3, Level 0.0
(C) Copyright IBM Corporation, 1990, 2004, All Rights Reserved.
Last Updated Dec  8 2004
TSM Api Verison 5.3.0

Command: Install TSM Client Service
Machine: SALVADOR(Local Machine)


Locating the Cluster Services ...
Veritas cluster ...running


Installing TSM Client Service:

       Machine          : SALVADOR
       Service Name     : TSM Scheduler CL_VCS02_ISC
       Client Directory : c:\program files\tivoli\tsm\baclient
       Automatic Start  : no
       Logon Account    : LocalSystem

The service was successfully installed.


Creating Registry Keys ...

Updated registry value 'ImagePath' .
Updated registry value 'EventMessageFile' .
Updated registry value 'TypesSupported' .
Updated registry value 'TSM Scheduler CL_VCS02_ISC' .
Updated registry value 'ADSMClientKey' .
Updated registry value 'OptionsFile' .
Updated registry value 'EventLogging' .
Updated registry value 'ClientNodeName' .
Updated registry value 'ClusterNode' .
```

Archived

```
Updated registry value 'ClusterGroupName' .

Generating registry password ...
Authenticating TSM password for node CL_VCS02_ISC ...

Connecting to TSM Server via client options file 'j:\tsm\dsm.opt' ...

Password authentication successful.

The registry password for TSM node CL_VCS02_ISC has been updated.


Starting the 'TSM Scheduler CL_VCS02_ISC' service ...

The service was successfully started.
```

Tip: If there is an error message, "*An unexpected error (-1) occurred while the program was trying to obtain the cluster name from the system*", it is because there is a .stale file present in Veritas cluster directory. Check the Veritas support Web site for an explanation of this file. We can delete this file and run the command again.

5. We stop the service using the Windows service menu before going on.
6. We move the resources to the second node, and run exactly the same commands as before (steps 1 to 3).

Attention: The Tivoli Storage Manager scheduler service names used on both nodes *must match*. Also remember to use the same parameters for the **dsmcutil** tool. *Do not forget* the clusternode **yes** and clustername options.

So far the Tivoli Storage Manager scheduler service is created on both nodes of the cluster with exactly the same name for each resource group. The last task consists of the definition for a new resource in the Service Group.

Creation of a resource for scheduler service in VCS

For a correct configuration of the Tivoli Storage Manager client, we define, for each Service Group, a new generic service resource. This resource relates to the scheduler service name created for this group.

Important: Before continuing, make sure you stop the service created in "Creation of the Tivoli Storage Manager Scheduler service" on page 972 on all nodes. Also make sure all the resources are on one of the nodes.

We use the VERITAS Application Configuration Wizard to modify the SG-ISC group that was created in “Creating the service group for the Administrative Center” on page 933, and include two new resources: a Generic Service and a Registry Replication.

1. Click **Start** →**Programs** →**VERITAS** →**VERITAS Cluster Service** →**Application Configuration Wizard**.
2. We review the welcome page in Figure 23-2 and click **Next**.

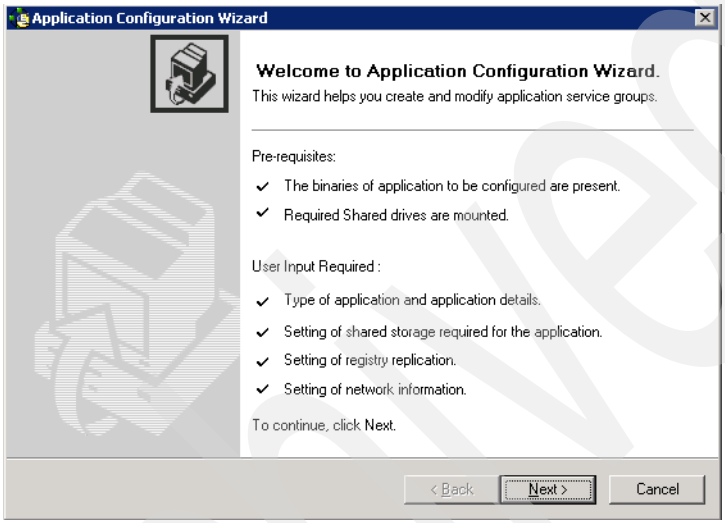

Figure 23-2 Starting the Application Configuration Wizard

Archived

3. We select **Modify service group** option as shown in Figure 23-3, select the CG-ISC group and click **Next**.

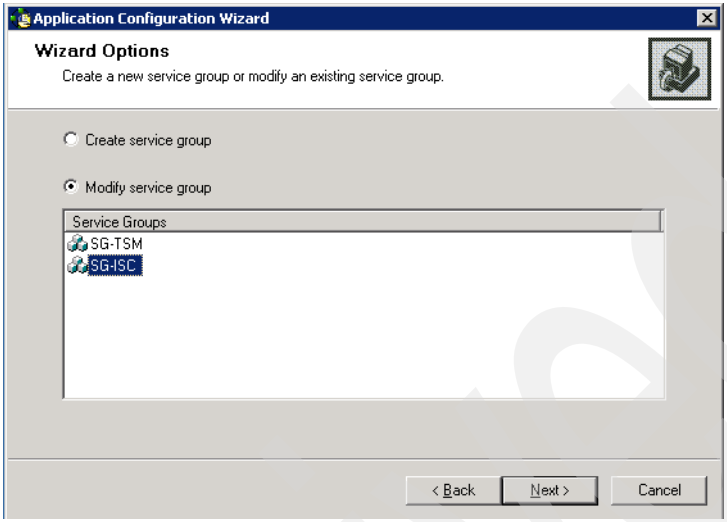

Figure 23-3 Modifying service group option

4. We receive a message that the group is not offline, but that we can create new resources, as shown in Figure 23-4. We click **Yes**.

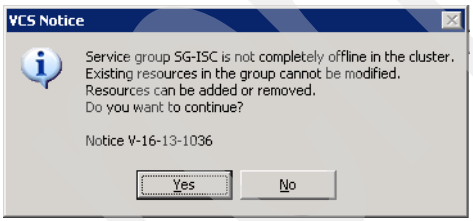

Figure 23-4 No existing resource can be changed, but new ones can be added

5. We confirm the servers that will hold the resources, as in Figure 23-5. We can set the priority between the servers moving them with the down and up arrows. We click **Next**.

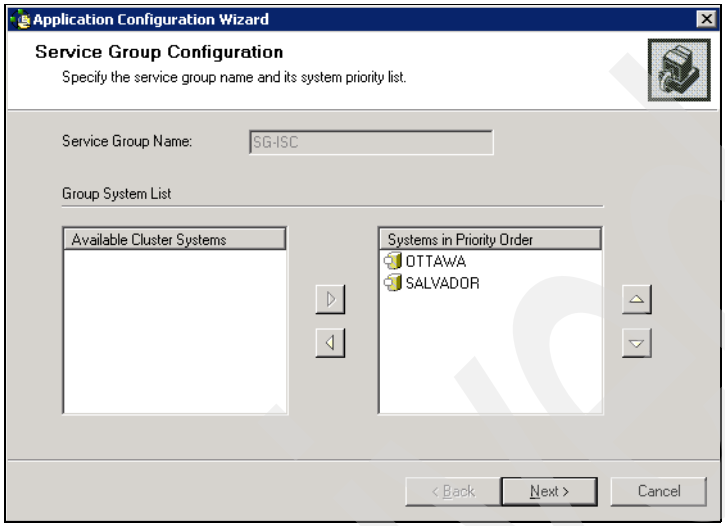

Figure 23-5 Service group configuration

6. The wizard will start a process of discovering all necessary objects to create the service group, as shown in Figure 23-6. We wait until this process ends.

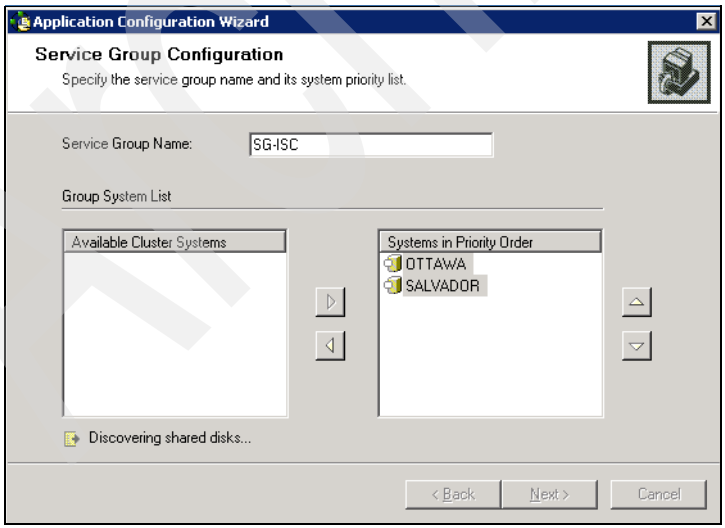

Figure 23-6 Discovering process

7. We then define what kind of application group this is. In our case there is one service: *TSM Scheduler CL_VCS02_ISC*. We choose **Generic Service** from the drop-down list in Figure 23-7 and click **Next**.

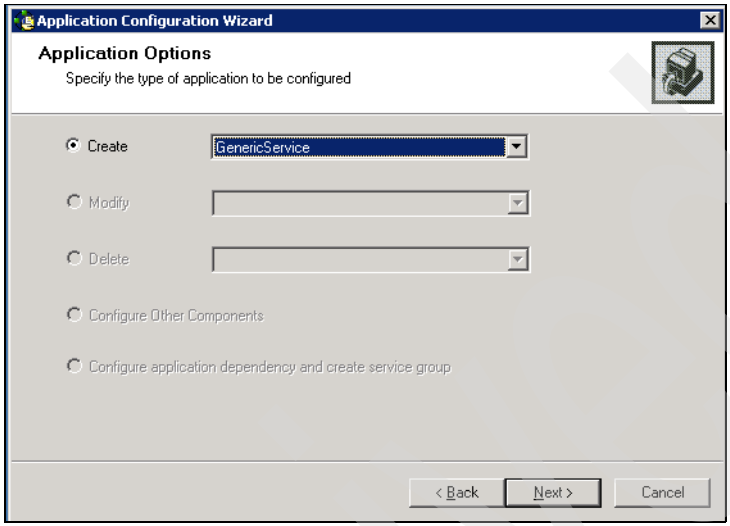

Figure 23-7 Choosing the kind of application

8. We click the button next to the **Service Name** line and choose the service *TSM Scheduler CL_VCS02_ISC* from the drop-down list as shown in Figure 23-8.

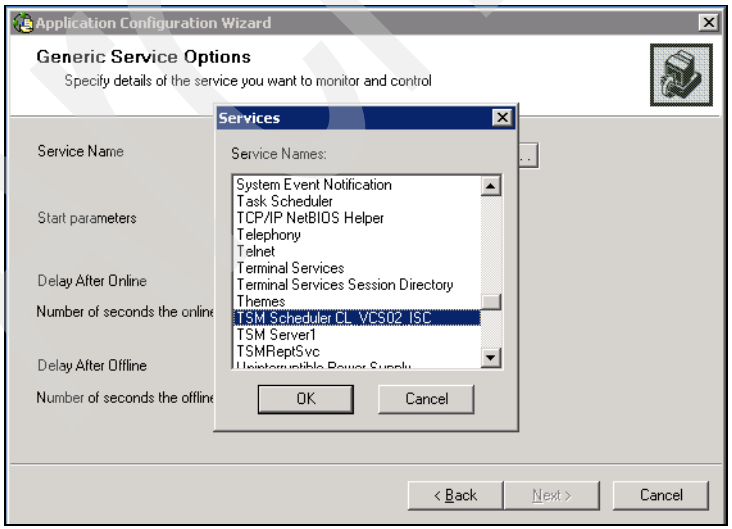

Figure 23-8 Choosing TSM Scheduler CL_VCS02_ISC service

9. We confirm the name of the service chosen and click **Next** in Figure 23-9.

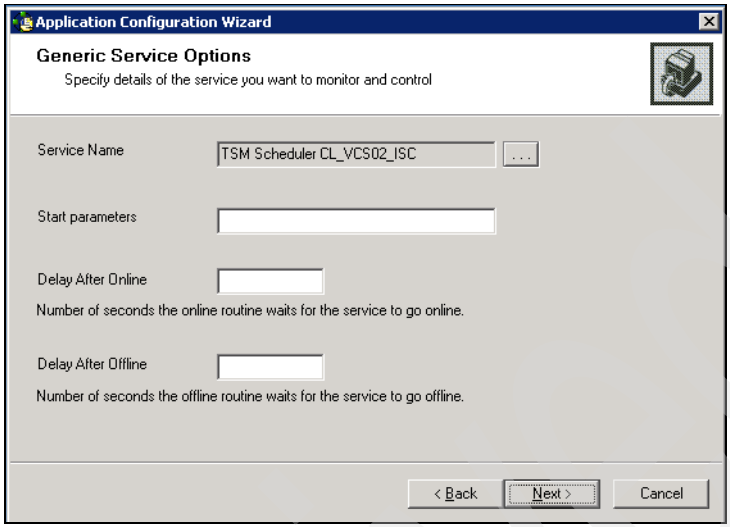

Figure 23-9 Confirming the service

10. In Figure 23-10 we choose to start the service with the LocalSystem account.

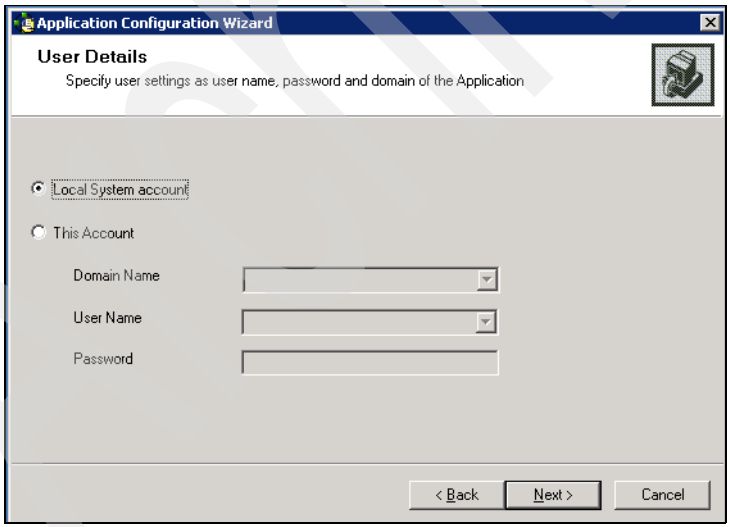

Figure 23-10 Choosing the service account

11. We select the drives that will be used by the Administration Center. We refer to Table 23-1 on page 968 to confirm the drive letters. We select the letters as in Figure 23-11 and click **Next**.

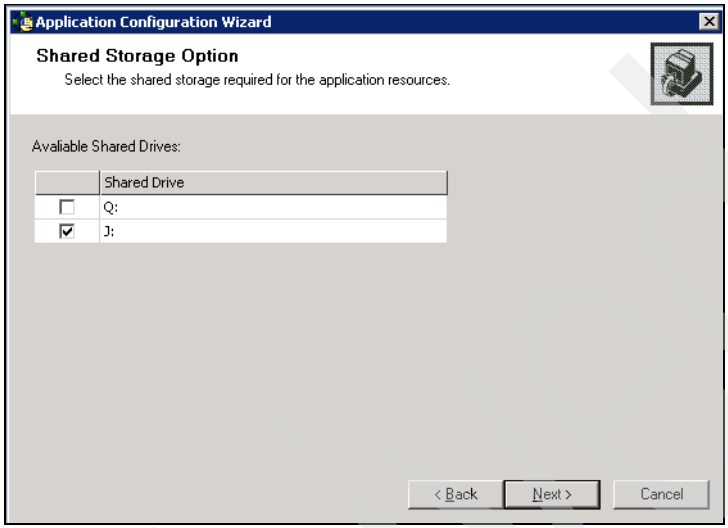

Figure 23-11 Selecting the drives to be used

12. We receive a summary of the application resource with the name and user account as in Figure 23-12. We confirm and click **Next**.

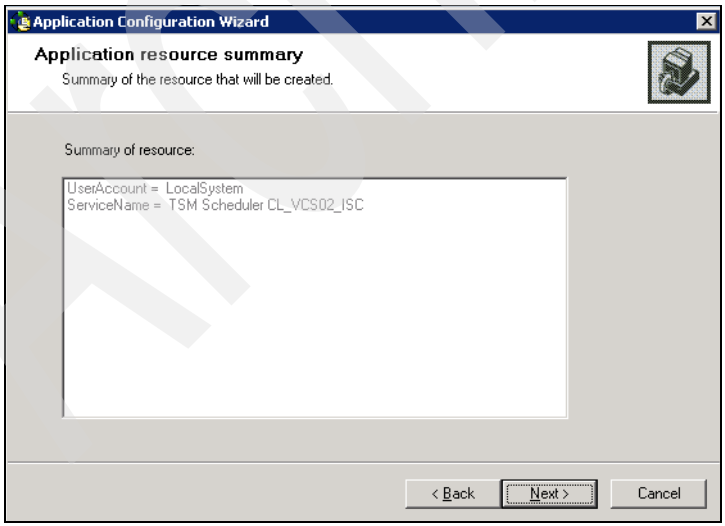

Figure 23-12 Summary with name and account for the service

13. We need one more resource for this group: Registry Replicator. So in Figure 23-13 we choose **Configure Other Components** and then click **Next**.

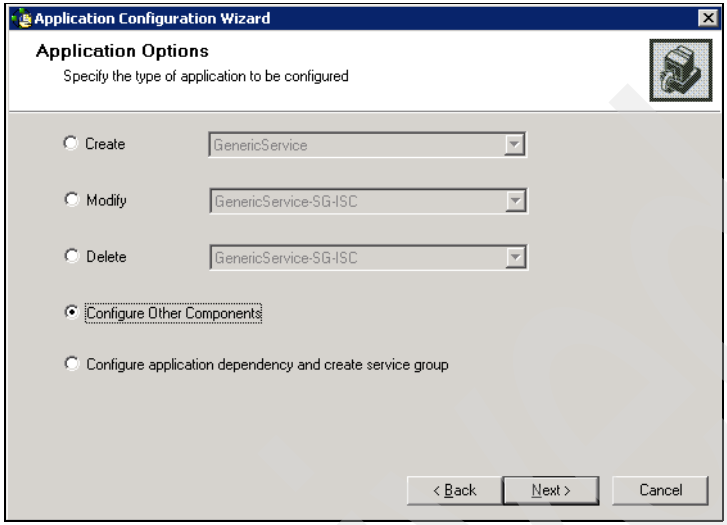

Figure 23-13 Choosing additional components

14. In Figure 23-14 we choose **Registry Replication Component** and leave checked the **Network Component** and **Lanman Component** and click Next. If we uncheck these last two, we receive a message saying the wizard would delete them.

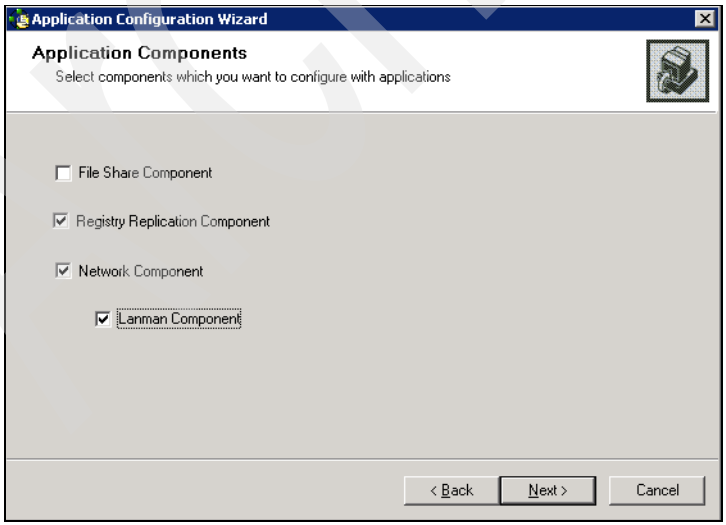

Figure 23-14 Choosing other components for Registry Replication

15. In Figure 23-15 we specify the drive letter that we are using to create this resource (J:) and then click **Add** to navigate through the registry keys until we have:

```
\HKLM\SOFTWARE\IBM\ADSM\CurrentVersion\BackupClient\Nodes\CL_VCS02_ISC>\TSM
SRV06
```

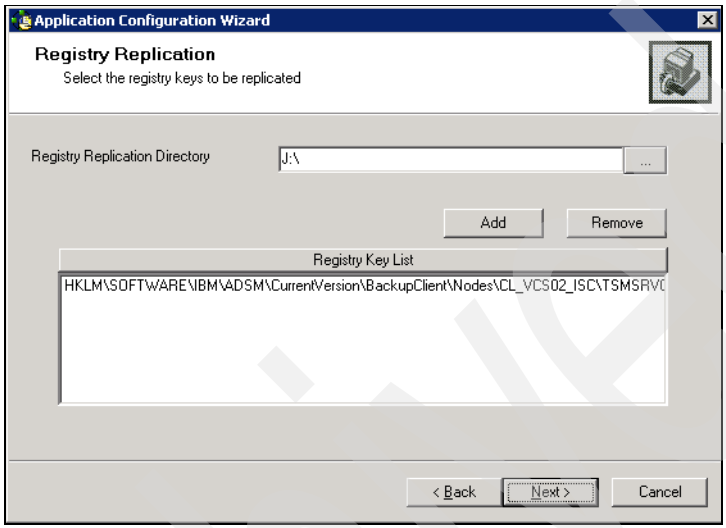

Figure 23-15 Specifying the registry key

16. In Figure 23-16 we click **Next**. This information is already stored in the cluster.

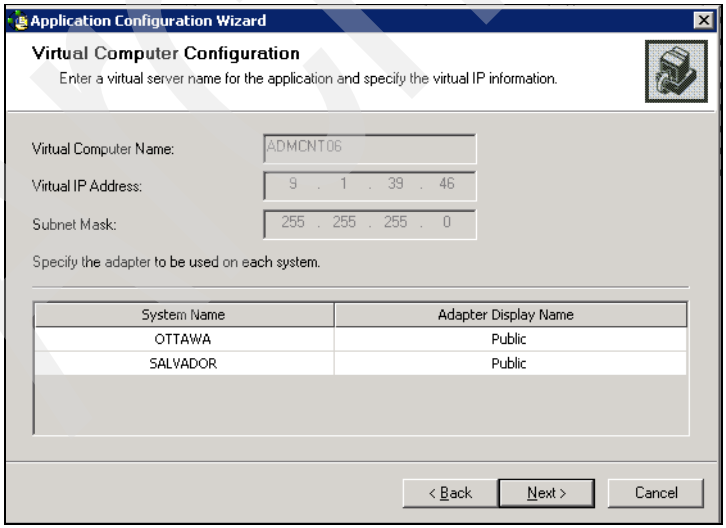

Figure 23-16 Name and IP addresses

17. We do not need any other resources to be configured. We choose **Configure application dependency and create service group** in Figure 23-17 and click **Next**.

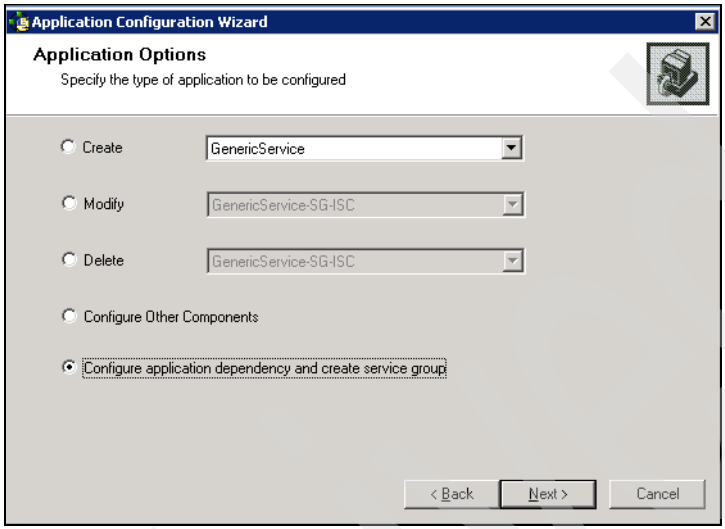

Figure 23-17 Completing the application options

18. We review the information presented in the summary, and pressing F2 we change the name of the service as shown in Figure 23-18 and click **Next**.

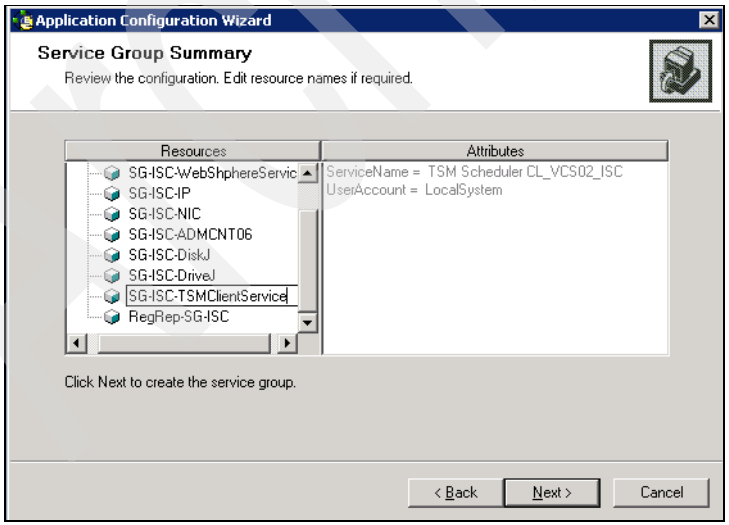

Figure 23-18 Service Group Summary

19. We confirm we want to create the service group clicking **Yes** in Figure 23-19.

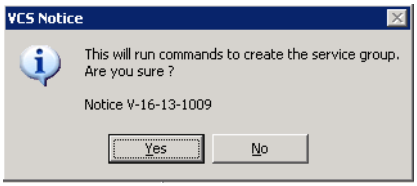

Figure 23-19 Confirming the creation of the service group

20. When the process completes, we uncheck the **Bring the service group online** option as shown in Figure 23-20. We need to confirm the dependencies before bringing this new resource online.

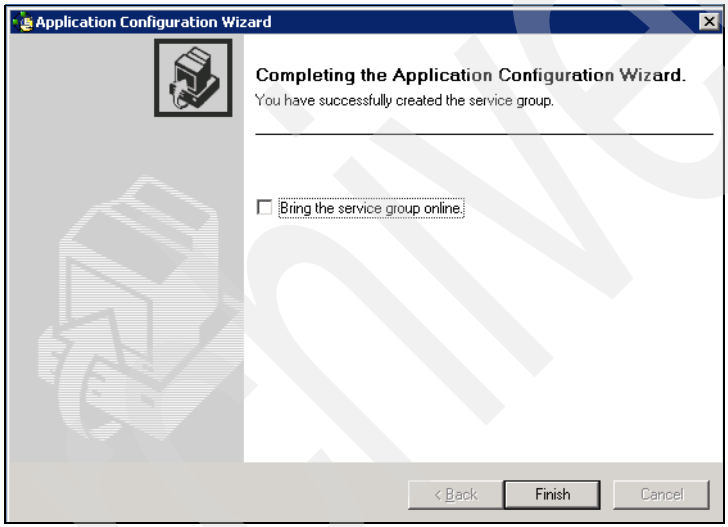

Figure 23-20 Completing the wizard

Archived

21. We adjust the links so that the result is the one shown in Figure 23-21, and then bring the resources online.

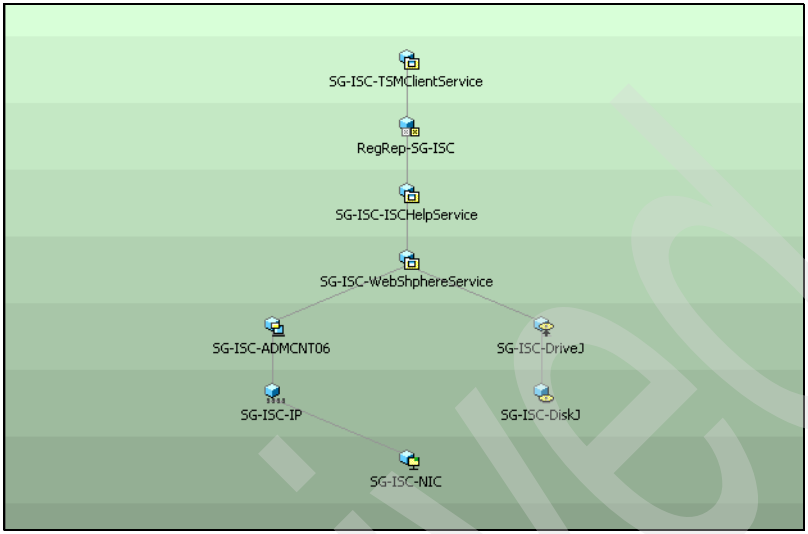

Figure 23-21 Link after creating the new resource

22. If you go to the Windows service menu, TSM Scheduler CL_VCS02_ISC service is started on OTTAWA, the node which now hosts this resource group.

23. We move the resources to check that Tivoli Storage Manager scheduler services successfully start on the second node while they are stopped on the first node.

Note: The TSM Scheduler CL_VCS02_ISC service must be brought online/offline using the Veritas Cluster Explorer, for shared resources.

Creating the Tivoli Storage Manager web client services

This task is not necessary if we do not want to use the Web client. However, if we want to be able to access virtual clients from a Web browser, we must follow the tasks explained in this section.

We create Tivoli Storage Manager Client Acceptor and Tivoli Storage Manager Remote Client Agent services on both physical nodes with the same service names and the same options.

1. We make sure we are on the server that hosts all resources in order to install the scheduler service.

2. We install the scheduler service for each group using the **dsmcutil program**. This utility is located on the Tivoli Storage Manager client installation path (c:\program files\tivoli\tsm\baclient).
3. In our lab we install one Client Acceptor service for our SG_ISC Service Group, and one Remote Client Agent service. When we start the installation the node that hosts the resources is OTTAWA.
4. We open a MS-DOS Windows command line and change to the Tivoli Storage Manager client installation path. We run the **dsmcutil tool** with the appropriate parameters to create the Tivoli Storage Manager client acceptor service for the group:

