IBM® Tivoli®

IBM Tivoli Storage Manager as a Data Protection Solution

- Provides solutions for common data protection challenges
- Includes business and technical challenges and solution matrices
- Describes use of the Tivoli Storage Manager Toolkit

Mary Lovelace
Gerd Becker
Rosane Langnor
Mikael Lindstrom
Pia Nymann
Felipe Peres
Norbert Pott
Julien Sauvanet
Gokhan Yildirim

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM Tivoli Storage Manager as a Data Protection Solution

August 2014

SG24-8134-00

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

First Edition (August 2014)

This edition applies to Version 7, Release 1, of IBM Tivoli Storage Manager family of products.

Contents

Notices ix
Trademarks x

Preface xi
Authors xi
Now you can become a published author, too! xiv
Comments welcome xiv
Stay connected to IBM Redbooks xiv

Part 1. Today's environment and challenges 1

Chapter 1. Tivoli Storage Manager as a data protection solution 3
1.1 The changing world 4
1.1.1 The world of data protection has changed 4
1.1.2 Big data 5
1.1.3 Virtualization 5
1.2 Tivoli Storage Manager can help you in this new world 6
1.2.1 Product positioning 6
1.2.2 Tivoli Storage Manager is not a "tape only solution" 7
1.2.3 Tivoli Storage Manager is more than backup software 7
1.2.4 Tivoli Storage Manager is easy to administer 8
1.3 How this book can help you 9
1.3.1 Book structure 9
1.3.2 Summary 10

Part 2. Overview of Tivoli Storage Manager family of products 11

Chapter 2. Tivoli Storage Manager and the family of products 13
2.1 Tivoli Storage Manager evolution 14
2.1.1 Tivoli Storage Manager as a data protection solution 14
2.1.2 Tivoli Storage Manager Basic Edition 15
2.1.3 Tivoli Storage Manager Extended Edition 17
2.1.4 Tivoli Storage Manager Operations Center 18
2.2 Data protection family of products and components 20
2.2.1 Tivoli Storage Manager for Mail 20
2.2.2 Tivoli Storage Manager for Databases 21
2.2.3 Tivoli Storage Manager for Virtual Environments: Data Protection for VMware . . 21
2.2.4 AvePoint DocAve Backup and Restore 23
2.2.5 Tivoli Storage FlashCopy Manager 23
2.2.6 Tivoli Storage Manager for Enterprise Resource Planning 24
2.2.7 Tivoli Storage Manager for Space Management 25
2.2.8 Tivoli Storage Manager Hierarchical Storage Management for Windows 25
2.2.9 Tivoli Storage Manager for Storage Area Networks 25
2.2.10 Tivoli Storage Manager for System Backup and Recovery 26
2.2.11 Tivoli Storage Manager for z/OS Media 26
2.2.12 Tivoli Storage Manager FastBack 27
2.2.13 Tivoli Storage Manager FastBack for Microsoft Exchange 27
2.2.14 Tivoli Storage Manager FastBack for Bare Machine Recovery 27
2.2.15 Tivoli Storage Manager FastBack for Workstations 28

2.2.16 Tivoli Storage Manager FastBack Center . . . 28
2.2.17 Tivoli Storage Manager Suite for Unified Recovery . . . 28
2.3 Tivoli Storage Manager recent enhancements . . . 29
2.3.1 Tivoli Storage Manager V6.4 enhancements . . . 29
2.3.2 Tivoli Storage Manager V7.1 enhancements . . . 30
2.4 Conclusion . . . 32

Chapter 3. Data protection with Tivoli Storage Manager . . . 33
3.1 Protect the data at the client side . . . 34
3.1.1 Operating system protection . . . 34
3.1.2 Flat file backup . . . 36
3.1.3 Backup considerations . . . 39
3.1.4 Microsoft Volume Shadow Copy Service . . . 42
3.1.5 Microsoft cluster support wizard . . . 44
3.1.6 Virtual environment specifics . . . 44
3.1.7 Archive data . . . 45
3.1.8 Application or database backup . . . 46
3.1.9 Hierarchical storage management . . . 46
3.1.10 Protect data storage devices owning data . . . 46
3.2 Recover data from the client . . . 51
3.2.1 Restore data . . . 51
3.2.2 Retrieve . . . 52
3.2.3 Recall . . . 52
3.3 Data protection at server side . . . 52
3.3.1 Disaster recovery with traditional methods . . . 52
3.3.2 Disaster recovery with node replication . . . 57
3.3.3 Server-side data deduplication . . . 57
3.3.4 Storage hierarchy, storage pool features . . . 58
3.3.5 Tiered storage . . . 64
3.3.6 Advanced Tiered storage concept with virtual tape libraries . . . 66
3.3.7 Policy management . . . 67
3.3.8 Protecting data with Tivoli Storage Manager security . . . 70
3.3.9 Scripting for administration task automation . . . 72
3.3.10 Automatic client deployment . . . 72
3.3.11 Reporting analytics . . . 73
3.3.12 Storage area network (SAN) environments . . . 74
3.4 Tivoli Storage Manager server database management . . . 75
3.4.1 Data Protection with Tivoli Storage Manager server database . . . 75
3.4.2 Tivoli Storage Manager daily operation . . . 76

Part 3. Solution to challenges using Tivoli Storage Manager . . . 79

Chapter 4. Tivoli Storage Manager challenge matrix . . . 81
4.1 Solution matrix . . . 82
4.1.1 Setting up the matrix . . . 83
4.2 Reading the matrix . . . 85
4.2.1 Challenge: Data reduction . . . 85
4.2.2 Challenge: Data availability . . . 87
4.2.3 Challenge: Security . . . 88
4.2.4 Challenge: Disaster recovery . . . 89
4.2.5 Challenge: Virtualization . . . 90
4.3 Summary . . . 91

Chapter 5. Tivoli Storage Manager server protection . . . 93

5.1 Protecting the Tivoli Storage Manager server infrastructure 94
5.2 Protecting the Tivoli Storage Manager server database 95
5.2.1 Database mirroring 95
5.2.2 Database backup and restore 96
5.3 Additional server infrastructure protection 106
5.3.1 Protecting the Tivoli Storage Manager database log files 106
5.3.2 Protecting the volume history file 107
5.3.3 Protecting the device configuration file 107
5.3.4 Protecting the server options file 108
5.3.5 Protecting information about the database and recovery logs 109
5.3.6 Protecting the Secure Sockets Layer digital certificate file 109
5.3.7 Disaster Recovery Manager: Protecting the disaster recovery plan 110
5.4 Summary 110

Chapter 6. Tivoli Storage Manager Technologies and Solutions 111
6.1 Providing a set of data protection and restore tools 112
6.2 Disk-to-disk data protection solution using deduplication 113
6.2.1 Benefits deduplication provides 114
6.2.2 Solution architecture 114
6.2.3 Disk-to-disk data protection 116
6.2.4 Solution description 116
6.2.5 How Tivoli Storage Manager performs deduplication 117
6.2.6 Data reduction and data deduplication 117
6.2.7 Server-side and client-side deduplication 118
6.2.8 Prerequisites for configuring Tivoli Storage Manager deduplication 119
6.2.9 Exceptions to using Tivoli Storage Manager deduplication 119
6.2.10 Compare Tivoli Storage Manager deduplication and appliance deduplication . 120
6.2.11 Factors when deciding between Tivoli Storage Manager and appliance deduplication 121
6.2.12 Conditions for effective use of Tivoli Storage Manager deduplication 122
6.2.13 Traditional Tivoli Storage Manager architectures compared with deduplication architectures 123
6.2.14 Resource requirements for Tivoli Storage Manager deduplication 123
6.2.15 Database and log size requirements 123
6.2.16 Tivoli Storage Manager database log size estimation 124
6.2.17 Estimating capacity for deduplicated storage pools 125
6.2.18 Estimating storage pool capacity requirements 125
6.2.19 Database I/O requirements 125
6.2.20 Hardware requirements for Tivoli Storage Manager client deduplication 126
6.2.21 Processor 126
6.2.22 Memory 127
6.2.23 Solution guidelines 127
6.2.24 Deciding between client and server deduplication 127
6.3 Data protection solution using node replication feature 128
6.3.1 Node replication configuration considerations 129
6.3.2 Planning for incremental replication 129
6.3.3 Challenges and benefits the solution addresses 132
6.3.4 Solution architecture 132
6.3.5 Use scenarios 135
6.3.6 Hardware and software requirements 135
6.3.7 Database requirements for node replication 135
6.4 Tivoli Storage Manager together with ProtecTIER 135
6.5 Tivoli Storage Manager as a virtual appliance 136

6.5.1 Challenges the solution addresses 136
6.5.2 Solution architecture 137
6.5.3 Solution description 138
6.5.4 Use scenarios 138
6.5.5 Hardware and software requirements 139

Chapter 7. Protecting your data with Tivoli Storage Manager 141
7.1 Common virtualization challenges 142
7.2 Virtualization: VMware Data Protection solution 144
7.2.1 Sample solution and benefits 145
7.2.2 Solution architecture 146
7.2.3 Solution description 151
7.2.4 Use scenarios 154
7.2.5 Hardware and software requirements 158
7.2.6 Additional information 158
7.3 Virtualization: Hyper-V host-based backup 158
7.3.1 Summary of the client benefits 158
7.3.2 Solution architecture 158
7.3.3 Solution description 160
7.3.4 Use scenarios 161
7.3.5 Hardware and software requirements 166
7.3.6 References 167
7.4 Virtualization: In-guest backup 167
7.4.1 Summary of the client benefits 167
7.4.2 Solution architecture 167
7.4.3 Use scenarios 168
7.4.4 Software requirements 170
7.4.5 Tivoli Storage Manager as a comprehensive solution for a virtualized system. . 171
7.5 Application and email servers 171
7.5.1 Tivoli Storage Manager for Databases: Data Protection for Oracle 172
7.5.2 Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL 185
7.5.3 Tivoli Storage Manager for Mail: Data Protection for Domino 193
7.5.4 Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server 202
7.5.5 Tivoli Storage Manager for Enterprise Resource Planning 212
7.5.6 Tivoli Storage FlashCopy Manager 222
7.6 Big data: Structured, very large databases 231
7.6.1 LAN-free backups to tape 233
7.6.2 LAN-free backups to VTL 238
7.6.3 Protect very large databases with Tivoli Storage FlashCopy Manager. 245
7.7 Big data: Unstructured 259
7.7.1 Progressive incremental backups 260
7.7.2 Journal-based backups 264
7.7.3 Image backups 266
7.7.4 Traditional file system backups and RTO 268
7.7.5 NDMP backups 269
7.7.6 Off-host backup using snapshot differencing 277
7.7.7 GPFS, SONAS and V7000 Unified 283
7.7.8 SnapMirror to Tape 290
7.8 Remote office backup 293
7.8.1 Remote office client attached to data center server 294
7.8.2 Remote office server connected to data center server 295
7.8.3 Remote office server connected to data center server with 3-site copy 297

7.9 Employee workstation backup 298
7.9.1 Tivoli Storage Manager Backup-Archive Client. 300
7.9.2 Fastback for Workstations. 305

Appendix A. Practical approach to creating a Tivoli Storage Manager solution . . . 311
How to put solutions together. 312
Approaching the solution 312
Tivoli Storage Manager policy-based dynamic tiering 321

Appendix B. Hierarchical storage management (HSM) 323

Appendix C. VSS and Tivoli Storage Manager related product concepts 327

Appendix D. Additional material 331
Locating the Web material 331
Using the Web material 331
System requirements for downloading the Web material 331
Downloading and extracting the Web material 332

Related publications 333
IBM Redbooks 333
Other publications 333
Online resources 334
Help from IBM 334

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 grant 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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

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

COPYRIGHT LICENSE:

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

Trademarks

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

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

Active Cloud Engine®
AIX®
AS/400®
Cognos®
DB2®
developerWorks®
Domino®
DS8000®
FICON®
FlashCopy®
Global Technology Services®
GPFS™
HACMP™
HyperFactor®
IBM®
Informix®
Lotus®
Lotus Notes®
Notes®
OS/390®
PowerHA®
PowerVM®
ProtecTIER®
pureScale®
Redbooks®
Redbooks (logo) ®
RS/6000®
Storwize®
System Storage®
System x®
System z®
Tivoli®
Tivoli Storage Manager FastBack®
Workload Partitions Manager™
XIV®
z/OS®

The following terms are trademarks of other companies:

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

Ultrium, the LTO Logo and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

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

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

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

Preface

When you hear IBM® Tivoli® Storage Manager, the first thing that you typically think of is data backup. Tivoli Storage Manager is the premier storage management solution for mixed platform environments.

Businesses face a tidal wave of information and data that seems to increase daily. The ability to successfully and efficiently manage information and data has become imperative. The Tivoli Storage Manager family of products helps businesses successfully gain better control and efficiently manage the information tidal wave through significant enhancements in multiple facets of data protection.

Tivoli Storage Manager is a highly scalable and available data protection solution. It takes data protection scalability to the next level with a relational database, which is based on IBM DB2® technology. Greater availability is delivered through enhancements such as online, automated database reorganization.

This IBM Redbooks® publication describes the evolving set of data-protection challenges and how capabilities in Tivoli Storage Manager can best be used to address those challenges. This book is more than merely a description of new and changed functions in Tivoli Storage Manager; it is a guide to use for your overall data protection solution.

