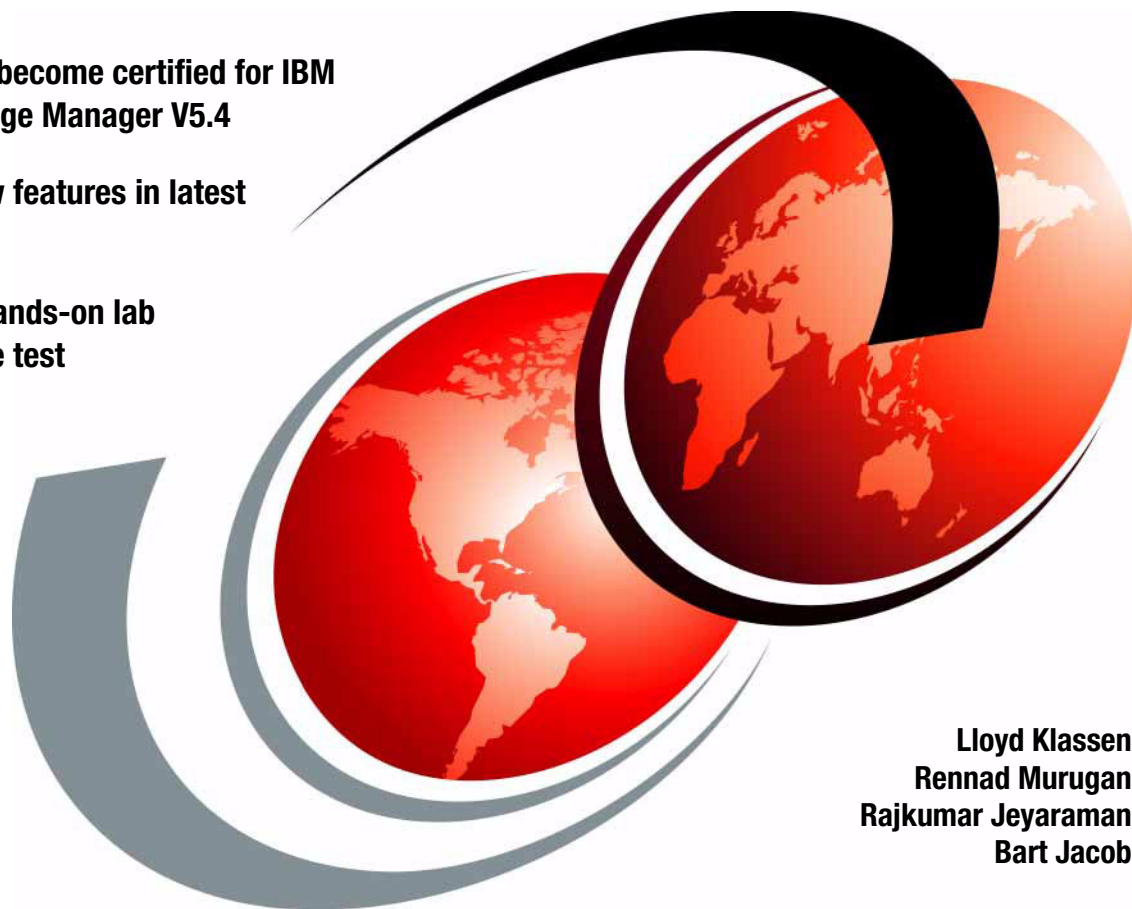**Tivoli**® software

IBM

Certification Study Guide: IBM Tivoli Storage Manager 5.4

Prepare to become certified for IBM Tivoli Storage Manager V5.4

Covers new features in latest version

Contains hands-on lab and sample test

Lloyd Klassen
Rennad Murugan
Rajkumar Jeyaraman
Bart Jacob

Redbooks

ibm.com/redbooks

International Technical Support Organization

Certification Study Guide: IBM Tivoli Storage Manager 5.4

June 2007

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

First Edition (June 2007)

This edition applies to Version 5, Release 4 of IBM Tivoli Storage Manager.

This document created or updated on May 22, 2007.

Contents

Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

ix

Trademarks

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

Redbooks (logo) ®	DS8000™	POWER5™
Eserver®	Enterprise Storage Server®	Redbooks®
iSeries®	FlashCopy®	RISC System/6000®
i5/OS®	GPFS™	RS/6000®
pSeries®	HACMP™	S/390®
z/OS®	Informix®	System i™
z/VM®	IBM®	System p™
zSeries®	Lotus Notes®	System z™
AIX 5L™	Lotus®	System Storage™
AIX®	MVS™	SysBack™
AS/400e™	NetView®	SANergy®
AS/400®	Notes®	Tivoli Enterprise™
Domino®	OpenPower™	Tivoli Enterprise Console®
DB2 Universal Database™	OS/390®	Tivoli®
DB2®	OS/400®	TotalStorage®
DPI®	PartnerWorld®	WebSphere®
DS6000™	POWER™	

The following terms are trademarks of other companies:

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

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

Snapshot, Network Appliance, Data ONTAP, and the Network Appliance logo are trademarks or registered trademarks of Network Appliance, Inc. in the U.S. and other countries.

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

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

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

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

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

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

Preface

This IBM Redbooks publication is a study guide for IBM Tivoli® Storage Manager Version 5.4 and is aimed at individuals who want to get IBM Certifications in this specific product.

The IBM Tivoli Storage Manager Certification, offered through the Professional Certification Program from IBM, is designed to validate the skills required of technical professionals who work in the implementation of the IBM Tivoli Storage Manager Version 5.4 product.

This book provides a combination of theory and practical experience needed for a general understanding of the subject matter. It also provides sample questions that will help in the evaluation of personal progress and provide familiarity with the types of questions that will be encountered in the exam.

This publication does not replace practical experience, nor is it designed to be a stand-alone guide for any subject. Instead, it is an effective tool that, when combined with education activities and experience, can be a very useful preparation guide for the exam.

The team that wrote this book

Figure 1   l. to r. Lloyd Klassen, Rennad Murugan, Bart Jacob (Rajkumar Jeyaraman not pictured)

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

Lloyd Klassen is an IT Specialist with IBM Integrated Technology Delivery Canada. He has over 7 years of experience designing, implementing and supporting backup / recovery and SAN environments. His area of expertise is with Tivoli Storage Manager, IBM disk and tape subsystems, SAN Volume Controller, Brocade and Cisco SAN's. He is an IBM certified Tivoli Storage Manager Deployment Professional.

Rennad Murugan is an IT Specialist for the Integrated Delivery Centre in South Africa, supporting IBM local and International accounts.He has been working for IBM since 2004. His area of expertise includes 24 by 7 Operational support, Tivoli Storage Manager planning, implementation, troubleshooting, as well as playing a role in IBM Storage and Data Protection solutions. He is an IBM Certified Deployment Professional: Tivoli Storage Manager V5.3.

Rennad holds a Bachelor of Commerce degree with majors in Information Systems and Management from Rhodes University in Grahamstown, South Africa.

Rajkumar Jeyaraman is a Senior IT consultant with Cognizant Technology Solutions, USA. He has over 9 years of experience in managing and administering IT infrastructure for various clients of Cognizant. His area of expertise includes implementation and administration of various mission critical applications running on IBM AIX, MS Windows, and Solaris platforms. He has also worked extensively on storage management products like IBM Tivoli Storage Manager and ARCServe. He is a certified Tivoli Storage Management Deployment Professional.

Rajkumar holds a Master of Science degree in Software Engineering from BITS Pilani, India.

Bart Jacob is a Senior Consulting IT Specialist at IBM Corp - International Technical Support Organization, Austin Center. He has over 25 years of experience providing technical support across a variety of IBM products and technologies, including communications, object-oriented software development, and systems management. He joined the ITSO in 1989, where he has been writing IBM Redbooks™ and creating and teaching workshops around the world on a variety of topics. He holds a Masters degree in Numerical Analysis from Syracuse University.

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

Charlotte Brooks
International Technical Support Organization, San Jose Center

Howard Martin,
TSM Development, IBM Tucson

Judy Green
Tivoli Certification, IBM Austin

Become a published author

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

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

ibm.com/redbooks

► Send your comments in an e-mail to:

redbooks@us.ibm.com

► Mail your comments to:

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

1

Certification overview

This chapter provides an overview of the skill requirements needed to obtain an IBM Advanced Technical Expert certification. The following chapters are designed to provide a comprehensive review of specific topics that are essential for obtaining the certification:

► IBM Professional Certification Program
► IBM Tivoli® Storage Manager 5.4 Implementation Certification
► Recommended study resources

     1

1.1 IBM Professional Certification Program

Having the right skills for the job is critical in the growing global marketplace. IBM Professional Certification, designed to validate skill and proficiency in the latest IBM solution and product technology, can help provide that competitive edge. The IBM Professional Certification Program Web site is available at:

http://www.ibm.com/certify/index.shtml

The Professional Certification Program from IBM offers a business solution for skilled technical professionals seeking to demonstrate their expertise to the world.

The program is designed to validate your skills and demonstrate your proficiency in the latest IBM technology and solutions. In addition, professional certification may help you excel at your job by giving you and your employer confidence that your skills have been tested. You may be able to deliver higher levels of service and technical expertise than non-certified employees and move on a faster career track. Professional certification puts your career in your control.

This is the way for skilled IT professionals to demonstrate their expertise to the world. It validates your skills and demonstrates your proficiency in the latest IBM technology and solutions.

The certification requirements are tough, but not impossible. Certification is a rigorous process that differentiates you from everyone else.

The mission of IBM Professional Certification is to:

► Provide a reliable, valid, and fair method of assessing skills and knowledge.
► Provide IBM with a method of building and validating the skills of individuals and organizations.
► Develop a loyal community of highly skilled certified professionals who recommend, sell, service, support, or use IBM products and solutions.

The Professional Certification Program from IBM has developed certification role names to guide you in your professional development. The certification role names include IBM Certified Specialist, IBM Certified Solutions/Systems Expert, and IBM Certified Advanced Technical Expert for technical professionals who sell, service, and support IBM solutions.

For technical professionals in application development, the certification roles include IBM Certified Developer Associate and IBM Certified Developer. IBM Certified Instructor certifies the professional instructor.

The Professional Certification Program from IBM provides you with a structured program leading to an internationally recognized qualification. The program is designed for flexibility by allowing you to select your role; prepare for and take tests at your own pace; and, in some cases, select from a choice of elective tests best suited to your abilities and needs. Some roles also offer a shortcut by giving credit for a certification obtained in other industry certification programs.

You may be a network administrator, systems integrator, network integrator, solution architect, solution developer, value-added reseller, technical coordinator, sales representative, or educational trainer. Regardless of your role, you can start charting your course through the Professional Certification Program from IBM today.

1.1.1 Benefits of certification

Certification is a tool to help objectively measure the performance of a professional on a given job at a defined skill level. Therefore, it is beneficial for individuals who want to validate their own skills and performance levels, their employees, or both. For optimum benefit, the certification tests must reflect the critical tasks required for a job, the skill levels of each task, and the frequency by which a task needs to be performed. IBM prides itself on designing comprehensive, documented processes that ensure that IBM certification tests remain relevant to the work environment of potential certification candidates.

In addition to assessing job skills and performance levels, professional certification may also provide such benefits as:

► For employees:

 – Promotes recognition as an IBM certified professional
 – Helps to create advantages in interviews
 – Assists in salary increases, corporate advancement, or both
 – Increases self esteem
 – Provides continuing professional benefits

► For employers:

 – Measures the effectiveness of training
 – Reduces course redundancy and unnecessary expenses
 – Provides objective benchmarks for validating skills
 – Makes long-range planning easier
 – Helps to manage professional development
 – Aids as a hiring tool
 – Contributes to competitive advantage
 – Increases productivity, morale, and loyalty

► For Business Partners and consultants:

  – Provides independent validation of technical skills

  – Creates competitive advantage and business opportunities

  – Enhances prestige of the team

  – Contributes to IBM requirements for various IBM Business Partner programs

Specific benefits may vary by country (region) and role. In general, after you become certified, you should receive the following benefits:

► Industry recognition

  Certification may accelerate your career potential by validating your professional competency and increasing your ability to provide solid, capable technical support.

► Program credentials

  As a certified professional, you receive an e-mail with your certificate of completion and the certification mark associated with your role for use in advertisements and business literature. You may also request a hardcopy certificate, which includes a wallet-size certificate.

  The Professional Certification Program from IBM acknowledges the individual as a technical professional. The certification mark is for the exclusive use of the certified individual.

► Ongoing technical vitality

  IBM Certified professionals are included in mailings from the Professional Certification Program from IBM.

1.1.2 Tivoli Software Professional Certification

Tivoli's professional certification program offers certification testing that sets the standard for qualified product consultants, administrators, architects, and partners.

The program also offers an internationally recognized qualification for technical professionals who are seeking to apply their expertise in today's complex business environment. The program is designed for those who implement, buy, sell, service, and support Tivoli solutions and wish to deliver higher levels of service and technical expertise.

Whether you are a Tivoli customer, partner, or technical professional wishing to put your career on the fast track, you can start your journey to becoming a Tivoli Certified Professional today.

Benefits of being Tivoli certified

Tivoli Certification has the following benefits:

▶ For the individual:

 – IBM Certified certificate and use of logos on business cards

> **Note:** Certificates are sent by e-mail; however, a paper copy of the certificate and a laminated wallet card can also be requested by sending an e-mail to certify@us.ibm.com.

 – Recognition of your technical skills by your peers and management

 – Enhanced career opportunities

 – Focus for your professional development

▶ For the business partner:

 – Confidence in the skills of your employees

 – Enhanced partnership benefits from Business Partner Program

 – Higher rates for billing out your employees

 – Stronger customer proposals

 – Demonstration of the depth of technical skills available to prospective customers

▶ For the customer:

 – Confidence in the services professionals handling your implementation

 – Ease of hiring competent employees to manage your Tivoli environment

 – Enhanced return on investment (ROI) through more thorough integration with Tivoli and third-party products

 – Ease of selecting a Tivoli Business Partner that meets your specific needs

Certification checklist

Here is the Certification checklist:

1. Select the certification you would like to pursue.

2. Determine which tests are required by reading the certification role description.

3. Prepare for the test, using the following resources:

   – Test objectives
   – Recommended educational resources
   – Sample/assessment test
   – Other reference materials
   – Opportunities for experience

   > **Note:** These resources are available from each certification description page, as well as from the Test information page.

4. Register to take a test by contacting one of our worldwide testing vendors:

   – Thomson Prometric
   – Pearson VUE (Virtual University Enterprises)

   > **Note:** When providing your name and address to the testing vendor, be sure to specify your name exactly as you would like it to appear on your certificate.

5. Take the test. Be sure to keep the Examination Score Report provided upon test completion, as your record of taking the test.

   > **Note:** After you take the test, the results and demographic data (such as name, address, e-mail, and phone number) are sent from the testing vendor to IBM for processing (please allow two to three days for transmittal and processing). Once all the tests required for a certification are passed and received by IBM, your certificate will be issued.

6. Repeat steps three through five until all required tests are successfully completed for the certification. If there are additional requirements (such as another vendor certification or exam), please follow the instructions on the certification description page to submit these requirements to IBM.

7. Once you meet the requirements, you will be sent an e-mail asking you to accept the terms of the IBM Certification Agreement.

8. Upon your acceptance, you receive an e-mail with the following deliverables:

   – A Certification Certificate in .pdf format, which can be printed in either color or black and white.

   – A set of graphic files containing the IBM Professional Certification mark associated with the certification achieved.

   – Guidelines for the use of the IBM Professional Certification mark.

9. To avoid unnecessary delay in receiving your certificate, please ensure that your current e-mail is on file by keeping your profile up to date. If you do not have an e-mail address on file, your certificate will be sent by postal mail.

   Once you receive a certificate by e-mail, you may also contact IBM at certify@us.ibm.com to request that a hardcopy certificate be sent by postal mail.

> **Note:** IBM reserves the right to change or delete any portion of the program, including the terms and conditions of the IBM Certification Agreement, at any time without notice. Some certification roles offered through the IBM Professional Certification Program require recertification.

1.2 IBM Tivoli Storage Manager 5.4 Certification

In this section, we categorize the certification process.

> **Important:** IBM offers a promotion code below, which is good for a 15 percent discount on the indicated Tivoli certification exams if taken at any Thomson Prometric or Pearson VUE testing centers.
>
> ► Code: 15T899
> ► Percent off: 15 percent
> ► Valid for exams: 000-899
> ► Expires: Code is valid as long as exam is available

1.2.1 Job description and target audience

An IBM Certified Deployment Professional - IBM Tivoli Storage Manager V5.4 is an individual who has demonstrated the ability to implement and support IBM Tivoli Storage Manager. It is expected that this person is able to perform the following tasks independently a majority of the time and, in some situations, take leadership and provide mentoring to peers. It is expected that this person will be able to perform these tasks with limited assistance from peers, product documentation, and vendor support services.

1.2.2 Key areas of competency

This person is expected to perform the following tasks independently:

► Describe IBM Tivoli Storage Manager's features, functions, and benefits.

► Understand the customer's requirements, design the proper IBM Tivoli Storage Manager solution (including policy definitions, schedules, devices,

and communication requirements), and configure IBM Tivoli Storage Manager (including selecting the appropriate IBM Tivoli Storage Manager products).

► Install IBM Tivoli Storage Manager; set up remote clients and local storage devices; specify backup schedules; authorize users; and perform performance tuning, problem determination, and resolution.

► Protect the customer's IBM Tivoli Storage Manager server through database and storage pool backups.

1.2.3 Required prerequisites

To attain the IBM Certified Deployment Professional - IBM Tivoli Storage Manager V5.4 certification, candidates must pass one test.

The required prerequisites are:

► The IBM Tivoli Storage Manager V5.4 exam is platform independent and applies to heterogeneous environments. Candidates *must* have background experience sufficient to prepare for training and testing on storage solutions. The following qualifications are requirements for success:

► Advanced backup methodologies (full, incremental, differential, image, progressive, on site, off site, etc.)

► Basic operating systems methodologies

► Basic networking methodologies

► Basic storage device knowledge (optical, disk, tape, NAS, etc.)

► Basic hardware knowledge (server, CPU, memory, etc.)

1.2.4 Core requirement

To be certified you must select Test 899: IBM Certified Deployment Professional - IBM Tivoli Storage Manager V5.4 implementation. This test has the following components and characteristics:

► Test 899 Objectives
► Test 899 Sample test
► Test 899 Recommended educational resources
► Approximate number of questions: 69
► Duration in minutes: 75
► Format: Multiple choice
► Required passing score: 73 percent or 51 correct answers

1.2.5 Test 899 Objectives

In this section we review the Test 899 Objectives.

Section 1: Planning

This section of the test has these requirements:

► Given a qualified, feasible customer service level agreement, and business and legal data retention requirements, apply these agreements to the policy workshop recommendations so that a qualified, validated data retention plan for TSM is produced. With emphasis on performing the following steps:

- – Determine feasibility of customer's service level agreement request with current version of TSM.
- – Apply Business and legal data retention requirements to policy retention requirements.
- – Review policy retention recommendations.
- – Produce a data retention plan.

► Given the SLAs and design requirements, evaluate the features of TSM to design a solution which meets requirements while minimizing cost. With emphasis on performing the following steps:

- – Interpret business SLAs and retention requirements into technical requirements.
- – Translate technical requirements into features of TSM.
- – Evaluate the cost of each of the features in TSM which can be applied to the technical requirements.
- – Design a TSM solution which meets the requirements with the minimum cost.

► Given the knowledge of a proposed TSM server environment, a working internet connection and an internet browser, compare the characteristics of the proposed environment with system requirements so that compliance is confirmed. With emphasis on performing the following steps:

- – Open internet browser.
- – Navigate to IBM web site
- – Navigate to IBM Tivoli Storage Manager Product web page.
- – Select server requirements.
- – Select entry for appropriate server platform.
- – Review requirements.

– Compare proposed server environment to requirements.

► Given client hardware, network protocols and applications, discover customer hardware, identify OS software, determine network connectivity and locate applications running, so that a list of hardware, operating systems software and OS configuration requirements for TSM is produced. With emphasis on performing the following steps:

– Discover customer's hardware to be used with TSM.

– Identify Operating systems used on discovered hardware.

– Determine network connectivity and protocols available.

– Locate applications running on hardware to be used with TSM.

► Given the configuration management information about the customer environment, identify the systems and applications in the environment so that a list of required TSM nodes can be determined. With emphasis on performing the following steps:

– Identify the systems and applications in the environment.

– Map the systems and application to TSM client nodes.

► Given customer requirements and a knowledge of the TSM product family, identify special application requirements to design an appropriate TSM solution. With emphasis on performing the following steps:

– Gather customer requirements.

– Obtain knowledge of IBM TSM related products.

– Compare TSM family products with customer requirements.

– Recommend appropriate TSM family.

► Given hardware, operating systems, and network connectivity requirements, evaluate bare machine recovery so that the most appropriate methodology can be determined.

– Determine customer's hardware to be used with TSM. With emphasis on performing the following steps:

– Identify Operating systems used on hardware.

– Determine network connectivity and protocols available.

– Determine the most appropriate bare machine recovery method for the customer's environment.

► Given the SLAs, design requirements, existing hardware, and hardware preferences, evaluate the infrastructure and design goals, so that a TSM solution which meets requirements and maps to specific hardware is produced. With emphasis on performing the following steps:

- – Interpret network infrastructure to determine if it's sufficient to meet needs.
- – Determine if a SAN solution should be deployed for TSM and the systems where it should be deployed.
- – Evaluate the disk infrastructure of the customer to determine what's appropriate for the TSM server and client.
- – Evaluate TSM server hardware to determine what's appropriate for the TSM server.
- – Design a TSM solution which meets the requirements with a given hardware.

► Given customer requirements and a knowledge of the TSM v5.4, identify features to best address the customer requirements so that an appropriate TSM solution is designed. With emphasis on performing the following steps:

- – Gather customer requirements.
- – Obtain knowledge of IBM TSM v5.4 server features.
- – Identify TSM v5.4 server features which meet customer requirements.
- – Recommend appropriate TSM v5.4 server and client features.

► Given the customer's hardware, total amount of data to be saved, percentage of daily change rate for the data, the network connectivity available, and policy recommendations, evaluate the infrastructure and requirements, so that TSM backup methods for client data can be recommended. With emphasis on performing the following steps:

- – Determine customer's available hardware.
- – Identify total amount of data to be saved.
- – Identify percentage of data that is changed daily.
- – Determine network connectivity and protocols available.
- – Review TSM policy recommendations for client servers.

► Given the SLAs, design requirements, network infrastructure, amounts of data, and types of data, determine TSM storage, feature and database configuration so that the number and location of TSM server instances in a particular design is planned.
Interpret network infrastructure to determine where it would be appropriate to place TSM servers. With emphasis on performing the following steps:

- – Determine the number of copy pools and versions required in TSM to meet design requirements.
- – Determine the features of TSM to be used to meet design requirements.
- – Estimate the size of TSM server databases in the environment.

► Given customer requirements, knowledge of hardware components and a basic knowledge of TSM design, identify appropriate hardware components needed so that an appropriate TSM solution can be determined. With emphasis on performing the following steps:

– Gather customer requirements.

– Obtain knowledge of hardware characteristics.

– Obtain a basic knowledge of TSM solution design.

– Apply solution design and hardware component knowledge to customer requirements. Recommend appropriate hardware to implement the TSM solution.

Section 2: Installation

This section of the test has the following requirements:

► Given the network and hardware that meet requirements for TSM, determine and verify the configuration of existing hardware, so that TSM software can be installed and configured. With emphasis on performing the following steps:

– Verify latest TSM 5.4 system requirements on the Tivoli support Web site.

– Ensure that the required operating system is installed.

– Ensure that the required browser is installed.

– Test the system for network connectivity.

– Verify disk and volume or directory space requirements are set up correctly.

– Verify characteristics of tape device and verify that it is set up correctly.

► Given a correctly set up system and network which meet the requirements, and access to the TSM code, follow the documented TSM server installation steps for Windows so that the TSM server is installed on the system. With emphasis on performing the following steps:

– Read appropriate documentation and Support Flashes on the IBM Web site.

– Insert TSM server CD or run downloaded code.

– Select demonstration language.

– Click Install Products.

– Select TSM server.

– Accept License Agreement.

– Select setup type (Complete or Custom).

- – Click Install in the Ready to Install the Program window.

- – Click Install.

- – When the Install Shield Wizard Completed dialog appears, click finish.

▶ Given a correctly set up system and network which meet the requirements, and access to the TSM code, follow the documented TSM server installation steps for AIX so that the TSM server is installed on the system. With emphasis on performing the following steps:

- – Read appropriate documentation and Support Flashes on the IBM Web site.

- – Log in as root user.

- – Insert TSM server CD or run downloaded code.

- – In the Common Desktop Environment, access SMIT in the System_Admin folder of the Application Manager and make selections according to the level of your AIX operating system

- – Click the SMIT icon and select Software Installation and Maintenance.

- – List software to install.

- – Set both Automatically Install Required Software and Accept New License Agreement to Yes.

- – Select the file sets you want to install.

- – If you are working in a language other than US English, you must also install your specific language pack in addition to the required U.S. English language pack.

- – Click Install.

- – Exit SMIT when the installation is complete.

▶ Given that the environment requirements are met and the TSM server is correctly installed, work through the Install Products dialog so that the license files are installed. With emphasis on performing the following steps:

- – From the Install Products dialog, click Tivoli Storage Manager Server Licenses.

- – In the Setup Type dialog, select Complete or Custom.

- – In the Ready to Install the Program window, click Install.

- – When the Install Shield Wizard Completed dialog appears, click Finish.

▶ Given that the environment requirements are met and the TSM server is correctly installed, work through the Install Products dialog so that the language packs are installed. With emphasis on performing the following steps:

- – From the Install Products dialog, click Tivoli Storage Manager Server Language Packs.

- – Select the appropriate language to install.

- – Click Install.

► Given that the environment requirements are met and the TSM server is correctly installed, work through the Install Products dialog so that the device drivers are installed. With emphasis on performing the following steps:

- – From the Install Products dialog, click Tivoli Storage Manager Device Driver.

- – Complete the Customer Information Panel.

- – In the Setup Type dialog, select Complete or Custom.

- – In the Ready to Install the Program window, click Install. When the Install Shield Wizard @Completed dialog appears, click Finish.

► Given that the environment requirements are met and the TSM server is correctly installed, perform the installation steps for the Integrated Solutions Console so that the Integrated Solutions Console is operating. With emphasis on performing the following steps:

- – Insert the Integrated Solutions Console CD into your CD drive.

- – Start one of the installation wizards and follow the instructions.

- – Check the box to accept the license agreement.

- – Change the Integrated Solutions Console user ID and password.

- – Verify installation by accessing the console through a supported Web browser: http:// :8421/ibm/console.

► Given that the environment requirements are met and the TSM server and the Integrated Solutions Console are correctly installed, perform the installation steps so that the TSM Administration Center can manage the Tivoli Storage Manager servers. With emphasis on performing the following steps:

- – Start the Tivoli Storage Manager Version 5.4 servers.

- – Give each server a unique name.

- – Use the Administration Center CD or image to install the Administration Center on the same system as the Integrated Solutions Console.

- – Start one of the installation wizards and follow the instructions.

- – Check the box to accept the license agreement.

- – Verify installation by accessing the console through a supported Web browser: http:// :8421/ibm/console.

- Add connections for all the Tivoli Storage Manager servers managed by the Administration Center.

► Given a correctly installed LVSA, enable OFS and IMAGE backup options so that OFS and IMAGE type backups can be performed on a correctly installed TSM BA client. With emphasis on performing the following steps:

- Determine which drives will need to be backed up using Image Backup.

- Start the client installation Wizard.

- In the "Setup Type" panel, select the "Custom" radio button.

- In the "Custom Setup" panel select "Image Back up" and "Open File Support".

- Click the "Next" button. Then click the "Install" button.

- Monitor Installation.

- If failure, review TSM error pop up dialog box to determine error.

- View the install Wizard Completed pop up Box and verify successful install.

► Given a correctly installed BA Client, implement and configure JBB so that file systems are backed up and data is sent to the server. With emphasis on performing the following steps:

- Determine which file system and files to use JBB back-up.

- Determine where to put the directory, where the journal database files will be written.

- Determine max size of the journal DB.

- Determine user notification filters.

- Determine any security options (Account that will has permissions to activate JBB).

- Start the TSM BA Client GUI.

- Click on the "Utilities" menu item and select "Setup Wizard".

- Check the box next to "Help me configure the TSM Journal Engine".

- Select "Install a new TSM Journal Engine".

- Step through the rest of the setup Wizard to configure all JBB parameters.

- Start the Journal Service.

- Perform an incremental backup to activate JBB.

- Perform another incremental backup to very JBB has been activated and a JBB backup was performed.

- View the JBB DB to verify data is available to the server.

► Given that the operating system and TSM server are running, run the setup wizard selecting the appropriate options so that the TSM Client Options and Administrative client are installed and functioning correctly. With emphasis on performing the following steps:

– Start the client GUI installation program.

– In the "Setup Type" panel, select the "Custom" radio button.

– In the "Custom Setup" panel select "Administrative Client Command Line files".

– Click the "Next" button. Then click the "Install" button.

– Monitor installation.

– If failure, review TSM error pop up dialog box to determine error.

– View the install Wizard Completed pop up Box and verify successful install.

– Open the Administrative command line Client and log on to the TSM Server.

► Given a working TSM environment, install the Storage Agent and configure TSM server and client, so that LAN-free backups can be performed. With emphasis on performing the following steps:

– Configure the TSM server (library manager or database server)

– Define domain that includes a LAN-free Storage pool

– Install storage agent software for each LAN-free client

– Define storage agent on server (using "define server" command)

– Customize the file dsmsta.opt

– Configure the storage agent using "dsmsta setstorageserver" this is usually done via creating a script.

– Inspect the resulting device configuration file (dev.cfg)

– Define paths to storage agent drives on the server

– Register node on TSM server, and assign it to the domain which goes directly to the @LAN-free tape storage pool

– Set up TSM client with the correct LAN-free options

– Enablelanfree yes

– Lanfreecommmethod tcpip|namedpipe|sharedmem

– Lanfreetcpserv xyz.sanjose.ibm.com

– Lanfreetcpp 1500

– Lanfreeshmport 1580

- – Run TSM server command "Validate Lanfree" to validate LAN-free setup was successful.
  - – If the validation test fails review log files.

▶ Given an existing TSM server and operational reporting requirements, install and configure the TSM management console and operational reporting so that reports can be sent to appropriate TSM management staff. With emphasis on performing the following steps:

  - – Interpret reporting requirements and map them to TSM Operational reporting.
  - – Install the TSM Console.
  - – Configure Operational Reporting Daily Reports to run against your TSM servers.
  - – Configure Operating Reporting e-mails and events to be sent to administrators covering failures and other pertinent data.

Section 3: Configuration

This section of the test has these requirements:

▶ Given TSM server is installed, modify the dsmserv.opt and issue server option set commands so that server options are set appropriately. With emphasis on performing the following steps:

  - – Configure options in dsmserv.opt which are appropriate to a customer's environment.
  - – Configure options using server set commands which are appropriate to a customer's environment.

▶ Given a correctly configured TSM server environment, a correctly configured Integrated Solutions Console and an Internet connection, add a TSM server connection to the Administration Center to provide centralized access to the TSM server. With emphasis on performing the following steps:

  - – Log on to Internet browser.
  - – Navigate to ISC Web page.
  - – Select Storage Devices.
  - – Select Add Server Connection . . . option in Select Action box.
  - – Provide server connection information and select OK.

▶ Given a working TSM environment, TSM client options file and knowledge of customer environment, configure client and server security-related options so that TSM server and client operates in a custom security designed environment. With emphasis on performing the following steps:

- – Determine security requirements before installing TSM server client.

- – Determine if a firewall will be needed.

- – Determine who will be performing backup.

- – Determine if password encryption be used.

- – Determine is SSL a requirement.

- – Determine if the client or the server be behind a proxy server or firewall.

- – Determine if TSM Ports need to be turned on or turned off for proper operation.

- – Determine your environment Allow only admin backup and restores over certain ports.

- – Determine where passwords files need to be stored (passworddir).

- – Determine retention policy as related to security.

- – Determine how to preserve file dates related to security.

► Given the customers hardware, media type available, network connectivity being used, policy recommendations, and the TSM Commands and processes necessary, configure tape library, drives and device class so that the library is ready for use. With emphasis on performing the following steps:

- – Determine customer's hardware available.

- – Identify media types available with customer's hardware.

- – Identify Operating Systems to be used with hardware.

- – Determine network connectivity and protocols available.

- – Review TSM policy recommendations for client servers.

- – Configure library, drives, path and device classes.

► Given a functioning TSM server, examine number of files to be backed up, so that appropriately sized database and log volumes can be created. With emphasis on performing the following steps:

- – Determine how many files will be backed up, and multiply that by 400 bytes for initial, 200 for each copy (for example, copypool) and 200 for each addition version in a primary.

- – Define database volumes (most TSM databases are limited to 100GB for performance).

- – Define log volumes (size limit is 13 GB).

- – Determine whether to select option to automatically create DB volumes as needed.

- – Determine whether to select option to automatically create Log volumes as needed.

► Given a functioning TSM server and sufficient storage space, define a primary storage pool so that it is configured and ready for use. With emphasis on performing the following steps:

- – Determine what device class to use

- – Determine what type of storage pool to create.

- – Determine a name for the storage pool

- – Give the pool a max scratch parameter.

- – Set the access mode.

- – Set the nextinchain.

- – Set migration threshold.

- – Specify number of migration processes.

► Given a functioning TSM environment where client data has been backed up to the primary storage pools, configure copy storage pools, so that primary storage pool backups can be processed. With emphasis on performing the following steps:

- – Determine which primary pool the copypool will be created for.

- – Determine the format of the primary pool as it must be the same data format as the source storage pool (NATIVE, NONBLOCK, or any of the NDMP formats for a regular copy storage pool).

- – Run define stgpool command providing name of pool and appropriate parameters.
  Specify the device class
  Specify the pooltype as copy

- – Determine if you will be collocating the data in the copy storage pool.

- – Specify the access mode, readonly, readwrite, or unavailable.

- – Specify the maxscratch value.

- – Set reclamation threshold.

- – Set reusedelay.

- – Set offsite reclaim limit to specify the number of reclamation processes allowed.

- – Add volumes.

► Given an installed TSM server environment, configure active-data pools so that storage pools are designed and configured appropriately. With emphasis on performing the following steps:

- – Determine appropriate use for active-data pools in the TSM solution.

- – Determine where active-data pools will be stored in the TSM solution.

- – Determine which TSM clients will use active-data pools.

- – Log on to Internet browser.

- – Navigate to ISC Web page.

- – Select Storage Devices.

- – Select TSM Server and choose View Storage Pools option in the Select Action box.

- – In newly opened Storage Pool window, select Create a Storage Pool.

- – Enter storage pool information (specifying Sequential Access and Active Data) and select OK.

- – Complete configuration wizard.

► Given a functioning TSM server and client, and an existing policy domain and policy set, create a management class so that backup and archive copy groups can be added. With emphasis on performing the following steps:

- – Determine for which policy domain and policy set you will create the management class.

- – Use the define mgmtclass command, specifying the policy domain, the policy set, and the name of the management class.

► Given multiple TSM server instances, configure enterprise management and server-to-server communications so that TSM policy and objects are standardized and TSM administration is centralized. With emphasis on performing the following steps:

- – Determine whether enterprise management is appropriate for a given solution.

- – Configure server-to-server communications.

- – Configure TSM objects to be centrally managed using enterprise management.

► Given an existing TSM server, SAN or LAN infrastructure, and a supported NDMP device, install and configure TSM support for the NDMP device so that the device can be backed up and restored. With emphasis on performing the following steps:

- – Determine the location of the NAS backups.

- – Set up the tape library for TSM operations.

- – Configure TSM policy and Storage Pools for NDMP operations.

- – Schedule the NDMP backups.

► Given a correctly installed TSM server, configure TSM client options so that the TSM client backup-archive functions are working properly. With emphasis on performing the following steps:

  – Determine what dsm.opt options are appropriate

  – Launch the TSM client GUI or JAVA GUI.

  – From the Utilities menu, select Setup Wizard. This will launch the Tivoli Storage Manager Client Configuration Wizard.

  – On the first panel of the Tivoli Storage Manager Client Configuration Wizard, make sure Help me configure the TSM Backup Archive Client is selected.

  – Follow the prompts to finish creating your options file. Most options can be set up using the Setup Wizard.

  – include/exclude

  – client communications

  – performance tuning

  – Alternatively Options can set manually by editing the DSM.opt file (Windows) DSM.SYS and DSM.OPT (AIX).

► Given a correctly installed and configured TSM environment, install and configure the TSM backup-archive client in the cluster environment so that Cluster volumes can be backed up successfully. With emphasis on performing the following steps:

  – Verify Cluster environment.

  – Identify the cluster group/s to manage.

  – Determine where backup operations should be performed (for example, Which node all nodes).

  – Register backup-archive client node definitions on the Tivoli Storage Manager Server for each cluster group.
    Define the Domain option as required.

  – Configure the Client Options file (DSM.OPT) for each cluster group. (for example, cluster yes). The DSM.opt file must reside on the shared volumes. Use the Cluster name for the TSM node name.

  – Validate the setup on each node.

  – On the Administrative command line (on the TSM client machine) issue command dsmc query session

  – Enter password if prompted

  – Enable Tivoli Storage Manager web client to access the Cluster environment.

- – Edit the DSM.opt file as required

- – Set up any scheduling for cluster node backup.

► Given a correctly installed and configured TSM environment, configure the BA client so that files are backed up using adaptive subfile backup. With emphasis on performing the following steps:

- – Configure the server to allow this type of backup
  On the server issue: SET SUBFILE CLIENT

- – Set the subfilebackup, subfilecachepath and subfilecachesize options in your client options file "dsm.opt

- – set lthe subfilebackup option to yes (subfilebackup yes)

- – subfilecachepath (/mydir/cachedir)

- – subfilecachesize (500M)

- – Add an exclude.dir option to your client options file (dsm.opt), which excludes the directory containing the subfile cache information.

- – Exclude.Dir (/mydir/cachedir)

- – Perform an incremental backup.

- – Verify backup completed.

- – Change a couple file's data this will also change the size.

- – Perform a second incremental backup.

- – Verify only changes to files in step 6 were backed up.

► Given a functioning TSM environment, configure Disaster Recovery Management, so that the server and the client data can be recovered in the event of a disaster. With emphasis on performing the following steps:

- – Determine which storage pools will be managed by disaster recovery manager.

- – Determine which copy storage pools will be managed by disaster recovery manager.

- – Set the identifier for replacement volumes using the set drmplanvpostfix command.

- – Specify a prefix for the names of the recovery instructions source files in the recovery plan using the set drminstrprefix command.

- – Specify a prefix to the path name of the recovery plan file using the set drmplanprefix command.

- – Set the number of days after creation that a disaster recovery plan file stored on a target server will expire, using the set drmrpfexpiredays command.

- – Create recovery instructions files.
- – Use the define machine command with adsmserver=yes to define machines to be recovered, their location, and the priority of recovery.
- – Issue the define machnodeassociation to associate one or more client nodes with a machine.
- – Issue the insert machine command to add machine characteristics and recovery instructions to the database.
- – Use define recoverymedia command to specify bootable media needed to reinstall the Operating system of a client machine, and associate the machines with media.
- – Use the query drmstatus command to verify the status of your plan.
- – Use the set drmcmdfilename c:\drm\orm\exec.cmd to set the file name where executable commands can be stored.
- – Use the query drmedia or move drmedia to generate executable commands and store them in a file.

► Given a TSM environment with Disaster Recovery Management in use, configure tape rotation, so that media is utilized efficiently. With emphasis on performing the following steps:

- – Use set commands to change default states or locations, if desired, for offsite recovery media management.
- – Determine state of recovery media using the query drmedia command.
- – Mark the volumes as unavailable to Tivoli Storage Manager and give them to the courier.
- – Indicate the movement of volumes whose current state is MOUNTABLE by issuing the following command: move drmedia * wherestate=mountable.
- – Issue the prepare command after the move drmedia command has been issued to create your Disaster Recovery Plan file.
- – Send all DR media, and plan file in softcopy and hardcopy, offsite.
- – For all volumes in the NOTMOUNTABLE state, DRM updates the volume state to COURIER and the volume location according to the SET DRMCOURIERNAME. If the SET command has not yet been issued, the default location is COURIER.
- – When the vault location confirms receipt of the volumes, issue the MOVE DRMEDIA command in the COURIER state.
- – To display a list of volumes that contain valid data at the vault, issue the following command: query drmedia wherestate=vault.

- Issue the query drmedia * wherestate=vaultretrieve command to determine non-virtual volumes that can be returned from the vault. Update the location as needed until the volumes are returned.

- Update the state of these tapes and check them into the library.

Section 4: Performance Tuning and Problem Determination

This section has these requirements:

► Given an existing TSM server with a tape drive, label and check in tapes so that the TSM server has scratch volumes. With emphasis on the following steps:

- Label Tape media for the TSM server.

- Check-in tape media to the TSM server.

► Given a correctly functioning Administration Center and client machine, navigate through the Policy Domains and Client Nodes wizard so that client nodes can be accessed and managed with the TSM Administration Center. With emphasis on the following steps:

- Log on to the TSM Administration Center.

- In the server's Policy Domains table, click the name of a domain.

- In the domain's properties portlet, click Client Nodes.

- In the Client Nodes table, click Select Action, select Create a Client Node.

- Enter the client name.

- Enter contact information.

- Enter the client's Web address

- Select a client option set.

- Select a client schedule.

- Verify the client nodes can be accessed by the Administration Center.

► Given a working TSM environment where there is the need to have one node restore data backed up from another node, create a node access so that one node now has been granted the authority to restore files for another specified node. With emphasis on the following steps:

- Issue the query access command to determine node access.

- Set up server to allow client to client node access.

- Use the GUI to create a node access list.

- Alternatively on the command line client you can use the set access command.

- ▶ Given a correctly installed and connected TSM environment, perform backup so that successful installation and configuration of the TSM environment can be confirmed. With emphasis on the following steps:
  - Start the TSM backup archive client GUI.
  - Click Backup.
  - Select objects to be backed up.
  - Select type of backup (selective, incremental,)
  - Click Backup.
  - Monitor backup progress.
  - Review the detailed status report and error code to confirm backup completion.
  - If failure, review TSM error log to determine error.

- ▶ Given a correctly installed and connected TSM environment, perform a file restore so that successful installation and configuration of the TSM environment can be confirmed. With emphasis on the following steps:
  - Start the TSM backup archive client GUI.
  - Click Restore.
  - Identify files to restore.
  - Determine version of the file to restore.
  - Select restore options.
  - Click Restore.
  - Monitor progress.
  - Review the detailed status report and error code to confirm restore completion.
  - If failure, review TSM error log to determine error.

- ▶ Given a correctly installed and connected TSM environment, perform archive so that successful installation and configuration of the TSM environment can be confirmed. With emphasis on the following steps:
  - Start the TSM backup-archive GUI.
  - Click Archive.
  - Select objects to be archived.
  - Enter description.
  - Select archive options
  - Click Archive.

- – Monitor archive progress.

- – Review the detailed status report and error code to confirm archive completion.

- – If failure, review TSM error log to determine error.

► Given a correctly installed and connected TSM environment, perform a retrieve of an archived file so that successful installation and configuration of the TSM environment can be confirmed. With emphasis on the following steps:

- – Start the TSM backup archive client GUI.

- – Click Retrieve.

- – Identify files to retrieve.

- – Select retrieve options.

- – Determine destination of the archived file.

- – Click Retrieve.

- – Monitor progress.

- – Review the detailed status report and error code to confirm restore completion.

- – If failure, review TSM error log to determine error.

► Given the customers policy requirements, client backup and archive data volume, available storage pool size being used and the amount of TSM Server time available for maintenance, create administrative schedules so that server maintenance is automated. With emphasis on the following steps:

- – Identify media types available with customer's hardware.

- – Identify client backup and archive volume and storage pool sizes.

- – Review TSM policy recommendations for client servers.

- – Determine customers backup time available and TSM server maintenance time available.

► Given a correctly configured TSM server environment, configure activity log retention option so that the activity log is automatically pruned. With emphasis on the following steps:

- – Query the current value for activity log retention.

- – Determine the change value.

- – Update the value for activity log retention.

► Given the customers hardware and media types available, database size and estimated threshold, configure TSM database backup trigger so that database is automatically backed up. With emphasis on the following steps:

- – Identify media types available with customer's hardware.

- – Determine TSM Server maintenance time available.

- – Define backup trigger.

- – Set logmode to rollforward.

► Given a functioning TSM server, set migration thresholds, so that the flow of data through the storage pool is optimized. With emphasis on the following steps:

- – Determine appropriate migration thresholds for your storage pools.

- – Set migration threshold on the storage pool using the HIGHMIG and LOWMIG options on the DEFINE STGPOOL or UPDATE STGPOOL commands identifying a high threshold, a low threshold, and the next storage pool in the chain.

- – Optionally use the update stgpool command to manually force migration.

- – Use the MIGDELAY to delay migration a specified number of days.

► Given a correctly installed BA Client, implement and configure a schedule so that files are backed up automatically to insure that data is available as needed. With emphasis on the following steps:

- – Determine what needs a scheduled backup, archive, restore and retrieve.

- – Determine how often files need to be backed up, archived, retrieved or restored.

- – Determine what kind of schedule is required (polling or prompted).

- – Define a policy domain on the server.

- – Assign the TSM client (that is scheduled) to the policy domain just defined.

- – Define a schedule name on the Server.

- – Define an association on the server between the TSM client node and the schedule name just created.

- – Invoke the schedule.

- – Verify Schedule log file has all schedule occurrences listed.

► Given a correctly installed TSM environment, generate a backup set then perform a query of the backupset so that the correct information was generated. With emphasis on the following steps:

- – Determine what files need to be included in a backup set.

- – Determine location for storing the backup set. (portable media, TSM server Client workstation)

- – Create/define a device class.

- – Perform an incremental backup/selective backup of the files to be included in the backupset.

- – Generate a backupset on the server.

- – Assign the backupset to the device class.

- – Assign name and description.

- – Assign retention limit.

- – Specify location of backupset (default is server).

- – Issue a Query Backupset command on the server (or client admin command line) to verify backup set got properly generated.

- – Restore Backupset.

► Given a correctly configured TSM server environment, review messages from the TSM server activity log so that errors and issues in the TSM environment can be detected. With emphasis on the following steps:

- – Select Server Maintenance in left column.

- – Select TSM Server and choose Server Properties option in the Select Action menu.

- – Select Activity Log Tab.

- – Update filter information and select update table.

- – Review messages in the resulting table.

► Given a TSM environment with backup and recovery events scheduled, monitor these events so that the administrator in charge is aware of the status of all backup and recovery events. With emphasis on the following steps:

- – Verify that drives are online.

- – Verify that database and recovery log volumes are online and synchronized.

- – Check the status of disk volumes.

- – Check that scratch volumes are available.

- – Check the access state of the tape volumes.

- – Check database and recovery log statistics.

- – Verify that scheduled database backups completed successfully.

- – Check the activity log for error.

- – Install and use Operational Reporting to monitor events.

► Given a Tivoli Storage Manager environment where Disaster Recovery Management is completely configured and in use, perform the proper steps,

in the correct order, so that a complete Disaster Recovery Plan file is generated. With emphasis on the following steps:

– Complete all client backups.

– Complete all storage pool backups.

– Complete the database backup.

– Run the move drmedia command to update the DR plan and checkout media.

– Run the prepare command to generate the DR Plan file.

►

Given a Tivoli Storage Manager environment where Disaster Recovery Management is completely configured and in use, and files are being regularly sent off-site, perform a test of the DR Plan so that the recoverability of the server and client files can be verified. With emphasis on the following steps:

– Deliver DR materials to a designated recovery site.

– Break out the disaster recovery plan and the DR media and follow the instructions, testing any scripts for automation of recovery where possible.

– If instructions are insufficient or data is not recoverable, take steps to correct these issues and try the test again.

– Continue testing and correcting until you successfully recover the TSM server, storage pools, and client data.

– Conduct a lessons learned discussion and update DR Plan.

– Create a schedule for testing the Disaster Recovery Plan at least once or twice each year, and allow time during each test to update the plan so until a successful recovery is achieved.

► Given an existing TSM server, administrative client, and Operational Reporting, extract capacity planning information so that information on current capacity usage for the TSM server can be determined. With emphasis on the following steps:

– Determine information from the TSM environment which can help an organization do capacity planning.

– Extract capacity planning information from the summary table and TSM Operational Reporting.

– Produce Capacity Management Reports which can help size the TSM server.

► Given a configured TSM server, policy definitions for administrators, nodes and policies, execute TSM commands so that server information and data can be exported and imported. With emphasis on the following steps:

– Identify TSM Policy definitions.

– Identify system administrators.

– Identify client nodes.

– Identify policy domains.

– Issue import/export command.

Section 5: Administration

This section has these requirements:

► Given the customers hardware, TSM maintenance level, configured Operating system, correct hardware device drivers, and current firmware levels, validate the installation process so that installation problems are detected. With emphasis on the following steps:

– Identify operational status of customer's hardware and operating system.

– Identify TSM Server availability.

– Review TSM activity log, Client event log, OS error logs.

– Identify Device drivers and firmware levels.

► Given a correctly configured TSM server environment, navigate to the directory where TSM instance is installed so that logs can be reviewed to aid in problem determination. With emphasis on the following steps:

– Log on to the server hosting the TSM server software instance.

– Navigate to the installation directory for the server instance.

– Review the various logs for this TSM server instance.

► Given a correctly configured TSM server environment, review error messages from the TSM server activity log so that corrective action can be taken. With emphasis on the following steps:

– Query the activity log for the period of interest.

– Identify errors, warnings and suspicious items.

– Determine the cause aforementioned items.

– Determine corrective actions to take.

► Given an issue which needs to be resolved, determine the appropriate TSM resources so that a solution can be provided. With emphasis on the following steps:

– Understand resources available to help diagnose issues.

– When an issue occurs in the environment, determine where to look for help in discovering the solution.

▶ Given the customers hardware, network environment, and Operating system, determine the cause of failure or lack of connectivity so that corrective action can be taken. With emphasis on the following steps:

– Identify operational status of customer's hardware and operating system.

– Identify TSM Server availability.

– Review TSM activity log, Client event log, OS error logs.

– Verify TSM communication-related information in the options file.

– Identify communication devices and connectivity methods.

▶ Given the customers hardware and Operating system, determine the cause of failure of a specific storage-related device so that the failure can be corrected. With emphasis on the following steps:

– Identify operational status of customer's hardware and operating system.

– Identify TSM Server availability.

– Review TSM device class, tape drive, library configuration and attached devices.

– Review TSM activity log, Client event log, OS error logs.

▶ Given the customers available hardware and connectivity and TSM Server availability, check the TSM Activity log, and Client error and schedule logs, to locate messages and errors so that the cause of backup failures can be corrected. With emphasis on the following steps:

– Identify operational status of customer's hardware and operating system.

– Identify TSM Server availability.

– Review TSM activity log, Client event log, OS error logs and client schedule logs.

– Correct issues with client backups.

▶ Given a basic performance issue in TSM, adjust TSM server and client parameters so that better performance is achieved and client SLAs are met. With emphasis on the following steps:

– Investigate current performance of the TSM client and server.

– Determine where performance falls outside of agreed upon SLAs.

– Make changes to the TSM client and server solution to improve performance to meet SLAs, including changes to the dsmserv.opt, dsm.opt, and dsm.sys.

▶ Given a TSM server, determine issues with bad or corrupt media and attempt to repair or restore them so that the faults with the media are corrected. With emphasis on the following steps:

- – Determine if any tape of disk media in the TSM server is bad or corrupted.
- – Attempt to recover or repair the media with the audit vol command.
- – If media is destroyed, attempt to restore the data with a restore vol or restore pool command.

► Given a Tivoli Storage Manager environment where the amount of tape media being used is increasing faster than the increase in the amount of data being stored, determine the issues that are causing the inefficient use of media and correct it so that tape rotation is regulated and media usage is optimized. With emphasis on the following steps:

- – Verify that all collocated client data really needs to be collocated.
- – Run an audit lib to make sure the database information on the volumes is correct.
- – Ensure that tape reclamation and expiration are running regularly.
- – Check to see if data retention is on where it should be off.
- – Check the reclamation threshold to see if it's appropriate.
- – Check the reusedelay to see if it's set too high.
- – Check the offsite reclaim limit for the storage pool to see if it's too low.
- – Check the RETMIN parameter in the archive copy group.
- – Check the RETVER parameter on the archive copy group also.
- – Verify deletion hold.
- – Clean tape heads regularly to keep media from wearing out too soon, thus keeping tape media in the rotation longer.

1.3 Recommended resources for study

Courses and publications are offered to help you prepare for the certification tests. The courses are recommended, but not required, before taking a certification test. If you wish to purchase Web-based training courses or are unable to locate a Web-based course or classroom course at the time and location you desire, please feel free to contact one of our delivery management teams at:

► Americas:

tivamedu@us.ibm.com

► EMEA:

tived@uk.ibm.com

► Asia-Pacific:

tivtrainingap@au1.ibm.com.

> **Note:** Course offerings are continuously being added and updated. If you do not see the courses below listed in your geographical location, please contact the delivery management team.

1.3.1 Courses

This section provides information about pricing, scheduling, and course registration. Course names and numbers vary depending on the education delivery arm used in each region. Please refer to the Tivoli software education Web site to find the appropriate course and education delivery vendor for each geography.

General training information can also be found at this Web site:

http://ibm.com/training

Course: IBM Tivoli Storage Manager 5.4 Implementation and Administration Part One

This five-day instructor-led course provides hands-on practice in implementing a data management strategy with IBM Tivoli Storage Manager 5.4. The course begins with installing IBM Tivoli Storage Manager 5.4 Server, backup/archive clients, and administrative clients. Once the basic installation is complete, students gain practice with using the commands needed to install and configure a tape library locally attached to IBM Tivoli Storage Manager server and to prepare media for use in a device library.

After devices are configured, students create simple storage pools and subsequently learn to configure hierarchies of storage pools to manage the flow of data from disk to tape media. Next, storage pools are mapped to the data sent by client nodes through IBM Tivoli Storage Manager policies that provide the rules to implement data management service level agreements. Students define new policy domains, policy sets, management classes, backup copy groups, and archive copy groups.

To authorize administrators to perform specific storage tasks, students register administrators and grant the privileges required to support specified responsibilities. Students back up and restore client data, as well as manage long-term and off-site data management through archiving and retrieving client data.

After learning to back up, restore, archive, and retrieve data, students practice automating client operations by creating new IBM Tivoli Storage Manager schedules. To monitor client operations, students subsequently learn to configure and manage storage events. Finally, students will learn how to protect their IBM Tivoli Storage Manager server by automating the backup and restore of the IBM Tivoli Storage Manager databases and storage pools from off-site data storage pools.

The excercises are available for Windows® 2003 and AIX® 5.3. The lecture material covers all supported Intel® and UNIX® Platforms.