```
dsmcutil inst cad /name:"TSM Client Acceptor CL_VCS02_ISC"
/clientdir:"c:\Program Files\Tivoli\tsm\baclient" /optfile:j:\tsm\dsm.opt
/node:CL_VCS02_ISC /password:itsosj /clusternode:yes /clustername:CL_VCS02
/autostart:no /httpport:1584
```

5. After a successful installation of the client acceptor for this resource group, we run the **dsmcutil** tool again to create its remote client agent partner service typing the command:

```
dsmcutil inst remoteagent /name:"TSM Remote Client Agent CL_VCS02_ISC"
/clientdir:"c:\Program Files\Tivoli\tsm\baclient" /optfile:j:\tsm\dsm.opt
/node:CL_VCS02_ISC /password:itsosj /clusternode:yes /clustername:CL_VCS02
/startnow:no /partnername:"TSM Client Acceptor CL_VCS02_ISC"
```

Important: The client acceptor and remote client agent services *must* be installed with the same name on each physical node on the VCS, otherwise failover will not work.

6. We move the resources to the second node (SALVADOR) and repeat steps 1-5 with the same options.

So far the Tivoli Storage Manager web client services are installed on both nodes of the cluster with exactly the same names. The last task consists of the definition for new resource on the Service Group. But first we go to the Windows Service menu and stop all the web client services on SALVADOR.

Creating a generic resource for the Client Acceptor service

For a correct configuration of the Tivoli Storage Manager web client we define a new generic service resource for each Service Group. This resource will be related to the Client Acceptor service name created for this group.

Important: Before continuing, we make sure we stop all services created in "Creating the Tivoli Storage Manager web client services" on page 985 on all nodes. Also we make sure all resources are on one of the nodes.

We create the Generic Service resource for Tivoli Storage Manager Client Acceptor CL_VCS02_ISC using the Application Configuration Wizard with the following parameters as shown in Figure 23-22. We do not bring it online before we change the links.

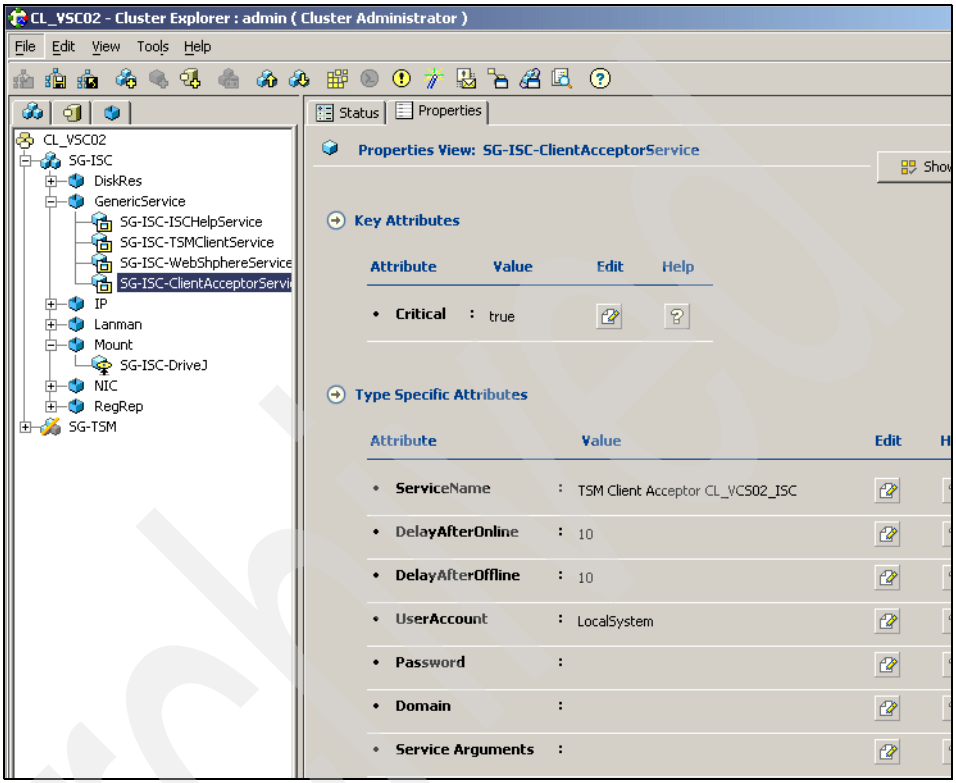

Figure 23-22 Client Acceptor Generic service parameters

7. After changing the links to what is shown in Figure 23-23, we bring the resource online and then switch the group between the servers in the cluster to test.

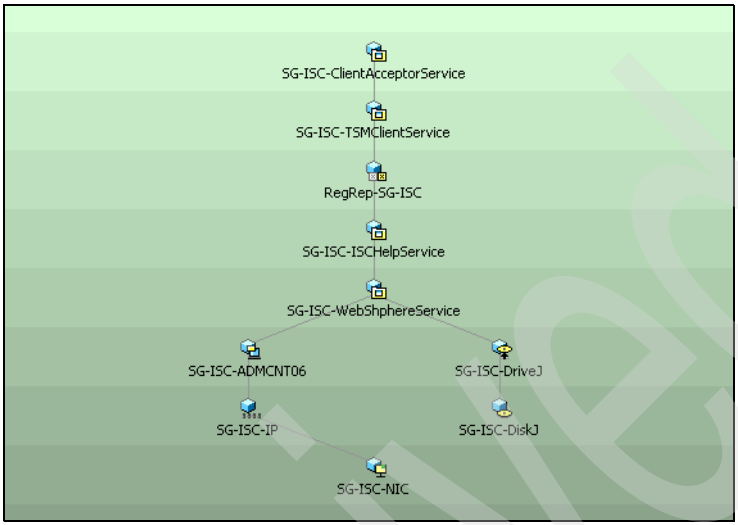

Figure 23-23 Final link with dependencies

Note: The Tivoli Storage Manager Client Acceptor service must be brought online/offline using the Cluster Explorer, for shared resources.

23.6 Testing Tivoli Storage Manager client on the VCS

In order to check the high availability of Tivoli Storage Manager client on our lab environment, we must do some testing.

Our objective with these tests is to know how Tivoli Storage Manager client can respond, on a VCS environment, after certain kinds of failures that affect the shared resources.

For the purpose of this section we use a Tivoli Storage Manager server installed on an AIX machine: TSMSRV03.

Our Tivoli Storage Manager virtual client for testing is CL_VCS02_ISC.

Archived

23.6.1 Testing client incremental backup

Our first test consists of an incremental backup started from the client.

Objective

The objective of this test is to show what happens when a client incremental backup is started for a virtual node on the VCS, and the client that hosts the resources at that moment suddenly fails.

Activities

To do this test, we perform these tasks:

1. We open the Veritas Cluster Explorer to check which node hosts the resource Tivoli Storage Manager scheduler for CL_VCS02_ISC.
2. We schedule a client incremental backup operation using the Tivoli Storage Manager server scheduler and we associate the schedule to CL_VCS02_ISC nodename.
3. A client session starts on the server for CL_VCS02_ISC and Tivoli Storage Manager server commands the tape library to mount a tape volume as shown in Figure 23-24.

```
IBM Tivoli Storage Manager
                          *  (SESSION: 25)
02/19/2005 09:21:38  ANR2017I Administrator ADMIN issued command: QUERY EVENT *
                          *  (SESSION: 25)
02/19/2005 09:22:03  ANR0406I Session 26 started for node CL_VCS02_ISC (WinNT)
                          (Tcp/Ip 9.1.39.188(3743)). (SESSION: 26)
more...   (<ENTER> to continue, 'C' to cancel)

02/19/2005 09:22:03  ANR0403I Session 26 ended for node CL_VCS02_ISC (WinNT).
                          (SESSION: 26)
02/19/2005 09:22:03  ANR0406I Session 27 started for node CL_VCS02_ISC (WinNT)
                          (Tcp/Ip 9.1.39.188(3744)). (SESSION: 27)
02/19/2005 09:22:16  ANR2017I Administrator ADMIN issued command: QUERY EVENT *
                          *  (SESSION: 25)
02/19/2005 09:22:19  ANR0406I Session 28 started for node CL_VCS02_ISC (WinNT)
                          (Tcp/Ip 9.1.39.188(3745)). (SESSION: 28)
02/19/2005 09:23:32  ANR8439I SCSI library LIBLTO is ready for operations.
                          (SESSION: 27)
02/19/2005 09:23:39  ANR2017I Administrator ADMIN issued command: QUERY SESSION
                          (SESSION: 25)
02/19/2005 09:23:42  ANR2017I Administrator ADMIN issued command: QUERY MO
                          (SESSION: 25)
02/19/2005 09:23:42  ANR8379I Mount point in device class LTO is waiting for
                          the volume mount to complete, status: WAITING FOR VOLUME.
                          (SESSION: 25)
02/19/2005 09:23:42  ANR8334I          1 matches found. (SESSION: 25)
```

Figure 23-24 A session starts for CL_VCS02_ISC in the activity log

4. When the tape volume is mounted the client starts sending files to the server, as we can see on its schedule log file shown in Figure 23-25.

```
dsmsched - Notepad
File  Edit  Format  Help
Executing scheduled command now.
02/19/2005 09:29:47 --- SCHEDULEREC OBJECT BEGIN INCR_BACKUP 02/19/2005 09:21:24
02/19/2005 09:29:47 Incremental backup of volume '\\cl_vcs02\j$'
02/19/2005 09:30:51 Directory-->                    0 \\cl_vcs02\j$\ [Sent]
02/19/2005 09:30:51 Directory-->                    0 \\cl_vcs02\j$\CmdFileCache [Sent]
02/19/2005 09:30:51 Directory-->                    0 \\cl_vcs02\j$\code [Sent]
02/19/2005 09:30:51 Directory-->                    0 \\cl_vcs02\j$\Program Files [Sent]
02/19/2005 09:30:51 Directory-->                    0 \\cl_vcs02\j$\RECYCLER [Sent]
02/19/2005 09:30:51 Directory-->                    0 \\cl_vcs02\j$\System Volume Information [s
02/19/2005 09:30:51 Directory-->                    0 \\cl_vcs02\j$\TSM [Sent]
02/19/2005 09:30:51 Directory-->                    0 \\cl_vcs02\j$\TSM_Images [Sent]
02/19/2005 09:30:51 Directory-->                    0 \\cl_vcs02\j$\code\isc [Sent]
02/19/2005 09:30:51 Normal File-->                263 \\cl_vcs02\j$\code\GSKit-README.txt [Sent]
02/19/2005 09:30:51 Normal File-->        1,205,657 \\cl_vcs02\j$\code\GSKit.zip [Sent]
02/19/2005 09:30:51 ANS1114I Waiting for mount of offline media.
02/19/2005 09:30:56 Retry # 1  Directory-->                    0 \\cl_vcs02\j$\ [Sent]
02/19/2005 09:30:56 Retry # 1  Directory-->                    0 \\cl_vcs02\j$\CmdFileCache [Ser
02/19/2005 09:30:56 Retry # 1  Directory-->                    0 \\cl_vcs02\j$\code [Sent]
02/19/2005 09:30:56 Retry # 1  Directory-->                    0 \\cl_vcs02\j$\Program Files [Se
02/19/2005 09:30:56 Retry # 1  Directory-->                    0 \\cl_vcs02\j$\RECYCLER [Sent]
02/19/2005 09:30:56 Retry # 1  Directory-->                    0 \\cl_vcs02\j$\System Volume Inf
02/19/2005 09:30:56 Retry # 1  Directory-->                    0 \\cl_vcs02\j$\TSM [Sent]
02/19/2005 09:30:56 Retry # 1  Directory-->                    0 \\cl_vcs02\j$\TSM_Images [Sent]
02/19/2005 09:30:56 Retry # 1  Directory-->                    0 \\cl_vcs02\j$\code\isc [Sent]
02/19/2005 09:30:56 Retry # 1  Normal File-->                263 \\cl_vcs02\j$\code\GSKit-README
02/19/2005 09:32:44 Retry # 1  Normal File-->        1,205,657 \\cl_vcs02\j$\code\GSKit.zip [S
02/19/2005 09:32:52 Normal File-->       25,388,669 \\cl_vcs02\j$\code\TSM530C_WIN32.exe [Sent
02/19/2005 09:32:53 Normal File-->       10,762,389 \\cl_vcs02\j$\code\TSMCON5300_WIN.exe [Ser
02/19/2005 09:32:54 Normal File-->       12,426,660 \\cl_vcs02\j$\code\TSMLNG5300_CHS_WIN_337.
```

Figure 23-25 CL_VCS02_ISC starts sending files to Tivoli Storage Manager server

Note: Notice in Figure 23-25 the name of the filespace used by Tivoli Storage Manager to store the files in the server (`\\cl_vcs02\j$`). If the client is correctly configured to work on VCS, the filespace name always starts with the cluster name. It does not use the local name of the physical node which hosts the resource at the time of backup.

5. While the client continues sending files to the server, we force a failure in the node that hosts the shared resources. The following sequence takes place:

   a. The client loses its connection with the server temporarily, and the session terminates. The tape volume is dismounted from the tape drive as we can see on the Tivoli Storage Manager server activity log shown in Figure 23-26.

```
IBM Tivoli Storage Manager
                        (SESSION: 25)
02/19/2005 09:26:23  ANR0480W Session 28 for node CL_VCS02_ISC (WinNT)
                      terminated - connection with client severed. (SESSION:
                      28)
02/19/2005 09:26:25  ANR0514I Session 27 closed volume 030AKK. (SESSION: 27)
02/19/2005 09:26:25  ANR0480W Session 27 for node CL_VCS02_ISC (WinNT)
                      terminated - connection with client severed. (SESSION:
                      27)
02/19/2005 09:27:25  ANR8325I Dismounting volume 030AKK - 1 minute mount
                      retention expired.
02/19/2005 09:27:31  ANR8336I Verifying label of LTO volume 030AKK in drive
```

Figure 23-26 Session lost for client and the tape volume is dismounted by server

b. In the Veritas Cluster Explorer, the second node tries to bring the resources online.

c. After a while the resources are online on this second node.

d. When the scheduler resource is online, the client queries the server for a scheduled command, and since it is still within the startup window, the incremental backup restarts and the tape volume is mounted again such as we can see in Figure 23-27 and Figure 23-28.

```
IBM Tivoli Storage Manager
tsm: TSMSRV03>Q EV * *
Scheduled Start      Actual Start         Schedule Name Node Name     Status
-------------------- -------------------- ------------- ------------- ---------
02/19/2005 07:41:08                       TEST_SCHED    CL_HACMP03_C- Missed
                                                         LIENT
02/19/2005 09:21:24  02/19/2005 09:22:03  INCR_BACKUP   CL_UCS02_ISC  Restarted
tsm: TSMSRV03>_
```

Figure 23-27 The event log shows the schedule as restarted

```
IBM Tivoli Storage Manager
                    (SESSION: 30)
more...   (<ENTER> to continue, 'C' to cancel)

02/19/2005 09:28:46  ANR0406I Session 31 started for node CL_UCS02_ISC (WinNT)
                      (Tcp/Ip 9.1.39.187(3521)). (SESSION: 31)
02/19/2005 09:28:52  ANR8468I LTO volume 030AKK dismounted from drive DRLTO_1
                      (/dev/rmt0) in library LIBLTO. (SESSION: 27)
02/19/2005 09:28:58  ANR0406I Session 32 started for node CL_UCS02_ISC (WinNT)
                      (Tcp/Ip 9.1.39.187(3523)). (SESSION: 32)
02/19/2005 09:29:43  ANR8337I LTO volume 030AKK mounted in drive DRLTO_2
                      (/dev/rmt1). (SESSION: 31)
02/19/2005 09:29:43  ANR0511I Session 31 opened output volume 030AKK. (SESSION:
                      31)
02/19/2005 09:29:52  ANR2017I Administrator ADMIN issued command: QUERY EVENT *
```

Figure 23-28 The tape volume is mounted again for schedule to restart backup

Archived

6. The incremental backup ends without errors as shown on the schedule log file in Figure 23-29.

```
dsmsched - Notepad
File  Edit  Format  Help
02/19/2005 09:47:51 Normal File-->          348,160 \\cl_vcs02\j$\TSM_Images\TSM_BA_Client
02/19/2005 09:47:52 Normal File-->        1,241,088 \\cl_vcs02\j$\TSM_Images\TSM_BA_Client
02/19/2005 09:47:52 Normal File-->           98,304 \\cl_vcs02\j$\TSM_Images\TSM_BA_Client
02/19/2005 09:47:57 Successful incremental backup of '\\cl_vcs02\j$'

02/19/2005 09:47:57 --- SCHEDULEREC STATUS BEGIN
02/19/2005 09:47:57 Total number of objects inspected:   20,239
02/19/2005 09:47:57 Total number of objects backed up:   19,744
02/19/2005 09:47:57 Total number of objects updated:         0
02/19/2005 09:47:57 Total number of objects rebound:         0
02/19/2005 09:47:57 Total number of objects deleted:         0
02/19/2005 09:47:57 Total number of objects expired:         0
02/19/2005 09:47:57 Total number of objects failed:          0
02/19/2005 09:47:57 Total number of bytes transferred:     1.20 GB
02/19/2005 09:47:57 Data transfer time:                   96.78 sec
02/19/2005 09:47:57 Network data transfer rate:        13,002.18 KB/sec
02/19/2005 09:47:57 Aggregate data transfer rate:       1,830.70 KB/sec
02/19/2005 09:47:57 Objects compressed by:                   0%
02/19/2005 09:47:57 Elapsed processing time:           00:11:27
02/19/2005 09:47:57 --- SCHEDULEREC STATUS END
02/19/2005 09:47:57 --- SCHEDULEREC OBJECT END INCR_BACKUP 02/19/2005 09:21:24
02/19/2005 09:47:57 Scheduled event 'INCR_BACKUP' completed successfully.
02/19/2005 09:47:57 Sending results for scheduled event 'INCR_BACKUP'.
02/19/2005 09:47:57 Results sent to server for scheduled event 'INCR_BACKUP'.
```

Figure 23-29 Schedule log shows the backup as completed

7. In the Tivoli Storage Manager server event log, the schedule is *completed* as we see in Figure 23-30.

```
IBM Tivoli Storage Manager
tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start         Schedule Name  Node Name      Status
-------------------- -------------------- -------------- -------------- ----------
02/19/2005 07:41:08                       TEST_SCHED     CL_HACMP03_C-  Missed
                                                          LIENT
02/19/2005 09:21:24  02/19/2005 09:22:03  INCR_BACKUP    CL_VCS02_ISC   Completed

tsm: TSMSRV03>_
```

Figure 23-30 Schedule completed on the event log

Results summary

The test results show that, after a failure on the node that hosts the Tivoli Storage Manager scheduler service resource, a scheduled incremental backup started on one node of a Windows VCS is restarted and successfully completed on the other node that takes the failover.

This is true if the startup window used to define the schedule is not elapsed when the scheduler services restarts on the second node.

The backup restarts from the point of the last committed transaction in the Tivoli Storage Manager server database.

Archived

23.6.2 Testing client restore

Our second test consists of a scheduled restore of certain files under a directory.

Objective

The objective of this test is to show what happens when a client restore is started for a virtual node on the VCS, and the client that hosts the resources at that moment fails.

Activities

To do this test, we perform these tasks:

1. We open the Veritas Cluster Explorer to check which node hosts the Tivoli Storage Manager scheduler resource.
2. We schedule a client restore operation using the Tivoli Storage Manager server scheduler and we associate the schedule to CL_VCS02_ISC nodename.
3. In the event log the schedule reports as started. In the activity log a session is started for the client and a tape volume is mounted. We see all these events in Figure 23-31 and Figure 23-32.

```
IBM Tivoli Storage Manager
------ ------ ------ ------ ------ ------ ------ -------- --------------------
    40 Tcp/Ip IdleW  1.6 M    2.5 K    492 Admin WinNT    ADMIN
    42 Tcp/Ip SendW    0 S   84.1 M    747 Node  WinNT    CL_VCS02_ISC
    43 Tcp/Ip Run      0 S      127    124 Admin WinNT    ADMIN

tsm: TSMSRV03>q mo
ANR8330I LTO volume 030AKK is mounted R/W in drive DRLTO_1 (/dev/rmt0), status:
IN USE.
ANR8334I         1 matches found.

tsm: TSMSRV03>q mo
ANR8330I LTO volume 030AKK is mounted R/W in drive DRLTO_1 (/dev/rmt0), status:
IN USE.
ANR8334I         1 matches found.

tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name     Status
-------------------- -------------------- ------------- ------------- ---------
02/19/2005 07:41:08                       TEST_SCHED    CL_HACMP03_C- Missed
                                                         LIENT
02/19/2005 09:21:24  02/19/2005 09:22:03  INCR_BACKUP   CL_VCS02_ISC  Completed
02/19/2005 10:08:00  02/19/2005 10:09:01  RESTORE       CL_VCS02_ISC  Started

tsm: TSMSRV03>_
```

Figure 23-31 Scheduled restore started for CL_MSCS01_SA

Archived

```
IBM Tivoli Storage Manager
                          (SESSION: 40)
02/19/2005 10:09:01  ANR0406I Session 41 started for node CL_VCS02_ISC (WinNT)
                      (Tcp/Ip 9.1.39.187(3697)). (SESSION: 41)
02/19/2005 10:09:01  ANR0403I Session 41 ended for node CL_VCS02_ISC (WinNT).
                      (SESSION: 41)
more...   (<ENTER> to continue, 'C' to cancel)

02/19/2005 10:09:01  ANR0406I Session 42 started for node CL_VCS02_ISC (WinNT)
                      (Tcp/Ip 9.1.39.187(3698)). (SESSION: 42)
02/19/2005 10:09:29  ANR8337I LTO volume 030AKK mounted in drive DRLTO_1
                      (/dev/rmt0). (SESSION: 42)
02/19/2005 10:09:29  ANR0510I Session 42 opened input volume 030AKK. (SESSION:
                      42)
02/19/2005 10:10:17  ANR0407I Session 43 started for administrator ADMIN
                      (WinNT) (Tcp/Ip 9.1.39.188(1416)). (SESSION: 43)
02/19/2005 10:10:18  ANR2017I Administrator ADMIN issued command: QUERY SESSION
                       (SESSION: 43)
02/19/2005 10:10:22  ANR2017I Administrator ADMIN issued command: QUERY MO
                      (SESSION: 43)
02/19/2005 10:10:22  ANR8330I LTO volume 030AKK is mounted R/W in drive DRLTO_1
                      (/dev/rmt0), status: IN USE. (SESSION: 43)
02/19/2005 10:10:22  ANR8334I         1 matches found. (SESSION: 43)
02/19/2005 10:10:35  ANR0405I Session 40 ended for administrator ADMIN (WinNT).
                      (SESSION: 40)
02/19/2005 10:11:08  ANR2017I Administrator ADMIN issued command: QUERY MO
```

Figure 23-32 A session is started for restore and the tape volume is mounted

4. The client starts restoring files as we can see on the schedule log file in Figure 23-33.

```
dsmsched - Notepad
File  Edit  Format  Help
Executing scheduled command now.
02/19/2005 10:16:45 --- SCHEDULEREC OBJECT BEGIN RESTORE 02/19/2005 10:08:00
02/19/2005 10:16:45 Restore function invoked.

02/19/2005 10:16:46 ANS1247I Waiting for files from the server... ** Interrupted **
02/19/2005 10:17:52 ANS1114I Waiting for mount of offline media.
02/19/2005 10:17:52 Restoring              0 \\cl_vcs02\j$\code\isc [Done]
02/19/2005 10:17:52 Restoring            263 \\cl_vcs02\j$\code\GSKit-README.txt [Done]
02/19/2005 10:17:52 Restoring      1,205,657 \\cl_vcs02\j$\code\GSKit.zip [Done]
02/19/2005 10:17:55 Restoring     25,388,669 \\cl_vcs02\j$\code\TSM530C_WIN32.exe [Done]
02/19/2005 10:17:56 Restoring     10,762,389 \\cl_vcs02\j$\code\TSMCON5300_WIN.exe [Done]
02/19/2005 10:17:58 Restoring     12,426,660 \\cl_vcs02\j$\code\TSMLNG5300_CHS_WIN_337.exe [Do
02/19/2005 10:17:59 Restoring     12,444,265 \\cl_vcs02\j$\code\TSMLNG5300_CHT_WIN_337.exe [Do
02/19/2005 10:18:00 Restoring     12,865,309 \\cl_vcs02\j$\code\TSMLNG5300_DEU_WIN_337.exe [Do
02/19/2005 10:18:02 Restoring     12,826,564 \\cl_vcs02\j$\code\TSMLNG5300_ESP_WIN_337.exe [Do
02/19/2005 10:18:04 Restoring     12,881,222 \\cl_vcs02\j$\code\TSMLNG5300_FRA_WIN_337.exe [Do
02/19/2005 10:18:12 Restoring     12,770,448 \\cl_vcs02\j$\code\TSMLNG5300_ITA_WIN_337.exe [Do
```

Figure 23-33 Restore starts in the schedule log file

Archived

5. While the client is restoring the files, we force a failure in the node that hosts the scheduler service. The following sequence takes place:

   a. The client loses temporarily its connection with the server, the session is terminated and the tape volume is dismounted as we can see on the Tivoli Storage Manager server activity log shown in Figure 23-34.