Authors

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

Mary Lovelace is a Consulting IT Specialist at the International Technical Support Organization. She has experience with IBM in large systems, storage and storage networking product education, system engineering and consultancy, and system support. She has written IBM Redbooks publications about Tivoli Storage Productivity Center, Tivoli Storage Manager, Scale Out Network Attached Storage, and IBM z/OS® storage products.

Gerd Becker is a Project Manager for EMPALIS GmbH, an IBM Business Partner in Germany. He has more than 30 years of IT experience, including over 15 years of experience with storage management products such as DFSMS and Tivoli Storage Manager. His areas of expertise include IBM Tivoli Storage Manager implementation projects and education at customer sites, including mainframe environments (IBM OS/390®, VSE, VM, and Linux for zSeries). He holds several certifications, including technical and sales, and is an IBM Tivoli Certified Instructor. He has developed and taught several storage classes for IBM Education Services in Germany, Switzerland, and Austria, and for qSkills in Nuernberg. He has been Chairman of the Guide Share Europe (GSE) Storage - User Group for more than six years. Gerd is an author of these IBM Redbooks publications: *IBM Tivoli Storage Manager Version 5.3 Technical Guide,* SG24-6638; and *Certification Guide Series: IBM Tivoli Storage Manager V6.1,* SG24-7781. Gerd participated in the beta test for Tivoli Storage Manager from Version 5.5 through 6.4, and participated in the EAP Program for Tivoli Storage Manager Version 7.1. Gerd is a member of the Tivoli Storage Manager Advisory Council.

Rosane Langnor is an IBM Certified IT Specialist, working in Storage Lab Services in Brazil. She has fifteen years of IT experience and has been working with IBM since 1999. She has been working with Tivoli Storage Manager solutions since 2006, mainly as a global resource for the IBM Global Technology Services® organization, developing, implementing and supporting backup solutions. Rosane is an IBM Certified Deployment Professional for Tivoli Storage Manager and has authored several Redbooks publications about Tivoli Storage Manager: *IBM Tivoli Storage Manager in a Clustered Environment*, SG25-6679; and *Deployment Guide Series: IBM Tivoli Storage Manager Express,* SG24-7033.

Mikael Lindstrom is an IBM Certified Consulting IT Specialist, working as a Nordic Lead Engineer and Storage Architect in the Global Storage Specialty Service Area at Delivery Technology & Engineering. Mikael is developing, executing, and implementing Global technology strategy and architecture for Global Technology Services with focus on Tivoli Storage Manager solutions and best practices. He has 13 years of IT experience and has worked with IBM since 2006 in various leadership roles. Mikael has authored several Redbooks publications and also published several IBM Tivoli Storage Manager white papers at IBM developerWorks®.

Pia Nymann is an IBM Certified Senior IT Specialist working in Systems Lab Services in Denmark, which is part of IBM Systems Technology Group. She has 15 years of experience in the IT industry and has several years of Tivoli Storage software experience including designing and implementing Tivoli Storage Manager Solutions. This includes working on various platforms and applications. She used her skills during that time by supporting and educating many clients about the Tivoli Storage Manager software. Pia has worked with various areas of the storage management discipline, including Tivoli Productivity Center, storage analysis, and service management. This is her first Redbooks residency.

Felipe Peres is an Advisory Product Service Specialist for IBM Technical Support Services in Brazil. He has four years of experience with Tivoli Storage Manager working as L1 Support. His main areas of expertise include supporting and managing Tivoli Storage Manager environments and all components, and Tivoli Data Protection and Fastback components. In addition to working as remote support with Tivoli Storage Manager L1, his other activities include local support, implementing Tivoli Storage Manager solutions, and creating and teaching Tivoli Storage Manager workshops. He has more than 10 Tivoli Storage Manager Certifications and he is an IBM Certified Solution Advisor for Tivoli Storage Solutions V3, IBM Certified ADP for IBM Service Management Tivoli Storage Management V4, Tivoli Storage Productivity Center Deployment Professional, and Fastback Deployment Professional. He holds a degree in System Analysis, with a focus on data security, from São Paulo State Technological College (FATEC).

Norbert Pott is an IBM Tivoli Storage Manager Support Specialist in Germany. He works for the Tivoli Storage Manager back-end support team, providing support to customers worldwide. He has 32 years of experience with IBM, over 23 years of experience in IT, and more than 16 years of experience with the Tivoli Storage Manager product, starting with ADSM Version 2.1.5. His areas of expertise include Tivoli Storage Manager client development skill and in-depth knowledge of problem determination. He is an author of several Redbooks publications: *IBM Tivoli Storage Manager Version 5.3 Technical Workshop Presentation Guide*, SG24-6774; *IBM Tivoli Storage Manager Implementation Guide*, SG24-5416; *IBM Tivoli Storage Management Concepts*, SG24-4877; *IBM Tivoli Storage Manager Versions 5.4 and 5.5 Technical Guide*, SG24-7447; *Tivoli Storage Manager V6.1 Technical Guide*, SG24-7718; and *Certification Study Guide Series: Tivoli Storage Manager V6.1*, SG24-7781.

Julien Sauvanet is an IBM Certified Expert IT Specialist, working in the French IBM Global Technology Services organization. He has more than 10 years of experience with Tivoli Storage Manager and other related storage products, and he is an IBM Certified Deployment Professional for Tivoli Storage Manager. His areas of expertise include the Tivoli Storage Manager product family, and enterprise disk and tape subsystems. He is also a subject matter expert in system and storage virtualization (such as IBM PowerVM®, VMware, IBM Storwize® and IBM ProtecTIER® for storage products). He is an author of this Redbooks publication: *Tivoli Storage Manager for Virtual Environments - Data Protection for VMware Deployment Guide,* REDP-4991. He also is an author of *IBM Tivoli Storage Manager server and datamover running in a Virtual Environment Sample Solution* found in the developerWorks Tivoli Storage Manager Wiki.

Gokhan Yildirim is a Principal Specialist of Storage Team. He works for Garanti Bank, one of the largest banks in Turkey. He has 28 years of IT experience and has worked for Garanti Bank since 1997. Gokhan started his career as a Programmer for System 38, IBM AS/400®, i Series RPG, Mainframe PL/1, PC Systems Visual Basic, and C. He has more than ten years of Tivoli Storage Manager Administration knowledge with an IBM Certified Administrator for Tivoli Storage Manager certification, and has expertise in other related products like tape systems, disk systems, and virtual tape libraries (VTLs).

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

Ella Buslovich
Larry Coyne
Diane Sherman
Erica Wazewski
Debbie Willmschen
International Technical Support Organization

Jason Basler
Dave Cannon
Colin Dawson
Fraser Macintosh
David McClelland
Kathy Mitton
Dominic Mueller-Wicke
Ian Smith
Jim Smith
Daniel Wolfe
Justin Youngblood
Chris Zaremba
IBM Tivoli Storage Manager

Tommy Hueber
SWG Lab Services Managing Consultant

Stephane Criachi
Concat AG

Richard Spurlock
Cobalt Iron

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

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 publications in one of the following ways:

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

  ibm.com/redbooks

- Send your comments in an email 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

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Part 1

Today's environment and challenges

In this part, we describe data growth and the challenges it creates.

This part contains the following chapter:

- Chapter 1, "Tivoli Storage Manager as a data protection solution" on page 3

1

Tivoli Storage Manager as a data protection solution

Today, the changes in our world include the instrumentation, interconnectedness, and intelligence of our environments. Such changes combine to produce a massive glut of new useful information, from new and existing sources, with new ways to use it. These pressures typify a few of the storage challenges that we dealt with for some time now, but on a whole new scale.

In this chapter, we look at the type of data that is generated today and how Tivoli Storage Manager is used as a data protection solution to manage that data.

1.1 The changing world

Organizations are storing and using more data than ever before. Data volume is growing exponentially, and government regulations and competitive pressures are increasing, forcing organizations to retain more data for longer periods of time.

As data grows, storage administrators are challenged to complete backup operations within established backup windows. More data in the backup system means a longer amount of time is necessary to recover when something goes wrong, increasing the risks.

One solution to all this data growth has been to buy more storage, as the cost of the storage itself decreases over time. But the cost of housing, powering, cooling, and managing these devices is exploding.

And, of course, businesses are always changing. Storage administrators need to adapt to any number of changes in their environments, from bringing new technologies, applications and data sources online, to complying with new corporate and government data management mandates.

1.1.1 The world of data protection has changed

For many organizations, protecting and retaining data is *table stakes*, a function so basic and important that business cannot move forward without it. On a smarter planet, where the gathering and use of data is increasingly instrumented, interconnected, and intelligent, the loss of data and interruptions to data-driven processes are not options for a growing percentage of organizations. Data is becoming more valuable, so organizations want to keep more of it, and for longer periods of time.

Greater efficiency and effectiveness in core data management functions (backup, archiving and ensuring continuous operations) can have a direct and positive impact on the business.

Expectations and requirements for data protection have changed dramatically in recent years:

- Where it was once possible to operate with backup windows that lasted eight hours or more, those windows have shrunk or even disappeared.
- Where backups used to be daily (and the loss of 24 hours worth of data was acceptable) backups for core applications are now much more frequent, with far less tolerance for data loss.
- The IT infrastructure has become much more complex and distributed, with different types of operating systems and applications. Where all data once was treated the same, many users and applications now have unique restore requirements.

1.1.2 Big data

Every day, we create 2.5 quintillion bytes of data, so much that 90% of the data in the world today has been created in the last two years alone. This data comes from everywhere: sensors used to gather climate information, posts to social media sites, digital pictures and videos, purchase transaction records, and cell phone GPS signals to name a few. This data is big data, also shown visually in Figure 1-1.

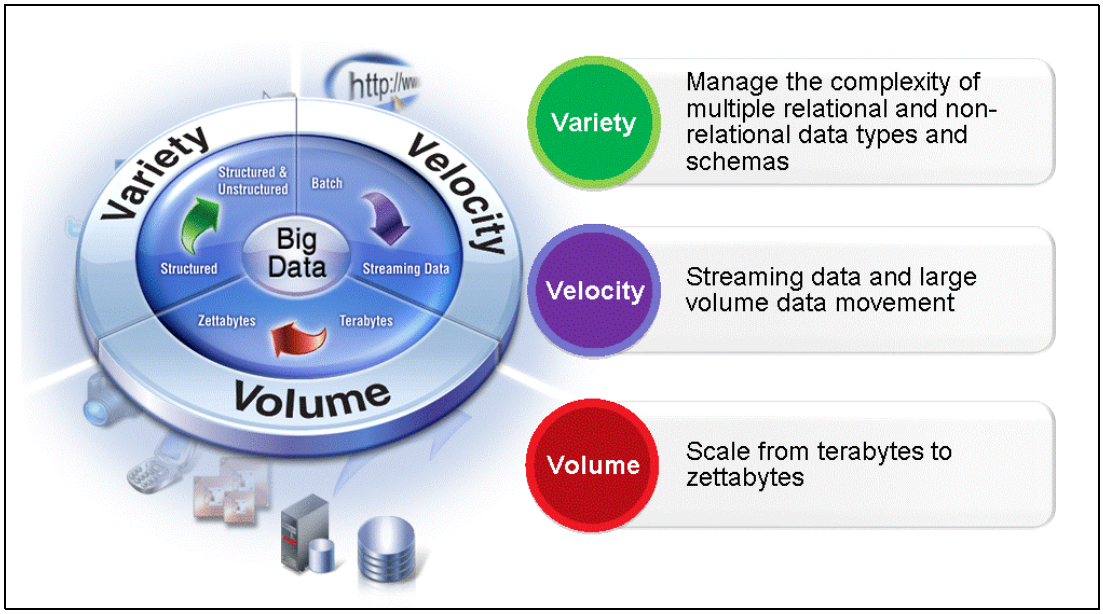

Figure 1-1 IBM characterizes big data by its volume, velocity and variety

In this book, we look at the challenges the three V's impose on the backup/archive environment. We cover the ways that you can use Tivoli Storage Manager to provide effective and efficient protection of your data (protecting the right data at the right times), helping you to not increase the amount of storage you have.

1.1.3 Virtualization

More enterprises are using virtualization as a means to reduce cost and improve efficiency. Virtualization is taking place in many areas of the IT Infrastructure such as servers, storage, and networking. The backups also need to follow this trend and administrators need to change to a smarter way of doing these backups.

In this book, we also cover how Tivoli Storage Manager features can be used to run smarter backups of your smarter infrastructure

1.2 Tivoli Storage Manager can help you in this new world

Today's data storage managers are being asked to work miracles: back up massive amounts of data without impacting operations, restore lost data almost instantaneously, protect different sets of data in accordance with different security and retention policies, and do it all on a shrinking budget.

You ask yourself how you can improve your backup process through these ways:

- Improved performance
- Reduced time of execution
- Improved data integrity
- Reduced physical space to store backup data
- Reduced LAN/WAN bandwidth consumption

A huge contributor to data growth is the repeated duplication of large amounts of data every time you perform a full backup. In an IBM holistic approach, one option is to avoid data growth from unnecessary data duplication, by only backing up data that has changed since the last backup. Another option is to determine what types of data you have and categorize it so that you can manage it most effectively, by moving less frequently accessed data to lower-cost tiers of storage, and by automatically moving older data to the correct tier of storage and deleting data that you no longer need or want. This can shorten your backup cycles and improve application performance. Finally, you can compress and deduplicate the data that you put into your data protection and retention systems.