This course should be followed by IBM Tivoli Storage Manager 5.4 Implementation and Administration Part Two.

Course duration: 5 days

Author Comment: Need to establish if there is a course listing out yet.
The URL for Part one is as follows:
http://www-306.ibm.com/software/tivoli/education/H323002F21863Q31.html
as there is no course number present.

Course number: TSM for Windows -> TS573 (US); TS57G (WW) TSM for AIX -> TSX73 (US); TS57G (WW)

Course: IBM Tivoli Storage Manager 5.4 Implementation and Administration Part Two

This five-day course is designed and intended to provide in-depth administration training with IBM Tivoli Storage Manager 5.4 Extended Edition, which requires a skill level that is beyond the Implementation and Administraton Part 1 course. Through the use of reporting tools, administrators will be able to control their TSM environment, and determine which actions to take. These tasks will vary according to the action we are trying to accomplish. The course includes installing and configuring multiple servers, server-to-server and virtual volumes. The students will participate in group problem determine exercises, and be presented with options for a tapeless backup, optimal data management, and tuning the backup-archive client and the TSM 5.4 server for optimal performance.

Course duration: 5 days

<div style="border: 2px solid red;">

Author Comment:The URL Link for the above course is as follows:
http://www-306.ibm.com/software/tivoli/education/R583421F56278Y99.html
as there is no course number present

</div>

Course number: Refer to Author comment above

Geo education page: Worldwide schedules are available at the Tivoli software education Web site:

http://www.ibm.com/software/tivoli/education/

IBM PartnerWorld® "You Pass We Pay": This course is approved for IBM PartnerWorld You-Pass, We-Pay; for more information please visit:

http://www.ibm.com/partnerworld/pwhome.nsf/weblook/index.html?OpenDocument

Self Study: IBM Tivoli Storage Manager V5.4 product information and related links

Detailed information and documentation about IBM Tivoli Storage Manager V5.4 can be found at the Web site listed below. Some of the documents contained on this Web page include guides, datasheets, features, advantages, benefits and flyers. Reading these documents in addition to hands-on experience and skills with the product will help prepare a candidate for certification testing.

IBM Tivoli Storage Manager V5.4
For more information go to:

http://www.ibm.com/software/tivoli/products/storage-mgr/

1.3.2 Publications

IBM Tivoli Storage Manager manuals and Redbooks are useful tools for preparing to take Test 899.

IBM Tivoli Storage Manager product manuals
You may want to refer to the following manuals:

► *IBM Tivoli Storage Manager for AIX Installation Guide Version 5.4*, SC32-0134-00

► *IBM Tivoli Storage Manager for AIX Administrator's Guide Version 5.4*, SC32-0117-00

► *IBM Tivoli Storage Manager for Windows Installation Guide Version 5.4*, SC32-0138-00

► *IBM Tivoli Storage Manager for Windows Administrator's Guide Version 5.4*, SC32-0121-00

► *IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide V5.4*, SC32-0145-00

► *IBM Tivoli Storage Manager for Windows Backup-Archive Clients Installation and User's Guide Version 5.4*, SC32-0146-00

To obtain the online publications for IBM Tivoli Storage Manager, visit this Web site:

http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp

IBM Tivoli Storage Redbooks

> **Note:** Some of the IBM Redbook Publications listed below still refer to IBM Tivoli Storage Manager Version 5.3, the content in these books are still valid for IBM Tivoli Storage Manager 5.4 and is very useful

The following are Tivoli Storage related Redbooks:

► *IBM Tivoli Storage Management Concepts,* SG24-4877-04

This IBM Redbook publication describes the features and functions of IBM Tivoli Storage Manager. It introduces Tivoli storage management concepts for those new to storage management, in general, and to IBM Tivoli Storage Manager, in particular.

This easy-to-follow guide gives a broad understanding of IBM Tivoli Storage Manager software, the key technologies to know, and the solutions available to protect your business. It offers a broad understanding of how IBM Tivoli Storage Manager will work in heterogeneous environments including Windows, UNIX/Linux®, OS/400®, and z/OS® platforms, and with such applications as DB/2, Oracle®, Lotus® Domino®, Exchange, SAP®, and many more.

The book introduces storage management software by explaining the concepts, architecture, and systems management features of IBM Tivoli Storage Manager and showing available complementary products. It will help you design solutions to protect data holdings from losses ranging from those caused by user error to complete site disasters.

► *IBM Tivoli Storage Manager Implementation Guide*, SG24-5416-03

This IBM Redbook publication describes how to integrate, install, configure, and operate IBM Tivoli Storage Manager Version 5.3 in heterogeneous environments.

You will learn how to implement and operate IBM Tivoli Storage Manager. You should already have a conceptual understanding of IBM Tivoli Storage Manager. We will show you how to set up and implement the software, covering basic and advanced topics for Windows, AIX, and Linux based operating system platforms.

We demonstrate how to handle all important tasks to protect your business: Planning, client and server installation, operations, performance considerations, SAN environments, NDMP, as well as explaining many more features. This practical guide is intended for the following audiences: System administrators, new to IBM Tivoli Storage Manager, who are asked to commence a basic IBM Tivoli Storage Manager implementation for the very first time, and administrators who want to learn more about the basic and advanced components and their implementation.

► *IBM Tivoli Storage Manager Version 5.3 Technical Guide,* SG24-6638-00

This IBM Redbook publication presents an overview of IBM Tivoli Storage Manager Version 5.3, giving detailed descriptions of the changes provided in this new release. This book also covers the cumulative changes in the releases after Version 5.1.

This book is intended for customers, consultants, IBM Business Partners, and IBM and Tivoli staff who are familiar with earlier releases of IBM Tivoli Storage Manager and who want to understand what is new in Version 5.3. It should be used in conjunction with the manuals and readme files provided with the products and is not intended to replace any information contained therein.

► *Certification Study Guide: IBM Tivoli Storage Manager Version 5.3,* SG24-7127-00

This IBM Redbook publication is a study guide for IBM Tivoli Storage Manager Version 5.3 and is meant for those who want to achieve IBM Certifications for this specific product.

The IBM Tivoli Storage Manager Certification, offered through the Professional Certification Program from IBM, is designed to validate the skills required of technical professionals who work in the implementation of the IBM Tivoli Storage Manager Version 5.3 product.

This book provides a combination of theory and practical experience needed for a general understanding of the subject matter. It also provides sample questions that will help in the evaluation of personal progress and provide familiarity with the types of questions that will be encountered in the exam.

► *Deployment Guide for IBM Tivoli Storage Manager Version 5.3,* SG24-6761-00

Deploying a storage management solution for a small or medium-size business begins with a thorough analysis of the existing business and IT environment. The next step after fully understanding the organization is planning and sizing. This is the key stage, because deployment success depends on how well you plan and size.

This IBM Redbook takes a step-by-step approach to deploying a storage management solution based on IBM Tivoli Storage Manager. The first two chapters introduce overall planning and sizing issues. Then we focus on the detailed installation and configuration of Tivoli Storage Manager, and offer some deployment scenarios and troubleshooting of storage management implementation.

Readers should be familiar with the following topics:

- Storage management concepts

- Network topologies

- Distributed systems architectures and configuration

► *Deployment Guide Series:IBM Tivoli Storage Manager Version 5.4,* SG24-7379-00

The objective of this IBM Redbook publication is to provide broad instructions about deploying the IBM Tivoli Storage Manager to various environments.

Readers must have general knowledge about communication network architecture and design, basic sizing of the Tivoli Storage Manager database, and basic pool management of TSM servers.

This document is intended to be read and used by pre-sales systems engineers and services personnel to build customized deployment of the Tivoli Storage Manager. A significant amount of knowledge of Tivoli Storage Manager is expected, and ideally the reader should have attended TSM basic and advanced training classes.

► Using the IBM System Storage N Series with IBM Tivoli Storage Manager, SG24-7243-00.

This IBM Redbook gives detailed descriptions and setup instructions, practical examples and best practises for backing up the IBM System Storage N series using IBM Tivoli Storage Manager.

2

Introduction to IBM Tivoli Storage Manager

In this chapter we provide an overview of IBM Tivoli Storage Manager concepts. This includes a high-level technical introduction to IBM Tivoli Storage Manager, its architecture, and base concepts and key changes in version 5.4. However this chapter dicusses part of the first section of the IBM Tivoli Storage Manager 5.4 Implementation Certification test objectives.

. The following topics are discussed:

► IBM Tivoli Storage Manager overview
► IBM Tivoli Storage Manager architecture
► Features of IBM Tivoli Storage Manager, Express Edition
► Features of IBM Tivoli Storage Manager, Basic Edition
► Features of IBM Tivoli Storage Manager Extended Edition
► IBM Tivoli Storage Manager Data complementary products
► What's new in Version 5.4

2.1 IBM Tivoli Storage Manager

Tivoli Storage Manager is the number one product of choice to accomplish an efficient and effective enterprise wide storage solution in many companies. It provides a solution for data protection, archiving, disaster recovery planning, space management, database and application protection bare machine recovery, and record retention. More than 44 operating platforms are supported, using a consistent graphical user interface.

Tivoli Storage Manager provides:

► Centralized administration for data and storage management

► Fully automated data protection

► Efficient management of information growth

► High-speed automated server recovery

► Full compatibility with hundreds of storage devices, and local area network (LAN), wide area network (WAN), and storage area network (SAN) infrastructures

► Optional customized backup solutions for major groupware, enterprise resource planning (ERP) applications, and database products

Tivoli Storage Manager is the premier choice for complete storage management in mixed platform environments.

2.1.1 IBM Tivoli Storage Manager overview

In the world today, data has become the key asset of companies and one of its most important competitive differentiating factors. Temporary inaccessibility or the complete loss of data has a huge financial impact, and can drive companies out of business. The inability to manage data can have a negative impact on a company's profitability and can limit their ability to grow.

Tivoli Storage Manager protects an organization's data against hardware failures and other errors by storing backup and archive copies of data in offline storage. It can scale to protect hundreds of computers ranging from notebooks (mobile computers) to mainframes, running a variety of different operating systems (OSs), connected via the Internet, WANs, LANs or SANs. Centralized Web-based management, smart data move-and-store techniques, and comprehensive policy-based automation work together to minimize data protection administration costs and the impact on both computers and networks.

Storing, protecting, and managing data growth are now among the major challenges of today's businesses, requiring solutions that go beyond traditional backup and recovery solutions.

The base functions provided by Tivoli Storage Manager and its complementary products are as follows:

► Data protection (periodic backup and restore as well as disaster recovery):

 – In operational backup and restore of data, the backup process creates a copy of the data to protect against the operational loss or destruction of file or application data. The customer defines how often to backup (frequency) and how many copies (versions) to hold.

    The restore process places the backup copy of the data back into a customer-designated system or workstation.

 – Disaster recovery refers to all the activities required to organize, manage, and automate the recovery process from a major loss of IT infrastructure and data across the enterprise. It includes processes to move data off-site into a secure vault location, to rebuild IT infrastructure, and to reload data successfully in an acceptable time frame.

► Data resource management (vital record retention, archive, and retrieval):

 – The archive process creates a copy of a file or a set of files representing an endpoint of a process for long-term storage. Files can remain on the local storage media or can be deleted. The customer controls how long (via the retention period) an archive copy is to be retained.

 – The retrieval process locates the copies within the archival storage and places them back into a customer-designated system or workstation.

► Space management (or Hierarchical Storage Management):

    This process provides the automatic and transparent movement of operational data from the user system disk space to a central storage repository. If the user accesses this data, it is dynamically and transparently restored to the client storage.

The solution is network based, which means that these functions are available to the whole network environment. Administration costs are minimized by centralization of all of the management of Tivoli Storage Manager components.

Tivoli Storage Manager is available in three editions; Express, Basic Edition, and Extended Edition.

2.1.2 IBM Tivoli Storage Manager architecture

IBM Tivoli Storage Manager is a client server architecture, as seen in Figure 2-1.

IBM Tivoli Storage Manager clients are the workstations, file servers, mobile computers, and other machines that must have their data protected. IBM Tivoli Storage Manager client software is installed on these systems.
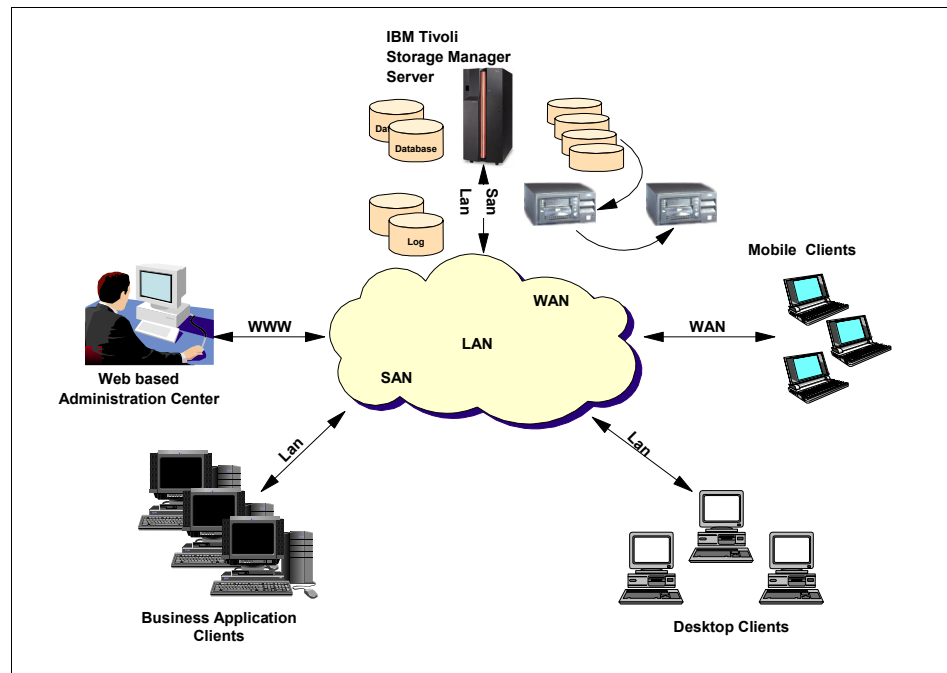

Figure 2-1   IBM Tivoli Storage Manager architecture

Tivoli Storage Manager is based on a relational database and transaction log. The database and transaction log track metadata, what is backed up, where it is stored, and what are the policies, schedules, and administrators.

The transaction log enables a two-phase commit, which protects the integrity of the database and allows for interrupted backups and restores to be restarted. The relational database empowers Tivoli Storage Manager to perform tasks that are not possible when you use a flat file master catalog to track metadata. For example, the relational database can:

► Move data from one type of storage pool to another
► Retroactively update backed-up data when a policy changes
► Track individual files
► Schedule any type of client or administrative process
► Reclaim expired dead space on tapes

The Tivoli Storage Manager client sends its data to the Tivoli Storage Manager server either by the LAN or by the SAN. Most backups occur through schedules,

but clients can perform on demand backups whenever they want. Clients can also perform their own restores. Tivoli Storage Manager has a new Administration Center that was introduced in V5.3 and can be installed either on the same machine as the Tivoli Storage Manager Server or on a separate machine.

The actual data that the client sends is stored in the storage pools. Tivoli Storage Manager is unique in the fact that the storage pools can form a storage hierarchy made up of more than 500 supported devices. This allows for flexibility, longevity, and, most important, fast backups and fast restores.

Most businesses back up their data initially to disk storage. This allows for hundreds of clients to back up at the same time. Then, based on policies, the data migrates in a fashion that expedites restores to tape or CD. When the data migrates, all data belonging to one client is moved together to the next pool. By keeping all of that data together, restores are faster because not as much tape positioning is required. This migration process can also accommodate movement to collocated tapes, which further expedites restores by just having one user's data on them.

The environment can be firewall protected, but you will still want to be able to use the GUI interfaces of a Tivoli Storage Manager client. Tivoli Storage Manager allows individual configuration of nearly every TCP port that it uses for communication:

► TCP/IP port:

   To enable the Backup/Archive client, command line Administrative Client, and the scheduler to run outside a firewall, the port specified by the tcpport server option must be opened by the firewall administrator. This port is set on the client and the server using the tcpport option. The setting must be the same on the client and server. The default TCP/IP port is 1500.

► TCP/IP ports for the remote workstation:

   The two TCP/IP ports for the remote workstation client must be opened. Use the WEBPORTS option in the remote workstations option file to specify these ports. If you do not specify the values for the Web ports option, the default zero (0) causes TCP/IP to randomly assign two free port numbers.

2.1.3 IBM Tivoli Storage Manager server

One of the principal architectural components of the IBM Tivoli Storage Manager server is its built-relational database. The IBM Tivoli Storage Manager database was especially designed for the task of managing data, and it implements zero-touch administration.

All policy information, logging, authentication and security, media management, and object inventory are managed through this database.

Most of the fields are externalized through IBM Tivoli Storage Manager high-level administration commands; SQL SELECT statements; or, for reporting purposes, by using an ODBC driver. Obviously, this database is fully protected with software mirroring, roll-forward capability, and its own management and online backup and restore functions.

For storing the managed data, the IBM Tivoli Storage Manager server manages a storage repository. The storage repository can be implemented in a hierarchy using any combination of supported media or magnetic or optical disk, tape, and robotic storage devices, which are locally connected to the server system or are accessible through a SAN. To take advantage of SAN technology, the IBM Tivoli Storage Manager server has features implemented that dynamically share SAN-connected, automated tape library systems among multiple IBM Tivoli Storage Manager servers, as well as provide (as an option) local area network, or LAN-free and server-free backup.

2.1.4 IBM Tivoli Storage Manager Backup/Archive client

Data is sent to the IBM Tivoli Storage Manager server using the IBM Tivoli Storage Manager Backup/Archive client and complementary Tivoli and non-IBM/Tivoli products. These products work together with the IBM Tivoli Storage Manager server base product to ensure that the data you need to store is managed as defined.

The IBM Tivoli Storage Manager Backup/Archive client, included with the server, provides the operational backup and archival function. The client implements the patented progressive backup methodology, adaptive sub-file backup technology, and unique record retention methods for backup and archive functions.

The Backup/Archive clients are implemented as multi-session clients, which means that they are able to take advantage of the multi-threading capabilities of modern operating systems.

2.1.5 IBM Tivoli Storage Manager storage agent

The IBM Tivoli Storage Manager storage agent supports LAN-free backup solutions using a SAN infrastructure. The storage agent dynamically shares SAN connected tape libraries and disks with the IBM Tivoli Storage Manager server, and it has the ability to write and read a large amount of client data directly to and from server-owned storage media. This provides a great opportunity for lowering the backup window, reducing the traffic on the LAN, and reducing the utilization of the IBM Tivoli Storage Manager server.

2.1.6 IBM Tivoli Storage Manager Express Edition

IBM Tivoli Storage Manager Express is a new product aimed at two market segments: the small to medium business with a less sophisticated IT environment, or the enterprise department that does not require the full suite of Tivoli Storage Manager features.

IBM Tivoli Storage Manager Express provides a subset of Tivoli Storage Manager features, focusing on backup and recovery for between 5 to 20 client machines. The features of IBM Tivoli Storage Manager Express are:

► Easy installation: IBM Tivoli Storage Manager Express takes less than one hour to install, configure, and start running backups.

► Simplified administration GUI: a new GUI simplifies administration and operational reporting is integrated. Client software deployment is also included.

► Fully upgradeable: up to IBM Tivoli Storage Manager Extended Edition.

► Disk-based incremental backup: client backups are done to disk storage pools on the IBM Tivoli Storage Manager Express server. You have the option to use tape devices for longer term retention or off-site backups.

► Simplified tape management: use of traditional methods such as Grandfather/Father/Son backup sets simplify tape rotation, and all tape management is fully automated.

► Automatic configuration: clients are automatically configured with scheduled backups using industry best practices.

IBM Tivoli Storage Manager Express supports:

► Windows 2003 as the platform for the Tivoli Storage Manager Express server
► From 5 to 20 client systems
► A database size of up to 20 GB
► LAN-based systems and devices
► MS Exchange and SQL Server optional backup
► LTO, DLT, 4 mm DDS, and Sony 8 mm AIT devices

For more information about IBM Tivoli Storage Manager Express, please refer to the IBM Redbook, IBM Tivoli Storage Manager Express Deployment Guide, and visit the following Web site:

http://www.ibm.com/software/tivoli/products/storage-mgr-express/

2.1.7 IBM Tivoli Storage Manager Basic Edition

IBM Tivoli Storage Manager Basic Edition contains a rich set of features and provides the core functions of backup, recovery, and archive management.

► Progressive backup methodology

  Saves time and storage space by backing up only new files and modified files. The progressive backup feature uses Tivoli Storage Manager's own relational database to track data wherever it is stored, delivering a direct one-step file restore. Progressive backup eliminates the requirement for traditional full-plus-incremental or full-plus-differential backup and restore procedures, commonly used by other storage management products.

► Tape resource sharing

  Enables multiple Tivoli Storage Manager servers to use the same tape library and drives, optimizing tape hardware asset utilization.

► Network-free rapid recovery

  Supports high-speed client data recovery directly from tape or optical devices. Recovery time is minimized by eliminating the use of network and central server resources.

► Dynamic multithreaded transfer

  Permits multiple clients to simultaneously transfer data to and from the same Tivoli Storage Manager server. Performance is boosted to more than three times the rate of a single-threaded session. The higher speed is achieved by transparently optimizing the number of data transfer sessions, based on available system resources.

► Adaptive differencing technology

  Changes the way data is backed up from the client. Using adaptive differencing, data is transferred to the server either by byte, block, or file level, based on the size of the file being backed up, and the portion of the file that has changed since the last backup. Adaptive differencing technology supports all connectivity strategies, including LANs, WANs, SANs, Internet, and dial-up connections. Adaptive differencing was initially designed with mobile computer users in mind. However, other users with a requirement to minimize data transmitted over the network can also benefit from the technology.

► Enterprise administration

  Simplifies centralized control across multiple Tivoli Storage Manager implementations without sacrificing network performance. Tivoli Storage Manager V5.4 is also powered by the Integrated Solutions Console (ISC) which provides a task-based GUI interface to Tivoli Storage Manager administrative tasks.

► Clustering

Tivoli Storage Manager includes enhanced support for IBM High Availability Cluster Multi-Processing (HACMP™), Microsoft® Cluster Services (MSCS), Novell Cluster Services (NCS) and VERITAS Cluster Services (VCS) on Windows.

Tivoli Storage Manager V5.4 has improved the support for Small Computer System Interface (SCSI) and fibre-attached tape device failover on Windows and UNIX, and support for Storage Agents, Library Managers, and Library Clients as cluster members.

► LAN-free data transfer

An optional module for Tivoli Storage Manager effectively exploits SAN environments by moving data transfers from the communication network to a SAN. Communication bandwidth availability is therefore improved, increasing service levels for users and customers.

► Hierarchical Storage Management

An optional module for Tivoli Storage Manager automatically and transparently moves unused data files from online disk storage to offline tape storage. In the event that a file is accessed after it has been moved to offline storage, Tivoli Storage Manager transparently recalls the file.

► Library and device support

Tivoli Storage Manager Basic Edition supports libraries with up to three tape drives and up to 40 cartridge capacity. Larger libraries can be accommodated but with only three devices and 40 slots enabled.

You can find more information about IBM Tivoli Storage Manager Basic Edition at the Web site:

http://www.ibm.com/software/tivoli/products/storage-mgr/

2.1.8 IBM Tivoli Storage Manager Extended Edition

The Extended Edition of IBM Tivoli Storage Manager expands on the features and possibilities of the Basic Edition described in the previous section.

Tivoli Storage Manager Extended Edition adds disaster recovery planning capability for the server, Network Data Management Protocol (NDMP) control for network-attached storage (NAS) filers, and support for larger capacity tape libraries and more tape drives.

You can find more information at:

http://www.ibm.com/software/tivoli/products/storage-mgr-extended/

Disaster Recovery Manager

The Disaster Recovery Manager (DRM) component of Tivoli Storage Manager Extended Edition provides disaster recovery for the Tivoli Storage Manager server and assists with disaster recovery for clients.

DRM offers various options to configure, control, and automatically generate a disaster recovery plan (DRP) file. The plan contains the information, scripts, and procedures required to automate restoration and help ensure quick recovery of data after a disaster. The scripts contain the commands necessary to rebuild the Tivoli Storage Manager server.

One of the key features of Tivoli Storage Manager and DRM is the ability to track media in all possible states, such as on-site, in transit, or in a vault. The media movement features of DRM assist greatly with the daily tasks of sending disaster recovery media off-site, and receiving expired media on-site for reuse. With these features the system administrator can quickly locate all available copies of data.

DRM functions help maintain business continuity by:

► Establishing and helping to automate a thorough server DRP: clients can then subsequently restore their data from the server if required, and can continue their daily backup procedures.

► Ensuring that vital site-specific information is available in the same plan.

► Automating vital recovery steps to return the Tivoli Storage Manager server and backup environment to normal operation.

► Managing and identifying off-site media required for recovery.

► Tracking and reporting destroyed systems in the event of a disaster.

► Storing client configuration information and assigning client recovery priorities.

With DRM you can recover at an alternate site, on a replacement system with a different hardware configuration, and with people who are not familiar with the applications. The DRP can be periodically tested to certify the recoverability of the server. The DRP can, and must, be recreated easily every day so that it stays up to date. Figure 2-2 illustrates the main functions of DRM.
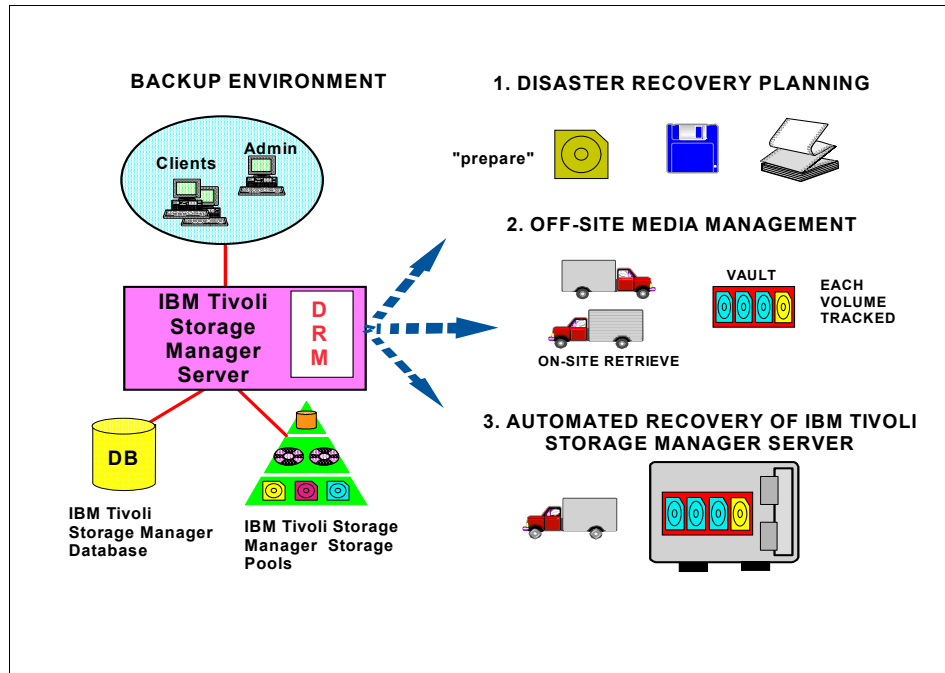
Figure 2-2   IBM Tivoli DRM functions

During a disaster, these are some commonly encountered errors:

► A DRP does not exist.
► The DRP was not tested, or if it was, it is now out of date.
► The testing team's skills were not sufficient to perform and evaluate testing.
► The disk volume definitions for the recovery site are not known.
► The location of recovery tapes is not known.
► It is not known which tapes are to be applied first.

DRM keeps track of all the vital information required to rebuild the Tivoli Storage Manager environment, such as:

► The current server configuration information and its location
► The current Tivoli Storage Manager server database volumes (size, location, number)
► The recovery sequence
► The currency of the DRP
► The server and client machines configurations
► The people to be contacted in the event of a disaster

► The location of the recovery media and the organization or persons responsible

► The point in time (PIT) to which the environment can be restored

During recovery from a disaster, DRM automates the following procedures to restore the Tivoli Storage Manager servers:

► Restore Tivoli Storage Manager server's key option files

► Copy files from alternate locations to production locations

► Initialize the Tivoli Storage Manager database and log volumes

► Match sizes and locations of the Tivoli Storage Manager database and log volumes

► Automatically launch restoration of the Tivoli Storage Manager database

► Track media required and availability

► Register installed Tivoli Storage Manager server features and return the server state to a valid license configuration

► Update Tivoli Storage Manager volume catalog information, including whether volumes have been destroyed during the disaster

► Rebuild Tivoli Storage Manager hierarchical storage configuration

► Restore destroyed volumes from those available where possible

► Recreate customer backup environment

A detailed description, recovery scenario, and recovery plan built with DRM can be found in Disaster Recovery Strategies with Tivoli Storage Management, SG24-6844. Also, recommendations and examples of using DRM to store client machine information in the DRM plan file for use during a client disaster recovery are given in the same book.

In summary, DRM will systematically rebuild the Tivoli Storage Manager server environment and ensure that current application data for the entire enterprise is available for recovery. This can all be done automatically from a single scripted command.

NDMP support for Network Attached Storage

For NAS devices, Tivoli Storage Manager Extended Edition uses NDMP to perform high-performance, scalable backups and restores. NDMP-based backups and restores minimize network traffic and transfer data outboard of the Tivoli Storage Manager client and server. NDMP enables a full and differential file system image backup and restore of Network Appliance™ file servers with OS Data ONTAP® 7.1 or higher, and EMC Celerra systems. Multiple backup and

restore operations can be performed simultaneously. General NDMP support also allows other NAS vendors to certify integration with Tivoli Storage Manager.

The NDMP backup and restore features are fully integrated with Tivoli Storage Manager Extended Edition server and client. No extra software is required on the server, client, or NAS appliance. When doing backups and restores, the NAS device and the Tivoli Storage Manager server and client all have specific roles, as shown in Figure 2-3.
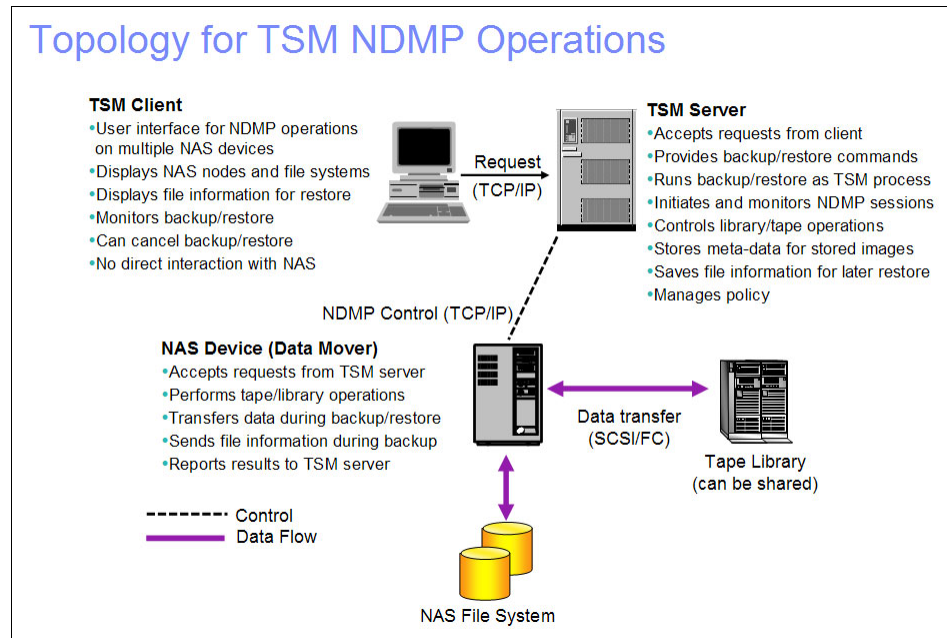

Figure 2-3   Topology for NDMP using IBM Tivoli Storage Manager

Tivoli Storage Manager Extended Edition offers the ability to do file-level and full/differential file system image backups and restore of servers that support the NDMP protocol. Now you can back up directly to the Tivoli Storage Manager hierarchy and also implement DRM as it now supports NAS storage. Multiple backup and restore operations can be performed in parallel.

During backup and restore operations, data flows directly between the tape drive and the NAS appliance. NDMP for NAS backup uses either an SCSI-attached tape device local to the NAS appliance, or a SAN-attached SCSI or Automated Cartridge System Library Software (ACSLS) device that can be shared with the Tivoli Storage Manager server. Library robotics can be controlled directly by the Tivoli Storage Manager server or by passing SCSI commands via an NAS file server.

Drives must be supported by both the NAS appliance and the NAS OS. Drives can be dedicated to NDMP operations from a single NAS file server or can be shared. Multiple NAS appliances can share SAN-attached shared tape resources if backups are performed via the same Tivoli Storage Manager server. Drives can be also shared with LAN-free backup/restore operations, provided that the library is controlled directly by the Tivoli Storage Manager server.

Extended library and drive support

Tivoli Storage Manager Extended Edition supports larger tape libraries, thus removing the 40-cartridge limit for library capacity, and allowing more than three tape drives within a single library

Database and application online protection

This feature protects a wide range of application data via the protection of the underlying databases and application management systems holding that data. This module automates data protection tasks and allows database and application servers to continue running their primary applications while they back up and restore data to and from offline storage.

Data retention

IBM Tivoli Storage Manager for Data Retention helps manage and simplify the retrieval of the ever-increasing amount of data that organizations must retain for strict records retention regulations. Many regulations demand archiving of records, e-mails, design documents, and other data for many years, in addition to requiring that the data is not changed or deleted.

Bare Machine Recovery

This feature backs up and automatically restores the operating system structures required to rebuild the operating system, as well as data files. It schedules regular operating system backups so tha a recovery brings back the latest information.

2.2 IBM Tivoli Storage Manager complementary products

IBM Tivoli Storage Manager can be integrated with several optional applications that together form a powerful integrated storage management solution. These include:

- ▶ IBM Tivoli Storage Manager for Space Management
- ▶ IBM Tivoli Storage Manager for HSM for Windows
- ▶ IBM Tivoli Storage Manager for Storage Area Networks

- ► IBM Tivoli Storage Manager for System Backup and Recovery
- ► IBM Tivoli Storage Manager for Data Protection product family
- ► IBM Tivoli Continuous Data Protection for Files
- ► IBM System Storage^{TM} Archive Manager

For a full product listing, visit:

http://www.ibm.com/software/tivoli/products/

2.2.1 IBM Tivoli Storage Manager for Space Management

IBM Tivoli Storage Manager for Space Management provides hierarchical storage management (HSM) to automatically migrate rarely-accessed files to alternative storage, without disrupting the most frequently used files in local storage. Migrated files are automatically and transparently recalled to primary storage when required by applications or users. Administrators and users are freed from manual file system maintenance tasks, and more online disk space is available for more important active data. Tivoli Storage Manager for Space Management can also help defer the requirement to purchase additional disk storage for clients, by making optimal use of available client storage.

Tivoli Storage Manager for Space Management offers increased scalability and performance via parallel migrations, improved candidate search, and optimized synchronization between the IBM Tivoli Storage Manager server and the HSM client.

IBM Tivoli Storage Manager for Space Management complements both IBM Tivoli Storage Manager and IBM Tivoli Storage Manager Extended Edition, and is supported on AIX, HP/UX, Solaris™, and Linux.

2.2.2 IBM Tivoli Storage Manager for HSM for Windows

IBM Tivoli Storage Manager for HSM for Windows is a new product that provides hierarchical storage management functionality to the Windows platform. As with IBM Tivoli Storage Manager for Space Management, HSM for Windows automatically migrates rarely-accessed files to alternative storage, without disrupting the most frequently used files in local Windows file systems. Similarly, migrated files are automatically and transparently recalled to their original location when required by applications or users.

HSM for Windows allows various levels of granularity for migration of files. Files can be migrated individually, and file systems can be partially or fully migrated, based on a comprehensive set of policy options.

IBM Tivoli Storage Manager for Space Management complements both IBM Tivoli Storage Manager and IBM Tivoli Storage Manager Extended Edition.

2.2.3 IBM Tivoli Storage Manager for Storage Area Network

IBM Tivoli Storage Manager for Storage Area Network enables your SAN-connected Tivoli Storage Manager servers and client computers to make maximum use of their direct network connection to storage. This software extension enables both servers and client computers to make the bulk of their backup/restore and archive/retrieve data transfers over the SAN instead of the LAN, either directly to tape or to the Tivoli Storage Manager disk storage pool. This ability greatly reduces the impact of data protection on the LAN while also reducing CPU utilization on both client and server.

For computers running Windows, some SAN configurations allow specific SAN devices to perform data movements directly to and from some tape devices, further reducing client and server CPU utilization.

Tivoli Storage Manager for Storage Area Networks complements and coexists with the standard library-sharing functionality of both Basic and Extended editions of the Tivoli Storage Manager server.

The core functions of IBM Tivoli Storage Manager for Storage Area Network are:

► LAN-free backup/restore
► SAN-connected tape library

2.2.4 IBM Tivoli Storage Manager for Backup and Recovery

IBM Tivoli Storage Manager for System Backup and Recovery (SysBack™) provides a flexible backup method for AIX systems to help protect data and provide bare machine recovery capabilities. It offers a comprehensive system backup, restore, and reinstallation tool. SysBack is a simple-to-use and highly effective tool. Any feature may be executed from either the AIX command line or by using the SMIT menu interface.

For Windows platforms, bare machine recovery can be achieved with the TSM Backup/Archive client's Automated System Recovery capability.

In addition, Windows, Sun, and Linux bare machine recovery can be done with Cristie Bare Machine Recovery. This integrates directly with Tivoli Storage Manager to provide operating system recovery for these platforms.

2.2.5 IBM Tivoli Storage Manager for data protection

IBM Tivoli Storage Manager provides data protection for a wide variety of applications, databases, mail, and hardware, ensuring that data is safe and secure no matter where it is located or how it is stored. These products interface directly with the applications using their backup-certified utilities and interfaces, simplifying online backup and restore procedures. These products are described in the following sections.

IBM Tivoli Storage Manager for Application Servers

IBM Tivoli Storage Manager for Application Servers (formerly Tivoli Data Protection for WebSphere® Application Servers) is a software module that works with IBM Tivoli Storage Manager to better protect the infrastructure and application data and improve the availability of WebSphere Application Servers.

IBM Tivoli Storage Manager for Databases

IBM Tivoli Storage Manager for Databases is a software module designed to work with IBM Tivoli Storage Manager to protect a wide range of application data through the protection of the underlying database management systems holding that data. IBM Tivoli Storage Manager for Databases exploits the various backup-certified utilities and interfaces provided for Oracle, Microsoft SQL Server, and Informix®.

This same functionality is included in the IBM DB2® Universal Database™ package and Informix Dynamic Server, enabling them to work directly with IBM Tivoli Storage Manager without the need to buy any additional modules.

IBM Tivoli Storage Manager for Hardware

IBM Tivoli Storage Manager for Hardware improves the data protection of business-critical databases and Enterprise Resource Planning (ERP) applications that require 24x7 availability.

IBM Tivoli Storage Manager for Mail

IBM Tivoli Storage Manager for Mail is a software module for IBM Tivoli Storage Manager that automates the data protection of e-mail servers running either Lotus Domino or Microsoft Exchange. This module utilizes the application programming interfaces (API) provided by e-mail application vendors to perform online backups without shutting down the e-mail server and improve data-restore performance. As a result, it can help protect the growing amount of new and changing data that should be securely backed up to help maintain Domino and Exchange application availability 24x7, 365 days a year.

IBM Tivoli Storage Manager for Enterprise Resource Planning

IBM Tivoli Storage Manager for Enterprise Resource Planning is a software module that works with IBM Tivoli Storage Manager to better protect infrastructure and application data and improve the availability of SAP R/3® servers.

IBM Tivoli Storage Manager for Advanced Copy Services

IBM Tivoli Storage Manager for Advanced Copy Services (formerly known as IBM Tivoli Storage Manager for Hardware) is an optional software module for AIX that integrates with Tivoli Storage Manager Extended Edition. Tivoli Storage Manager for Advanced Copy Services protects mission-critical data that must be available 24x7, and integrates hardware- and software-based snapshot capabilities with Tivoli Storage Manager and its Data Protection components for DB2 UDB, Oracle, and mySAP™.

Tivoli Storage Manager for Advanced Copy Services supports a wide range of hardware:

► IBM Enterprise Storage Server® (ESS)
► IBM DS6000™
► IBM DS8000™
► SAN Volume Controller (SVC) and all IBM and non-IBM devices supported by the SVC. For a complete list, see:
http://www.ibm.com/systems/storage/software/virtualization/svc/interop.html

Tivoli Storage Manager for Advanced Copy Services also provides the following functionality:

► IBM FlashCopy® support for ESS for Oracle
► FlashCopy support for ESS for DB2
► FlashCopy support for ESS for mySAP on DB2 UDB
► FlashCopy support for ESS for mySAP on Oracle
► Snapshot™ support for DS8000, DS6000 and SVC for DB2 UDB
► Snapshot support for DS8000, DS6000 and SVC for Oracle
► Snapshot support for DS8000, DS6000 and SVC for mySAP on DB2 UDB
► Snapshot support for DS8000, DS6000 and SVC for mySAP on Oracle
► Multiple snapshot versions managed by Tivoli Storage Manager policy
► Coordinated FlashCopy backup of multi-partition DB2 UDB databases distributed across multiple host systems

Support of FlashCopy and snapshot functionality allows for "Zero Impact" backups and instant recovery. Data transfer to the Tivoli Storage Manager server is handled from a separate storage server, allowing the primary production data to remain online and undisturbed.

IBM Tivoli Storage Manager for Copy Services

IBM Tivoli Storage Manager for Copy Services is a new optional module for Windows that integrates with Tivoli Storage Manager or Tivoli Storage Manager Extended Edition. It is designed to leverage Microsoft's Volume Snapshot Services (VSS) on Windows 2003. Tivoli Storage Manager for Copy Services provides similar functionality to Tivoli Storage Manager for Advanced Copy Services, but supports Windows VSS and Microsoft Exchange Server 2003 only.

The features of Tivoli Storage Manager for Copy Services are:

► Single command-line interface (CLI) for performing legacy and VSS snapshot backup, restore, and query operations

► Single GUI for performing legacy and VSS snapshot backup, restore, and query operations

► Support for both hardware and software VSS providers that strictly adhere to Microsoft VSS provider requirements

► Support for a clustered Exchange environment

Full and Copy backup types are supported, with granularity at the Exchange "Storage Group" level. backups are managed by Tivoli Storage Manager policies and can be stored on the Tivoli Storage Manager server, local disks or both. Different policies can be assigned for the different storage locations and backup types (Full or Copy). As with Tivoli Storage Manager for Advanced Copy Services, zero impact backups and instant recovery allow the primary production data to remain online and undisturbed. Data movement to Tivoli Storage Manager storage can be off-loaded to a secondary machine via a VSS hardware provider that supports transportable shadow copy volumes.

2.2.6 IBM Tivoli Continuous Data Protection for Files

According to industry surveys, almost 70% of corporate data exists on notebooks (mobile computers) or desktop machines, and less than 8% of it is backed up regularly. For notebook, desktop, and file server machines that contain important, critical, or sensitive data that is constantly being updated, a typical 24-hour backup cycle may not be sufficient to provide adequate data protection. The addition of Tivoli Continuous Data Protection for Files provides a client machine with the capability of being able, transparently in real time, to back up a file to a Tivoli Storage Manager server as soon as the file is saved. Files that are backed up by this method are managed in the same ways as other corporate data by the Tivoli Storage Manager server.

Tivoli Continuous Data Protection for Files was developed with notebook (mobile computer) and desktop users in mind, but can be applied to any client with a high rate of change of data on its file systems.

Tivoli Continuous Data Protection for Files provides clients with true point-in-time recoverability. It is supported on AIX, Solaris, Linux, and Windows platforms. For more information, see:

http://www.ibm.com/software/tivoli/products/continuous-data-protection/

2.2.7 IBM System Storage Archive Manager

IBM System Storage™ Archive Manager facilitates compliance with regulatory requirements. It helps manage and simplify the retrieval of the ever increasing amount of data that organizations must retain for strict records retention regulations. Many of the regulations demand the archiving of records, e-mails, design documents and other data for many years, in addition to requiring that the data is not changed or deleted.

IBM Tivoli Storage Manager's existing policy-based data management capabilities help organizations meet many of the regulatory requirements of various government and industry agencies. But some new regulations require additional safeguards on data retention. IBM System Storage Archive Manager provides data retention policies that help meet these new regulations.

Data retention protection

IBM System Storage Archive Manager makes the deletion of data before its scheduled expiration extremely difficult. Short of physical destruction to storage media or server, or deliberate corruption of data or deletion of the Archive Manager database, Archive Manager does not allow data on the storage managed by the IBM System Storage Archive Manager server to be deleted before its scheduled expiration date. Content management and archive applications can apply business policy management for ultimate expiration of archived data at the appropriate time.

Features and functions

IBM System Storage Archive Manager hierarchical storage capabilities provides policies, so that data is stored on the type of media that best meets that data's longevity, access speed, and cost needs.

Movement of the data from one media type to another (as media requires change, or as new types of media become available) is achieved by migration. Migration automates moving the data to help ensure data longevity, and also allows for data to be stored on the type of media that best meets its speed of access and cost needs.

► Expiration policies: expire the data when it is no longer required, thus freeing up the storage media, and providing cost effectiveness.

► Off-site data protection is standard: off-site copies can be created onto any of the hundreds of types of media supported, and like the primary copy, is policy managed to allow for expiration.

► Archive client program: permits users to archive files from their workstations or file servers to archive retention protected storage, and also retrieve archived copies of files to their local workstations.

► Expiration and deletion suspension: allows you to place an unconditional hold on data. It means that data cannot be deleted or modified until the deletion hold is released.

► Event-based retention management: data is retained based, subject to a time interval that is calculated after a retention-initiating event occurs. The data can then not be deleted until the time limit has expired. For example, you can specify to keep records for a particular employee for one year after the employee leaves the organization.

► Data retention protection: data is not deleted until the retention criteria for the object is satisfied.

For more information, visit the Web page:

http://www.ibm.com/software/tivoli/products/storage-mgr-data-reten/

2.3 What's new in version 5.4

The following features and functions have been added to IBM Tivoli Storage Manager Version 5.4.

2.3.1 Collocate of active data

In this new release, you can create a new type of storage pool for storing active versions of client backup data (archive and space-managed data are not allowed). These new pools, which are called active-data pools, provide a number of benefits as faster restores because when associated with a sequential-access disk (FILE) device type, FILE volumes are not required to be mounted, so the server does not need to position past inactive files that do not need to be restored, eliminating the requirement for disk staging pools and this reduces the size of the storage pool as only active versions are stored.

When using removable media, such as tape or optical, this allows you to reduce the number of volumes for on-site and off-site storage and fewer tapes are mounted in case of failure or disaster.

2.3.2 Hardware tape drive encryption

IBM tape device encryption is now supported for 3592 generation 2 and LTO4 drives through the method of application encryption. When enabled, Tivoli Storage Manager handles encrypting and decrypting data on tapes, according to specifications set when defining the device class. Tape device encryption provides security for data on individual tapes and protects sensitive information that might be transported off-site. The data is encrypted utilizing a stronger level of encryption by requiring 256-bit Advanced Encryption Standard (AES) encryption keys. Keys are passed to the drive by a key manager in order to encrypt and decrypt data.

2.3.3 NDMP enhancements (LAN Tivoli Storage Manager hierarchy integration)

When backing up NAS data in Tivoli Storage Manager, you now have the following new features:

► NDMP: Filer to Server backup

Storage of NAS backup data in the Tivoli Storage Manager hierarchy is now available for NAS backups over the LAN. Separate hierarchies of data are maintained for NAS data directed to Tivoli Storage Manager native pools. Data may be continued to be directed to existing NAS destinations allowing centralization of tape resources and exploiting the capacity of Tivoli Storage Manager Storage hierarchy.

Backups from NAS devices to the Tivoli Storage Manager server retain the same table of contents (TOC) creation and directory backup and restore functions as in prior releases.

► NDMP: Back-end Data Movement Support for NDMP Images

Tivoli Storage Manager now supports the backing up of primary storage pools that contain NDMP-generated backup images. NDMP tape-to-tape function now provides support for the following operations:

– Storage-pool backup and restore and volume restore
– DRM function now supports NAS storage pool

Data is stored in special, NDMP-formatted storage pools, which can be backed up to storage media that can be moved off-site for protection in case of an on-site disaster.

You also have the ability to manage snapshots backup for MySap data and restore from local snapshot version of from Tivoli Storage Manager Storage hierarchy, this is included in Tivoli Storage Manager for Advanced Copy Services.

2.3.4 Backup set enhancements

For backup sets, which is a collection of a node's data, you now have the ability to:

► Generate backup sets to a specific PIT

► Improve the tracking of backup sets by the Tivoli Storage Manager Server

► Stack multiple nodes' data on a single set of backup set volumes

► Enable clients to display contents of a backup set and allow the selection of individual files for restore

► Support image data in backup sets

► Create a backup set TOC, enabling clients to display contents of a backup set and allow the selection of individual files for restore

Group Nodes for Generating Backup Sets

In this release, you can use the new set of NODEGROUP commands to group client nodes for backup set generation. A node group is a group of client nodes that are acted upon as a single entity. If one or more node groups is specified, the server generates a backup set for each node group and places all of the backup sets together on a single set of output volumes, providing self describing tapes for off-site vaulting, tape retention and rotation.

2.3.5 Data shredding

Support for shredding is the destruction of deleted data to make it difficult to discover and reconstruct that data later. Tivoli Storage Manager now supports shredding data in random-access disk storage pools. You can perform shredding of sensitive data either automatically or manually.

After client data has been deleted, it might still be possible to recover it. For sensitive data, this condition is a potential security exposure. Shredding the deleted data increases the difficulty of discovering and reconstructing the data later.

2.3.6 Administration Center updates

The following updates were made in the Administration Center interface in order to achieve functional equivalence with the 5.2 GUI:

► Update administrator password to a group of selected servers
► Performance fixes, Admin API rewritten for streaming mode
► Command Line applet usability fixes: improve the "expert user" experience

Added support for:

► AUDIT LIBRARY command
► AUDIT VOLUME command
► QUERY NODEDATA and MOVE NODEDATA command
► QUERY MEDIA and MOVE MEDIA

Administration Center updates to support the new function in the following:

► Use the Tivoli Storage Manager server as the NDMP Tape Server (network filer-to-Tivoli Storage Manager server backup)

► Create media for off-site vaulting from NDMP-generated images

► Destruction of expired data (data shredding) for random-access disk

► Tape Drive Encryption

► Optical Device Support on Linux

► Administration Center tutorials translated

2.3.7 Tivoli Storage Manager Express upgrade to Tivoli Storage Manager enterprise

Tivoli Storage Manager Express users can now upgrade to Tivoli Storage Manager enterprise. This upgrade allows you to take advantage of one or more of the advanced features available in Tivoli Storage Manager. After upgrading, you can either continue using the Tivoli Storage Manager Express interface (CLIENTVIEW=EXPRESS), or you can use the standard Tivoli Storage Manager interface, which contains the enhancements necessary to support the backup set functionality found in Tivoli Storage Manager Express.

2.3.8 Basic Vista support

Tivoli Storage Manager supports backing up and restoring system files on Windows Vista™. ASR is not supported at this time.

2.3.9 Query Disk Space in a file device class

The new command QUERY DIRSPACE allows you to display the amount of total and available disk space for each directory in a FILE device class.

2.3.10 Dropped support for Windows 2000

Beginning with this release of Tivoli Storage Manager, we have dropped support for Windows 2000. The supported Windows OSs are Windows XP, 2003 and Windows Vista.

2.3.11 New schedule options

Three new options: srvprepostscheddisabled, srvprepostssnapdisabled, and schedrestretrdisabled, are now supported, in addition to the existing schedcmddisabled option. These options help to prevent Tivoli Storage Manager administrators from executing inadvertent or malicious operations on Tivoli Storage Manager client nodes.

2.3.12 Reducing the memory used during the incremental backup

The memory efficient diskcachemethod option uses local disk space, which reduces the amount of memory required to complete an incremental backup.

The diskcachelocation option specifies the location where the disk cache database is created if the option memoryefficientbackup=diskcachem is set during an incremental backup.

> **Note:** For all the information given here, you can find a more detailed version in the Administrator's Guide. Also, refer to the book *IBM Tivoli Storage Management Concepts*, SG24-4877 for more in depth information.

2.3.13 Summary

In this chapter we have discussed the overview, architecture, complementary products, and new functions of Tivoli Storage Manager. Now we move on to the planning and sizing for deployment.

3

Requirements and IBM Tivoli Storage Manager functionality

In this chapter we discuss the prerequisites of implementing IBM Tivoli Storage Manager 5.4. These include current hardware and software requirements. We also describe some of the features, functions, benefits, and advantages of Storage Manager.

This chapter provides the following:

► What you should know about the first section of the test objectives
► Hardware and software prerequisites
► Basic components of IBM Tivoli Storage Manager
► Functions, features, and benefits of IBM Tivoli Storage Manager

3.1 What to know about the first section

In this chapter, we continue to discuss the first section of the IBM Tivoli Storage Manager 5.4 Implementation Certification test objectives, which covers customer requirements and IBM Tivoli Storage Manager functionality. At the end of this section, you will able to complete the following tasks:

► Identify the current customer environment.

► Determine whether Storage Manager supports the current environment.

► Identify customer Storage Manager server platform preferences.

► Identify the scope of their needs.

► Determine the current and expected operating parameters and constraints.

► Determine a customer's functional requirements for IBM Tivoli Storage Manager.

► Determine features, benefits, and advantages of IBM Tivoli Storage Manager.

3.2 Tivoli Storage Manager System requirements

In this section we discuss the hardware and software requirements required to install and configure IBM Tivoli Storage Manager Extended Edition Version 5.4.

3.2.1 Windows environment requirements

On Windows, the Tivoli Storage Manager server requires the following hardware and software:

► Server software requirements:

– Microsoft Windows 2003 Server

– Standard, Enterprise or Datacenter Edition

– 64bit Microsoft Windows 2003 Server Service Pack (SP)1

– Enterprise or Datacenter Server

– Microsoft Windows 2003 Server x64

– Standard, Enterprise or Datacenter Edition

► Client software requirements:

– Microsoft Windows 2003 Server

– Web server, Enterprise, Datacenter Edition and Storage Server

- 32bit or 64 bit as appropriate, all SPs supported
- Windows XP Professional (32 bit and 64 bit) SP2 and up
- Windows 2003 Server R2, Enterprise Server R2, Datacenter R2, and Storage Server R2 (x32 or x64 as appropriate, toleration only)
- Citrix Presentation Server for Windows 2003 (32bit)
- Windows Preinstallation Environments (Windows PE) (32bit only) for recovery scenarios
- Windows Vista, toleration only

► Web browser requirements

You require a Web browser to log on and use the console. The browser can be installed on the same system or a separate system. The browsers that are supported are:

- Microsoft Internet Explorer 6.x (Windows systems only)
- Mozilla 1.4 or later
- Netscape 6.2 and Netscape 7

Your browser must support the server codepage. If not, the panels may be unreadable. If your browser meets these requirements but does not correctly display an IBM Tivoli Storage Manager Webbased interface, consider trying a different browser.