```
IBM Tivoli Storage Manager
                          (SESSION: 43)
02/19/2005 10:11:08  ANR8330I LTO volume 030AKK is mounted R/W in drive DRLTO_1
                      (/dev/rmt0), status: IN USE. (SESSION: 43)
02/19/2005 10:11:08  ANR8334I          1 matches found. (SESSION: 43)
02/19/2005 10:11:11  ANR2017I Administrator ADMIN issued command: QUERY EVENT *
                      *  (SESSION: 43)
02/19/2005 10:11:27  ANR0514I Session 42 closed volume 030AKK. (SESSION: 42)
02/19/2005 10:11:27  ANR0480W Session 42 for node CL_VCS02_ISC (WinNT)
                      terminated - connection with client severed. (SESSION:
                      42)
02/19/2005 10:11:39  ANR2017I Administrator ADMIN issued command: QUERY ACTLOG
                      (SESSION: 43)
02/19/2005 10:12:27  ANR8325I Dismounting volume 030AKK - 1 minute mount
                      retention expired.
02/19/2005 10:12:27  ANR8336I Verifying label of LTO volume 030AKK in drive
                      DRLTO_1 (/dev/rmt0). (SESSION: 42)
02/19/2005 10:12:38  ANR2017I Administrator ADMIN issued command: QUERY ACTLOG
                      (SESSION: 43)
more...   (<ENTER> to continue, 'C' to cancel)
```

Figure 23-34 Session is lost and the tape volume is dismounted

   b. In the Veritas Cluster Explorer, the second node starts to bring the resources online.

   c. The client receives an error message in its schedule log file such as we see in Figure 23-35.

```
dsmsched - Notepad
File  Edit  Format  Help
02/19/2005 10:17:52 Restoring       1,205,657 \\cl_vcs02\j$\code\GSKit.zip [Done]
02/19/2005 10:17:55 Restoring      25,388,669 \\cl_vcs02\j$\code\TSM530C_WIN32.exe [Done]
02/19/2005 10:17:56 Restoring      10,762,389 \\cl_vcs02\j$\code\TSMCON5300_WIN.exe [Done]
02/19/2005 10:17:58 Restoring      12,426,660 \\cl_vcs02\j$\code\TSMLNG5300_CHS_WIN_337.exe [Do
02/19/2005 10:17:59 Restoring      12,444,265 \\cl_vcs02\j$\code\TSMLNG5300_CHT_WIN_337.exe [Do
02/19/2005 10:18:00 Restoring      12,865,309 \\cl_vcs02\j$\code\TSMLNG5300_DEU_WIN_337.exe [Do
02/19/2005 10:18:02 Restoring      12,826,564 \\cl_vcs02\j$\code\TSMLNG5300_ESP_WIN_337.exe [Do
02/19/2005 10:18:04 Restoring      12,881,222 \\cl_vcs02\j$\code\TSMLNG5300_FRA_WIN_337.exe [Do
02/19/2005 10:18:12 Restoring      12,770,448 \\cl_vcs02\j$\code\TSMLNG5300_ITA_WIN_337.exe [Do
02/19/2005 10:18:14 Restoring      12,885,559 \\cl_vcs02\j$\code\TSMLNG5300_JPN_WIN_337.exe [Do
02/19/2005 10:18:15 Restoring      12,543,008 \\cl_vcs02\j$\code\TSMLNG5300_KOR_WIN_337.exe [Do
02/19/2005 10:18:17 Restoring      12,718,897 \\cl_vcs02\j$\code\TSMLNG5300_PTB_WIN_337.exe [Do
02/19/2005 10:18:33 ANS2820E An interrupt has occurred. The current operation will end and the
client will shut down.
```

Figure 23-35 The restore process is interrupted in the client

   d. After a while the resources are online on the second node.

   e. When the Tivoli Storage Manager scheduler service resource is again online and queries the server, if the startup window for the scheduled operation is not elapsed, the restore process restarts from the beginning, as we can see on the schedule log file in Figure 23-36.

```
dsmsched - Notepad
File  Edit  Format  Help

02/19/2005 10:20:44 --- SCHEDULEREC QUERY BEGIN
02/19/2005 10:20:44 --- SCHEDULEREC QUERY END
02/19/2005 10:20:44 Next operation scheduled:
02/19/2005 10:20:44 ------------------------------------------------------------
02/19/2005 10:20:44 Schedule Name:         RESTORE
02/19/2005 10:20:44 Action:                Restore
02/19/2005 10:20:44 Objects:               j:\code\*
02/19/2005 10:20:44 Options:               -subdir=yes -replace=yes
02/19/2005 10:20:44 Server Window Start:   10:08:00 on 02/19/2005
02/19/2005 10:20:44 ------------------------------------------------------------
02/19/2005 10:20:44
Executing scheduled command now.
02/19/2005 10:20:44 --- SCHEDULEREC OBJECT BEGIN RESTORE 02/19/2005 10:08:00
02/19/2005 10:20:44 Restore function invoked.

02/19/2005 10:20:45 ANS1247I Waiting for files from the server... ** Interrupted **
02/19/2005 10:22:38 ANS1114I Waiting for mount of offline media.
02/19/2005 10:22:38 Restoring              0 \\cl_vcs02\j$\code\isc [Done]
02/19/2005 10:22:38 Restoring            263 \\cl_vcs02\j$\code\GSKit-README.txt [Done]
02/19/2005 10:22:38 Restoring      1,205,657 \\cl_vcs02\j$\code\GSKit.zip [Done]
02/19/2005 10:22:41 Restoring     25,388,669 \\cl_vcs02\j$\code\TSM530C_WIN32.exe [Done]
02/19/2005 10:22:42 Restoring     10,762,389 \\cl_vcs02\j$\code\TSMCON5300_WIN.exe [Done]
02/19/2005 10:22:43 Restoring     12,426,660 \\cl_vcs02\j$\code\TSMLNG5300_CHS_WIN_337.exe [Dc
02/19/2005 10:22:44 Restoring     12,444,265 \\cl_vcs02\j$\code\TSMLNG5300_CHT_WIN_337.exe [Dc
02/19/2005 10:22:45 Restoring     12,865,309 \\cl_vcs02\j$\code\TSMLNG5300_DEU_WIN_337.exe [Dc
02/19/2005 10:22:46 Restoring     12,826,564 \\cl_vcs02\j$\code\TSMLNG5300_ESP_WIN_337.exe [Dc
```

Figure 23-36 Restore schedule restarts in client restoring files from the beginning

f. The event log of Tivoli Storage Manager server shows the schedule as restarted:

```
IBM Tivoli Storage Manager
02/19/2005 10:12:27  ANR8325I Dismounting volume 030AKK - 1 minute mount
                      retention expired.
02/19/2005 10:12:27  ANR8336I Verifying label of LTO volume 030AKK in drive
                      DRLTO_1 (/dev/rmt0). (SESSION: 42)
02/19/2005 10:12:38  ANR2017I Administrator ADMIN issued command: QUERY ACTLOG
                      (SESSION: 43)
more...   (<ENTER> to continue, 'C' to cancel)

tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start         Schedule Name Node Name     Status
-------------------- -------------------- ------------- ------------- ---------
02/19/2005 07:41:08                       TEST_SCHED    CL_HACMP03_C- Missed
                                                         LIENT
02/19/2005 09:21:24  02/19/2005 09:22:03  INCR_BACKUP   CL_VCS02_ISC  Completed
02/19/2005 10:08:00  02/19/2005 10:09:01  RESTORE       CL_VCS02_ISC  Restarted

tsm: TSMSRV03>_
```

Figure 23-37 Schedule restarted on the event log for CL_MSCS01_ISC

6. When the restore completes, we can see the final statistics in the schedule log file of the client for a successful operation as shown in Figure 23-38.

```
dsmsched - Notepad
File  Edit  Format  Help
02/19/2005 10:23:36 Restoring      29,065,149 \\cl_vcs02\j$\code\TSMSTA5301_WIN.exe [Done]
02/19/2005 10:23:48
Restore processing finished.
02/19/2005 10:23:49 --- SCHEDULEREC STATUS BEGIN
02/19/2005 10:23:49 Total number of objects restored:        17
02/19/2005 10:23:49 Total number of objects failed:           0
02/19/2005 10:23:49 Total number of bytes transferred:    373.55 MB
02/19/2005 10:23:49 Data transfer time:                    30.97 sec
02/19/2005 10:23:49 Network data transfer rate:        12,350.26 KB/sec
02/19/2005 10:23:49 Aggregate data transfer rate:       2,065.25 KB/sec
02/19/2005 10:23:49 Elapsed processing time:            00:03:05
02/19/2005 10:23:49 --- SCHEDULEREC STATUS END
02/19/2005 10:23:49 --- SCHEDULEREC OBJECT END RESTORE 02/19/2005 10:08:00
02/19/2005 10:23:49 --- SCHEDULEREC STATUS BEGIN
02/19/2005 10:23:49 --- SCHEDULEREC STATUS END
02/19/2005 10:23:49 Scheduled event 'RESTORE' completed successfully.
02/19/2005 10:23:49 Sending results for scheduled event 'RESTORE'.
02/19/2005 10:23:49 Results sent to server for scheduled event 'RESTORE'.
```

Figure 23-38 Restore completes successfully in the schedule log file

Results summary

The test results show that after a failure on the node that hosts the Tivoli Storage Manager client scheduler instance, a scheduled restore operation started on this node is started again on the second node of the VCS when the service is online.

This is true if the startup window for the scheduled restore operation is not elapsed when the scheduler client is online again on the second node.

Also notice that the restore is not restarted from the point of failure, but started from the beginning. The scheduler queries the Tivoli Storage Manager server for a scheduled operation, and a new session is opened for the client after the failover.

23.7 Backing up VCS configuration files

There is a VERITAS tool named **hasnap** that can be used to back up and restore configuration files. This tool can be used in addition to the normal Tivoli Storage Manager backup-archive client. This is a valuable tool to use before making any changes to the existing configuration.

Archived

Archived

24

VERITAS Cluster Server and the IBM Tivoli Storage Manager Storage Agent

This chapter describes the use of Tivoli Storage Manager for Storage Area Network (also known as the *Storage Agent*) to back up the shared data of our Windows 2003 VCS using the LAN-free path.

© Copyright IBM Corp. 2005. All rights reserved.

24.1 Overview

The functionality of Tivoli Storage Manager for Storage Area Network (Storage Agent) is described in “IBM Tivoli Storage Manager for Storage Area Networks V5.3” on page 14.

In this chapter we focus on the use of this feature applied to our Windows 2003 VCS environment.

24.2 Planning and design

There are different types of hardware configurations that take advantage of using the Storage Agent for LAN-free backup in a SAN.

We must carefully plan and design our configuration, always referring to the compatibility and support requirements for Tivoli Storage Manager for Storage Area Network to work correctly.

In our lab we use IBM disk and tape Fibre Channel attached storage devices supported by LAN-free backup with Tivoli Storage Manager.

24.2.1 System requirements

Before implementing Tivoli Storage Manager for Storage Area Network, we download the latest available software levels of all components and check supported hardware and software configurations. For information, see:

http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html

In order to use the Storage Agent for LAN-free backup, we need:

- A Tivoli Storage Manager server with LAN-free license.
- A Tivoli Storage Manager client or a Tivoli Storage Manager Data Protection application client.
- A supported Storage Area Network configuration where storage devices and servers are attached for storage sharing purposes.
- If you are sharing disk storage, Tivoli SANergy must be installed. Tivoli SANergy Version 3.2.4 is included with the Storage Agent media.
- The Tivoli Storage Manager for Storage Area Network software.

Archived

24.2.2 System information

We gather all the information about our future client and server systems and use it to implement the LAN-free backup environment according to our needs:

We need to plan and design carefully things such as:

- Name conventions for local nodes, virtual nodes, and Storage Agents
- Number of Storage Agents to use, depending upon the connections
- Number of tape drives to be shared and which servers will share them
- Segregate different types of data:
  - Large files and databases to use the LAN-free path
  - Small and numerous files to use the LAN path
- TCP/IP addresses and ports
- Device names used by Windows 2003 operating system for the storage devices

24.3 Lab setup

Our Tivoli Storage Manager clients and Storage Agents for the purpose of this chapter are located on the same Veritas Windows 2003 Advanced Server Cluster we introduce in “Installing the VERITAS Storage Foundation HA for Windows environment” on page 879.

Refer to Table 21-1 on page 881, Table 21-2 on page 882, and Table 21-3 on page 882, for details of the cluster configuration: local nodes, virtual nodes, and Service Groups.

We use TSMSRV03, an AIX machine, as the server because Tivoli Storage Manager Version 5.3.0 for AIX is, so far, the only platform that supports high availability Library Manager functions for LAN-free backup.

Archived

24.3.1 Tivoli Storage Manager LAN-free configuration details

Figure 24-1 shows our LAN-free configuration:

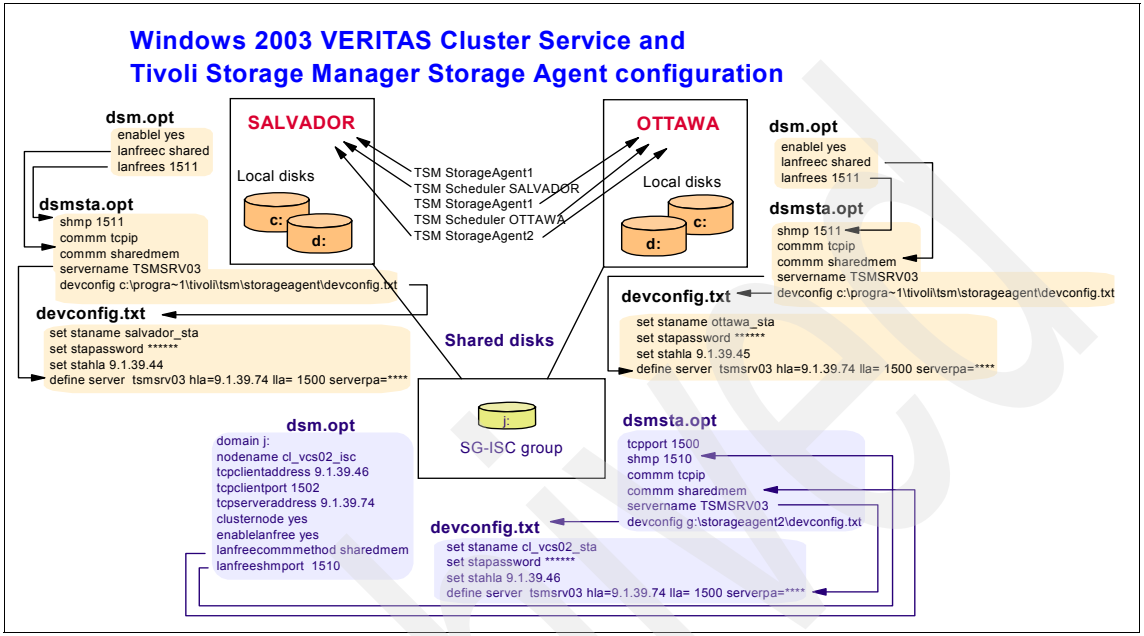

Figure 24-1 Clustered Windows 2003 configuration with Storage Agent

For details of this configuration, refer to Table 24-1, Table 24-2, and Table 24-3 below.

Archived

Table 24-1 LAN-free configuration details

Node 1	
TSM nodename	SALVADOR
Storage Agent name	SALVADOR_STA
Storage Agent service name	TSM StorageAgent1
dsmsta.opt and devconfig.txt location	c:\program files\tivoli\tsm\storageagent
Storage Agent high level address	9.1.39.44
Storage Agent low level address	1502
Storage Agent shared memory port	1511
LAN-free communication method	sharedmem
Node 2	
TSM nodename	OTTAWA
Storage Agent name	OTTAWA_STA
Storage Agent service name	TSM StorageAgent1
dsmsta.opt and devconfig.txt location	c:\program files\tivoli\tsm\storageagent
Storage Agent high level address	9.1.39.45
Storage Agent low level address	1502
Storage Agent shared memory port	1511
LAN-free communication method	sharedmem
Virtual node	
TSM nodename	CL_VCS02_TSM
Storage Agent name	CL_VCS02_STA
Storage Agent service name	TSM StorageAgent2
dsmsta.opt and devconfig.txt location	j:\storageagent2
Storage Agent high level address	9.1.39.46
Storage Agent low level address	1500
Storage Agent shared memory port	1510
LAN-free communication method	sharedmem

Archived

Table 24-2 TSM server details

TSM Server information	
Server name	TSMSRV03
High level address	9.1.39.74
Low level address	1500
Server password for server-to-server communication	password

Our SAN storage devices are described in Table 24-3.

Table 24-3 SAN devices details

SAN devices	
Disk	IBM DS4500 Disk Storage Subsystem
Tape Library	IBM LTO 3582 Tape Library
Tape drives	IBM 3580 Ultrium 2 tape drives
Tape drive device names for Storage Agents	drlto_1: mt0.0.0.2 drlto_2: mt1.0.0.2

24.4 Installation

For the installation of the Storage Agent code, we follow the steps described in “Installation of the Storage Agent” on page 332.

IBM 3580 tape drives also need to be updated. Refer to “Installing IBM 3580 tape drive drivers in Windows 2003” on page 381 for details.

24.5 Configuration

The installation and configuration of the Storage Agent involves three steps:

1. Configuration of Tivoli Storage Manager server for LAN-free.
2. Configuration of the Storage Agent for local nodes.
3. Configuration of the Storage Agent for virtual nodes.

Archived

24.5.1 Configuration of Tivoli Storage Manager server for LAN-free

The process of preparing a server for LAN-free data movement is very complex, involving several phases.

Each Storage Agent must be defined as a server in the Tivoli Storage Manager server. For our lab, we define one Storage Agent for each local node and another one for the cluster node.

In 7.4.2, “Configuration of the Storage Agent on Windows 2000 MSCS” on page 339 we show how to set up server-to-server communications and path definitions using the new administrative center console. In this chapter we use the administrative command line instead.

The following tasks are performed in the AIX Server TSMSRV03, where we assume the clients for backup/archive using LAN are already existent:

1. Preparation of the server for enterprise management. We use the following commands:

```
set servername tsmsrv03
set serverpassword password
set serverhladress 9.1.39.74
set serverlladdress 1500
```

2. Definition of the Storage Agents as servers. We use the following commands:

```
define server salvador_sta serverpa=itsosj hla=9.1.39.44 lla=1500
define server ottawa_sta serverpa=itsosj hla=9.1.39.45 lla=1500
define server cl_vcs02_sta serverpa=itsosj hla=9.1.39.46 lla=1500
```

3. Change of the node’s properties to allow either LAN or LAN-free movement of data:

```
update node salvador datawritepath=any datareadpath=any
update node ottawadatawritepath=any datareadpath=any
update node cl_vcs02_tsm datawritepath=any datareadpath=any
```

4. Definition of the tape library as shared (if this was not done when the library was first defined):

```
update library liblto shared=yes
```

5. Definition of paths from the Storage Agents to each tape drive in the Tivoli Storage Manager server. We use the following commands:

```
define path salvador_sta drlto_1 srctype=server desttype=drive
library=liblto device=mt0.0.0.2
define path salvador_sta drlto_2 srctype=server desttype=drive
library=liblto device=mt1.0.0.2
define path ottawa_sta drlto_1 srctype=server desttype=drive library=liblto
device=mt0.0.0.2
```

```
define path ottawa_sta drlto_2 srctype=server desttype=drive library=liblto
device=mt1.0.0.2
define path cl_vcs02_sta drlto_1 srctype=server desttype=drive
library=liblto device=mt0.0.0.2
define path cl_vcs02_sta drlto_2 srctype=server desttype=drive
library=liblto device=mt1.0.0.2
```

6. Defintion of the storage pool for LAN-free backup:

```
define stgpool spt_bck lto pooltype=PRIMARY maxscratch=4
```

7. Definition/update of the policies to point to the storage pool above and activation of the policy set to refresh the changes. In our case we update the backup copygroup in the standard domain:

```
update copygroup standard standard standard type=backup dest=spt_bck
validate policyset standard standard
activate policyset standard standard
```

24.5.2 Configuration of the Storage Agent for local nodes

As mentioned before, we set up three Storage Agents: one local for each node (SALVADOR_STA and OTTAWA_STA) and one for the TSM Group of the cluster (CL_VCS02_STA).

The configuration process differs whether it is local or cluster. Here we describe the tasks we follow to configure the Storage Agent for local nodes.

Updating dsmsta.opt

Before we start configuring the Storage Agent, we need to edit the dsmsta.opt file located in c:\program files\tivoli\tsm\storageagent.

We change the following line, to make sure it points to the whole path where the device configuration file is located:

```
DEVCONFIG C:\PROGRA~1\TIVOLI\TSM\STORAGEAGENT\DEVCONFIG.TXT
```

Figure 24-2 Modifying devconfig option to point to devconfig file in dsmsta.opt

Note: We need to update the dsmsta.opt because the service used to start the Storage Agent does not use as default path for the devconfig.txt file the installation path. It uses as default the path where the command is run.

Using the management console to initialize the Storage Agent

The following steps describe how to initialize the Storage Agent:

1. We open the Management Console (**Start** →**Programs** → **Tivoli Storage Manager** → **Management Console**) and click **Next** on the welcome menu of the wizard.
2. We provide the Storage Agent information: name, password, and TCP/IP address (high level address) as shown in Figure 24-3.

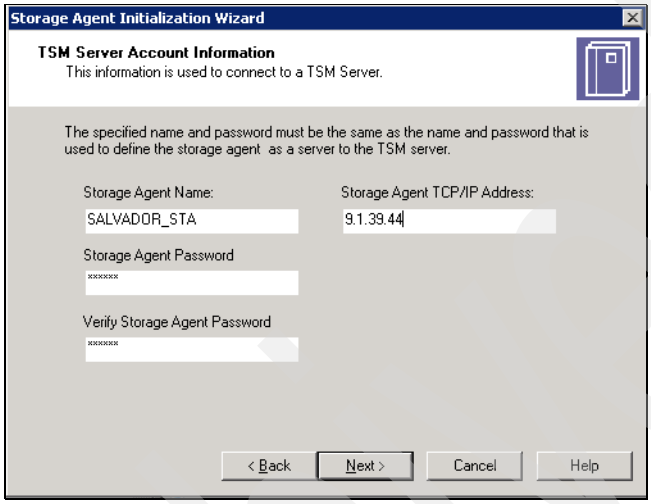

Figure 24-3 Specifying parameters for the Storage Agent

3. We provide all the server information: name, password, TCP/IP, and TCP port, as shown in Figure 24-4, and click **Next**.

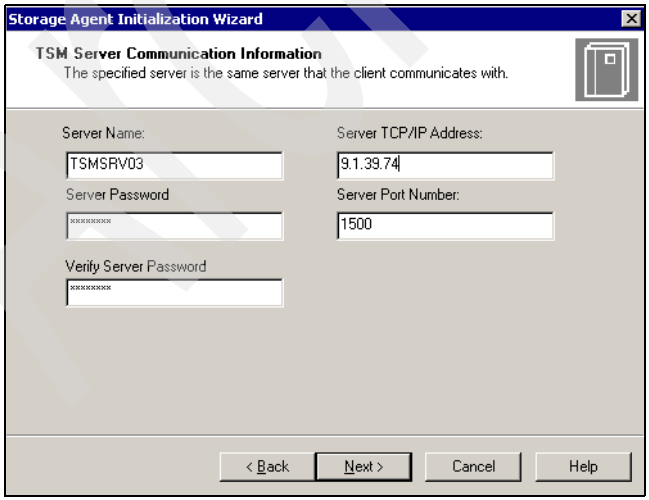

Figure 24-4 Specifying parameters for the Tivoli Storage Manager server

4. In Figure 24-5, we select the account that the service will use to start. We specify the administrator account here, but we could also have created a specific account to be used. This account should be in the administrator's group. We type the password and accept the service to start automatically when the server is started, we then click **Next**.

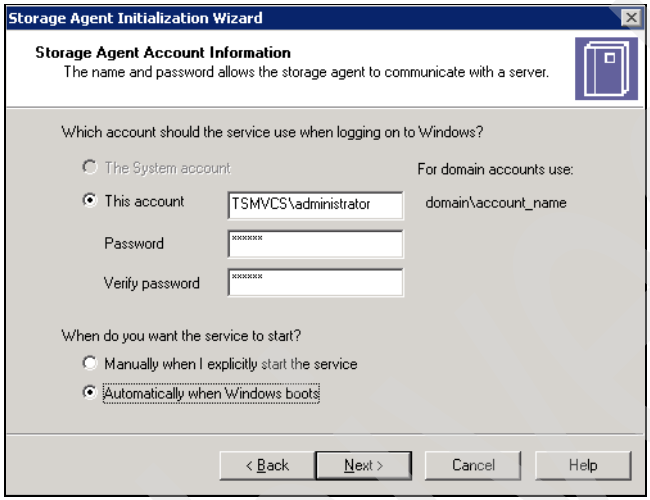

Figure 24-5 Specifying the account information

5. We click **Finish** when the wizard is complete.
6. We click **OK** on the message that says that the user has been granted rights to log on as a service.
7. The wizard finishes, informing you that the Storage Agent has been initialized (Figure 24-6). We click **OK**.

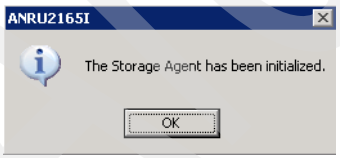

Figure 24-6 Storage agent initialized

8. The Management Console now displays the Tivoli Storage Manager StorageAgent1 service running in Figure 24-7.

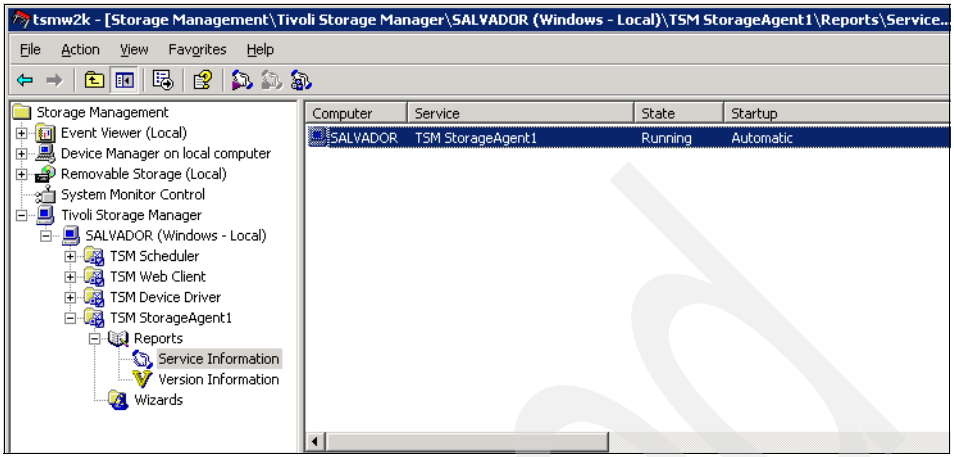

Figure 24-7 StorageAgent1 is started

9. We repeat the same steps in the other server (OTTAWA).

This wizard can be re-run at any time if needed, from the Management Console, under **TSM StorageAgent1** →**Wizards**.

Updating the client option file

To be capable of using LAN-free backup for each local node, we include the following options in the dsm.opt client file.