1.2.1 Product positioning

Tivoli Storage Manager can help you build a smarter storage management infrastructure that can help you to cope with all of these challenges. It will help you in these ways:

- Reduce your capital and operational costs by reducing your storage requirements.
- Improve your application availability and service levels by reducing downtime.
- Mitigate the risks associated with losing data in a rapidly changing environment.

The goal is to protect the organization's data from failures. Protection is provided across a wide range of operating systems and applications, running on hardware as different as notebooks and mainframes. Figure 1-2 on page 7 shows the IBM family of products that helps you cope with today's challenges on storage management. These are covered in more details in the next chapters.

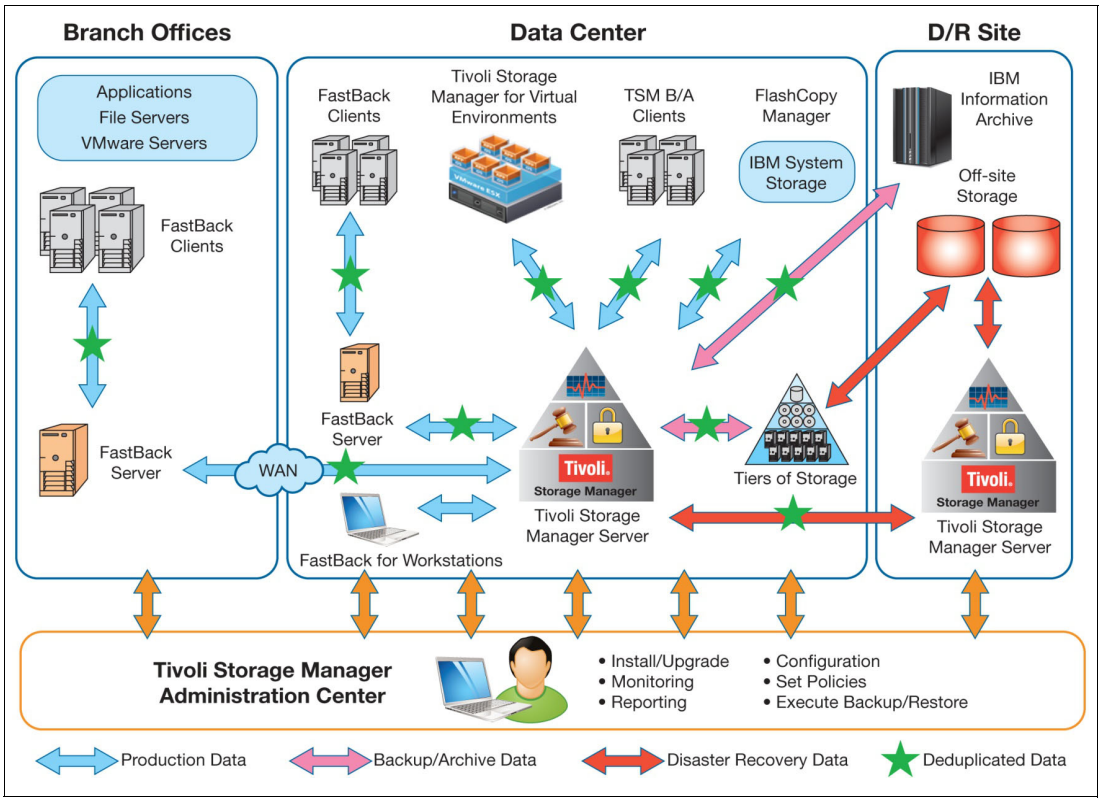

Figure 1-2 IBM unified recovery management portfolio

1.2.2 Tivoli Storage Manager is not a “tape only solution”

Much is being said about tapeless backup solutions as the trend for data protection. This is not a new idea. Since its first version, Tivoli Storage Manager has used disk as the primary storage for backups, together with tapes.

More often, as disk storage technology evolves, other solutions become possible and are being incorporated to the product. Deduplication and snapshots are examples of Tivoli Storage Manager features that take advantage of the disk subsystem architecture.

However, tapes cannot be discarded. Tape technology continues to evolve, increasing its speed and capacity. It is still the storage device suitable for many applications and business requirements. It is best suited for the backup of large sequential files, such as big databases and digital images, and to store data that needs to be kept for a long period of time to meet regulatory requirements.

We discuss several scenarios of how Tivoli Storage Manager uses disk and tape for a better solution.

1.2.3 Tivoli Storage Manager is more than backup software

Tivoli Storage Manager is the premier storage management solution for mixed platform environments. IBM Tivoli Storage Manager helps businesses manage and control their business data by delivering a single point of control and administration for data backup and recovery.

This advanced, highly scalable product helps increase the efficiency of your IT operations and helps cut costs related to storage management, by providing a wide range of data protection, recovery management, movement, retention, reporting, and monitoring capabilities using policy-based automation.

Businesses face a tidal wave of information and data that seems to increase daily. The ability to successfully and efficiently manage information and data has become imperative. The IBM Tivoli Storage Manager family of products helps businesses successfully gain better control and efficiently manage the information tidal wave with significant enhancements in multiple facets of data protection.

From its inception, Tivoli Storage Manager has been a highly scalable and available data protection solution. Tivoli Storage Manager takes data protection scalability to the next level with a new relational database, based on IBM DB2 technology. Greater availability is delivered through enhancements such as online, automated database reorganization. In addition, the increased scalability and the ability to use the latest in server technology helps deliver increased performance of backup and recovery processes for business critical recovery point objective (RPO) and recovery time objective needs

1.2.4 Tivoli Storage Manager is easy to administer

Tivoli Storage Manager has several tools that help with administration of it. You can configure and manage IBM Tivoli Storage Manager server through web interfaces and with wizards to help guide you through common configuration tasks and advanced management tasks.

- You need to log in only once to access multiple Tivoli Storage Manager servers from a single interface.
- You can easily monitor the health of your storage environment.
- You can filter and sort storage objects, such as client nodes and library volumes.
- You can use wizards to more easily perform complex tasks, such as these:
  - Creating schedules to perform client node and administrative operations
  - Creating a server maintenance script to perform database and storage pool backup, migration, expiration, and reclamation
  - Configuring storage devices; a comprehensive wizard helps you create a library, add drives, check in media volumes, and create storage pools
  - Configuring Version 6.1 or later server on the local or a remote UNIX system
  - Scheduling a client deployment for client nodes in multiple domains

Another important tool is the IBM Tivoli Storage Manager monitoring and reporting feature that can be installed on IBM AIX®, Linux on IBM System x®, and Microsoft Windows platforms, but can monitor a Tivoli Storage Manager server running on any platform.

You can view the historical reports to determine if any issues or trends need attention, such as uncontrolled growth over time. You can also view workspaces that are being monitored to see the Tivoli Storage Manager server IDs, database size, agent status, client node status, scheduled events, and so on.

The reporting component, sometimes referred to as Tivoli Common Reporting, reports on the retrieved historical data. IBM Tivoli Monitoring acts as a monitoring application that provides workspaces for you to monitor near real-time information.

Note: If an IBM Tivoli Monitoring infrastructure exists, you can reuse Tivoli Storage Manager licences for IBM Tivoli Monitoring Tivoli Storage Manager agents.

1.3 How this book can help you

We want this book to be a valuable source of information for you to plan your data protection environment. The next chapters in the book guide you through Tivoli Storage Manager features and components that are part of the solutions presented in the last part or the book.

1.3.1 Book structure

The book is divided into three parts.

- Part 1, “Today’s environment and challenges” on page 1 The first part of this book contains the chapter you are reading now.
  - Chapter 1, “Tivoli Storage Manager as a data protection solution” on page 3 discusses the type of data that is generated today and how Tivoli Storage Manager is used as a data protection solution to manage that data. It is an introduction to the remainder of the book.
- Part 2, “Overview of Tivoli Storage Manager family of products” on page 11 consists of two chapters. It provides a brief description of the Tivoli Storage Manager family of products and how the Tivoli Storage Manager family has evolved to continue meeting your business needs.
  - Chapter 2, “Tivoli Storage Manager and the family of products” on page 13 covers the history of the Tivoli Storage Manager and its family of products. A brief summary of the enhancements in Tivoli Storage Manager V6.4 and Tivoli Storage Manager V7.1 is included. If you are familiar with the product you might want to only review the release updates in 2.3.1, “Tivoli Storage Manager V6.4 enhancements” on page 29 and 2.3.2, “Tivoli Storage Manager V7.1 enhancements” on page 30.
  - Chapter 3, “Data protection with Tivoli Storage Manager” on page 33 provides a brief summary of techniques and concepts that Tivoli Storage Manager product provides. If you are familiar with the product, you might want to go directly to Chapter 4, “Tivoli Storage Manager challenge matrix” on page 81.
- Part 3, “Solution to challenges using Tivoli Storage Manager” on page 79 has several chapters that contain the matrices and solutions provided by Tivoli Storage Manager family of products that meet various data protection challenges of today’s world.
  - Chapter 4, “Tivoli Storage Manager challenge matrix” on page 81 introduces the solution matrix that is central to this book. We explain how you can use the matrix to find the solution for the data protection challenge you must meet. We use the matrix to reference you from this chapter to solutions and Tivoli Storage Manager toolkit features that we document in this book.
  - Chapter 5, “Tivoli Storage Manager server protection” on page 93 discusses how protecting your Tivoli Storage Manager server infrastructure involves more than just taking a backup copy of your database. When implementing a Tivoli Storage Manager data protection solution, one of the most important issues is server infrastructure protection. In addition to the storage pool volumes, the server database and its infrastructure files must be protected against corruption or loss.

– Chapter 6, “Tivoli Storage Manager Technologies and Solutions” on page 111 discusses several Tivoli Storage Manager data protection solutions. One of the great features Tivoli Storage Manager offers is data deduplication. We explain how to implement a Tivoli Storage Manager solution using data deduplication. Next we describe how, through a disk-to-disk data protection solution using node replication, to enlarge the architecture by adding new Tivoli Storage Manager components, thus new functionality. Finally, we discuss Tivoli Storage Manager in a virtual environment, and how you can implement a fully virtualized data protection solution based on Tivoli Storage Manager products.

– Chapter 7, “Protecting your data with Tivoli Storage Manager” on page 141 provides various data protection solutions based on common challenges in today’s data protection world. The solutions are based on the components available in the Tivoli Storage Manager toolkit and reflect the information in the challenges matrix the team created, that is the core of the information in this book.

1.3.2 Summary

Our goal for the reader of this book is that, when you have gone through the material in this book and studied the solutions we describe, you will be able to build a data protection solution that meets your unique requirements by using the Tivoli Storage Manager family of products.

As you go through the material in this book, keep in mind the following product highlights, based on Tivoli Storage Manager 6.3, described in the ESG Lab Validation Report[1]:

- Tivoli Storage Manager progressive incremental backup, combined with client- and source-side data deduplication technology provided an impressive 95% data reduction factor (more than 19:1) over just 11 days of backups.
- DB2 continues to provide an enterprise-class, scalable back end for Tivoli Storage Manager. ESG Lab confirmed that Tivoli Storage Manager can scale to more than four billion objects under management while providing higher performance and more functionality.
- Tivoli Storage Manager offers hot Tivoli Storage Manager server disaster recovery, and demonstrated the ability to replicate deduplicated data sets from a live Tivoli Storage Manager server to a hot standby server. Tivoli Storage Manager 6.3 provided seamless and transparent restores for clients after failover.
- Tivoli Storage Manager showed many improvements to ease of use and management. The Tivoli Integrated Portal provided a common repository for all Tivoli Storage Manager interfaces, with a “single pane of glass” for management of an enterprise Tivoli Storage Manager environment. ESG Lab also tested automated client deployment, with no scripting or manual commands needed.

When you decide to use the Tivoli Storage Manager product, you can rely on 20 years of data protection technologies that, with current versions, allow you to implement data protection solutions for your current and future requirements.

[1] http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=WH&infotype=SA&appname=SWGE_TI_SG_USEN&htmlfid=TIL14047USEN&attachment=TIL14047USEN.PDF (*IBM Tivoli Storage Manager 6.3: Deduplication, Remote Site Recovery, and Scalability*, November 2012)

Part 2

Overview of Tivoli Storage Manager family of products

In this part of the book, we provide an overview of the Tivoli Storage Manager and the Tivoli Storage Manager family of products. Tivoli Storage Manager products provide backup, archive, recovery, space management, database and application protection, and bare machine recovery and disaster recovery capabilities. We look at how the Tivoli Storage Manager family has evolved to continue to meet your business needs. A summary of the features and functions, introduced in Tivoli Storage Manager V6.4 and V7.1, is included.

This part contains the following chapters:

- Chapter 2, “Tivoli Storage Manager and the family of products” on page 13
- Chapter 3, “Data protection with Tivoli Storage Manager” on page 33

2

Tivoli Storage Manager and the family of products

In this chapter, we look at the history of the IBM Tivoli Storage Manager and its family of products. We discuss how these products meet your business needs. A brief summary of the enhancements in Tivoli Storage Manager V6.4 and Tivoli Storage Manager V7.1 is included.

2.1 Tivoli Storage Manager evolution

The current Tivoli Storage Manager began as ADSTAR Distributed Storage Manager (ADSM). It has evolved in response to new technologies in the market. Figure 2-1 shows the release timeline for the various versions of ADSM and then Tivoli Storage Manager, up to the current version.