► Server hardware requirements

- Intel Pentium® compatible processor or multiprocessor based computer
- Intel Itanium® processorbased computer
- An AMD64 or Intel EM64T or compatible processorbased computer
- At least 110 MB of free disk storage (for a typical installation)
- At least 256 MB of memory

► Client hardware requirements

- Intel Pentium compatible processor or multiprocessor based computer
- Intel Itanium supporting ACPI 2.0 64bit fixed tables
- An AMD64 or Intel EM64T or compatible processorbased computer
- At least 100 MB of free disk storage (for a typical installation)
- At least 128 MB of memory

► One of the following communication protocols is required:

- TCP/IP

–   Named Pipes

–   Shared Memory Protocol

3.2.2 UNIX environment requirements

On UNIX, the Tivoli Storage Manager server requires the following hardware and software.

AIX environment

►   Server software requirements

–   AIX 5.2 or AIX 5.3 (64bit only)

►   Client software requirements

–   AIX 5.2 or AIX 5.3 (32 or 64bit)

►   Server hardware requirements:

–   RISC System p™ or i 64bit

–   At least 200 MB of free disk storage (for a typical installation)

–   1 GB of memory

►   Communication protocol:

–   TCP/IP connection

–   Shared Memory

►   Client hardware requirements:

–   IBM RS/6000® or RISC System p or i

–   At least 50 MB of free disk storage (for a typical installation)

–   128 MB of memory

–   Communication protocol:

►   TCP/IP connection

►   Shared Memory

Hewlett-Packard UX environment

►   Server software requirements:

–   HewlettPackard (HP) PARISC: 11iv2 (64bit only) with the most current maintenance levels installed.

–   HP Itanium: 11iv2 (64bit only) with the most current maintenance levels installed.

- ► Client software requirements:
  - – HP PARISC: machine must have operating system (OS) 11iv2 (32 or 64bit)
  - – The HP Itanium 2 machine must have OS 11iv2 (64bit only)
- ► Server hardware requirements:
  - – HP PARISC 64bit 9000 Series or later machine
  - – HP Integrity Server (Itanium)
  - – 100 MB of free disk storage for install, 60 MB for temporary space
  - – 256 MB of memory
- ► Communication protocol:
  - – TCP/IP
  - – Shared Memory (default)
- ► Client hardware requirements:
  - – HP PARISC 9000 Series 700 or 800 or later machine
  - – HP Integrity Server (Itanium)
  - – 100 MB of free disk storage
  - – 256 MB of memory
  - – Communication protocol: TCP/IP

Sun Solaris SPARC environment

- ► Server software requirements:
  - – Sun Solaris 9 (64 bit, only on SPARC servers)
  - – Sun Solaris 10 (64 bit, SPARC or x86 servers)
- ► Client software requirements:
  - – Sun Solaris 9 (32 or 64 bit)
  - – Sun Solaris 10 (32 or 64 bit)
- ► Server hardware requirements:
  - – Any Sun system which supports the Sun4u or Sun4v architecture
  - – 128 MB of free disk storage in the install directory
  - – 512 MB of main memory
  - – Additional temporary disk space is required during the transfer
- ► Communication protocol:
  - – TCP/IP

- – Shared Memory

► Client hardware requirements:

- – A SPARCstation or compatible workstation

- – 100 MB of free disk storage

- – 128 MB of main memory

► Communication protocol:

- – TCP/IP

- – Shared Memory communication with Solaris server only

Sun Solaris x86_64 environment

► Server software requirements:

- – Sun Solaris 9

- – Sun Solaris 10 (64 bit, SPARC or x86 servers)

- – Client software requirements:

- – Solaris 10 (32bit or 64bit)

- – For Java™ GUI Java JRE™ 1.4.x, where x=1 or higher

- – For Web GUI Java JRE 1.4.1

- – Mozilla 1.4

► Server hardware requirements:

- – An AMD64 or Intel EM64T or compatible processor

- – 128 MB of free disk storage in the install directory

- – 512 MB of main memory

- – Additional temporary disk space is required during the transfer

► Communication protocol:

- – TCP/IP

- – Shared Memory

► Client hardware requirements:

- – An X86 based PC architecture (such as Pentium) or later, or x8664

- – AMD64/EM64T processor

- – 100 MB of free disk storage

- – 128 MB of main memory

- – Communication protocol: TCP/IP

3.2.3 Linux environment requirements

On Linux, the Tivoli Storage Manager server requires the following hardware and software.

Linux X86

► Server software requirements:
    – Red Hat Enterprise Linux 4 (AS, WS, ES)
    – SUSE Linux Enterprise Server (SLES) 9 and 10
    – Asianux 2.0
    – Red Flag DC 5.0
    – Miracle Linux 4.0
    – Haansoft Linux 2006
► Additional software
    – V2.3.3, or later of the GNU C libraries installed on the target machine
► Client software requirements:
    – Red Hat Enterprise Linux 4 (AS, WS, ES)
    – SUSE Linux Enterprise Server (SLES) 9 and 10
    – Asianux 2.0
    – Novell OES
    – For Java GUI: Java JRE 1.4.x, where x=1 or higher
    – For Web GUI: Java JRE 1.4.1
    – Mozilla 1.4
► Server hardware requirements:
    – An i686 compatible symmetric multiprocessor (SMP) or uniprocessor (UP)
    – 128 MB of RAM
    – 256 MB hard disk space (includes full install, 9 MB recovery log, 17 MB database); add 6 MB for each additional language support package
► Client hardware requirements:
    – An X86 based PC architecture (such as Pentium) or later, or
    – AMD64/EM64T processor
    – 128 MB of RAM
    – 40 MB of hard disk space

► Communication protocol:

  – TCP/IP, standard with Linux

  – Shared Memory Protocol (with Tivoli Storage Manager V5.4 Linux x86 Client)

Linux X86_64 (AMD Opteron architecture)

► Server software requirements:

  – Red Hat Enterprise Linux 3 or higher

  – Red Hat Enterprise Server 4 or higher

  – SUSE LINUX Enterprise Server 9

  – SUSE LINUX Enterprise Server 10

  – Asianux 2.0

► Client software requirements:

  – Red Hat Enterprise Linux 4 (AS, WS, ES)

  – SUSE Linux Enterprise Server (SLES) 9 and 10

  – Asianux 2.0

  – Novell OES

  – For Java GUI: Java JRE 1.4.x, where x=1 or higher

  – For Web GUI: Java JRE 1.4.1

  – Mozilla 1.4

► Server hardware requirements:

  – An x8664 SMP, or UP (AMD Opteron or Intel Xeon® processor)

  – 256 MB of RAM

  – 256 MB hard disk space (includes full install, 9 MB recovery log, 17 MB database); add 6 MB for each additional language support package

► Client hardware requirements:

  – An X86 based PC architecture (such as Pentium) or later, or

  – AMD64/EM64T processor

  – 128 MB of RAM

  – 40 MB of hard disk space

► Communication protocol:

  – TCP/IP, standard with Linux

  – Shared Memory Protocol (with Tivoli Storage Manager V5.4 Linux x86

– Client)

Linux System z

For Tivoli Storage Manager Server on IBM Linux System z™:

► Server software requirements:
  – Red Hat Enterprise Linux 4
  – SUSE Linux Enterprise Server 9 and 10
  – Additional software
  – V2.3.3, or later of the GNU C libraries
► Client software requirements:
  – Red Hat Enterprise Linux 4
  – SUSE Linux Enterprise Server 9 and 10
► Server hardware requirements:
  – A System z server with either native 64bit LPARS or 64bit VM guests
  – 128 MB of RAM
  – 72 MB hard disk space (includes full install, 9 MB recovery log, 17 MB database). Add 6 MB for each additional language support package
► Client hardware requirements:
  – An IBM zSeries® 800 or 900, in 32 bit and 64 bit mode
  – Communication Protocol:
  – TCP/IP, standard with Linux
  – Shared Memory Protocol

Linux System p and i

For Tivoli Storage Manager Server on Linux System p (ppc64 architecture):

► Server software requirements:
  – SUSE Linux Enterprise Server 9 and 10 (supported on Power5 processors only)
  – Red Hat Enterprise Linux 4 (supported on Power5 processors only)
  – Asianux 2.0 Red Flag DC 5.0 and Haansoft Linux 2006

Up-to-date information about hardware requirements can be found at:

http://www.ibm.com/systems/p/linux/

  – Additional software:

- V2.3.3, or later of the GNU C libraries installed on the target machine

► Client software requirements:

– SUSE Linux Enterprise Server 9 and 10

– Red Hat Enterprise Linux 4

– Additional software:

- Java JRE 1.4.1+ and 1.5

► Server hardware requirements:

– One of the following IBM systems is required:

– IBM System p or IBM System i™, with IBM POWER™ processors, as supported by Linux distribution

– 128 MB of RAM

– 128 MB hard disk space (includes full install, 9 MB recovery log, 17 MB database). Add 6 MB for each additional language support package

► Client hardware requirements:

– Linux on Power

– IBM OpenPower™

– JS20 Blades

– System p partition

– System i partition

– 128 MB of RAM

– 80 MB hard disk space

► Communication protocol:

– TCP/IP, which comes standard with Linux

– Shared Memory Protocol (with Linux System p or System i Tivoli Storage Manager Server only)

3.2.4 Supported devices

This section lists the devices supported by IBM Tivoli Storage Manager. When you want to add a new attachment to the IBM Tivoli Storage Manager server, you should visit the IBM Tivoli Storage Manager Web sites for the most recent updates.

The devices and links are:

► IBM AIX, HP, SUN, and Windows

http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devi
ces_for_AIXHPSUNWIN.html

► iSeries

http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devi
ces_for_iSeries.html

► Linux

http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devi
ces_for_Linux.html

IBM Tivoli Storage Manager, complementary products, and add-on products
continuously evolve to better support your ever-changing environment. For more
information and currently supported platforms and operating systems, visit:

http://www.ibm.com/software/tivoli/products/storage-mgr/platforms.html

3.3 IBM Integrated Solutions Console and Administration Center requirements

The Administration Center is installed as an IBM Integrated Solutions Console
(ISC) component. The ISC enables you to install components provided by
multiple IBM applications and access them from a single interface.

The Tivoli Storage Manager server can require a large amount of memory,
network bandwidth, and processor resources. In most cases, the server performs
best when other applications are not installed on the same system. If the system
meets the combined requirements for the server and the Administration Center
(for example, it has at least two GB of physical memory), it can support both
applications.

You might plan to use the Administration Center to manage an environment with
a large number of servers or administrators; if so, then you must consider
installing the Administration Center on a separate system.

3.3.1 Software and system requirements

These are the software requirements for various platforms:

Windows

► Windows 2000 Professional, Server with SP4, Windows 2000 Advanced
► Server with SP4.

► Windows 2003 Standard, Windows 2003 Enterprise, Windows 2003

► Datacenter.

► CPU speeds: Speeds of late model, midrange to highend servers are recommended. Pentium 800 MHz or equivalent at a minimum. Production environments must consider the Pentium 4 processor at 1.4GHz or higher.

► Physical memory: 512 MB or more for runtime installations.

► File system: NT file system (NTFS) is recommended.

AIX:

► AIX 5.2 with ML1 + APAR IY44183

► AIX 5.2 with ML5 on power 32bit and 64bit

► AIX 5.3 with ML1 on power 32bit and 64bit

► RS/6000 at 450MHz at a minimum; production environments must consider higher speeds

► Physical memory: 512 MB or more

Linux on x86 32bit:

► CPU speeds of late model, midrange to highend servers are recommended

► Pentium 800MHz or equivalent at a minimum. Production environments should consider the Pentium 4 processor at 1.4GHz or higher

► Physical memory: 512 MB or more for runtime installations

► SUSE LINUX Enterprise Server (SLES) 8 Powered by/ United Linux 1.0 on Intel

► SUSE LINUX Enterprise Server 9

► Red Hat Enterprise Linux 3 update 3

► Red Hat Enterprise Linux 4

Linux on IBM System p 64bit

► CPU speeds: pSeries® models that support Linux (64bit support only).

► Power4 minimum, Power5 recommended

► Physical memory: 512 MB minimum, 1 GB recommended

► SUSE LINUX Enterprise Server (SLES) 8

► SUSE LINUX Enterprise Server 9

► Red Hat Enterprise Linux AS 3 update 3

► Red Hat Enterprise Linux AS 4

Linux on IBM System z
- ► CPU speeds: G5 minimum, G6 recommended
- ► Physical memory: 512 MB minimum, one GB recommended
- ► SUSE LINUX Enterprise Server (SLES) 8
- ► SUSE LINUX Enterprise Server 9
- ► Red Hat Enterprise Linux AS 3 update 3
- ► Red Hat Enterprise Linux AS 4

Solaris:
- ► Solaris 8 and Solaris 9 SPARC on SPARC 32bit and 64bit
- ► Ultra™ 60 at 450MHz at a minimum; Sun Blade 2000 workstation at 1 GHz or higher is recommended
- ► Physical memory: 512 MB or more

3.3.2 Hardware requirements

The machine hosting the Administration Center and ISC requires the following capabilities:

- ► To install the console on a system for the first time, the user requires 982 MB to satisfy the installation program disk space check. An additional 679 MB is also required in the /tmp directory (which is checked during the installation), and 500 MB for the completed installation.

- ► Double the user's physical memory is required. At a minimum, this must be at least equal to the physical memory.

- ► Network adapter and connection to a physical network that can carry IP packets; for example, Ethernet, tokenring, and ATM Static IP address.

- ► Configured fully qualified host name. ISC must be able to resolve an IP address from its fully qualified host name. To verify that the host name is configured correctly, the user can issue the ping command from a command line.

- ► The IBM Tivoli Storage Manager Administration Center Web interface for the server and a Web client interface for client machines require a Java Swingcapable (at JRE 1.4.1) Web browser:
    - – MS Internet Explorer 6.0, or later with Java Plugin 1.4.1
    - – Mozilla 1.6, or later

3.4 Planning the implementation

Depending on the given customer environment, planning should be for a clean implementation.

The checklist in this section can help ensure that the customer's functional requirements match the IBM Tivoli Storage Manager capabilities for the correct implementation.

To determine these requirements, take the following steps:

Step 1: Obtain the following data retention parameters from the customer to understand their backup needs:

► Number of backup versions

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

Use this number as a basis to group your datastorage requirements into management classes.

► Backup file retention period

Determine the number of days you want to keep a backup version of a file (other than the current version). There are two options:

– Keep the backup version for a specified number of days.

– Specify NOLIMIT, which implies that you want IBM Tivoli Storage Manager to retain all backup versions (other than the most recent version), indefinitely.

> **Note:** The current version is kept indefinitely by default.

► Number of deleted versions

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

backup versions of the file for deletion and just keep the most current backup version.

► Last deleted file version retention period

Determine the number of days you want to keep the last backup version of a file after it has been deleted from the client. There are two options:

– Keep the last backup version for a set number of days.

– Specify NOLIMIT, which implies that you want to keep the backup version indefinitely.

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

► Archive retention period

Determine how long you want to keep a file that is archived. Many sites set up a limited number of data groups with standard archive retention periods, such as 7 days, 31 days, 180 days, 365 days, or 7 years.

Nonstandard requests for archive retention periods are slotted into the retention period group closest to the request without being smaller. This reduces management complexity at the expense of keeping some data longer than actually required. If every nonstandard request is honored, the number of groups quickly becomes unmanageable. On the other hand, you can use the backup set feature to retain all nonstandard backup requirements, or even just use backup sets instead of archive.

► Off-site copies

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

► On-site collocation

Determine whether you want to use on-site collocation. IBM Tivoli Storage Manager uses collocation to dedicate the minimum number of tapes required to hold all of one client's files. Collocation reduces elapsed time for multiple file restores and full client restores. Collocating by client allows as many clients to be restored simultaneously as you have tape drives. If you have stringent restore requirements and sufficient tape drives, then collocation makes good sense.

► Image backup retention

Determine how long you want to keep an image backup. Consider your restore time frame and balance the criticality of a full file space restore, compared with single or a small number of files. Image backups can be very useful for quick restores of large file systems. However, this process will take longer, depending on how many changes there have been to the file system since the last image backup. We recommend keeping at least one weekly image for small servers and a monthly image for bigger servers (or more frequently if change rates in the file system are high).

► Backup set retention

A backup set execution creates a copy of the client node's previously backed up active files and stores them on sequential media. This has an impact on the number of tapes that you may need, especially if you want to retain those backup sets for long periods of time or even if you want to have one copy on site and one for off-site purposes.

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

Step 2: Analyze and choose the server platform.

Make sure that customer's platforms are fully supported by IBM Tivoli Storage Manager Version 5.4 (see 3.2, "Tivoli Storage Manager System requirements" on page 66) and size the IBM Tivoli Storage Manager server.

There is a very little information to guide you in selecting a IBM Tivoli Storage Manager platform. At the same time, the risk of choosing an inappropriate platform size increases as the size of the IBM Tivoli Storage Manager implementation increases. Small IBM Tivoli Storage Manager implementations are at less risk, and the incremental cost to scale up or down is small. Many sites start small and grow into larger systems.

The IBM Tivoli Storage Manager server is processing intensive. Processing is a function of the number of files to manage and how your platform processes input/output (I/O). A large number of small files uses more processing resources than a small number of large files. As the number of files and the amount of data to be moved increases, each backup, migration, storage pool copy, and expiration process will use more processing resources to maintain the database entries. Therefore, IBM Tivoli Storage Manager takes advantage of multiple processors.

I/O is the major part of IBM Tivoli Storage Manager processing. Memory is used to cache database entries, among other things. As the number of files being

managed increases (and thus the database size increases), the amount of memory that IBM Tivoli Storage Manager requires increases.

Step 3: Determine network load.

Network topologies such as Ethernet, Token Ring, Fiber Distributed Data Interface (FDDI), and asynchronous transfer mode (ATM) all work well with IBM Tivoli Storage Manager. Each has its strengths and weaknesses. In general, choose the fastest network topology you can afford; otherwise, you should consider some performance functions such as compression for better utilization in slow networks. Alternatively, the new SAN technologies can be implemented to take full advantage of tape library resources, freeing part of the ordinary network consumption to business applications. Network is the most important consideration when determining the IBM Tivoli Storage Manager solution. Table 3-1 gives the estimated network throughputs.

To estimate the speed of the network required, use the following calculation. Calculate the total amount of data to be transferred during a backup window. For each client:

1. Multiply the GB changed per backup field by (1 - Data compression field).

2. Sum this number for all clients to obtain the total data to be transferred.

3. Divide the total data transferred by the number of hours in your backup window to obtain the transfer rate in GB/hr.

Table 3-1   Network throughputs

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

For example, the customer has:

– Two DB2 database servers, each with 5 GB needs fully backed up

– Two Windows 2003 File servers each with 60 GB and 5 percent daily change rate

– One hour backup window for all servers

– Limited budget

Using the data in this table, calculate the network speed required as follows:

- (5+5) + (2*60*.05) = 16 GB.
- Given that the backup window is 1 hour, then 16/1 = 16 GB/hr.

For this environment Ethernet (100 MB/sec) will be the proper choice, since the customer has a limited budget.

Step 4: Determine the IBM Tivoli Storage Manager database and recovery log by identifying the sizing of each.

The database and the recovery log are closely related. The IBM Tivoli Storage Manager database contains information about client nodes, users, policies, schedules, and storage pools. The database also contains pointers to the locations of all backup versions and client files in the IBM Tivoli Storage Manager storage pools. The client data itself is stored in storage pools, not in the database.

Using the IBM Tivoli Storage Manager database, customers can define storage management policies around business needs for individual clients or groups of clients. Client data attributes such as storage destination, number of versions, and retention period can be assigned at the individual file level and stored in the database.

The IBM Tivoli Storage Manager database also ensures reliable storage management processes. To maintain data integrity, the database uses a recovery log to roll back any changes and updates made if a storage transaction is interrupted before it completes. This is known as a two-phase commit.

Also, both the IBM Tivoli Storage Manager database and recovery log can be mirrored for availability, providing automatic volume switching after a media failure. In the unlikely event of an IBM Tivoli Storage Manager database recovery, operators can restore the database to the exact point of a failure by rolling the recovery log forward after restoring from the latest database backup.

► Database sizing

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

Backup sizing calculates the size of the IBM Tivoli Storage Manager database holding backup entries. Archive sizing calculates the size of the IBM Tivoli Storage Manager database holding archive entries. Use either or both

depending on the data you will be storing in IBM Tivoli Storage Manager. If both are used, then add the calculated database sizes together to arrive at the total database size. If you are planning to use backup sets, remember that these are not tracked in the IBM Tivoli Storage Manager database, so they do not impact our calculation.

The calculation is based upon two types of numbers:

– Actual number of file versions backed up

– Percentage of the data backed up (when actual number is not known)

Based on the Number of files method

i.   Sum the value in the Number of Files Backed Up field for all clients, leaving out any fields containing estimates in GB.

ii. Multiply this number by the value in the Number of Versions Kept field to obtain the total number of files backed up.

iii. Multiply this number by 600 bytes to obtain the bytes used in the database for all known files backed up

iv. If copy storage pools are used (and we strongly recommend this), multiply the total number of files backed up and calculated in step ii by 200 bytes to obtain the bytes for known copy storage pool files.

v.   Add the total bytes for backed up files (calculated in steps i through iii) to the total bytes for copy storage pool files to obtain the total bytes calculated for the database.

vi. Calculate 135 percent of total bytes calculated for the database to obtain the database size. This is for overhead and for growth.

Example:

Scenario

Total number of servers to be backed up = 3

Number of files to be backed up on each server =50000

Number of versions to keep = 3

Calculating the total size of the database is as follows

i.   First calculate the total number of files to be backed up TSM i.e 50000x3=150000

ii. Multiply this value by the number of versions will give the actual number of files stored in TSM server any given time

150000x3=450000 no. of files

iii. Multiply this value by 600 bytes (standard occupancy size for a primary storage files) will give number of bytes occupied in the database

450000x600= 270000000 bytes

iv. Assume copy storage pools are also used then multiply 450000 (the total number of files including total number of versions) by 200 bytes

45000x200= 90000000 bytes

v. Sum the above 2 result will give the total space occupancy in the database.

270000000 + 90000000 = 360000000 bytes of space required for backing up 3 clients

vi. 135 percent of the above data is 486000000 bytes (for capacity planning) = 463.49MB of Databas size

Note: Divide number of bytes by 1024 will give you the size in Kilobytes and further dividing this gives the result in Megabytes.

Example where number of files not given:

Scenario

Backup of 40 file servers with 45GB of data each

Retention backup policy of 10 versions

One copy pool for off-site media

100 kbytes average file size for all backup clients

Overhead that requires up to 20% in additional space for TSM database

Calculate the estimated TSM database size as follows

i. Calculate the amount of data backedup

40 x 45GB = 1800GB

ii. Convert this to Megabytes

1800GB x 1024 = 1843200MB

iii. It was mentioned that the average file size is 100kbytes, convert this to megabytes as well

100 / 1024 = 0.09765625MB

iv. Calculate the number of files

1843200MB / 0.09765625MB = 18874368 files

v. Multiply this value by 600 bytes (standard occupancy size for a primary storage files) will give number of bytes occupied in the database. Multiply this number by the number of versions for primary pool files.

18874368 files x 600 bytes x 10 = 113246208000 bytes

vi. Seeing we require copy storage pools, multiply the number of files by 200 bytes (standard occupancy size for copy storage files) and multiply this by the number of versions for copy pool files.

18874368 x 200 bytes x 10 = 37748736000 bytes

vii. Add the two totals together

113246208000 bytes + 113246208000 bytes = 150994944000 bytes

viii.Convert this to Gigabytes

150994944000 / 1024 /1024 /1024 = 140.625GB

ix. Add 20% overhead

140.625GB x 1.20 = 168.75GB

The TSM database size is now 168.75GB

► Recovery log sizing

The size of the recovery log depends on the amount of data changed between IBM Tivoli Storage Manager database backups. The greater the amount of data, the larger the recovery log needs to be. Either a full or an incremental IBM Tivoli Storage Manager database backup (in roll-forward mode) resets the recovery log back to empty. If the recovery log fills up completely, IBM Tivoli Storage Manager stops and you have to manually increase the size of the recovery log by using the `dsmfmt` command to create additional log space.

To estimate the size of the recovery log, multiply the database size by the percentage of data that changes each backup cycle. Double this number to allow for two backup cycles to occur without a database backup. This is a starting point for the recovery log. Also, the number of concurrent client sessions and using the roll-forward mode should be considered calculations of the size of recovery log.

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

As with the database, we recommend using the IBM Tivoli Storage Manager mirroring function for the recovery log instead of a hardware mirror or operating system mirror. You should use the `dsmfmt` command to format volume mirrors.

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

Step 5: Planning the storage pool.

Using IBM Tivoli Storage Manager, you can configure storage pools to provide the best combination of performance throughput and data permanence. The import and export feature helps you to move storage pool data among the different platforms.

IBM Tivoli Storage Manager has two types of storage pools:

► Primary storage pools

When a client node backs up, archives, or migrates data, the data is stored in a primary storage pool. When a user tries to restore, retrieve, or export file data, the requested file is obtained from a primary storage pool if possible. Primary storage pool volumes are always located on-site.

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

Active-data storage pools are primary storage pools that can use either squential disk or tape drives to store the data.

► Copy storage pools

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

A copy storage pool can use sequential access storage (for example, tape, optical, or FILE device classes). Copy storage pools can also be created remotely on another IBM Tivoli Storage Manager server, thus providing electronic vaulting.

Copy pools are not part of the storage migration hierarchy. Files are not migrated to or from copy storage pools. There are two ways to store files in a copy storage pool:

– Copy the primary storage pool to a copy storage pool using the BACKUP STGPOOL command.

– Simultaneously write to copy storage pools during client data transfer activity.

IBM Tivoli Storage Manager can simultaneously store a client's files to each copy storage pool specified for the primary storage pool where the client's files are written. This is illustrated in Figure 3-1. The simultaneous write to the copy pools only takes place during a backup or archive from the client (in other words, when the data enters the storage pool hierarchy). It does not take place during data migration from a Tivoli Space

Manager client or on a LAN-free backup from a storage agent. As many as 10 copy storage pools can be specified for each primary storage pool.
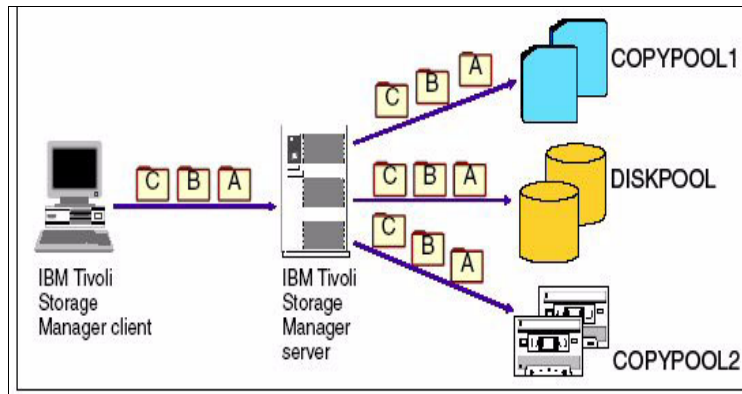
Figure 3-1   Simultaneous write

This function reduces the overall copy storage pool window and places copies of data off site faster.

Step 6: Estimate the size of the total IBM Tivoli Storage Manager disk required.

IBM Tivoli Storage Manager requires disk to operate. The disk holds the database, logs, and usually the primary storage pools.

You should size the disk subsystem for growth as follows because IBM Tivoli Storage Manager implementations grow substantially. Choose a disk model that meets your present estimated needs and has room for expansion. Multiple I/O paths and hot-swappable components should also be considered, both for performance and availability.

► Backup disk sizing

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

   a. Multiply the GB changed per backup by (1 - the Data compression rate) to obtain the total bytes transferred for each client.

   b. Sum the total bytes transferred for all clients to obtain the total bytes transferred per backup cycle.

   c. Add 15 percent to the total bytes transferred per backup cycle to obtain the storage pool size. This allows for variations in the size and number of files per backup.

▶ Archive disk sizing

In most environments, archive disk sizing is less critical than for backup. Archives may run less frequently than backups, and on weekends when the overall workload is lighter.

To increase the storage pool to hold archive data as well, follow these steps:

a. Group all machines that require simultaneous archive operations during one common time frame (for example, every month). Select the biggest group to obtain the peak archive size.

b. Take 10 percent of the peak archive size to obtain the archive storage size.

▶ Image disk sizing

If you are planning to use image backups, you would consider sizing the disk storage pool to hold the file spaces you want to back up. This is because image backups are single objects and, therefore, the server will require that size for storing data.

▶ Device configuration table and volume history file sizing

The device configuration table and the volume history table also require disk space, but typically this is a small amount. The device configuration table has entries for defined device classes and definitions for drives and libraries. All volumes used by IBM Tivoli Storage Manager are tracked in the volume history database, including the volume identifier for the database backups. The volume history is periodically copied to a volume history file that you specify with the VOLUMEHISTORY option in the dsmserv.opt file.

▶ Total disk

The considerations you should include when you are calculating total disk space are:

– IBM Tivoli Storage Manager code (depended on platform)
– IBM Tivoli Storage Manager database
– IBM Tivoli Storage Manager recovery log
– Primary storage pools
– Device configuration table and volume history table
– Other (RAID, Operating system)

Step 7: Determine tape considerations and sizing.

Most IBM Tivoli Storage Manager systems use tape as the ultimate destination for data storage. There are a variety of tape subsystems available for most IBM Tivoli Storage Manager server platforms.

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

Optical disks are an alternative to tape, but we recommend tape because it is generally faster and more convenient to work with.

► Tape devices

Tape drives come in all sizes, including, but not limited to, digital linear tape (DLT), super DLT (SDLT), linear tape open (LTO), IBM TotalStorage® 3590 Tape Drive, and other device types. Each type of drive has different data capacity, performance, cost, and reliability characteristics. Although data capacity and cost per megabyte stored are important, reliability is much more important.

In general, tape drives where the tape touches the read/write heads, such as 4 mm and 8 mm, tend to be less reliable (and slower) than tape drives where the tape does not touch the read/write heads, such as IBM TotalStorage 3580 model L33 Tape Drive and 3590. We recommend an automated tape library with a bar code reader for all but the smallest of IBM Tivoli Storage Manager implementations.

► Backup tapes

To determine the number of on-site tape volumes required for backup operations, perform the following calculations.

– If this is a sequential storage pool (tape device), multiply the primary storage pool size by the number of backup versions to obtain the versions pool size.

– Add the sum of all Total storage used fields for each client to the versions pool size to obtain the tape pool size.

– Divide the tape pool size by the device capacity to obtain the number of tape cartridges required.

– Add 50 percent to cater to tapes that are in filling status to obtain the total cartridges required for on-site tapes.

– If using collocation, normally there should be at least as many tape cartridges as there are clients. Consider tape native capacity as a rounding factor.

Step 8: Identify the IBM Tivoli Storage Manager administrator ID

An IBM Tivoli Storage Manager administrator controls IBM Tivoli Storage Manager resources. There can be numerous administrators with varying levels of authority. It is possible to use the Web backup-archive client to perform backup, restore, archive, and retrieve operations on the behalf of other users using a Web browser. Help desk personnel can use the IBM Tivoli Storage Manager

Web Client to perform these client tasks for their end users without having to log on to the client machine.

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

4

Solution architecture

This chapter provides the required information about developing a Tivoli Storage Manager solution based on the data gathered from the customer's needs and the environment that customer has.

In this chapter, we discuss the following topics:

► What to know about the first section of the test objectives.
► Design Server and client configuration.
► Determine storage pool hierarchy.
► Evaluate the network infrastructure.

4.1 What to know about the first section

In this chapter, we continue to discuss the first section of the Tivoli Storage Manager 5.4 Implementation Certification Test objectives, which is planning. This includes customer storage management and hardware requirements, network capabilities, and designing a solution that meets the customer's needs. At the end of this section, you will able to accomplish these tasks:

► Document assumptions inherent in the customer requirements and environment.

► Determine the platform, location, and number of servers.

► Determine the size and mirroring of the Storage Manager server database and logs.

► Determine the storage pool hierarchy and size of the storage pools and distribution of the storage pool levels in the enterprise.

► Determine the disaster backup and recovery plan, including the use of electronic vaulting.

► Determine the protocols to be used for each environment.

► Determine the service level agreement (SLA) the solution will be required to meet.

► Determine the appropriate number of servers given the environment.

► Determine the hardware requirements to meet the solution design and evaluate whether existing hardware will meet the requirements of the solution.

► Determine which Storage Manager features and licenses are required.

► Determine data types (for example, databases, network attached storage, and mail servers).

► Evaluate the customer's network infrastructure (for example, LAN, WAN, SAN, mobile).

4.2 Tivoli Storage Manager architecture considerations

Now that you have gathered information on the total client environment, you can make decisions about the architecture of the Tivoli Storage Manager server environment.

4.2.1 Platform and server considerations

In this section, we examine the platform and server considerations. A Tivoli Storage Manager server runs on several platforms. How do you choose one

platform over another? With only minor differences, a Tivoli Storage Manager server provides the same functionality on every platform. The differences relate to capacity, cost, installation, operation, supported devices, and installed user base.

Installed user base

The number of Tivoli Storage Manager servers installed for a particular platform is a consideration. IBM ships new functions for the most popular Tivoli Storage Manager server platforms (AIX, MVS, and Windows 2003) first. The more popular a platform, the more customers use the software, creating a lower probability that you will find a unique problem.

Cost

Cost is further divided into platform costs and Tivoli Storage Manager software license costs. Platform costs include the cost of acquiring the hardware and software to run the platform exclusive of the Tivoli Storage Manager software license. It ranges from very low for Windows 2003 to extremely high for IBM Multiple Virtual Storage (MVS).

Tivoli Storage Manager license costs vary considerably, with Windows 2003 costs being the lowest, followed by IBM AIX, HP-UX, Sun Solaris, and iSeries costs. These are one-time product purchases. MVS licenses are available for a one-time charge or as a monthly license fee. You may want to calculate the break-even point.

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

Capacity

A Tivoli Storage Manager server can manage numerous clients and an unlimited amount of data. However, the platform on which the Tivoli Storage Manager server software runs can limit what Tivoli Storage Manager can administer. Various platforms have different capacities in regard to the processing power they can deliver to Tivoli Storage Manager, the number of devices it can attach, and the throughput it can deliver.

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

Platform installation

The Tivoli Storage Manager server code installation varies by platform in the specifics, but generally follows a similar procedure. Installation on Windows 2003 can be easier due to the Windows 2003 wizards that have been provided. Installation of the Tivoli Storage Manager server on other platforms is not difficult for an administrator familiar with the platform.

Operation

Operation of a platform varies from almost completely automatic on Windows 2003, to very complex on MVS, with the UNIX and iSeries platforms somewhere in the middle.

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

Supported devices

There are a wide variety of supported devices on the Windows 2003 and UNIX platforms, including disk drives, tape drives, optical drives, automated and virtual tape libraries. MVS and iSeries are limited to their standard choice of devices, but these devices generally have tremendous capacity.

A concern with the smaller platforms is the ability to attach the required amount of devices as the environment grows. On larger platforms this concern is usually not as significant.

4.2.2 Server recommendations

There are a number of considerations to be taken when selecting a server for the Tivoli Storage Manager architecture.

Choice of platform

If you already have a z/OS system or an iSeries system with capacity available, these platforms are a good choice. Otherwise, you should probably eliminate them from consideration due to their complexity, cost, or specialized nature.

If you have experience with Windows 2003, and the Tivoli Storage Manager implementation is small and will remain small, choose Windows 2003.

Server system size

We have already explained the detailed server system sizing parameters in Chapter 3, "Requirements and IBM Tivoli Storage Manager functionality" on page 65.

In addition to processing, I/O, and memory, multiple servers should also be a consideration depending on how large the environment is. Multiple Tivoli Storage Manager servers can be configured to provide some redundancy and disaster recoverability in the event of a Tivoli Storage Manager server outage.

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

For a large, enterprise-wide business intelligence complex, a dedicated Tivoli Storage Manager server (either on the same system or a different one) might be the best solution. In installations where network connectivity is slow or expensive, placing a Tivoli Storage Manager server close to the clients usually makes sense.

Multiple servers increase costs. Two small servers may be more expensive than one larger server of the same power. Where one automated tape library may be enough, multiple servers may require multiple automated libraries. Every Tivoli Storage Manager server requires a Tivoli Storage Manager server license. For example, a customer with a large central site and many remote locations with one small network and a slow link to the central site must consider one storage server for the central site and all remote sites.

Communication protocol

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

Client-server communications options

By default, the server uses these methods of communication:

► TCP/IP options

    TCP/IP is the most common communication method and possibly the easiest to set up. Certain functions, such as server-to-server communications,

server-prompted mode, and the Web clients, require TCP/IP. A sample setting
is:

```
commmethod tcpip
tcpport 1500
tcpwindowsize 8
tcpnodelay no
```

▶ Named Pipes Options

The Named Pipes communication method is ideal when running the server
and client on the same Windows machine because Named Pipes support is
internal to the Windows base system. Named Pipes require no special
configuration. This is an example of a Named Pipes setting:

```
commmethod namedpipe
namedpipename \.\pipe\adsmpipe
```

▶ HTTP Options

Use the HTTP communication method for the client Web interface. HTTP and
communications require a TCP/IP connection. This is a sample HTTP setting:

```
commmethod http
httpport 1581
```
The Integrated Solutions Console and the TSM Administration Center
software needs to be installed to use the server administrative Web interface.

Shared Memory option

Tivoli Storage Manager supports shared memory communication between the
AIX Storage Manager server and client are on the same machine. To use shared
memory, TCP/IP must be installed on the machine. Here is a sample shared
memory setting:

```
commmethod sharedmem
shmport 1510
```

4.2.3 Connecting with Tivoli Storage Manager across a firewall

The Tivoli Storage Manager server and clients can work across a firewall.
However, the server can also securely manage client backup and restore
operations and administrative functions across a firewall.

To enable clients to work across a firewall, configure the firewall to open the
ports that the server and clients need. Ports need to be opened only if functions
beyond those of server-prompted scheduling are required.

Server-prompted scheduling can be handled with the SESSIONINIT=
SERVERONLY and HLADDRESS and LLADDRESS options of the REGISTER
NODE or UPDATE NODE commands.

4.2.4 Client architecture

Tivoli Storage Manager is a client-server program. The client product, which must be installed on the machine you want to back up, is responsible for sending and receiving data to and from the Tivoli Storage Manager server. The backup-archive client has two distinct features: Backup and archiving.

All backup-archive clients are implemented as multi-session clients, which means that they are able to exploit the multithreading capabilities of modern operating systems. This enables the running of backup and archive operations in parallel to maximize the throughput to the server system.

Each of these features has a complementary function as well: Restore and retrieve. Restore allows users to recover any data that has been backed up previously. Regarding the individual implementation of Tivoli Storage Manager, the user is able to recover from different versions of the lost data. Retrieve enables users to request formerly archived data so that this data is accessible again.

Depending on the client platform, the backup-archive client has a graphical, command line, or Web user interface. Many platforms provide all three interfaces.

The command line interface is useful for experienced users, and it allows generation of backup or restore scripts for scheduled execution. The graphical interface is designed for end-user ease of use for ad hoc backups and restores and does not require any Storage Manager skills.

The Web client is especially useful for those clients, such as NetWare, where no native GUI is available, or for performing remote backup/restore operations, such as in a help desk environment.

Configuration and options files

Configuration files and options files are used to specify one or more servers and communication options for backup and restore services. The file can include authorization options, backup and archive processing options, scheduling options, and, where applicable, Tivoli Storage Manager for Space Management options.

On a UNIX platform, the Tivoli Storage Manager options reside in two options files: The client system options file (dsm.sys) and the client options file (dsm.opt). On other platforms, the options file (dsm.opt) contains all options. The end user sets up these files when the Tivoli Storage Manager backup-archive client is first installed on the user's workstation. For the Windows platforms you should define client options using the DEFINE CLIENTOPT command.

The minimum configuration parameters for successful communication are:

► Communication protocol

  The client and server use the same type.

► Tivoli Storage Manager server address

  This address identifies the correct Tivoli Storage Manager server to use.

► Nodename

  This is the name by which the Tivoli Storage Manager server knows the client. This information is required to allow access for this client. The node name and password are set up on the server, and if different from the machine name, they also must be added in the client options file.

4.3 Storage hierarchy

A data storage policy can be very simple or very complex. The simplest policy has one set of rules that apply to all of the data in an organization, while the most complex policy has a unique set of rules for each individual file. Most installations fall somewhere between these two extremes.
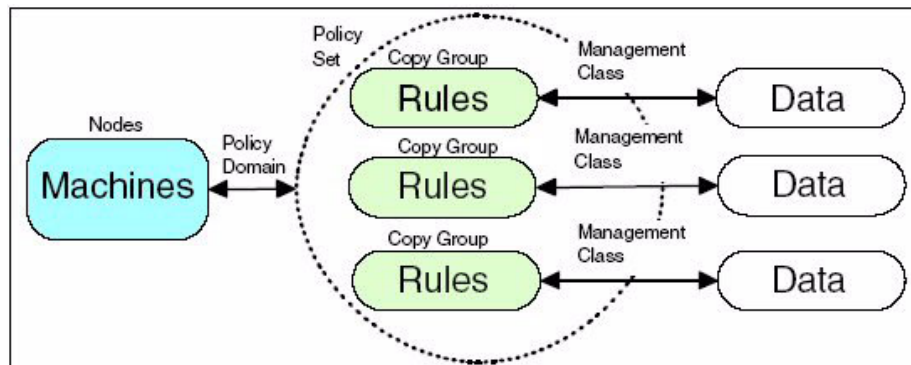
Figure 4-1   Data storage policy relationships and resources

Tivoli Storage Manager has entities that group and organize the resources and define relationships between them. A machine, or node in Tivoli Storage Manager terminology, is grouped together with other nodes into a policy domain. The domain links the nodes to a policy set that consists of management classes. A management class contains rules called copy groups that it links to the data. When the data is linked to particular rules, it is said to be bound to the management class that contains the rules.

4.3.1 Policy components

The hierarchical structure of the Tivoli Storage Manager policy components is shown in Figure 4-2. The policy domain must exist before the policy set, and so on. Examining the diagram leads to the realization that most of the data storage policy components exist solely to provide flexibility in our configuration or to serve as containers for rules.
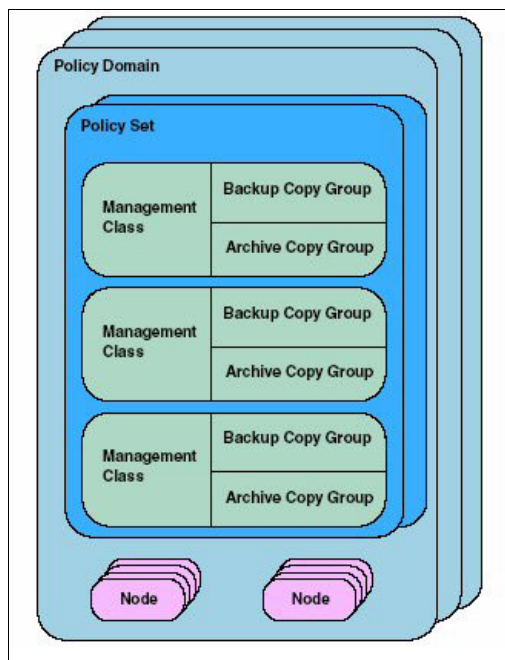
Figure 4-2   Data storage policy components

4.3.2 Copy groups

The copy groups consist of rules used to govern the retention of data. There are two types of copy groups: A backup copy group that holds the rules for backup data, and an archive copy group that holds the rules for archive data.

While these two copy groups serve different purposes, they also share common ground. They both specify where to store the data sent to them from backup-archive operations. The copy group destination parameter specifies a valid primary storage pool to hold the backup or archive data. The copy group bridges the gap between data files and storage pools, as illustrated in Figure 4-3. This diagram shows different types of data flowing through the copy groups and into the storage pools. Note that there is not necessarily a one-to-one

relationship between copy groups and storage pools. It is possible to have just one storage pool that is the destination for all of the copy groups.
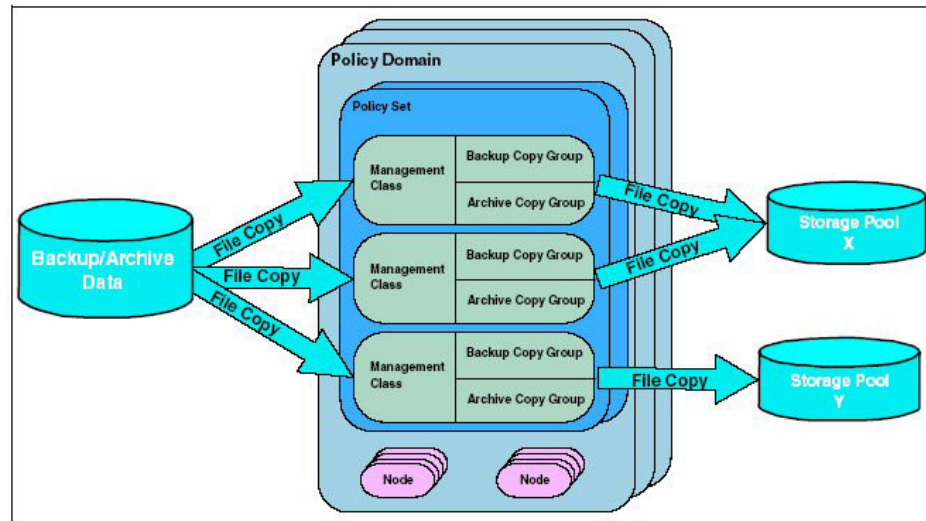
Figure 4-3   Copy group data flow

Both copy groups also need to know what to do with files that are modified during a backup or archive operation.

The copy serialization parameter provides four options:

► The shrstatic setting specifies that a file will not be backed up if it is modified during backup, but multiple attempts will be made to back up the file.

► The static setting specifies that a file will not be backed up if it is modified during backup and no additional attempts will be made.

► The shrdynamic setting specifies that a file will be backed up if it is modified during backup but multiple attempts will be made to back it up without modification first. If that cannot be done, then the file will be backed up anyway.

► The dynamic setting specifies that a file will be backed up even if it is modified during backup. There is no preliminary attempt to back up the file unmodified; it is backed up on the first attempt.

4.3.3 Backup copy group

The backup copy group is concerned with two logical objects: The file and the file copy. A file is the actual data on a client node, while a file copy is a point-in-time

copy of the file stored on the server. In other words, the Tivoli Storage Manager server contains file copies and nodes contain files.

A file can be in one of two possible states: Existing or deleted. When we talk about an existing file on a node, we mean a file that has been previously backed up and still exists on the node. A deleted file is a file that has been previously backed up and subsequently deleted from the node. This simple concept is important when discussing data storage rules.

A file copy can be in one of three states:

► Active: The most current copy of the file
► Inactive: The previous copy or version of the file
► Expired: A copy to be removed from the Tivoli Storage Manager server

A backup file copy is set to the expired state when it no longer conforms to the rules specified in the backup copy group.

The retention periods that we set in our rules specify the length of time to retain inactive file copies. It is important to note that there is no retention period for active file copies; they exist as long as the file exists on the node.

Whether the file exists on the node will affect which rules are used to expire the file copies. If the file exists, the following two backup copy group parameters are in effect:

► verexists

   This parameter specifies the number of file copies, or versions, to keep. This number includes active and inactive file copies.

► retextra

   This parameter specifies how long to keep inactive file copies. When a file changes from active to inactive, it will be kept for *retextra* days and then removed. It is important to note that the retention period starts when the file copy becomes inactive, not from its original backup date.

If the file has been deleted, the active file copy will be made inactive. At this point, there are only inactive file copies for this data in the Tivoli Storage Manager server, and the following parameters apply:

► verdeleted

   This parameter specifies the number of file copies to keep after the file has been deleted.

► retonly

This parameter specifies how long to maintain the last file copy of the data. This is the number of days to keep the last copy only and does not apply to other inactive file copies.

4.3.4 Archive copy group

The archive copy group works with entire archives as single unique entities, so it has fewer rules. There is only ever one copy of a particular archive, so we do not have to worry about rules to manage versioning. We still have to specify the retention period for the archive object and that is done with the retver setting. It specifies the number of days to retain the archive copy from the day of the archive operation.

4.3.5 Management class

The management class serves two purposes: It contains copy groups and associates data to them. A management class must contain at least one copy group. This may be a backup or archive copy group, or it may contain both a backup and an archive copy group. Figure 4-4 shows the basic structure of a management class with both copy groups defined. It also illustrates how a management class links the backup-archive data to the rules defined in a copy group. The link is very granular and can be assigned to a single file or groups of files. When a file is linked to a management class, it is said to be bound to the management class.
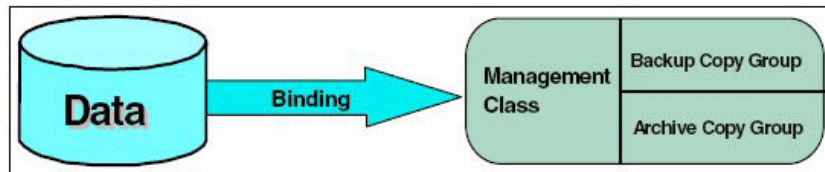
Figure 4-4   Management class

There is a special instance of a management class called the default management class. There is only one for each logical grouping of nodes (policy domain), and it contains the rules that you want used for your data unless you explicitly bind it to another management class. Therefore, there are two ways of binding data to a management class: Default and explicit. Unless an object is explicitly defined, the default management class is used.

Binding your backup data to different management classes enables you to manage different types of files with different sets of rules. Backup data is bound to a management class using the include option of the Tivoli Storage Manager client include-exclude list.

4.3.6 Policy set

A policy set is a group of management classes. There can be multiple policy sets within a policy domain, but only one of them can be active at a time. In other words, the active policy set contains the only group of management classes that can be bound to the data within the domain. Management classes in other policy sets are not available unless you activate the policy set that contains them. Each time a new management class is defined in the policy set, you must validate and activate the current policy set.

A policy set can contain many management classes but only one of them can be the default. The basic structure of a policy set is shown in Figure 4-5 and illustrates that the policy set is used primarily for flexibility. It allows us to group management classes and assign one of them as a default for the policy domain.
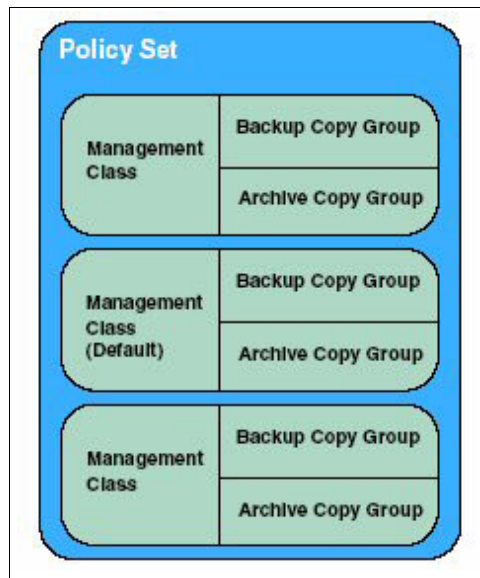
Figure 4-5   Policy set

A policy domain enables you to logically group the machines in your organization according to:

► Default policy

  This policy is the default set of rules to apply to the clients. The rules define the storage management policy including how many copies of data to keep and how long to keep them. The default management class of the active policy set contains the default rules applied to the domain.

► Administrative control

Access™ to the client data and the policy rules can be restricted to certain administrators by allowing or disallowing the administrators access to the policy domain.

4.4 Planning disaster recovery management

One of the most valuable business assets is the critical data that resides on the computer systems throughout the company. The recovery of this data needs to be a primary focus of the disaster recovery plan.

Tivoli Storage Manager, Extended Edition delivers disaster recovery with its Disaster Recovery Manager (DRM) function. This feature can also be licensed to your base product. It offers various options for configuring, controlling, and automatically generating a disaster recovery plan, whenever the storage manager database is updated. This plan contains the information, scripts, and procedures needed to automate restoration and help ensure quick recovery of data after a disaster.

DRM also manages and tracks the media on which data is stored, whether on-site, in transit, or in a vault, so that your data can be easily located if disaster strikes. The scripts can help you document your basic IT recovery strategy, the steps to rebuild your core systems, and the critical machines that you must recover.

Figure 4-6 on page 104 shows a typical scenario for DRM.
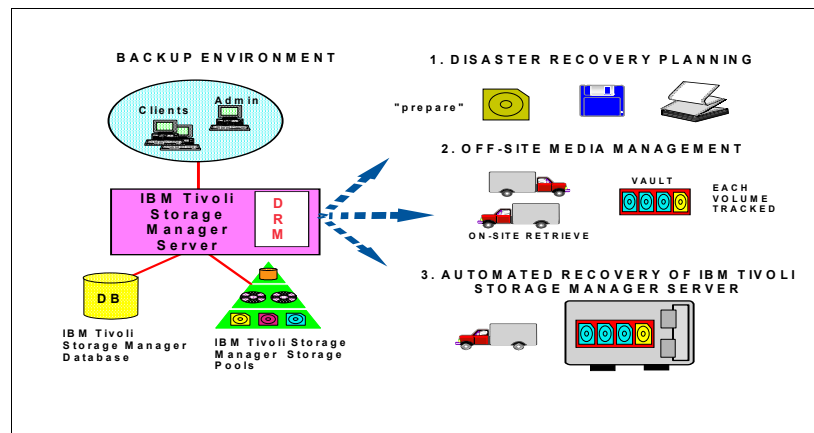
Figure 4-6   Disaster recovery flow

Disaster recovery management is accomplished with Tivoli Storage Manager through the following actions:

► Backing up client data to the Tivoli Storage Manager server

► Backing up the server database to removable media and storing the media off site

► Backing up the primary storage pools and storing the media off site

► Using the disaster recovery plan file to assist with the Tivoli Storage Manager server recovery

► Optionally using LAN-free recovery options, such as backup sets where available and appropriate to improve recovery

► Optionally using virtual volumes to save data, recovery plan files, and database information electronically to an alternate Tivoli Storage Manager server

There are three types of data objects that Tivoli Storage Manager and DRM monitor and track:

► Tivoli Storage Manager server database backups

The heart of Tivoli Storage Manager, database backup, is vital for server recovery.

► Copy storage pool data

As Tivoli Storage Manager backs up the clients, new data is stored in the primary pools. For off-site storage requirements, use copy storage pools. The BACKUP STORAGEPOOL command copies all new primary storage pool files to the copy storage pool. This ensures that the copy storage pool is an up-to-date reflection of your most recent backup. Each time the primary pool is backed up to the copy storage pool (and this is recommended to be performed daily), the newly generated tapes should be sent off site.