```
ENABLELANFREE yes
LANFREECOMMMETHOD sharedmem
LANFREESHMPORT 1511
```

We specify the 1511 port for Shared Memory instead of 1510 (the default), because we will use this default port to communicate with the Storage Agent associated to the cluster. Port 1511 will be used by the local nodes when communicating to the local Storage Agents.

Instead of the options specified above, you also can use:

```
ENABLELANFREE yes
LANFREECOMMMETHOD TCPIP
LANFREETCPPORT 1502
```

Restarting the Tivoli Storage Manager scheduler service

To use the LAN-free path, it is necessary, after including the lanfree options in dsm.opt, to restart the Tivoli Storage Manager scheduler service. If we do not restart the service, the new options will not be read by the client.

Archived

24.5.3 Configuration of the Storage Agent for virtual nodes

In order to back up shared disk drives on the cluster using the LAN-free path, we can use the Storage Agent instance created for the local nodes. Depending upon the node that hosts the resources at that time, it will be used for one local Storage Agent or another one.

This is the technically supported way of configuring LAN-free backup for clustered configurations. Each virtual node in the cluster should use the local Storage Agent in the local node that hosts the resource at that time.

However, in order to also have high-availability for the Storage Agent, we configure a new Storage Agent instance that will be used for the cluster.

> **Attention:** This is *not* a technically supported configuration but, in our lab tests, it worked.

In the following sections we describe the process for our TSM Group, where a TSM Scheduler generic service resource is located for backup of the j: shared disk drive.

Using the dsmsta setstorageserver utility

Instead of using the management console to create the new instance, we use the **`dsmsta`** utility from an MS-DOS prompt. The reason to use this tool is because we have to create a new registry key for this Storage Agent. If we use the management console, we would use the default key, StorageAgent1, and we need a different one.

To achieve this goal, we perform these tasks:

1. We begin the configuration in the node that hosts the shared disk drives.
2. We copy the storageagent folder (created at installation time) from c:\program files\tivoli\tsm onto a shared disk drive (j:) with the name “storageagent2”.
3. We open a Windows MS-DOS prompt and change to j:\storageagent2.
4. We change the line devconfig in the dsmsta.opt file to point to j:\storageagent2\devconfig.txt.
5. From this path, we run the command we see in Figure 24-8 to create another instance for a Storage Agent called StorageAgent2. For this instance, the option (dsmsta.opt) and device configuration (devconfig.txt) files will be located on this path.

```
C:\WINDOWS\system32\cmd.exe

J:\storageagent2>dsmsta -k StorageAgent2 setstorageserver myname=CL_VCS02_STA my
pa=itsosj myhla=9.1.39.46 servername=TSMSRV03 serverpassword=password hla=9.1.39
.74 lla=1500
ANR0900I Processing options file J:\storageagent2\dsmsta.opt.
ANR7800I DSMSERV generated at 12:45:40 on Jan 28 2005.

Tivoli Storage Manager for Windows
Version 5, Release 3, Level 0.1

*************************************************************
* Patch 5.3.0.1 TSM Server (Driver 0000, Build 05028.1245) *
*                                                           *
*************************************************************
Licensed Materials - Property of IBM

(C) Copyright IBM Corporation 1990, 2004.
All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

ANR4726I The ICC support module has been loaded.
ANR1432I Updating device configuration information to defined files.
ANR1433I Device configuration information successfully written to j:\storageag-
ent2\devconfig.txt.
ANR2119I The SERVERNAME option has been changed in the options file.
ANR0467I The SETSTORAGESERVER command completed successfully.
```

Figure 24-8 Installing Storage Agent for LAN-free backup of shared disk drives

Attention: Notice in Figure 24-8 the new registry key used for this Storage Agent, StorageAgent2, as well as the name and IP address specified in the “myname” and “myhla” parameters. The Storage Agent name is CL_VCS02_STA, and its IP address is the IP address of the ISC Group. Also notice that when executing the command from j:\storageagent2, we make sure that the dsmsta.opt and devconfig.txt updated files are the ones in this path.

6. Now, from the same path, we run a command to install a service called “TSM StorageAgent2” related to the StorageAgent2 instance created in step 4. The command and the result of its execution are shown in Figure 24-9:

```
C:\WINDOWS\system32\cmd.exe

J:\storageagent2>install "TSM StorageAgent2" j:\storageagent2\dstasvc.exe admini
strator itsosj
Service installed
```

Figure 24-9 Installing the service attached to StorageAgent2

7. If we open the Tivoli Storage Manager management console in this node, we now can see two instances for two Storage Agents: the one we created for the local node, TSM StorageAgent1; and a new one, TSM Storage Agent2.This last instance is stopped, as we can see in Figure 24-10.

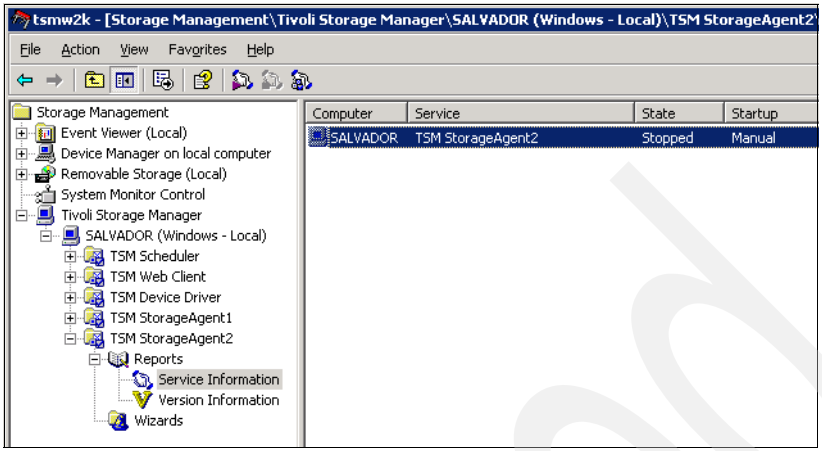

Figure 24-10 Management console displays two Storage Agents

8. We start the TSM StorageAgent2 instance by right-clicking and selecting **Start** as shown in Figure 24-11.

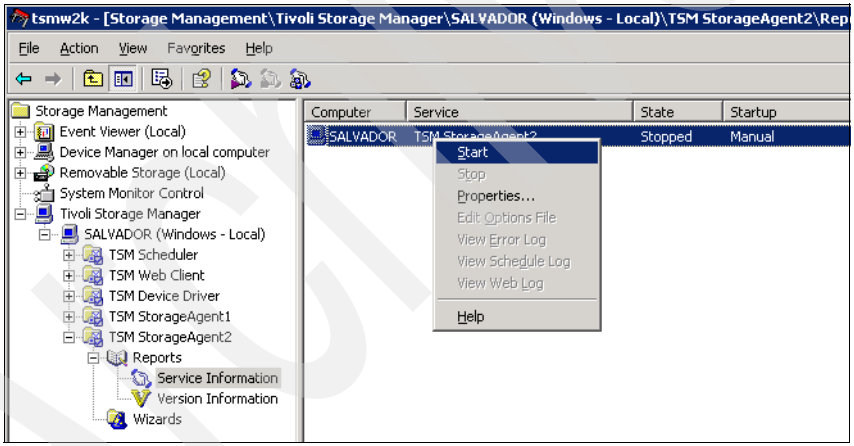

Figure 24-11 Starting the TSM StorageAgent2 service in SALVADOR

9. Now we have two Storage Agent instances running in SALVADOR:
   - TSM StorageAgent1: Related to the local node and using the dsmsta.opt and devconfig.txt files located in c:\program files\tivoli\tsm\storageagent.
   - TSM StorageAgent2: Related to the virtual node and using the dsmsta.opt and devconfig.txt files located in j:\storageagent2.
10. We stop the TSM StorageAgent2 and move the resources to OTTAWA.

Archived

11. In OTTAWA, we follow steps 3 to 6. After that, we open the Tivoli Storage Manager management console and we again find two Storage Agent instances: TSM StorageAgent1 (for the local node) and TSM StorageAgent2 (for the virtual node). This last instance is stopped and set to manual.
12. We start the instance by right-clicking and selecting **Start**. After a successful start, we stop it again.

Creating a resource in VCS service group

Finally, the last task consists of the definition of TSM StorageAgent2 service as a cluster resource, and make it go online before the TSM Scheduler for drive J.

1. Using the Application Configuration Wizard, we create a resource for the service TSM StorageAgent2 as shown in Figure 24-12.

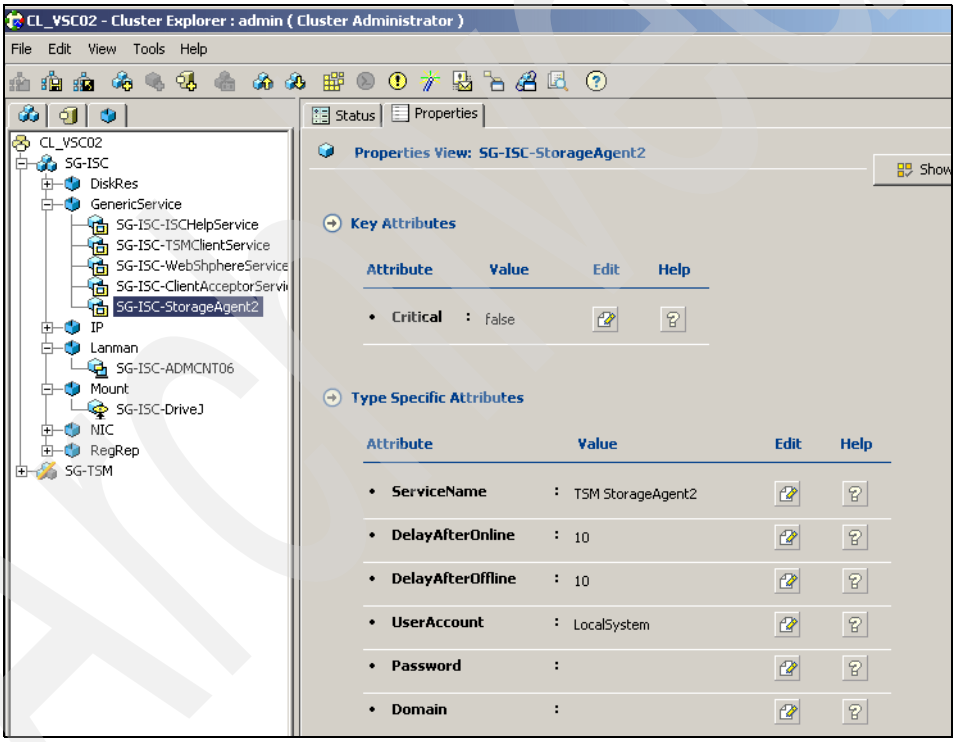

Figure 24-12 Creating StorageAgent2 resource

Important: The name of the service in Figure 24-12 must match the name we used to install the instance in both nodes.

2. We link the StorageAgent2 service in such a way that it comes online before the Tivoli Storage Manager Client Scheduler, as shown in Figure 24-13.

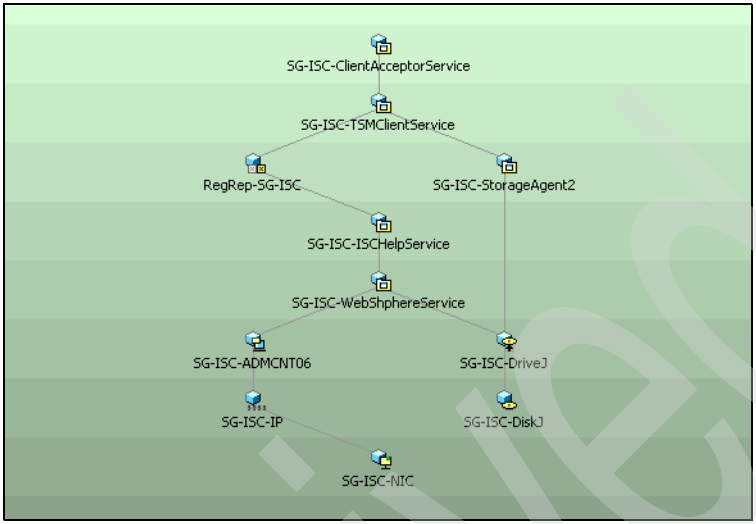

Figure 24-13 StorageAgent2 must come online before the Scheduler

3. We move the cluster to the other node to test that all resources go online.

Updating the client option file

To be capable of using LAN-free backup for the virtual node, we must specify certain special options in the client option file for the virtual node.

We open g:\tsm\dsm.opt and we include the following options:

```
ENABLELANFREE yes
LANFREECOMMMETHOD SHAREDMEM
LANFREESHMPORT 1510
```

For the virtual node, we use the default shared memory port, 1510.

Instead of the options above, you also can use:

```
ENABLELANFREE yes
LANFREECOMMMETHOD TCPIP
LANFREETCPPORT 1500
```

Restarting the Tivoli Storage Manager scheduler service

After including the LAN-free options in dsm.opt, we restart the Tivoli Storage Manager scheduler service for the Tivoli Storage Manager Group using the Cluster Explorer. If we do not restart the service, the new options will not be read by the client.

24.6 Testing Storage Agent high availability

The purpose of this section is to test our LAN-free setup for the clustering.

We use the SG_ISC Service Group (nodename CL_VCS02_ISC) to test LAN-free backup/restore of shared data in our Windows VCS cluster.

Our objective with these tasks is to know how the Storage Agent and the Tivoli Storage Manager Library Manager work together to respond, on a LAN-free client clustered environment, after certain kinds of failures that affect the shared resources.

Again, for details of our LAN-free configuration, refer back to Table 24-1 on page 1003, Table 24-2 on page 1004, and Table 24-3 on page 1004.

24.6.1 Testing LAN-free client incremental backup

First we test a scheduled client incremental backup using the SAN path.

Objective

The objective of this test is to show what happens when a LAN-free client incremental backup is started for a virtual node on the cluster using the Storage Agent created for this group (CL_VCS02_STA), and the node that hosts the resources at that moment suddenly fails.

Activities

To do this test, we perform these tasks:

1. We open the Veritas Cluster Manager console menu to check which node hosts the Tivoli Storage Manager scheduler service for SG_ISC Service Group.
2. We schedule a client incremental backup operation using the Tivoli Storage Manager server scheduler and we associate the schedule to CL_VCS02_ISC nodename.
3. We make sure that TSM StorageAgent2 and TSM Scheduler for CL_VCS02_STA are online resources on this node.

Archived

4. When it is the scheduled time, a client session for CL_VCS02_ISC nodename starts on the server. At the same time, several sessions are also started for CL_VCS02_STA for Tape Library Sharing and the Storage Agent prompts the Tivoli Storage Manager server to mount a tape volume. The volume 030AKK is mounted in drive DRLTO_1, as we can see in Figure 24-14.

```
IBM Tivoli Storage Manager

                        (/dev/rmt0). (SESSION: 255)
02/19/2005 12:51:58  ANR0409I (Session: 249, Origin: CL_VCS02_STA)  Session 13
                      ended for server TSMSRV03 (AIX-RS/6000). (SESSION: 249)
02/19/2005 12:51:58  ANR0409I Session 255 ended for server CL_VCS02_STA
                      (Windows). (SESSION: 255)
02/19/2005 12:51:58  ANR8337I (Session: 249, Origin: CL_VCS02_STA)  LTO volume
                      030AKK mounted in drive DRLTO_1 (mt0.0.0.4). (SESSION:
                      249)
02/19/2005 12:51:58  ANR0511I Session 253 opened output volume 030AKK.
                      (SESSION: 253)
02/19/2005 12:51:58  ANR0511I (Session: 249, Origin: CL_VCS02_STA)  Session 10
                      opened output volume 030AKK. (SESSION: 249)
02/19/2005 12:52:12  ANR2017I (Session: 249, Origin: CL_VCS02_STA)
                      Administrator ADMIN issued command: QUERY SESSION
                      (SESSION: 249)
02/19/2005 12:52:14  ANR2017I (Session: 249, Origin: CL_VCS02_STA)
                      Administrator ADMIN issued command: QUERY MO  (SESSION:
                      249)
02/19/2005 12:52:14  ANR8330I (Session: 249, Origin: CL_VCS02_STA)  LTO volume
                      030AKK is mounted R/W in drive DRLTO_1 (mt0.0.0.4),
                      status: IN USE. (SESSION: 249)
02/19/2005 12:52:14  ANR8334I (Session: 249, Origin: CL_VCS02_STA)          1
                      matches found. (SESSION: 249)
02/19/2005 12:52:22  ANR2017I (Session: 249, Origin: CL_VCS02_STA)
                      Administrator ADMIN issued command: QUERY SESSION
                      (SESSION: 249)
02/19/2005 12:52:27  ANR2017I (Session: 249, Origin: CL_VCS02_STA)
more...   (<ENTER> to continue, 'C' to cancel) _
```

Figure 24-14 Storage Agent CL_VCS02_STA session for Tape Library Sharing

5. The Storage Agent shows sessions started with the client and the Tivoli Storage Manager server TSMSRV03, and the tape volume is mounted. We can see all these events in Figure 24-15.

```
IBM Tivoli Storage Manager

tsm: CL_VCS02_STA>q mo
ANR8330I LTO volume 030AKK is mounted R/W in drive DRLTO_1 (mt0.0.0.4), status:
IN USE.
ANR8334I          1 matches found.

tsm: CL_VCS02_STA>q se

  Sess Comm.  Sess      Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State     Time    Sent   Recvd Type
------ ------ ------ -------- ------- ------ -----  -------- --------------------
     1 Tcp/Ip IdleW    3 S      694   1.5 K Serv-  AIX-RS/- TSMSRV03
                                              er     6000
     3 Tcp/Ip IdleW    0 S    2.8 K   3.0 K Serv-  AIX-RS/- TSMSRV03
                                              er     6000
     4 Tcp/Ip IdleW    0 S    6.3 K     253 Serv-  AIX-RS/- TSMSRV03
                                              er     6000
    10 Shared IdleW    0 S      377 102.8 M Node   WinNT    CL_VCS02_ISC
       Memo-
       ry
    11 Tcp/Ip Run      0 S   10.3 K  10.7 K Serv-  AIX-RS/- TSMSRV03
                                              er     6000
    14 Tcp/Ip Run      0 S   19.1 K     498 Admin  WinNT    ADMIN

tsm: CL_VCS02_STA>_
```

Figure 24-15 A tape volume is mounted and Storage Agent starts sending data

6. The client, by means of the Storage Agent, starts sending files to the drive using the SAN path as we see on its schedule log file in Figure 24-16.

```
dsmsched - Notepad
File  Edit  Format  Help
Executing scheduled command now.
02/19/2005 12:59:01 --- SCHEDULEREC OBJECT BEGIN INCR_BACKUP 02/19/2005 12:50:47
02/19/2005 12:59:01 Incremental backup of volume 'j:\code\*'
02/19/2005 12:59:13 Directory-->                  0 \\cl_vcs02\j$\code [Sent]
02/19/2005 12:59:13 Directory-->                  0 \\cl_vcs02\j$\code\adminc [Sent]
02/19/2005 12:59:13 Directory-->                  0 \\cl_vcs02\j$\code\isc [Sent]
02/19/2005 12:59:13 Normal File-->              263 \\cl_vcs02\j$\code\GSKit-README.txt [Sent]
02/19/2005 12:59:13 Normal File-->        1,205,657 \\cl_vcs02\j$\code\GSKit.zip [Sent]
02/19/2005 12:59:13 ANS1114I Waiting for mount of offline media.
02/19/2005 12:59:14 Retry # 1  Directory-->                  0 \\cl_vcs02\j$\code [Sent]
02/19/2005 12:59:14 Retry # 1  Directory-->                  0 \\cl_vcs02\j$\code\adminc [Sent]
02/19/2005 12:59:14 Retry # 1  Directory-->                  0 \\cl_vcs02\j$\code\isc [Sent]
02/19/2005 13:00:50 Retry # 1  Normal File-->              263 \\cl_vcs02\j$\code\GSKit-README.txt [
02/19/2005 13:00:50 Retry # 1  Normal File-->        1,205,657 \\cl_vcs02\j$\code\GSKit.zip [Sent]
02/19/2005 13:00:59 Normal File-->       25,388,669 \\cl_vcs02\j$\code\TSM530C_WIN32.exe [Sent]
02/19/2005 13:01:00 Normal File-->       10,762,389 \\cl_vcs02\j$\code\TSMCON5300_WIN.exe [Sent]
02/19/2005 13:01:03 Normal File-->       12,426,660 \\cl_vcs02\j$\code\TSMLNG5300_CHS_WIN_337.exe [S
02/19/2005 13:01:10 Normal File-->       12,444,265 \\cl_vcs02\j$\code\TSMLNG5300_CHT_WIN_337.exe [S
```

Figure 24-16 Client starts sending files to the server in the schedule log file

7. While the client continues sending files to the server, we force a failure in the node that hosts the resources. The following sequence takes place:

   a. The client and also the Storage Agent lose their connections with the server temporarily, and both sessions are terminated, as we can see on the Tivoli Storage Manager server activity log shown in Figure 24-17.

```
IBM Tivoli Storage Manager
                         shutdown. (SESSION: 249)
02/19/2005 12:54:28  ANR7835I (Session: 249, Origin: CL_VCS02_STA)  The server
                      thread 9 (tid 2852) terminated in response to server
                      shutdown. (SESSION: 249)
02/19/2005 12:54:28  ANR7835I (Session: 249, Origin: CL_VCS02_STA)  The server
                      thread 10 (tid 3512) terminated in response to server
                      shutdown. (SESSION: 249)
02/19/2005 12:54:28  ANR7835I (Session: 249, Origin: CL_VCS02_STA)  The server
                      thread 11 (tid 3520) terminated in response to server
                      shutdown. (SESSION: 249)
02/19/2005 12:54:28  ANR7835I (Session: 249, Origin: CL_VCS02_STA)  The server
                      thread 12 (tid 3596) terminated in response to server
                      shutdown. (SESSION: 249)
02/19/2005 12:54:28  ANR7835I (Session: 249, Origin: CL_VCS02_STA)  The server
                      thread 13 (tid 3480) terminated in response to server
                      shutdown. (SESSION: 249)
02/19/2005 12:54:28  ANR7835I (Session: 249, Origin: CL_VCS02_STA)  The server
more...   <<ENTER> to continue, 'C' to cancel>

                      thread 14 (tid 3424) terminated in response to server
                      shutdown. (SESSION: 249)
02/19/2005 12:54:28  ANR7835I (Session: 249, Origin: CL_VCS02_STA)  The server
                      thread 16 (tid 3404) terminated in response to server
                      shutdown. (SESSION: 249)
02/19/2005 12:54:28  ANR0991I (Session: 249, Origin: CL_VCS02_STA)  Storage
                      agent shutdown complete. (SESSION: 249)
02/19/2005 12:54:28  ANR0480W Session 253 for node CL_VCS02_ISC (WinNT)
                      terminated - connection with client severed. (SESSION:
                      253)
```

Figure 24-17 Sessions for Client and Storage Agent are lost in the activity log

   b. In the Veritas Cluster Manager console, the second node tries to bring the resources online after the failure on the first node.

c. The schedule log file in the client receives an error message (Figure 24-18).

```
dsmsched - Notepad
File  Edit  Format  Help
02/19/2005 13:01:58 Normal File-->        29,065,149 \\cl_vcs02\j$\code\TSMSTA5301_WIN.exe [Sent]
02/19/2005 13:01:58 Normal File-->         7,515,003 \\cl_vcs02\j$\code\adminc\AdminCenter.war [Sent]
02/19/2005 13:01:58 Directory-->                   0 \\cl_vcs02\j$\code\isc\RuntimeExt [Sent]
02/19/2005 13:01:58 Directory-->                   0 \\cl_vcs02\j$\code\isc\tutorial [Sent]
02/19/2005 13:01:58 Directory-->                   0 \\cl_vcs02\j$\code\isc\wps [Sent]
02/19/2005 13:02:11 ANS2820E An interrupt has occurred. The current operation will end and the
client will shut down.
```

Figure 24-18 Backup is interrupted in the client

d. The tape volume is still mounted on the same drive.

e. After a short period of time, the resources are online.

f. When the Storage Agent CL_VCS02_STA and the scheduler are again online, the tape volume is dismounted by the Tivoli Storage Manager server from the drive and it is mounted in the second drive for use of the Storage Agent, such as we show in Figure 24-19.

```
IBM Tivoli Storage Manager
02/19/2005 12:59:50  ANR2017I (Session: 270, Origin: CL_VCS02_STA)
                      Administrator ADMIN issued command: QUERY SESSION
more...   (<ENTER> to continue, 'C' to cancel)

                      (SESSION: 270)
02/19/2005 12:59:52  ANR2017I (Session: 270, Origin: CL_VCS02_STA)
                      Administrator ADMIN issued command: QUERY MO  (SESSION:
                      270)
02/19/2005 12:59:52  ANR8379I (Session: 270, Origin: CL_VCS02_STA)  Mount point
                      in device class LTO is waiting for the volume mount to
                      complete, status: WAITING FOR VOLUME. (SESSION: 270)
02/19/2005 12:59:52  ANR8334I (Session: 270, Origin: CL_VCS02_STA)          1
                      matches found. (SESSION: 270)
02/19/2005 13:00:08  ANR8337I LTO volume 030AKK mounted in drive DRLTO_2
                      (/dev/rmt1). (SESSION: 278)
02/19/2005 13:00:08  ANR0409I (Session: 270, Origin: CL_VCS02_STA)  Session 16
                      ended for server TSMSRV03 (AIX-RS/6000). (SESSION: 270)
02/19/2005 13:00:08  ANR0409I Session 278 ended for server CL_VCS02_STA
                      (Windows). (SESSION: 278)
02/19/2005 13:00:08  ANR8337I (Session: 270, Origin: CL_VCS02_STA)  LTO volume
                      030AKK mounted in drive DRLTO_2 (mt1.0.0.4). (SESSION:
                      270)
02/19/2005 13:00:09  ANR0511I Session 276 opened output volume 030AKK.
                      (SESSION: 276)
02/19/2005 13:00:09  ANR0511I (Session: 270, Origin: CL_VCS02_STA)  Session 13
                      opened output volume 030AKK. (SESSION: 270)
02/19/2005 13:00:21  ANR2017I (Session: 270, Origin: CL_VCS02_STA)
                      Administrator ADMIN issued command: QUERY MO  (SESSION:
more...   (<ENTER> to continue, 'C' to cancel) _
```

Figure 24-19 Tivoli Storage Manager server mounts tape volume in second drive

g. Finally, the client restarts its scheduled incremental backup if the startup window for the schedule has not elapsed, using the SAN path as we can see in its schedule log file in Figure 24-20.

```
dsmsched - Notepad
File  Edit  Format  Help
Executing scheduled command now.
02/19/2005 13:05:58 --- SCHEDULEREC OBJECT BEGIN INCR_BACKUP 02/19/2005 12:50:47
02/19/2005 13:05:58 Incremental backup of volume 'j:\code\*'
02/19/2005 13:06:19 Normal File-->       327,709,515 \\cl_vcs02\j$\code\isc\C8241ML.exe  ** Unsuccess
02/19/2005 13:06:19 ANS1114I Waiting for mount of offline media.
02/19/2005 13:13:03 Retry # 1  Normal File-->       327,709,515 \\cl_vcs02\j$\code\isc\C8241ML.exe [S
02/19/2005 13:13:03 Normal File-->            20,763 \\cl_vcs02\j$\code\isc\dsminstall.jar [Sent]
02/19/2005 13:13:03 Normal File-->         6,484,490 \\cl_vcs02\j$\code\isc\ISCAction.jar [Sent]
02/19/2005 13:13:03 Normal File-->         6,621,305 \\cl_vcs02\j$\code\isc\ISCCheck.jar [Sent]
02/19/2005 13:13:10 Normal File-->        44,275,581 \\cl_vcs02\j$\code\isc\ISCRuntime.jar [Sent]
02/19/2005 13:13:11 Normal File-->            81,360 \\cl_vcs02\j$\code\isc\LICENSE.TXT [Sent]
02/19/2005 13:13:11 Normal File-->                12 \\cl_vcs02\j$\code\isc\media.inf [Sent]
02/19/2005 13:13:11 Normal File-->            14,001 \\cl_vcs02\j$\code\isc\README [Sent]
02/19/2005 13:13:11 Normal File-->            23,984 \\cl_vcs02\j$\code\isc\README.INSTALL [Sent]
02/19/2005 13:13:17 Normal File-->        40,229,284 \\cl_vcs02\j$\code\isc\setupISC.exe [Sent]
02/19/2005 13:13:18 Directory-->                   0 \\cl_vcs02\j$\code\isc\RuntimeExt\eclipse [Sent]
```

Figure 24-20 The scheduled is restarted and the tape volume mounted again

8. The incremental backup ends successfully as we can see on the final statistics recorded by the client in its schedule log file in Figure 24-21.