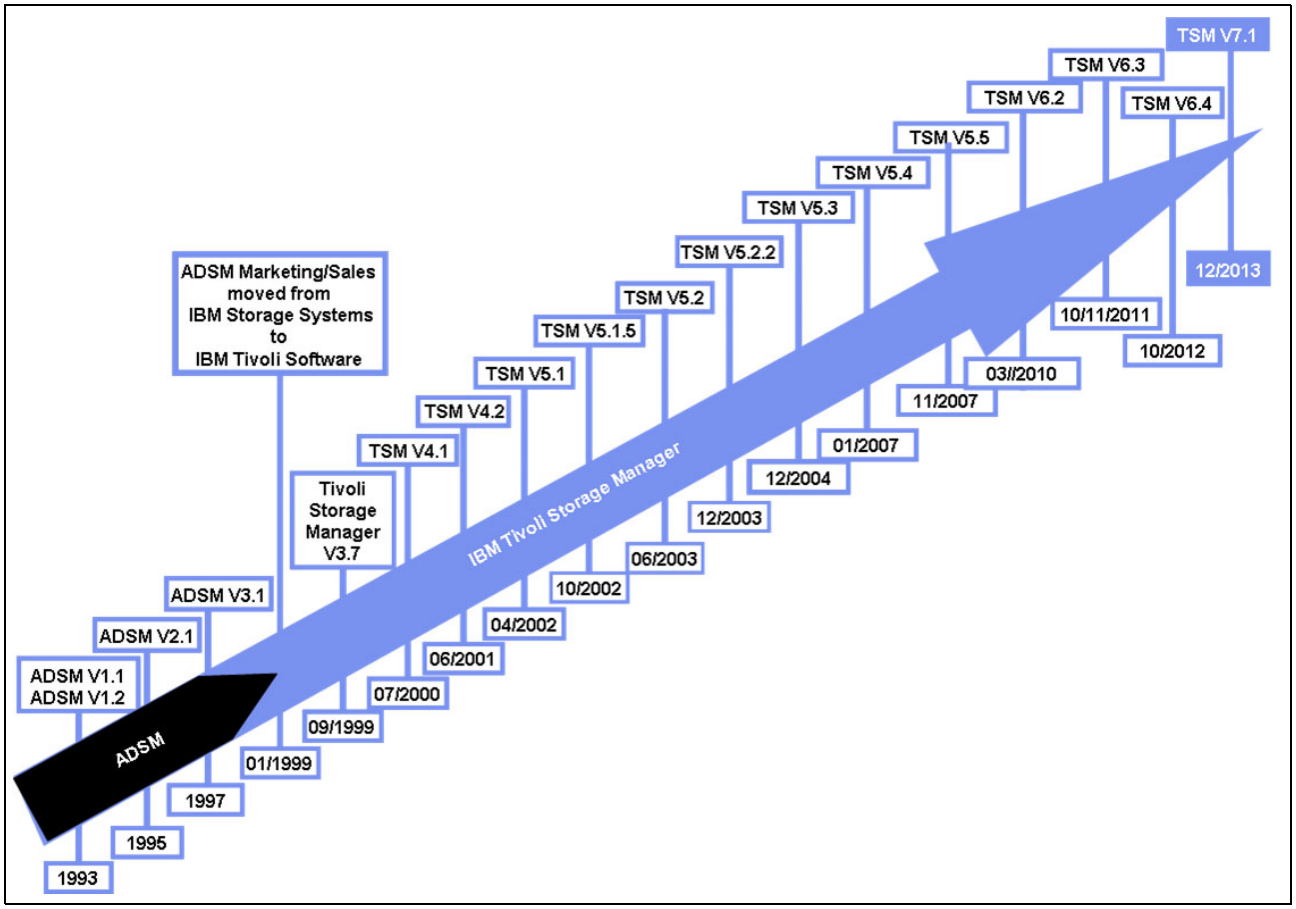

Figure 2-1 Tivoli Storage Manager evolution timeline

2.1.1 Tivoli Storage Manager as a data protection solution

Tivoli Storage Manager protects the data of an organization against hardware failures and other errors by storing backup and archive copies of data in online or offline storage. It can scale to protect thousands of computers, ranging from notebooks to mainframes, running various operating systems, which are connected from WAN, LAN, or SAN. 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. Optional modules enable business-critical applications that must run 24x7x365 to use Tivoli Storage Manager centralized data protection with no interruption to their services.

Tivoli Storage Manager provides centralized, automated data protection to help reduce the risks that are associated with data loss. This highly scalable software helps you manage more data with less infrastructure and simplified administration. You can save money, improve

service levels, and comply with data retention regulations. Tivoli Storage Manager includes the following Data protection features:

- Backup and recovery management

  Tivoli Storage Manager automates data backup, restore, and archive-retrieve functions. It centralizes data management operations. A single administrator interface allows configuration, monitoring, reporting, and backup and recovery execution across the entire complex IT environment.

- Hierarchical storage management

  Hierarchical storage management enables policy-based management of file backup and archiving, with automatic migration of data between tiers of storage. This feature reduces storage media requirements and administrative costs associated with managing data.

- Scalability

  A single Tivoli Storage Manager server manages as many as four billion data objects.

- Advanced data reduction

  This feature combines progressive incremental backup, source and target data deduplication, compression, and tape management. This advanced technology cuts data storage costs, decreases environmental requirements and simplifies administration.

2.1.2 Tivoli Storage Manager Basic Edition

Tivoli Storage Manager Basic Edition is the flagship product in the Tivoli Storage Manager family. It contains a rich set of features and provides the core functions of backup, recovery, and archive management:

- Progressive incremental backup methodology

  Saves time and storage space by backing up only new and modified files. The progressive incremental backup feature uses the Tivoli Storage Manager relational database to track data wherever it is stored, delivering a direct one-step file restore. Progressive incremental 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 library sharing

  Enables multiple Tivoli Storage Manager servers to use the same tape library and drives. This feature optimizes tape hardware asset utilization.

- Enterprise administration

  Simplifies centralized control across multiple Tivoli Storage Manager implementations without sacrificing network performance.

- Node replication

  Data belonging to backup-archive client nodes is initially stored on source replication servers. The data is then replicated to target replication servers. Node replication provides the ability to incrementally replicate a node's data to a remote target Tivoli Storage Manager server for disaster recovery purposes. True incremental replication only replicates directories and files that do not exist on the target server. It allows for tiering and grouping of recovery data classes, and deletes data on the target server that has been deleted on the source server. It can recover client data directly from the hot standby server (unlike virtual volumes). Node replication can be performed from both tape and disk devices. It is not dependent on any specific storage type.

To help reduce WAN bandwidth requirements, the replicated data can be sent deduplicated, and the transfer can be scheduled. Multiple Tivoli Storage Manager servers can be replicated to a single target server.

Additional information and guidelines for node replication are at the following location:

https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/Guidelines%20for%20node%20replication

- Clustering

  Tivoli Storage Manager includes enhanced support for IBM Power High Availability Cluster, Microsoft Cluster for Windows, and VERITAS Cluster Services (VCS) on Windows.

  Tivoli Storage Manager 6.1 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 backup and restore

  An optional module for Tivoli Storage Manager effectively uses storage area network (SAN) environments by moving data transfers from the communication network to a SAN. Communication bandwidth availability is therefore improved, increasing service levels for clients.

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

- Hierarchical Storage Management

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

- Data deduplication

  Data deduplication is a method of eliminating redundant data in sequential-access disk primary, copy, and active-data storage pools. One unique instance of the data is retained on storage media, and redundant data is replaced with a pointer to the unique data copy. In addition to reducing overall storage costs, another goal of deduplication is to reduce the overall amount of time that is required to retrieve data by letting you store more data on disk, rather than on tape. Tivoli Storage Manager V6.2 and later releases offer server-side and client-side data deduplication.

- Library and device support

  Tivoli Storage Manager supports libraries with up to three tape drives and up to 40 cartridge capacity. Larger libraries can be accommodated, either with only three devices and 40 slots enabled, or by upgrading to Tivoli Storage Manager Extended Edition.

  You can find more information about IBM Tivoli Storage Manager at this website:

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

2.1.3 Tivoli Storage Manager Extended Edition

Tivoli Storage Manager Extended Edition is an enhanced data protection software for expediting disaster recovery preparation. This edition provides the following functions:

- Enhanced backup, restore and archive capabilities.
- Exceptional scalability and performance with advanced disaster recovery functionality.
- Data protection and disaster recovery functionality for more than 500 disk, tape, and virtual tape storage devices from many vendors. There are no limitations on the number of drives or slots.
- Extended library and drive support and supports larger tape libraries.
- Disaster recovery planning capability.
- Network Data Management Protocol (NDMP) data protection for network-attached storage (NAS) filers.

You can find more information at this website:

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

Tivoli Storage Manager Disaster Recovery Manager

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

Disaster Recovery Manager offers various options to configure, control, and automatically generate a disaster recovery plan file. The plan contains the information, scripts, and procedures required to automate restoration and help ensure quick and reliable 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 Disaster Recovery Manager is the ability to track media in various states, such as onsite, in transit, or in a vault. The media movement features of Disaster Recovery Manager assist greatly with the daily tasks of sending disaster recovery media offsite, and receiving expired media onsite for reuse. With these features, the system administrator can quickly locate all available copies of data.

Disaster Recovery Manager helps maintain business continuity by taking care of these functions:

- Establishing and helping to automate a thorough server disaster recovery plan, 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 offsite 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 Disaster Recovery Manager, you can recover at an alternate site, on a replacement system with another hardware configuration

2.1.4 Tivoli Storage Manager Operations Center

IBM Tivoli Storage Manager Operations Center is a graphical user interface (GUI), introduced in IBM Tivoli Storage Manager v6.4.1 and improved in version 7.1 with new features (as shown in Figure 2-2). It provides an advanced visualization dashboard, built-in analytics, and integrated workflow automation features that dramatically simplify backup administration.

Figure 2-2 Operations Center welcome page

The web-based dashboard (Figure 2-3 on page 19) aggregates information about the backup environment into a unified display. This information includes the systems, virtual machines and applications being backed up, backup servers, and storage repositories.

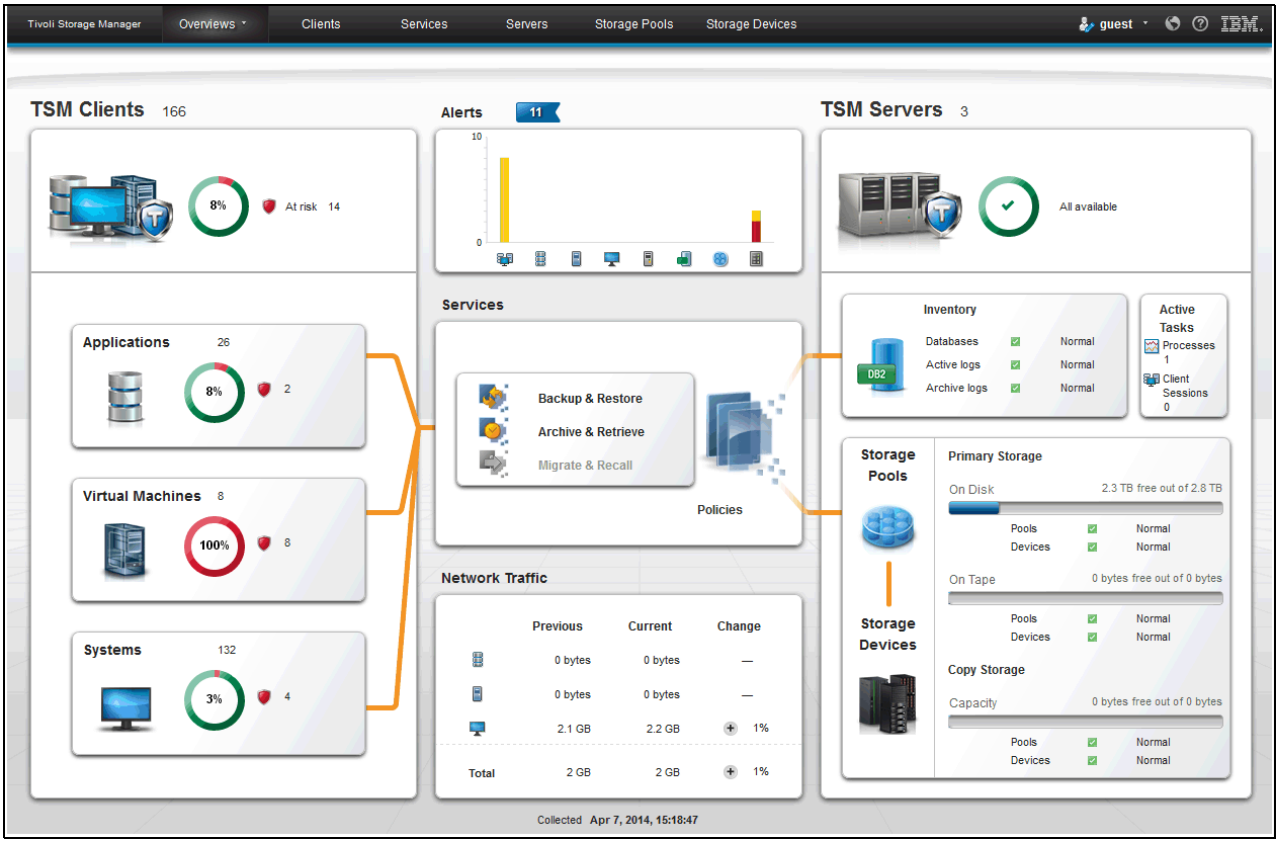

Figure 2-3 Operation Center dash board

With the visibility that Tivoli Storage Manager Operations Center provides, you can more easily manage backup tasks running across multiple backup servers. You can see how well those tasks were carried out. The health and capacity of backup systems, and also task fulfillment, is evident at a glance. This enables you to see whether an overnight backup went smoothly, or which backups require investigation and possible corrective action. The web-based GUI can help you more easily accomplish all of this anytime, anywhere (for example, from a desktop computer, a laptop, or even a browser-enabled mobile device).