► Tivoli Storage Manager Server option file

4.4.1 Server recovery steps

DRM simplifies the disaster recovery planning process for the Tivoli Storage Manager server by generating a recovery plan file that is based on a pre-defined recovery strategy. The recovery plan file contains the information and procedures necessary to help restore the key components of the Tivoli Storage Manager server. The content of the plan file includes:

► Installation-specific server recovery instructions

► A list of Tivoli Storage Manager database backup and copy storage pool volumes required to perform the recovery, including the off-site location where the volumes reside

► Devices required to read the database backup and copy storage pool volumes

► Space requirements for the Tivoli Storage Manager database and recovery log

► A copy of Tivoli Storage Manager configuration files

► Shell scripts and Tivoli Storage Manager macros for performing server database recovery and primary storage pool recovery

4.4.2 Creation of up-to-date disaster recovery plan

To make the creation and maintenance of the server disaster recovery plan easier, the `prepare` command automatically queries the required information from the Tivoli Storage Manager server and creates the recovery plan file. The `prepare` command can be scheduled using the Tivoli Storage Manager central scheduling capabilities.

Off-site recovery media management

Knowing the location of off-site recovery media is critical to the successful implementation of a disaster recovery management plan.

The off-site recovery media management function provides:

► Determination of which database and copy storage pool volumes need to be moved off site and back on site

► Automatic ejection of volumes from an automated library

► Tracking of the media location and state in the Tivoli Storage Manager database

The command `query drmedia` determines which tapes should be taken off site. This function allows database backup volumes and copy storage pool volumes to be treated as logical collections that are selected to move off site for safekeeping and on-site for use. The reclamation of off-site volumes includes the capability to specify the number of days to retain a Tivoli Storage Manager database backup series. After the expiration interval is reached, the data on the media is no longer considered to be valid. The media can then be reused.

Client recovery information

DRM enables the machine information needed to help recover the Tivoli Storage Manager clients stored in the Tivoli Storage Manager database. This information includes:

- ► Tivoli Storage Manager client machine location, machine characteristics, and recovery instructions

- ► Business priorities associated with the Tivoli Storage Manager client machines

- ► Description, location, and volume and diskette labels of Tivoli Storage Manager client boot media

- ► Centralized management of the disaster recovery process

5

Installing and configuring a solution

This chapter provides the detailed information for installing and configuring a Storage Manager solution in a customer environment. In this chapter we introduce administrative standards for Storage Management environment, client backup and restore types, schedule and automation function, and detailed implementing tips for disaster recovery management.

This chapter discusses the following topics:

► What to know about the second and third sections of the test objectives
► Installing Storage Server
► Installing ISC and Administration Center
► Adding client node
► Basic client functions
► Defining storage pools

5.1 What to know about the second and third sections

In this chapter we discuss the second and third sections of the Tivoli Storage Manager 5.4 Implementation Certification Test objectives. At the end of these sections, you should be familiar with the following activities:

► Define client and administrative standards (such as naming conventions, password conventions, option files, installation directories, and Web enablement).

► Determine Storage Manager communications options and parameters.

► Determine additional functions needed for special data types (such as databases, network attached storage, and mail servers).

► Identify tasks based on detailed design.

► Install the Tivoli Storage Manager code.

► Apply systems maintenance code.

► Use client and server configuration tools.

► Customize Tivoli Storage Manager option files.

► Install and customize the ISC and Administration Center

► Test the Tivoli Storage Manager client/server communications.

► Define the hardware to Tivoli Storage Manager.

► Define the storage hierarchy.

► Format and define the database, recovery log, and disk storage pool volumes.

► Test the hardware through Storage Manager (format a volume and define to storage Manager).

► Define data policies.

► Configure storage pools and volumes.

► Customize the server, including licensing.

► Register/customize the administrator with proper authority.

► Register and customize clients, including Web clients and Web administration.

► Use the Health Monitor and other Administrator Center functions.

► Set up schedules and automation.

► Define tape sharing (library sharing, LAN-free).

► Set up enterprise administration.

► Establish log mirrors, database mirrors, and storage pool copies.

- ► Test basic client functions (backup/restore, archive/retrieve).
- ► Test advanced client functionality (backup sets, image backup, mobile, LAN-free).
- ► Identify performance contention areas (bottlenecks).

5.2 Tivoli Storage Manager Server installation

In this section, we discuss the installation of the server and client code that will be used to implement Tivoli Storage Manager. We provide implementation checklists and describe the various planning considerations you will need to take into account to get the best results.

Tivoli Storage Manager server and client code fixes and enhancements are released regularly. You must check the following resources for the latest available fixes: Tivoli Storage Manager Web site, readme files, and other FTP services.

5.2.1 Basic installation

Installation process can be done either with an operating system utility (for example, IBM AIX smitty or Windows setup wizard) or the command line interface. Within the basic installation, you create the Tivoli Storage Manager database, log volumes, standard policy domain and backup policies, and device drivers. Your final task is to register the administrative and backup and archive client.

5.2.2 Customization

The Tivoli Storage Manager server has a number of options and settings that control its operation. Tivoli Storage Manager uses the options specified in the options file at server start-up. You can specify many server settings using an administrative interface.

Options file

Tivoli Storage Manager provides a server options file with a set of default options to start the server. The supplied file contains information on what options and option values can be specified. You can display the current server options using the `query options` command.

Table 5-1 shows the option files and their locations.

Table 5-1   Server option file locations

Platform	File location
AIX	/usr/tivoli/tsm/server/bin/dsmserv.opt
Linux	/opt/tivoli/tsm/server/bin/dsmserv.opt
HP-UX	/opt/tivoli/tsm/server/bin/dsmserv.opt
MVS	As specified by DDname OPTIONS
Solaris	/opt/tivoli/tsm/server/bin/dsmserv.opt
Windows 2003	C:\Program Files\Tivoli sm\server1\dsmserv.opt

Although the number of server options is very large, only a small number need to be changed for each server.

These options fall into the following categories:

► Communication
► Server storage
► Client server
► Site-dependent options
► Database and recovery log
► Data transfer
► Message
► Event logging
► Security and licensing
► Miscellaneous

Make sure that the following options are set before starting any configuration on your Tivoli Storage Manager server:

► COMMmethod
► TCPport
► VOLUMEHistory
► DEVCONFig
► MAXSESSIONS

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

Settings

Tivoli Storage Manager provides default server run-time settings. These settings are stored in the Tivoli Storage Manager database and persist across server restarts. You specify the server settings through the administrative interface.

You are able to alter your server name by using the following command: `set servername new_name`. You can also display the current server level and settings with the `query status` command.

5.2.3 Licensing

Tivoli Storage Manager, Extended Edition includes several features such as disaster recovery manager, large libraries and NDMP that are included with the product. Other Tivoli Storage Manager licensed products such as Tivoli Storage Manager for Mail, Tivoli Storage Manager for Databases, Tivoli Storage Manager for ERP, Tivoli Storage Manager for Hardware, Tivoli Storage Manager for Space Management / HSM do not need to be registered to the Tivoli Storage Manager server.

Registering licensed features

If you bought a base Tivoli Storage Manager server license, you can obtain licenses for licensed features, and register those licenses by specifying the license files using the `register license` command.

You can register any or all of these features. For example, if you want to license the server for extended edition, the command you should issue is:

```
register license file=tsmee.lic
```

Saving your licenses

When license registration is complete, the licenses are stored in a file named NODELOCK in the server start directory. You may need to register your licenses again for any of the following reasons:

► The server is corrupted.
► The server has moved to a different machine.
► The NODELOCK file is destroyed or corrupted.

Monitoring licenses

There are two commands that monitor the license registration on your Tivoli Storage Manager system:

```
query license
```

Use this command to display details of your current licenses and determine licensing compliance.

`audit license`

Use this command to compare the current configuration with the current licenses.

5.2.4 User management and security

Administrators, client nodes, and client sets provide user management and security.

Administrators

A Tivoli Storage Manager administrator manages Tivoli Storage Manager resources on the server such as storage pools, devices, and data management policies. An administrator may also be responsible for backup and restore of client data. The number of administrators and their level of privileges will vary according to environment.

There are two ways to create a Tivoli Storage Manager administrator ID: The `register node` and `register admin` commands. The `register admin` command is used to explicitly create an administrator ID with certain defined privileges. The `register node` command automatically creates an administrator ID with the same name as the node and owner access privilege to the node.

In an emergency situation, administrative commands can be issued from the server console to correct situations, such as forgotten administrator passwords and locked IDs.

Privileges are granted to an administrator through the `grant authority` command. You need system privileges to issue this command. These are the administrative user privileges:

► System

An administrator with system privilege can perform any Tivoli Storage Manager administrative task. All other privileges are included when an administrator is assigned system privilege.

► Policy

An administrator can have either unrestricted or restricted policy privilege. An administrator with unrestricted policy privilege can manage the backup and archive policy definitions (for example, management classes, copy groups, and schedules) for client nodes assigned to any policy domain. When new policy domains are defined to the server, an administrator with unrestricted policy privilege is automatically authorized to manage the new policy domains.

An administrator with restricted policy privilege can perform the same operations as an administrator with unrestricted policy privilege, but only for specified policy domains.

► Storage

An administrator can have either unrestricted or restricted storage privilege. An administrator with unrestricted storage privilege has the authority to manage the Tivoli Storage Manager database, recovery log, and all storage pools. An administrator with unrestricted storage privilege cannot define or delete storage pools.

Administrators with restricted storage privilege can manage only those storage pools to which they are authorized. They cannot manage the Tivoli Storage Manager database or recovery log.

► Operator

Administrators with operator privilege control the immediate operation of the Tivoli Storage Manager server and the availability of storage media.

► Analyst

An administrator with analyst privilege can issue commands that reset the counters that track server statistics, but otherwise can perform only query commands.

► Node

Administrators with node privilege can remotely access a Web backup-archive client and perform backup and restore actions on that client using an administrative user ID and password.

Client nodes

Every client node has to be registered and assigned a password to identify itself against its designated server, and to prevent access to the other users' data.

To simplify administration and automation, the client password is usually stored locally on the client using this option:

```
passwordaccess generate
```

This allows the password to authenticate itself against the server. This prevents situations where users forget their passwords and are unable to operate basic functions of Tivoli Storage Manager.

The password is encrypted before being stored. When the password expires, the Tivoli Storage Manager server and client negotiate a new random password according to the configured password rules. The client will then re-encrypt this password and store it locally.

With Tivoli Storage Manager, the client system can encrypt its data during backup or archive using standard Advanced Encryption Standard (AES) 128-bit encryption. If you use the AES 128-bit encryption feature to encrypt your data during backup or archive, data is encrypted during network transfer and remains encrypted in the storage pool. You must have the encryption key in order to restore or retrieve the data. Data Encryption Standard (DES) 56-bit encryption can be used however AES 128-bit is the default.

Client option sets

One way of maintaining client nodes is to centrally define options that the clients will use during backup, restore, archive, and retrieve operations. This is done using client option sets.

A Tivoli Storage Manager client session has a set of options that are used during the backup, archive, restore, or retrieve processes. Options can be specified in two ways:

► Client options file
► Client options set

The first is mandatory, while the second is optional.

The client options file is a configuration file (or files, in the case of UNIX clients) that is local to each Tivoli Storage Manager client. It contains entries of valid client options with an associated value. It also contains include-exclude file specifications.

A client option set is a set of Tivoli Storage Manager client options stored in the Tivoli Storage Manager database. It is used in conjunction with a client options file. An option set can be associated with one or more clients, but a client can be associated with only one option set.

You use client option sets for ease of administration. Management of an environment where the number of clients is growing and the number of options is increasing can be complex. The use of client option sets centralizes the management of those options and clients. It is easier to update a client options set once than to perform the same update to the local client options file on each node.

5.2.5 Defining a new database volume

Tivoli Storage Manager tracks all volumes defined to the database as one logical object. For performance reasons, you can divide the database volumes into different disks, so that the actual read and write operations are balanced across the disks available. For example, Figure 5-1 shows a database that consists of

four volumes (file01 through file04). The volumes could exist on the same disk, or be spread across any number of actual disk drives. Although the volumes are separate files from an operating system perspective, Tivoli Storage Manager tracks the database as a single logical image. The database and recovery log files information can be found in the dsmserv.dsk file.
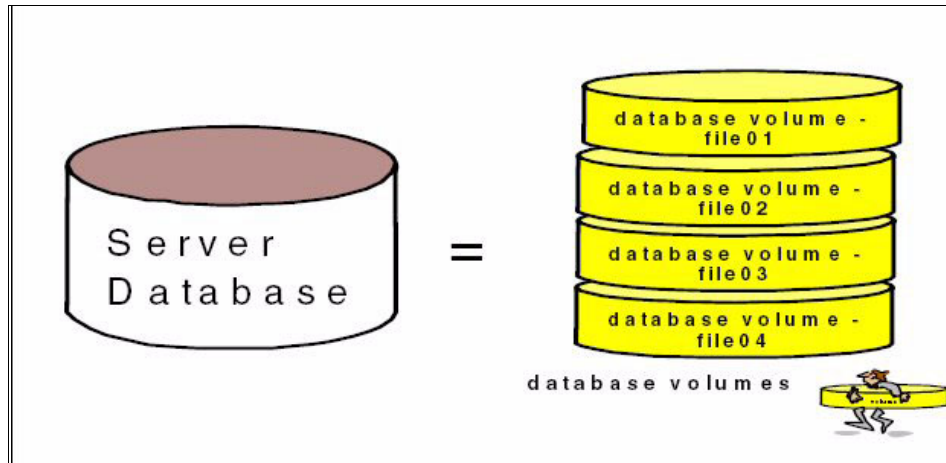
Figure 5-1   Tivoli Storage Manager Server database

There are two methods of successfully allocating a new database volume. The most common and easiest to use is to allocate a new Tivoli Storage Manager database volume by using the `define dbvolume` command and the FORMATSIZE option from either the Tivoli Storage Manager console or an administrative client.

The second method is to use the DSMFMT utility command to allocate the volume file and then the `define dbvolume` command to define the database volume to Tivoli Storage Manager.

5.2.6 Defining new recovery log volume

The recovery log contains information about updates that have not yet been committed to the database. For example, when a client is running backups, all transactional data that controls this operation is first written to the recovery log.
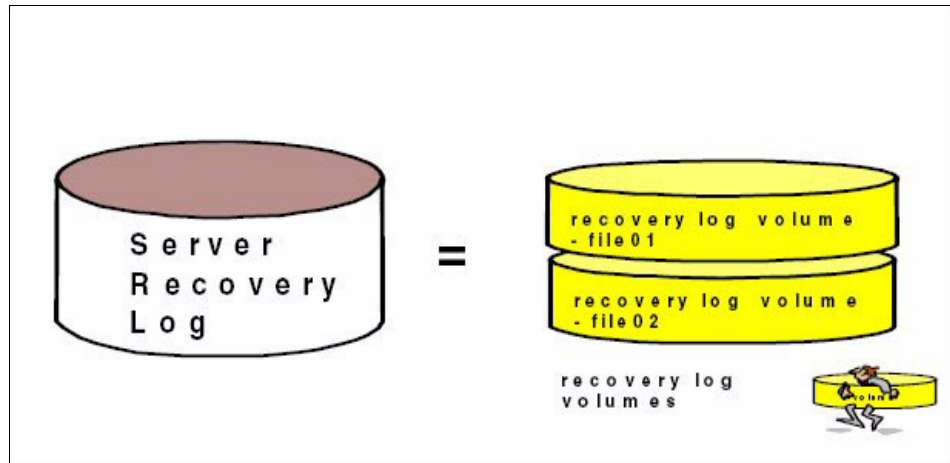
Figure 5-2   Tivoli Storage Manager Recovery log and its volumes

For performance reasons, you can split the recovery log volumes into different disks, so that the actual read and write operations are balanced across the disks available. Figure 5-2 shows a recovery log that consists of two volumes (file01 and file02). The volumes could exist on the same disk, or be spread across any number of actual disk drives. Although the volumes are separate files from an operating system point of view, Tivoli Storage Manager tracks the recovery log as a single logical image.

There are two methods of successfully allocating a new recovery log volume. The most common and easiest method to use is to allocate a new Tivoli Storage Manager recovery log volume by using the `define logvolume` command with the FORMATSIZE option from either the Tivoli Storage Manager console or an administrative client.

The second method is to use the DSMFMT utility command to allocate the volume file and then the `define logvolume` command to define the recovery log volume to Tivoli Storage Manager.

Setting the log mode to ROLLFORWARD

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

The two log modes are:

▶ NORMAL

Tivoli Storage Manager saves only those records needed to restore the database to the point of the last backup (point-in-time recovery). Tivoli Storage Manager deletes any unnecessary records from the recovery log.

► ROLLFORWARD

Tivoli Storage Manager saves all recovery log records that contain changes made to the database since the last time it was backed up.

Example 5-1   set logmode command

```
tsm > set logmode rollforward
```

Defining database backup trigger

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

Tivoli Storage Manager uses the settings you specify with this command only when the log mode is set to ROLLFORWARD, which you previously configured with the `set logmode` command (see Example 5-1). With the `define dbbackuptrigger` command, you specify the percentage of the assigned capacity of the recovery log that can be used before Tivoli Storage Manager begins a backup of the database.

Setting the expansion trigger

Tivoli Storage Manager lets you fully automate the process of increasing the database and recovery log through these triggers:

► Database space trigger

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

The MAXIMUMSIZE limit for the database is 543,184 MB (530 GB). A value of zero, (0) or omitting this parameter disables checking for maximum size. If later on you need to change the trigger definition, you can use the `update spacetrigger` command.

Whenever the system detects that the database needs extra space, it triggers the expansion by the `query actlog` command.

► Recovery log space trigger

To define a new space trigger for the recovery log in the /tsm/log/primary directory (Tivoli Storage Manager generates the volume names), use the `define spacetrigger log` command.

You can use a MAXIMUMSIZE from 9 MB through 13,564 MB (13 GB) for the recovery log. A value of zero (0) or omitting this parameter disables checking for maximum size. If later on you need to change the trigger definition, you can use the `update spacetrigger` command.

Whenever the system detects that the recovery log needs extra space, it will allocate additional space based on the spacetrigger parameters. If the log size is greater than 12 GB, the spacetrigger settings will not add more space to the log.

5.2.7 Mirroring

You can prevent the loss of the database or recovery log due to a hardware failure on a single drive, by mirroring drives. Mirroring simultaneously writes the same data to multiple disks. However, mirroring does not protect against a disaster or a hardware failure that affects multiple drives or causes the loss of the entire system. While Tivoli Storage Manager is running, you can dynamically start or stop mirroring and change the capacity of the database.

► Mirroring offers the following benefits:

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

► However, there are also costs:

– Required direct access storage device for mirrored volumes doubles
– Performance decreases

Database mirroring

When you first install Tivoli Storage Manager, it does not mirror the database. To mirror the database, you must first create new database volumes using the DSMFMT utility with the `dsmfmt -db` command.

The next step is to define all database volumes in Tivoli Storage Manager with the `define dbcopy` command.

After a database volume copy is defined, Tivoli Storage Manager synchronizes the volume copy with the original volume. You can request information about mirrored database volumes by repeating the `query dbvolume` command.

Recovery log mirroring

When you first install Tivoli Storage Manager, it does not mirror the recovery log. To mirror the recovery log, you must first create the new recovery log volumes by using the DSMFMT utility with the `dsmfmt -log` command.

The next step is to define all recovery log volumes in Tivoli Storage Manager using the `define logcopy` command. After a volume copy is defined, Tivoli Storage Manager synchronizes the volume copy with the original volume. You can request information about mirrored recovery log volumes with the `query logvolume` command.

5.2.8 Database backup

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

The DBBACKUPTRIGGER command starts database backups as needed. However, it is a good idea to run a *full* database backup, because the trigger may not start in the next few hours. You should also backup the database at least once a day as a scheduled operation. The triggered database backups are mainly regarded as a "safety valve" used in case of unusual activity on the server.

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

▶ Full backup (TYPE=FULL)

This backup copies the entire Tivoli Storage Manager database.

▶ Incremental backup (TYPE=INCREMENTAL)

This backup copies only those database pages that have been added or changed since the last time the database was backed up. This operation reduces recovery log size.

▶ DBSnapshot (TYPE=DBSNAPSHOT)

This backup is a full snapshot database backup. The entire contents of the database are copied, and a new snapshot database backup is created without interrupting or clearing the existing full and incremental backup series for the database. Database snapshot stores the volume information in the volume history file. The benefits of using database snapshot backup are:

– Quicker recovery
– May be used in disaster recovery situations
– It is independent of the full / incremental database backup rotation

5.2.9 Data Policies

The Tivoli Storage Manager data management is governed through one or more policy domains, policy sets, management classes and copy groups.

► Policy Domain: A policy domain is a group of policy users with one or more policy sets that manages the data and storage resources. Each Tivoli Storage Manager client is associated to a policy domain. The following command will define a policy domain called AIX.

```
define domain AIX
```

► Policy Set: A policy set contains one or more management classes that exist for a policy domain. While a number of policy sets can be created, only one can be active at any given time. The following command will define a policy set named POLICY1 for the AIX policy domain:

```
define policyset AIX POLICY1
```

► Management Class: A management class defines how the data will be managed. There can be multiple archive or backup management classes for a policy domain. The following command will define the management class SYSTEM to the AIX domain and POLICY1 policy set:

```
define mgmgtclass AIX POLICY1 SYSTEM
```

► Copy Groups: Copy groups control the data retention of the backup versions and archive copies. Copy groups also control which storage pool will be used for each management class. The following command will create a backup copy group for the AIX policy domain, POLICY1 policy set, SYSTEM management class:

```
define copygroup AIX POLICY1 SYSTEM type=backup destination=backuppool
verexists=2 verdeleted=1 retextra=30 retonly=60
```

► The policy definitions can be checked to verify that the policy set is complete. The following command will verify the policy set definitions:

```
validate policyset AIX POLICY1
```

► The following command will activate the policyset if no significant errors are found:

```
activate policyset AIX POLICY1
```

► The following command will set the SYSTEM management class as the default for the AIX domain policy, POLICY1 domain.

```
assign defmgmtclass AIX POLICY1 SYSTEM
```

5.3 Installing ISC and configuring Administration Center

This section provides instructions for installing the ISC and Administration Center on Windows machines. Starting with the IBM Tivoli Storage Manager Version 5.4, you can only access IBM Tivoli Storage Manager via the Administration Client or Administration Center.

Version 6.0.1.1 of the Integrated Solutions Console is required to use the Administration Center. This version of the Integrated Solutions Console is to be installed prior to installing the Administration Center. From the Administration Center on any platform, all Tivoli Storage Manager Servers on any platform can be accessed.

> **Note:** The Integrated Solutions Console and the Administration Center can only manage Tivoli Storage Manager server 5.3.0.0 and later.

Access to Tivoli Storage Manager server V5.4 from a browser is no longer supported. If you try this, you will receive the message shown in Figure 5-3.

Not Supported

ANR4747W The web administrative interface is no longer supported. Begin using the Integrated Solutions Console instead.
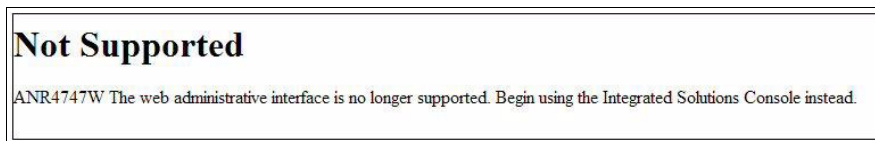
Figure 5-3   Web GUI is not supported

Installing and configuring the ISC-Admin Center consists of the following tasks:

► Installing the ISC

► Installing the Administration Center

► Adding the Tivoli Storage Manager server to the Admin Center

► Operating and configuring Tivoli Storage Manager server from the Admin Center

Installation the IBM Integrated Solution Center (ISC)

Follow these steps (Figure 5-4):

1. Insert the ISC CD, or download code from Internet.

2. Run `setupISC.exe`.
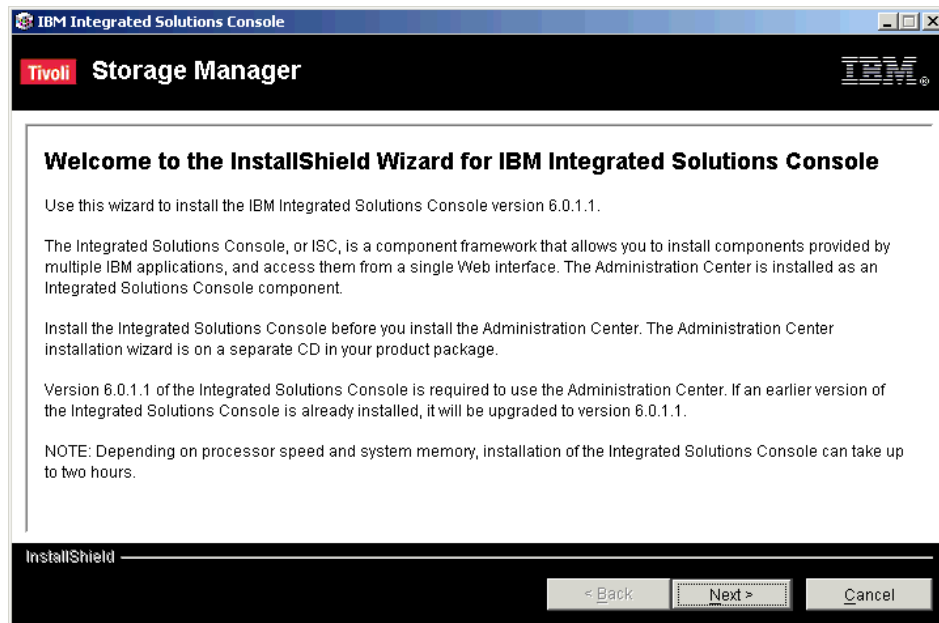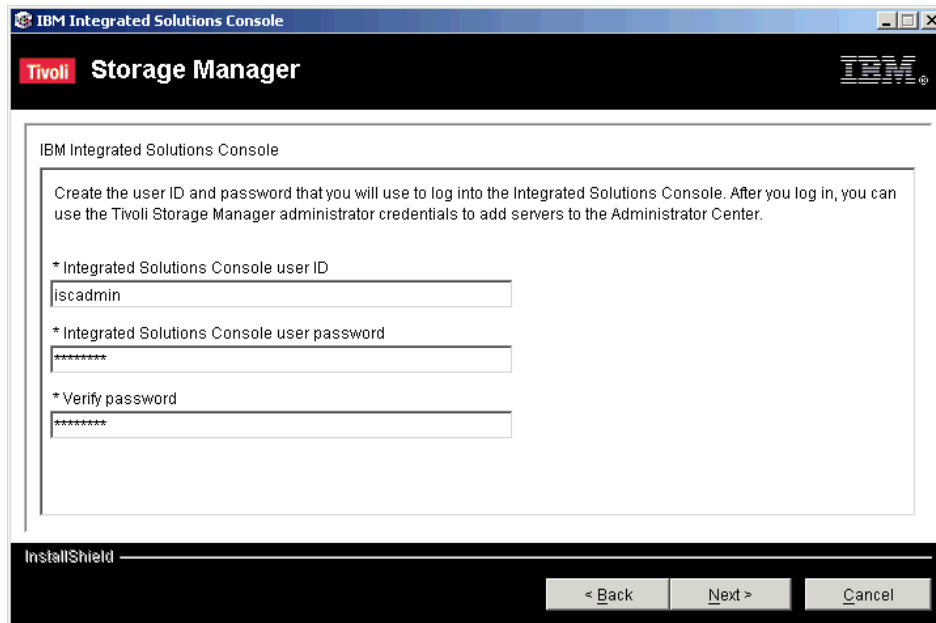
Figure 5-4   ISC installation welcome panel

3. Click **Next**. Accept the License Agreement and click **Next**.

4. Specify the directory and port you want to use to install ISC (Default C:\Program Files\IBM\ISC601). The default HTTP port is 8421, the default HTTPS port is 8422.

5. Specify the console user ID and password (default iscadmin) (See Figure 5-5 on page 125.

Figure 5-5   User ID and password creation.

6. The installation time can vary depending on the server's capacity. It can take up to 2 hours.

7. You will see the installation process panel as shown in Figure 5-6 on page 126.
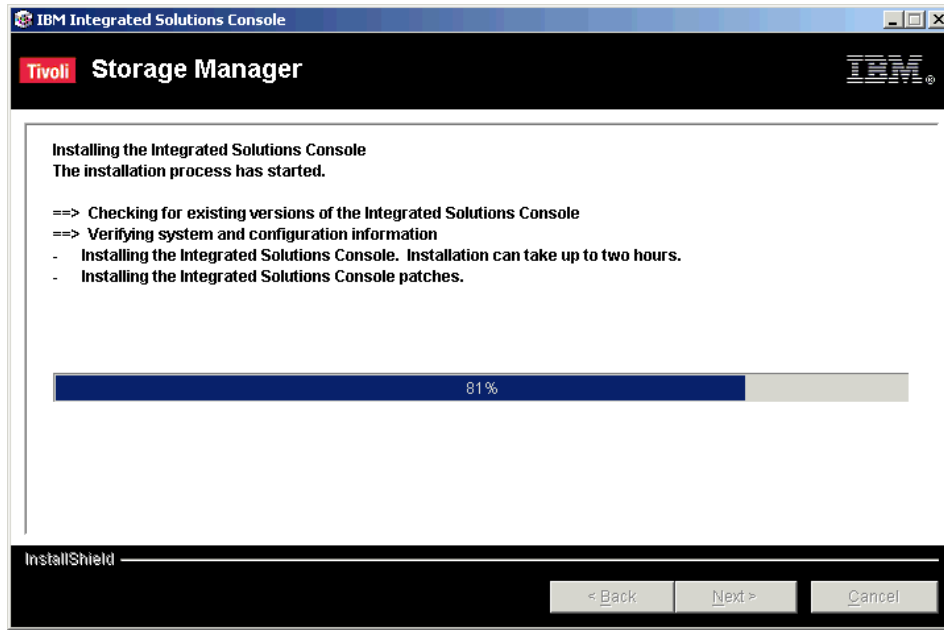
Figure 5-6   Installation starts

8. You will get the following message and URL for ISC when the installation completes, as shown in Figure 5-7.
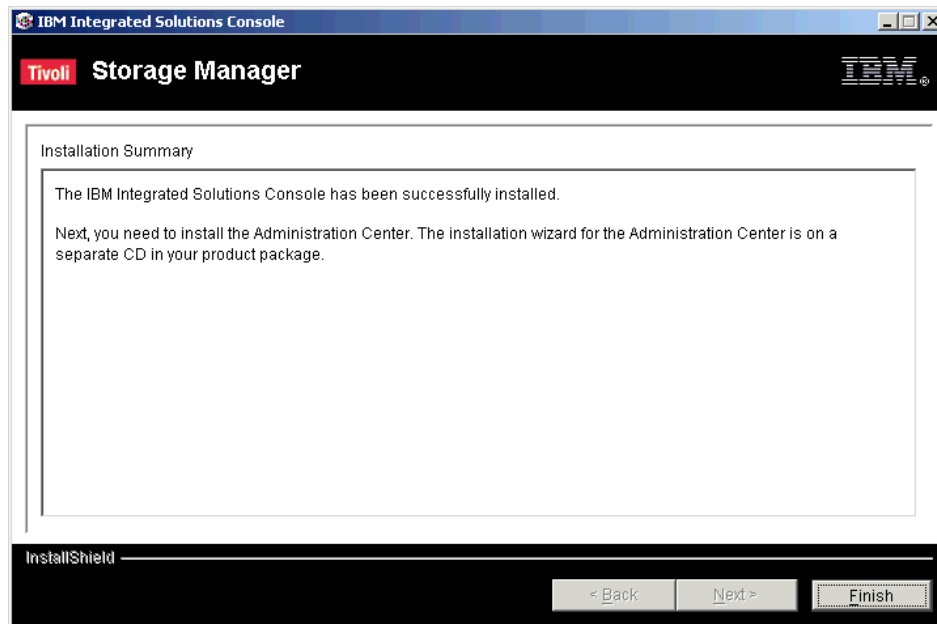
Figure 5-7   Installation completes

9. To verify the installation (Figure 5-8), open the browser and reach:

    `http://localhost:8421/ibm/console`
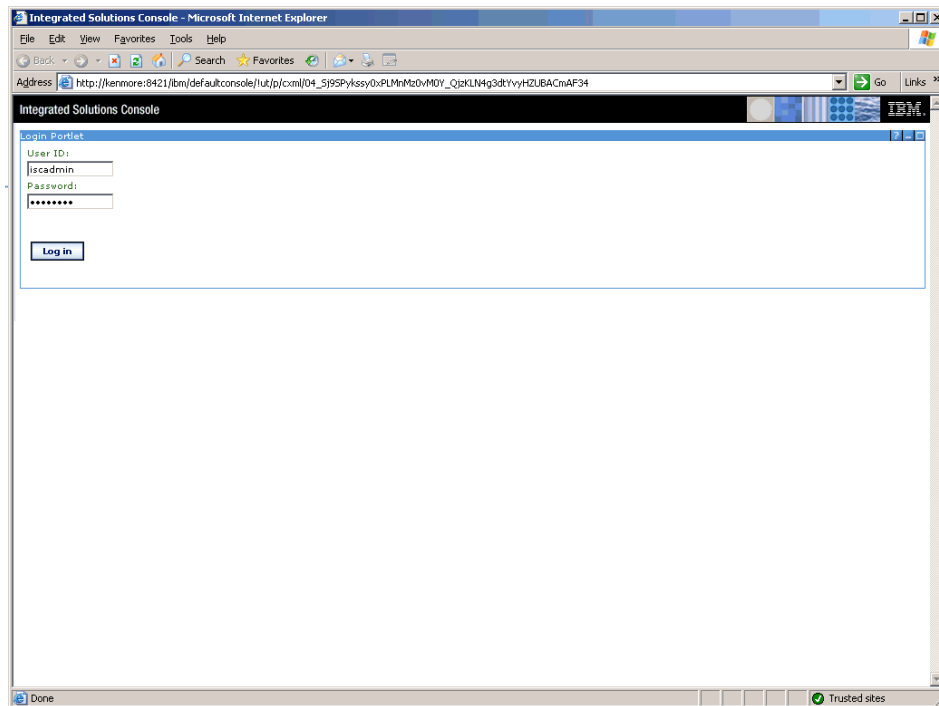
Figure 5-8   Starting the ISC console.

10. To run the ISC, enter the user ID and password that you defined in installation step 5.

With this step we have successfully installed the ISC, and now we are going to install the Administration Center.

Installation of Administration Center

Follow these steps (Figure 5-9 on page 129):

1. Insert the Administration center CD, or download code from Internet.

    Downloaded code is in tar format. Extract it to a temporary directory.

2. Run setupAC.exe.
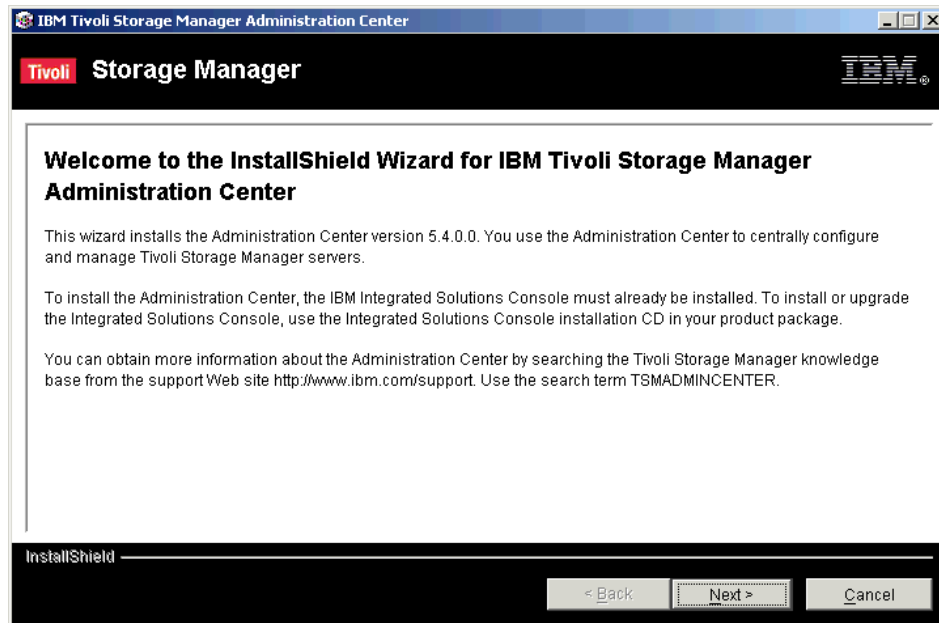
Figure 5-9   Administration Center welcome panel

3. Click **Next**. Accept the License Agreement and click **Next**.

4. The installation will ask whether ISC has been installed on the default paths. If the information is not correct, then select **I would like to update the information** and click **Next**.

5. Specify the directory where you installed ISC (Default C:\Program Files\IBM\ISC601).

6. Specify the Web administration port (default port is 8421).

7. Specify a console user ID (default iscadmin) and password.

8. Select the version for which you require additional language support as shown in Figure 5-10 on page 130. If English is being used, make no selection and press **Next**.

Figure 5-10   Language Support Panel

9. You will get the following message and the URL for ISC.
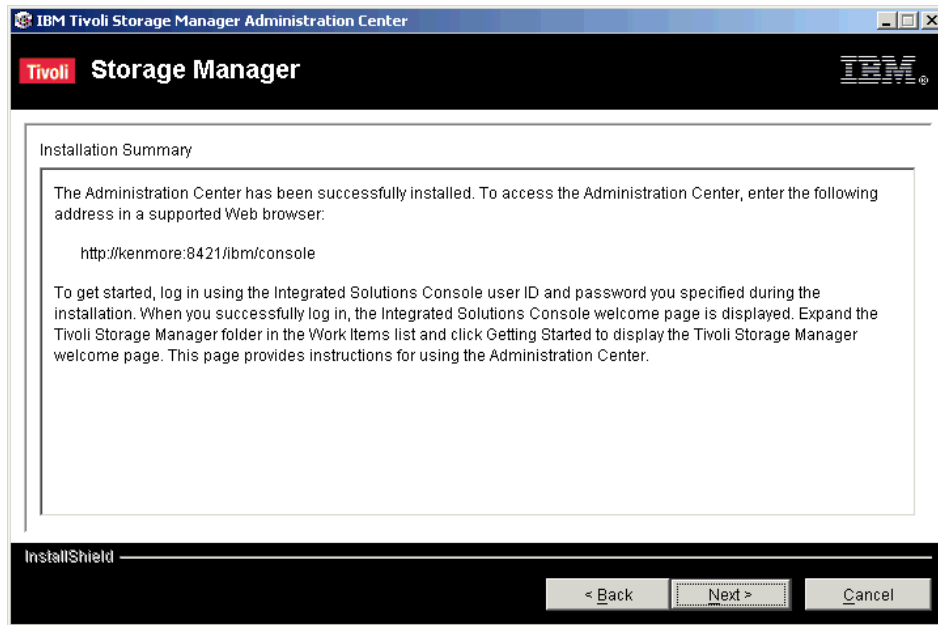
Figure 5-11   Administration Center installation completes

10.Click **Next**, and you will be transferred to the ISC login page.

To get started, log in using the Integrated Solutions Console user ID and password you created during the installation. When you successfully log in, the Integrated Solutions Console welcome page is displayed, as shown in Figure 5-12.
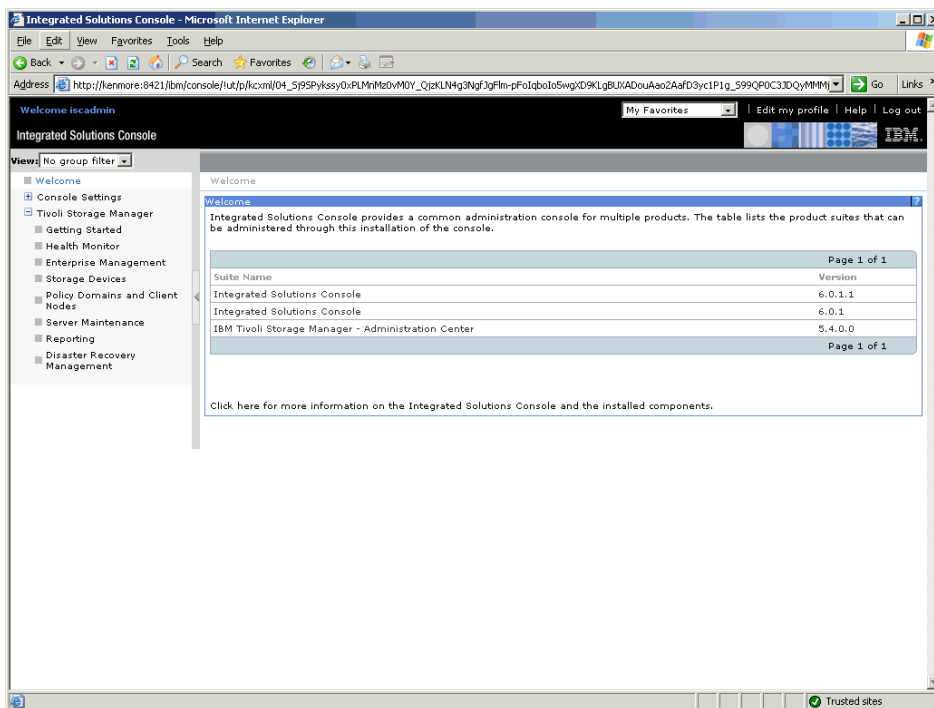
Figure 5-12   ISC login automatically start

11.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.

5.3.1 Configuring the Administration Center

After installing ISC and Administration Center, Tivoli Storage Manager servers should be added to the portal. Only Tivoli Storage Manager 5.3 and later releases can be added.

In the following sections we explain the configuration steps for the Administration Center.

Adding a TSM server to the Administration Center

Follow these steps (Figure 5-13 on page 133).

1. Open a browser and go to http://host:8421/ibm/console.

2. Enter a console user ID and password and click **login**.

3. Click **Tivoli Storage Manager**.

4. Click **Storage Devices**.

5. From the Servers window, select **Action** → **Add Server Connection** → **Go**.

6. Enter the necessary information, such as admin name, password, server address, and server port:

   a. If this is the first installation, admin name and password are ADMIN.

   b. In order help the monitor to check your Tivoli Storage Manager server, during installation, the ADMIN_CENTER admin name has been created. You have to unlock the Admin.

7. The health monitor default password is ADMIN_CENTER. Change the ADMIN_CENTER password. You can also change the refresh time for the health monitor. The default is 10 minutes.
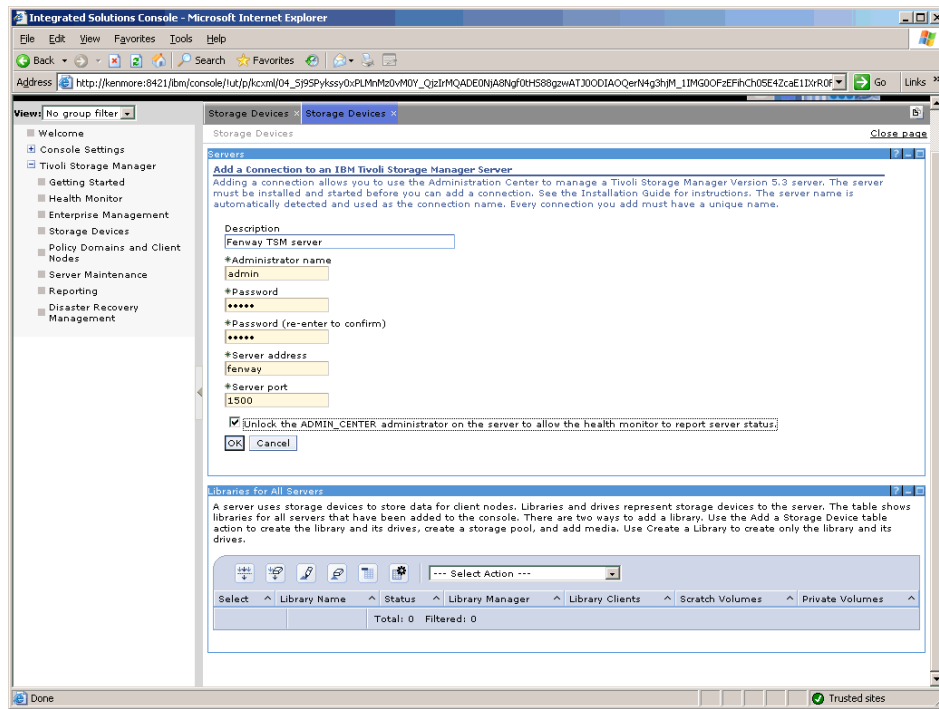
Figure 5-13   Adding a Tivoli Storage Manager server to Admin Center

8. Your Tivoli Storage Manager server is now ready for the configuration.

Checking the Health of Tivoli Storage Manager server

Follow these steps (Figure 5-14 on page 135):

1. Click **Health Monitor**.

2. Just under the Health column, you may see the following messages:

    –   Normal (green): The server is running and the health monitor identified no
        problems.

    –   Unknown (blue): The server could not be contacted. Possible causes
        include these:

        •   The server is not currently running.

        •   Network problems prevent communications with the server.

        •   The administrator name that the health monitor uses,
            ADMIN_CENTER, is locked or does not exist on the server.

        •   There were internal errors in the health monitor.

        Click the server name to get more information about possible causes.

    –   Warning or Needs Attention (yellow): The health monitor detected
        conditions that might develop into significant problems. You need to take
        preventive actions for the server's database, storage devices, or both.

        Click the server name for details.
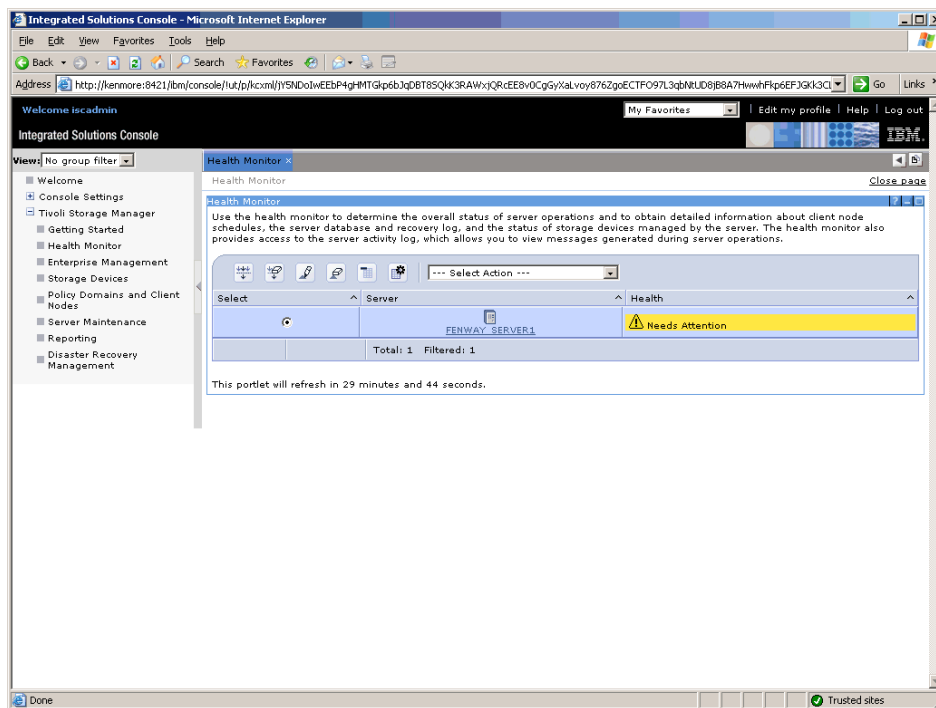
Figure 5-14   Health monitor panel

– Critical or Error (red): The health monitor detected significant problems in the server's database, storage devices, or both. The problems prevent the server from operating normally, as shown in Figure 5-15 on page 136.

The database or recovery log is out of space, or a library is not functioning. Click the server name for details.

3. Click the server name to get more information. Figure 5-15 shows the Health Monitor details.
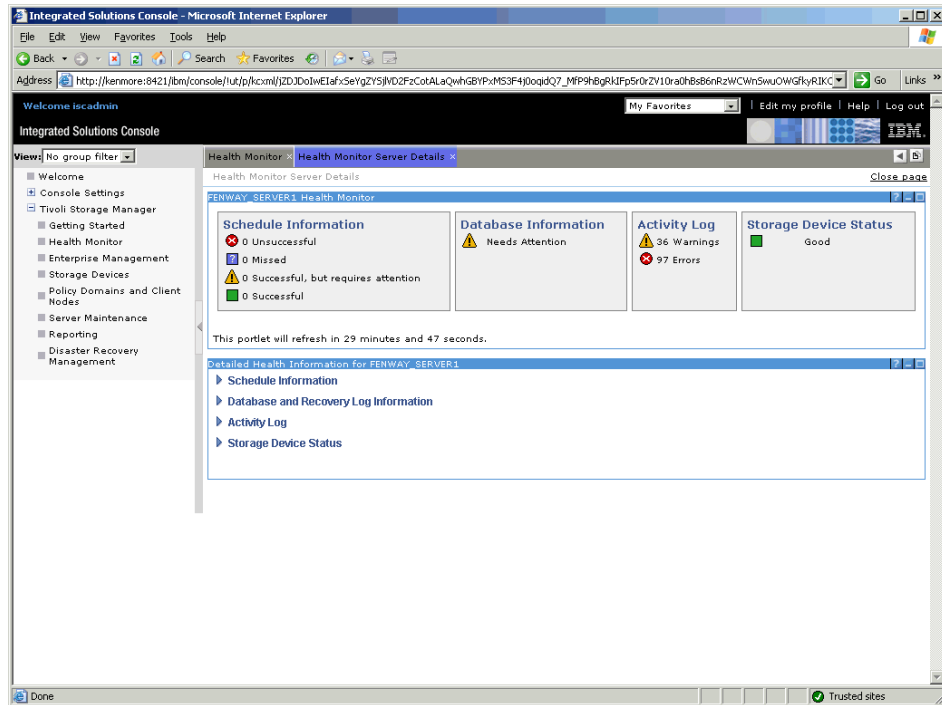
Figure 5-15   Detailed health information

4. This window will give you the overall status of your Tivoli Storage Manager server's schedule information, database and recovery log information, activity log, and storage device status. Identify the problem area and select the appropriate Detailed Report:

   a. In this page, in the Database Information column, there is a Need Attention flag. Choose **Database and Recovery Log Information**.

   b. A Detailed Health report for Database and Recovery log takes our attention to following point: Cache Hit Ratio is less than its ideal value. Change the bufferpoolsize and wait for the next refresh. When the Health Monitor refreshes its information, you will see that the status of the Tivoli Storage Manager server is now good. You can also get this ratio by issuing the `query db f=d` command.

Adding database/recovery volumes

These are the steps to follow (Figure 5-16 on page 137):

1. Select **Storage Devices**.

2. Click the server name to get the Server Properties windows.

3. Select **Database and Log**.

4. Select the **Add Volume** task from Select Action on top of the Database Section and click **Go**.

5. Fill in the Database Volume name and capacity, and click **OK**.
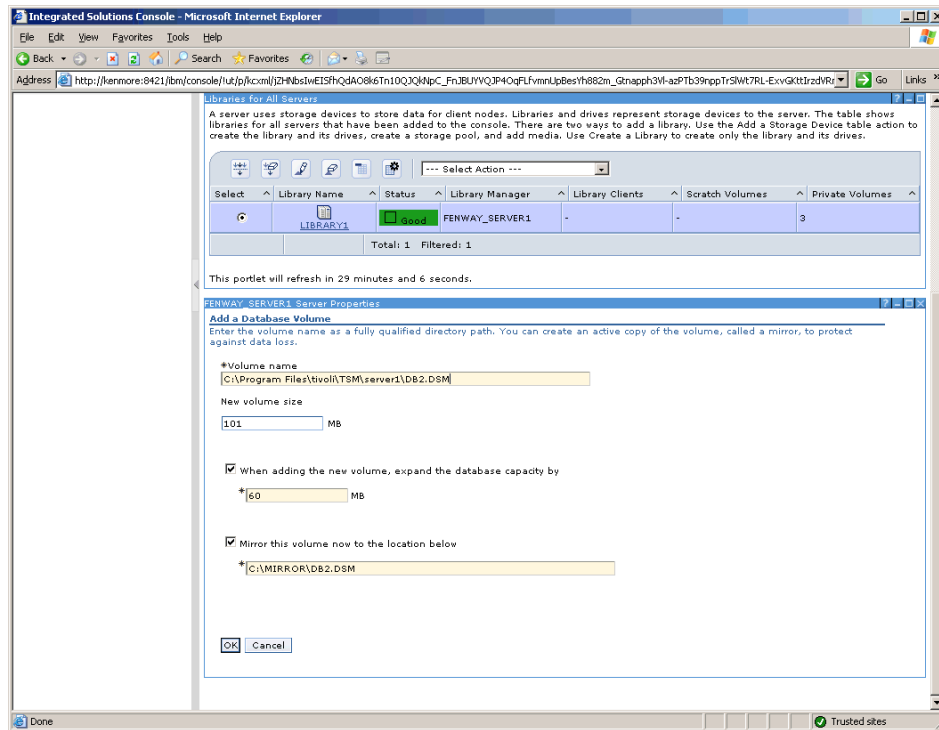
Figure 5-16   Adding a database volume

6. With this release, it is possible to extend the database while creating it.

> **Note:** To create the volume using something other than the system default path, fill in the Volume Name field with the fully qualified directory path name.

7. The health monitor in Figure 5-17 on page 138 shows that the server is healthy.
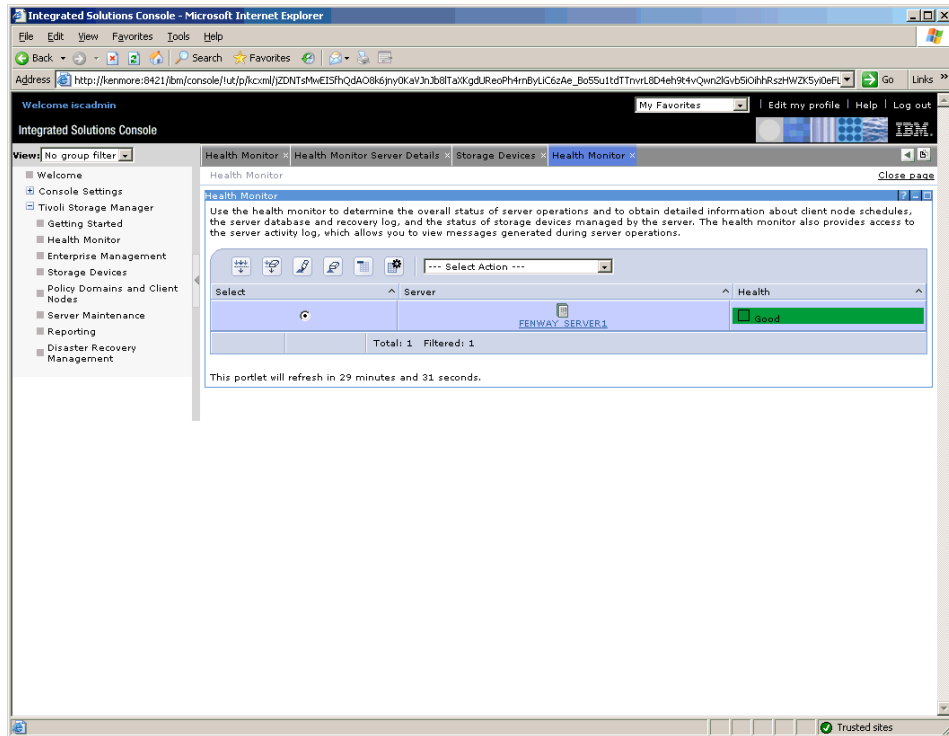
Figure 5-17   Health monitor shows the server is healthy

5.4 Backup-archive client installation

Once Tivoli Storage Manager server is installed, you can install the backup and archive client.