```
dsmsched - Notepad
File  Edit  Format  Help
02/19/2005 13:14:44 Normal File-->               286 \\cl_vcs02\j$\code\isc\wps\migration\templates\x
02/19/2005 13:14:44 Normal File-->               175 \\cl_vcs02\j$\code\isc\wps\migration\templates\x
02/19/2005 13:14:44 Normal File-->           110,861 \\cl_vcs02\j$\code\isc\wps\migration\templates\x
02/19/2005 13:14:44 Normal File-->             6,869 \\cl_vcs02\j$\code\isc\wps\migration\templates\x
02/19/2005 13:14:48 Successful incremental backup of '\\cl_vcs02\j$\code\*'

02/19/2005 13:14:48 --- SCHEDULEREC STATUS BEGIN
02/19/2005 13:14:48 Total number of objects inspected:    1,822
02/19/2005 13:14:48 Total number of objects backed up:    1,799
02/19/2005 13:14:48 Total number of objects updated:          0
02/19/2005 13:14:48 Total number of objects rebound:          0
02/19/2005 13:14:48 Total number of objects deleted:          0
02/19/2005 13:14:48 Total number of objects expired:          0
02/19/2005 13:14:48 Total number of objects failed:           0
02/19/2005 13:14:48 Total number of bytes transferred:    963.56 MB
02/19/2005 13:14:48 LanFree data bytes:                   963.43 MB
02/19/2005 13:14:48 Data transfer time:                   400.95 sec
02/19/2005 13:14:48 Network data transfer rate:         2,460.84 KB/sec
02/19/2005 13:14:48 Aggregate data transfer rate:       1,862.60 KB/sec
02/19/2005 13:14:48 Objects compressed by:                    0%
02/19/2005 13:14:48 Elapsed processing time:            00:08:49
02/19/2005 13:14:48 --- SCHEDULEREC STATUS END
02/19/2005 13:14:48 --- SCHEDULEREC OBJECT END INCR_BACKUP 02/19/2005 12:50:47
02/19/2005 13:14:48 Scheduled event 'INCR_BACKUP' completed successfully.
02/19/2005 13:14:48 Sending results for scheduled event 'INCR_BACKUP'.
02/19/2005 13:14:49 Results sent to server for scheduled event 'INCR_BACKUP'.
```

Figure 24-21 Backup ends successfully

Results summary

The test results show that, after a failure on the node that hosts both the Tivoli Storage Manager scheduler as well as the Storage Agent shared resources, a scheduled incremental backup started on one node for LAN-free is restarted and successfully completed on the other node, also using the SAN path.

This is true if the startup window used to define the schedule is not elapsed when the scheduler service restarts on the second node.

The Tivoli Storage Manager server on AIX resets the SCSI bus when the Storage Agent is restarted on the second node. This permits us to dismount the tape volume from the drive where it was mounted before the failure. When the client restarts the LAN-free operation, the same Storage Agent commands the server to mount again the tape volume to continue the backup.

Restriction: This configuration, with two Storage Agents started on the same node (one local and another for the cluster) is not technically supported by Tivoli Storage Manager for SAN. However, in our lab environment, it worked.

Note: In other tests we made using the local Storage Agent on each node for communication to the virtual client for LAN-free, the SCSI bus reset did not work. The reason is that the Tivoli Storage Manager server on AIX, when it acts as a Library Manager, can handle the SCSI bus reset only when the Storage Agent name is the same for the failing and recovering Storage Agent.

In other words, if we use local Storage Agents for LAN-free backup of the virtual client (CL_VCS02_ISC), the following conditions must be taken into account: The failure of the node SALVADOR means that all local services will also fail, including SALVADOR_STA (the local Storage Agent). VCS will cause a failover to the second node where the local Storage Agent will be started again, but with a different name (OTTAWA_STA). It is this discrepancy in naming which will cause the LAN-free backup to fail, as clearly, the virtual client will be unable to connect to SALVADOR_STA.

Tivoli Storage Manager server does not know what happened to the first Storage Agent because it does not receive any alert from it, so that the tape drive is in a RESERVED status until the default timeout (10 minutes) elapses. If the scheduler for CL_VCS02_ISC starts a new session before the ten-minute timeout elapses, it tries to communicate to the local Storage Agent of this second node, OTTAWA_STA, and this prompts the Tivoli Storage Manager server to mount the same tape volume.

Since this tape volume is still mounted on the first drive by SALVADOR_STA (even when the node failed) and the drive is RESERVED, the only option for the Tivoli Storage Manager server is to mount a new tape volume in the second drive. If either there are not enough tape volumes in the tape storage pool, or the second drive is busy at that time with another operation, or if the client node has its maximum mount points limited to 1, the backup is cancelled.

24.6.2 Testing client restore

Our second test is a scheduled restore using the SAN path.

Objective

The objective of this test is to show what happens when a LAN-free restore is started for a virtual node on the cluster, and the node that hosts the resources at that moment suddenly fails.

Activities

To do this test, we perform these tasks:

1. We open the Verirtas Cluster Manager console to check which node hosts the Tivoli Storage Manager scheduler resource.
2. We schedule a client restore operation using the Tivoli Storage Manager server scheduler and associate the schedule to CL_VCS02_ISC nodename.
3. We make sure that TSM StorageAgent2 and TSM Scheduler for CL_VCS02_ISC are online resources on this node.
4. When it is the scheduled time, a client session for CL_VCS02_ISC nodename starts on the server. At the same time several sessions are also started for CL_VCS02_STA for Tape Library Sharing and the Storage Agent prompts the Tivoli Storage Manager server to mount a tape volume. The tape volume is mounted in drive DRLTO_1. All of these events are shown in Figure 24-22.

```
IBM Tivoli Storage Manager
                         started for node CL_VCS02_ISC (WinNT) (Shared Memory).
                         (SESSION: 270)
02/19/2005 13:14:10  ANR0408I Session 300 started for server CL_VCS02_STA
                      (Windows) (Tcp/Ip) for storage agent.  (SESSION: 300)
02/19/2005 13:14:10  ANR0408I (Session: 270, Origin: CL_VCS02_STA)  Session 55
                      started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      storage agent.  (SESSION: 270)
02/19/2005 13:14:10  ANR0415I Session 300 proxied by CL_VCS02_STA started for
                      node CL_VCS02_ISC. (SESSION: 300)
02/19/2005 13:14:11  ANR0408I Session 301 started for server CL_VCS02_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 301)
more...   (<ENTER> to continue, 'C' to cancel)

02/19/2005 13:14:11  ANR0408I (Session: 270, Origin: CL_VCS02_STA)  Session 56
                      started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 270)
02/19/2005 13:14:11  ANR0409I (Session: 270, Origin: CL_VCS02_STA)  Session 56
                      ended for server TSMSRV03 (AIX-RS/6000). (SESSION: 270)
02/19/2005 13:14:11  ANR0409I Session 301 ended for server CL_VCS02_STA
                      (Windows). (SESSION: 301)
02/19/2005 13:14:11  ANR0408I Session 302 started for server CL_VCS02_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 302)
02/19/2005 13:14:11  ANR0408I (Session: 270, Origin: CL_VCS02_STA)  Session 57
                      started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
more...   (<ENTER> to continue, 'C' to cancel)

                      library sharing.  (SESSION: 270)
02/19/2005 13:14:41  ANR8337I LTO volume 030AKK mounted in drive DRLTO_1
                      (/dev/rmt0). (SESSION: 302)
02/19/2005 13:14:41  ANR0409I (Session: 270, Origin: CL_VCS02_STA)  Session 57
                      ended for server TSMSRV03 (AIX-RS/6000). (SESSION: 270)
02/19/2005 13:14:41  ANR0409I Session 302 ended for server CL_VCS02_STA
                      (Windows). (SESSION: 302)
02/19/2005 13:14:41  ANR8337I (Session: 270, Origin: CL_VCS02_STA)  LTO volume
                      030AKK mounted in drive DRLTO_1 (mt0.0.0.4). (SESSION:
```

Figure 24-22 Starting restore session for LAN-free

5. The client starts restoring files as we can see on the schedule log file in Figure 24-23.

```
dsmsched - Notepad
File  Edit  Format  Help
Executing scheduled command now.
02/19/2005 13:21:52 --- SCHEDULEREC OBJECT BEGIN RESTORE 02/19/2005 13:13:46
02/19/2005 13:21:52 Restore function invoked.

02/19/2005 13:21:53 ANS1247I Waiting for files from the server... ** Interrupted **
02/19/2005 13:21:56 ANS1114I Waiting for mount of offline media.
02/19/2005 13:23:35 Restoring              0 \\cl_vcs02\j$\code\adminc [Done]
02/19/2005 13:23:35 Restoring              0 \\cl_vcs02\j$\code\isc [Done]
02/19/2005 13:23:35 Restoring            263 \\cl_vcs02\j$\code\GSKit-README.txt [Done]
02/19/2005 13:23:35 Restoring      1,205,657 \\cl_vcs02\j$\code\GSKit.zip [Done]
02/19/2005 13:23:36 Restoring     25,388,669 \\cl_vcs02\j$\code\TSM530C_WIN32.exe [Done]
02/19/2005 13:23:36 Restoring     10,762,389 \\cl_vcs02\j$\code\TSMCON5300_WIN.exe [Done]
02/19/2005 13:23:37 Restoring     12,426,660 \\cl_vcs02\j$\code\TSMLNG5300_CHS_WIN_337.exe
02/19/2005 13:23:37 Restoring     12,444,265 \\cl_vcs02\j$\code\TSMLNG5300_CHT_WIN_337.exe
02/19/2005 13:23:40 Restoring     12,865,309 \\cl_vcs02\j$\code\TSMLNG5300_DEU_WIN_337.exe
```

Figure 24-23 Restore starts on the schedule log file

6. While the client is restoring the files, we force a failure in the node that hosts the resources. The following sequence takes place:

   a. The client CL_VCS02_ISC and the Storage Agent CL_VCS02_STA lose both temporarily their connections with the server, as shown in Figure 24-24.

```
IBM Tivoli Storage Manager
                         thread 10 (tid 3320) terminated in response to server
                         shutdown. (SESSION: 270)
02/19/2005 13:16:28  ANR7835I (Session: 270, Origin: CL_VCS02_STA)  The server
                         thread 11 (tid 3324) terminated in response to server
                         shutdown. (SESSION: 270)
02/19/2005 13:16:28  ANR7835I (Session: 270, Origin: CL_VCS02_STA)  The server
                         thread 12 (tid 3028) terminated in response to server
                         shutdown. (SESSION: 270)
02/19/2005 13:16:28  ANR7835I (Session: 270, Origin: CL_VCS02_STA)  The server
                         thread 13 (tid 3336) terminated in response to server
                         shutdown. (SESSION: 270)
02/19/2005 13:16:28  ANR7835I (Session: 270, Origin: CL_VCS02_STA)  The server
                         thread 15 (tid 1824) terminated in response to server
                         shutdown. (SESSION: 270)
02/19/2005 13:16:28  ANR7835I (Session: 270, Origin: CL_VCS02_STA)  The server
                         thread 16 (tid 3552) terminated in response to server
                         shutdown. (SESSION: 270)
02/19/2005 13:16:28  ANR7835I (Session: 270, Origin: CL_VCS02_STA)  The server
                         thread 17 (tid 3540) terminated in response to server
                         shutdown. (SESSION: 270)
02/19/2005 13:16:28  ANR0991I (Session: 270, Origin: CL_VCS02_STA)  Storage
                         agent shutdown complete. (SESSION: 270)
more...   (<ENTER> to continue, 'C' to cancel)

02/19/2005 13:16:29  ANR3605E Unable to communicate with storage agent.
                         (SESSION: 272)
02/19/2005 13:16:29  ANR0480W Session 300 for node CL_VCS02_ISC (WinNT)
                         terminated - connection with client severed. (SESSION:
                         300)
02/19/2005 13:16:30  ANR0480W Session 299 for node CL_VCS02_ISC (WinNT)
                         terminated - connection with client severed. (SESSION:
                         299)
```

Figure 24-24 Both sessions for Storage Agent and client are lost in the server

   b. The tape volume is still mounted on the same drive.

   c. After a short period of time the resources are online on the other node of the VCS.

d. When the Storage Agent CL_VCS02_STA is again online, as well as the TSM Scheduler service, the Tivoli Storage Manager server resets the SCSI bus and dismounts the tape volume as we can see on the activity log in Figure 24-25.

```
IBM Tivoli Storage Manager
02/19/2005 13:21:02  ANR8331I LTO volume 030AKK is mounted R/W in drive DRLTO_1
                      (/dev/rmt0), status: DISMOUNTING. (SESSION: 304)
02/19/2005 13:21:02  ANR8334I          1 matches found. (SESSION: 304)
02/19/2005 13:21:02  ANR2997W The server log is 88 percent full. The server
                      will delay transactions by 3 milliseconds.
02/19/2005 13:21:04  ANR2017I Administrator ADMIN issued command: QUERY SESSION
                       (SESSION: 304)
02/19/2005 13:21:06  ANR2017I Administrator ADMIN issued command: QUERY RESTORE
                       (SESSION: 304)
02/19/2005 13:21:36  ANR8779E Unable to open drive /dev/rmt0, error number=16.
                      (SESSION: 306)
02/19/2005 13:21:53  ANR0408I (Session: 307, Origin: CL_VCS02_STA)  Session 13
                      started for server TSMSRV03 (AIX-RS/6000) (Tcp/Ip) for
                      library sharing.  (SESSION: 307)
02/19/2005 13:21:53  ANR0408I Session 317 started for server CL_VCS02_STA
                      (Windows) (Tcp/Ip) for library sharing.  (SESSION: 317)
02/19/2005 13:21:53  ANR0409I Session 317 ended for server CL_VCS02_STA
                      (Windows). (SESSION: 317)
02/19/2005 13:21:53  ANR0409I (Session: 307, Origin: CL_VCS02_STA)  Session 13
                      ended for server TSMSRV03 (AIX-RS/6000). (SESSION: 307)
more...   (<ENTER> to continue, 'C' to cancel)

02/19/2005 13:21:55  ANR8336I Verifying label of LTO volume 030AKK in drive
                      DRLTO_1 (/dev/rmt0). (SESSION: 302)
02/19/2005 13:21:57  ANR0406I Session 318 started for node CL_VCS02_ISC (WinNT)
                      (Tcp/Ip 9.1.39.187(1402)). (SESSION: 318)
```

Figure 24-25 The tape volume is dismounted by the server

e. The client (if the startup window for the schedule is not elapsed) re-establishes the session with the Tivoli Storage Manager server and the Storage Agent for LAN-free restore. The Storage Agent prompts the server to mount the tape volume as we can see in Figure 24-26.

```
IBM Tivoli Storage Manager
tsm: CL_VCS02_STA>q se

  Sess Comm.  Sess      Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State     Time    Sent   Recvd Type
------ ------ ------   ------ -------- ------ -----  -------- ----------------
     1 Tcp/Ip IdleW     0 S    3.5 K   4.7 K Serv-  AIX-RS/- TSMSRV03
                                               er     6000
     3 Tcp/Ip IdleW     0 S    4.6 K     253 Serv-  AIX-RS/- TSMSRV03
                                               er     6000
     5 Tcp/Ip Run       0 S    5.4 K  38.3 K Serv-  AIX-RS/- TSMSRV03
                                               er     6000
    10 Tcp/Ip Run       0 S    4.6 K     247 Admin  WinNT    ADMIN
    16 Shared MediaW   50 S      402     424 Node   WinNT    CL_VCS02_ISC
       Memo-
       ry
    17 Tcp/Ip Run       0 S      897   1.7 K Serv-  AIX-RS/- TSMSRV03
                                               er     6000
    19 Tcp/Ip IdleW    49 S      462     264 Serv-  AIX-RS/- TSMSRV03
                                               er     6000

tsm: CL_VCS02_STA>q mo
ANR8379I Mount point in device class LTO is waiting for the volume mount to
complete, status: WAITING FOR VOLUME.
ANR8334I          1 matches found.

tsm: CL_VCS02_STA>_
```

Figure 24-26 The Storage Agent waiting for tape volume to be mounted by server

7. In Figure 24-27, the event log shows the schedule as restarted.

```
IBM Tivoli Storage Manager

tsm: TSMSRV03>q ev * *

Scheduled Start      Actual Start         Schedule Name  Node Name      Status
-------------------- -------------------- -------------- -------------- ---------
02/19/2005 07:41:08                       TEST_SCHED     CL_HACMP03_C-  Missed
                                                          LIENT
02/19/2005 12:50:47  02/19/2005 12:51:17  INCR_BACKUP    CL_VCS02_ISC   Completed
02/19/2005 13:13:46  02/19/2005 13:14:08  RESTORE        CL_VCS02_ISC   Restarted

tsm: TSMSRV03>_
```

Figure 24-27 Event log shows the restore as restarted

8. The client starts the restore of the files from the beginning, as we see in its schedule log file in Figure 24-28.

```
dsmsched - Notepad
File  Edit  Format  Help
Executing scheduled command now.
02/19/2005 13:29:42 --- SCHEDULEREC OBJECT BEGIN RESTORE 02/19/2005 13:13:46
02/19/2005 13:29:42 Restore function invoked.

02/19/2005 13:29:43 ANS1247I Waiting for files from the server... ** Interrupted **
02/19/2005 13:29:46 ANS1114I Waiting for mount of offline media.
02/19/2005 13:31:55 Restoring              0 \\cl_vcs02\j$\code\adminc [Done]
02/19/2005 13:31:55 Restoring              0 \\cl_vcs02\j$\code\isc [Done]
02/19/2005 13:31:55 Restoring            263 \\cl_vcs02\j$\code\GSKit-README.txt [Done]
02/19/2005 13:31:55 Restoring      1,205,657 \\cl_vcs02\j$\code\GSKit.zip [Done]
02/19/2005 13:31:55 Restoring     25,388,669 \\cl_vcs02\j$\code\TSM530C_WIN32.exe [Done]
02/19/2005 13:31:56 Restoring     10,762,389 \\cl_vcs02\j$\code\TSMCON5300_WIN.exe [Done]
02/19/2005 13:31:57 Restoring     12,426,660 \\cl_vcs02\j$\code\TSMLNG5300_CHS_WIN_337.exe
02/19/2005 13:31:59 Restoring     12,444,265 \\cl_vcs02\j$\code\TSMLNG5300_CHT_WIN_337.exe
02/19/2005 13:32:00 Restoring     12,865,309 \\cl_vcs02\j$\code\TSMLNG5300_DEU_WIN_337.exe
02/19/2005 13:32:00 Restoring     12,826,564 \\cl_vcs02\j$\code\TSMLNG5300_ESP_WIN_337.exe
02/19/2005 13:32:01 Restoring     12,881,222 \\cl_vcs02\j$\code\TSMLNG5300_FRA_WIN_337.exe
02/19/2005 13:32:01 Restoring     12,770,448 \\cl_vcs02\j$\code\TSMLNG5300_ITA_WIN_337.exe
02/19/2005 13:32:01 Restoring     12,885,559 \\cl_vcs02\j$\code\TSMLNG5300_JPN_WIN_337.exe
02/19/2005 13:32:02 Restoring     12,543,008 \\cl_vcs02\j$\code\TSMLNG5300_KOR_WIN_337.exe
```

Figure 24-28 The client restores the files from the beginning

9. When the restore is completed, we can see the final statistics in the schedule log file of the client for a successful operation as shown in Figure 24-29.

Archived