Tivoli Storage Manager Operations Center enables administrators to view the backup history for individual systems, focusing on key error messages. It proactively detects problems and brings them to the attention of the backup administrator through the GUI. It also provides essential information needed for resolution.

Tivoli Storage Manager Operations Center enables almost any IT team member to handle everyday backup administration tasks, presenting information in an easy-to-understand format. In many cases, non-specialists can investigate and resolve backup issues, which enables these tasks to be accomplished faster and frees experts to focus on new projects.

The solution includes a built-in service for assigning problems among support teams, significantly enhancing team collaboration in order to speed, simplify and lower the cost of backup problem resolution. For more complex issues that still require an expert backup administrator to solve, Tivoli Storage Manager Operations Center can help drive a quick and effective resolution process. A pop-up command builder can be used from the web dashboard. This command builder enables experts to take necessary corrective actions quickly, then updates the dashboard to reflect the fix status. This helps improve backup, archive, and restore problem-resolution times, and to reduce complexity.

2.2 Data protection family of products and components

Figure 2-4 shows the current Tivoli Storage Manager product family. Each component is described in this section.

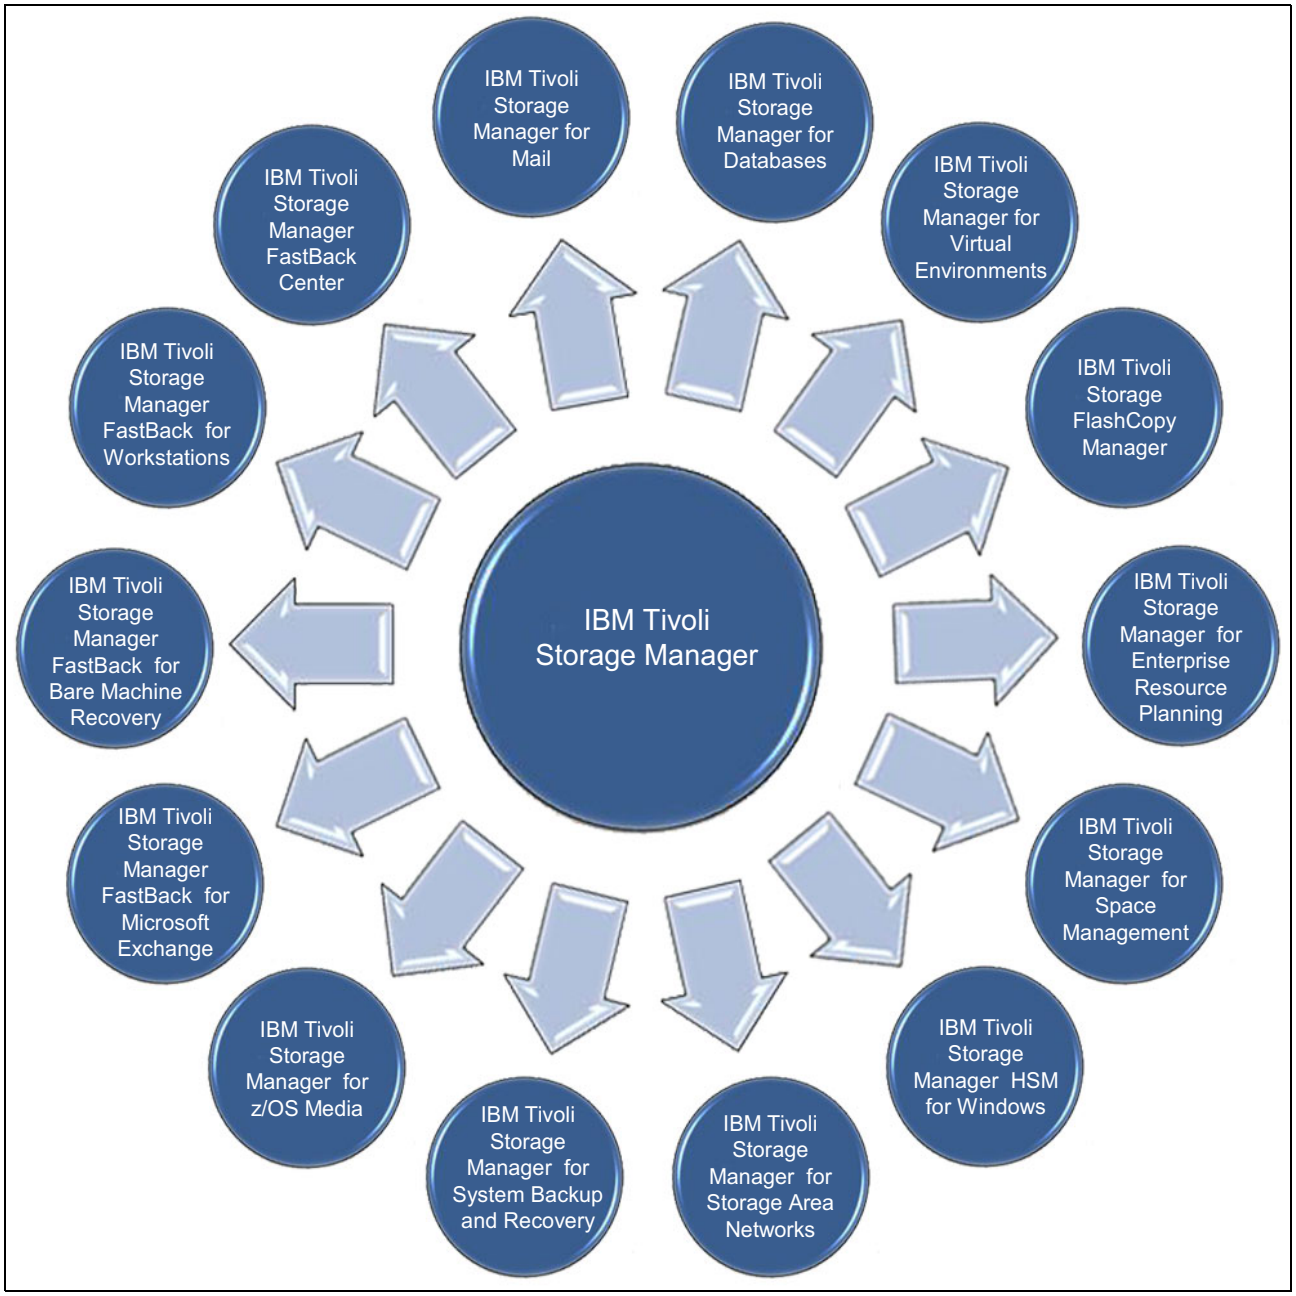

Figure 2-4 Tivoli Storage Manager product family

2.2.1 Tivoli Storage Manager for Mail

Tivoli Storage Manager for Mail protects data on email servers running IBM Lotus® Domino® or Microsoft Exchange. This software module for IBM Tivoli Storage Manager enables data protection of your mail databases while they are online. It automates data protection, enables hot backups without shutting down the application and improves data restore performance.

APIs provided by email application vendors are used to perform online hot backups and restores. Microsoft backup APIs are used to create a copy of the Exchange server storage group databases and the associated transaction logs. Tivoli Storage Manager for Mail can produce different types of backups specified by the Microsoft backup APIs: full, incremental, differential, copy, and database copy backups.

Features of Tivoli Storage Manager for Mail are as follows:

- Provides item-level recovery of Microsoft Exchange email objects. Search capabilities allow you to recover a user's mailbox or individual items within the mailbox.
- Helps to protect the growing amount of new and changing data that should be securely backed up to help maintain continuous availability.
- Exploits the Lotus Domino transaction logging feature, which enables the capture and logging of database changes, resulting in less-frequent full backups.

2.2.2 Tivoli Storage Manager for Databases

Tivoli Storage Manager for Databases helps to protect Oracle and Microsoft SQL data. It provides greater backup and restore options for Oracle and SQL servers by taking advantage of Oracle and SQL Server backup-certified utilities and interfaces. Applies IBM Tivoli Storage Manager automated data protection tasks to running database servers. The Oracle Automated Storage Management (ASM) feature available in Oracle 10g and later releases is supported.

Tivoli Storage Manager for Databases includes Data Protection for Oracle, which interfaces with Oracle Recovery Manager (RMAN), to support Oracle backup and restore utilities. It also includes Data Protection for Microsoft SQL Server, which enables users to perform online backups of SQL databases to Tivoli Storage Manager servers.

Note: DB2 and IBM Informix® are not included in Tivoli Storage Manager for Databases because these IBM database products include Tivoli Storage Manager agents.

You can continue running primary applications on your database servers while backup jobs are running. You can restore data to and from offline storage using automated tasks, utilities, and interfaces. This software performs online, consistent, and centralized backups to help you avoid downtime, protect vital enterprise data and minimize operational costs.

2.2.3 Tivoli Storage Manager for Virtual Environments: Data Protection for VMware

Tivoli Storage Manager for Virtual Environments - Data Protection for VMware provides, advanced data protection and flexible recovery options for VMware vSphere ESX and ESXi servers. Nondisruptive virtual machine backup is simplified and streamlined. VMware vSphere 4 and vSphere 5 environments are supported.

Tivoli Storage Manager for Virtual Environments protects VMware vSphere and vCloud virtual machines by offloading backup workloads to a centralized server and enabling near-instant recovery. It enables you to protect data without the need for a traditional backup window. You can protect the massive amounts of information that virtual machines generate without impacting the physical resources of the VMware server. Multiple virtual machines are supported with one Tivoli Storage Manager agent. LAN-free data transfer from the VMware server's storage to the backup server is supported. This preserves bandwidth for other users.

The management of the backup and restore processes for virtual machines is simplified. Tivoli Storage Manager for Virtual Environments provides an easy-to-use GUI that you can access from within the VMware vCenter. Tivoli Storage Manager for Virtual Environment - Data Protection for VMware V7.1 introduces a standalone web-based user interface in addition to the vCenter plug-in, to perform the same operation. This interface has the advantage of being web-based and compatible with the VMware vSphere 5 and later web-based management interface. It allows backup administrators (or VMware administrators) to create a backup policy or restore a virtual machine with just a few clicks.

Tivoli Storage Manager for Virtual Environments integrates with and extends the role of Tivoli Storage Manager in providing data protection services to virtualized applications in production environments. These include backup and recovery, online database and application protection, disaster recovery, data reduction, and bare machine recovery.

Overhead is eliminated with centralized vStorage backup. The VMware vStorage APIs for Data Protection technology is supported, which simplifies and optimizes data protection. This technology conducts operations on the backup and management servers rather than the virtual machines, greatly reducing system overhead and any disruption that backups might cause to virtualized applications.

A broad range of virtualized applications are supported and may be combined with application-aware software such as IBM Tivoli Storage Manager for Databases or IBM Tivoli Storage Manager for Mail to provide application-consistent snapshot backups.

VMware vStorage APIs for Data Protection are used, including block-level incremental forever backups based on VMware's Changed Block Tracking that take a nondisruptive snapshot at the virtual machine image level. The need for traditional backup windows is eliminated by continuously capturing data changes at the block level.

Tivoli Storage Manager for Virtual Environments retrieves data from image-level backups. It provides flexible recovery options for file-level, volume-level or virtual machine (VM) image-level recovery using a single backup of a virtual machine image.

In a file restore operation, the administrator launches the Tivoli Storage Manager for Virtual Environments restore on the vStorage backup server, accesses a point-in-time view of the data in the storage pool using Tivoli Storage Manager and performs a drag-and-drop of the desired files. The file restore operation can also be initiated from within the virtual guest environment. This feature is only available for Windows and Linux operating systems.

For a full disk volume restore, Tivoli Storage Manager for Virtual Environments mounts the point-in-time snapshot for that volume (mounted to the recovery volume) and makes it available immediately to users and applications. The actual data recovery happens in the background. All requests to write to and read from the volume are handled first, providing full, near-normal performance during the recovery process. This feature is only available for Windows and Linux operating systems.

VM image-level restores provide the recovery of not only the data, but the virtualized computing environment, from operating systems, applications and patches to upgrades and custom configurations.

2.2.4 AvePoint DocAve Backup and Restore

AvePoint DocAve Backup and Restore is a policy-based backup and recovery solution. It allows you to restore your Microsoft SharePoint business data and content after almost any kind of business interruption.

It provides the following functions:

- Performs backup and recovery of Microsoft SharePoint Portal 2003 and Microsoft Office SharePoint Server 2007 and SharePoint Server 2010 environments
- Restores portals, top level sites, subsites and individual document libraries, attachments, lists, folders, areas, and sub areas
- Schedules full, incremental, or differential backup at the site-level, subsite-level and item-level
- Preserves all meta-data versions
- Integrates with the Tivoli Storage Manager server so that you can create synchronous or asynchronous copies of SharePoint data for offsite protection
- Includes an easy-to-use browser-based graphical user interface (GUI).