5.4.1 Client components

Each client has two major components that help you protect your important data:

▶ Software components

These are the software programs and customization files that you must have in place to use Tivoli Storage Manager. The most important are the client interfaces. Each is designed so that you can perform all client operations from the one that best suits your needs. For successful interaction with the server, you must configure some basic parameters in a client options file.

▶ Operation components

When you use the Tivoli Storage Manager interface to back up or archive a file, it sends a copy of the file and its associated attributes to the Tivoli Storage Manager server. The Tivoli Storage Manager client can perform backup and archive operations to send data to its designated server.

5.4.2 Client interfaces

There are three types of Tivoli Storage Manager client interfaces: Native backup-archive client, Web client, and API client.

Native backup-archive client

This client is installed on every Tivoli Storage Manager client node and provides the local interface to back up and restore operations on that node. It offers a command line and a GUI interface to back up and restore files from that system only.

Web client

The Web client provides a remote interface to back up and restore operations on a Tivoli Storage Manager client node. This client is particularly suited for help desk operations.

API client

The backup-archive client is the software component that is installed on each machine that you need to protect with backups. After installing the client code and customizing how it should interact with the Tivoli Storage Manager server, you have a working machine ready to send and receive data.

5.4.3 Code installation

The procedure for client code installation is generally the same across all platforms except for some minor differences. The best way to distribute Tivoli Storage Manager client software and updates is to use an automated software distribution application.

For more information on installing and configuring the backup and archive client, refer to the *IBM Tivoli Storage Manager Backup-Archive Clients Installation and User Guide* for your specific platform.

Tivoli Storage Manager client code fixes and enhancements are released on a regular basis. The fixes are available from this IBM Web site:

http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html

5.4.4 Customization

In the following sections we describe the various settings for the backup-archive client and the Web backup-archive client.

Backup-archive client

This section summarizes the steps to follow to customize your client installation:

► Setting environment variables

It is highly recommended that these variables are added to your client environment:

– PATH: This is the default search path that the operating system uses to locate executable files. Set this to include the fully qualified paths of the Tivoli Storage Manager client directories.

– DSM_CONFIG: Tivoli Storage Manager uses this environment variable to locate the client options file dsm.opt. It points to the client user options file for users who create their own personalized options file.

– DSM_DIR: Tivoli Storage Manager uses this environment variable to locate all other client files. It points to the executable file DSMTCA, the resource files, and the dsm.sys file (on UNIX only).

– DSM_LOG: This points to the directory where you want the dsmerror.log, dsmwebcl.log, and dsmsched.log files to reside. The error log file contains information about any errors that occur during processing. The error log is intended to help you diagnose severe errors. The Web log file contains the errors that occur when the client is accessed through the Web interface. The schedule log contains the output of the scheduled command.

► Defining client options files

The backup-archive client has at least one configuration file, which is divided into the following parts:

– Communication options

The following is a list of communication options you need to set in the dsm.opt file:

• COMMMETHOD: This option specifies the communication method you are using to provide connectivity for client-server communication. The Tivoli Storage Manager server configuration file (dsmserv.opt) must have specified the same communication method with a valid port address, so that the server can accept client requests.

• TCPSERVERADDRESS: This option specifies the TCP/IP address for a Tivoli Storage Manager server. This is either the IP address or the hostname of the Tivoli Storage Manager server.

- TCPPORT: This option specifies a TCP/IP port address for a Tivoli Storage Manager server. By default, this value is 1500.

- TCPWINDOWSIZE: This option specifies the size, in kilobytes, of the TCP/IP sliding window for your client node. This setting is highly operating system specific. You must only use the allowed values for your TCP/IP implementation. This parameter can be customized at both server and client side.

- TCPBUFFSIZE: This option specifies the size, in kilobytes, of the Tivoli Storage Manager internal TCP/IP communication buffer. This setting is highly operating system specific. You must only use the allowed values for your TCP/IP implementation.

- NODENAME: This option assigns a new name to your client node if you do not want to use the default. The default is the hostname of your client machine, which we recommend using.

> **Note:** The recommended settings for the communications options can be found in the IBM Tivoli Storage Manager Performance Tuning Guide.
>
> `http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp`

- Operational options

  - PASSWORDACCESS: This option specifies whether you want your Tivoli Storage Manager password automatically generated or always manually prompted for. We highly recommend setting PASSWORDACCESS to GENERATE.

  - REPLACE: This option specifies what you want Tivoli Storage Manager to do when it restores files that already exist on your workstation. This option applies to the restore and retrieve commands only.

  - SUBDIR: This option specifies whether you want Tivoli Storage Manager to include subdirectories of named directories. This option applies, for example, to selective, restore, archive, and retrieve.

  - TAPEPROMPT: This option specifies whether to wait for a tape to mount if it is required for a backup, archive, restore, or retrieve process, or to prompt you for a choice.

  - SCHEDMODE: This option specifies whether you want to use the client-polling mode (your client node periodically asks the Tivoli Storage Manager server for scheduled work), or the server-prompted mode (the server contacts your client node when it is time to start a scheduled operation). All communication methods can use the client polling mode, but only TCP/IP can use the server prompted mode.

- Include-exclude options

On all non-UNIX platforms, all options reside in the client options file, dsm.opt (Preferences file for Macintosh), which resides by default in the Tivoli Storage Manager client directory.

On UNIX, you can set options in three different files:

- Client system options file (dsm.sys): In the client system options file, a root user sets options that are required to establish communication with a Tivoli Storage Manager server, and options that authorize users on your workstation to use Tivoli Storage Manager services.

- Client user options file (dsm.opt): In the default client user options file, a root user can set options that determine which Tivoli Storage Manager server your client node contacts and that specify the formats to use for date, time, and numbers. A root user can also set options that affect backup, archive, restore, and retrieve processing. In addition, users can also create their own personalized client user options file if they want to use different options. Users can overwrite an option contained in a client user options file by entering a different value for the option with an appropriate Tivoli Storage Manager command.

- Include-exclude options file: In the include-exclude options file, a root user can set options to exclude specific files from backup services. In addition, a root user can set options to associate specific files with different management classes. This file can be called any legal file name and be located in any directory, provided that you reference its full path name in the dsm.sys file.

▶ Defining include-exclude lists

The include-exclude options must be placed as part of the client options file dsm.opt on non-UNIX platforms, or in the dsm.sys file or a separate include-exclude file on UNIX. Also, on UNIX, you need to add the name of the file in your client system options file dsm.sys with the keyword INCLEXCL as follows:

```
INCLEXCL /tsm/server1/inclexcl.file
```

The INCLUDE option specifies files within a broad group of excluded files that you want to include for backup services. You also use this option to assign a management class either to specific files or to all files to which you have not already assigned a specific management class and for which you do not want Tivoli Storage Manager to use the default management class.

The EXCLUDE option excludes files from backup services. When you back up files, any files you exclude are not considered for backup. Include-exclude options are checked from the bottom to the top of the list until a match is found. If so, the processing stops and checks whether the option is INCLUDE or EXCLUDE. If the option is INCLUDE, the file is backed up using the

assigned management class (Example 5-2 on page 143). If the option is EXCLUDE, the file is not backed up.

Example 5-2   Sample include-exclude

```
Given the following include/exclude list:

    exclude /.../core
    include /home/gibes/options.scr
    include /home/ross/arianna
    include /home/gibes/doodads/drivers
    exclude /home/gibes/*
    include /core/data/.../* sysdata
    exclude.dir /home/ross

In this example all directories and files in /core/data will be backed up. It
will use the sysdata management class. The exclude /home/gibbs/* statement will
backup all directories in the path. The exclude.dir statement does not backup
any files in the directory structure.
```

5.4.5 Establishing the session

The Tivoli Storage Manager backup-archive client node is required to register with the Tivoli Storage Manager server. Once registered, the Tivoli Storage Manager client can start its communication with the server by first completing a sign-on process. This sign-on process requires the use of a password that, when coupled with the node name of the client, ensures proper authorization when it connects to the server.

Multi-session

Tivoli Storage Manager exploits the multithreading capabilities of modern operating systems by transparently initiating multiple backup-archive or restore/retrieve sessions on the client where necessary for rapid processing and data transfers between the client and the server.

The administrator and the user each have controls to influence the number of sessions that a client can start. On the server, the global setting MAXSESSIONS limits the total number of sessions of any kind that may be present. The client node setting MAXNUMMP, in its server definition, controls how many mount points (for sequential devices such as tape drives) a client may allocate. Finally, the RESOURCEUTILIZATION setting in the client option file increases or decreases the ability of the client to create multiple sessions.

Consolidating Multiple Clients under a Single Client Node

Backups of multiple nodes that share storage can be consolidated to a common target node name on the Tivoli Storage Manager server. For example, several nodes in a GPFS™ cluster, NODE_1, NODE_2, and NODE_3, can back up to the same node (NODE_OLIV) on the server. This is useful when the machine responsible for performing the backup may change over time, such as with a cluster. Consolidating shared data from multiple machines under a single name space on the Tivoli Storage Manager server means that the directories and files can be easily found when restore operations are required. Backup time can be reduced and clustered configurations can store data with proxy node support.

Client nodes can also be configured with proxy node authority to support many of the systems that support clustering failover. By granting client nodes proxy node authority to another node, you gain the ability to back up, archive, migrate, restore, recall, and retrieve shared data on multiple clients under a single node name on the Tivoli Storage Manager server. When authorized as agent nodes, Tivoli Storage Manager nodes and HSM clients can be directed to back up or restore data on behalf of another node (the target node).

5.4.6 Transaction

All data sent to Tivoli Storage Manager storage during a backup or archive session is done within the bounds of a transaction. That means that not every single file is sent to the server separately. Tivoli Storage Manager combines multiple files in one transaction to reduce overhead and to increase performance.

The size of a transaction is controlled by the server setting TXNGROUPMAX, which specifies the maximum number of client files that can comprise a single transaction, which increases the number of files that are backed up. The client setting TXNBYTELIMIT specifies the maximum number of bytes that can be sent. The limit that is reached first determines the complete transaction.

> **Note:** The recommended settings for the transaction options can be found in the IBM Tivoli Storage Manager Performance Tuning Guide.
>
> http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp

5.5 Setting up storage pools

When you are defining a new device for the Tivoli Storage Manager server, the device must first be configured for the operating system.

Tivoli Storage Manager treats physical devices as two types: Random access devices and sequential devices, and these are connected locally to the server.

▶ Random access devices:

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

– To store the database and recovery log.

– To store client data that has been backed up or archived from client nodes. The client data is stored in disk storage pools.

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

– Random access volumes

– Storing data sequentially using a FILE device class

▶ Sequential storage devices:

For Tivoli Storage Manager, sequential storage devices are tape libraries, optical devices, WORM devices and active data storage pools.

To configure a device for Tivoli Storage Manager, six components must be configured. These can be thought of in two component classes, physical and logical. These components must be configured as follows:

▶ Physical:

– The library: AUTOMATED or MANUAL
– The drive (or drives, if there are more than one in your library)
– Tape or disk path

▶ Logical:

– The device class: Here you specify the device type and associate the device class name with the device type and the library name. Here the device type is from a Tivoli Storage Manager defined list, and the command is common to all Tivoli Storage Manager server platforms, for example, 4MM, DLT, FILE or LTO.

– The storage pool: This specifies the device class to use for the storage pool that you are creating and includes specific storage pool rules for handling data.

– The volumes: This is where you label the volumes and either define them as scratch volumes or as private volumes belonging to a particular storage pool.

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

5.5.1 Defining a physical library

First, you must define a library, even if you are using a single tape drive. The `define library m8mm` command defines a manual library named M8MM. There is no device name associated with a manual library, because the library part is considered virtual.

5.5.2 Defining a physical drive

After defining the library, you should define a drive for the library to use, utilizing the device number that the operating system uses to recognize it. The drive therefore must be configured for that operating system.

The `define drive m8mm drv0 device=/dev/mt0` command defines a drive named DRV0 in the manual library, named M8MM. The physical tape drive identifier is /dev/mt0, because it is an 8 mm tape unit that uses the Tivoli Storage Manager built-in tape device drivers.

5.5.3 Defining a device logical class

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

To define a device class named C8MM that uses a predefined Tivoli Storage Manager device type of 8 MM in the library named M8MM, use this command:

```
define devclass c8mm devtype=8mm library=m8mm format=8500 estcap=5g
mountretention=5
```

When a tape is idle in a drive in the library, it remains there for 5 minutes (MOUNTRETENTION=5). It uses a tape format of 8500 with a capacity of 5 GB on a 112 m tape.

5.5.4 Defining a device physical path

Use this command to define a path from a source to a destination. A path provides access to a destination from a source. You must define the source and destination before you can define a path.

For example, if a path is required between a server and a drive, you must first issue the DEFINE DRIVE command and then issue the DEFINE PATH command. A path must be defined after you issue the DEFINE DRIVE command in order to make the drive usable by Tivoli Storage Manager software.

The following command defines a library path:

```
define path ITSM_Server M8MM src=server autodetect=yes libr=M8MM
dev=/dev/rmt0.smc desttype=library
```

The following command creates a tape path:

```
define path ITSM DRV0 src=server autodetect=yes libr=M8MM dev=/dev/rmt0
desttype=drive
```

5.5.5 Defining a device logical class

A device class can be considered a software driver. It is defined so that Tivoli Storage Manager can communicate with the physical drive correctly. The following command defines a device class named C8MM that uses a predefined Tivoli Storage Manager device type of 8MM in the library named M8MM:

```
define devclass c8mm devtype=8mm library=m8mm format=8500 estcap=5g
mountretention=5
```

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

Random storage pools

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

Sequential storage pools

A disk can simulate a sequential device, such as tape using the FILE device class. If you are using a single tape drive and want to perform reclamation, then the pool type you specify as the reclaim storage pool must be a primary sequential storage pool. The FILE device class is required when active data storage pools are created.

A copy pool must also use sequential access storage and can be defined on a disk using this device type. To configure a sequential storage pool, you must first define a FILE device class, and tell Tivoli Storage Manager which directory to use and what the maximum size of this file can be.

Determine the amount of disk space needed for either purposes:

► For *backup* storage pools, provide enough disk space to support sufficient daily incremental backups.

► For *archive* storage pools, provide sufficient space for a user to archive a file system of moderate size without causing migration from the disk storage pool to occur.

▶ For *active-data* storage pools, provide sufficient disk or tape space that allows for the daily rate of change for all the servers that will use this storage pool.

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

```
define devclass cfile devtype=file directory=/tsm/stgpool/ maxcapacity=12m
```

The following command lists the capacity of the storage pools:

```
query stg
```

Active data storage pool

An active-data storage pool only contains active versions of client backup data. Archive or HSM data are not allowed to use this storage pool. Active-data storage pools can use three device types, sequential access disk (using the FILE device class), tape / optical devices or sequential volumes on another Tivoli Storage Manager server.

Copy storage pool

A copy storage pool can use only sequential access storage (for example, tape, optical, or file device classes). The device class you define must represent the type of device you are using for your copy pool.

5.5.6 Shredding

When data is deleted on random disk or tape devices, it is not physically removed from the media. Shredding provides the capability to ensure that the deleted space is overwritten so that the deleted data is not easily reconstructed. Shredding can be set to run automatically or manually, it does not occur immediately after the data is deleted, the deleted space is not available to be reused until the shredding process has completed. The storage pools used for shredding need to be be configured.

The following command will overwrite the deleted data area three times.

```
define stgpool shred4 disk shred=3
```

The *shredding automatic* option needs to be added to the dsmserv.opt file.

To run manually, no changes need to be done to the dsmserv.opt file. The following command will shred the deleted data area for 60 minutes:

```
shred data duration=60
```

5.5.7 Defining primary storage pools

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

Primary storage pool for a manual library

Once you have defined the device class, you can define the storage pool you are going to use with the library.

The following command defines a storage pool named TAPEDATA using the device class named C8MM:

```
define stgpool tapedata c8mm highmig=100 maxscratch=10
```

The high migration parameter is set to 100 to stop migration because there is no NEXTSTGPOOL for this storage pool to migrate to. The MAXSCRATCH value of 10 means that when 10 scratch tapes have been used, the storage pool indicates that it is full.

Primary storage pool for a disk

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

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

```
define stgpool diskdata disk description="Data Storage" high=70 low=30
```

Primary storage pool for sequential disk

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

The following command defines a storage pool named DISKFILE using a device class named CFILE:

```
define stgpool diskfile cfile maxscratch=100 highmig=100
```

Because it simulates a sequential storage pool, the MAXSCRATCH parameter must be included, and because there is no next storage pool in this case, migration is disabled.

The following command defines an active data storage pool named DATA.

```
define stgpool DATA CFILE pooltype=activedata
```

5.5.8 Defining storage pool logical volumes

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

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

Manual library

There are two methods of labeling a volume in a manual library. The first is a one-step process that uses the `label libvolume` command.

The second method uses the DSMLABEL utility to create a volume and the Tivoli Storage Manager `define volume` command to define the volume to a storage pool.

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

Disk storage pool

Disk volumes can be defined to a Tivoli Storage Manager storage pool using either a one-step or two-step process. This section describes the one-step process.

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

```
define volume diskdata /tsm/stgpool/diskdata1.dsm formatsize=50
```

The two-step process requires preparing the disk volume using the DSMFMT utility and then defining the volume to the storage pool using the `define volume` command.

5.5.9 Audit volume

The `audit volume` command checks for inconsistencies between the database references for what files are stored on a volume, and what is actually found to be

stored on the volume. The volume must be mounted in a tape drive so that the contents can be read.

This differs from the **query content** command, which displays only the database view of the what is stored and does not require the volume to be mounted.

The **audit volume** command is especially useful when integrity errors are detected on a volume, as you can determine what, if any, files on the volume are still readable and take appropriate action to restore it from a copy storage pool.

5.5.10 Backup storage pool

The **backup stgpool** command backs up a primary storage pool to a copy pool. If the primary storage pool is a sequential storage pool, then this command requires two physical drives. The input volume from the primary storage pool will be mounted in one, and the output volume from the copy storage pool will be mounted in the other. The two drives can be in the same or different tape libraries. To back up a random access storage pool to a sequential storage pool requires only one drive.

The following command backs up a storage pool named TAPEDATA to a copy pool named OFFDATA.

```
backup stgpool tapedata offdata
```

5.5.11 Checking in library volumes

The **checkin libvolume** command checks in a library volume that has been physically placed in the library, so that it can be seen by Tivoli Storage Manager.

The following command checks in volume DAN001 as a scratch volume:

```
checkin libv 3583lib dan001 status=scratch
```

5.5.12 Checking out library volumes

There are three parts to this operation:

1. Check out the library volume.

    This command checks out the library volume named DAN001 from a library named 3583lib.

    ```
    checkout libv 3583lib dan001
    ```

2. Check for outstanding mount requests.

3. Reply to the mount request.

5.5.13 Migration

Migration is forced by updating the threshold parameters for migration on the storage pool to a value lower than its current utilization level. This will cause migration to start automatically.

To clear all data from one pool to the next storage pool, you change the high and low values to 0. The important information about migration is that the level of granularity is by filespace.

5.5.14 Data movement

The `move data` command moves the data from one volume to another within the same storage pool (this requires that two drives be available if it is a sequential storage pool) or to another storage pool.

As stated previously, collocation keeps client files close together, maybe even in one volume, which may increase restore performance. In large environments, collocation usually uses a great deal of tape, so it can be an uneconomical choice for them. Moreover, numerous mounts are required during migration and reclamation. However, collocation can be used in large environments to decrease restore time. Tivoli Storage Manager enables you to move data down to a level of a client's filespaces using the MOVE NODEDATA command. This offers a chance to collocate data for specified nodes only on demand so the next restore would finish faster.

Using the same command, Tivoli Storage Manager also enables data movement between different storage pools. Unlike migration, selected client data can be moved rather than the entire storage pool. Moving data to a storage pool offers fast and parallel access, such as a disk storage pool. So multiple clients can restore their data faster using multiple sessions.

5.5.15 Reclamation

Reclamation, which is shown in Figure 5-18 on page 153, is used to free complete tape (or optical) volumes in sequential storage pools. Because Tivoli Storage Manager keeps a defined number of versions of files when it performs incremental backups, the oldest copy of a file (beyond the defined number of versions to keep) gets marked for expiry. This file is deleted when the next expiration occurs.

For volumes that are not marked as off-site volumes, reclamation can occur only after the volume is full and then begins to empty because of file deletion. The reclamation process is very device intensive; it typically requires at least two available drives in the library.
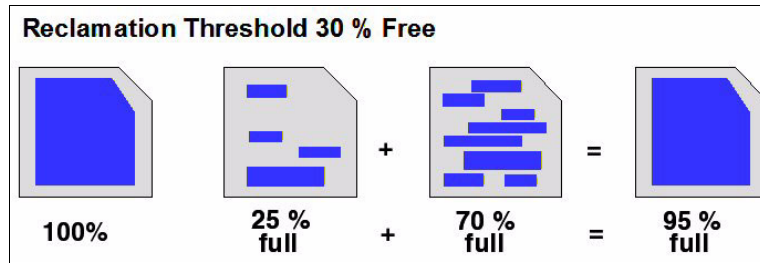
Figure 5-18   Reclamation of tapes

For off-site volumes, reclamation can occur regardless of whether the volume has ever been filled. An off-site volume is eligible for reclamation when the percentage of unused space in the volume is greater than the reclaim parameter value. The unused space includes space that has never been used in the volume and space that has become empty because of file deletion.

Single drive reclamation

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

The storage pool specified as the reclaim storage pool must be a primary sequential storage pool on the system.

Reclamation of off-site volumes

Tivoli Storage Manager cannot physically move the data from one off-site volume to another because the volumes are in a vault and are not available in the library.

Tivoli Storage Manager manages reclamation for an off-site copy pool by obtaining the active files from a primary storage pool or from an on-site volume of a copy pool. These files are then written to a new volume in the copy pool, and the database is updated. A message is then issued that an off-site volume was reclaimed.

5.5.16 Reduce restore times

Storage pool configuration can reduce restore times using the following techniques:

- ► Collocation (minimizes the number of tape volumes used to store client data)
- ► Disk caching (restores data from disk even if it has already been migrated)
- ► Moving data to fast access storage pools or consolidating data before restoring them

5.5.17 Collocation

Collocation is a process whereby the server attempts to keep all files belonging to a client node on a minimal number of sequential access storage volumes.

Collocation has four options to use:

- ► Collocation by group allows you to define a group of client nodes to use to use the fewest number of volumes. This could be used when you have high capacity tapes and the client nodes archive / backup data is small in comparison to the volume capacity.
- ► Collocation by node allows you to put each client node's backup data on the fewest number of volumes. This is the same as putting COLLOCATE=YES in previous Tivoli Storage Manager versions.
- ► Collocation by filespace allows you to put each filespace for a client node on the fewest number of volumes.
- ► Setting COLLOCATE=NO will disable collocation.

To collocate data when files from different client nodes are mixed in the same storage pool, set collocation to group or node when you define or update a sequential storage pool. This reduces the number of volume-mount operations required when users restore or retrieve many files from the storage pool at the same time that you increase mounts during a migration process. Collocation (Figure 5-19 on page 155) therefore improves access time for these operations.
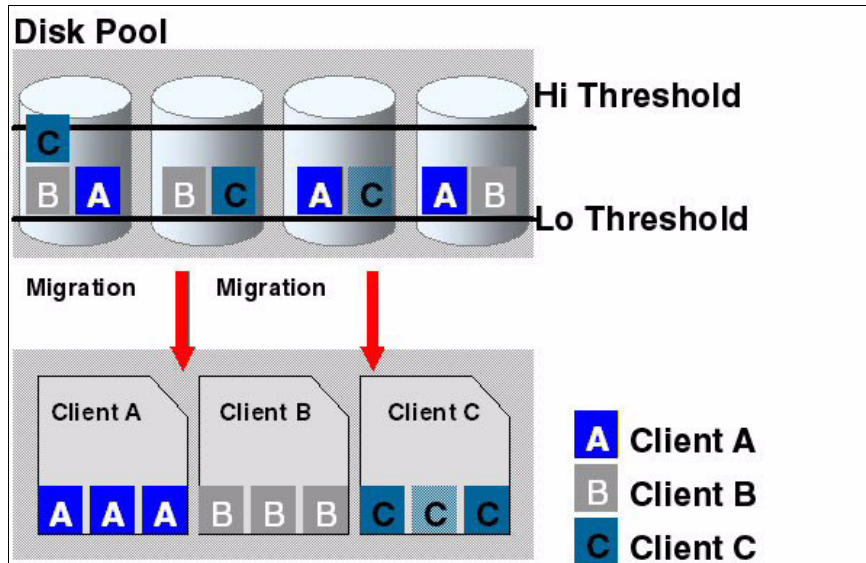
Figure 5-19   Storage pool collocation

5.5.18 Collocation by group

The utilization of tapes is improved when using collocation by providing for collocation of a group of nodes whose data will be collocated together on sequential media. Collocation can reduce the number of volume mounts that are required when a large number of files are restored from a sequential-access storage pool by minimizing the number of tapes on which the data is stored. As tapes provide larger and larger capacity, they can be more efficiently used when data from multiple nodes is stored on them while still preserving the benefits of collocation. The efficiency of Tivoli Storage Manager internal data-transfer operations is improved by transferring all nodes in the group together. This can result in less database scanning, with fewer tape passes and reduced tape wear for file processing.

If no groups and members are defined, it works the same as using collocation by node. Collocation by group can yield the following benefits: Reduce unused tape capacity by allowing more collocated data on individual tapes; minimize mounts of target volumes; and minimize database scanning and reduce tape passes for sequential-to-sequential transfer. For newly defined storage pools, the default storage pool collocation setting is now GROUP.
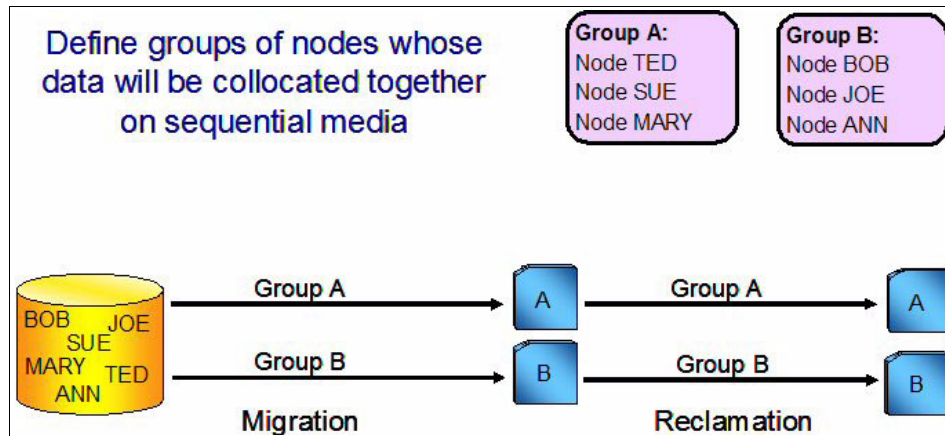
Figure 5-20   Collocation by groups

The following commands are used to define a collocation group and collocation member;

DEFINE COLLOCGROUP group_name

DEFine COLLOCMEMBER group_name node_name,node_name, etc.

5.5.19 Query content

The `query content` command queries the contents of a volume and lists all the files contained on it.

5.5.20 SQL commands

Tivoli Storage Manager uses a database that accepts structured query language (SQL) commands. Tivoli Storage Manager provides three system catalog tables:

▶ SYSCAT.TABLES: Contains information about all tables that can be queried.

▶ SYSCAT.COLUMNS: Describes the columns in each table.

▶ SYSCAT.ENUMTYPES: Defines the valid values for each enumerated data type.

For select operations, you can simply use the TABLES, COLUMNS, and ENUMTYPES to display the database structure:

▶ SELECT * FROM TABLES
▶ SELECT * FROM COLUMNS
▶ SELECT * FROM ENUMTYPES

There are many commands that can be generated in this way to create a customized query to the Tivoli Storage Manager database. Some examples of these commands are:

▶ `select * from volumes where error_state='YES'`

This command lists all volumes from the database table VOLUMES that are in an error condition. In our example, this command failed because we did not have any volumes in an error condition.

▶ `select node_name, client_version, client_relase from nodes`

This command lists all the Tivoli Storage Manager clients with their Tivoli Storage Manager version and release.

▶ `select node_name,count(*) from backups group by node_name`

This command lists how many files have been backed up from each node. If you replace BACKUPS with ARCHIVES, you see the number of files archived from each node.

5.6 Backup

Tivoli Storage Manager can back up both files and raw logical volumes. When backing up files, the Tivoli Storage Manager server database keeps a list of all files and their attributes (for example, time, date, size, access control lists, and extended attributes). At each file backup operation, this list is compared to the current file system on the client workstation to determine new, deleted, and changed files. Raw logical volumes are treated as separate entities, and the management class policy is applied to the entire image as a whole.

During a backup, the client first establishes a session with the Tivoli Storage Manager server. After that, it sends the data using the transaction controls.

Tivoli Storage Manager stores a number of backup versions for each file or object on each client node. If and when the number of versions stored on the server exceeds the number set by the Tivoli Storage Manager administrator, older versions are deleted as newer versions are made. When you back up files, Tivoli Storage Manager also backs up all related directory information and access information.

There are two types of file backups: Incremental and selective. An incremental backup creates backups of files, directories, or subdirectories that are new or have changed since the last incremental backup. A selective backup creates backups of specific files or entire directories unconditionally.

A file-level backup can be extended for WAN-connected clients using sub-file backup. Typically, the amount of bandwidth that these clients use to connect to

the Tivoli Storage Manager server is small. Using sub-file backup, only the parts of a file that have changed are transferred to the server.

Another method of backup is called image or volume backup. In this case the backup process does not distinguish between single files, but sends the specified volume as one single object to the Tivoli Storage Manager server.

5.6.1 Traditional LAN and WAN backup

In a traditional LAN and WAN environment, the Tivoli Storage Manager backup and archive client or application reads data from locally attached disks and sends it over the LAN to the Tivoli Storage Manager backup server, as shown in Figure 5-21. The server receives the data then writes it out to its storage pool (tape, disk, or optical media) based on predefined policies and server configuration. Data is read and written by both the Tivoli Storage Manager client and Tivoli Storage Manager server machines. In addition, control information is also sent over the LAN to the Tivoli Storage Manager server.
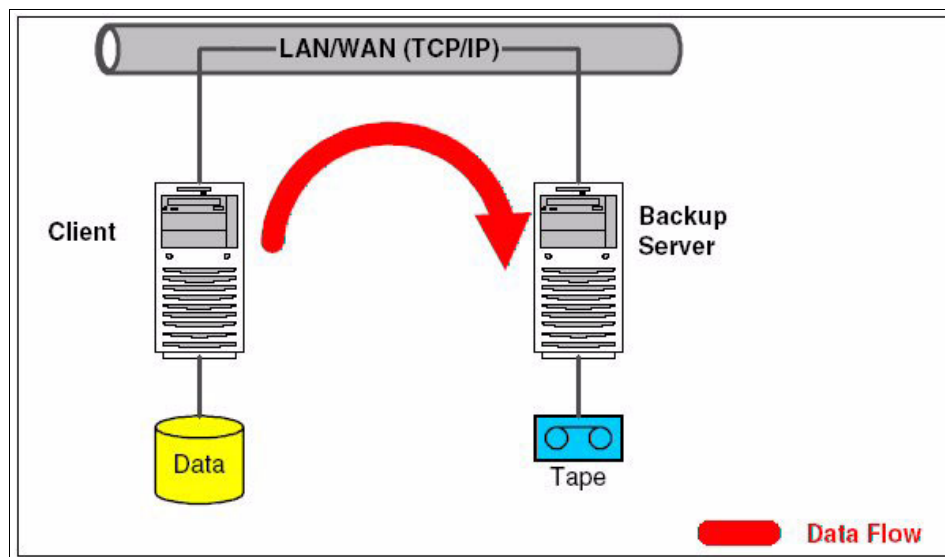
Figure 5-21   LAN - Wan backup

5.6.2 Progressive backup

Tivoli Storage Manager is unique in offering an incremental or progressive backup methodology for backing up client data. This approach can remove the need for periodic full dumps because only the changed files are backed up. This

can have significant benefits in backup time, least number of tapes used, reduced network traffic, size of backup servers, and manageability.

The incremental backup operation is a full scan of the client's file systems, which backs up all files and other information (and only those things) necessary to ensure that the Tivoli Storage Manager inventory matches the current state of the client's storage. This means that the first time this is run on a new client, everything is backed up. Each time after this, only new and changed files are sent.

During the incremental backup, the client queries the Tivoli Storage Manager server so that it knows what files are currently stored. The client uses this information to:

- ▶ Back up new files.
- ▶ Back up files whose contents have changed.
- ▶ Expire backup versions on the server for files that were deleted from the workstation.

5.6.3 Selective backup

During a selective backup, Tivoli Storage Manager sends copies of the files to the server even if they have not changed since the last backup. This is useful if you want to back up multiple files that are not in the same directory structure, regardless of their actual status in the Tivoli Storage Manager server. It may also apply where you want to enforce a complete backup.

However, remember that versioning still applies. If you back up a file multiple times when it has not changed, this will result in having multiple copies of exactly the same file on the server, instead of a number of different versions of the file.

This more or less defeats the purpose of Tivoli Storage Manager version control. To avoid that, use the incremental backup technique command to back up only changed and new files. Typically, selective backup will only be used in special circumstances.

5.6.4 Logical volume backup

With Tivoli Storage Manager, you can back up a file system or raw logical volume as a single object from your client machine. The Tivoli Storage Manager client accomplishes this by dynamically loading an image plug-in utility that sends the object to the server with the Tivoli Storage Manager API.

This capability is currently available for the IBM AIX, HP-UX, Solaris, Linux, and Windows 2003/XP clients and can be used in a logical volume whether or not

there is an associated file system. This will ensure a clean backup. Windows clients require configuring by an additional service that is available with the Tivoli Storage Manager client.

5.6.5 Image backup

An image backup is a block-by-block copy of data from the Tivoli Storage Manager client to the Tivoli Storage Manager backup server. One important function of an image restore is to accelerate recovery in a disaster recovery scenario.

With image backup, the Tivoli Storage Manager server does not track individual files in the file system image. File system images are tracked as individual objects, and the management class policy will be applied to the file system image as a whole.

Online image backup is possible for Windows 2003 and Linux clients by using the Tivoli Storage Manager snapshot function.

5.6.6 Open file backup

Some files in your system might be in use when you try to back them up. These are called open files. The Tivoli Storage Manager client for Windows 2003/XP comes with a Logical Volume Snapshot Agent (LVSA) that performs a snapshot backup of files that are open (or locked) by other applications.

Windows XP and Windows Server 2003 have the Microsoft Volume Shadow-Copy Service (VSS) that is able to perform online backup of in-use files using snapshot. The snapshot allows the backup to be made from a point-in-time copy that matches the file system at the time the snapshot is taken. Subsequent changes to the file system are not included in the backup.

While Tivoli Storage Manager attempts to back up open files, this is not always possible when the LVSA is not installed or in use. Some files are open exclusively for the application that opened them. If Tivoli Storage Manager encounters such a file, it cannot read it for backup purposes. If you are aware of such file types in your environment, you should exclude them from backup to avoid seeing error messages in the log file.

In UNIX clients, if the file is in use, Tivoli Storage Manager acts as follows. First, the Tivoli Storage Manager server rejects the backup request, then it backs up the files while they are in use, then it tries until the retry parameter is satisfied.

5.6.7 Adaptive subfile backup

As the number of mobile computers approaches 20 percent of the PC install base, many central support organizations will need to provide storage management services for their mobile and remote workers in their WAN environment.

Mobile and remote computers have limited access to the infrastructure that serves the rest of the company. Some limitations include being attached to the corporate network with reduced bandwidth, limited connect time, and minimal assistance to perform the backup.

This limited access both increases the criticality of storage management services and limits the applicability of traditional methods and policies. Tivoli Storage Manager helps resolve these problems with its adaptive sub-file backup feature, shown in Figure 5-22, which reduces the amount of data transferred while backing up changed files.
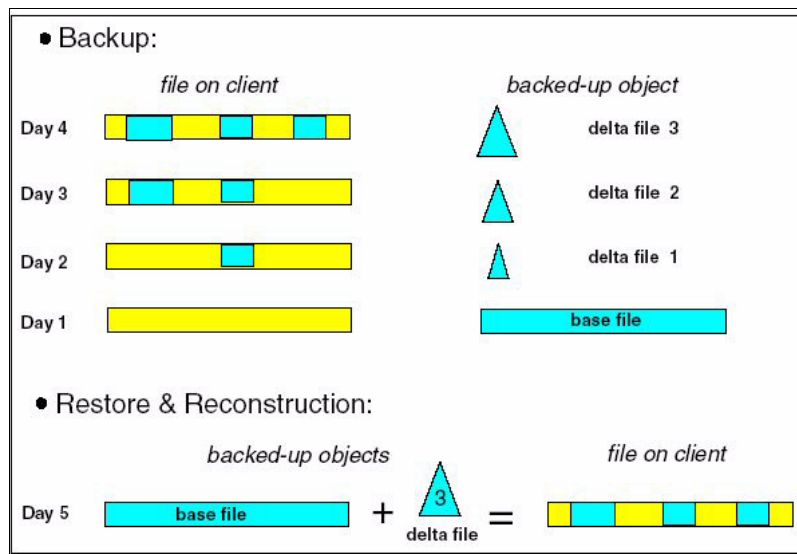
Figure 5-22   Adaptive subfile backup

Using this feature, the backup-archive client (Web client, command line, or GUI) backs up only the changed portion of a file, either on byte level or on block level, instead of transferring the whole file to the server every time. The changed file portion is backed up as a differential backup relative to the last complete backup of the file (base or reference file). It is called a *delta file*. All changes since the last complete backup of the file are included in this delta file. In the case of a restore, this allows for the restore of the whole file by restoring only two sub-file components, one delta file and the last complete backup of the whole file.

The adaptive sub-file backup, as well as the restore of a file consisting of a base file and the delta file, is completely transparent to the end user. All necessary file data separations or reconstructions happen under the backup-archive client. Also, all other Tivoli Storage Manager features, such as policy management or fault-tolerant backup and restore, still fully apply.

Adaptive sub-file backup is used for incremental as well as for selective backup. It is aware of multithreading and will work together with client data compression and encryption. This option is valid for Windows platforms only.

5.6.8 Journal-based backup

Journal-based backup provides an alternative to traditional progressive incremental backup, which under certain circumstances may dramatically increase overall backup performance. Journal-based backup is supported on all AIX and Windows clients.

The main difference between journal-based backup and progressive incremental backup is the method by which the list of backup candidate objects is derived. The backup candidate list specifies objects for a particular file system that are to be backed up, expired, or updated on the Tivoli Storage Manager server by a Tivoli Storage Manager backup-archive client.

Progressive incremental backup derives the backup candidate list by building and comparing the list of active, previously backed-up objects stored on the Tivoli Storage Manager server with the list of objects currently residing in the local file system.

5.6.9 Group backup

A group backup enables you to create a consistent point-in-time backup of a group of files that is managed as a single logical entity as follows:

► All objects in the group are assigned to the same management class.

► Existing exclude statements for any files in the group are ignored.

► All objects in the group are exported together.

► All objects in the group are expired together as specified in the management class.

The group backup function also supports differential and full backup. You usually restore the entire group to get a consistent point-in-time restore, but Tivoli Storage Manager supports single file restore from a group as well.

The group backup function is similar to the well-established archive function of Tivoli Storage Manager.

5.6.10 SAN (LAN-free) backup

SAN technology provides an alternative path for data movement between the Tivoli Storage Manager client and the server. Shared storage resources (such as disk and tape) are accessible to both the client and the server through the SAN. Data is off-loaded from the LAN and from the server processor, which can create greater scalability. Tivoli SANergy® must be used to send LAN-free backups to a disk storage pool.

LAN-free backups decrease the load on the LAN by introducing a storage agent. The storage agent can be perceived as a small Tivoli Storage Manager server (without a database or recovery log) that is installed and run on the Tivoli Storage Manager client machine. The storage agent handles the communication with the Tivoli Storage Manager server over the LAN but sends the data directly to SAN-attached tape devices, relieving the Tivoli Storage Manager server from the actual I/O transfer. A LAN-free backup environment is shown in Figure 5-23.
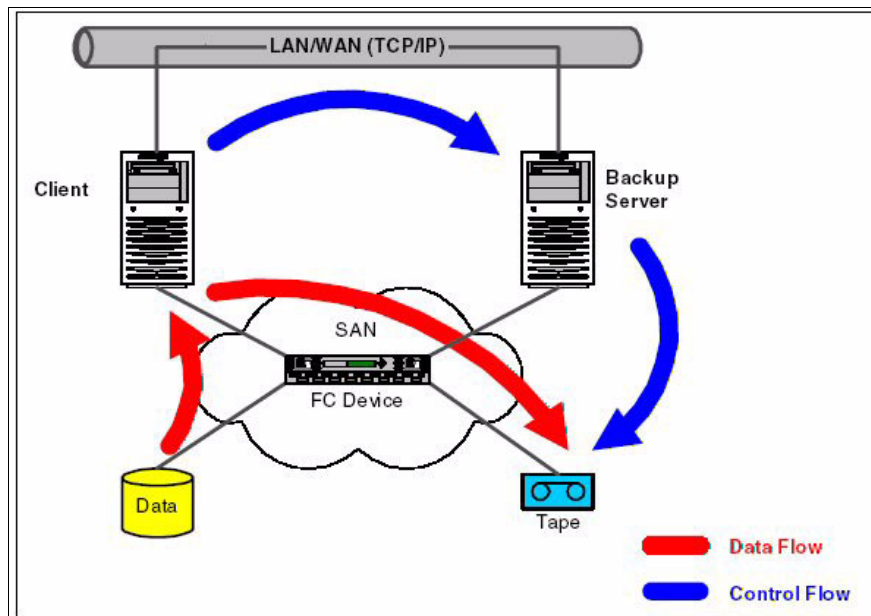
Figure 5-23   SAN backup

5.6.11 Split-mirror/point-in-time copy backup

A split-mirror/point-in-time backup occurs when a copy volume generated by operating system mirroring or a hardware-assisted instant copy function (found in many of today's high-end storage systems) is backed up to a Tivoli Storage Manager server, as shown in Figure 5-24.
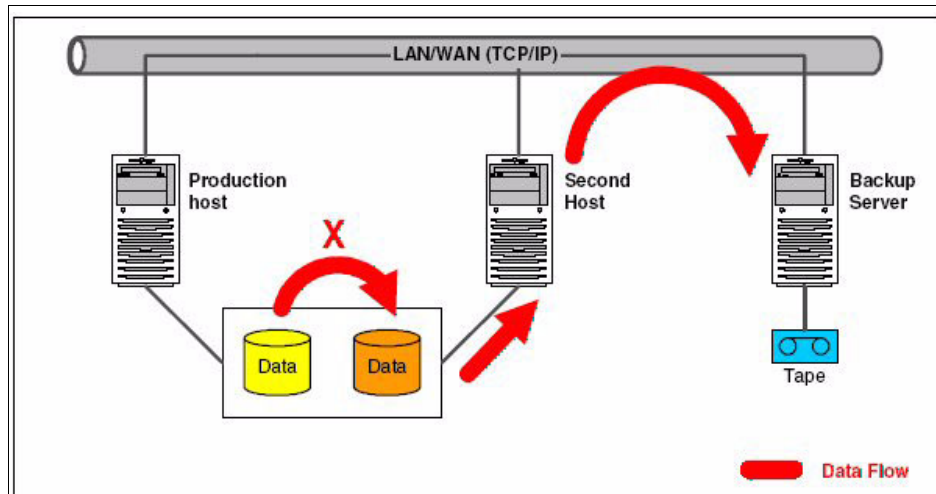
Figure 5-24   Point-in-time backup

Such a backup method almost eliminates the backup-related performance impact on the production host. This approach is facilitated and automated with the Tivoli Storage Manager for hardware components by coupling the FlashCopy function of IBM Enterprise Storage Server or SAN Volume Controller with Tivoli Storage Manager and its database protection capabilities for DB2, Oracle, and SAP R/3 databases.

5.6.12 NAS and NDMP backup

The IBM Network Attached Storage (NAS) products have been pre-installed with the Tivoli Storage Manager client, which enables an existing or planned Tivoli Storage Manager server to back up the data in the NAS system, as explained in Figure 5-25 on page 165.
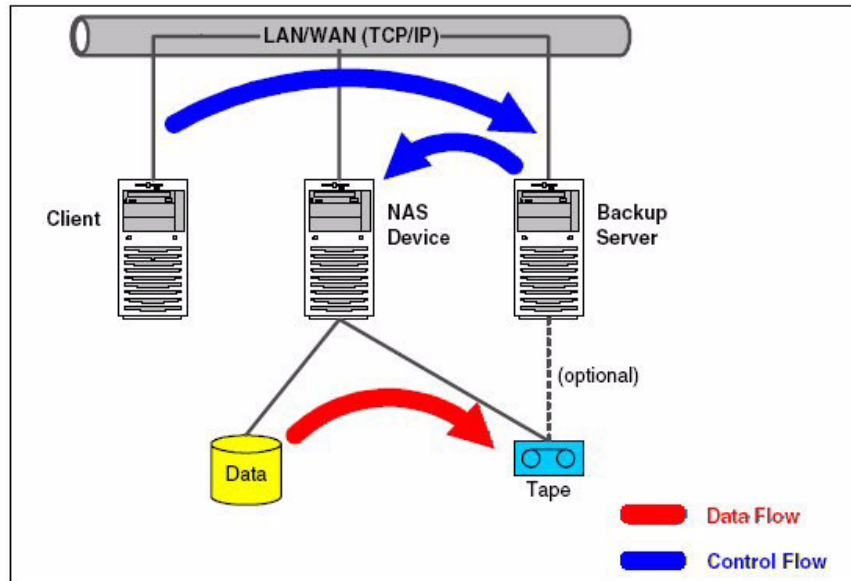
Figure 5-25   NAS and NDMP backup

This backup client is designed to provide file-level and sub-file level backup and restore functionality. A Tivoli Storage Manager environment integrated with NAS systems can be used to manage a broad range of data storage, recovery, and availability functions across the entire computing infrastructure.

Based on the Tivoli Storage Manager server's configuration, the final destination of the NAS appliance backup may be located either in the disk storage of the Tivoli Storage Manager server or an attached tape subsystem. Data is sent direct to the storage pool but the table of contents or TOC is sent to the Tivoli Storage Manager server. Automated scheduling to backup the NAS images can then be configured from the Tivoli Storage Manager server.

As with a standard Tivoli Storage Manager client, an option file is created and stored on the NAS system. Although a Network Data Management Protocol (NDMP) backup is usually started and controlled by a Tivoli Storage Manager server, the Tivoli Storage Manager Web client is also able to initiate and control an NDMP backup or restore. Using a TOC, the Tivoli Storage Manager Web client provides file-level access to the TOC so that it becomes browsable.

5.6.13 Active and inactive file versions

One of the most important concepts in Tivoli Storage Manager data management is the difference between an active backup version and an inactive backup version.

Assume a new file is created on your workstation. The next time you run a backup operation (say, Monday at 9 p.m.), Tivoli Storage Manager server stores this file. This copy of the file is known as the *active* version. When you run an incremental backup again (say, Tuesday at 9 p.m.), Tivoli Storage Manager uses this active version already stored to check back with your workstation to determine whether the file has changed since the last backup. If it has, it is backed up again. This version now becomes the active version and the copy from Monday becomes an *inactive* version.

The most recent backed-up version of the file is always the active version, as long as it still exists on the original client. Tivoli Storage Manager will keep storing a new active version and making the previous active version inactive, up to the limit of the total number of versions defined to be retained in the management class. Once this limit is exceeded, the oldest inactive version is deleted from Tivoli Storage Manager storage and can no longer be restored.

Tivoli Storage Manager controls the retention of its active and inactive versions of a file that exist on a client machine by using two criteria defined in the Management Class:

► How many versions: The parameter that controls the number of backup versions is called *verexist*. This may be set to a specific number or to *unlimited*.

► How long to keep: The *retextra* parameter controls how much time must elapse before an inactive file version is considered expired. This parameter controls how long to retain all remaining inactive files and may be set to a specific number of days or to *nolimit*, which means they will never be expired.

For a file deleted on a client machine, Tivoli Storage Manager uses different criteria:

► How many files: The parameter that controls the number of inactive backup versions is called *verdeleted*. This number is normally less than or equal to the number you have for *verexist*.

► How long to keep files: The *retextra* parameter controls how much time must elapse before an inactive file version is considered expired. This parameter controls how long to retain all remaining inactive files except for the last one and may be set at a specific number of days or to *nolimit*, which means they will never be expired.

> ► How long to retain the last file: The retonly parameter controls the last inactive copy of a file. As files get expired by retextra, you can configure Tivoli Storage Manager to manage the last inactive copy differently, so that you can keep that file for a longer period of time. It may be set at a specific number of days or to *nolimit*.

> ► Typically, configure *retonly* to be either the same value or longer than retextra because it functions as a grace period before expiring the file.

5.7 Archive

The Tivoli Storage Manager archive function stores selected files unconditionally on the server according to the applicable management class limits. *Unconditionally* means that there is no version limit and they will be retained for the defined time period regardless of whether they are deleted on the client.

Archived files are useful if you want to take a snapshot of particular files, or if you want to delete files to free space, yet still have the ability to retrieve them if required. It is common to have legislative requirements to archive business records for long periods of time, and the archive function is ideal for this purpose.

5.7.1 Packages

Archive packages are groups of files archived together with a common description. The system automatically supplies a description consisting of the time and date stamp, but you can override this with your own meaningful description. This description is used for searching and selecting of archive packages to retrieve.

You can add to an existing package on a subsequent archive operation by supplying an existing archive package description. You can retrieve individual files within a package and delete files from a package.

5.8 Backup set

You can generate a copy of your client's most recent backup from the Tivoli Storage Manager server onto sequential media. This is accomplished with the `generate backupset` command, which copies all active file versions of the file set from server storage onto the media. This copy of the backup, also called backup set or portable backup, is self-contained. It can be used independently of Tivoli Storage Manager to restore client data from a locally attached device that can also read this media, such as a CD-ROM.

This technique provides the Tivoli Storage Manager client with rapid recovery and no server and network dependency. You can also transfer the backup set from one server to another by generating the backup set on the source server, then transporting the backup set volume and defining it to the destination server, assuming both servers have the same media type,

5.9 Restore

To restore a file, a directory, or even the whole machine, you need to know two things: What you want to restore (file name, directory); and, optionally, from when (the point in time), and if you want to restore a file other than the most recent one. You do not need to know where the data actually is. When you request a file, Tivoli Storage Manager obtains the location of that particular file version from its database.

To restore files, specify the directories or selected files, or select the files from a list or GUI window. By default, only active file versions will be available for selection; however, inactive versions can be specified easily. You can restore files to their original location or specify a different directory. Collision options control whether existing files of the same name are replaced. No additional skill is needed to restore a file.

5.9.1 Restartable restore

When you are running a normal file-restore operation, Tivoli Storage Manager keeps track of the files that you have already restored, so that it can retry the operation should you have any network problems. If the restore operation terminates prematurely for any reason, such as network failure, the session state remains in the database so that it can be restarted from the last completed transaction.

This procedure (see Figure 5-26 on page 169) is called restarting the restore. Restarting the restore prevents you from having to re-send files that were already restored to the client before the session aborted.
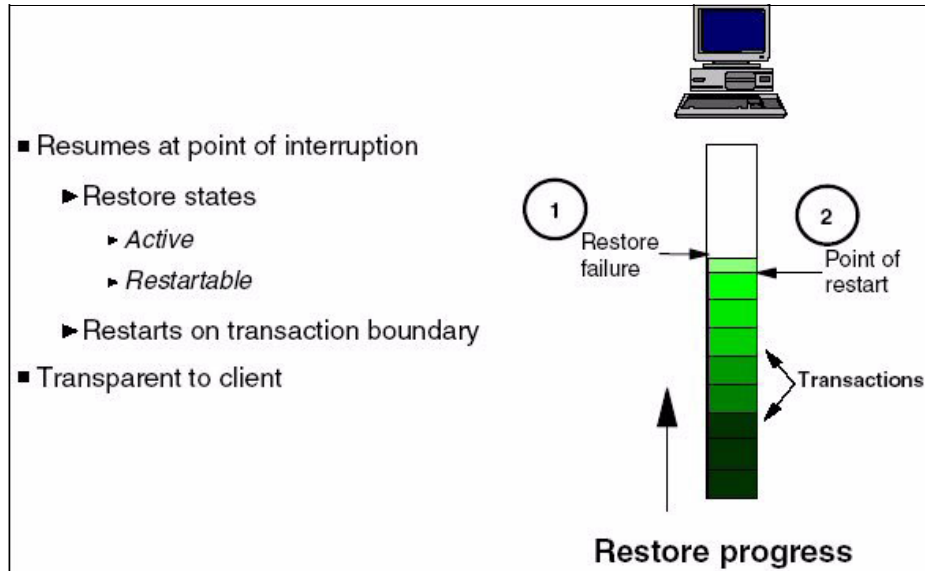
Figure 5-26   Restartable restore

5.9.2 Point-in-time restore

A point-in-time operation restores the specified objects to the state that existed at a specific date and time. A point-in-time restore is supported at the file space, directory, or file level. You must specify a sufficiently long retention period in the management class for this to occur.

To provide a point-in-time restore capability for, say, up to one month previously, set the verexists and verdeleted parameters to nolimit and retextra, and retonly to at least that number of days. With this method, it does not matter how many times the files change in the restorable period because you will always have enough versions stored to perform the restore.