```
dsmsched - Notepad
File  Edit  Format  Help
02/19/2005 13:35:11 Restoring          110,861 \\cl_vcs02\j$\code\isc\wps\migration\template
02/19/2005 13:35:11 Restoring            6,869 \\cl_vcs02\j$\code\isc\wps\migration\template
02/19/2005 13:35:11
Restore processing finished.
02/19/2005 13:35:13 --- SCHEDULEREC STATUS BEGIN
02/19/2005 13:35:13 Total number of objects restored:     1,821
02/19/2005 13:35:13 Total number of objects failed:           0
02/19/2005 13:35:13 Total number of bytes transferred:     1.00 GB
02/19/2005 13:35:13 LanFree data bytes:                    1.00 GB
02/19/2005 13:35:13 Data transfer time:                   15.80 sec
02/19/2005 13:35:13 Network data transfer rate:        66,849.17 KB/sec
02/19/2005 13:35:13 Aggregate data transfer rate:      3,191.04 KB/sec
02/19/2005 13:35:13 Elapsed processing time:           00:05:31
02/19/2005 13:35:13 --- SCHEDULEREC STATUS END
02/19/2005 13:35:13 --- SCHEDULEREC OBJECT END RESTORE 02/19/2005 13:13:46
02/19/2005 13:35:13 --- SCHEDULEREC STATUS BEGIN
02/19/2005 13:35:13 --- SCHEDULEREC STATUS END
02/19/2005 13:35:13 Scheduled event 'RESTORE' completed successfully.
02/19/2005 13:35:13 Sending results for scheduled event 'RESTORE'.
02/19/2005 13:35:13 Results sent to server for scheduled event 'RESTORE'.
```

Figure 24-29 Final statistics for the restore on the schedule log file

Attention: Notice that the restore process is started from the beginning. It is not restarted.

Results summary

The test results show that after a failure on the node that hosts the Tivoli Storage Manager client scheduler instance, a scheduled restore operation started on this node using the LAN-free path is started again from the beginning on the second node of the cluster when the service is online.

This is true if the startup window for the scheduled restore operation is not elapsed when the scheduler client is online again on the second node.

Also notice that the restore is not restarted from the point of failure, but started from the beginning. The scheduler queries the Tivoli Storage Manager server for a scheduled operation and a new session is opened for the client after the failover.

Restriction: Notice again that this configuration, with two Storage Agents in the same machine, is not technically supported by Tivoli Storage Manager for SAN. However, in our lab environment it worked. In other tests we made using the local Storage Agents for communication to the virtual client for LAN-free, the SCSI bus reset did not work and the restore process failed.

Archived

Archived

Part 7

Appendixes

In this part of the book, we describe the Additional Material that is supplied with the book.

© Copyright IBM Corp. 2005. All rights reserved.

Archived

A

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/SG246679

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, SG246379.

Using the Web material

The additional Web material that accompanies this redbook includes the following files listed in Table A-1.

© Copyright IBM Corp. 2005. All rights reserved.

Table A-1 Additional material

File name	Description
sg24_6679_00_HACMP_scripts.tar	This file contains the AIX scripts for HACMP and Tivoli Storage Manager as shown and developed in this IBM Redbook.
sg24_6679_00_TSA_scripts.tar	This file contains the Red Hat scripts for IBM System Automation for Multiplatforms and Tivoli Storage Manager as shown and developed in this IBM Redbook.
sg24_6679_00_VCS_scripts.tar	This file contains the AIX scripts for Veritas Cluster Server and Tivoli Storage Manager as shown and developed in this IBM Redbook.
corrections.zip	If it exists, this file contains updated information and corrections to the book.

Requirements for downloading the Web material

You should have 1 MB of free disk space on your computer.

How to use the Web material

Create a subdirectory (folder) on your workstation, and if applicable, unzip the contents of the Web material zip file into this folder.

Archived

Glossary

A

Agent A software entity that runs on endpoints and provides management capability for other hardware or software. An example is an SNMP agent. An agent has the ability to spawn other processes.

AL See arbitrated loop.

Allocated storage The space that is allocated to volumes, but not assigned.

Allocation The entire process of obtaining a volume and unit of external storage, and setting aside space on that storage for a data set.

Arbitrated loop A Fibre Channel interconnection technology that allows up to 126 participating node ports and one participating fabric port to communicate. See also Fibre Channel Arbitrated Loop and loop topology.

Array An arrangement of related disk drive modules that have been assigned to a group.

B

Bandwidth A measure of the data transfer rate of a transmission channel.

Bridge Facilitates communication with LANs, SANs, and networks with dissimilar protocols.

C

Client A function that requests services from a server, and makes them available to the user. A term used in an environment to identify a machine that uses the resources of the network.

Client authentication The verification of a client in secure communications where the identity of a server or browser (client) with whom you wish to communicate is discovered. A sender's authenticity is demonstrated by the digital certificate issued to the sender.

Client-server relationship Any process that provides resources to other processes on a network is a server. Any process that employs these resources is a client. A machine can run client and server processes at the same time.

Console A user interface to a server.

D

DATABASE 2 (DB2) A relational database management system. DB2 Universal Database is the relational database management system that is Web-enabled with Java support.

Device driver A program that enables a computer to communicate with a specific device, for example, a disk drive.

Disk group A set of disk drives that have been configured into one or more logical unit numbers. This term is used with RAID devices.

Archived

© Copyright IBM Corp. 2005. All rights reserved.

E

Enterprise network A geographically dispersed network under the backing of one organization.

Enterprise Storage Server® Provides an intelligent disk storage subsystem for systems across the enterprise.

Event In the Tivoli environment, any significant change in the state of a system resource, network resource, or network application. An event can be generated for a problem, for the resolution of a problem, or for the successful completion of a task. Examples of events are: the normal starting and s ping of a process, the abnormal termination of a process, and the malfunctioning of a server.

F

Fabric The Fibre Channel employs a fabric to connect devices. A fabric can be as simple as a single cable connecting two devices. The term is often used to describe a more complex network utilizing hubs, switches, and gateways.

FC See Fibre Channel.

FCS See Fibre Channel standard.

Fiber optic The medium and the technology associated with the transmission of information along a glass or plastic wire or fiber.

Fibre Channel A technology for transmitting data between computer devices at a data rate of up to 1 Gb. It is especially suited for connecting computer servers to shared storage devices and for interconnecting storage controllers and drives.

Fibre Channel Arbitrated Loop A reference to the FC-AL standard, a shared gigabit media for up to 127 nodes, one of which can be attached to a switch fabric. See also arbitrated loop and loop topology. Refer to American National Standards Institute (ANSI) X3T11/93-275.

Fibre Channel standard An ANSI standard for a computer peripheral interface. The I/O interface defines a protocol for communication over a serial interface that configures attached units to a communication fabric. Refer to ANSI X3.230-199x.

File system An individual file system on a host. This is the smallest unit that can monitor and extend. Policy values defined at this level override those that might be defined at higher levels.

G

Gateway In the SAN environment, a gateway connects two or more different remote SANs with each other. A gateway can also be a server on which a gateway component runs.

GeoMirror device (GMD) The pseudo-device that adds the geo-mirroring functionality onto a file system or logical volume.

H

Hardware zoning Hardware zoning is based on physical ports. The members of a zone are physical ports on the fabric switch. It can be implemented in the following configurations: one to one, one to many, and many to many.

HBA See host bus adapter.

Host Any system that has at least one internet address associated with it. A host with multiple network interfaces can have multiple internet addresses associated with it. This is also referred to as a server.

Host bus adapter (HBA) A Fibre Channel HBA connection that allows a workstation to attach to the SAN network.

Hub A Fibre Channel device that connects up to 126 nodes into a logical loop. All connected nodes share the bandwidth of this one logical loop. Hubs automatically recognize an active node and insert the node into the loop. A node that fails or is powered off is automatically removed from the loop.

IP Internet protocol.

J

Java A programming language that enables application developers to create object-oriented programs that are very secure, portable across different machine and operating system platforms, and dynamic enough to allow expandability.

Java runtime environment (JRE) The underlying, invisible system on your computer that runs applets the browser passes to it.

Java Virtual Machine (JVM) The execution environment within which Java programs run. The Java virtual machine is described by the Java Machine Specification which is published by Sun Microsystems. Because the Tivoli Kernel Services is based on Java, nearly all ORB and component functions execute in a Java virtual machine.

JBOD Just a Bunch Of Disks.

JRE See Java runtime environment.

JVM See Java Virtual Machine.

L

Local GeoMirror device The local part of a GMD that receives write requests directly from the application and distributes them to the remote device.

Local peer For a given GMD, the node that contains the local GeoMirror device.

Logical unit number (LUN) The LUNs are provided by the storage devices attached to the SAN. This number provides you with a volume identifier that is unique among all storage servers. The LUN is synonymous with a physical disk drive or a SCSI device. For disk subsystems such as the IBM Enterprise Storage Server, a LUN is a logical disk drive. This is a unit of storage on the SAN which is available for assignment or unassignment to a host server.

Loop topology In a loop topology, the available bandwidth is shared with all the nodes connected to the loop. If a node fails or is not powered on, the loop is out of operation. This can be corrected using a hub. A hub opens the loop when a new node is connected and closes it when a node disconnects. See also Fibre Channel Arbitrated Loop and arbitrated loop.

LUN See logical unit number.

LUN assignment criteria The combination of a set of LUN types, a minimum size, and a maximum size used for selecting a LUN for automatic assignment.

LUN masking This allows or blocks access to the storage devices on the SAN. Intelligent disk subsystems like the IBM Enterprise Storage Server provide this kind of masking.

M

Managed object A managed resource.

Managed resource A physical element to be managed.

Management Information Base (MIB) A logical database residing in the managed system which defines a set of MIB objects. A MIB is considered a logical database because actual data is not stored in it, but rather provides a view of the data that can be accessed on a managed system.

MIB See Management Information Base.

MIB object A MIB object is a unit of managed information that specifically describes an aspect of a system. Examples are CPU utilization, software name, hardware type, and so on. A collection of related MIB objects is defined as a MIB.

N

Network topology A physical arrangement of nodes and interconnecting communications links in networks based on application requirements and geographical distribution of users.

N_Port node port A Fibre Channel-defined hardware entity at the end of a link which provides the mechanisms necessary to transport information units to or from another node.

NL_Port node loop port A node port that supports arbitrated loop devices.

O

Open system A system whose characteristics comply with standards made available throughout the industry, and therefore can be connected to other systems that comply with the same standards.

P

Point-to-point topology Consists of a single connection between two nodes. All the bandwidth is dedicated for these two nodes.

Port An end point for communication between applications, generally referring to a logical connection. A port provides queues for sending and receiving data. Each port has a port number for identification. When the port number is combined with an Internet address, it is called a socket address.

Port zoning In Fibre Channel environments, port zoning is the grouping together of multiple ports to form a virtual private storage network. Ports that are members of a group or zone can communicate with each other but are isolated from ports in other zones. See also LUN masking and subsystem masking.

Protocol The set of rules governing the operation of functional units of a communication system if communication is to take place. Protocols can determine low-level details of machine-to-machine interfaces, such as the order in which bits from a byte are sent. They can also determine high-level exchanges between application programs, such as file transfer.

Archived

R

RAID Redundant array of inexpensive or independent disks. A method of configuring multiple disk drives in a storage subsystem for high availability and high performance.

Remote GeoMirror device The portion of a GMD that resides on the remote site and receives write requests from the device on the local node.

Remote peer For a given GMD, the node that contains the remote GeoMirror device.

S

SAN See storage area network.

SAN agent A software program that communicates with the manager and controls the subagents. This component is largely platform independent. See also subagent.

SCSI Small Computer System Interface. An ANSI standard for a logical interface to computer peripherals and for a computer peripheral interface. The interface utilizes a SCSI logical protocol over an I/O interface that configures attached targets and initiators in a multi-drop bus topology.

Server A program running on a mainframe, workstation, or file server that provides shared services. This is also referred to as a host.

Shared storage Storage within a storage facility that is configured such that multiple homogeneous or divergent hosts can concurrently access the storage. The storage has a uniform appearance to all hosts. The host programs that access the storage must have a common model for the information on a storage device. You need to design the programs to handle the effects of concurrent access.

Simple Network Management Protocol (SNMP) A protocol designed to give a user the capability to remotely manage a computer network by polling and setting terminal values and monitoring network events.

Snapshot A point in time copy of a volume.

SNMP See Simple Network Management Protocol.

SNMP agent An implementation of a network management application which is resident on a managed system. Each node that is to be monitored or managed by an SNMP manager in a TCP/IP network, must have an SNMP agent resident. The agent receives requests to either retrieve or modify management information by referencing MIB objects. MIB objects are referenced by the agent whenever a valid request from an SNMP manager is received.

SNMP manager A managing system that executes a managing application or suite of applications. These applications depend on MIB objects for information that resides on the managed system.

SNMP trap A message that is originated by an agent application to alert a managing application of the occurrence of an event.

Software zoning Is implemented within the Simple Name Server (SNS) running inside the fabric switch. When using software zoning, the members of the zone can be defined with: node WWN, port WWN, or physical port number. Usually the zoning software also allows you to create symbolic names for the zone members and for the zones themselves.

SQL Structured Query Language.

Storage administrator A person in the data processing center who is responsible for defining, implementing, and maintaining storage management policies.

Storage area network (SAN) A managed, high-speed network that enables any-to-any interconnection of heterogeneous servers and storage systems.

Subagent A software component of SAN products which provides the actual remote query and control function, such as gathering host information and communicating with other components. This component is platform dependent. See also SAN agent.

Subsystem masking The support provided by intelligent disk storage subsystems like the Enterprise Storage Server. See also LUN masking and port zoning.

Switch A component with multiple entry and exit points or ports that provide dynamic connection between any two of these points.

Switch topology A switch allows multiple concurrent connections between nodes. There can be two types of switches, circuit switches and frame switches. Circuit switches establish a dedicated connection between two nodes. Frame switches route frames between nodes and establish the connection only when needed. A switch can handle all protocols.

T

TCP See Transmission Control Protocol.

TCP/IP Transmission Control Protocol/Internet Protocol.

Topology An interconnection scheme that allows multiple Fibre Channel ports to communicate. For example, point-to-point, arbitrated loop, and switched fabric are all Fibre Channel topologies.

Transmission Control Protocol (TCP) A reliable, full duplex, connection-oriented, end-to-end transport protocol running on of IP.

W

WAN Wide Area Network.

Z

Zoning In Fibre Channel environments, zoning allows for finer segmentation of the switched fabric. Zoning can be used to instigate a barrier between different environments. Ports that are members of a zone can communicate with each other but are isolated from ports in other zones. Zoning can be implemented in two ways: hardware zoning and software zoning.

Archived

Other glossaries:

For more information on IBM terminology, see the IBM Storage Glossary of Terms at:

http://www.storage.ibm.com/glossary.htm

For more information on Tivoli terminology, see the Tivoli Glossary at:

http://publib.boulder.ibm.com/tividd/glossary/termsmst04.htm

Archived

Archived

Abbreviations and acronyms

ABI	Application Binary Interface
ACE	Access Control Entries
ACL	Access Control List
AD	Microsoft Active Directory
ADSM	ADSTAR Distributed Storage Manager
AFS®	Andrew File System
AIX	Advanced Interactive eXecutive
ANSI	American National Standards Institute
APA	All Points Addressable
API	Application Programming Interface
APPC	Advanced Program-to-Program Communication
APPN	Advanced Peer-to-Peer Networking
ARC	Advanced RISC Computer
ARPA	Advanced Research Projects Agency
ASCII	American National Standard Code for Information Interchange
ATE	Asynchronous Terminal Emulation
ATM	Asynchronous Transfer Mode
AVI	Audio Video Interleaved
BDC	Backup Domain Controller
BIND	Berkeley Internet Name Domain
BNU	Basic Network Utilities
BOS	Base Operating System
BRI	Basic Rate Interface
BSD	Berkeley Software Distribution
BSOD	Blue Screen of Death
BUMP	Bring-Up Microprocessor
CA	Certification Authorities
CAD	Client Acceptor Daemon
CAL	Client Access License
C-SPOC	Cluster single point of control
CDE	Common Desktop Environment
CDMF	Commercial Data Masking Facility
CDS	Cell Directory Service
CERT	Computer Emergency Response Team
CGI	Common Gateway Interface
CHAP	Challenge Handshake Authentication
CIDR	Classless InterDomain Routing
CIFS	Common Internet File System
CMA	Concert Multi-threaded Architecture
CO	Central Office

Archived

© Copyright IBM Corp. 2005. All rights reserved.

CPI-C	Common Programming Interface for Communications
CPU	Central Processing Unit
CSNW	Client Service for NetWare
CSR	Client/server Runtime
DAC	Discretionary Access Controls
DARPA	Defense Advanced Research Projects Agency
DASD	Direct Access Storage Device
DBM	Database Management
DCE	Distributed Computing Environment
DCOM	Distributed Component Object Model
DDE	Dynamic Data Exchange
DDNS	Dynamic Domain Name System
DEN	Directory Enabled Network
DES	Data Encryption Standard
DFS™	Distributed File System
DHCP	Dynamic Host Configuration Protocol
DLC	Data Link Control
DLL	Dynamic Load Library
DS	Differentiated Service
DSA	Directory Service Agent
DSE	Directory Specific Entry
DNS	Domain Name System
DTS	Distributed Time Service
EFS	Encrypting File Systems
EGID	Effective Group Identifier
EISA	Extended Industry Standard Architecture
EMS	Event Management Services
EPROM	Erasable Programmable Read-Only Memory
ERD	Emergency Repair Disk
ERP	Enterprise Resources Planning
ERRM	Event Response Resource Manager
ESCON®	Enterprise System Connection
ESP	Encapsulating Security Payload
ESS	Enterprise Storage Server
EUID	Effective User Identifier
FAT	File Allocation Table
FC	Fibre Channel
FDDI	Fiber Distributed Data Interface
FDPR	Feedback Directed Program Restructure
FEC	Fast EtherChannel technology
FIFO	First In/First Out
FIRST	Forum of Incident Response and Security
FQDN	Fully Qualified Domain Name
FSF	File Storage Facility
FTP	File Transfer Protocol
FtDisk	Fault-Tolerant Disk
GC	Global Catalog
GDA	Global Directory Agent
GDI	Graphical Device Interface

Archived

GDS	Global Directory Service
GID	Group Identifier
GL	Graphics Library
GSNW	Gateway Service for NetWare
GUI	Graphical User Interface
HA	High Availability
HACMP	High Availability Cluster Multiprocessing
HAL	Hardware Abstraction Layer
HBA	Host Bus Adapter
HCL	Hardware Compatibility List
HSM	Hierarchical Storage Management
HTTP	Hypertext Transfer Protocol
IBM	International Business Machines Corporation
ICCM	Inter-Client Conventions Manual
IDE	Integrated Drive Electronics
IDL	Interface Definition Language
IDS	Intelligent Disk Subsystem
IEEE	Institute of Electrical and Electronic Engineers
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IIS	Internet Information Server
IKE	Internet Key Exchange
IMAP	Internet Message Access Protocol
I/O	Input/Output
IP	Internet Protocol
IPC	Interprocess Communication
IPL	Initial Program Load
IPsec	Internet Protocol Security
IPX	Internetwork Packet eXchange
ISA	Industry Standard Architecture
iSCSI	SCSI over IP
ISDN	Integrated Services Digital Network
ISNO	Interface-specific Network Options
ISO	International Standards Organization
ISS	Interactive Session Support
ISV	Independent Software Vendor
ITSEC	Initial Technology Security Evaluation
ITSO	International Technical Support Organization
ITU	International Telecommunications Union
IXC	Inter Exchange Carrier
JBOD	Just a Bunch of Disks
JFS	Journaled File System
JIT	Just-In-Time
L2F	Layer 2 Forwarding
L2TP	Layer 2 Tunneling Protocol
LAN	Local Area Network
LCN	Logical Cluster Number

Archived

LDAP	Lightweight Directory Access Protocol
LFS	Log File Service (Windows NT®)
LFS	Logical File System (AIX)
LFT	Low Function Terminal
JNDI	Java Naming and Directory Interface
LOS	Layered Operating System
LP	Logical Partition
LPC	Local Procedure Call
LPD	Line Printer Daemon
LPP	Licensed Program Product
LRU	Least Recently Used
LSA	Local Security Authority
LTG	Local Transfer Group
LUID	Login User Identifier
LUN	Logical Unit Number
LVCB	Logical Volume Control Block
LVDD	Logical Volume Device Driver
LVM	Logical Volume Manager
MBR	Master Boot Record
MDC	Meta Data Controller
MFT	Master File Table
MIPS	Million Instructions Per Second
MMC	Microsoft Management Console
MOCL	Managed Object Class Library
MPTN	Multi-protocol Transport Network
MS-DOS	Microsoft Disk Operating System
MSCS	Microsoft Cluster Server
MSS	Maximum Segment Size
MSS	Modular Storage Server
MWC	Mirror Write Consistency
NAS	Network Attached Storage
NBC	Network Buffer Cache
NBF	NetBEUI Frame
NBPI	Number of Bytes per I-node
NCP	NetWare Core Protocol
NCS	Network Computing System
NCSC	National Computer Security Center
NDIS	Network Device Interface Specification
NDMP	Network Data Management Protocol
NDS	NetWare Directory Service
NETID	Network Identifier
NFS	Network File System
NIM	Network Installation Management
NIS	Network Information System
NIST	National Institute of Standards and Technology
NLS	National Language Support
NNS	Novell Network Services
NSAPI	Netscape Commerce Server's Application
NTFS	NT File System

Archived

NTLDR	NT Loader
NTLM	NT LAN Manager
NTP	Network Time Protocol
NTVDM	NT Virtual DOS Machine
NVRAM	Non-Volatile Random Access Memory
NetBEUI	NetBIOS Extended User Interface
NetDDE	Network Dynamic Data Exchange
OCS	On-Chip Sequencer
ODBC	Open Database Connectivity
ODM	Object Data Manager
OLTP	OnLine Transaction Processing
OMG	Object Management Group
ONC	Open Network Computing
OS	Operating System
OSF	Open Software Foundation
OU	Organizational Unit
PAL®	Platform Abstract Layer
PAM	Pluggable Authentication Module
PAP	Password Authentication Protocol
PBX	Private Branch Exchange
PCI	Peripheral Component Interconnect
PCMCIA	Personal Computer Memory Card International Association
PDC	Primary Domain Controller
PDF	Portable Document Format
PDT	Performance Diagnostic Tool
PEX	PHIGS Extension to X
PFS	Physical File System
PHB	Per Hop Behavior
PHIGS	Programmer's Hierarchical Interactive Graphics System
PID	Process Identification Number
PIN	Personal Identification Number
PMTU	Path Maximum Transfer Unit
POP	Post Office Protocol
POSIX	Portable Operating System Interface for Computer Environment
POST	Power-On Self Test
PP	Physical Partition
PPP	Point-to-Point Protocol
PPTP	Point-to-Point Tunneling Protocol
PReP	PowerPC® Reference Platform
PSM	Persistent Storage Manager
PSN	Program Sector Number
PSSP	Parallel System Support Program
PV	Physical Volume
PVID	Physical Volume Identifier
QoS	Quality of Service
RACF®	Resource Access Control Facility

Archived

RAID	Redundant Array of Independent Disks
RAS	Remote Access Service
RDBMS	Relational Database Management System
RFC	Request for Comments
RGID	Real Group Identifier
RISC	Reduced Instruction Set Computer
RMC	Resource Monitoring and Control
RMSS	Reduced-Memory System Simulator
ROLTP	Relative OnLine Transaction Processing
ROS	Read-Only Storage
RPC	Remote Procedure Call
RRIP	Rock Ridge Internet Protocol
RSCT	Reliable Scalable Cluster Technology
RSM	Removable Storage Management
RSVP	Resource Reservation Protocol
SACK	Selective Acknowledgments
SAK	Secure Attention Key
SAM	Security Account Manager
SAN	Storage Area Network
SASL	Simple Authentication and Security Layer
SCSI	Small Computer System Interface
SDK	Software Developer's Kit
SFG	Shared Folders Gateway
SFU	Services for UNIX
SID	Security Identifier
SLIP	Serial Line Internet Protocol
SMB	Server Message Block