For more information, see the website:

http://www-03.ibm.com/software/products/en/tivostormana/#docave

2.2.5 Tivoli Storage FlashCopy Manager

IBM Tivoli Storage FlashCopy® Manager delivers high levels of data protection for business-critical applications and databases through snapshot backup and restore capabilities that are integrated with the applications and databases. IBM Tivoli Storage FlashCopy Manager provides application-aware backups and restores by using the advanced snapshot technologies of storage systems.

Three versions of Tivoli Storage FlashCopy Manager are available:

- Tivoli Storage FlashCopy Manager for UNIX or Linux
- Tivoli Storage FlashCopy Manager for Windows
- Tivoli Storage FlashCopy Manager for VMware

Within these versions, data protection of the following applications can be configured with Tivoli Storage FlashCopy Manager:

- IBM DB2, IBM DB2 with SAP
- Oracle, Oracle with SAP
- Microsoft Exchange and Microsoft SQL Server

In addition, other applications can be supported on IBM AIX, HP-UX, Linux, Solaris, and Microsoft Windows platforms with script customizing.

These capabilities are achieved through the use of advanced storage hardware snapshot technology to help create a high performance, low impact application data protection solution. Tivoli Storage FlashCopy Manager is easy to install, configure, deploy, and seamlessly integrates with various storage systems such as: IBM System Storage® DS8000®; IBM XIV® Storage System products; IBM Storwize V7000 and Storwize V3700; IBM System Storage N series and NetApp systems; IBM System Storage SAN Volume Controller; and all IBM and non IBM devices supported by the SAN Volume Controller.

For a complete list, see the following website:

http://www.ibm.com/systems/storage/software/virtualization/svc/interop.html

The following list identifies the storage solutions that are natively supported with the Tivoli Storage FlashCopy Manager software:

- IBM XIV Storage System
- IBM Storwize V7000
- IBM System Storage SAN Volume Controller
- IBM System Storage DS8000
- IBM System Storage N series
- NetApp systems

By using Rocket Software and the Rocket Device Adapter Pack for IBM Tivoli Storage FlashCopy Manager for UNIX and Linux, new storage solutions can be used with Tivoli Storage FlashCopy Manager, such as EMC Symmetrix VMAX/DMX (TimeFinder Snapshot and Clone).

Rocket Device Adapter Pack (RDAP) for IBM Tivoli Storage FlashCopy Manager is a plug-in for IBM Tivoli Storage FlashCopy Manager. With RDAP for Tivoli Storage FlashCopy Manager, you can exploit Tivoli Storage FlashCopy Manager advanced snapshot-based data protection capabilities for application data that resides on EMC Symmetrix storage devices. RDAP for Tivoli Storage FlashCopy Manager supports TimeFinder/Snap and VP Snap, and mount support with device access masking ability.

For more information, see the following website:

http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&htmlfid=897/ENUS213-567&appname=USN

For more information about storage solutions supported for Tivoli Storage FlashCopy Manager see this website:

http://www.ibm.com/support/docview.wss?uid=swg21455924

In addition to those devices, Tivoli Storage FlashCopy Manager V3.2 and later on Windows supports any storage system that is Microsoft Volume Shadow Copy Services (VSS) capable by using the VSS system provider or a VSS hardware provider that strictly adheres to the Microsoft VSS Provider interface.

2.2.6 Tivoli Storage Manager for Enterprise Resource Planning

IBM Tivoli Storage Manager for Enterprise Resource Planning (ERP) V6 includes Data Protection for SAP. It 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 and SAP HANA in-memory databases. As a SAP certified product, Tivoli Storage Manager for Enterprise Resource Planning efficiently and consistently provides automated data protection for mySAP and SAP R/3 environments. You can improve the availability of your SAP database servers and reduce your administration workloads. Very large SAP databases can be backup efficiently.

Tivoli Storage Manager for Enterprise Resource Planning integrates with database-specific utilities of IBM DB2 for Linux, UNIX and Windows, Oracle, and SAP BR*Tools.

2.2.7 Tivoli Storage Manager for Space Management

Tivoli Storage Manager for Space Management automatically moves inactive data between storage tiers from expensive to less expensive media and frees online disk space for important active data. Inactive data is moved to reclaim online disk space for important active data. The movement of seldom-used files to and from less expensive tiers of storage is automated.

Tivoli Storage Manager for Space Management contains pre-migration tools that send a copy of the file to be migrated to the Tivoli Storage Manager server prior to migration, making storage space available and allowing for scheduled data transfer over the network.

You can predefine the stub file size. This feature eliminates the recall of the entire file for programs that only browse the first part of a file.

It is more integrated when used with the IBM General Parallel File System (GPFS™) policy engine. Files managed by Tivoli Storage Manager for Space Management can be moved from one storage pool (such as fast RAID) to another storage pool (such as slower SATA disks). The file path remains the same, including the file name, directory and file system name.

2.2.8 Tivoli Storage Manager Hierarchical Storage Management for Windows

Tivoli Storage Manager Hierarchal Storage Management for Windows provides hierarchical storage management (HSM) with a policy-based management system for migrating rarely used files on Windows file servers, economically and transparently.

Tivoli Storage Manager HSM for Windows controls disk storage requirements by automatically migrating rarely used files. It moves low-activity or inactive files to a hierarchy of lower-cost storage options. You gain an efficient storage capacity management solution that is based on rules-driven migration policies. Tivoli Storage Manager HSM for Windows has the following functions:

- Enhances user productivity by offering rapid access to archived data. It also helps maximize available disk space for users.
- Minimizes backup times and resource usage by focusing on active files.
- Migrates Microsoft Windows files based on file name, dates of creation, last access and modification. It also transparently recalls files to the original host system as needed.
- Takes advantage of an IBM Tivoli Storage Manager server for a hierarchy of storage tiers and reliable backup and recovery functions.
- Provides automatic threshold migration, which helps maintain a certain amount of free space on protected file systems.

2.2.9 Tivoli Storage Manager for Storage Area Networks

Tivoli Storage Manager for Storage Area Networks optimizes storage network connections for Tivoli Storage Manager servers and client computers by enabling LAN-free backup and recovery. Storage area network connections are used for LAN-free backup and recovery.

Tivoli Storage Manager for Storage Area Networks works with client computers and servers to make data transfers over a storage area network (SAN). It helps SAN-connected backup clients and Tivoli Storage Manager servers optimize their direct network connections to storage. As a result server CPU utilization is reduced.

Tivoli Storage Manager for Storage Area Networks provides LAN-free backup and restore, which removes data transfer from the LAN. This enables high-performance backup and restore and minimizes network traffic, improving application performance and transaction response times on the Tivoli Storage Manager server. It supports a SAN-connected tape library, multiple servers can share same tape library and tape drives on a SAN.

2.2.10 Tivoli Storage Manager for System Backup and Recovery

Tivoli Storage Manager for System Backup and Recovery offers system backup, restore and reinstallation, plus bare machine recovery capabilities for IBM AIX systems. It provides utilities to create backup scripts and schedules for easier task automation for businesses of any size.

Several types of backups are available: full system (installation image), volume group, file system, file or directory, and raw logical volume.

Tivoli Storage Manager for System Backup and Recovery offers the following functions:

- Supports configuration of network ports to communicate through firewalls, provides pass-thru support for AIX operating system language locales.
- Allows for the storage of backup objects into an IBM Tivoli Storage Manager server, Tivoli Storage Manager for System Backup and Recovery can back up any non-rootvg data.
- Integrates network boot and installation features in support of IBM RS/6000® SP systems. It also provides Offline Mirror Backup options to enable simultaneous user and system access to active copies of data.

2.2.11 Tivoli Storage Manager for z/OS Media

Tivoli Storage Manager for z/OS Media allows access to storage devices attached by IBM FICON® on a z/OS system.

IBM Tivoli Storage Manager for z/OS Media V6.3 and IBM Tivoli Storage Manager Extended Edition for z/OS Media V6.3 enable IBM Tivoli Storage Manager V6.3 servers, running on AIX and Linux on IBM System z®, to access various FICON attached tape or DASD resources on z/OS.

Tivoli Storage Manager provides access to SCSI and Fibre Channel storage; through Tivoli Storage Manager for z/OS Media, it also provides access to FICON attached storage. Tivoli Storage Manager for z/OS Media receives data from and sends data to a Tivoli Storage Manager V6.3 server, performing I/O to tape or DASD using sequential file access on behalf of the Tivoli Storage Manager V6.3 server.

Tivoli Storage Manager for z/OS Media provides improved sequential-access file implementation and supports the same storage devices as Tivoli Storage Manager for z/OS V5.5 server. It also interacts with z/OS Data Facility Storage Management Subsystem (DFSMS) and the Tape Management System in the same way as Tivoli Storage Manager for z/OS V5.5 server.

The database of a Tivoli Storage Manager V5 server that is running on a z/OS system can be migrated to a V6.3 server that runs on AIX or Linux on System z. After the upgrade, z/OS users can continue to access data stored on tape volumes whose contents are accessed by using FICON attached storage devices through the new Tivoli Storage Manager for z/OS Media connector.

2.2.12 Tivoli Storage Manager FastBack

IBM Tivoli Storage Manager FastBack® provides continuous data protection and recovery management for Microsoft Windows and Linux servers. It provides smarter data protection and near-instant recovery in data centers, remote offices, and small to mid-sized enterprises. Tivoli Storage Manager FastBack provides the following functions:

- Helps to eliminate the need for traditional backup windows by continuously capturing data changes at the block level with extremely low overhead on the protected systems.
- Schedules automated data transfers based on flexible, policy-based settings to help administrators meet data protection and recovery requirements on a per-application basis.
- More effectively uses your available bandwidth through built-in capabilities such as block-level incremental backup, data deduplication, bundling of small files and industry-standard compression.
- Supports a unified recovery management strategy across your entire enterprise through tight integration with IBM Tivoli Storage Manager.

2.2.13 Tivoli Storage Manager FastBack for Microsoft Exchange

Tivoli Storage Manager FastBack for Microsoft Exchange enables users and administrators to recover Microsoft Exchange data objects more quickly. It can restore individual mail messages, attachments, calendar entries, contacts, and tasks. Tivoli Storage Manager FastBack for Microsoft Exchange integrates with Active Directory and Exchange Server security to help limit unauthorized access to backup and restore systems. Now you can increase productivity by reducing recovery time from hours or days to just minutes.

Microsoft Exchange data objects can be recovered from virtually any Microsoft Exchange database, even corrupt databases. It enables recovery of objects that were previously considered unrecoverable, such as deleted mail messages or address books that were lost because of synchronization errors.

Microsoft Exchange recovery is optimized by applying it at a granular level to any individual data object or group of objects, including individual mail messages, contact lists, tasks, or calendar entries. This software also supports public folders.

The downtime that is associated with data recovery is minimized, thus improving service levels. The software supports Microsoft Exchange 2003, 2007, and 2010. Objects can be restored directly to an Exchange Server or sent to a user-defined destination using Simple Mail Transfer Protocol (SMTP).

2.2.14 Tivoli Storage Manager FastBack for Bare Machine Recovery

You can recover Windows or Linux server operating system volume quickly from a disaster or catastrophic server failure with Tivoli Storage Manager FastBack for Bare Machine Recovery. Organizations can perform bare machine recovery at a local office, data center, or central recovery site. Tivoli Storage Manager FastBack for Bare Machine Recovery facilitates the fast and easy migration of workloads from old hardware or standalone servers to new hardware platforms.

For Windows systems, it provides the flexibility of recovering to comparable hardware, to dissimilar hardware, or to a virtual machine using VMware or Microsoft Virtual Hyper V.

This feature provide near-instant access to applications and data, while full recovery occurs in the background. It helps to protect remote or branch offices with a cost-effective disaster recovery and business resiliency strategy that requires minimal standby hardware.

2.2.15 Tivoli Storage Manager FastBack for Workstations

Tivoli Storage Manager FastBack for Workstations provides near real-time continuous data protection for desktop and laptop computers. It provides central management of end points, including auto-discovery, pushing updates, configuring settings, and initiating backups.

This product is specifically designed for desktop and laptop computers. It provides continuous protection for key corporate information also. By backing up your most important files as they are saved. It does not wait for a scheduled interval to back up files.

Tivoli Storage Manager FastBack for Workstations helps minimize recovery times and maximize the granularity of recoverable data by providing user-initiated recovery. It helps avoid an impact on employee productivity following a data loss or system failure.

2.2.16 Tivoli Storage Manager FastBack Center

Tivoli Storage Manager FastBack Center offers full-featured data protection and recovery for midsized businesses and remote office systems on Windows and Linux.

Tivoli Storage Manager FastBack Center is a convenient bundle of three IBM software products: Tivoli Storage Manager FastBack, IBM Tivoli Storage Manager FastBack for Bare Machine Recovery, and IBM Tivoli Storage Manager FastBack for Microsoft Exchange.

2.2.17 Tivoli Storage Manager Suite for Unified Recovery

Tivoli Storage Manager Suite for Unified Recovery provides an enterprise-wide unified recovery management solution with extensive data protection providing simplified licensing and management. This bundle of Tivoli Storage Manager and Tivoli Storage Manager FastBack products enables you to deploy any of 10 solution components, and in any location and quantity. The simplified license measures only the amount of data being managed.