Figure 5-27 on page 170 shows the point-in-time restore feature.
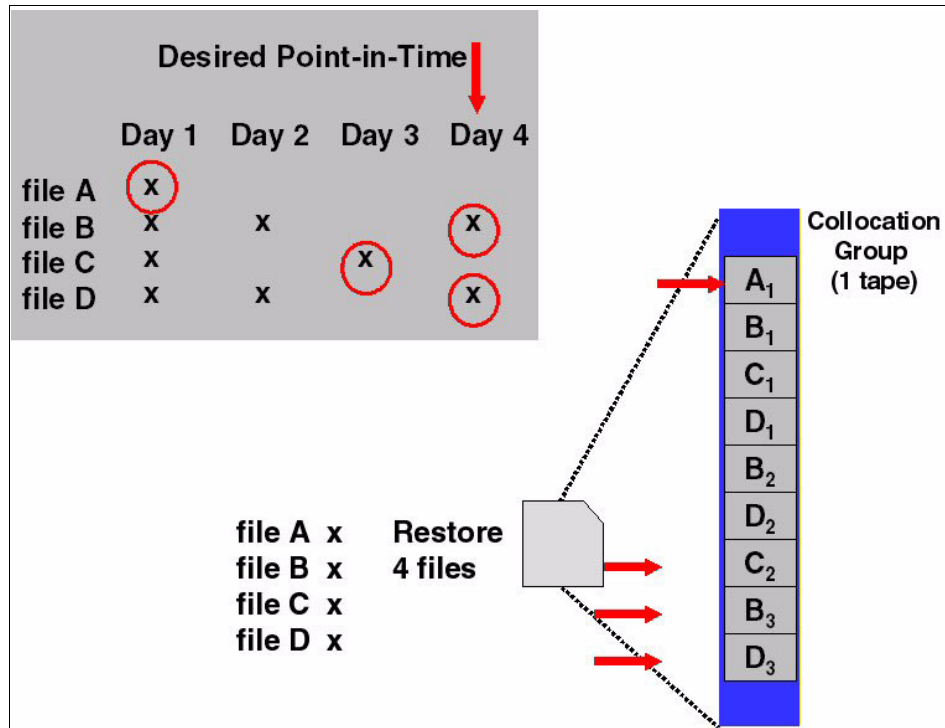
Figure 5-27   Point-in-time restore

Perform incremental backups to support a point-in-time restore. During an incremental backup, the client notifies the server when files are deleted from a client file space or directory. Selective and incremental-by-date backups do not notify the server about deleted files. Run incremental backups at a frequency consistent with possible restore requirements.

5.9.3 Multi-session restore

Multi-session restore allows the backup-archive clients to perform multiple restore sessions for no-query restore operations, increasing the speed of restores. This is similar to multiple backup session support.

Multi-session restore uses the mount points available on the server. If the data to be restored is on several tapes, there are sufficient mount points available, and the restore is done using the no-query restore protocol, then multiple sessions can be used to restore the data.

Note: Multisession restore is enabled by the *resourceutilization* parameter.

5.9.4 Backup set restore

Restoration of a backup set can be performed on a complete filespace level or by selecting individual files. You can restore from a backup set using either:

► Server-based backup set restore from Tivoli Storage Manager server

► Local backup set restore directly through the Tivoli Storage Manager client from locally attached devices without contacting the server

To restore from the backup set, the node name of the client must match the one that was defined in the backup set. If using LAN-free restore, the backup volumes are mounted by the backup-archive client through normal operating system device drivers and file system media such as CD-ROM, Jaz, Zip, and disk.

5.10 Retrieve

The `retrieve` command obtains copies of archived files from the Tivoli Storage Manager server. You can specify either selected files or entire directories to retrieve archived files. The description option enables you to search for the descriptions assigned to the files when they were archived; you may decide to replace the files into the same directory from which they were archived, or into a different directory.

Again, archive packages are groups of files archived together with a common description. You can search for a specific file within a package, or even retrieve the whole package. When you retrieve one or multiple files, Tivoli Storage Manager locates the volume where both directories and files are, so you do not need to know which tape holds the data.

5.11 Scheduling

In a typical production environment, the backup and other operations that protect the client data should be scheduled, so that we can be sure they regularly execute and can see if or when something goes wrong. Tivoli Storage Manager provides you with a client scheduling interface, which interacts with the server's Central Scheduler for this purpose.

If you use Tivoli Storage Manager's own central scheduling, the administrator defines appropriate schedules on the server to perform the Tivoli Storage Manager tasks automatically. Central scheduling is a cooperative effort between the server and each client node in that the client must run its own scheduler process so that the client and server can contact each other to run the scheduled operation correctly.

Scheduled operations are recorded centrally in the event log at the server. You can view which schedules ran successfully, which were missed, and which are scheduled to run in the future. You can create an exception reporting list to view only those schedules that failed.

The client scheduling process normally should be configured to start automatically each time the client boots to avoid missing schedule execution and compromising data security. There are two methods used to control how the client and server make contact to run a schedule: Client polling and server prompted. The client polling and server prompted methods can be configured in the client node configuration files or defining client optionsets on the Tivoli Storage Manager server. The client node schedules can run using one to two methods. First, the classic schedule will run the action type on the specified day, Sunday - Saturday, weekday or weekend. The enhanced schedule will run the action type on a specific day (1 - 31) of a specific month or a specific week of a month.

The classic and enhanced scheduling criteria can also be used for administrative schedules.

5.12 Compression

You have the option to specify that each client should compress its files or other objects before sending them to the Tivoli Storage Manager server. Compression is available for both backup and archive operations. Enabling client compression will decrease the network traffic between client and server (because you are sending a smaller quantity of data) at the expense of requiring more client processing resources to perform the operation. Therefore, the decision to enable client compressions must be made individually for each configuration.

If you are using client compression, then the client also automatically decompresses any objects that are sent back to it from the server when the reverse restore or retrieve operation is requested. Objects that are compressed also ultimately take up less storage space in the Tivoli Storage Manager server disk storage pools, reducing resource requirements. If compression is enabled on tape devices, the size of the compressed data could grow larger than the size of the compressed data.

5.13 Encryption

The Tivoli Storage Manager client implements an encryption function, which encrypts data before it is sent to the Tivoli Storage Manager server. This helps secure backed up data during transmission, and it means that the data stored on

the Tivoli Storage Manager server cannot be read by any malicious administrators.

The function uses either AES 128-bit or DES 56-bit encryption. The user can choose which files are subject to encryption via include/exclude processing. The encryption uses a very simple key management system, which means that the user either must remember the encryption key password during restore or store it locally on the client system. The encryption processing is the last task on the client system before the data is sent to the server; other client operations such as compression happen before encryption is done. Encryption works for both backup and archive operations.

Current tape drive technology has encyption capabilities with the IBM LTO4 generation and TS1120 tape drives. The tape drives are supported to work with Tivoli Storage Manager. The encryption strengths for tape drives are AES128, AES192 or AES256. IBM offers Encryption Key Manager (EKM) to manage the encryption keys.

Unlike the Tivoli Storage Manager user password, the encryption key password is case-sensitive. If the password is lost or forgotten, the encrypted data cannot be decrypted, which means that the data is lost.

5.14 Windows issues

This section discusses Windows-specific issues as they relate to Tivoli Storage Manager.

5.14.1 Windows System Services

The Windows System Services contains the services required for the system to operate. The services can be backed up individually however it is recommended to backup all the services to ensure that the server can be restored to a consistent state. The components that are included are:

▶ Background Intelligent Transfer
▶ Event log (not applicable for Windows Vista)
▶ Removable Storage Management
▶ Cluster Database
▶ Remote Storage
▶ Terminal Server Licensing
▶ Windows Management Instrumentation
▶ Internet Information Services
▶ DHCP Database
▶ Wins Database

5.14.2 Windows System State

A System State is a collection of files or databases that represent a logical entity and help the system achieve a consistent state. System objects can be part of a larger, distributed entity known as System State.

You can back up Windows 2003 and Windows Vista system objects together or individually. Microsoft recommends that all system objects be backed up together to maintain a consistent system state. These are valid system objects:

- ▶ Active Directory® (not applicable for Windows Vista)
- ▶ System Volume
- ▶ Certificate server database
- ▶ COM+ database
- ▶ Windows Registry
- ▶ System and boot files
- ▶ ASR (Windows XP / 2003 only)

To back up and restore all Windows system components, except the Windows Registry, a user must have administrator privileges. A member of the backup operators group can back up and restore the Windows Registry.

5.14.3 Automated System Recovery

Automated System Recovery (ASR) is a restore feature of Windows XP Professional and Windows Server 2003 that provides a framework for saving and recovering the Windows XP or Windows Server 2003 operating state in the event of a catastrophic system or hardware failure. Tivoli Storage Manager creates the files required for ASR recovery and stores them on the Tivoli Storage Manager server.

The goal of ASR, as stated by Microsoft, is to return the operating system to the point of the last backup. ASR is not used to recover application or user data. Such data is recovered by normal Tivoli Storage Manager restore procedures after successful completion of ASR recovery. ASR is a two-phase process:

1. Windows installs a temporary operating system image using the original operating system media.

2. Windows invokes Tivoli Storage Manager to restore the system volume and system state information.

5.15 SAN environment

This section describes how to perform LAN-free client operations such as backups and restores. This includes instructions for configuring the Tivoli

Storage Manager server and client to use both disk and tape as LAN-free storage destinations. It also presents important considerations for using LAN-free data transfers, either to disk or tape.

Tivoli Storage Manager provides a number of SAN solutions. Many other solutions are in development and will be released as SANs mature. These solutions address the need for efficient and reliable data protection:

► Tape library sharing

   The tape resource sharing feature of Tivoli Storage Manager allows administrators to pool tape resources for use by many Storage Manager servers running on heterogeneous platforms. This can improve backup and recovery performance and tape hardware asset utilization.

► LAN-free client data transfer

   At the direction of the Tivoli Storage Manager server, tape storage pools are dynamically allocated to Tivoli Storage Manager clients. This enables backup information to be sent across the SAN directly from the client to the server storage pools. The data path completely bypasses the LAN and the Tivoli Storage Manager server. LAN-free data transfer can significantly improve the data transfer speed. If your LAN environment is congested during the backup window, this is an excellent alternative.

5.15.1 Tape library sharing

SAN technology allows Tivoli Storage Manager to share its libraries. Multiple Tivoli Storage Manager servers can dynamically share the volume and tape drive resources of one connected tape library. The hosts can maintain high-speed connections to the same devices through the SAN fabric. Backup and restore applications benefit immediately from this, and the effect is pronounced for environments with large amounts of data to back up over shrinking windows of time and constrained LAN bandwidth.

5.15.2 LAN-free client data transfer on sequential media

In this section, we discuss LAN-free client data transfer solutions using Tivoli Storage Manager and storage agent. In particular, we discuss these topics:

► Setup and configuration of Tivoli Storage Manager server for LAN-free client data transfer

► Setup and configuration of storage agent for LAN-free client data transfer

► Operations and considerations when using client data transfer

Server setup for LAN-free client data transfer

Before setting up LAN-free data transfer, you should have a Tivoli Storage Manager 5.4 server configured and running, since it must be able to communicate with any tape device through the SAN. The storage agent software version must be the same as the Tivoli Storage Manager server or N-2. If the Tivoli Storage Manager server was at 5.4.2.0, the storage agent needs to be at 5.4.0.0 or higher but not above 5.4.2.0. You can have any server platform running.

In a Tivoli Storage Manager SAN environment, the Tivoli Storage Manager server is referred to as the library manager. The library manager processes all mount requests from all library clients (storage agents).

The library must be defined as a *shared* library or the Storage Agent will not be able to communicate with it.

You also need to have labeled volumes checked in your library, a defined policy that uses the storage pool you defined as a backup or archive destination, and a client node registered to the appropriate policy domain.

Installing and configuring the storage agent

To use the storage agent, you must have a Tivoli Storage Manager 5.4 server properly configured for your SAN attached library.

A copy group definition in the appropriate management class, which is what you use for LAN-free backups, should point to the tape storage device connected in the SAN. You can use disks for LAN-free backup. You must also install the Tivoli Storage Manager backup-archive client and API on the server on which you wish to perform LAN-free data transfer.

Configure the storage agent

The storage agent uses the dsmserv.opt options file when the application starts. Ensure that the DEVCONFIG option is set. The following statement shows how the devconfig.out file is used to store device configuration information:

```
DEVCONFIG devconfig.out
```

To establish communication between the storage agent and the Tivoli Storage Manager server, issue the DSMSTA SETSTORAGESERVER command. The VALIDATE LANFREE command will check that the LAN-free configuration is configured properly.

Verifying LAN-free data movement

When the Tivoli Storage Manager server, client, and storage agent are completely configured, you can perform LAN-free backups, restore, archives, or

retrieves. You have several options for verifying that the client transfer was LAN-free:

► Issue the query session command from the storage agent console. This should show you that the bytes received are increasing to the total amount of data being backed up. A query session from the Tivoli Storage Manager server will show that there is a little amount of bytes received for the same node because this refers to the metadata.

► Issue the command below to verify that the storage agent has initiated communication with the Tivoli Storage Manager server. If you get an output from the commands in the example, then the client session established is LAN-free.

```
query actlog search=storage_agent_name msgno=8337
```

► You can also use this command:

```
query actlog search=storage_agent_name
```

LAN-free data transfer considerations

Tivoli Storage Manager LAN-free data transfer effectively removes backup-archive data traffic from your LAN. It improves performance specifically when you are backing up large files.

We recommend using LAN-free functions to back up large databases and file system images because there is a significant performance improvement. This also allows you to fully utilize your LAN for other application and communication processing.

For smaller file backups, you can use LAN-free capabilities, but performance improvement may not be significant. There are times when LAN data transfer is faster than LAN-free data transfer for small files.

5.16 Reporting utility of Tivoli Storage Manager

The Tivoli Storage Manager environment includes a series of tasks that must be performed regularly. To perform these tasks efficiently requires a set of resources, such as space in storage pools or in the Tivoli Storage Manager database, and tape drives and volumes. Scheduled operations must complete in a timely manner and without failures. In large environments, the number of operations can be quite high, and managing them effectively can be quite complex. In order to manage this complex environment we need a reporting features that show us that tasks are complete.

The daily summary report or overview report is the most basic type of report. From this report you can verify that Tivoli Storage Manager is performing tasks

correctly and view a summary of all failed operations, such as missed backups or server errors.

The daily summary report should be as short as possible, to facilitate reading and understanding the information. It should provide the following information:

► Files that failed during backup
► Amount of data transferred to the Tivoli Storage Manager server
► Space usage trends on the Tivoli Storage Manager server
► Removable media (tape) errors
► Server error messages
► Information on the server

The Tivoli Storage Manager server has three main sources of information for creating reports:

► Database

   This is the most important source. The Tivoli Storage Manager database contains all server and most client definitions. It is the prime source of static information. When database information is requested and given, it is in the form of snapshot or point-in-time data; trends cannot be seen.

► Activity log

   The Tivoli Storage Manager activity log contains all server messages for the past several days. The activity log displays historical information; you can see the progress of operations over time.

   With Tivoli Storage Manager, client messages can be logged as events in the activity log. Commands are available for selecting all, none, or a subset of events to be logged.

► Accounting log

   Tivoli Storage Manager accounting can be activated with a server command. Once started, records are automatically collected and the log is written to a file called dsmaccnt.log in the server instance directory. The dsmaccnt.log file contains one record for each client session that terminates.

The Tivoli Storage Manager client has three main sources of information:

► Client error log:

   On the client node, Tivoli Storage Manager writes error information to the dsmerror.log file, usually for situations where the client did not succeed in contacting the server. By default it is written to the Tivoli Storage Manager client installation directory; however, it can be stored on any drive or directory accessible to the client (for example, a shared network drive) by setting a client option.

► Scheduler log:

- ► Also on the client node is the dsmsched.log file, which contains information for all scheduled operations, such as the name of files that are backed up or archived, failures and errors, and backup summary statistics.

- ► Web client log:

- ► All web client messages are written to the dsmwebcl.log file. By default it will be written to the same directory as the client error and scheduler logs.

5.16.1 Client and server event reporting

To take advantage of standard systems management interfaces, Tivoli Storage Manager can send client and server events to external interfaces. The supported interfaces are:

- ► Simple Network Management Protocol (SNMP)
- ► Managers such as NetView® for IBM AIX, CA Unicenter, or HP OpenView
- ► Tivoli Enterprise™ Console
- ► NetView for MVS
- ► Windows event log
- ► User-written exit
- ► Direct to a file

Interfaces that receive event data are called event receivers.

Each event message, whether client or server, can be enabled for any of the supported receivers. It is possible to enable one message or a group of messages for more than one receiver. As with client event logging, events are enabled for receivers by message number or severity.

5.16.2 SNMP heartbeat monitoring

SNMP can be used to monitor network elements from a central point. The monitored systems can send traps notifying the SNMP manager about events taking place on the local system. In addition, a heartbeat monitor can be established to determine whether managed Tivoli Storage Manager servers are still alive.

To enable Tivoli Storage Manager to take advantage of SNMP monitoring, it includes an interface for SNMP. The interface is distributed with the server in the form of an SNMP subagent. It is supported for the server running on IBM AIX, HP-UX, Solaris, and Windows. Communication between the server and the SNMP manager is established through one of these connection channels:

- ► Tivoli Storage Manager server <--> SNMP subagent <--> SNMP agent <--> SNMP manager
- ► Tivoli Storage Manager server <--> SNMP agent <--> SNMP manager

To enable communication between the SNMP subagent and the SNMP agent, the SNMP agent must support the Distributed Protocol Interface (DPI®).

5.16.3 SQL queries and ODBC interface

Tivoli Storage Manager provides an SQL interface that supports queries to its internal database. The interface is read-only and includes a SELECT command and an open database connectivity (ODBC) driver for Windows XP / 2003 / Vista.

SELECT command

The SQL interface represents Tivoli Storage Manager information in the form of relational tables with rows and columns that can be accessed by the SELECT command, which uses SQL syntax that complies with the SQL92/93 standard. This command can only be issued from the administrative command line interface.

Because SQL processing uses database resources, long-running or very complicated select statements can slow server performance significantly. Therefore, resource-intensive queries display a confirmation message, offering the possibility to abort the query before executing it.

ODBC driver

ODBC is a standard interface between SQL database engines and front-end applications. Using ODBC, products such as Lotus Approach, Lotus 123, Microsoft Access and Microsoft Excel® can be used to graphically construct SQL queries, which are then dispatched to the database (in this case the Tivoli Storage Manager database). The select statement results are returned in tabular form and can be processed to be displayed as charts or tables. The Tivoli Storage Manager ODBC driver only ships as part of the Win32® client package. The driver is compliant with the ODBC 2.5 application programming interface (API).

5.16.4 Operational reporting

The Tivoli Storage Manager operational reporting feature automates your daily monitoring tasks, and generates reports that can tell you if a server requires attention. You can run the operational reporting interface on Windows only, but it supports servers running on any supported operating system. See the readme file, included with the installation package, for a list of supported Tivoli Storage Manager servers.

Operational reporting views a Tivoli Storage Manager server as being in one of two states: Running Smoothly or Needs Attention. The state of a Tivoli Storage Manager server is automatically determined and a message can be sent to you

in an e-mail, posted to the Web, or displayed on the management console. The message provides access to a status report that can be customized.

If a server needs attention, the message describes any issues and provides recommendations for getting Tivoli Storage Manager running smoothly again.

Tivoli Storage Manager operational reporting is available as:

► Part of the Tivoli Storage Manager management console included when the Tivoli Storage Manager Windows server and storage agent packages are installed

► A stand-alone package as an independent component to be used if you do not run Windows Tivoli Storage Manager servers or storage agents, and do not anticipate doing so in the future

5.17 Clustering support

A key benefit of clustering is high availability. Resources on clustered servers act as highly available versions of unclustered resources. If a node (an individual computer) in the cluster is unavailable or too busy to respond to a request for a resource, the request is transparently passed to another node capable of processing it. Clients are therefore unaware of the exact locations of the resources they are using.

There are two clustering issues for Tivoli Storage Manager server and client clustering support:

► IBM high availability cluster multiprocessing (Tivoli Storage Manager HACMP)

HACMP detects system failures and manages Tivoli Storage Manager failover to a recovery processor with a minimal loss of end-user time. You can set up a Tivoli Storage Manager server on a system in a HACMP cluster so that, if the system fails, the Tivoli Storage Manager server will be brought back up on another system in the cluster. In both failover and fallback, it appears that the Tivoli Storage Manager server has crashed or halted and was then restarted.

► Microsoft Cluster Server (Tivoli Storage Manager MSCS)

Tivoli Storage Manager can be configured in a Tivoli Storage Manager MSCS high availability environment. The administrator uses the MSCS cluster administrator interface and Tivoli Storage Manager to designate cluster arrangements and define the Tivoli Storage Manager failover pattern. The systems are connected to the same disk subsystem, and they provide a high-availability solution that minimizes or eliminates many potential sources of downtime. Microsoft Cluster Server (MSCS) is software that helps

configure, monitor, and control applications and hardware components that are deployed on a Windows cluster.

5.17.1 HACMP and client configuration

The Tivoli Storage Manager backup-archive client itself (including the administrator, backup/archive, HSM, and API pieces) is supported for use in an HACMP cluster environment. This configuration enables scheduled client operations to continue processing in the event of a system failure on a redundant clustered failover server.

You can install the Tivoli Storage Manager client locally on each node of an HACMP environment. You can also install and configure the Tivoli Storage Manager Scheduler Service for each cluster node to manage all local disks and each cluster group containing physical disk resources.

If a scheduled incremental backup of a clustered volume is running on machine-a and a system failure causes a failover to machine-b, machine-b then reconnects to the server.

If a failover occurs during a user-initiated (that is, non-scheduled) client session, the Tivoli Storage Manager client starts on the node that is handling the takeover. This allows it to process scheduled events and provide Web client access.

5.17.2 MSCS and client configuration

The Tivoli Storage Manager client is supported in an MSCS Tivoli Storage Manager cluster environment. This configuration allows Tivoli Storage Manager scheduled client operations to continue processing in the event of a system failure on a redundant clustered failover server.

You can install the Tivoli Storage Manager client locally for both nodes. Then install and configure the scheduler for one of the nodes. Finally, fail the virtual node over and configure the scheduler for the second node.

6

Problem Determination, Adminstration and Performance Tuning

In this chapter, we describe the process of identifying problems while using IBM Tivoli Storage Manager. We also provide suggestions for solving them.

In this chapter, the following topics are discussed:

► What to know about the fourth and fifth section of the test
► Problem determination
► Tivoli Storage Manager error messages
► Hints and tips
► Performance tuning

6.1 What to know about the fourth and fifth section

Sections 4 and 5 of the IBM Tivoli Storage Manager 5.4 Certification test objectives emphasizes performance tuning, problem determination and administration. In this chapter, we discuss the processes involved with performance tuning, problem determination and administration. At the end of this chapter, you will be familiar with the following procedures:

► Verify that hardware is functioning correctly with other applications.
► Review OS logs.
► Review Tivoli Storage Manager logs.
► Ensure network connectivity.
► Ensure B/A client has connectivity/access to Tivoli Storage Manager server.
► Review change management process to view changes to the environment.
► Determine if a particular function is not working.
► Check maintenance levels.
► Determine which users are affected.
► Determine which file types are affected.
► Check for known problems.
► Identify the problem.
► Make changes.
► Test and monitor change effects.
► Document change.
► Identifying bottlenecks.
► Measuring application performance.

6.2 Problem determination

In the sections that follow, we discuss key topics relating to problem determination. This includes interpreting Tivoli Storage Manager error messages and reviewing logs.

6.2.1 Finding the problem

When you are presented with a problem, the first thing to determine is whether it is related to hardware or software. This includes the following tasks:

► Verify that the hardware is working properly with other applications.

► Review the operating system logs to determine if any OS-related problems occurred.

► Verify that Tivoli Storage Manager services are running.

If nothing seems to be wrong with the hardware and software, then you should take the following steps:

1. Review the Tivoli Storage Manager logs.

   – Activity log: Displays messages generated by the server

   – Event log: Displays the status of scheduled events

   – dsmerror.log: Contains information about any errors that occur during processing

   – dsmsched.log: Contains information about scheduled events

2. Verify that the backup-archive connects to the Tivoli Storage Manager server.

3. Determine if a particular function is not working.

4. Check the known problems.

6.2.2 Tivoli Storage Manager messages

Error messages can appear in the following places:

► Server console
► Administrative client
► Operator terminal
► Administrative GUI
► Backup-archive client
► Space-management client
► Application clients

Tivoli Storage Manager provides an activity log that can help the administrator track server activity and monitor the system. The activity log contains messages generated by the server, and is stored in the database. The server automatically deletes messages from the activity log after they have passed the specified retention period. Any messages sent to the server console are stored in the activity log.

The `query actlog` command displays messages generated by the server. Examples of the types of messages stored in the activity log include:

► When client sessions start or end
► When migration starts or ends
► When backed up files are expired from server storage
► Any output generated from background processes

Figure 6-1 shows the basic format of Tivoli Storage Manager messages.
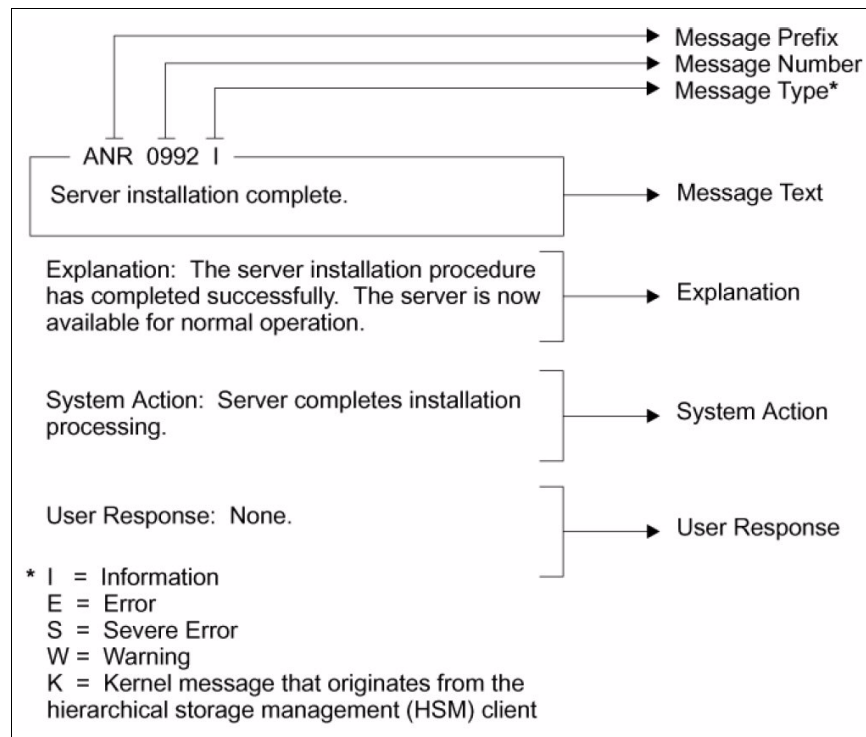
Figure 6-1   Tivoli Storage Manager Message Structure

The following formats are used for Tivoli Storage Manager messages:

► Messages that begin with the prefix ANE and range from 4000 to 4999 originate from the backup-archive client. These messages (or events) are sent to the server for distribution to various event-logging receivers.

► The client may send statistics to the server providing information about a backup or restore. These statistics are informational messages that may be enabled or disabled to the various event-logging receivers.

► Messages that begin with the prefix ANR originate from the server. These messages are related to common and platform-specific behavior.

► Messages that begin with the prefix ANS are from one of the following clients:

   – Administrative clients
   – Application program interface clients
   – Backup-archive clients
   – Space Manager (HSM) clients
   – Data Protection for Lotus Notes®

► Messages that begin with the prefix ACD are from Data Protection for Lotus Domino.

► Messages that begin with the prefix ACN are from Data Protection for Microsoft Exchange Server.

► Messages that begin with the prefix ACO are from Data Protection for Microsoft SQL Server.

► Messages that begin with the prefix ANU are from Data Protection for Oracle.

► Messages that begin with the prefix BKI are from Data Protection for R/3 for DB2 Universal Database and Data Protection for R/3 for Oracle.

► Messages that begin with the prefix EEO and are in range 0000-7999 are from Data Protection for Snapshot Devices for mySAP for Oracle.

► Messages that begin with the prefix EEO and range from 0000 to 9999 are from Data Protection for Disk Storage and SAN VC.

► Messages that begin with the prefix EEP and are in range 0000-7999 are from Data Protection for Snapshot Devices for mySAP for DB2 UDB.

► Messages that begin with the prefix IDS and are in range 0000-1999 are from Data Protection for Snapshot Devices for mySAP for Oracle.

► Messages that begin with the prefix IDS and are in range 0000-2999 are from Data Protection for Snapshot Devices for mySAP for DB2 UDB.

6.2.3 Problem determination steps

When a Tivoli Storage Manager problem occurs, these are the questions to ask:

► What is the problem?

► Where did it occur?

► When did it begin happening?

► What action was being performed?

► Were any messages issued?

   – Check the server activity log for error messages.

   – If error messages are in the server activity log, check 30 minutes before and after the time that the error message was issued. Often the problem encountered is actually a symptom of another problem and seeing the other error messages that were issued may help to isolate this.

   – Did the Explanation or User Response section of the Tivoli Storage Manager message offer any suggestions on how to resolve the problem?

► How frequently does this error occur?

▶ What do the system error logs say?

– On Windows: Check the application log.
– On IBM AIX and other UNIX platforms: Check the error report.

▶ Has anyone made changes in the environment that could affect Tivoli Storage Manager?

Administrators who may have made changes in a typical IT environment include:

– SAN Administrator
– Network Administrator
– Database Administrator
– Client or machine owners

5. What is in the Tivoli Storage Manager error logs?

The following Tivoli Storage Manager error logs should be reviewed:

– dsmserv.err

Server error file: The dsmserv.err file is located on the same machine as the server, typically in the server install directory. Note that the storage agent may also create a dsmserv.err file to report errors.

– dsmerror.log

Client error log: This is located on the same machine as the client.

– dsmsched.log

Client log for scheduled client operations: This is located on the same machine as the client.

– dsmwebcl.log
web client log: This is located on the same machine as the client.

– db2diag.log, db2alert.log, userexit.log

DB2 log files: These are useful when troubleshooting a problem during the backup process for a DB2 database using Tivoli Data Protection for DB2. These are located on the same machine where DB2 is installed. See the DB2 documentation for additional information about what is in these files and where they are located.

– tdpess.log

Default error log file: This file is used by the Data Protection for Enterprise Storage Server client.

– tdpexc.log

Default error log file: This is used by the Data Protection for Exchange client.

– dsierror.log

Default error log: This is used for the client API.

– tdpoerror.log

Default error log: This error log is for the Data Protection for Oracle client.

– tdpsql.log

Default error log: This is for the Data Protection for SQL client.

► Are the devices still accessible to the system and to Tivoli Storage Manager?

► Are there matching error messages or problem descriptions in the online Knowledge Base?

You should also test other operations to better determine the scope and impact of the problem. This may also help to determine if a specific sequence of events caused the problem.

The problem determination guide can also be found at this Web site:

`http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/topic/com.ibm.it smm.doc/0142/pdguide.pdf`

6.3 Hints and tips

This section provides suggestions for what to do should a problem occur.

6.3.1 Tivoli Storage Manager server tips

Some tips for addressing problems with the Tivoli Storage Manager server are:

► Examine any error messages that were issued.

Check for ANSnnnnx messages issued to the console, dsmsched.log, or dsmerror.log.

► Examine the server activity log messages for this session.

Check for the server activity log using QUERY ACTLOG for messages issued for this client session.

► For problems connecting to the server, check the following:

– Have the client communication options in the client option file been changed? If so, review these changes and try reverting to the previous values and retrying the connection.

- – Have the server communication settings been changed? If so, either update the client communication options to reflect the changed server values or else revert the server back to its original values.

- – Have any network settings been changed? For example, has the TCP/IP address for the client or server been changed? If there have been network changes, work with the network administrator to understand these changes and update the client, server, or both for these network changes.

► Check to see if client options changed.

  Changes to client options are not recognized by the client scheduler until the scheduler is stopped and restarted. Stop and restart the client scheduler.

► Check to see if the client is being run with the QUIET option.

  The QUIET processing option for the client suppresses messages.

► Verify INCLUDE/EXCLUDE syntax and ordering.

  The include and exclude processing can impact which files are sent to the server for a backup or archive operation.

► Verify that the correct Tivoli Storage Manager server is being used.

  In an environment with multiple Tivoli Storage Manager servers, a client may use different servers for different operations.

► Identify when and where the problem can occur.

  Client processing problems often only occur when performing specific operations, at certain times, or only on certain client machines.

► Check the server activity log.

  Check the server activity log for other messages 30 minutes before and after the time of the error.

► Check HELP for Tivoli Storage Manager messages issued for the problem.

  Check HELP for any messages issued by Tivoli Storage Manager.

► Try to reproduce the problem.

  If a problem can be easily or consistently reproduced, it may be possible to isolate the cause of the problem.

► Check to see if there is an error reading or writing to a device.

  If the problem is an error reading or writing data from a device, many systems and devices record information in a system error log. Examples of the system error log are *errpt* for IBM AIX, Event Log for Windows, and the System Log for IBM z/OS.

  If a device or volume being used by Tivoli Storage Manager is reporting some sort of error to the system error log, it is likely a device issue. The error

messages recorded in the system error log may provide enough information
to resolve the problem.

► Determine whether server options or settings been changed.

Changes to options in the server options file or configuration changes to the
server using the SET or UPDATE commands may cause operations that
previously succeeded to fail. Changes on the server to device classes,
storage pools, and policies may cause operations that previously succeeded
to fail.

► Determine if a scheduled client operation failed.

Scheduled client operations are influenced by the schedule definitions on the
server as well as the scheduling service (Tivoli Storage Manager Scheduler)
being run on the client machine itself.

► Determine if the server has run out of space.

The Tivoli Storage Manager server's primary function is to store data. If it runs
out of space in the database or storage pools, operations may fail.

To determine if the database is out of space, use this command:

`QUERY DATABASE`

More space should be defined if the "%Util" is at or near 100 percent.
Typically, if the database is running out of space, other server messages will
be issued indicating this. To add more space to the database, allocate one or
more new database volumes, and define them to the server using this
command:

`DEFINE DBVOL`

Then extend the available database space with this command:

`EXTEND DB`

6.3.2 Client password tips

A common error involving the client password is `Session rejected:`
`Authentication failure`. Generally, this error occurs when the password has
expired. However, it can also occur if either the server or the client has been
renamed.

If you receive this error during an interactive session, it is probable that the
password given is incorrect. The steps to take in this case are:

► Ask the Tivoli Storage Manager administrator to reset the node's password by
using the UPDATE NODE command.

► Issue a DSMC QUERY SESSION command, and when prompted, give the
new password.

If you are receiving this error during a non-interactive session such as Central Scheduling, ensure that the client option is PASSWORDACCESS GENERATE. This option causes the client to store the password locally. The password is encrypted and stored either in the registry for Windows clients or in a file named TSM.PWD for Macintosh, UNIX, and NetWare clients. Editing the registry or the TSM.PWD file is not recommended. Instead, follow these recommendations:

► Make sure that PASSWORDACCESS GENERATE is set in the option file.

► Issue a DSMC QUERY SESSION command. This command will force the locally stored password to be set.

   If this does not resolve the problem, update the node's password using the UPDATE NODE administrative command.

► Reissue the DSMC QUERY SESSION command, providing the new password.

   To see the password expiration setting for a particular node, issue the Tivoli Storage Manager QUERY NODE F=D admin command. Look for the Password Expiration Period field. If this field is blank, the default password expiration period of 90 days is in effect.

6.3.3 Client schedule

Use the following tips for performing client scheduling tasks:

► Checking the server activity log

   Narrow the query results down to the time frame surrounding the scheduled event. Begin the event log query at a time shortly before the start window of the scheduled event in question.

► Inspecting the client schedule log

   The Tivoli Storage Manager client keeps a detailed log of all scheduled activities. If queries of the server's activity log are not able to explain a failed scheduled event, the next place to check is the Tivoli Storage Manager client's local schedule log.

► Restarting the scheduler process on a remote machine

   When managing a large number of Tivoli Storage Manager clients running scheduler processes, it can be helpful to be able to start and stop the client service from a remote machine.

6.4 Performance tuning

This section provides suggestions on what to do to enhance performance.

6.4.1 Performance tuning tips:

- ► Determine the problem.

- ► Benchmark the performance of the application before making any changes.

- ► Establish the part of the application that requires performance to be tuned, in other words establish where the bottleneck is.

- ► Make changes to the application that would contribute to removing the bottleneck.

- ► Analyze the performance after changes have been made.

The performance guidelines for IBM Tivoli Storage Manager is located at the following link:

http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp

A
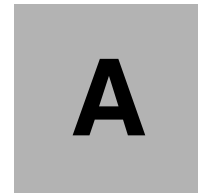
Lab exercises for Tivoli Storage Manager 5.4

A product installation guide and lab exercises for Tivoli Storage Manager Version 5.4 are included in this appendix. We recommend trying these exercises after you have reviewed each section.

Introduction to IBM Tivoli Storage Manager hands-on lab

Tivoli Storage Management protects an organization's data from hardware failures and other errors by storing backup and archive copies of data on offline storage. TSM has six basic capabilities:

► Backup/restore
► Archive/retrieve
► Space management
► Disaster recovery planning
► Bare machine recovery
► Online protection of database and applications

TSM centralizes the backup of an enterprise's computers. TSM runs on heterogeneous platforms. It utilizes a client/server architecture. The clients are all the machines whose data needs to be protected. The clients might be your organization's laptops, desktops, file servers, database servers, or development machines. The TSM server is usually a dedicated machine that can run on one of the eight different platforms we support. It is where the clients send their data to. The TSM server consists of three parts: The database and recovery log, which stores the meta information about the data being backed up; and the storage pools, which actually store the data. The storage pools can consist of any combination of disk, tape, and optical devices and are linked together in a hierarchy.

In the previous version (Version 5.3), the Administration Center was introduced. IBM Tivoli Storage Manager's new user interface consists of a number of different components that will assist the administrator in managing multiple IBM Tivoli Storage Manager servers within a single, integrated console.

The Integrated Solutions Console (ISC) builds on top of the WebSphere Application Server and WebSphere Portal base and includes lightweight versions of both in the ISC runtimes. It looks for common problems, actions, and subtasks across the range of ISC components. ISC on a lightweight portal infrastructure provides the ability to aggregate independent tasks and information into a single organized presentation.

Objectives

In this section, you will:

► Perform basic TSM backup and restore operations.
► Perform basic TSM archive and retrieve operations.
► Display and create TSM storage pools.
► Display and create TSM policies.
► Demonstrate the application of TSM policies on client backups.
► Display and expand the TSM server database.

► Perform basic Administration Center and Integrated Solutions operation.

This book does not cover installation of IBM Tivoli Storage Manager and its components. For detailed information about installing this product check with the following documents.

► *IBM Tivoli Storage Manager for AIX Installation Guide Version 5.4*, SC32-0134-00

► *IBM Tivoli Storage Manager for Windows Installation Guide Version 5.4*, SC32-0138-00

► *IBM Tivoli Storage Manager for Windows Backup-Archive Clients Installation and User's Guide V5.4,* SC32-0146-00

► *IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide V5.4, SC32-0145-00*

Lab 1: Basic file backup and restore

We are going to create data, back it up to TSM, change some files, back the data up again, delete it, and then restore it from TSM. You will be able to see the progressive backup in action.

Our goal is to back up the C:\CB directory and its files with IBM Tivoli Storage Manager. In this lab we have a directory called CB, and there are some files in it so you can create your own sample directory.

1. Double-click the TSM Client icon to start the TSM backup client.
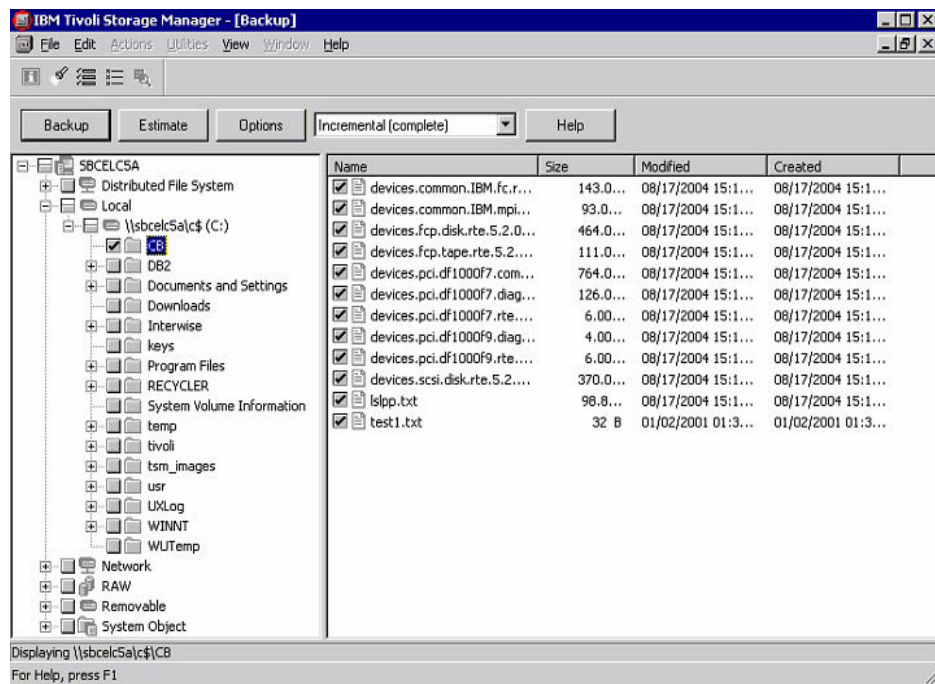
Figure A-1   IBM Tivoli Storage Manager Backup/Archive GUI

2. Click **Backup**.

3. Click the plus icon (**+**) next to Local and navigate down to C:\CB.

4. Click the first check box next to the CB directory (a tick should appear in it), and also click the actual folder icon next to the check box. Note that all the files appear as selected on the right-hand side. You could choose to individually exclude some files from backup by deselecting them individually.
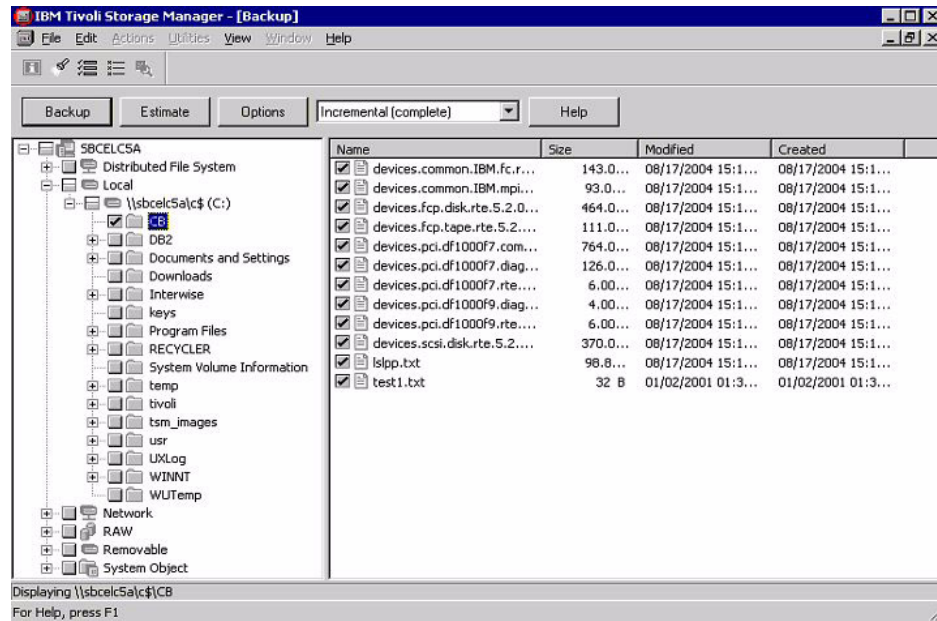
Figure A-2   Choosing a directory for backup

5. Click **Backup**.

6. When the backup finishes, see how many files were backed up and how much data was transferred. Because this is a new backup, all the files and the directory were backed up.
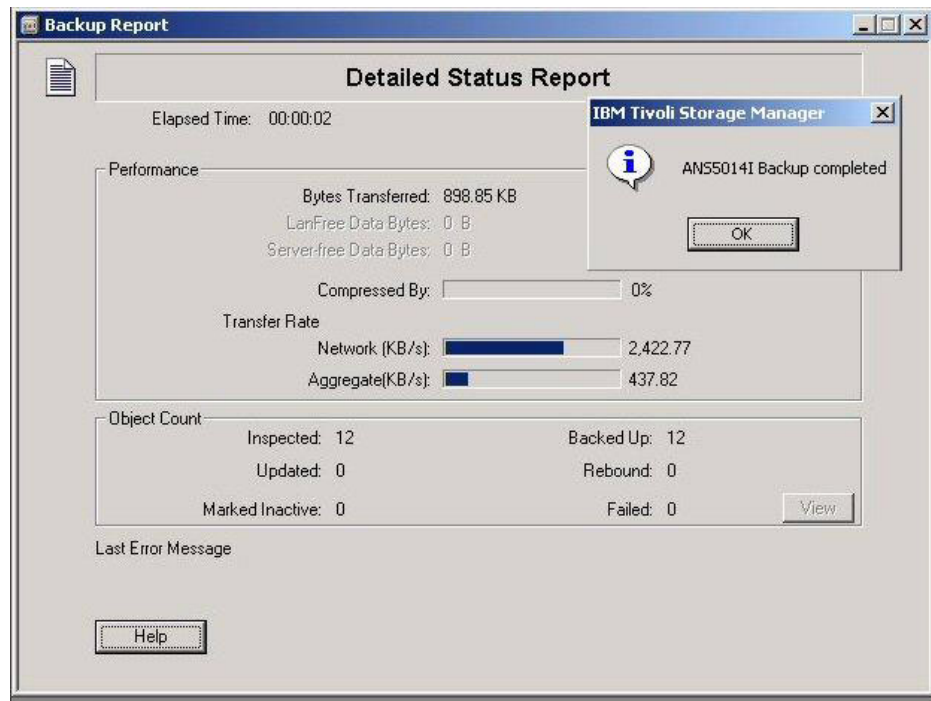
Figure A-3   Detailed status report for backup

7. Click **OK** to close the Backup completed window, then close the Detailed Status Report by clicking the **X** in the top right-hand corner. You will return to the main TSM client screen.

8. Now edit test1.txt in the CB directory. (Add your name to the bottom line.)

9. Back up the directory again using the same steps as before.

10. When the backup finishes, note how many files are backed up, and how much data was transferred; that is, only one file is backed up—the one that changed.
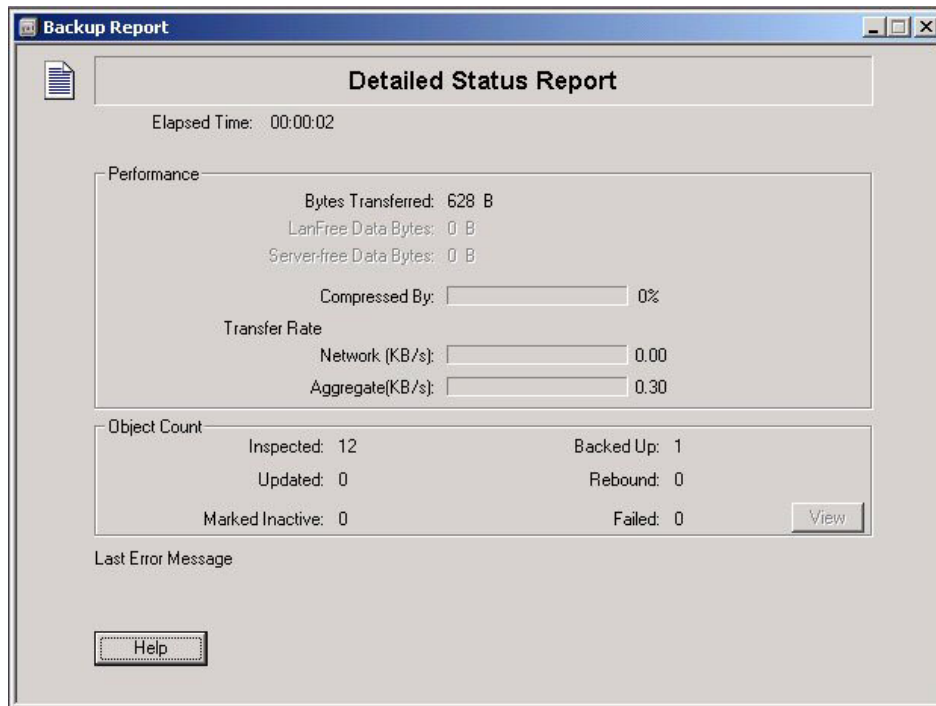
Figure A-4   Detailed report for incremental backup

11. TSM's progressive backup saves a lot of time and money by not re-backing up files unnecessarily, and by not having dependencies on a full backup.

12. Now delete some or all of the BFF extension files in the CB directory.

13. Re-run the backup of the whole directory. This time the backup report shows some files marked inactive. TSM recognizes deleted files and treats them differently, according to the policies you set up.
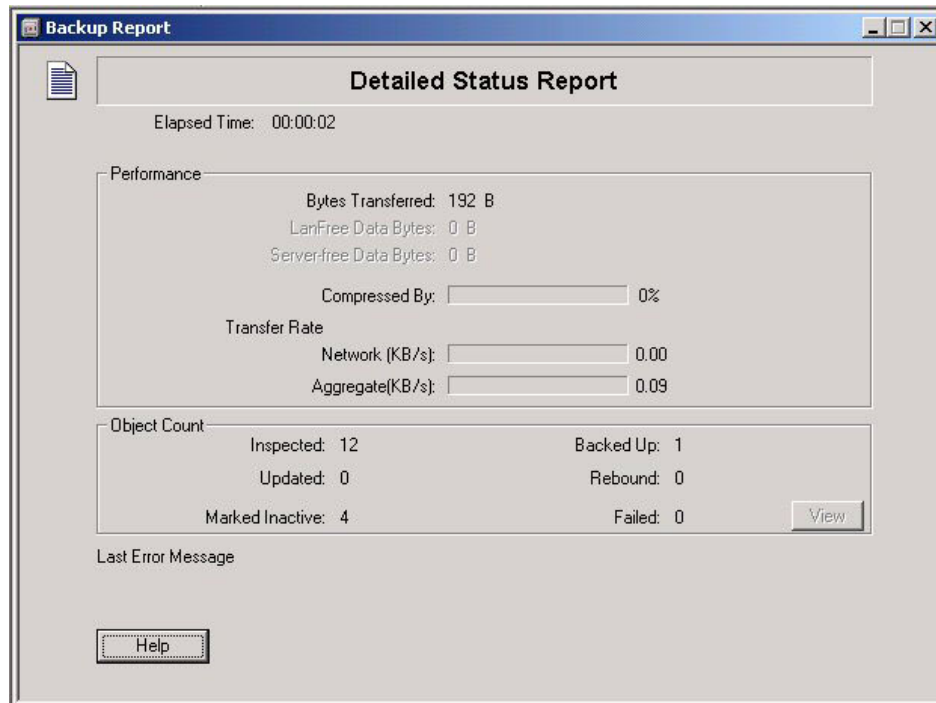
Figure A-5   Detailed status for Inactive files

14. Now we will restore those files. From the main panel, click **Restore**.

15. Drill down on File Level to the CB directory. Check the box and display the files. Note that only the files still existing on the real directory are displayed. What happened to the files we deleted?
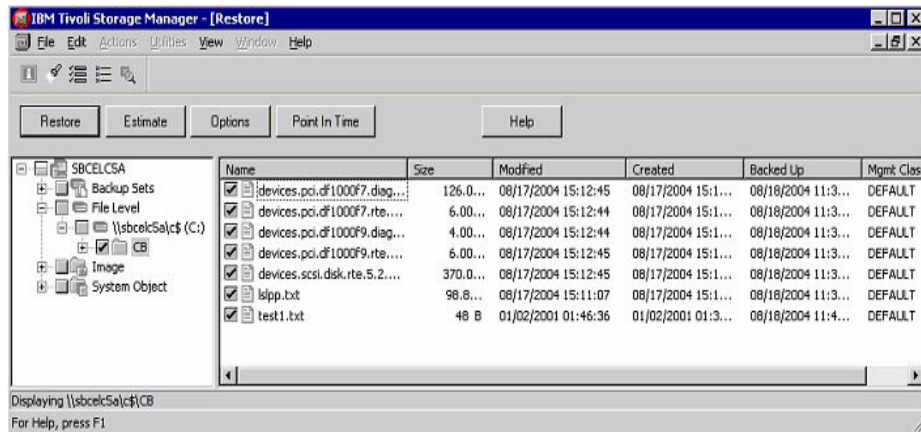
Figure A-6 Files selection for backup

16. The answer is that by, default, TSM only displays active files—these are the most recently backed up versions of files that still exist on the client. Extra backed up versions (like of the file that we modified and re-backed up) and deleted files on the client are not displayed in the default view.

17. To show all the files TSM has backed up, select **View** → **Display active/inactive files**. You have to re-drill down to the CB directory.
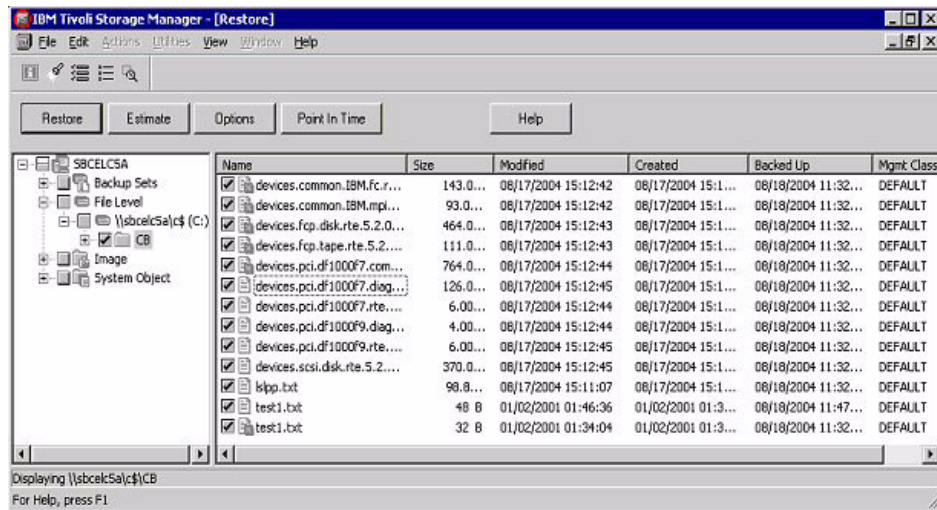
Figure A-7   Inactive files

18. Note the files with an x next to them. These correspond to the files you deleted (the BFF files), as well as the extra copy of the test1.txt file that you modified. You can see the different time stamps in the column Backed up.