SMIT	System Management Interface Tool
SMP	Symmetric Multiprocessor
SMS	Systems Management Server
SNA	Systems Network Architecture
SNAPI	SNA Interactive Transaction Program
SNMP	Simple Network Management Protocol
SP	System Parallel
SPX	Sequenced Packet eXchange
SQL	Structured Query Language
SRM	Security Reference Monitor
SSA	Serial Storage Architecture
SSL	Secure Sockets Layer
SUSP	System Use Sharing Protocol
SVC	Serviceability
TAPI	Telephone Application Program Interface
TCB	Trusted Computing Base
TCP/IP	Transmission Control Protocol/Internet Protocol
TCSEC	Trusted Computer System Evaluation Criteria
TDI	Transport Data Interface

Archived

TDP	Tivoli Data Protection
TLS	Transport Layer Security
TOS	Type of Service
TSM	IBM Tivoli Storage Manager
TTL	Time to Live
UCS	Universal Code Set
UDB	Universal Database
UDF	Universal Disk Format
UDP	User Datagram Protocol
UFS	UNIX File System
UID	User Identifier
UMS	Ultimedia Services
UNC	Universal Naming Convention
UPS	Uninterruptable Power Supply
URL	Universal Resource Locator
USB	Universal Serial Bus
UTC	Universal Time Coordinated
UUCP	UNIX to UNIX Communication Protocol
UUID	Universally Unique Identifier
VAX	Virtual Address eXtension
VCN	Virtual Cluster Name
VFS	Virtual File System
VG	Volume Group
VGDA	Volume Group Descriptor Area
VGSA	Volume Group Status Area
VGID	Volume Group Identifier
VIPA	Virtual IP Address
VMM	Virtual Memory Manager
VP	Virtual Processor
VPD	Vital Product Data
VPN	Virtual Private Network
VRMF	Version, Release, Modification, Fix
VSM	Virtual System Management
W3C	World Wide Web Consortium
WAN	Wide Area Network
WFW	Windows for Workgroups
WINS	Windows Internet Name Service
WLM	Workload Manager
WWN	World Wide Name
WWW	World Wide Web
WYSIWYG	What You See Is What You Get
WinMSD	Windows Microsoft Diagnostics
XCMF	X/Open Common Management Framework
XDM	X Display Manager
XDMCP	X Display Manager Control Protocol
XDR	eXternal Data Representation
XNS	XEROX Network Systems
XPG4	X/Open Portability Guide

Archived

Archived

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 1050. Note that some of the documents referenced here may be available in softcopy only.

- *IBM Tivoli Storage Manager Version 5.3 Technical Guide,* SG24-6638-00
- *IBM Tivoli Storage Management Concepts,* SG24-4877-03
- *IBM Tivoli Storage Manager Implementation Guide*, SG24-5416-02
- *IBM HACMP for AIX V5.X Certification Study Guide*, SG24-6375-00
- *AIX 5L Differences Guide Version 5.3 Edition*, SG24-7463-00
- *Introducing VERITAS Foundation Suite for AIX*, SG24-6619-00
- *The IBM TotalStorage NAS Gateway 500 Integration Guide*, SG24-7081-01
- *Tivoli Storage Manager Version 5.1 Technical Guide*, SG24-6554-00
- *Tivoli Storage Manager Version 4.2 Technical Guide*, SG24-6277-00
- *Tivoli Storage Manager Version 3.7.3 & 4.1: Technical Guide*, SG24-6110-00
- *ADSM Version 3 Technical Guide*, SG24-2236-01
- *Tivoli Storage Manager Version 3.7: Technical Guide*, SG24-5477-00
- *Understanding the IBM TotalStorage Open Software Family*, SG24-7098-00
- *Exploring Storage Management Efficiencies and Provisioning - Understanding IBM TotalStorage Productivity Center and IBM TotalStorage Productivity Center with Advanced Provisioning,* SG24-6373-00

Other publications

These publications are also relevant as further information sources:

Tivoli Storage Manager V5.3 Administrator's Guides

- *TSM V5.3 for HP-UX Administrator's Guide*, GC32-0772-03

© Copyright IBM Corp. 2005. All rights reserved.

- *TSM V5.3 for Windows Administrator's Guide*, GC32-0782-03
- *TSM V5.3 for Sun Solaris Administrator's Guide*, GC32-0778-03
- *TSM V5.3 for Linux Administrator's Guide*, GC23-4690-03
- *TSM V5.3 for z/OS Administrator's Guide*, GC32-0775-03
- *TSM V5.3 for AIX Administrator's Guide*, GC32-0768-03

Tivoli Storage Manager V5.3 Administrator's References

- *TSM V5.3 for HP-UX Administrator's Reference*, GC32-0773-03
- *TSM V5.3 for Sun Administrator's Reference*, GC32-0779-03
- *TSM V5.3 for AIX Administrator's Reference*, GC32-0769-03
- *TSM V5.3 for z/OS Administrator's Reference*, GC32-0776-03
- *TSM V5.3 for Linux Administrator's Reference*, GC23-4691-03
- *TSM V5.3 for Windows Administrator's Reference*, GC32-0783-03

Tivoli Storage Manager V5.3 Data Protection Publications

- *ITSM for Mail 5.3: Data Protection for Lotus Domino for UNIX, Linux, and OS/400 Installation and User's Guide*, SC32-9056-02
- *ITSM for Mail 5.3: Data Protection for Lotus Domino for Windows Installation and User's Guide*, SC32-9057-01

Tivoli Storage Manager V5.3 Install Guide

- *TSM V5.3 for AIX Installation Guide*, GC32-1597-00
- *TSM V5.3 for Sun Solaris Installation Guide*, GC32-1601-00
- *TSM V5.3 for Linux Installation Guide*, GC32-1599-00
- *TSM V5.3 for z/OS Installation Guide*, GC32-1603-00
- *TSM V5.3 for Windows Installation Guide*, GC32-1602-00
- *TSM V5.3 for HP-UX Installation Guide*, GC32-1598-00

Tivoli Storage Manager V5.3 Messages

- *TSM V5.3 Messages*, SC32-9090-02

Tivoli Storage Manager V5.3 Performance Tuning Guide

- *TSM V5.3 Performance Tuning Guide*, SC32-9101-02

Tivoli Storage Manager V5.3 Read This First

- *TSM V5.3 Read This First*, GI11-0866-06

Archived

Tivoli Storage Manager V5.3 Storage Agent User's Guides

- *TSM V5.3 for SAN for AIX Storage Agent User's Guide*, GC32-0771-03
- *TSM V5.3 for SAN for HP-UX Storage Agent User's Guide*, GC32-0727-03
- *TSM V5.3 for SAN for Linux Storage Agent User's Guide*, GC23-4693-03
- *TSM V5.3 for SAN for Sun Solaris Storage Agent User's Guide*, GCGC32-0781-03
- *TSM V5.3 for SAN for Windows Storage Agent User's Guide*, GC32-0785-03

Tivoli Storage Manager V5.3.0 Backup-Archive Clients

- *TSM 5.3 Using the Application Program Interface*, GC32-0793-03
- *TSM 5.3 NetWare Backup-Archive Clients Installation and User's Guide*, GC32-0786-05
- *TSM 5.3 UNIX and Linux Backup-Archive Clients Installation and User's Guide*, GC32-0789-05
- *TSM 5.3 Windows Backup-Archive Client Installation and User's Guide*, GC32-0788-05
- *TSM 5.3 for Space Management for UNIX and Linux User's Guide*, GC32-0794-03

Online resources

These Web sites and URLs are also relevant as further information sources:

- IBM Tivoli Storage Manager product page:

  http://www.ibm.com/software/tivoli/products/storage-mgr/

- IBM Tivoli Storage Manager information center:

  http://publib.boulder.ibm.com/infocenter/tivihelp/index.jsp?toc=/com.ibm.itstorage.doc/toc.xml

- IBM Tivoli Storage Manager product support:

  http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html

- IBM Tivoli Support:

  http://www.ibm.com/software/sysmgmt/products/support

- IBM Tivoli Support - Tivoli support lifecycle:

  http://www.ibm.com/software/sysmgmt/products/support/eos.html

- IBM Software Support Lifecycle - Tivoli Product lifecycle dates:

  http://www.ibm.com/software/info/supportlifecycle/list/t.html

Archived

- Tivoli Support - IBM Tivoli Storage Manager Supported Devices for AIX HPUX SUN WIN:

  http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_AIXHPSUNWIN.html

- Tivoli Support - IBM Tivoli Storage Manager Version Release Information:

  http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManagerVersionRelease.html

- IBM Tivoli System Automation for Multiplatforms:

  http://www.ibm.com/software/tivoli/products/sys-auto-linux/

- IBM Tivoli System Automation for Multiplatforms Version 1.2 Release Notes:

  http://publib.boulder.ibm.com/tividd/td/IBMTivoliSystemAutomationforMultiplatforms1.2.html

- Red Hat Linux:

  http://www.redhat.com/

- SUSE Linux:

  http://www.novell.com/linux/suse/index.html

- Microsoft Cluster Server General Questions:

  http://www.microsoft.com/ntserver/support/faqs/Clustering_faq.asp

- Guide to Creating and Configuring a Server Cluster under Windows Server 2003:

  http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/clustering/confclus.mspx

- VERITAS Clustering family of products:

  http://www.veritas.com/Products/www?c=subcategory&refId=150&categoryId=149

- VERITAS Software Support:

  http://support.veritas.com/

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Archived

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Archived

Archived

Index

Numerics
64-bit hardware 456, 744–745

A
Activity log 152, 156–159, 165, 213, 216, 218, 221, 223, 228, 278–279, 285, 287, 318–320, 323–324, 369–370, 375, 400, 404, 408, 412, 643–644, 646–647, 649–651, 665–666, 669, 671, 690–691, 693, 697, 950, 954, 956–957, 989–990, 993, 995, 1017, 1023
    informational message 159
activity log
    informational message 957
actlog 412, 495, 523, 583, 588, 691, 696, 872, 874
ADMIN_CENTER administrator 177, 239
Administration Center
    Cluster resources 633
    Installation 117
administration center
    Enterprise Administration 562, 564
Administration Center (AC) 13, 79, 92, 104, 112, 117, 173, 236, 427, 436, 438, 453–454, 464, 472–473, 478, 528, 531, 557, 562, 564, 567, 619, 621–624, 633, 639, 675, 720, 727, 729, 840, 842, 850, 933, 938, 944–945, 980
administrative interface 160, 164, 225, 227, 619, 626, 649, 651, 704, 960, 963
administrator ADMIN 870
administrator SCRIPT_OPERATOR 826–828, 834–835, 875
Agents 705
Aggregate data transfer rate 515, 876
AIX 5L
    5.1 424
    base operating system 714
    V5.3 419, 432
AIX 5L V5.3 441
AIX command
    line 448–449, 534, 731
    lscfg 725
    lslpp 432, 460, 749
    smitty installp 561, 798
    tail 771, 782, 811
AIX machine 239, 276, 316, 333, 378, 988, 1001
AIX patch 735
AIX server 239, 277, 489, 507, 512, 537, 541, 544, 551, 557, 572–574, 580, 586, 759–760, 782, 826, 832, 871, 874
allMediaLocation 466, 473, 622–623, 843
ANR0406I Session 1 784
ANR0916I TIVOLI STORAGE Manager 784
ANR0993I Server initialization 784
ANR1639I Attribute 828, 835, 870–871, 875
ANR2017I Administrator ADMIN 833
ANR2017I Administrator SCRIPT_OPERATOR 826–827, 834–835, 875
ANR2034E Select 875
ANR2560I Schedule manager 784
ANR2803I License manager 784
ANR2828I Server 784
ANR7800I DSMSERV 628, 680
ANS1809W Session 782
Application monitor 712
Application server 712
application server 31, 430, 465, 490, 493, 529, 534, 712, 717
atlantic lladdress 511, 569–570
atlantic root 724
attached disk device
    Linux scans 606
Attributes 707
automated fallover 5

B
Backup domain 250–251, 290–291, 530, 656, 968
backup file 486, 687, 754
Backup Operation 150, 211, 536, 538, 543, 548, 583–584, 620, 643, 870–872
backup storage pool
    command 649, 790
    failure 24
    operation 517, 787
    process 159, 224, 519, 647–649, 790, 957, 960
    tape 24
    task 159, 224, 956, 959
backup storage pool process 156, 159–160, 221,

© Copyright IBM Corp. 2005. All rights reserved.

224–225, 955, 957, 960
backup/archive client 675, 683, 965–966, 968–969
Installation 968
backup-archive GUI 252, 293, 969
Base fix 442
boot-time address 425
Bundled agents 705

C

case cluster 510
cd command 758
cdrom directory 621–623
change management 4
chvg command 440, 729–730
click Finish 42, 54, 58, 66, 70, 86, 89, 91, 102, 127, 129, 134, 138, 141, 189–190, 196, 200, 204, 247, 354, 364, 387, 395, 566, 679, 901, 912–913, 917, 930, 1008
click Go 176, 238, 342, 348, 562, 564, 567–568
click Next 41–42, 60–61, 65, 67, 69, 80, 83–84, 87, 89, 93–99, 102, 104–112, 115, 126, 128–133, 135–137, 140–141, 168–169, 171, 187–189, 191–194, 197, 199–200, 203, 232–233, 235, 333, 343–347, 349, 352–353, 362–363, 385–386, 393–394, 888–889, 891–893, 895–899, 910–912, 914–916, 920–921, 923–927, 933–934, 936–940, 975, 977–983, 1007–1008
84, 94, 127, 130, 132, 170, 188, 190, 192–193, 195, 197, 233, 244–245, 260–261, 270–271, 301–302, 310–311, 333, 352, 354, 363, 394
Client
enhancements, additions and changes 453
Client Accepter
Daemon 859, 863
client accepter 250–252, 254, 266–274, 290–291, 293, 295, 306–314, 532, 537, 544, 546, 658, 660, 857, 859, 968–969, 985–986, 988
Client Accepter Daemon (CAD) 859
Client Acceptor Daemon (CAD) 660–661
client backup 148, 150–151, 209–211, 213, 506–507, 537, 541, 544, 551, 640, 642, 781–782
Client Node 341–342, 373, 405, 528–530, 532, 561, 654–655, 658, 681, 1020
high level address 530, 656
low level address 530, 656
client node
communication paths 561
failover case 546
client restart 219, 278, 318, 370, 372, 401, 404, 665, 695, 1019–1020
client session 211, 277, 284, 317, 323, 367, 374, 398, 406, 537, 541, 546, 643, 660, 688, 828, 874, 949, 989, 1016, 1021
cluster 704
cluster address 430
local resolution 430
Cluster Administrator 42, 44, 59–60, 66–67, 70, 76, 123, 140, 143, 146–147, 150–151, 154–157, 161, 165, 167, 170, 172, 185, 202, 205, 207–208, 211, 213, 216–217, 219, 222, 225, 227, 231, 234, 236, 254, 257, 259, 264, 269, 273, 276, 278, 284, 286, 298, 300, 304, 310, 313, 316, 318, 323–324, 361–362, 365–367, 370, 373, 392, 396–398, 400, 406, 710
cluster command 501, 506, 511, 515, 517, 520, 524, 536, 540, 544, 550, 578, 584, 870
cluster configuration 9, 21, 78, 124, 132, 135, 138–140, 142, 181, 185, 197, 200–201, 203–205, 249, 290, 327, 333, 378, 422, 431, 464, 481, 624, 703–704, 708, 713, 715, 787, 842, 915, 1001
logical components 9
cluster configurations 708
Cluster group
name 251, 291
cluster group 31, 43, 47, 74, 130–131, 135, 140–141, 144, 150, 154, 156, 161, 165, 167, 170, 173, 192–193, 197–198, 202–203, 205–206, 208, 211, 216, 219, 222, 225, 227, 231, 234, 236, 242, 251, 253, 255, 257–259, 264, 266–269, 272, 291–292, 294, 296, 298, 300, 304, 306–309, 313, 333, 340, 366, 378, 397
Client Acceptor service 267, 307
new resource 259, 300
Option file 255, 296
scheduler service 257, 298
Tivoli Storage Manager Client Acceptor service 266, 306
Cluster Manager
GUI 766, 775–777, 808, 817–819
Web 738–739
cluster membership 705, 708, 771, 779–780, 783, 812, 822, 824
Cluster multi-processing 4
cluster name 19, 30, 42–43, 46, 73–74, 275, 278, 315, 318, 429, 443, 615, 881, 898, 974, 990
cluster node 9, 34, 49, 67, 124, 185, 383, 421–425, 430, 443, 445, 447, 455, 464, 486, 492, 513, 528,

544, 600, 613, 654, 664, 668, 707–708, 711, 713–714, 744, 1005
efficient use 9
following task 430
service interface 424
Tivoli Storage Manager server 528
cluster operation 478, 506, 511, 515, 517, 520, 524, 536, 540, 544, 550, 578, 584, 598, 782, 785, 788, 791–792, 825, 870, 896, 902
cluster resource 8, 91, 120, 124, 181, 361, 392, 421, 449, 481–482, 496, 619, 624, 629, 703, 711, 795, 1013
Cluster resources 705
Cluster Server
Version 4.0 running 701
cluster server 704–705, 708–709, 712, 716, 719–720, 731, 734, 740–742
cluster servers 704
cluster service 28, 35, 41–42, 44, 51, 59, 64, 68, 76, 482–483, 496, 499–500, 506, 511, 515, 517, 520, 524, 536, 540, 544, 550, 578, 584, 770, 773, 777, 781, 785, 788, 791, 810, 814, 820, 825, 831, 857, 870, 873, 920, 932–933, 975
cluster software 6–9, 17, 612, 794
clusternode yes 254–256, 259, 269, 295–297, 300, 309, 970–971, 974
command cp 771, 785, 811
engine log 785
command dsmserv 488, 756
Command Line
Backup/Archive Client Interface 659
command line 219, 435–436, 440, 443, 445, 448–449, 454, 456, 464, 478, 487, 489, 494, 506, 511, 562–563, 567–569, 584, 619, 621, 626, 675, 682, 714, 745, 753, 755, 763, 770, 772, 776, 778, 782, 810, 813, 819, 821, 831, 842, 872–873
same command 436
COMMMethod TCPip 569–570, 626, 679, 681, 799
completion state 787
concurrent access 4, 420
ConfigInput.admi nName 466, 622, 843
ConfigInput.admi nPass 466, 622, 843
ConfigInput.veri fyPass 466, 622, 843
configuration file 351, 363, 384, 394, 439, 454, 529, 532, 534, 558, 569, 603, 609–610, 618, 626, 630, 633–634, 655–656, 661, 679, 681, 684, 728, 730, 795, 798, 997
different path 529
different paths 655
disk volumes 454
configuration process 124, 185, 205, 254, 295, 350–351, 384, 676, 970, 1006
Copy Storage Pool 121, 156, 158–160, 182, 221–222, 224, 518, 647–648, 788, 907, 955–960
command q occupancy 958
primary tape storage pool 955
tape volume 159
tape volumes 160
valid volume 958
copy storage pool
SPCPT_BCK 955
tape volume 159, 224, 960
Tivoli Storage Manager 648
Copying VRTSvcsag.rte.bff.gz 741
cp startserver 490, 571, 573
cp stopserver 490, 571, 573
Custom agents 705
CUSTOMIZABLE Area 630–631, 633–634

D

Data transfer time 515, 830, 837, 876
database backup 160–161, 163–164, 225–227, 520, 522, 649–650, 785, 791, 960–963
command 225
operation 523, 791
process 161–162, 164, 225, 227, 523, 649–650, 792, 960–961, 963
Process 1 starts 961
task 961
volume 162–163, 522
datareadpath 383, 1005
David Bohm 759–760
DB backup
failure 24
default directory 528, 533, 571, 573, 654
Definition file
SA-nfsserver-tsmsta.def 684
detailed description 122, 183, 339, 381, 635, 637, 656, 707, 908
detailed information 494, 599, 618, 691, 902
devc 161, 225, 520, 649, 791
devconfig file 384, 1006
devconfig.txt file 360, 392, 557, 680, 798, 1006, 1012
default path 1006
devconfig.txt location 335, 379, 559, 796, 1003
device name 82, 89, 331, 337, 349–350, 381, 560,

568, 574, 907, 1001, 1004
disk 5
disk adapter 5
Disk channel 704
disk device 606–607, 609–610
persistent SCSI addresses 607
disk drive 107, 120, 181, 193, 351, 357–358, 389–390, 619, 906, 909, 1010–1011
disk resource 42, 44, 74, 78, 122, 130, 140, 183, 192, 202, 253, 271, 294, 311, 904, 929, 970
Disk Storage Pool 154, 452, 487–488, 515, 536, 618, 627, 633, 663, 756, 785, 907, 948, 952–953
Testing migration 952
Disk storage pool
enhancement 12
migration 645
disk storage pool
client backup restarts 643
DNS name 31, 47, 882
DNS server 28, 34, 50, 118, 180, 882–884, 944
DNS tab 33, 49
domain controller 28, 34, 50, 118, 180, 882–883
domain e 256, 297
domain j 255, 296, 971
Domain Name System (DNS) 28
domain Standard 872
downtime 4
planned 4
unplanned 4
drive library 384, 569, 628, 1005–1006
drop-down list 923–924, 936, 978
TSM Server1 service 924
dsm.sys file
stanza 841
dsmadmc command 456, 532, 745, 759–760, 799, 805
dsmcutil program 258, 266, 298, 306, 972, 986
dsmcutil tool 259, 266–267, 300, 306–307, 974, 986
same parameters 259, 300
dsmfmt command 487, 627, 755
dsmserv format
1 488, 627, 756
command 488, 627, 756
dsmserv.dsk file 488, 754
dsmsta setstorageserver
command 569, 679–680, 798
myname 569, 680, 798
utility 357, 389, 1010
Dynamic node priority (DNP) 426, 712, 717

E

Encrypting File System (EFS) 79, 242
engine_A 771, 773, 775, 777–780, 782–783, 785, 788, 791, 811, 814, 817, 819, 821–822, 824–825
Enhanced Scalability (ES) 711–712, 714–715, 718
Enterprise agents 705
Enterprise Management 175, 238, 383, 675–676, 1005
environment variable 488, 490, 613, 627, 680, 756, 857–858, 860
Error log
file 643
RAS 418
error message 34, 50, 62, 158, 162, 620, 645, 710, 858, 861, 884, 974, 995, 1018
errorlogretention 7 255, 296–297, 627, 971
Ethernet cable 505, 779–780, 822–823
event log 154, 216, 280–281, 287–288, 320, 325–326, 951–952, 991–993, 996, 1024
event trigger 710
example script 490, 532, 573
exit 0 493, 760–761, 804, 806, 858, 863–864
export DSMSERV_DIR=/usr/tivoli/tsm/StorageAgent/bin 569, 804
export LANG 758, 804

F

failover 5, 8, 78–79, 136, 154, 156, 165, 198–199, 215, 221, 229, 257, 269, 282–283, 289, 298, 309, 318, 321–322, 326, 377, 412, 629, 641, 645–646, 648, 654, 660, 665, 667, 669, 672, 687, 690, 695, 697, 700, 779, 783, 788–789, 791–792, 795, 822, 824, 829–833, 835, 837, 857, 859, 871, 873, 904, 909, 923, 952, 958, 992, 997, 1025
failover time 712
failure detection 5
fault tolerant systems 6
Fibre Channel
adapter 28, 606
bus 28
driver 600
fibre channel
driver 607
File System 79, 242, 607, 609, 619, 625, 658–659, 684, 720, 727–730, 784
file TSM_ISC_5300_AIX 465, 843

Archived

filesets 455–456, 458, 460–462, 464, 561, 732–734, 744–745, 747, 749–751, 753, 798
Filespace name 274–275, 314–315, 990
filesystems 428, 438–439, 454, 465, 480, 487
final smit 463, 752
final statistic 153, 218, 288, 326, 372, 376, 403, 411, 671, 997, 1019, 1024
first node 41, 59, 67, 80, 91–92, 102–104, 116–118, 123, 139, 150, 154, 159, 184, 201–202, 210, 215, 243, 248, 265, 274, 314, 332, 426, 435, 441, 448, 490, 571, 573, 621, 625, 634, 641, 645–647, 649, 651, 661, 664, 668, 675, 679, 684, 695, 887, 909, 919–920, 948, 952, 957, 985, 1017
  Administration Center installation 104
  backup storage pool process 159
  command line 435
  configuration procedure 123, 184
  diskhbvg volume 441
  example scripts 490
  local Storage Agent 675
  power cables 641
  Tivoli Storage Manager 123, 185
  Tivoli Storage Manager server 123, 184
first time 159, 607, 922, 955, 957
function CLEAN_EXIT 858, 860

G

GAB protocol 704
GABdisk 705
General Parallel File System (GPFS) 621, 626
generic applications 7
Generic Service 168, 170, 172, 231–232, 234–235, 254, 259–260, 262, 270–273, 295, 300–302, 310–313, 362, 393, 923, 936, 974–975, 978, 986–987
generic service
  application 923
  resource 168, 172, 231, 235, 254, 259–260, 265, 269–270, 277, 295, 300–301, 305, 309–310, 357, 362, 389, 393, 974, 986
grant authority
  admin class 489, 629, 757
  script_operator class 490, 757
Graphical User Interface (GUI) 704
grep dsmserv 486, 754
grep Online 772, 775, 777–778, 780–781, 813, 817, 819–820, 823–824
grep Z8 725
Group Membership Services/Atomic Broadcast (GAB) 704, 708

H

HACMP 704, 710
HACMP cluster 417, 443, 464, 486, 496, 505, 560, 584, 590, 711–713
  active nodes 713
  Components 714
  IP networks 711
  public TCP/IP networks 711
HACMP environment 420, 422, 528
  design conciderations 422
  Tivoli Storage Manager 528
HACMP event scripts 711
HACMP menu 715
HACMP V5.2
  installation 531
  product 555
HACMP Version
  4.5 718
  5.1 433
  5.2 433
hagrp 772, 775, 813, 817
Hardware Compatibility List (HCL) 29
hastatus 770–772, 775, 777, 780–781, 785, 788–789, 791, 810–814, 817, 819–820, 823–824, 831
hastatus command 770, 773, 789, 812, 814, 873
hastatus log 811
hastatus output 772, 775, 813, 817
heartbeat protocol 711
High Availability
  Cluster Multi-Processing 415, 417–425, 431–433, 435–436, 441–450, 703, 710–716
High availability
  daemon 708
  system 6
high availability 5–6, 703
High availability (HA) 3–7, 419–420, 595, 704, 708–709, 713, 715
High Availability Cluster Multi-Processing (HACMP) 417, 419–422, 424, 431–433, 436, 441–449, 710–716
High Availability Daemon (HAD) 708
High Available (HA) 419
Highly Available application 9, 422, 527, 531, 618, 653, 657, 701, 753, 839–840

Archived

Host Bus Adapter (HBA) 602, 611
http port 254, 296, 847, 970
httpport 1582 255, 296
HW Raid-5 20

I

IBM Tivoli Storage Manager 1, 12–14, 79–80, 92, 329, 452–454, 486–487, 555–556, 618–619, 627, 658–659, 681, 683, 754–755, 793–794, 903–904, 933, 965, 999
  Administration Center 14, 92
  Administrative Center 933
  backup-archive client 454
  Client 527
  Client enhancements, additions, and changes 453
  database 487, 755
  different high availability clusters solutions 1
  new features overview 452
  product 12
  Scheduling Flexibility 13
  Server 453, 754, 933
  Server enhancements, additions and changes 13, 453
  V5.3 12
  V5.3.0 933
  Version 5.3 12, 25, 415, 591, 701, 877
IBM Tivoli Storage Manager Client. see Client
IBM Tivoli Storage Manager Server. see Server
importvg command 440–441, 729
Include-exclude enhancement 14, 453
incremental backup 146–147, 149–150, 154, 208, 211, 276–277, 279, 281–283, 316–317, 319–320, 322, 367, 371–372, 398, 402–404, 506–507, 509, 533, 639–640, 643, 659, 663–664, 667, 682, 687, 694, 945–946, 948–949, 952, 989, 991–992, 1015, 1019
  local mounted file systems 659
  local mounted filesystems 533
  tape storage pool 663
installation path 80, 173, 236, 243, 245, 258, 266, 298–299, 306, 332, 351, 384, 468, 1006
installation process 80, 103, 106, 116–118, 122, 179, 183, 243, 332, 339, 381, 466, 473, 622, 757, 843, 857, 893
InstallShield wizard 80, 244, 466, 473, 622–623, 843
installvcs script 709
Instance path 558, 674, 795
Integrated Solution 92–93, 436, 438, 464–465, 492, 531, 621–622, 624, 720, 727, 729, 840, 842, 880, 933
  Installation 621
  installation process 622
  storage resources 438
  Tivoli Storage Manager Administration Center 464
Integrated Solution Console (ISC) 425, 427, 430, 528–533, 536, 557, 559, 564, 567, 569–570, 572, 577, 580, 586, 754, 757, 795–796, 799, 804–809, 829–831, 836
Integrated Solutions Console (ISC) 92, 96–97, 99, 102–103, 107–108, 110, 116–117, 120, 167, 170, 172–174, 181, 231, 234–237, 455, 464–465, 469–470, 472, 478, 489, 492, 619, 621–624, 633–639, 839–846, 849, 852–853, 857–858, 860, 863, 865–867, 870, 876, 933, 936, 939, 943–944
IP address 8, 30–31, 33–34, 42, 46–47, 49–50, 63, 78, 242, 346, 353, 358, 385, 390, 421, 424, 426, 429–430, 442, 565, 596–597, 613, 619, 629, 631, 634, 705, 711, 724, 763, 881–882, 904, 906, 927, 939–940, 966, 982, 1007, 1011
  Dynamic attributes 597
  Local swap 716
  other components 927, 940
  remote nodes 34, 50
IP app_pers_ip 809–810, 868
IP label 424–425, 429–430, 448
  1 429
  2 429
  move 424
IP network 5, 9, 427, 429, 711
ISC Help Service 31, 47, 103, 120, 181, 882, 906, 944
ISC installation
  environment 852
ISC name 120, 906
ISC service 116, 118, 167, 181, 231, 906
  default startup type 116
  name 120
  new resources 167, 231
ISCProduct.inst allLocation 466, 622–623, 843
itsosj hla 383, 1005

J

java process 492, 806

jeopardy 709

K

kanaga 427, 429–431, 437, 441
KB/sec 515, 830, 837, 876

L

lab environment 16, 29–30, 44, 46, 78, 136, 199, 249, 275, 289, 315, 372, 377, 404, 413, 528, 611, 618, 639, 654, 663, 739, 880, 904, 967, 988, 1020, 1025
Lab setup 118, 180, 455, 531, 560, 599, 619, 656, 797, 904, 967, 1001
LAN-free backup 330–331, 333, 337, 340, 342, 346–347, 350, 357–358, 366–367, 372, 378, 381, 384, 389–390, 397, 399, 403, 560, 570–571, 580, 590, 795, 797, 826, 828, 1000–1001, 1010, 1015
    high availability Library Manager functions 333, 378
    Storage Agent 330, 390
    tape volume 399
LAN-free client
    data movement 14
    incremental backup 367, 398, 1015
    system 578
LAN-free communication method 335, 379, 559, 796, 1003
lanfree connection 570, 799
    usr/tivoli/tsm/client/ba/bin/dsm.sys file 570
lanfree option 357, 366, 389, 1009
LAN-free path 329, 331, 351, 357, 365, 377, 389, 396, 412, 571, 673, 683, 699, 1001, 1009
LANFREECOMMMETHOD SHAREDMEM 356, 366, 388, 397, 1009, 1014
LANFREECOMMMETHOD TCPIP 356, 366, 388, 397, 1009, 1014
LANFREETCPPORT 1502 356, 388, 1009
Last access 660, 683, 829
last task 259, 269, 300, 309, 361, 365, 392, 396, 974, 986, 1013
Level 0.0 620, 627, 659, 680, 683
liblto device 383–384, 1005–1006
    =/dev/IBMtape0 628
    =/dev/IBMtape1 628
    =/dev/rmt1 489, 757
library inventory 163, 226, 962–963
    private volume 164, 227
    tape volume 164, 227
library liblto
    libtype 489, 628, 757
    RESETDRIVES 489
library LIBLTO1 569
library sharing 453, 688, 696, 833
license agreement 83, 94, 106, 333, 463, 611, 752, 844, 889
LIENT_NAME 826–828, 834–835, 875
Linux 12, 14, 17, 452, 454, 594–596, 598–603, 605–606, 610, 614
Linux distribution 594, 653
lla 383, 1005
lladdress 680–681, 798
local area network
    cluster nodes 9
local area network (LAN) 9, 14, 422, 999–1001, 1005–1006, 1009–1011, 1014–1015, 1019–1021, 1023, 1025
local disc 79–80, 91, 107, 252, 293, 331–332, 561, 607, 841, 909, 966, 969
    LAN-free backup 331
    local components 561
    system services 242
    Tivoli Storage Manager 909
    Tivoli Storage Manager client 969
local drive 147, 209, 252, 293, 640, 946, 969