It helps organizations of all sizes meet a wide range of data management challenges for complex, distributed infrastructures. You can deploy the advanced management tools you need for each of your individual data protection requirements without having to worry about individual product licenses. Tivoli Storage Manager Suite for Unified Recovery offers these functions:

- Provides extensive data protection for a wide range of systems including virtual machines, file servers, email, databases, mainframes and even desktops. This bundled solution allows you to use the right data protection tool for each of your requirements.
- Reduces costs and simplifies procurement and deployment with per terabyte capacity licensing. You can deploy any of 10 solution components, in any location and quantity, with a simplified license that measures only the amount of data being managed.
- Scales to meet the recovery needs of any size organization by managing up to four billion data objects on a single server. This solution supports more than 50 operating system versions and hundreds of server and storage devices.
- Manages the entire suite of products from a single user console, you can configure, manage, upgrade, report and monitor all 10 products from a single administration interface.

2.3 Tivoli Storage Manager recent enhancements

This section lists enhancements in Tivoli Storage Manager V6.4 and V7.1.

2.3.1 Tivoli Storage Manager V6.4 enhancements

Tivoli Storage Manager V6.4 contains many new features. Most of them are related to VMware and NetApp support.

Tivoli Storage Manager for Virtual Environment: Data Protection for VMware

The VMware enhancements in Tivoli Storage Manager V6.4 are as follows:

- Performance with progressive incremental backup for virtual environments, removing the requirement to perform periodic full backups in a VMware environment for Data Protection.
- Can now run incremental forever backups on hypervisors, which saves backup time and simplifies backup scheduling.
- Allows granular selection of domain objects, including host clusters, data store and wildcard characters.
- Backup performance is improved. Multiple backup processes are allowed to run in parallel. Parallelism is available only on the VM level, not on the disk level.
- A new configuration wizard helps to automate many of the configuration and deployment tasks of VMware Tivoli Storage Manager clients.
- For Data Protection for Microsoft SQL and Data Protection for Microsoft Exchange database backups, Tivoli Storage Manager V6.4 can truncate the SQL or Exchange transaction logs.
- Backup reporting from the vCenter plug-in, with sample reports is provided.
- Simplified deployment and configuration when integrating with VMware vCenter for Data Protection.
- Enhanced recovery performance and reporting for VMware environments.

NetApp

The Netapp enhancements in Tivoli Storage Manager V6.4 are as follows:

- Tivoli Storage Manager V6.4 supports snapshot-assisted progressive incremental backup. This is usually used with NetApp's SnapMirror replication to take backups from a remote SnapMirror site. Tivoli Storage Manager can now use NetApp snapshot facility to take a progressive incremental backup.
- Tivoli Storage Manager V6.4 supports snapshot-assisted progressive incremental backup operations now for vFilers (virtual controllers).

IBM Cognos reporting

The IBM Cognos® Business Intelligence Reporting Suite was introduced with Tivoli Storage Manager V6.3 to provide ready to use reports on Tivoli Storage Manager status.

Now Tivoli Storage Manager V6.4 provides more ready-to-use Cognos reports. Users can customize existing Cognos reports or rapidly create new reports with a drag-and-drop function.

SAP HANA database (ERP) support

SAP HANA uses in-memory databases that can be backed up using a new Tivoli Storage Manager for ERP component called Data Protection for SAP HANA. The entire backup process is automated into a single step and all the files are managed as a logical entity.

Data Protection for Microsoft Exchange Server

Data Protection for Microsoft Exchange Server enhancements in Tivoli Storage Manager V6.4 are as follows:

- Support for Exchange 2012 Database Availability Group (DAG) is enhanced, the management and automation of Exchange DAG active and passive database backups and restores are now fully integrated.
- Self-contained, application-consistent backup for Microsoft Exchange running on VMware
- Simplified deployment and configuration when integrating with VMware vCenter for Data Protection

Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server

The Tivoli Data Protection SQL enhancements in Tivoli Storage Manager V6.4 are as follows.

- Supports Microsoft SQL 2012 with AlwaysOn Availability Group, with automation for database backups and restores for a high availability configuration.
- Password management is now integrated with LDAPs to standardize control of access.
- Self-contained, application-consistent backup for MS-SQL running on VMware.
- Simplified deployment and configuration when integrating with VMware vCenter for Data Protection.

2.3.2 Tivoli Storage Manager V7.1 enhancements

Tivoli Storage Manager V7.1 has several enhancements. More information about all enhancements of the Tivoli Storage Manager V7.1 are at the following location:

http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1/index.jsp

Tivoli Storage Manager Server

Enhancements to Tivoli Storage Manager Server in V7.1 are as follows.

- Operations Center updates

  New features are available to help you manage your storage environment.
- Automatic client failover for restore operations from replicated servers

  A Tivoli Storage Manager Version V7.1 client can automatically fail over to a target replication server for restore and retrieve operations, if the source replication server is unavailable. This includes also some of the Data Protection modules (such as Data Protection for MS-SQL server or Data Protection for Oracle or Data Protection for Exchange).
- Immediate use of space that is added to the server database

  When you add space to the database, new database directories are now available for immediate use and parallel I/O performance is improved.
- File-space level collocation groups

Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server

You can recover Microsoft SQL databases from a VM backup. To complete this task, use both Tivoli Storage Manager for Virtual Environments and Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server.

Tivoli Storage Manager for Virtual Environments - Data Protection for VMware

Tivoli Storage Manager for Virtual Environments - Data Protection for VMware introduces several new features in V7.1.

- Simplified installation and configuration
- Recovery of individual Microsoft SQL databases from a VM backup
- Protection for VMware vCloud Director vApps
- Instant restore of a full-vm
- Protection for VM guests that host Microsoft Active Directory domain controller

Data Protection for Microsoft Exchange Server

In Tivoli Storage Manager V7.1 the following functions are provided in Data Protection for Microsoft Exchange Server.

- Microsoft Exchange Server 2013 support

  When you are planning to restore Exchange Server 2013 mailboxes, a requirement exists for MAPI clients to use RPC over HTTPS (also known as Outlook Anywhere). RPC over HTTP is no longer supported.

- Restoring mailboxes directly from Exchange database files

  With all Data Protection for Microsoft Exchange configurations, administrators can complete an individual mailbox restore for an `.edb` file that is stored on a disk. For administrators who only want to complete individual mailbox restores from an `.edb` file on disk, the Mailbox Restore Only configuration option is available. The `.edb` file can come from a backup that mounted read-write using Tivoli Storage Manager for Virtual Environments, Tivoli Storage Manager restore files, or an offline file system copy.

Tivoli Storage FlashCopy Manager

Tivoli Storage FlashCopy Manager V4.1 is the version associated with the Tivoli Storage ManagerV7.1 product family. Here we list the recent enhancements.

Tivoli Storage FlashCopy Manager for UNIX and Linux

Enhancements are as follows:

- Clone databases in an Oracle Data Guard standby environment.
- Protect applications on VMware virtual machines that run Linux guest operating systems.
- Protect databases in a DB2 IBM pureScale® environment.
- Support for network-attached storage in N series and NetApp environments.

Also, starting with this version of Tivoli Storage FlashCopy Manager, you can use Rocket Device Adapter for IBM Tivoli Storage FlashCopy Manager. This brings you the support of more storage vendors such as EMC disk bays.

For more information about Rocket Device Adapter software, go to the following address:

http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&htmlfid=897/ENUS213-567&appname=USN

Tivoli Storage FlashCopy Manager for Windows

Enhancements are almost all related to Microsoft Exchange data protection.

- Additional options for restoring personal storage folders (Restore Mail to Unicode PST file and Restore Mail to non-Unicode PST file)
- Group management, reports, and status reporting
- Managing remotely (Manage all of the Tivoli Storage FlashCopy Manager installations in your enterprise from a unique point)
- Microsoft Exchange Server 2013 support
- Restoring mailboxes on remote systems
- Restoring mailboxes directly from Exchange database files
- Automated failover for data recovery

Tivoli Storage FlashCopy Manager for VMware

Enhancements are as follows:

- Access Tivoli Storage FlashCopy Manager for VMware GUI by using a web browser GUI
- Instant restore processing (restore all data stores that are included in a backup and all virtual machines are registered)
- Tivoli Storage FlashCopy Manager for VMware integration with VMware vCenter Site Recovery Manager

2.4 Conclusion

As data and storage management becomes more sophisticated, it also becomes more complicated. Business needs change daily. Regulations exist about backup, archiving, disaster recovery, copy version, and retention; satisfying rules, laws, and regulations, is not easy for companies. The key to properly managing a complex storage environment is to approach it strategically.

In this book, we discuss many solutions by using a strategic approach that can make storage management easier, more efficient, and more effective.

3

Data protection with Tivoli Storage Manager

This chapter is intended for the reader who wants to have a brief refresh on techniques and concepts that are provided with the Tivoli Storage Manager product. If you are familiar with the product, you might want to go directly to Chapter 4, “Tivoli Storage Manager challenge matrix” on page 81.

Data protection has become a critical priority in business today. As our planet gets smarter, our computing systems must become smarter too. Systems must become more automated, adaptive, and robust. Data protection and retention solutions need to efficiently store and manage growing data volumes while keeping data secure and available. This section discusses various Tivoli Storage Manager tools and functions available to protect your data.

Data protection can be done at several different layers and places and for many different data types. This chapter discusses the Tivoli Storage Manager tools, explaining where and at what layer they are acting to effectively protect your data.

3.1 Protect the data at the client side

This section discusses the Tivoli Storage Manager tools to protect the data at the client side.

3.1.1 Operating system protection

Backups of various operating systems can be performed.

Windows operating system backup

Tivoli Storage Manager uses Microsoft Volume Shadow Copy Service (VSS) to back up all system state components as a single object, to provide a consistent point-in-time snapshot of the system state. System state consists of all startable system state and system services components.

Tivoli Storage Manager supports VSS on the supported Windows clients.

The system state backup consists of data from several VSS writers. When the files that belong to the System Writer change, an incremental backup is used for these files. Using incremental backup on the System Writer files reduces the amount of time to back up the system state. A full backup is used for all other system state data.

By default, a progressive incremental backup is used for the System Writer files in the system state. If you want to run a full backup of all system state data, you can specify `include.systemstate mc_name` in the client options file (`dsm.opt`), where `mc_name` is the name of the management class with copy group mode absolute.

Startable system state components include the following items:

- Automated System Recovery (ASR) writer
- Active Directory (domain controller only)
- System Volume (domain controller only)
- Certificate Server Database (domain controller only)
- COM+ database
- Windows Registry
- System and boot files

System services components include the following items:

- Background Intelligent Transfer Service (BITS)
- Removable Storage Management Database (RSM)
- Cluster Database (cluster node only)
- Remote Storage Service
- Terminal Server Licensing
- Windows Management Instrumentation (WMI)
- Internet Information Services (IIS) metabase
- Dynamic Host Configuration Protocol (DHCP)
- Windows Internet Name Service (WINS)

The list of startable system state and system services components are dynamic and can change depending on service pack and operating system features installed. Tivoli Storage Manager allows for the dynamic discovery and back up of these components.

Back up Automated System Recovery (ASR) files in preparation for recovering the Windows disk configuration information and system state in case a catastrophic system or hardware failure occurs.

On Windows Vista, Windows 7, and Windows Server 2008, the Tivoli Storage Manager backup-archive client backs up ASR data when the backup-archive client backs up the Windows system state. Use the **`backup systemstate`** command to back up ASR files on these operating systems.

Tivoli Storage Manager generates the ASR files in the `\adsm.sys\ASR` staging directory on the system drive of your local workstation and stores these files in the ASR file space on the Tivoli Storage Manager server.

See the following sources for more information:

- *Best Practices for Recovering Windows Server 2008, Windows Server 2008 R2, Windows 7, and Windows Vista*:

  http://ibm.co/1ptFRi7

- *Modified Instructions for Complete Restores of Windows Systems with the TSM Client: Bare Metal Restore (BMR), System State Restore, Windows System Object Restore*:

  http://www.ibm.com/support/docview.wss?uid=swg21164812

AIX system backup

On AIX platforms, *Tivoli Storage Manager for System Backup and Recovery* offers system backup, restore, and reinstallation, plus bare machine recovery capabilities for IBM AIX systems. It provides the following functions:

- Provides utilities to create backup scripts and schedules for easier task automation for businesses of any size.
- Enables the choice of several types of backups, these include full system (installation image), volume group, file system, file or directory and raw logical volume.
- Supports configuration of network ports to communicate through firewalls, provides pass-thru support for AIX operating system language locales.
- Allows for the storage of backup objects into an IBM Tivoli Storage Manager server, Tivoli Storage Manager for System Backup and Recovery can back up any non-rootvg data.
- Integrates network boot and install features in support of IBM RS/6000 SP systems, also provides Offline Mirror Backup options to enable simultaneous user and system access to active copies of data.

Other UNIX and Linux systems

Tivoli Storage Manager backup-archive client can be used to protect operating systems such as Linux, Solaris, or other UNIX systems. All operating system files are protected as any other files on the system, unless you specified exclude statements, so these operating system files are processed like any others.

To recover from a machine disaster, use a system restore utility that is appropriate for your platform or build a new replacement to prepare the recovery using the Tivoli Storage Manager backup-archive client.