19. Unselect the .txt file version that is not marked with an x, so that you will restore the previous version of it.

20. Click the **Options** box at the top.

21. Note the "Action for files that already exist" field. You can specify whether to overwrite automatically or to prompt when restoring files that already exist on the client. Select the **Replace** entry and click **OK** to exit.
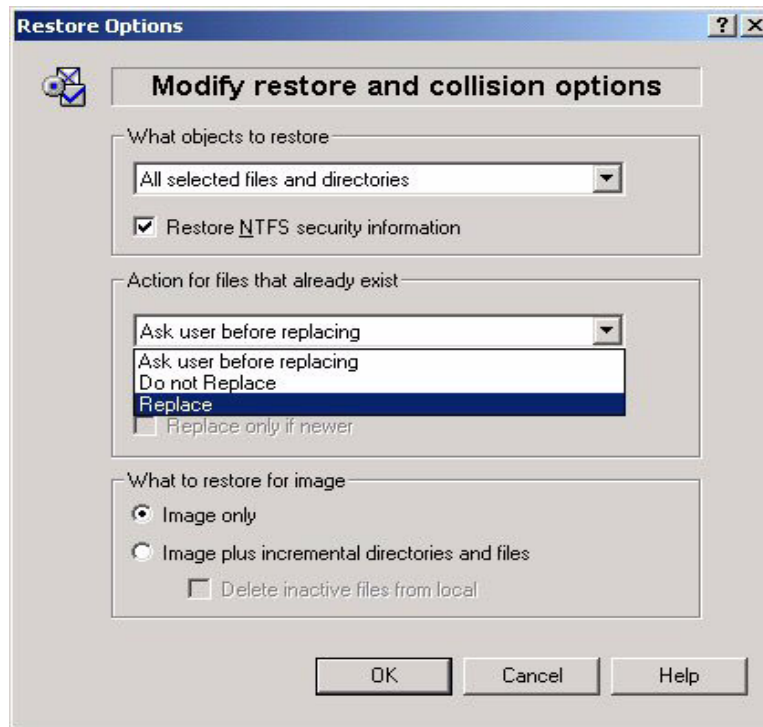
Figure A-8   Restore options

22. You can restore a file to the original or to a new location, and you can also decide what to do if the file already exists.

23. Accept the defaults and click **OK**.

24. Accept the default to restore to the original location. This is where you could specify to send the files to an alternative directory if required.
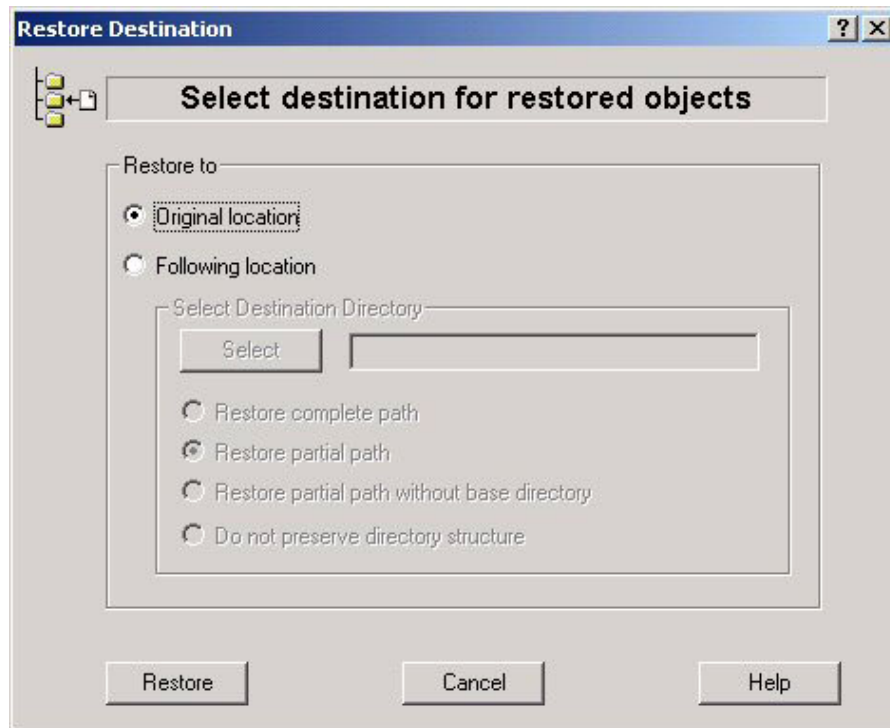
Figure A-9   Restore options

25. Click the **Restore** button to restore the files.

26. Close the Status Report.

27. Check that the files are back in the CB directory and see that the original version of the test1.txt file has been restored (that is, without your modifications).

28. Now let us see what happens with versions. The default policy of TSM is set to retain two versions of files that still exist on the client. (We will see this in more detail later.) We already have backed up two distinct versions of the test1.txt file (shown in steps 15 and 16). What will happen when we update the file again?

29. Again add your name to the last row in test1.txt. Re-back up the CB directory. Again, only the changed objects are backed up.

30. Check what has happened with the versions of test1.txt. Exit back to the main screen and click **Restore**. Select **View → Display active/inactive files**. Drill

down to the CB directory. Note there are still only two versions of test1.txt-one active and one inactive (because it is an extra version). You can see the backup dates by scrolling to the right. The oldest version has been expired from TSM, because our policy says to keep two versions. We can change this, of course, to a higher number.

Name	Size	Modified	Created	Backed Up
devices.common.IBM.fc.r...	143.0...	08/17/2004 15:1...	08/17/2004 15:1...	08/18/2004 15:24:52
devices.common.IBM.fc.r...	143.0...	08/17/2004 15:1...	08/17/2004 15:1...	08/18/2004 11:32:18
devices.common.IBM.mpi...	93.0...	08/17/2004 15:1...	08/17/2004 15:1...	08/18/2004 15:24:52
devices.common.IBM.mpi...	93.0...	08/17/2004 15:1...	08/17/2004 15:1...	08/18/2004 11:32:18
devices.fcp.disk.rte.5.2.0...	464.0...	08/17/2004 15:1...	08/17/2004 15:1...	08/18/2004 15:24:52
devices.fcp.disk.rte.5.2.0...	464.0...	08/17/2004 15:1...	08/17/2004 15:1...	08/18/2004 11:32:18
devices.fcp.tape.rte.5.2....	111.0...	08/17/2004 15:1...	08/17/2004 15:1...	08/18/2004 15:24:52
devices.fcp.tape.rte.5.2....	111.0...	08/17/2004 15:1...	08/17/2004 15:1...	08/18/2004 11:32:18
devices.pci.df1000f7.com...	764.0...	08/17/2004 15:1...	08/17/2004 15:1...	08/18/2004 15:24:52
devices.pci.df1000f7.com...	764.0...	08/17/2004 15:1...	08/17/2004 15:1...	08/18/2004 11:32:18
devices.pci.df1000f7.diag...	126.0...	08/17/2004 15:1...	08/17/2004 15:1...	08/18/2004 11:32:18
devices.pci.df1000f7.rte....	6.00...	08/17/2004 15:1...	08/17/2004 15:1...	08/18/2004 11:32:18
devices.pci.df1000f9.diag...	4.00...	08/17/2004 15:1...	08/17/2004 15:1...	08/18/2004 11:32:18
devices.pci.df1000f9.rte....	6.00...	08/17/2004 15:1...	08/17/2004 15:1...	08/18/2004 11:32:18
devices.scsi.disk.rte.5.2....	370.0...	08/17/2004 15:1...	08/17/2004 15:1...	08/18/2004 11:32:18
lslpp.txt	98.8...	08/17/2004 15:1...	08/17/2004 15:1...	08/18/2004 11:32:18
test1.txt	80 B	01/02/2001 07:3...	01/02/2001 01:3...	08/18/2004 17:26:41
test1.txt	32 B	01/02/2001 01:3...	01/02/2001 01:3...	08/18/2004 15:24:52

Figure A-10   Active/inactive files

You have completed Lab 1 on basic backup and restore. In the next lab, we see the basic steps of archive and retrieve functions.

Lab 2: Basic file archive and retrieve

We are going to work with the same directory, and we will use those files to archive to TSM and add more files to the existing archive, then finally retrieve those files from TSM.

1. If not already started, double-click the TSM Client icon to start the TSM backup client.

Figure A-11   Backup/Archive client GUI

2. Click **Archive**.

3. Click the plus sign (**+**) next to Local and navigate down to C:\CB.

4. Click the first check box next to the CB directory (a tick should appear in it), and also click the actual folder icon next to the check box. Note all of the files appear as selected on the right-hand side. You could choose to individually exclude some files from the archive by deselecting them individually.
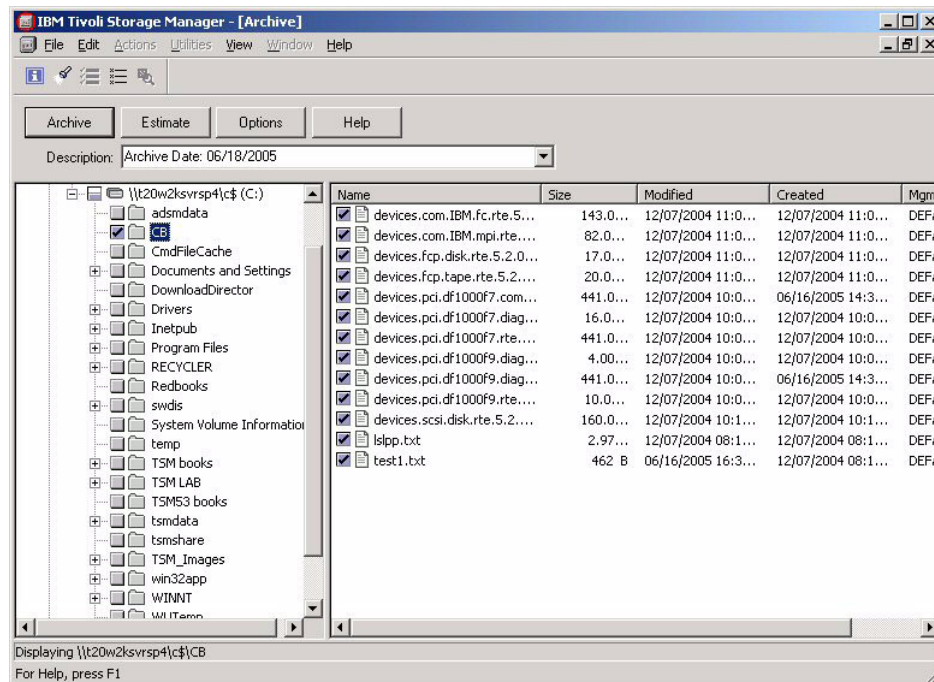
Figure A-12   Files selection for archive

5. Note the Description field; this is used to associate a meaningful description with the archive. TSM auto-generates a description based on the date, but you can edit/change this if you wish.

6. You may also check what the options are (click the **Options** button at the top); for example, you may delete the files after the archive completes.

7. For now, do not select any of these options.

8. Click **OK** to perform the archive. After the archive completes, the Detailed Status Report shows how many files were archived and how much data was transferred. Click **OK** on the Archive completed window, then close the Detailed Status Report by clicking the **X** in the right top corner. You will return to the main TSM Client GUI.
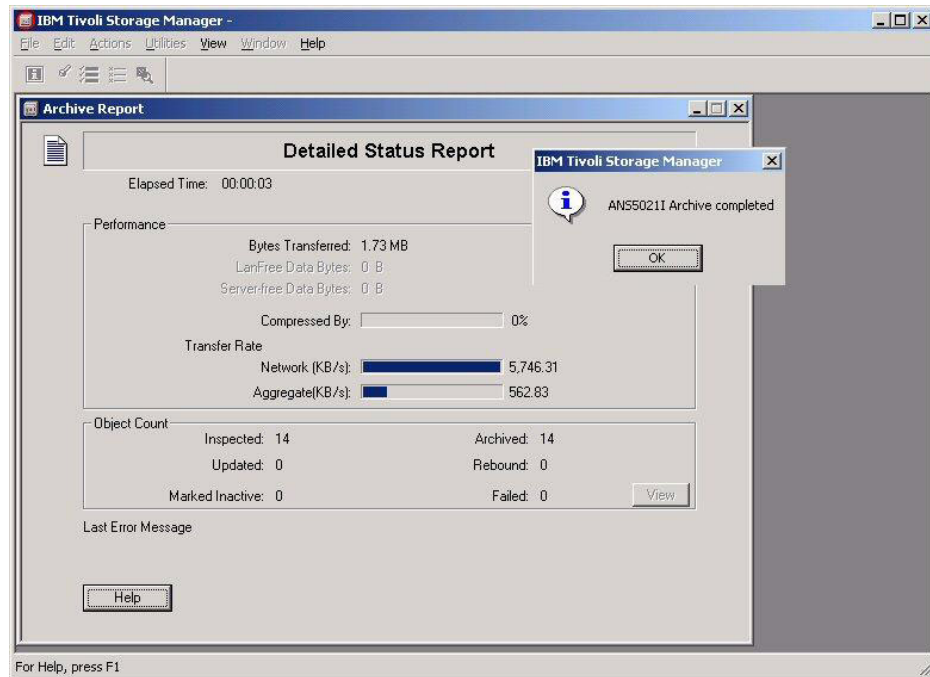
Figure A-13   Archive complete

9. Now we can add more files to an existing archive. Click the plus sign (**+**) next to Local and navigate down to C:\Redbooks.

> **Note:** If two different archives are done in the same day, using the default description, when retrieving you will find these files under the same day stamped archive. If you make an archive on a different day, in the retrieve menu you will find different dated archives.

10.Click the folder icon next to the Redbooks directory and select the files that start with IBM*.
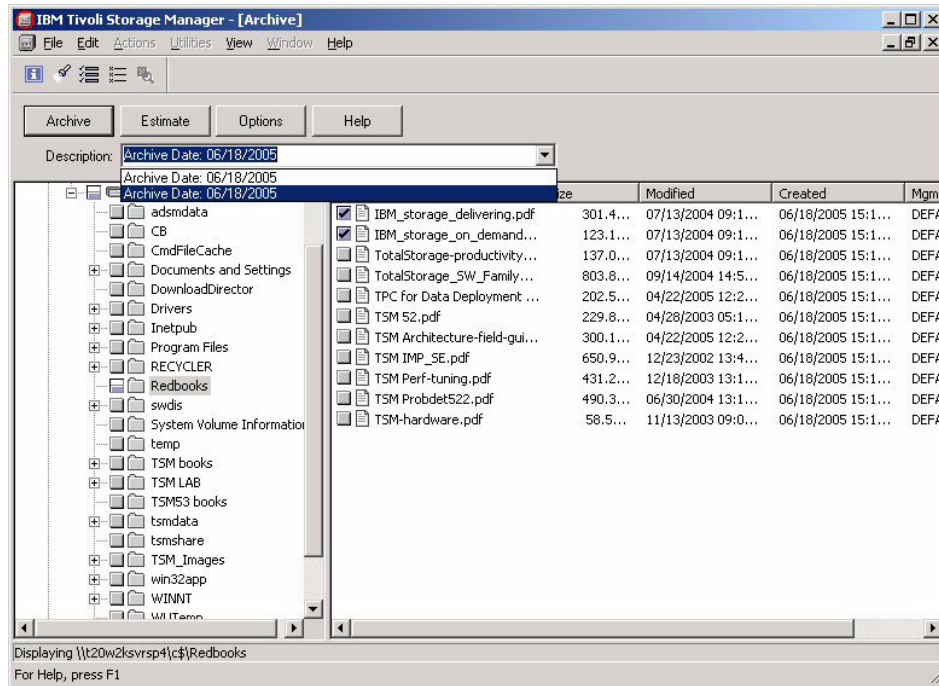
Figure A-14   Adding files to existing archive

11.Click **Archive**, and check again on the Detailed Status Report window how many files have been archived. Close the Detailed Status Report by clicking the **X** in the right top corner. You will return to the main TSM Client GUI.

12. In order to retrieve the files you have archived, click **Retrieve** from the GUI, and select the archived date. Then drill down to the Redbooks directory, click the gray box, and choose all the files. TSM allows you to retrieve all the files in an archive, or just some of them.
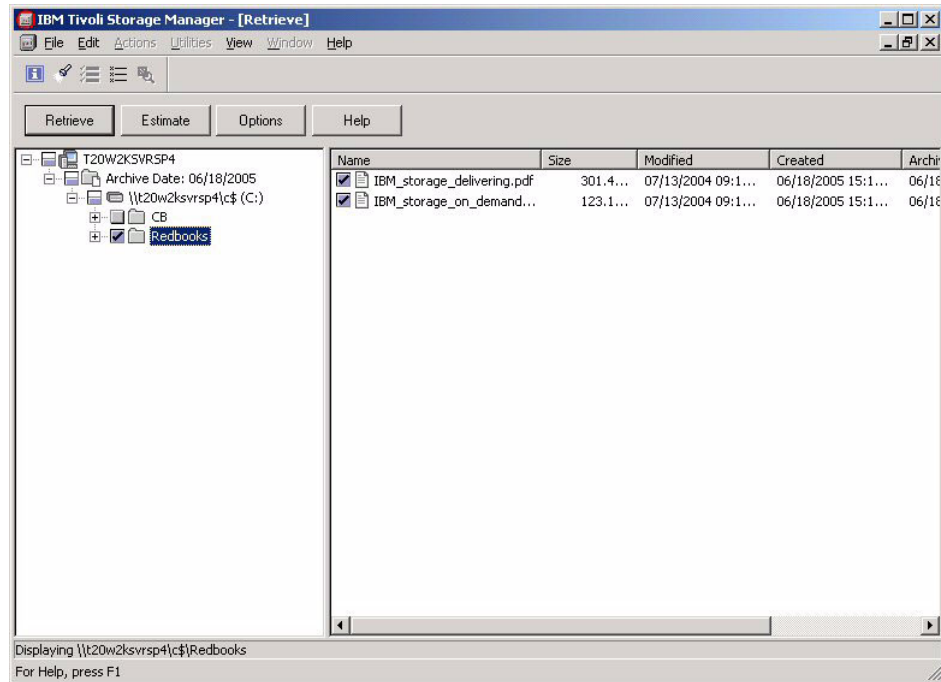
Figure A-15   Retrieve archives

13. Click **Options** to see the options if the files already exist. Click **OK** and click **Retrieve**.

14. A window will pop up and ask you for the destination for the retrieved objects. Choose **Following location** to C:\CB, then click **Retrieve**.
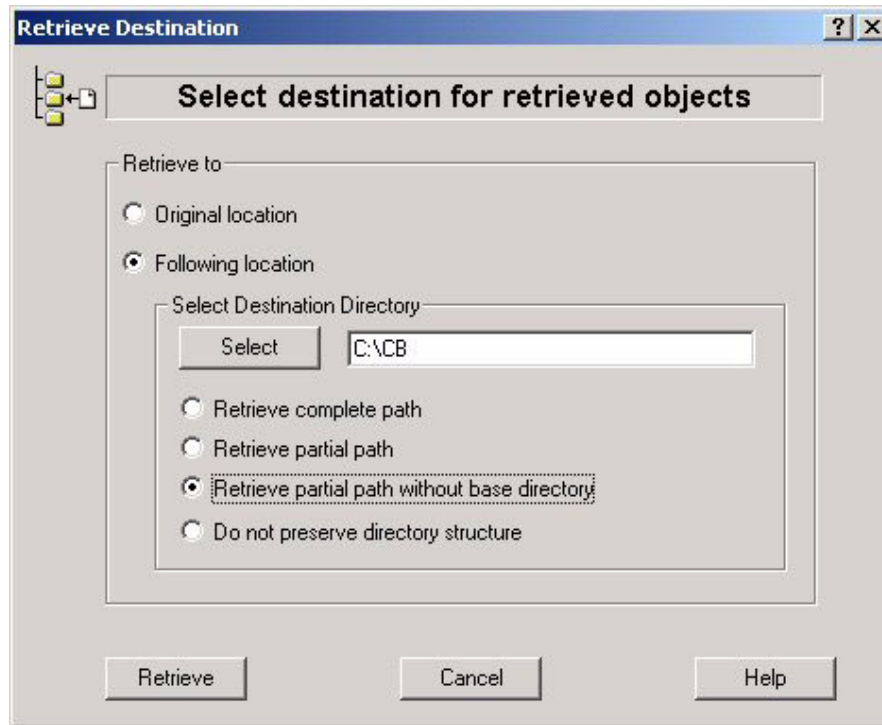
Figure A-16   Retrieve options

15. Click **OK** on the Retrieve completed window, then close the Detailed Status Report by clicking the **X** in the right top corner. You will return to the main TSM Client GUI.

16.After the retrieve completes, check the CB directory and see that the two files selected were correctly retrieved to this location.
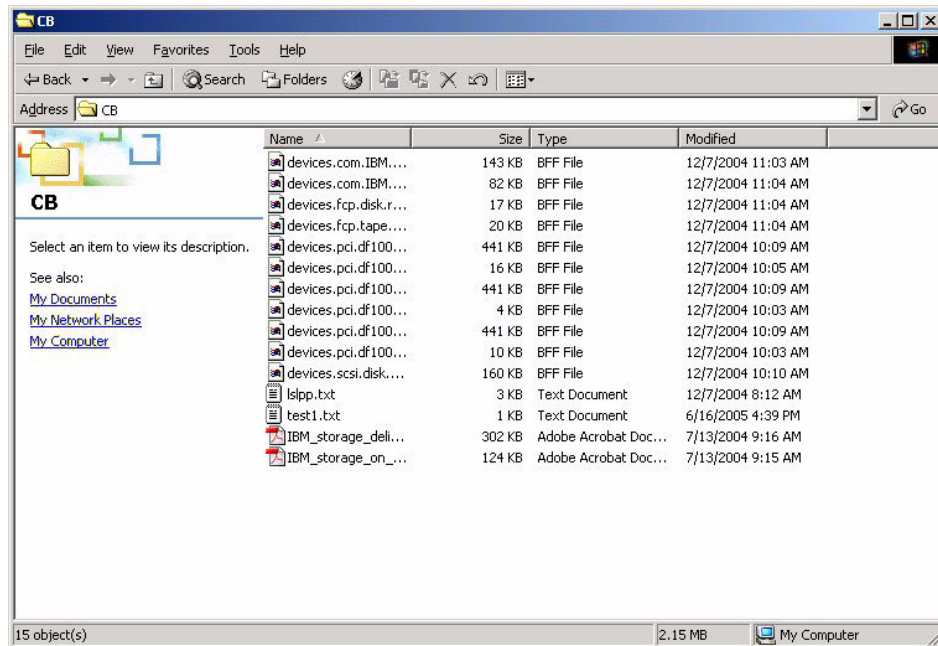
Figure A-17   Retrieved files

You have completed LAB 2 on archive and retrieve. Now we look at some TSM server functions.

Lab 3: Administration Center functions

IBM announced a new Administration Center and Integrated Solution Console (ISC) with the previous Version 5.3 of Tivoli Storage Manager.

Setup

You can log in to the Administration Center from your browser using the URL below.

http://localhost:8421/ibm/
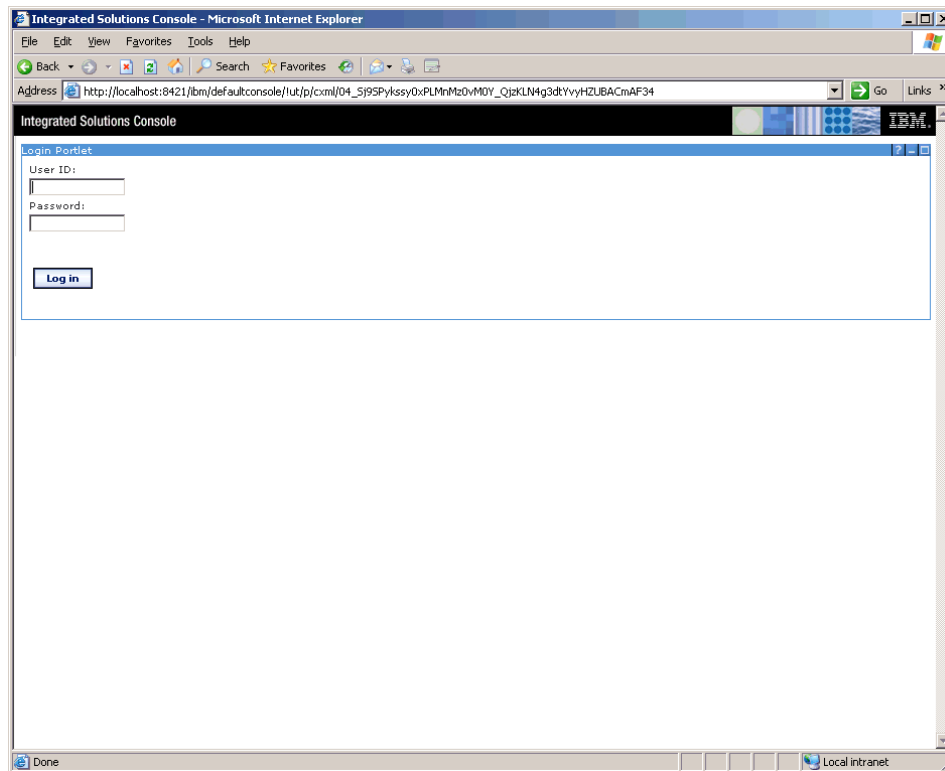
Figure A-18   ISC login screen

1. Log in to the ISC using ID iscadmin and your password.

2. You will see the welcome page, which also shows the version information.
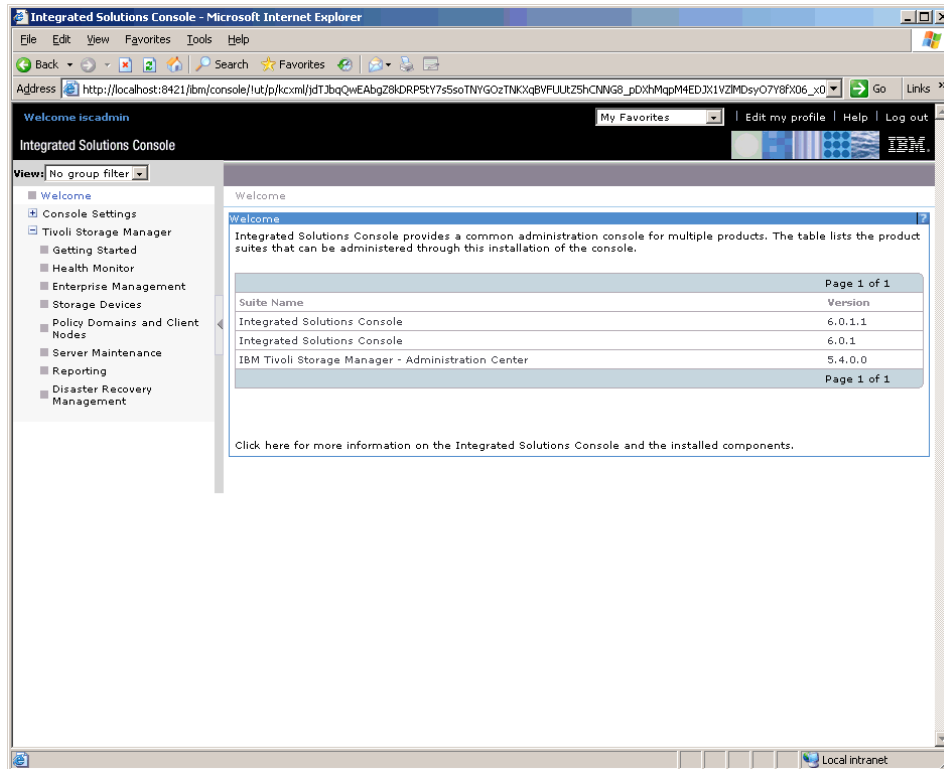
Figure A-19   ISC version information

3. Expand the Tivoli Storage Manager folder in the Welcome page and click **Getting Started** to display the Tivoli Storage Manager Administration Center welcome page. This page provides instructions for using the Administration Center.
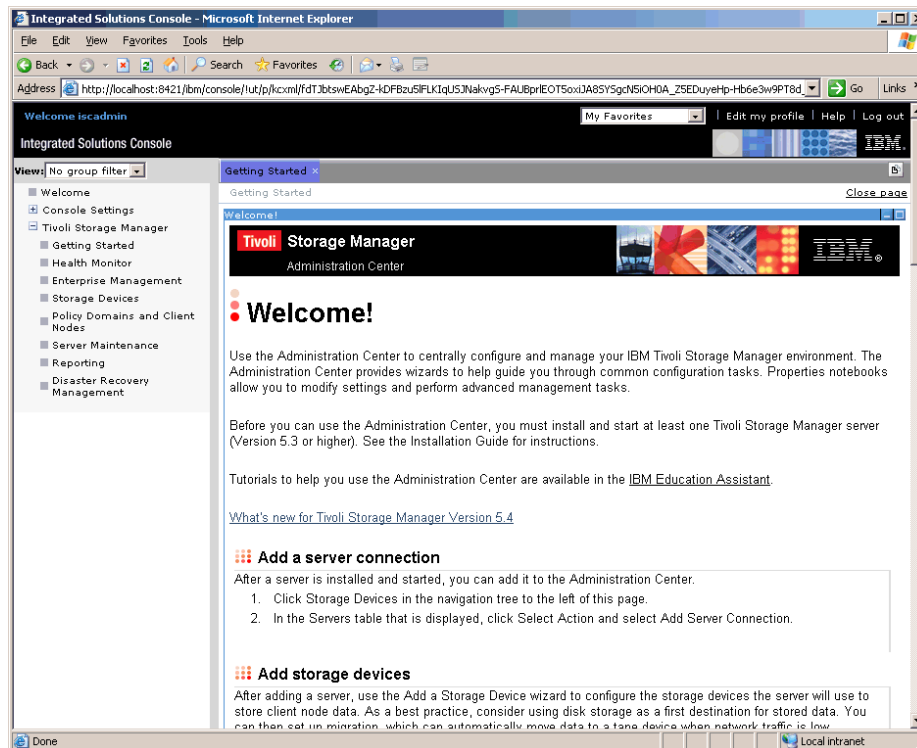
Figure A-20   ISC Administration Center Welcome screen

4. To add a server or modify the existing server connection click **Storage Devices**.

5. From the Action pull-down in the Server window you can select **Add Server Connection** or **Modify Server Connection**. We will choose Modify Server Connection. (We assume that the connection has already been set up.)

6. You can modify the administrator name and its password, as well as the Server IP address and server port number. Click **OK** to commit any changes or **Cancel** if none were made.
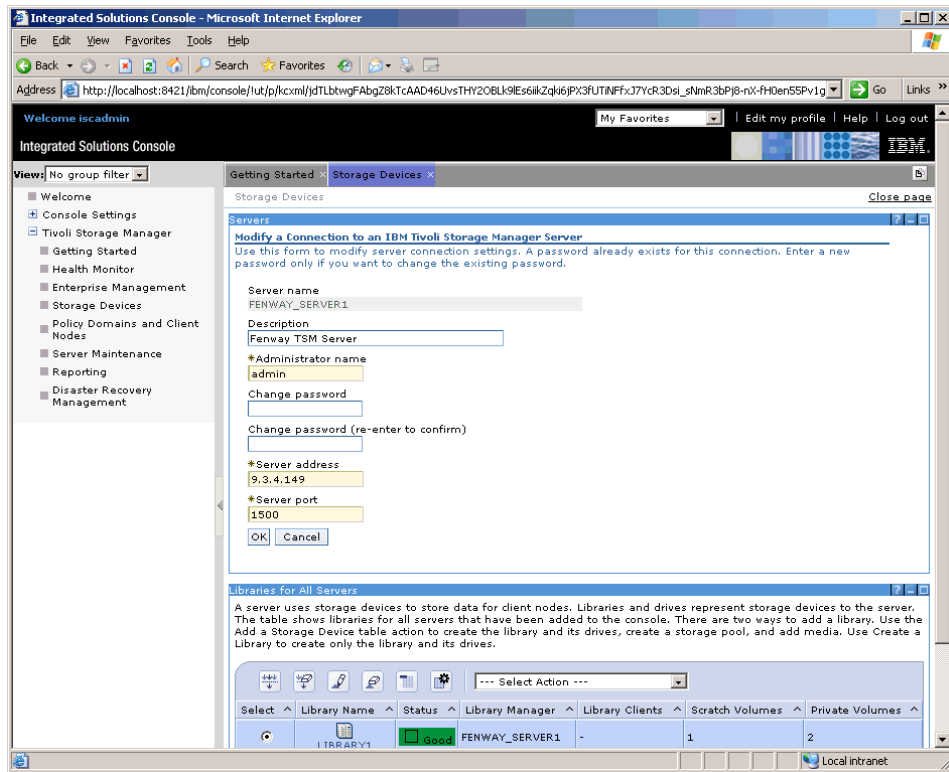
Figure A-21   Modify administrators

7. If you select Server Properties from the Action pull-down in the Server window, you can modify the TSM Server properties. Click **Cancel** to exit this screen.
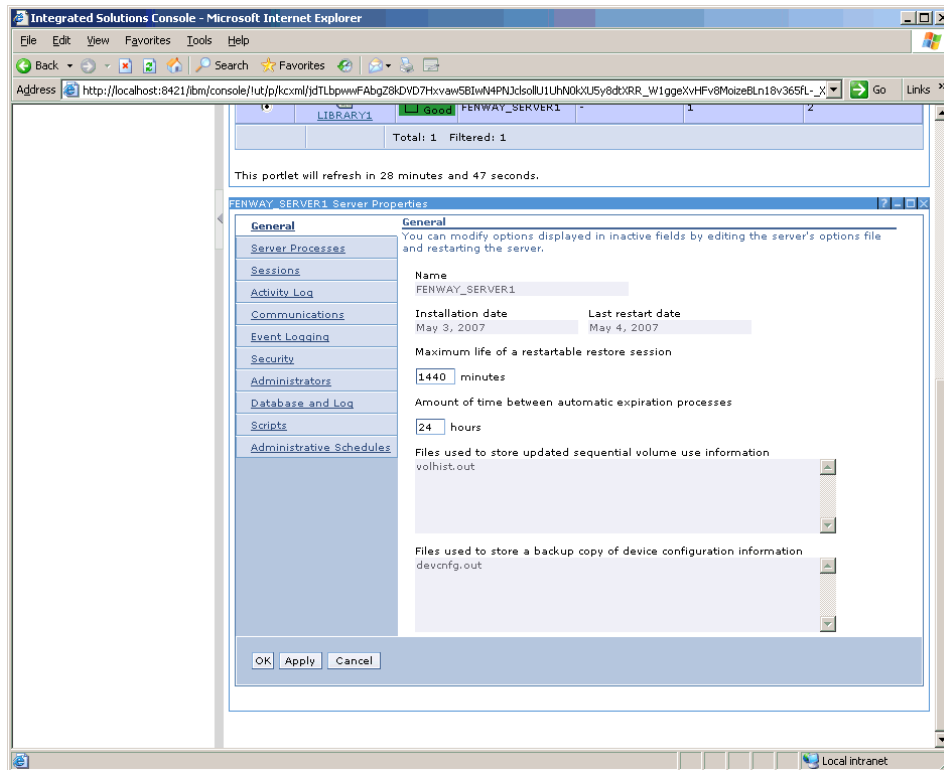
Figure A-22   Server properties

Health Monitor

After this brief introduction, let us start with the Health Monitor.

1. Click **Health Monitor** in the welcome page.
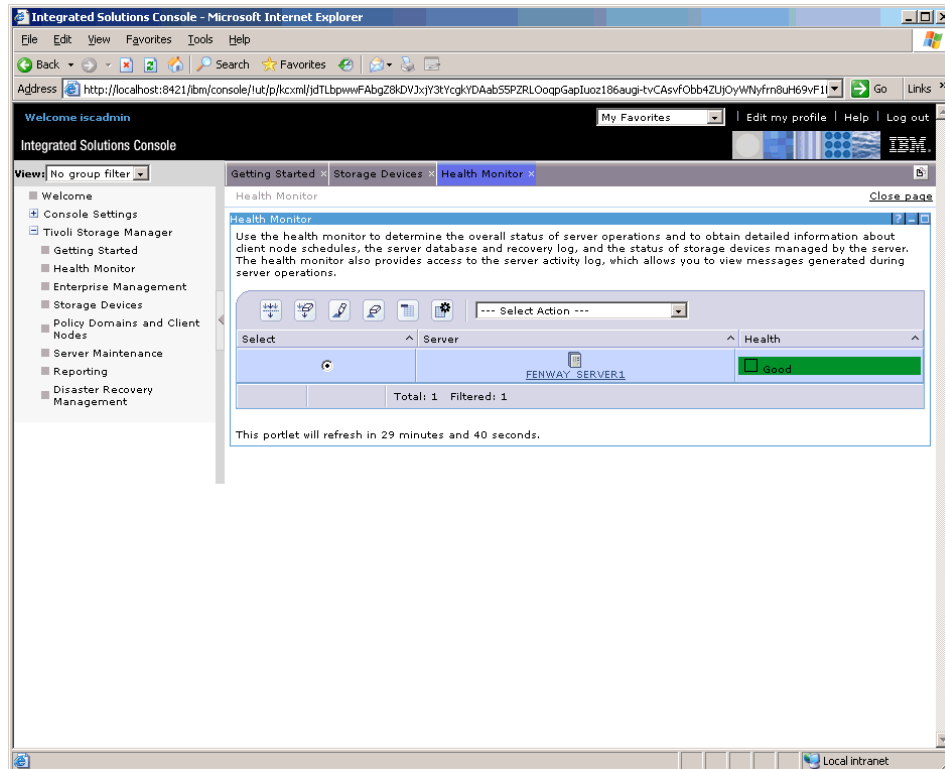
Figure A-23   Health Monitor

2. You can check your server's health. Just under the Health column you will see one of the following colors; the meanings of the colors are:

   – Green (Normal): The server is running.

   – Blue (Unknown): The server could not be contacted.

   – Yellow(Warning): Needs attention; there might be a significant problem.

   – Red (Error): There is a problem with the database, storage devices, or both.

3. From the Action pull-down in the Health Monitor window, select **View Health Details**.
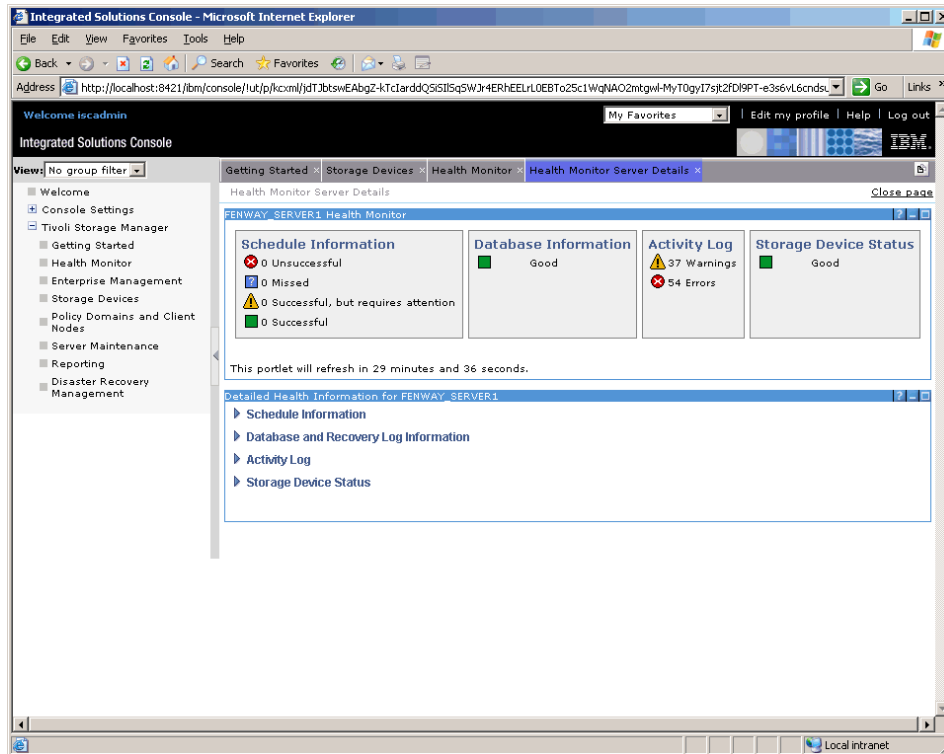
Figure A-24   Health monitor details

4. This screen gives you information about the schedules, TSM database, storage devices, and activity log. You can see the details by clicking each detailed information link.

5. Now we will see how TSM handles the files.

Storage pools

TSM stores backed up data in storage pools. These are reserved disk or tape allocations. You can have one or several storage pools; usually these are chained in a hierarchy. A typical hierarchy is to have a disk pool as the primary backup destination, and define a tape pool as the second level, so that when the disk pool fills up, TSM automatically copies files to the tape pool. In this way, backup operations can seamlessly continue, even when very large amounts of data are being backed up.

We now move to the storage pools by clicking *Storage Devices* from expanded TSM folder in the Welcome Page. Then select the **View Storage Pools** from the Action pull-down in the Server window.

Figure A-25   Storage pools

1. On this page, this TSM server has two storage pools. Note the capacity of each pool, and the Next field; this is used to set up the storage pool hierarchy.

2. Pool DISKPOOL shows 52.9 Pct Utilized (your display may show a different number). This is because up to now we are storing all files in this storage pool. How does TSM know where to put the files? We will see this in the policy lab.

3. Now we will define a new pool. From the Select Action menu choose the **Create a storage pool** option.

4. The name for the new storage pool will be TEST_POOL. Choose the Primary, random access option for this pool. Click **Next**.
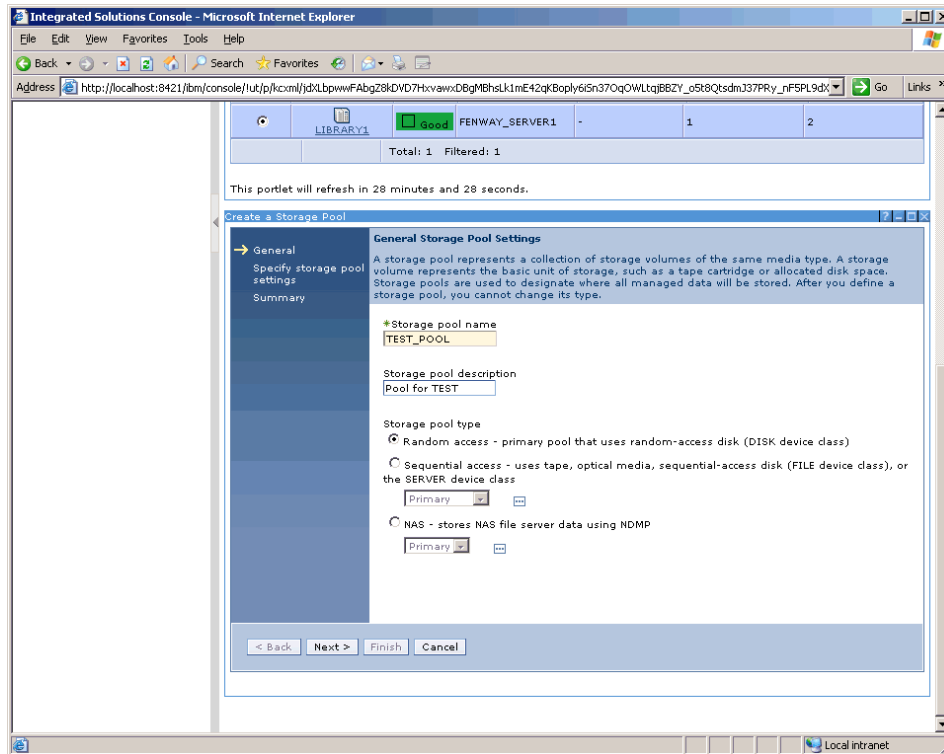
Figure A-26   Storage pool creation

5. Define the size of the TEST_POOL. Select **LTOPOOL1** for the next storage pool. This means that when TEST_POOL fills up, it will automatically overflow (migrate) to LTOPOOL1. Typically we would use a pool on a tape device to be the secondary destination. We specify a small size for the pool (1 MB) so that we will be able to see the migration process. Click **Next**.
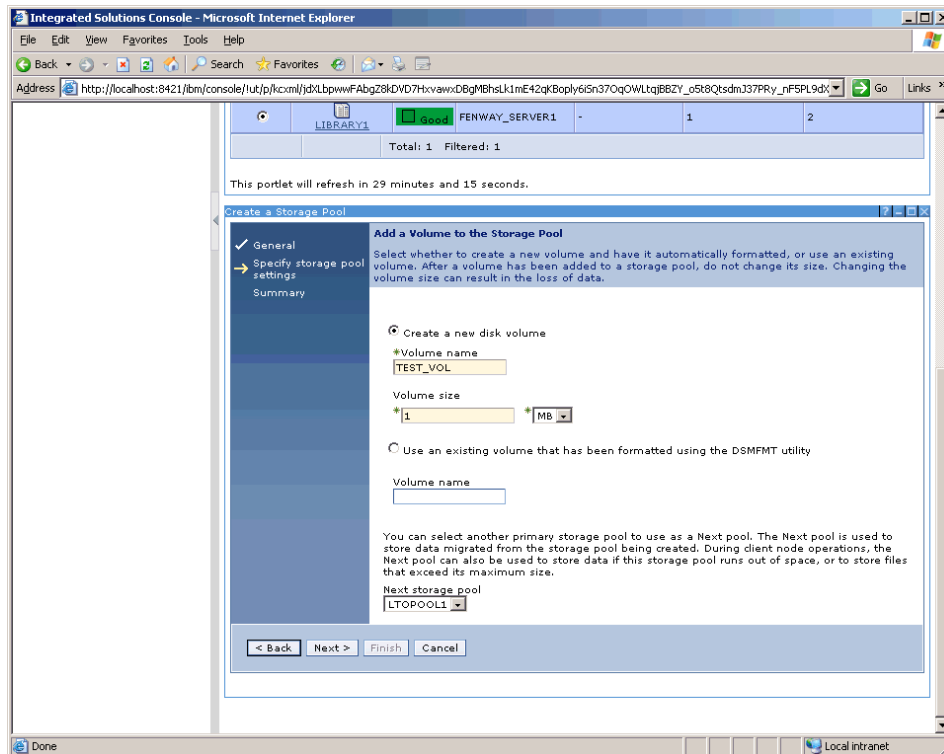
Figure A-27   Storage pool settings

6. The next screen summarizes the details of the new storage pool. Click **Finish** to complete the process.
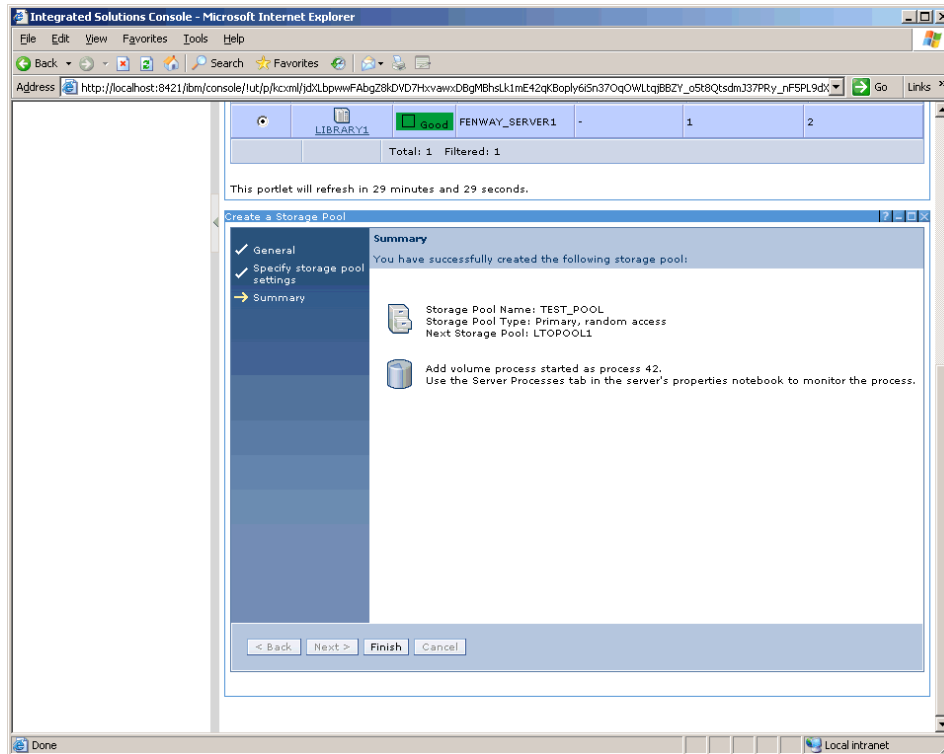
Figure A-28   Storage pool summary

7. To see the new storage pool, click **Storage Devices** in the Welcome page, select **View Storage Pools.**

Figure A-29   Storage pool view

8. We can see that TEST_POOL is created with the size of 1 MB and the next storage pool is LTOPOOL1. Click the **TEST_POOL** to see the detailed information and modify it.
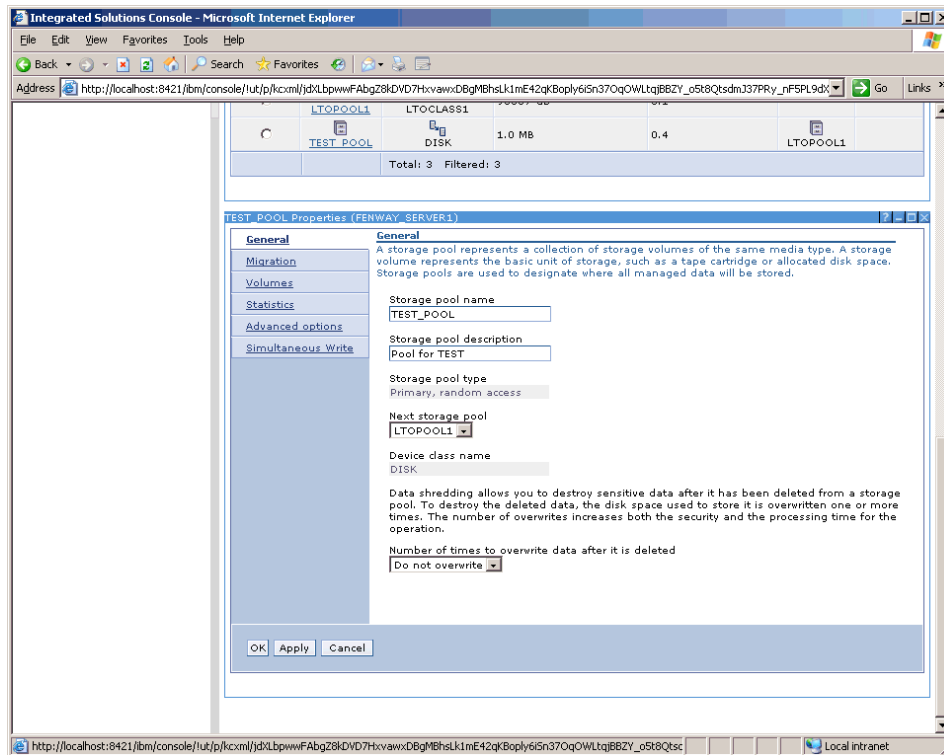
Figure A-30   Storage pool properties

9. You can also modify pool characteristics dynamically. Click the **Migration** link on the left side menu. Make sure the high migration threshold is set to 50, and the low migration threshold is set to 20, as shown. This means when the pool gets to 50 percent full, it will move files in the background to the pool defined as the next, which in this case is LTOPOOL1. Click **OK** to make the change.
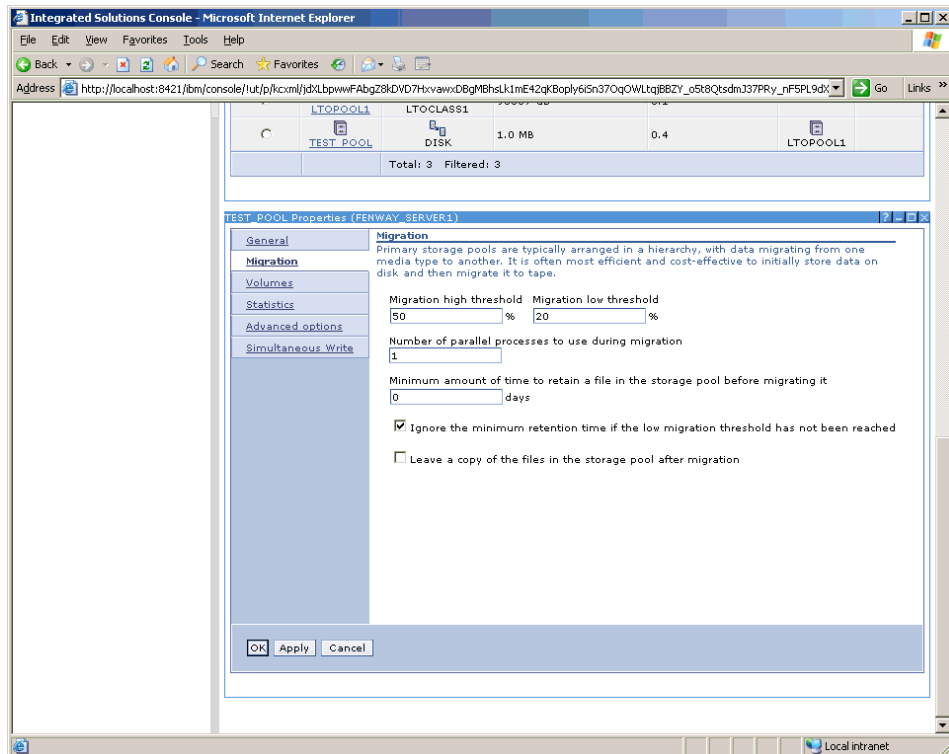
Figure A-31   Modify storage pool

10. To make the storage pool larger, we can add volumes. When we specified the initial size of the pool (1 MB) TSM automatically created a volume of that size. Click the **Volumes** link on the left side menu, and select the **Add Volume** action from the pull-down menu.
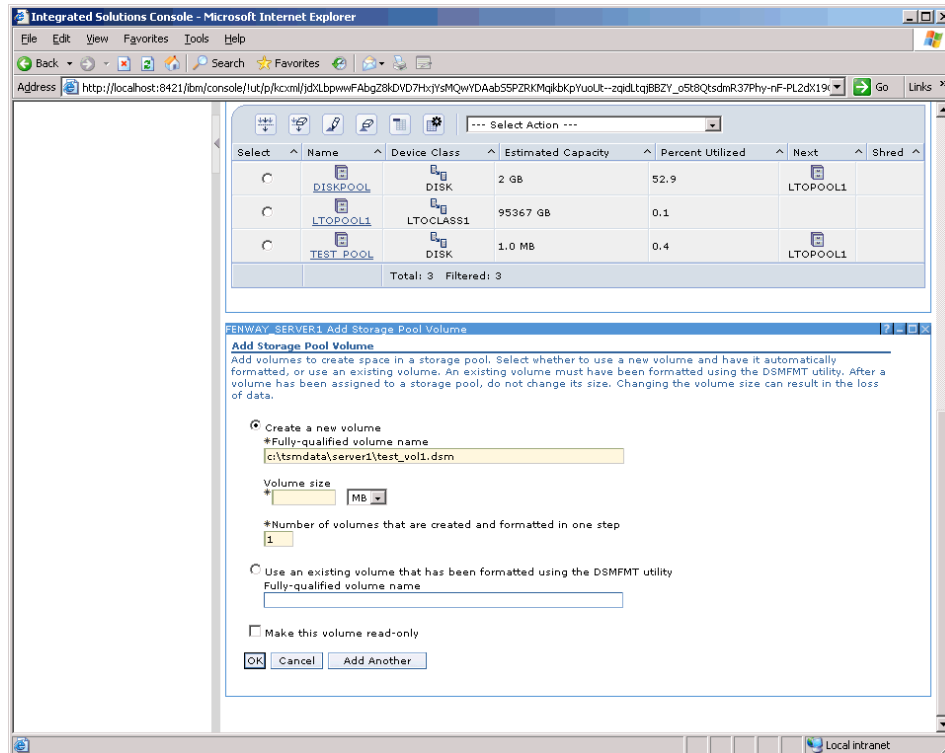
Figure A-32   Creating new volume

11. The volume list displays, showing the volume in the storage pool. If you do not see any volume, select **Refresh table** from the Select Action pull-down.
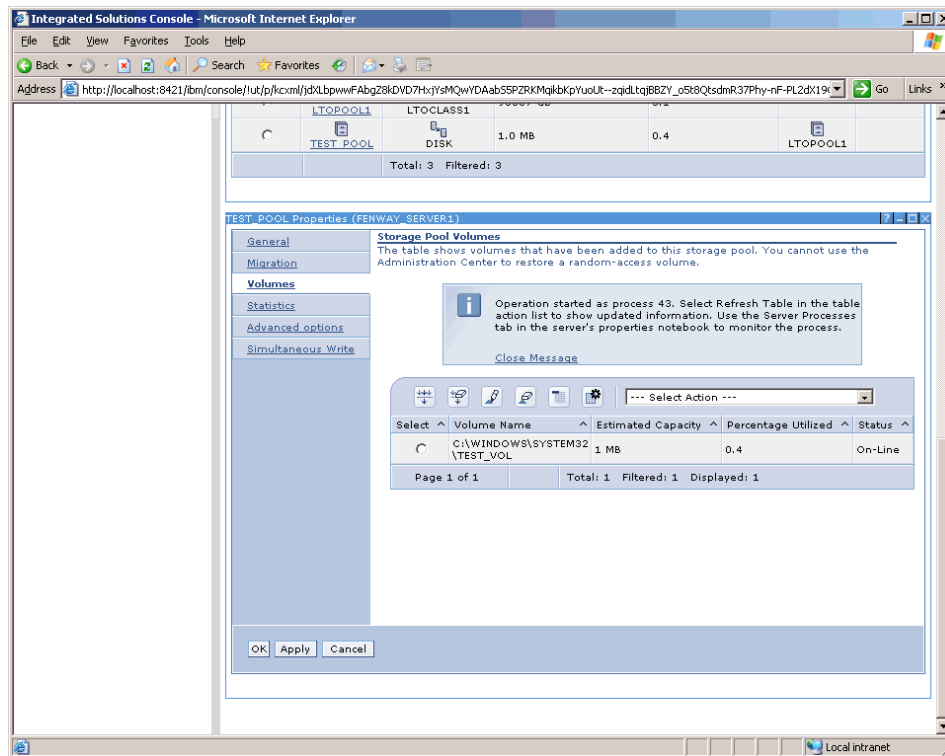
Figure A-33   Completing volume creation

12. We have completed the lab for storage pools. We will use this pool in our next backup by creating and associating another management class. This introduces us to TSM Policy.