local node 250, 265, 290, 305, 331, 333, 337, 351, 356–357, 378, 381, 384, 388–389, 654, 887, 968, 1001, 1006, 1009–1010
    configuration tasks 351
    LAN-free backup 356, 388
    local Storage Agent 357, 389
    Storage Agent 340, 383–384
    Tivoli Storage Manager scheduler service 265, 305
local resource 528, 654
local Storage Agent 352, 356–357, 388–389, 675, 794, 1009–1010, 1025
    RADON_STA 373
LOCKFILE 759–761, 805
log file 76, 600, 619, 643–645, 658, 660, 710, 715, 779, 822
LOG Mirror 20
logform command 439, 728, 730
Logical unit number (LUN) 605, 624, 721, 726
logical volume 418, 439, 441, 728, 730
login menu 173, 237
Low Latency Transport (LLT) 704, 709
lsrel command 637, 663, 686

Archived

lsrg 635–639, 661, 663, 684, 686
lssrc 483, 501, 506, 511, 515, 517, 520, 524, 536, 540, 544, 550, 578, 584, 870
lvlstmajor command 438, 440, 727, 730

M

machine name 34, 50, 613
main.cf 709
MANAGEDSERVICES option 857, 860
management interface base (MIB) 710
manpage 635, 637
manual process 160, 164, 225, 227, 649, 651, 960, 963
memory port 335, 366, 379, 397, 1003, 1014
Microsoft Cluster Server
  Tivoli Storage Manager products 25
Microsoft Cluster Server (MSCS) 25
migration process 155–156, 220–221, 517, 645–647, 953–955
mirroring 6
mklv command 439, 728, 730
mkvg command 438, 440, 727, 729
Mount m_ibm_isc 809–810, 868–869
mountpoint 619, 631, 634
MSCS environment 78–80, 118, 120, 242, 292
MSCS Windows environment 243, 332
MS-DOS 256, 258, 266, 297, 299, 306, 357, 389, 971–972, 986, 1010
Multiplatforms environment 661, 684
Multiplatforms setup 593
Multiplatforms Version 1.2
  cluster concept 593
  environment 591

N

ne 0 864
network 5
network adapter 5, 28, 33, 49, 431, 442, 597, 705, 711
  Properties tab 33, 49
Network channels 704
Network data transfer rate 515, 876
Network name 30–31, 46–47, 137, 143, 200, 202, 205, 242, 430, 448, 882, 966
Network partitions 709
Network Time Protocol (NTP) 600
next menu 138, 200, 245, 262, 271, 312, 353
Next operation 875
Next step 43, 74, 129, 191, 450, 678, 914
NIC NIC_en2 809–810, 868–869
NIM security 419
node 5
Node 1 30–31, 46–47, 335, 379, 429, 530, 559, 656, 796, 881–882, 1003
Node 2 30–31, 46–47, 335, 379, 429, 530, 559, 656, 796, 881–882, 1003
node CL_HACMP03_CLIENT 532, 536, 540, 544, 550
node CL_ITSAMP02_CLIENT 681
node CL_VERITAS01_CLIENT 828, 831, 835, 870–873, 875
  ANR0480W Session 407 875
Node Name 529, 655, 659, 683, 731, 829, 840, 969
node name
  first screen 731
nodename 250–256, 262, 275, 277, 284, 290–291, 293–297, 303, 315, 317, 323, 656, 659–660, 662, 664, 668, 682, 685, 687–688, 695–696, 968–971, 989, 993
nodenames 242, 253, 256, 297, 966, 969
Nodes 704
Nominal state 597, 637–639
non-clustered resource 528, 654
non-service address 424
Normal File 829

O

object data manager (ODM) 715
occu stg 159, 649, 958
offline medium 514, 645, 836
online resource 367, 373, 398, 406, 780–781, 823–824, 1015, 1021
Open File Support (OFS) 14, 454
operational procedures 7
option file 252–255, 293, 295–296, 528, 654, 969–971
  main difference 254, 295
output volume 368, 540, 544, 546, 548, 690, 786, 788, 790
  030AKK 870
  ABA990 786–787
  client session 546

P

password hladdress 511, 567, 569, 680, 798
physical node 253–254, 269, 294–295, 309, 842,

871, 970, 972, 985–986, 990
local name 278, 318
option file 295
same name 269, 309
separate directory 842
Tivoli Storage Manager Remote Client Agent services 266, 306
pid file 860–862
pop-up menu 176, 238
PortInput.secu reAdminPort 466, 622, 843
PortInput.webA dminPort 466, 473, 622–623, 843
primary node 465, 496, 498, 500–501, 510, 517, 519, 523, 534
cluster services 496
opt/IBM/ISC command 465
smitty clstop fast path 498
private network 705
process id 858, 861
processing time 515, 830, 837, 876
Public network
configuration 34, 49
IP address 30, 46, 881
property 72
public network 705, 711
PVIDs presence 729

Q

QUERY SESSION 494, 506, 512, 537, 540, 544, 551, 782, 825, 833, 870
ANR3605E 826, 833
Querying server 541, 829

R

RAID 5
read/write state 517, 520, 787, 790
README file 455, 744
readme file 431, 441
linux_rdac_readme 602
README.i2xLNX-v7.01.01.txt 600–601
recovery log 13, 79, 120–121, 132, 159–160, 181–182, 193, 224–225, 452, 486–488, 619, 626–627, 721, 754–756, 881, 906–907, 915, 920, 957, 959–960
Recovery RM 615
Recvd Type 782, 870, 872, 874
recvw state 541, 544, 551, 832, 874
Red Hat
Enterprise Linux 594, 599, 603
Linux 3.2.3 601
Redbooks Web site 1050
Contact us lii
Redundant Disk Array Controller (RDAC) 600, 602–604, 607, 885
register admin 489, 757
operator authority 489
Registry Key 262, 272, 303, 312, 357–358, 389–390, 982, 1010–1011
reintegration 5
Release 3 620, 627, 659, 680, 683, 829
Resource categories 706
On-Off 706
On-Only 706
Persistent 706
Resource Group
information 636
TSM Admin Center 120, 181
Resource group 712
Cascading 712
Cascading without fall back (CWOF) 712
Concurrent 713
Dynamic node priority (DNP) policy 712
node priority 712
nominal state 597
Rotating 712
resource group 713
Client Acceptor 267, 307
first node 426
first startup 634
initial acquisition 426
nominal state 637, 639
Persistent and dynamic attributes 636
resource chain 426
same name 974
same names 259, 269, 300, 309
same options 269, 309
scheduler service 257, 298
unique IP address 254, 296
web client services 253, 294
resource group (RG) 8–9, 23, 253–254, 256–257, 259, 265, 267, 269, 273–274, 294, 296–298, 300, 305, 307, 309, 314, 361, 365, 392, 396, 421, 424, 426, 478–479, 484, 496, 528–529, 535, 540, 544, 550, 562, 597, 618–619, 629–637, 639, 641, 643, 646–648, 650–651, 654–655, 657, 660–664, 668, 684, 686, 695, 707, 712–713, 715–717, 773, 777, 814, 820, 857, 859, 870, 880, 969–970, 972, 974, 985–986

Archived

resource online 154, 172, 235, 365–366, 396–397, 705–706, 810, 952, 984, 988
resource type 168, 232, 260, 270, 301, 310, 362, 393, 705–706, 711
multiple instances 705–706, 711
VERITAS developer agent 705
resources online 144–145, 151, 155, 157, 162, 165, 206, 213, 217, 264, 278, 304, 324, 370, 400, 486, 665, 669, 711, 713, 770, 930, 954, 956, 962, 985, 991, 995, 1017
Result summary 510, 515, 517, 519, 523, 526, 539, 543, 550, 554, 584, 590, 787, 792, 830, 837, 872, 876
Results summary 149, 154, 156, 160, 164, 167, 210, 215, 219, 221, 224, 227, 231, 283, 289, 322, 326, 372, 377, 404, 412, 645, 647, 649–650, 652, 667, 672, 694, 699, 948, 952, 955, 960, 963, 992, 997, 1019, 1025
Return code 154, 215, 218–219, 435, 494, 762, 951–952
rm archive.dsm 486, 626, 755
roll-forward mode 160, 225, 960
root@diomede bin 626, 680–681
root@diomede linux-2.4 601
root@diomede nfsserver 659–661, 683–684
root@diomede root 605, 609–610, 613–616, 624, 626–627, 631, 635–639, 661, 680, 684
rootvg 757, 804, 857, 863
rw 439, 728, 730

S

same cluster 80, 118, 179, 243, 248, 289, 332–333, 378
same command 166, 226, 230, 259, 300, 436, 650–651, 974
same name 133, 171, 195, 234, 257, 260, 269–270, 298, 301, 309–310, 346, 436, 972, 974, 986
same process 91, 140, 145, 172, 202, 206, 235, 268, 308, 351, 749, 909
same result 150, 154, 210, 215, 642, 645, 714, 948, 952
same slot 35, 51
same tape
drive 606
volume 155, 220, 373, 405, 954
same time 91, 367, 374, 398, 406, 409, 586, 688, 696, 713, 1016, 1021
multiple nodes 713
same way 630, 653, 675, 909
Clustered Storage Agent 675
second server 630
SAN Device
Mapping 611
SAN path 344, 367–368, 371–373, 399, 402, 404, 694, 1015, 1017, 1019, 1021
SAN switch 436, 561, 725
SA-tsmserver-rg 619, 632, 635–639, 641, 645–651
schedlogretention 7 255, 296–297, 971
Schedule log 277–278, 281–284, 286, 288, 317–318, 321–325, 643–645, 687, 692, 694–695, 698, 989, 992, 994–995, 997
file 151, 153–154, 213–214, 216–219, 283, 368–369, 372, 374, 376–377, 399, 402, 407, 411, 644, 664–665, 667–669, 671, 951, 994–995, 1017, 1019, 1022, 1024
Schedule Name 289, 326, 536, 825, 829, 831, 836, 873, 875
schedule webclient 532–533, 658–659
scheduled backup 24, 150, 211, 509, 642–643, 645, 664, 689, 948, 950–952, 960
scheduled client
backup 23, 150, 211, 642, 948
incremental backup 367, 540, 543, 1015
selective backup operation 536
scheduled command 279, 319, 876, 991
scheduled event 13, 154, 280, 320, 452, 645, 829–830, 876, 952
scheduled operation 286, 288–289, 324, 326, 377, 412, 510, 541, 544, 550, 584, 669, 672, 700, 830–831, 872–873, 995, 997, 1025
Tivoli Storage Manager server 326
scheduled time 216, 367, 374, 398, 406, 643, 664, 668, 687, 695, 1016, 1021
scheduler service 250–251, 253–254, 257–260, 262, 264–266, 270, 277, 283, 286, 290–291, 294–295, 298–302, 304–305, 310, 322, 324, 357, 361, 366–367, 370, 372, 375, 389, 392, 397–398, 400, 404, 408, 968–969, 972, 974, 985, 992, 995, 1009, 1014–1015, 1019, 1023
SCHEDULEREC OBJECT
End 837, 876
SCHEDULEREC Object 829–830, 836, 876
SCHEDULEREC QUERY End 875
SCHEDULEREC STATUS
End 830, 837, 876
SCHEDULEREC Status 830, 837, 876

scratch volume
021AKKL2 159
023AKKL2 957
SCSI address 605–607, 613
host number 607
only part 607
SCSI bus 370, 372, 376–377, 401, 404–405, 408, 413, 695, 1020, 1023, 1025
scsi reset 489, 556, 573, 582, 633, 683
second drive 150, 154, 210, 215, 350, 373, 405, 802–803, 948, 952, 1018, 1020
new tape volume 373, 405
second node 42, 67, 91–92, 116–118, 123, 139–140, 154, 156, 158–160, 164, 167, 184, 201–202, 205, 209, 219, 221, 223–224, 227, 231, 248, 259, 265, 269, 274, 283, 289, 300, 309, 314, 322, 326, 333, 365, 370, 372, 375, 377, 396, 401, 404, 409, 412, 435, 439–441, 445, 448, 464, 623–624, 641–642, 646–651, 668, 672, 675, 687, 691, 695, 698–699, 729, 731, 826, 842, 871, 887, 909, 919–920, 947, 952, 955, 957, 959–960, 963, 974, 985–986, 991–992, 995, 997, 1017, 1019–1020, 1025
Administration Center 116–117
Configuring Tivoli Storage Manager 919
diskhbvg volume group 441
incremental backup 209
initial configuration 140, 203
ISC code 116
local Storage Agent 675
PVIDs presence 439
same process 91
same tasks 333
scheduler service restarts 372, 404
scheduler services restarts 283, 322
server restarts 160, 224
Tivoli Storage Manager 139, 201–202
Tivoli Storage Manager restarts 209
tsmvg volume group 440
volume group tsmvg 729
Serv 825, 828, 832, 835–836
Server
enhancements, additions and changes 13, 453
server code
filesets 455, 744
installation 496
Server date/time 660, 683
server desttype 383, 489, 628, 757, 1005–1006
server instance 134, 140, 196, 626, 645, 647, 649–651, 909, 914, 916–920, 950, 952–956, 960, 962–963
server model 433–434
Server name 78, 120–121, 133, 181–182, 195, 337, 339–340, 493, 562–563, 677, 906–907, 917, 1004
server stanza 487, 492, 528, 533, 577, 626, 630, 659, 682, 755, 799, 841
server TSMSRV03 675, 681, 683, 798, 800, 825, 829
Server Version 5 531, 660, 683
Server Window Start 829, 836, 875
servername 532–533, 658–659, 680–681, 797–799, 805, 841
SErvername option 680, 759–760, 805
serverpassword password 340, 383, 1005
server-to-server communication 133, 195, 337, 381, 560, 562, 797, 917, 1004
Server password 337, 381
Service Group 23, 706–710, 712, 716–717, 720, 743, 753, 757, 763, 766–767, 770, 772–773, 775–778, 780–781, 785–786, 789–790, 811, 813, 817–818, 820, 840, 842, 857, 865, 867, 869, 871, 882, 920, 922, 933, 935, 966, 968, 970–972, 974, 976–977, 983–984, 986, 1001
configuration 865, 920, 933, 974
critical resource 822
IP Resource 763
manual switch 790
name 974, 986
NIC Resource 763
OFFLINE 817
sg_isc_sta_tsmcli 866–867
sg_tsmsrv 779, 822
switch 775, 817, 819
Service group 706
service group
new generic service resource 986
new resource 974
scheduler service 972
service group dependencies 707
service group type 706
Failover 706
Parallel 707
service name 120, 171, 181, 234, 250–251, 259, 262, 269, 271, 290–292, 300, 302, 309, 312, 335, 379, 924, 936, 939, 968, 974, 978, 986
serviceability characteristic 12, 452
set servername 340, 353, 383, 386

setupISC 465–466, 843
sg_isc_sta_tsmcli 798, 808, 810, 814, 817, 820–821, 865, 867–869
manual online 814
sg_tsmsrv 720, 763, 767, 769, 773, 775, 777–779, 783–784, 814, 817, 820–821, 823
potential target node 783
sg_tsmsrv Service Group
IP Resource 763
Mount Resource 764
Shared disc
Tivoli Storage Manager client 969
shared disc 32, 35–36, 40–41, 48, 51, 53, 58–59, 92, 167, 231, 242–243, 253–254, 294–295, 319, 331, 422, 454, 464, 469, 487, 528, 532, 534, 618–619, 621, 623–624, 626, 654, 658, 754, 756, 795, 798, 807, 840–842, 846, 884, 886, 920, 966, 969–970
also LAN-free backup 331
new instance 486
own directory 841
Shared external disk devices 704, 711
shared file system
disk storage pools files 627
shared resource 79, 146, 208, 275, 315, 367, 398, 639, 663, 945, 985, 988, 990, 1015, 1019
Shell script 629, 758–760, 804–805
simple mail transfer protocol (SMTP) 710, 715
simple network management protocol (SNMP) 710, 714–715
single point 4–7, 16–17, 423, 445, 704, 711, 843
single points of failure 4
single points of failure (SPOF) 4
single server 909
small computer system interface (SCSI) 704, 711, 718
Smit panel 458–459
smit panel 436, 459, 534, 747–748
smitty hacmp
fast path 481, 493
panel 501–503
SNMP notification 739
software requirement 136, 198, 422, 431, 599
split-brain 709
SPOF 4, 6
SQL query 574, 662, 685
STA instance 558, 674, 795
Start Date/Time 536, 825, 831, 873
start script 490, 493, 535, 546, 548, 551, 571–573, 577, 580, 582, 586–587, 757, 804, 826, 828, 833, 857–858, 863, 871, 874
StartAfter 598
StartCommand 654, 661–662, 684–685
startup script 528, 546
startup window 219, 279, 283, 286, 289, 319, 322, 324, 326, 372, 377, 404, 412, 584, 665, 668–669, 672, 695, 699, 991–992, 995, 997, 1019, 1023, 1025
stop script 487, 489, 491, 493, 535, 577–578, 758, 805, 859, 861, 864
Storage Agent 13, 15–16, 329–335, 337, 339–341, 345, 348, 351–358, 360, 365–372, 374–376, 378–379, 381, 383–385, 387–389, 392, 396–401, 404–410, 453, 489, 511–512, 514–515, 555–562, 564–565, 567–574, 577–580, 582–587, 590, 599–600, 614, 673–675, 677, 679, 681–684, 686–688, 690, 692, 694, 696–697, 793–799, 803–805, 807, 824, 826, 828, 833, 835, 841, 999–1000, 1002–1013, 1015–1023
appropriate information 352
CL_ITSAMP02_STA 690
CL_MSCS01_STA 368, 375
CL_MSCS02_STA 400
CL_VCS02_STA 1023
Configuration 331, 339, 383
configuration 331, 378, 798
correct running environment 572, 574
detail information 331
dsm.sys stanzas 799
high availability 398
high level address 335, 379, 559, 796, 1003
information 385
Installation 331–332
instance 357, 389, 558, 562, 573, 1010
Library recovery 514
local configuration 675
low level address 335, 379, 559, 796, 1003
name 335, 358, 372, 379, 385, 390, 405, 559, 796, 1003, 1011, 1020
new registry key 357, 389
new version 15
port number 346
related start 562
Resource configuration 683
server name 677
Server object definitions 564
service name 335, 379, 1003
software 331

Archived

successful implementation 365, 396
Tape drive device names 337
User 556, 560, 794, 797
Windows 2003 configuration 1002
Storage agent
CL_MSCS02_STA 401, 408
Storage Agents 705
Storage Area Network (SAN) 14, 16, 122, 329–330, 381, 673, 704, 794, 797, 824
Storage Area Networks
IBM Tivoli Storage Manager 674
Storage Certification Suite (SCS) 704
storage device 13–14, 35, 51, 330–331, 337, 452, 611, 800, 884, 945, 1000–1001, 1004
Windows operating system 331
Storage Networking Industry Association (SNIA) 602, 611
storage pool 79, 120–121, 132, 150–151, 154–160, 181–182, 193, 210, 215, 219–225, 340, 347, 373, 384, 405, 454, 487, 536, 540, 543, 627, 645, 649, 663, 755, 787, 907, 909, 915, 919, 959, 1006, 1020
backup 12, 150, 154–156, 222, 452, 517, 647–648, 789, 955
backup process 648
backup task 955
current utilization 787
file 488, 756, 920
SPC_BCK 518
SPCPT_BCK 156, 159, 222, 647
SPD_BCK 786–787
volume 625
storageagent 21, 24, 332, 335, 351, 357, 360, 379, 384, 389, 392, 558–559, 562, 569–570, 793, 795–796, 798–799, 804–806, 1003, 1006, 1010, 1012
subsystem device driver (SDD) 607
supported network type 705
SuSE Linux Enterprise Server (SLES) 594, 599
symbolic link 609–610, 661, 674, 684
Symmetric Multi-Processor (SMP) 601
sys atlantic 763–765, 807, 865–866
sys banda 763–765, 772, 775, 807, 813, 817, 865–866
system banda 771, 773, 775, 777–780, 783, 812, 814, 817, 820–822, 824
group sg_isc_sta_tsmcli 820
group sg_tsmsrv 777
System Management Interface Tool (SMIT) 709, 714–716
System zones 708
systemstate systemservices 290–291, 968

T

tape device 122, 136, 183, 198, 593, 605, 611, 629, 633, 725
shared SCSI bus 136, 198
Tape drive
complete outage 633
tape drive 79, 122, 184, 331, 337, 339, 348, 350, 381–382, 489, 517, 556, 560–561, 567, 580–581, 590, 606, 611, 628–629, 633, 651, 674, 691, 698, 794, 797, 908, 960, 990, 1001, 1004
configuration 489, 756
device driver 331
Tape Library 122, 155, 183, 220, 337, 339–340, 350, 367–368, 374, 381, 383, 398, 406, 489, 556, 606, 628, 724, 756, 794, 908, 953, 956, 959, 962, 989, 1004–1005, 1016, 1021
scratch pool 159, 224
second drive 350
Tape Storage Pool 121, 150, 154–156, 159, 182, 210, 215, 220–222, 224, 515, 517, 571, 633, 642, 645, 647–649, 663, 785, 787, 907, 948, 952–953, 955, 957
Testing backup 955
tape storage pool
Testing backup 647
tape volume 150, 154–156, 158–161, 163–164, 210, 215, 220–221, 224, 226–227, 367–368, 517, 519–520, 523, 582, 642, 645–646, 649, 651, 688, 690–693, 695–696, 698, 787, 790, 792, 948, 952–954, 956–963, 989–991, 993–995, 1016, 1018–1023
027AKKL2 962
028AKK 368
030AKK 690
status display 962
Task-oriented interface 12, 452
TCP Address 870, 875
TCP Name 828, 835, 871, 875
TCP port 177, 239, 254, 296, 353, 386, 970, 1007
Tcp/Ip 487, 704, 711, 755, 782, 784, 795, 825–828, 832–833, 835–836, 870–872, 874–875
TCP/IP address 346, 678
TCP/IP connection 645
TCP/IP property 33–34, 49–50
following configuration 33, 49

TCP/IP subsystem 5
tcpip addr 529, 558, 655, 674, 795, 840
tcpip address 529, 557, 563, 655
tcpip communication 557, 795
tcpip port 529, 558, 655, 674, 795, 840
TCPPort 1500 626
TCPPort 1502 569, 799
tcpserveraddress 9.1.39.73 255
tcpserveraddress 9.1.39.74 296, 971
test result 154, 215, 219, 283, 289, 322, 326, 372, 377, 404, 412, 645, 667, 672, 694, 699, 771, 811, 952, 992, 997, 1019, 1025
  historical integrity 771
test show 156, 160, 164, 167, 221, 224, 227, 231, 647, 649–650, 652, 955, 960, 963
testing 7
Testing backup 156, 221
Testing migration 154, 219, 645
tivoli 241–245, 248–249, 252–254, 256–260, 264–269, 274–279, 281–290, 292–295, 297–300, 304–309, 314–327, 329–333, 335, 337, 340–341, 350–351, 353, 356–357, 359–361, 365–373, 376–379, 381, 383–386, 389–390, 392, 396–398, 400–401, 404–408, 412–413, 451–456, 460, 464–465, 472, 478, 482, 486–490, 493, 495, 506–507, 510, 512, 514, 517, 519, 524, 903–905, 908–909, 911, 913–916, 918–920, 923, 925, 927, 933, 940, 945–950, 952–963, 965–966, 968–972, 974, 985–986, 988–990, 992–993, 995–997, 999–1001, 1003, 1005–1021, 1023, 1025
Tivoli Storage Manager (TSM) 242, 256, 297, 327, 451–456, 458, 460, 462, 464, 472, 478–480, 482, 486–490, 493, 495, 505, 507, 510, 513, 515–516, 518–520, 523, 526, 673–675, 679–683, 685–688, 690–691, 694, 696, 699, 743–745, 749, 753–757, 759–760, 762–763, 779, 781–782, 784–787, 789–790, 792
Tivoli Storage Manager Administration
  Center 453, 621, 624, 842
Tivoli Storage Manager Backup-Archive client 327
Tivoli Storage Manager Client
  Accepter 274, 314
  Acceptor CL_VCS02_ISC 987
  Acceptor Daemon 660
  Acceptor Polonium 252
  Acceptor Tsonga 293
  configuration 653, 657, 660
  Installation 531
  test 24
  Version 5.3 531
Tivoli Storage Manager client 241–246, 248, 252–253, 256, 258–259, 266, 270, 273–276, 284, 289, 292–294, 297–300, 306, 310, 314–316, 319, 323, 326–327, 653–655, 657–658, 660–663, 672, 681, 683–684, 839–842, 857, 867, 870, 873, 875
  acceptor service 266, 274, 306, 314, 986
  Cad 661, 684
  code 654
  command 831
  component 243
  directory 971
  environment variable 528, 654
  installation path 266, 306, 972, 986
  log 553
  node 528–529, 532
  node instance 529
  requirement 529
  resource 314, 654
  scheduler 572, 577
  service 274, 306, 314
  software 242, 289
  V5.3 527, 653, 657
Tivoli Storage Manager command
  line client 245
  q session 871
Tivoli Storage Manager configuration
  matrix 20
  step 629
  wizard 909
Tivoli Storage Manager database
  backup 791
Tivoli Storage Manager Group
  resource 143
Tivoli Storage Manager scheduler
  resource 373, 406
  service 257, 259, 265, 300, 357, 366, 389, 397, 1009, 1014
  service resource 305
Tivoli Storage Manager scheduler service 969, 972, 974, 992, 995
  installation 257, 298
  resource 300, 667, 669
Tivoli Storage Manager Server
  cluster 619
  resource 629, 633
  test 23
  V5.3 629, 657
Tivoli Storage Manager server 77–82, 86, 118, 120,

122–123, 129–132, 135, 139–140, 143, 145–147, 149–152, 154–160, 162, 164–167, 173, 175, 177, 179, 183–184, 191–194, 197, 201–202, 205, 207–211, 213, 215–217, 219–221, 223–225, 227–228, 230–231, 236, 238–240, 556–557, 561–564, 567, 571, 573–574, 577, 580–581, 584, 586, 617–620, 624–630, 633–635, 637–640, 642–645, 647, 649–652
Atlantic 782
Tivoli Storage Manager V5.3
address 17
server software 743
Tivoli System Automation 593–600, 606–607, 611–612, 614–615, 617–618, 621, 623–625, 629–631, 633–635, 653, 656, 673, 675, 684, 686
cluster 596, 598–599, 633
cluster application 661
configuration 635
decision engine 615
environment 618, 629, 654
fixpack level 612
Highly available NFS server 656, 661
Installation 593, 600
installation 657
manual 596
many resource policies 614
necessary definition files 634
NFS server 661
resource 661
Storage Agent 684
terminology 596
tie breaker disk 624
Tivoli Storage Manager client CAD 661
v1.2 596, 653, 657
v1.2 installation 657
Total number 515, 708, 830, 837, 876
trigger 710
TSM Admin Center 31, 47, 251, 265, 268, 274, 291, 294–295, 305, 308, 314, 492, 863
cluster group 167, 231
group 253, 255, 258, 296, 299
resource group 305, 314
Tivoli Storage Manager Client Acceptor service resource 314
TSM client 31, 47, 369–370, 400–401, 863
TSM Group 31, 43, 47, 120, 122, 124, 130, 140–141, 143–145, 181, 183, 185, 192, 203, 205–206, 251, 253, 255, 259, 265, 268, 274, 292, 294–296, 299, 305, 308, 314, 351, 357–358, 361, 366–367, 384, 389–390, 392, 397–398, 926, 1006, 1010
Cluster Administrator 122, 183
IP Address 206
IP address 358, 390
network name resources 140
Option file 256, 297
resource 206
scheduler service 259, 299
Server 143, 205
Tivoli Storage Manager scheduler service 366, 397
TSM Remote Client Agent
CL_MSCS01_QUORUM 251, 267
CL_MSCS01_SA 251, 268
CL_MSCS01_TSM 251, 269
CL_MSCS02_QUORUM 291, 307
CL_MSCS02_SA 291, 308
CL_MSCS02_TSM 292, 309
CL_VCS02_ISC 968, 986
Ottawa 968
Polonium 250
Radon 250
Salvador 968
Tsonga 291
TSM Scheduler 357, 366–367, 373, 389, 397–398, 406, 1010, 1013, 1015, 1021
CL_MSCS01_QUORUM 251, 258
CL_MSCS01_SA 251, 258, 265, 277–278
CL_MSCS01_TSM 251, 259, 265
CL_MSCS02_QUORUM 291, 299
CL_MSCS02_SA 291, 299, 305
CL_MSCS02_TSM 291, 299, 305
CL_MSCS02_TSM resource 318
CL_VCS02_ISC 968, 973, 978, 985
CL_VCS02_ISC service 978, 985
Ottawa 968
Polonium 250
Radon 250
resource 365–366, 396
Salvador 968
Senegal 290
Tsonga 291
TSM scheduler
service 254, 259, 295, 300, 357, 389, 1010, 1013
TSM Server
information 337, 1004
TSM server 31, 47, 82, 86, 120, 181, 207, 340, 369,

Archived

399, 430, 489, 516, 619, 626, 630–631, 634, 758–761, 804–805, 822, 882, 906, 918
TSM Storage Agent2 359, 390
TSM StorageAgent1 335, 359–360, 379, 387, 390, 392, 1003, 1009, 1011–1013
TSM StorageAgent2
generic service resource 362, 393
TSM userid 759–760, 805
TSM.PWD file 534, 659, 682, 842
tsm/lgmr1/vol1 1000 488, 756
tsmvg 438, 440, 727, 729
types.cf 709

U

Ultrium 1 560, 797
URL 472, 849, 856
user id 97, 110–111, 115, 132, 174, 194, 237, 341, 469, 492, 630, 659, 683, 916
usr/sbin/rsct/sapolicies/bin/getstatus script 645, 647, 649, 651, 664, 668, 687, 695
usr/tivoli/tsm/client/ba/bin/dsm.sys file 570, 759–760, 799, 805, 841

V

var/VRTSvcs/log/engine_A.log output 779–780, 822, 824
varyoffvg command 439–441, 728, 730
varyoffvg tsmvg 439–440, 728–729
VCS cluster
engine 713
network 704
server 711
software 731
VCS control 840
VCS WARNING V-16-10011-5607 779, 822
VERITAS Cluster
Helper Service 899
Server 703–704, 706–707, 710, 716, 718–720, 734, 740, 753, 793, 810, 839, 880, 887, 896, 902–903
Server 4.2 Administrator 902
Server Agents Developers Guide 705
Server environment 719
Server feature comparison summary 716
Server User Guide 707, 709
Server Version 4.0 infrastructure 701, 877
Server Version 4.0 running 701
Services 415
Veritas Cluster
Explorer 972, 985, 989, 991, 993, 995
Manager 757, 770, 945, 949–950, 953–956, 961, 1015, 1017
Manager configuration 857
Manager GUI 869
Server 1030
VERITAS Cluster Server 704
Veritas Cluster Server
Version 4.0 877
VERITAS Enterprise Administration (VEA) 887
VERITAS Storage Foundation
4.2 887
Ha 879–880
video
command line access 1029
unlock client node 1029
virtual client 150, 266, 276, 306, 316, 322, 372, 377, 405, 407, 413, 985, 1020, 1025
opened session 407
virtual node 251, 253–254, 284, 291, 294–295, 331, 333, 357, 378, 389, 530, 559, 656, 664, 668, 687, 695, 796, 841, 968–970, 989, 993, 1001, 1010
Storage Agent 357, 361
Tivoli Storage Manager Client Acceptor service 274
Tivoli Storage Manager scheduler service 265, 305
Web client interface 254, 295
Volume Group 418, 430, 438–441, 480, 720, 727–730, 865
volume spd_bck 489, 628, 756
vpl hdisk4 438, 727

W

web administration port
menu display 108
web client
interface 254, 295, 859, 970
service 253, 269, 294, 969, 985–986
Web material 1029
Web Site 1029
Web VCS interface 707
Web-based interface 92, 933
Windows 2000 25, 27–29, 31–32, 35, 41–42, 44, 79, 118, 122, 146, 167, 241–243, 248, 252, 262, 272, 275, 292, 327, 329, 331–333, 337, 339, 349, 367

Archived

IBM 3580 tape drive drivers 337
IBM tape device drivers 122
Windows 2000 MSCS 77, 79, 91, 118, 120, 242, 337, 946
Windows 2003 27–28, 44, 47, 51, 59, 61, 74, 79, 92, 179, 183, 208, 231, 241–243, 248, 289, 292, 303, 312, 315, 329, 331–332, 378, 381, 383, 398, 704, 879–882, 885, 999–1001
IBM 3580 tape drive drivers 381
IBM tape device drivers 183
Tivoli Storage Manager Client 242
Windows 2003 MSCS
setup 48
Windows environment 92, 879
clustered application 92

X

X.25 and SNA 711

Archived

Archived

Archived

Archived

IBM

Redbooks

IBM Tivoli Storage Manager in a Clustered Environment

(2.0” spine)
2.0” <-> 2.498”
1052 <-> 1314 pages

Archived

IBM Tivoli Storage Manager in a Clustered Environment

Learn how to build highly available Tivoli Storage Manager environments

Covering Linux, IBM AIX, and Microsoft Windows solutions

Understand all aspects of clustering

This IBM Redbook is an easy-to-follow guide that describes how to implement IBM Tivoli Storage Manager Version 5.3 products in highly available clustered environments.

The book is intended for those who want to plan, install, test, and manage the IBM Tivoli Storage Manager Version 5.3 in various environments by providing best practices and showing how to develop scripts for clustered environments.

The book covers the following environments: IBM AIX HACMP, IBM Tivoli System Automation for Multiplatforms on Linux and AIX, MicrosoftCluster Server on Windows 2000 and Windows 2003, VERITAS Storage Foundation HA on AIX, and Windows Server 2003 Enterprise Edition.

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

SG24-6679-00 ISBN 0738491144

Archived