Policy

TSM uses defined policies to control the number of versions that are kept for files, how long the versions are retained, and where the files are to be stored (storage pool). A policy domain consists of one or more management classes. A default policy and management class is created when you install TSM. This management class specified to keep two versions of active files (files that still exist on the client), and to store the files in the storage pool LTOPOOL1. We will create another management class that will keep three versions of active files, and that will store the files in the new pool we just created.

1. First of all we view the current policy domain information. Select the **Policy Domains and Client Node** link in the Welcome page. Then select the **VIEW POLICY DOMAINS**.
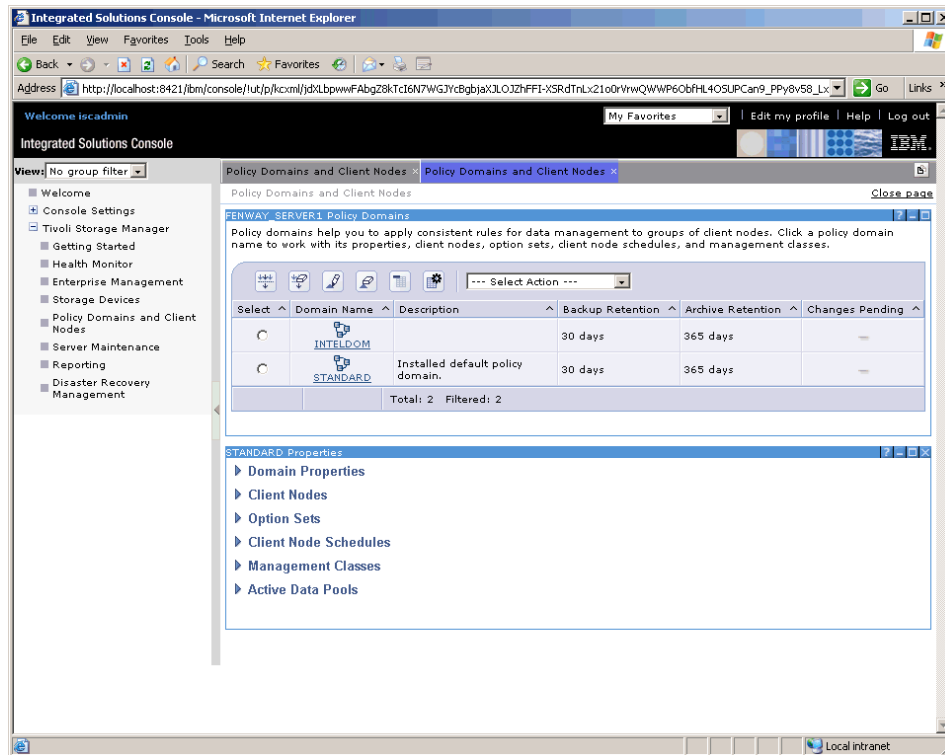
Figure A-34   Policy information

2. From the STANDARD Properties section let us see the detailed information.
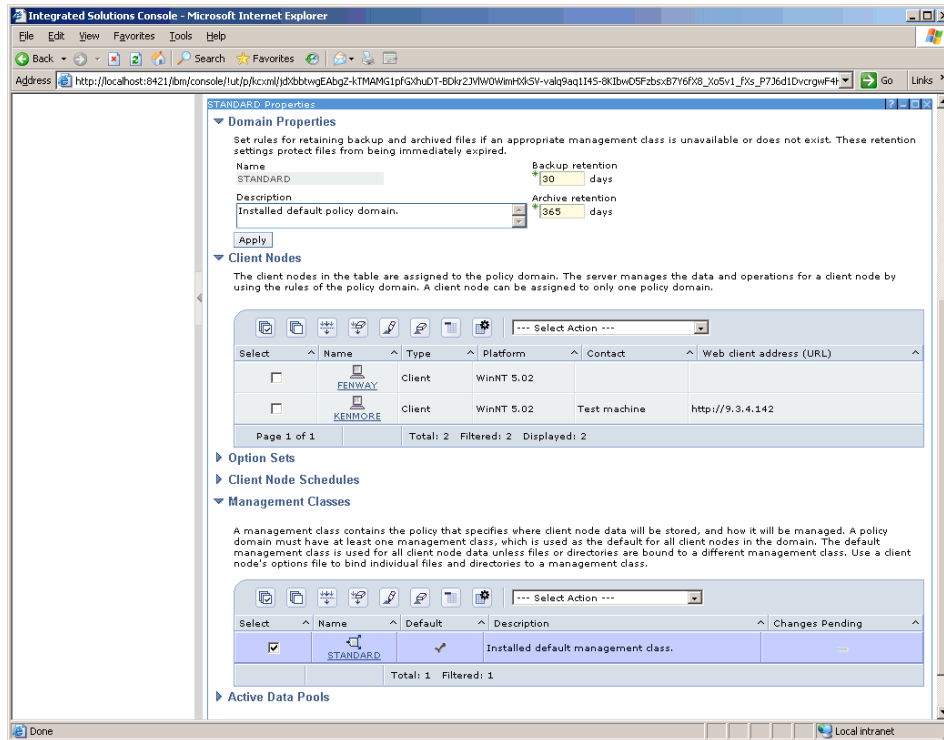
Figure A-35   Management class properties

3. For the standard domain policy:

– Backup retention: 30 days.

– Archive retention: 365 days.

– There are two clients attached to this policy: FENWAY and KENMORE

– Default management class: Standard; this means that this management class will be used unless we explicitly specify to use another.

4. Click **Standard Management Class**. You will see the number of different versions of a file to keep (2), number of days to keep extra copies of existing files, number of versions of deleted files, and number of days to keep the last file copy. You can also see the destination storage pool, which is LTOPOOL1.
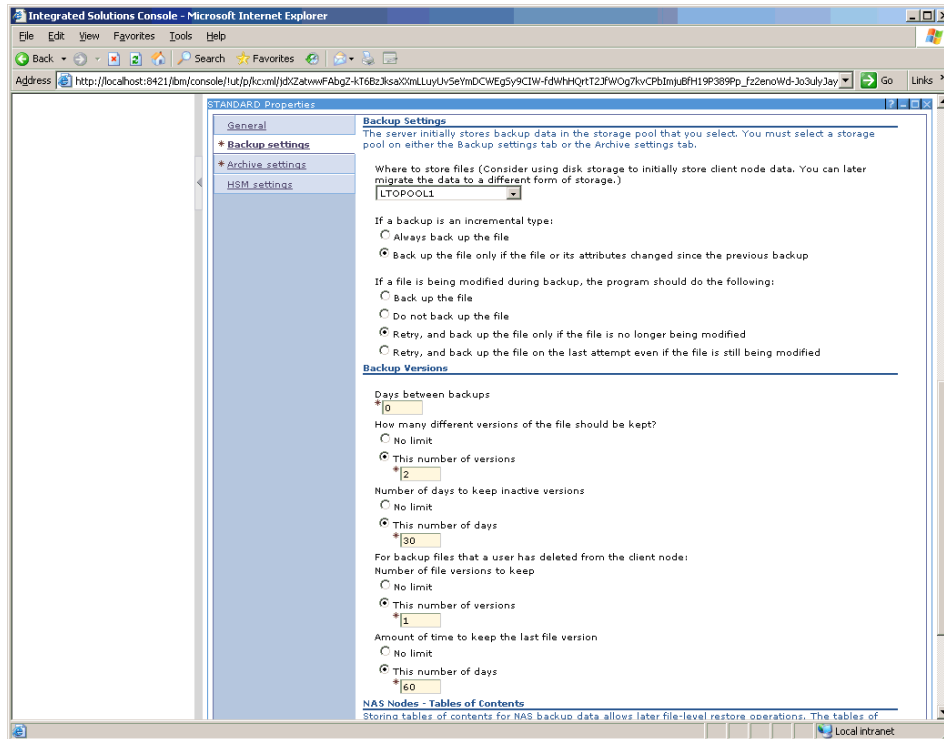
Figure A-36   Backup settings

5. You will define a new management class with its own backup copy group. In the Management Classes section, select **Create Management Class** from the Actions pull-down. Give it a name and description as shown; you can leave the rest of the fields as defaults.

– Management class name: TESTCLASS.

– Migration destination: LTOPOOL1.

– Different versions to keep: 3. Leave the rest as default.

– Copy destination: TEST_POOL and choose retry and backup the file on last attempt, even if the file is still been modified.

6. On the General tab, enter the management class's name as `TESTCLASS`, and you can write the description of this management class.
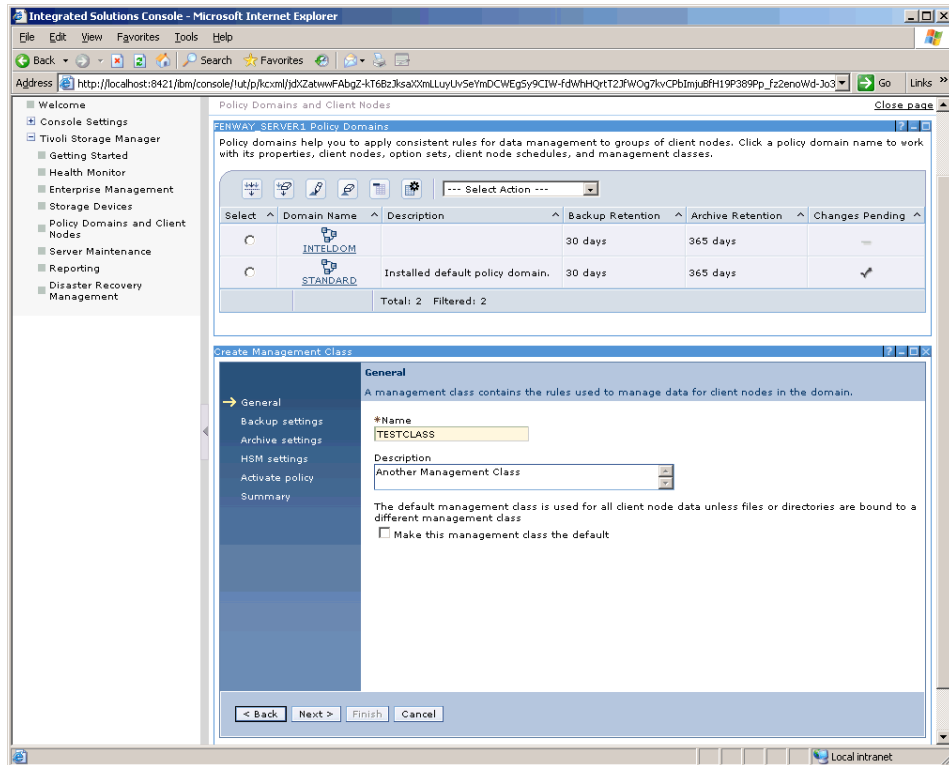
Figure A-37   Creating management class

7. In the Backup section choose the storage pool where this management class will store files. Click **Next** to continue with the backup version numbers.
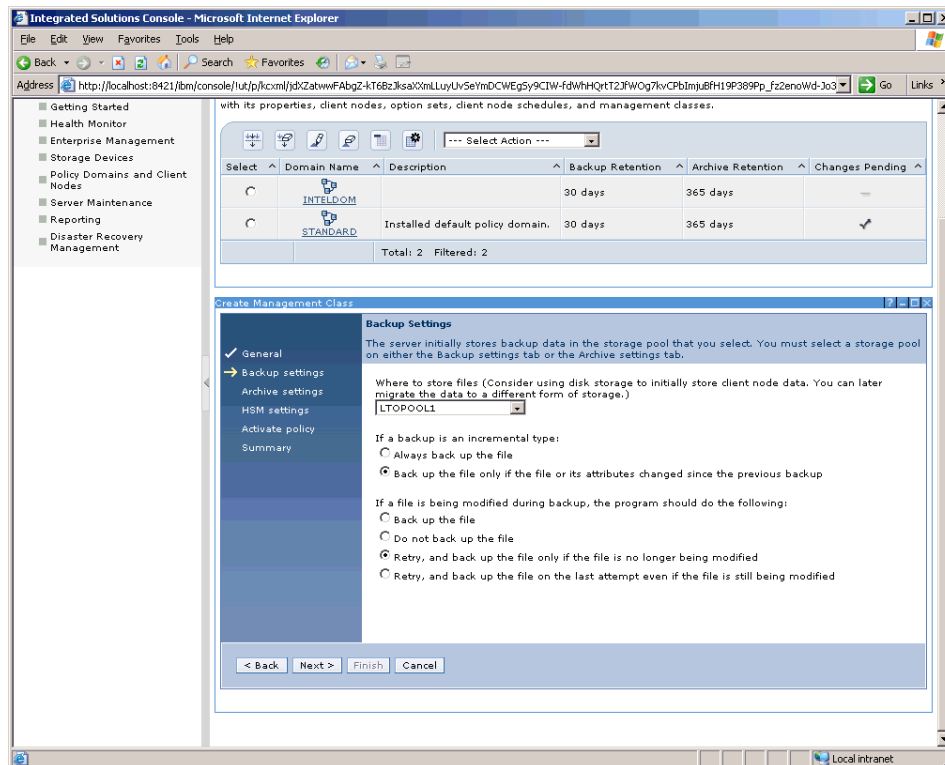
Figure A-38   Backup settings

8. For backup key in:

   – Number of versions: 3
   – This number of days: 30

   For Deleted versions key in:

   – This number of versions: 1
   – This number of days: 60
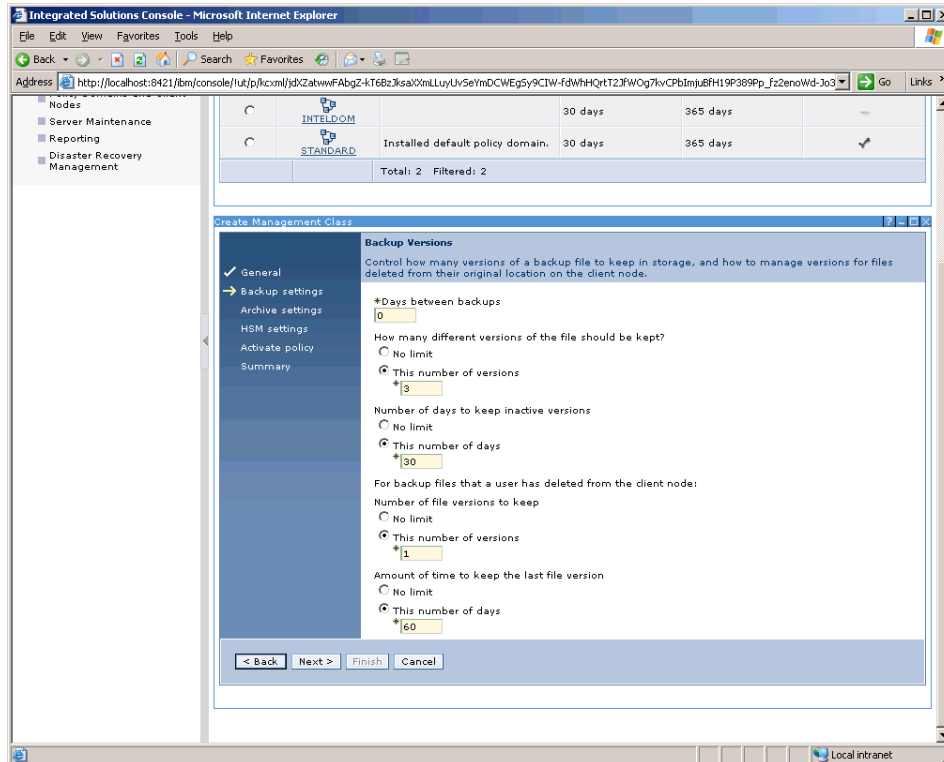
Figure A-39   Backup settings

9. For Archive settings enter `TEST_POOL` as a copy destination.
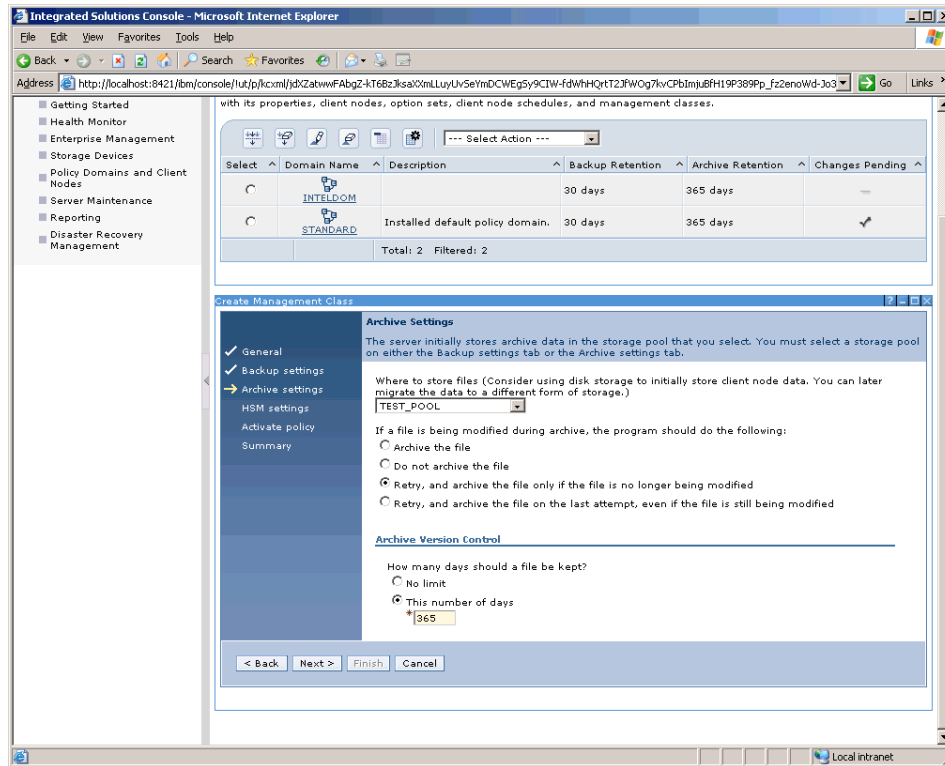
Figure A-40   Archive settings

10. Activate the policy set. Click **Next**.

11. Click **Next** until you check the summary page and click **Finish** to complete the creation of the new management class..
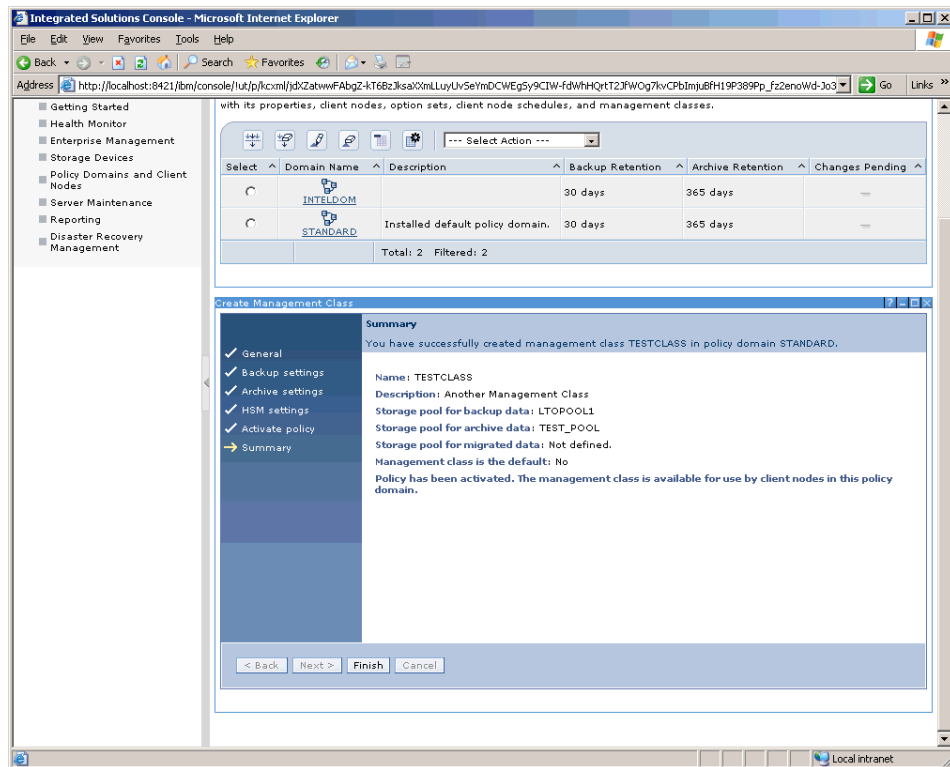
Figure A-41   Management class summary

12. The new class shows up. Experienced administrators of Tivoli Storage Manager are probably familiar with policy sets, which contain the management classes within a policy domain. The Administration Center eliminates the need to validate or activate a policy set when you make changes to management classes. In the Administration Center, you work only with the active policy set. Any changes that you make to management classes are immediately activated, without additional effort on your part.

13. We have now automatically activated our changes to the policy set. You can also change your default management class from one to another by choosing from the Management Classes section, clicking the desired class, and choosing from the pull-down list Make Management Class default. Do not do this now, or the next sections may not execute as expected.
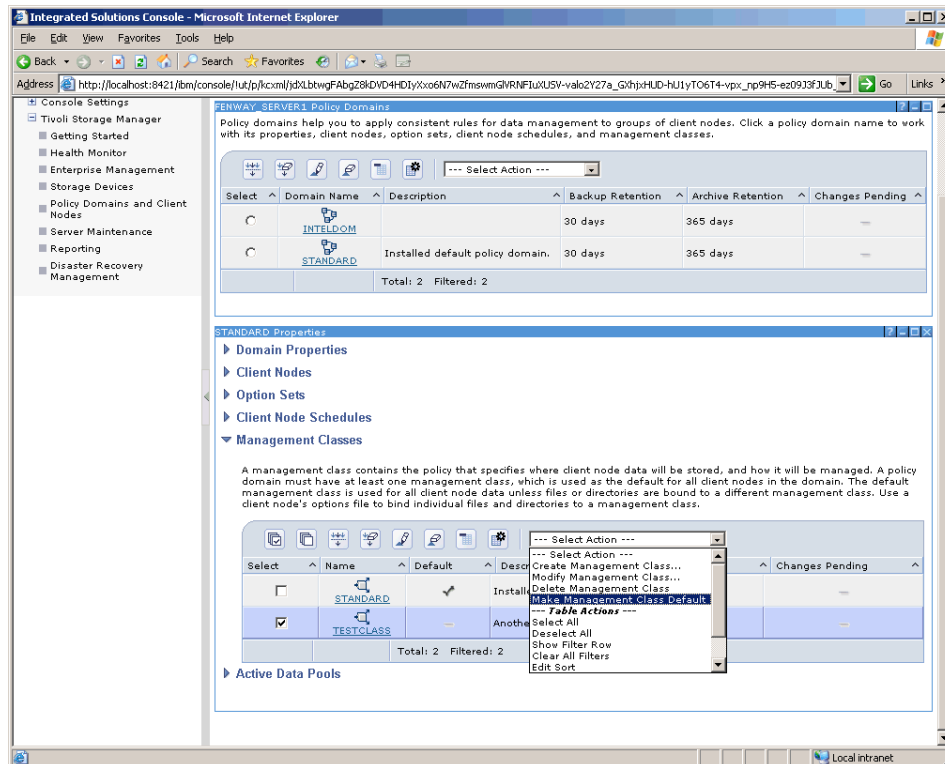
Figure A-42   Make the management class default

14. You have finished the policy domain lab. We now return to the client to see how our new policy takes effect.

TSM server database

For the TSM server database:

1. On the Health Monitor screen click the detailed information for the database and recovery log.
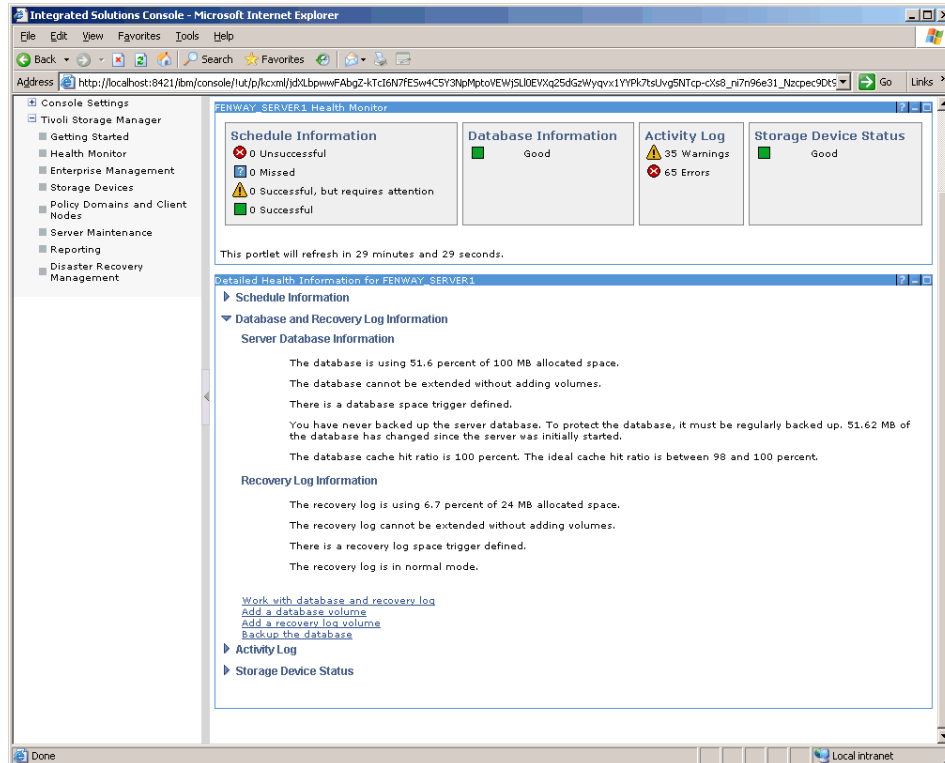
Figure A-43   Server database properties

2. This section displays the server database and recovery log information. Click the **Work with database and recovery log** link.
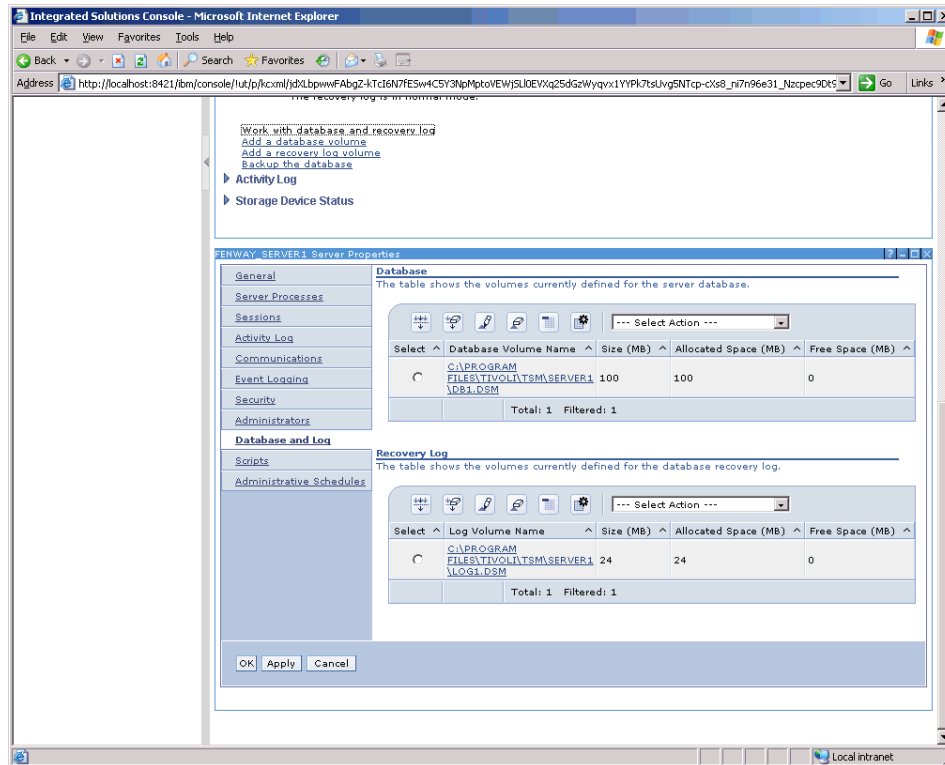
Figure A-44   Database and recovery log

3. This screen shows the current size, allocated space, and free space of the database and recovery log. You can see more details by selecting **Database properties** from the pull-down menu.
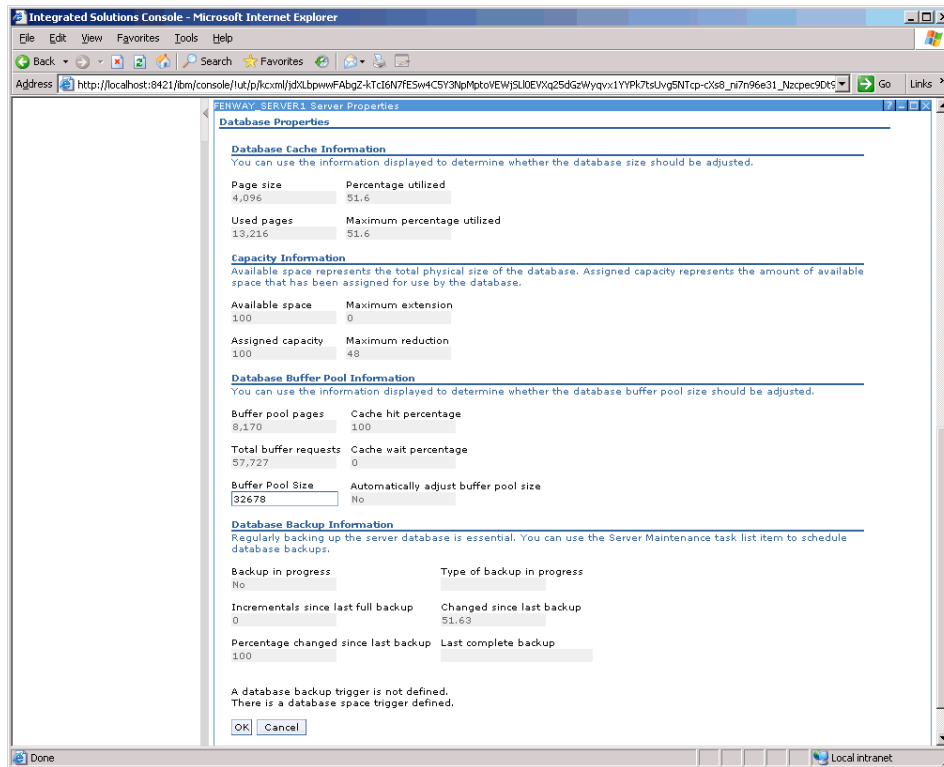
Figure A-45   Database properties

4. We can see that the TSM database has a 100MB capacity and is 51.6percent utilized.

5. Let us add a volume to the database. From the Select Action pull-down menu, select **Add Volume**.

6. Complete the fields as shown on the screen and click **OK** to execute.
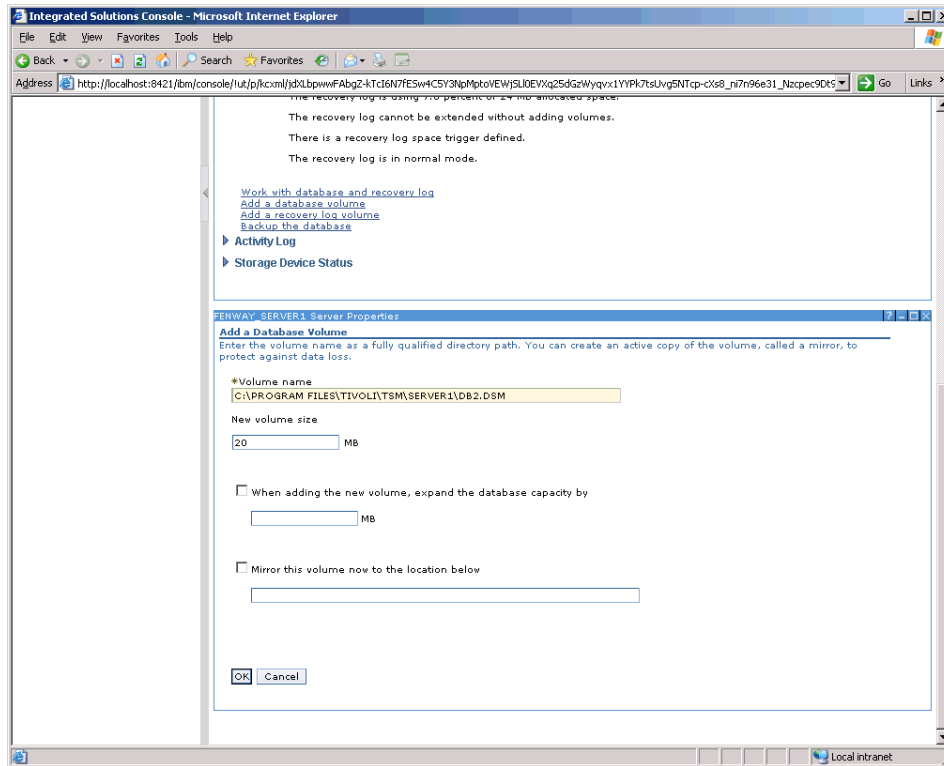
Figure A-46   Adding database volume

7. Then it will return the database information page and you can see the new volume has been added.

Figure A-47   Viewing capacity

8. As shown above, you see that the 20 MB volume is added but not allocated yet. To be able to use the new capacity we must now extend the database. Select **Extend or Reduce** from the action pull-down menu, specify the maximum expansion amount (20 MB), and click **OK** to complete the action.
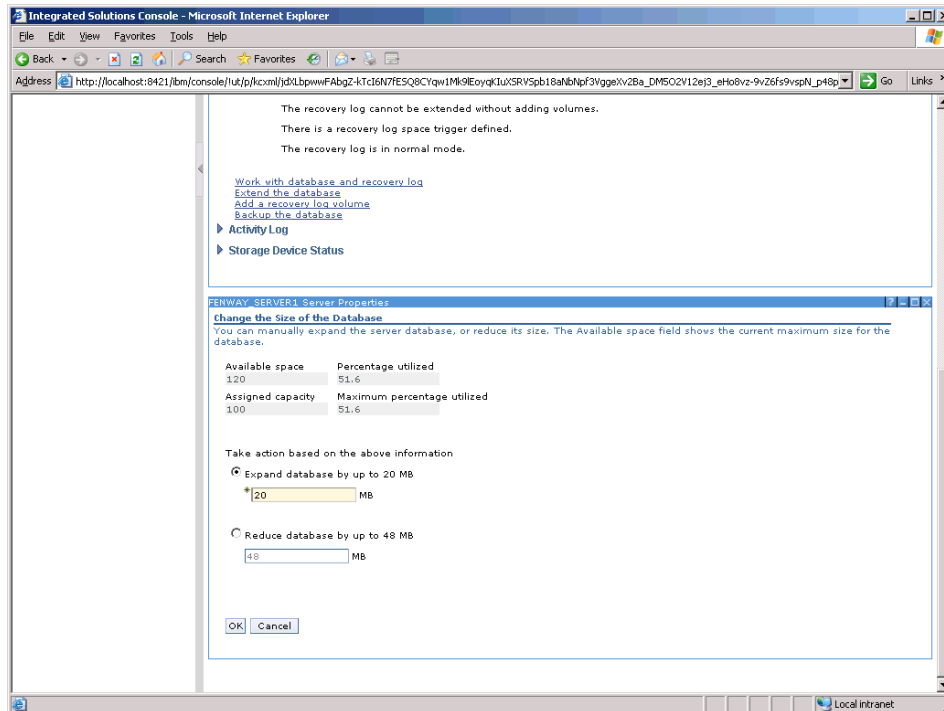
Figure A-48   Extending database

9. When you check the database volume details again you will see that all the available space in the new volume is allocated to the TSM database, and the total allocated size is 120MB.
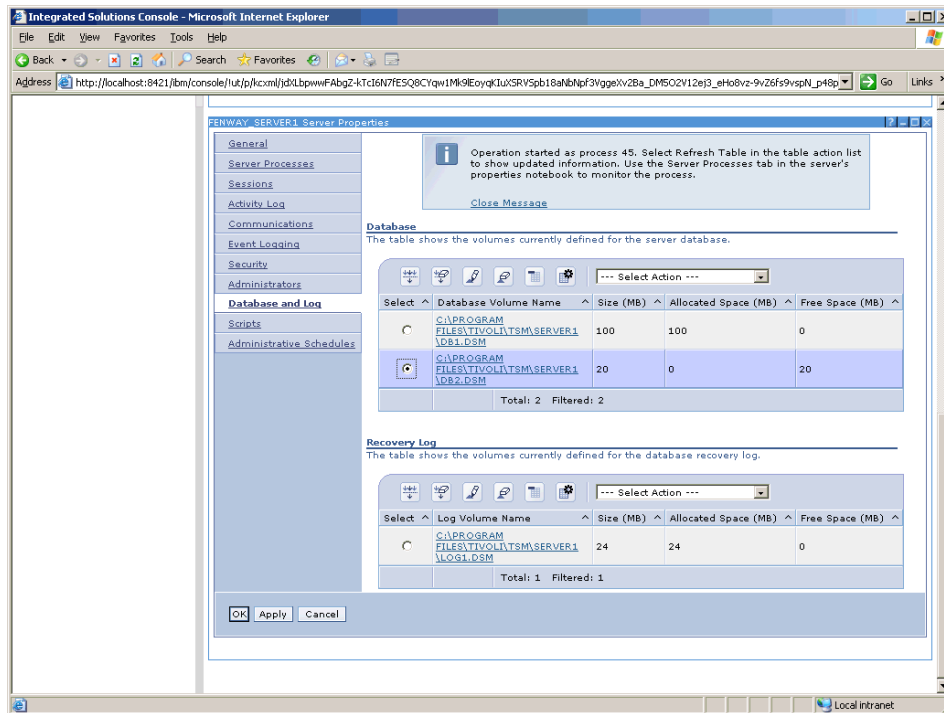
Figure A-49   Extending volume

10. Notice that after the expansion of the database volume, the utilization percentage is significantly reduced from 51.6 percent to 3.8 percent and capacity is 43MB.
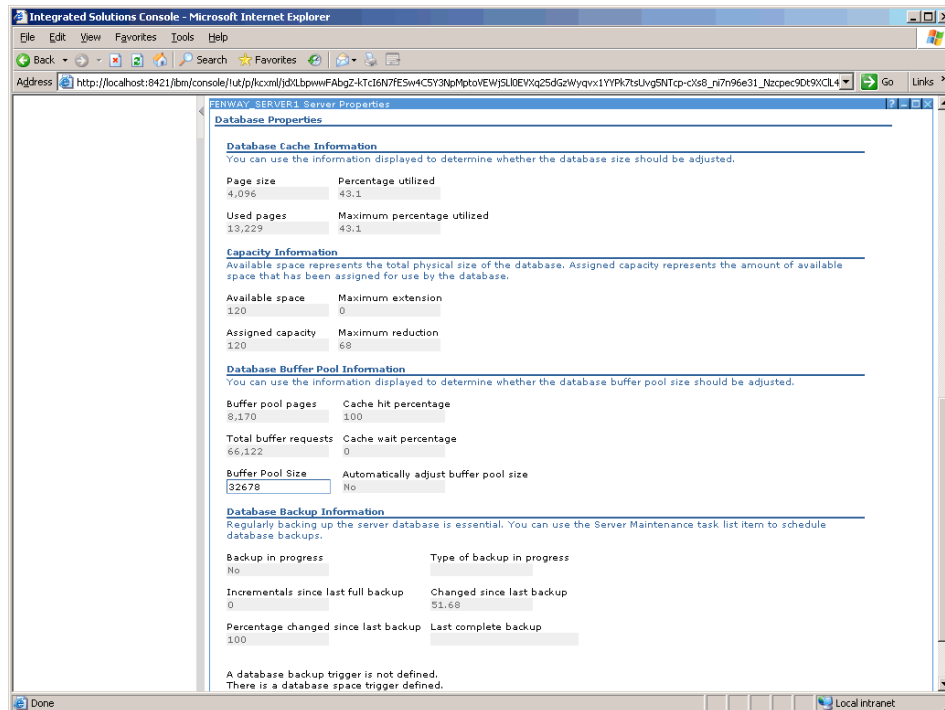
Figure A-50   Database properties

11. Optionally, you can do the same operations (view, add volumes, extend) for the recovery log, and you can also set up 2-way and 3-way mirroring for the database and recovery log. This makes the TSM system more highly available. These are left as an exercise.

This concludes our hands-on TSM LAB Version 5.4.

B

Sample test questions

This sample test is designed to give the candidate an idea of the content and format of the questions that will be on the certification exam. Performance on the sample test is NOT an indicator of performance on the certification exam and this should not be considered an assessment tool.

Sample test questions

1. What are two reasons Simultaneous Writes to multiple storage pools would be used? (Choose two.)

   a. decreases backup window
   b. improves LAN-free performance
   c. accommodates copy of data off-site faster
   d. allows more efficient use of tape cartridges
   e. reduces overall copy storage pool times

2. New management has adopted an IBM Tivoli Storage Manager configuration. They would like to understand the current retention period settings on the systems so that they can better tune their service level agreements.

   What command will they issue to find their current TSM retention policies?

   a. query mgmt f=d
   b. query policy f=d
   c. query ret f=d
   d. query copy f=d

3. The customer requires database data to be restored quickly, but also has data in files that do not require quick restore. Which storage pool hierarchy best meets these needs?

   a. non-database data to a disk pool, database data to a common tape pool
   b. two disk storage pools; one with collocation, the other without collocation
   c. one disk storage pool and two tape storage pools; one tape pool with collocation, the other a non-collated tape pool
   d. two disk pools; one migrating to a collocated tape pool, the other migrating to a non-collocated tape pool

4. How should you determine if the operating system type and version proposed for use by the IBM Tivoli Storage Manager (TSM) server is supported by TSM 5.4?

   a. check the TSM Concepts Redbook
   b. check the TSM Administrator's Guide
   c. check the IBM/Tivoli support Web site
   d. check with the operating system vendor

5. When using the IBM Tivoli Storage Manager (Storage Manager) Device Install wizard, what information is needed to install the library?

   a. libraryname
   b. element address
   c. library device type
   d. World Wide Name

6. What is the default installation directory for the Tivoli Storage Manager server on AIX?

   a. /usr/tivoli/tsm/server/bin
   b. /opt/tivoli/tsm/server/bin
   c. /etc/tsm/server/bin
   d. /opt/IBM/tivoli/tsm/server/bin

7. A customer has an IBM Tivoli Storage Manager (Storage Manager) server with 1 GB of RAM and an LTO tape library, a one hour backup window, and 300 GB of data with a 20% rate of change. What is the minimum network infrastructure necessary to accomplish the backup?

   a. Gigabit Ethernet
   b. ATM 155 Mb/sec
   c. Token Ring 16Mb/sec
   d. 2 Gigabit Fibre channel

8. What is the most important consideration when installing the Integrated Solutions Console (ISC)?

   a. memory
   b. number of clients
   c. server platforms
   d. disk space

9. Which IBM Tivoli Storage Manager (TSM) product must be used if LAN-Free data transfer from a backup-archive client is to be used?

   a. TSM for SAN
   b. TSM for ACS
   c. TSM for SANergy
   d. TSM for Server-Free

10. Which account is created during an installation of the IBM Tivoli Storage Manager (Storage Manager) Server?

    a. Administrator
    b. Root
    c. Admin
    d. Server_Console

11. Which IBM Tivoli Storage Manager (Storage Manager) Disaster Recovery Manager (DRM) command identifies which tapes should be taken off site from storage pool COPYPOOL?

    a. query drvolume * search=offsite
    b. query droffsite begindate=today-1
    c. query drmedia wherestate=mountable
    d. query actlog search=offsite begindate=today-1

12. A tape has been dropped by an operator and the operations manager is concerned that the data on the tape may have been corrupted. How can the information on the tape be validated?

    a. use the audit volume command from the IBM Tivoli Storage Manager (Storage Manager) server
    b. use the audit tape command from the Storage Manager server
    c. this cannot be done, the data must be restore from a copy storage pool
    d. use the Tape Audit and Recovery Tool (TART) from the Storage Manager Server

13. Remote and mobile clients using the Windows operating system often suffer from severe network bandwidth restrictions.

    Which two methodologies should be considered to offset this situation and reduce backup times? (Choose two.)

    a. data encryption
    b. client compression
    c. prompted scheduling
    d. memoryefficient backup
    e. adaptive subfile differencing backup

14. A hard drive is replaced on a Windows workstation, and the data is restored from IBM Tivoli Storage Manager (Storage Manager). The user calls and

reports the corruption of one database, but everything else is running correctly. What is a reason this situation occurred?

a. Restored files were backed up as shared dynamic.
b. The access permissions on the file were not restored.
c. The Storage Manager client was not installed correctly.
d. The backup db type=full command was not used.

15. Which protocol should be used to improve performance when the Linux Server and Client are on the same machine?

a. TCPIP
b. Named Pipes
c. NETBUI
d. Shared Memory

16. Which command is used to register new IBM Tivoli Storage Manager licenses?

a. Add License
b. Create License
c. Register License
d. #Include <License>

17. A customer has many remote and local clients supported by a Tivoli Storage Manager server. To reduce backup pool space, they would like to exclude the scheduled backups of all test data.

What is the easiest and fastest way to exclude this data?

a. Use admin schedule
b. Update every client option file
c. Use CLIENTOPT and CLOPTSET
d. Change the global Storage Manager server option to exclude.dir

18. In an environment where the average file size is less than 5KB, what are three benefits of file aggregation? (Choose three.)

a. Improved tape utilization
b. Reduced tape seek time
c. Improved performance of single file restore
d. Improved performance of migration processes
e. Improved performance of incremental backups

      f.   Reduced overhead of IBM Tivoli Storage Manager database updates

19. In preparation for a disaster recovery test, an IBM Tivoli Storage Manager (TSM) administrator runs Query DRMedia.

    What is the result of the administrator's action?

    a. The TSM server restores the latest database backup before it crashed.
    b. A TSM database backup and a copy of storage pool volumes are generated.
    c. The TSM database backup and a copy of storage pool volumes are modified.
    d. Information about the TSM database backup and copy storage pool volumes is displayed.

20. A customer has a requirement to back up a Network Appliance Network Attached Storage device (NAS). How can the tape library be connected in order to backup this device using NDMP?

    a. The tape library is dedicated to the Tivoli Storage Manager server.
    b. The tape library is connected to both the NAS device and the Tivoli Storage Manager server.
    c. The tape library is connected to the Tivoli Storage Manager client which can NFS mount the NAS device.
    d. The tape library is connected to the Tivoli Storage Manager client, which can NFS mount the NAS device and the Tivoli Storage Manager server.

21. What are three reasons for creating a disk storage pool? (Choose three.)

    a. Ability to do a faster restore
    b. High frequency of backups
    c. Large number of tape drives
    d. More clients than tape drives
    e. Enables sequential disk write-backs
    f.   Takes advantage of disk multistreaming

22. As an IBM Tivoli Storage Manager (TSM) administrator, you have configured a Disaster Recovery Plan using the Disaster Recovery Manager (DRM) function of TSM. You now need to perform a recovery.

    Where is the information on the type of devices needed to restore the TSM database and client data located?

a. Recovery Plan file
b. Volume History file
c. TSM Administrator's Guide
d. Query drmedia command output

23. What is the best way to determine if a client backup utilized the LAN or SAN to transfer its data in Tivoli Storage Manager V5.4?

a. Run the q actlog command
b. Run the q event * * command
c. Check the SAN Data Gateway log
d. Run the q sessions type=san command

24. It has been determined that the Tivoli Storage Manager recovery log has filled up and the Tivoli Storage Manager Server has stopped. How can the Tivoli Storage Manager server be brought back online?

a. Delete the log files and restart the server.
b. Use the Log File Trimmer Utility (LOTU) to reset the logs.
c. Use the Tivoli Storage Manager Database wizard to extend the logs.
d. Use the dsmfmt command to create additional log space and issue a dsmserv extend log.

25. An IBM Tivoli Storage Manager (TSM) backup-archive client backup failure occurs.

Which location is the best place to find the cause or indication of the failure?

a. IBM Tivoli website
b. TSM Client dsmerror.log
c. TSM Client activity log
d. Windows Registry file

Answer Key:

1. C, E

2. D

3. D

4. C

5. C

6. A

7. A

8. A

9. A

10. D

11. C

12. A

13. B, E

14. A

15. D

16. C

17. C

18. D, E, F

19. D

20. B

21. A, B, D

22. A

23. A

24. D

25. B

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

For information about ordering these publications, see "How to get Redbooks" on page 258. Note that some of the documents referenced here may be available in softcopy only.

► IBM Tivoli Storage Management Concepts, SG24-4877-04

► Deployment Guide Series: IBM Tivoli Storage Manager Version 5.4, SG24-7379-00

► IBM Tivoli Storage Manager Implementation Guide, SG24-5416-03

► IBM Tivoli Storage Manager Version 5.3 Technical Guide, SG24-6638-00

► Certification Study Guide: IBM Tivoli Storage Manager Version 5.3, SG24-7127-00

Other publications

These publications are also relevant as further information sources:

► IBM Tivoli Storage Manager for AIX Installation Guide Version 5.4, SC32-0134-00

► IBM Tivoli Storage Manager for AIX Administrator's Guide Version 5.4, SC32-0117-00

► IBM Tivoli Storage Manager for Windows Installation Guide Version 5.4, SC32-0138-00

► IBM Tivoli Storage Manager for Windows Administrator's Guide Version 5.4, SC32-0121-00

► IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide V5.4, SC32-0145-00

► IBM Tivoli Storage Manager for Windows Backup-Archive Clients Installation and User's Guide Version 5.4, SC32-0146-00

257

Online resources

These Web sites are also relevant as further information sources:

► Tivoli Storage Manager product documentation

`http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp`

► IBM Tivoli Storage Manager Express

`http://www.ibm.com/software/tivoli/products/storage-mgr-express/`

► IBM Tivoli Storage Manager Basic Edition information

`http://www.ibm.com/software/tivoli/products/storage-mgr/`

► IBM Tivoli Storage Manager Extended Edition information

`http://www.ibm.com/software/tivoli/products/storage-mgr-extended/`

► IBM Training

`http://www.ibm.com/software/sw-training/`

How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

IBM

Redbooks

Certification Study Guide: IBM Tivoli Storage Manager 5.4

(1.5" spine)
1.5"<-> 1.998"
789 <->1051 pages

IBM

Redbooks

Certification Study Guide: IBM Tivoli Storage Manager 5.4

(1.0" spine)
0.875"<->1.498"
460 <-> 788 pages

IBM

Redbooks

Certification Study Guide: IBM Tivoli Storage Manager 5.4

(0.5" spine)
0.475"<->0.875"
250 <-> 459 pages

IBM

Redbooks

Certification Study Guide: IBM Tivoli Storage Manager 5.4

(0.2"spine)
0.17"<->0.473"
90<->249 pages

Certification Study Guide: IBM Tivoli Storage Manager 5.4

(0.1"spine)
0.1"<->0.169"
53<->89 pages

IBM

Redbooks

Certification Study Guide: IBM Tivoli Storage Manager 5.4

IBM

Redbooks

Certification Study Guide: IBM Tivoli Storage Manager 5.4

IBM®

Certification Study Guide: IBM Tivoli Storage Manager 5.4

Redbooks®

Prepare to become certified for IBM Tivoli Storage Manager V5.4

Covers new features in latest version

Contains hands-on lab and sample test questions

This IBM Redbooks publication is a study guide for IBM Tivoli® Storage Manager Version 5.4 and is aimed at individuals who want to get IBM Certifications in this specific product.

The IBM Tivoli Storage Manager Certification, offered through the Professional Certification Program from IBM, is designed to validate the skills required of technical professionals who work in the implementation of the IBM Tivoli Storage Manager Version 5.4 product.

This book provides a combination of theory and practical experience needed for a general understanding of the subject matter. It also provides sample questions that will help in the evaluation of personal progress and provide familiarity with the types of questions that will be encountered in the exam.

This publication does not replace practical experience, nor is it designed to be a stand-alone guide for any subject. Instead, it is an effective tool that, when combined with education activities and experience, can be a very useful preparation guide for the exam.