TBMR

The backup-archive client can be used together with the third-party product TBMR for recovery. Cristie TBMR is powerful, easy-to-use software that integrates with IBM Tivoli Storage Manager. It provides rapid, automatic machine recovery to an identical state following failure of physical hardware or corruption of an operating system.

TBMR helps users of Tivoli Storage Manager recover their operating systems directly from a Tivoli Storage Manager backup. No additional backup is required and users can invoke the

power of Tivoli Storage Manager to help recover their operating system to any point in time that is provided by Tivoli Storage Manager. Multiple servers can be recovered simultaneously and support includes dissimilar hardware, allowing recovery to different server models or brands.

Key features of Cristie TBMR are as follows:

- Support for dissimilar hardware that allows recovery to different brands, models, or to virtual environments
- Incremental backup support that uses Tivoli Storage Manager to help recover operating systems to any point in time and reduce data storage so operating system files are not backed up twice
- Automatic hardware and driver detection
- Simultaneous recovery of multiple servers
- Individual basic partition restore that leaves other partitions intact
- Dynamic disk volume support
- Backup versioning
- System state backup

A description of the TBMR product is in the IBM announcement letter:

http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=an&subtype=ca&supplier=872&letternum=ENUSAP09-0011#

3.1.2 Flat file backup

Flat file backup is also known as unstructured data. It means that the files can be backed up without taking care of any consistency problems, as we would do for databases or applications.

Progressive incremental backup

One of the key differentiators between Tivoli Storage Manager and other data protection products is the *progressive incremental backup methodology*. The term *progressive incremental* is sometimes called *incremental forever*. When this technology is used, Tivoli Storage Manager backs up only new or changed files. Some Tivoli Storage Manager backups do not use this technology, for example, Network Data Management Protocol (NDMP).

Tivoli Storage Manager tracks all of the backups at a file level. It has no concept of a full backup with dependent incrementals or differentials. Because of the Tivoli Storage Manager powerful relational database, it does not require periodic full backups.

This methodology reduces network and storage resource consumption and lowers the overall cost of storage management. Tivoli Storage Manager's file level progressive backup methodology is far superior to other traditional backup methods such as full-plus-incremental or full-plus-differential, because progressive incremental backups are never redundant.

Image backup

An image backup is a block-by-block copy, single object backup of a volume (typically a UNIX file system or raw logical volume, or Windows drive) on a Tivoli Storage Manager client. Being able to restore an entire volume as one object can lead to faster recoveries. Image backup is available at the time of writing on AIX, HP-UX, Oracle Solaris, Linux, and Windows client platforms.

Review the specific image backup requirements for each platform in their associated installation guides found at this website:

http://publib.boulder.ibm.com/infocenter/tsminfo/v6/index.jsp

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 is applied to the file system image as a whole.

An image backup provides the following benefits:

- Backs up file systems containing a large number of files faster than a full file system incremental back up.
- Improves the speed with which Tivoli Storage Manager restores file systems containing many small files.
- Conserves resources on the server during backups because only one database entry is required for the image.
- Provides a point-in-time picture of your logical volume, which might be useful if your enterprise needs to recall that information.
- Restores a corrupt file system or raw logical volume. Data is restored to the same state it was when the last logical volume backup was performed.

To restore an image backup of a volume, the Tivoli Storage Manager client must be able to obtain an exclusive lock on the volume being restored.

Journal-based backup

Journal-based backup provides an alternative to traditional progressive incremental backup, which under certain circumstances can dramatically increase overall backup performance.

As the name already implies, journal-based backups have no effect on archive processing. The main difference between journal-based backup and progressive incremental backup is the method in which the list of backup candidate objects is derived.

A 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 backup-archive client.

The progressive incremental backup operation 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 on the local file system.

The server list is obtained over the network and the local list is obtained by scanning the local file system. Objects that exist in the local list but do not exist in the server list are added as backup candidates to the candidate list.

Objects that exist in both lists but differ in some way (such as attributes, policy, and size) are also added as backup candidates unless only the Tivoli Storage Manager database attributes differ, in which case they are added as attribute update candidates. Objects that exist in the server list but not in the local list are added as expiration candidates.

Figure 3-1 on page 38 describes how the journal engine keeps track of changes in the file system and communicates with the client during backup. You can place the journal databases for different disks on a single volume or separate volumes.

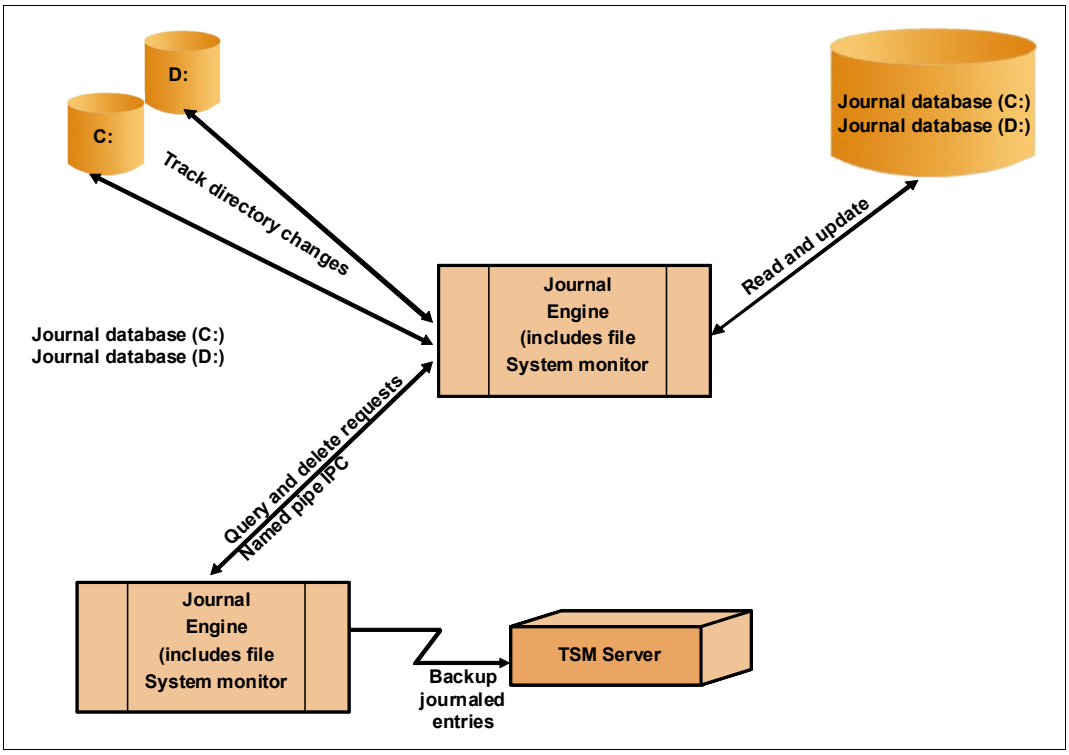

Figure 3-1 Tivoli Storage Manager journal-based backup

The Tivoli Storage Manager backup-archive client obtains the backup candidate list by contacting the journal-based backup daemon. This is a local background process that manages and maintains a journal database of change activity for each file system being journaled.

Journal-based backup can improve incremental backup performance in most environments. With journal-based backup, the client does not scan the local file system or obtain information from the server to determine which files to process. As such, journal-based backup reduces network traffic between the client and server. The backup-archive client, as always, still sends data (files) to the Tivoli Storage Manager server and, as has always been the case, the Tivoli Storage Manager server stores file details and location in the Tivoli Storage Manager database.

Because the backup-archive client does not carry out the initial metadata conversation, the backup-archive client does not have to sit idle. The backup-archive client can begin sending the files to the Tivoli Storage Manager server as soon as the journal-based backup is initiated. This means faster backup times and less backup-archive client idle time.

Previously, this file list construction and processing could severely impact any Tivoli Storage Manager backup-archive clients that were memory-bound or CPU-bound.

On supported platforms, journal-based backups aid all types of backups (progressive incremental backup, selective backup, adaptive subfile backup) by basing the backups on a list of changed files. It significantly reduces the amount of time required for backup as the files eligible for backup are known before the backup operation begins.

Adaptive subfile backup

A basic problem that remote and mobile users face today is connecting to storage management services by using modems with limited bandwidth or poor line quality. This creates a need for users to minimize the amount of data they send over the network, as well as the time that they are connected to the network.

To help address this problem, you can use subfile backups. When a client's file has been previously backed up, any subsequent backups are typically made of the portion of the client's file that has changed (a subfile), rather than the entire file. A base file is represented by a backup of the entire file and is the file on which subfiles are dependent. If the changes to a file are extensive, a user can request a backup on the entire file. A new base file is established on which subsequent subfile backups are dependent.

The adaptive subfile backup function is still available but is being superseded by the client side deduplication function to reduce the data amount during the backup and archive processing. For more information about client side deduplication, see "Client-side data deduplication" on page 40.

3.1.3 Backup considerations

In this section, we look at functions that affect client backup.

Compression

You have the option to specify whether 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 it sends a smaller quantity of data) at the expense of requiring client CPU resources to perform the operation.

Therefore, the decision to enable client compression must be made individually for each configuration. If using client compression, then the client will also automatically decompress 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 storage pools, reducing resource requirements.

If you do not enable client compression, the files will be sent at their full size to the Tivoli Storage Manager server. Most sequential storage devices, like tape drives, can perform hardware compression. If this is the case, you still can get the benefit of reduced space required in the storage pool. If a client has already-compressed the files, then a compression-enabled tape drive normally will not be able to compress them further. Compression rates vary considerably depending on the type of data presented.

As adaptive subfile backup, or client deduplication, compression can be used to back up nodes with a limited bandwidth.

For further data reduction, you can enable client-side data deduplication and compression together. Each extent is compressed before it is sent to the server. Compression saves space, but it increases the processing time on the client workstation.

Include-exclude lists

You can create an include-exclude list to exclude a specific file or groups of files from backup services, and to assign specific management classes to files. Tivoli Storage Manager backs up any file that is not explicitly excluded.

You might only need to back up particular drives or file systems and exclude all others. Or you might want to exclude from backup any files with a `.TMP` extension. You might need a different management policy for certain critical files (such as spreadsheets, documents, and email messages), than for ordinary files that you can easily get back if necessary (such as Internet files and temporary files).

The management class concept gives Tivoli Storage Manager granular control over how and where each file is backed up. The include-exclude list is a set of references local to each client that controls which files are backed up and what management class is used. If you do not select an explicit management class, Tivoli Storage Manager uses a designated default management class.

The include-exclude list is a very powerful tool for specifying exactly what a client should back up, if any additional processing should be applied and where the data is stored.

Client-side data deduplication

Client-side deduplication processes the redundant data during the backup or archive process on the host system where the source data is located. The net results of deduplication are virtually the same as with server-side deduplication, except that the storage savings are realized immediately, since only the unique data needs to be sent to the server in its entirety. Data that is duplicate requires only a small signature to be sent to the Tivoli Storage Manager server. Client-side duplication is especially effective when it is important to conserve bandwidth between the Tivoli Storage Manager client and server. For greater data reduction client side deduplication can also be combined with client compression. Server-side deduplication will also effectively work with data that has already been compressed by the client, because the server now understands the client compression format.

Client deduplication cache

Although it is necessary for the backup client to “check in” with the server to determine whether a chunk is unique or a duplicate, the amount of data transfer is small. The client must query the server for each chunk of data that is processed. The overhead associated with this query process can be reduced substantially by configuring a cache on the client, which allows previously discovered chunks on the client (during the backup session) to be identified without a query to the Tivoli Storage Manager server. For the backup-archive client (including VMware backup,) you should always configure a cache when using client-side deduplication.

For applications that use the Tivoli Storage Manager API, the deduplication cache should not be used because of the potential for backup failures that are caused by the cache being out of sync with the Tivoli Storage Manager server. If multiple, concurrent Tivoli Storage Manager client sessions are configured (such as with a Tivoli Storage Manager for VMware vStorage backup server), there must be a separate cache that is configured for each session.

Data encryption

To improve the security of stored data, the backup-archive client implements an optional encryption function, which allows for encrypting 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 is encrypted and thus is unreadable by any malicious intruders.

For the strongest possible encryption, you can specify to use 128-bit Advanced Encryption Standard (AES) data encryption, with the encryptiontype option. The user can choose which files are subject to encryption through include and exclude processing. The encryption uses a 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.

Tivoli Storage Manager also supports a key management method where the server stores the encryption key. When `encryptkey=generate` is specified, an encryption key password is dynamically generated when the Tivoli Storage Manager client begins a backup or archive.

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 backup as well as for archive. Tivoli Storage Manager will not attempt to deduplicate client-encrypted data.

LAN-free data movement

Tivoli Storage Manager for Storage Area Networks works with client computers and servers to make data transfers over a storage area network (SAN). It helps SAN-connected backup clients and Tivoli Storage Manager servers optimize their direct network connections to storage. This is called *LAN-free data movement*.

LAN-free data movement makes LAN bandwidth available for other uses and decreases the load on the Tivoli Storage Manager server, allowing it to support a greater number of concurrent client connections.

The key component of Tivoli Storage Manager for Storage Area Networks is the *storage agent*. You install the storage agent on a client system that shares storage resources with the Tivoli Storage Manager server.

The storage agent can support several clients while installed on only one of the clients. You can point a Tivoli Storage Manager client to an off-host storage agent. To do this, you do not need to install the storage agent except on the machine that is writing to the SAN.

Figure 3-2 shows SAN data movement with the `lanfreecommmethod` option.

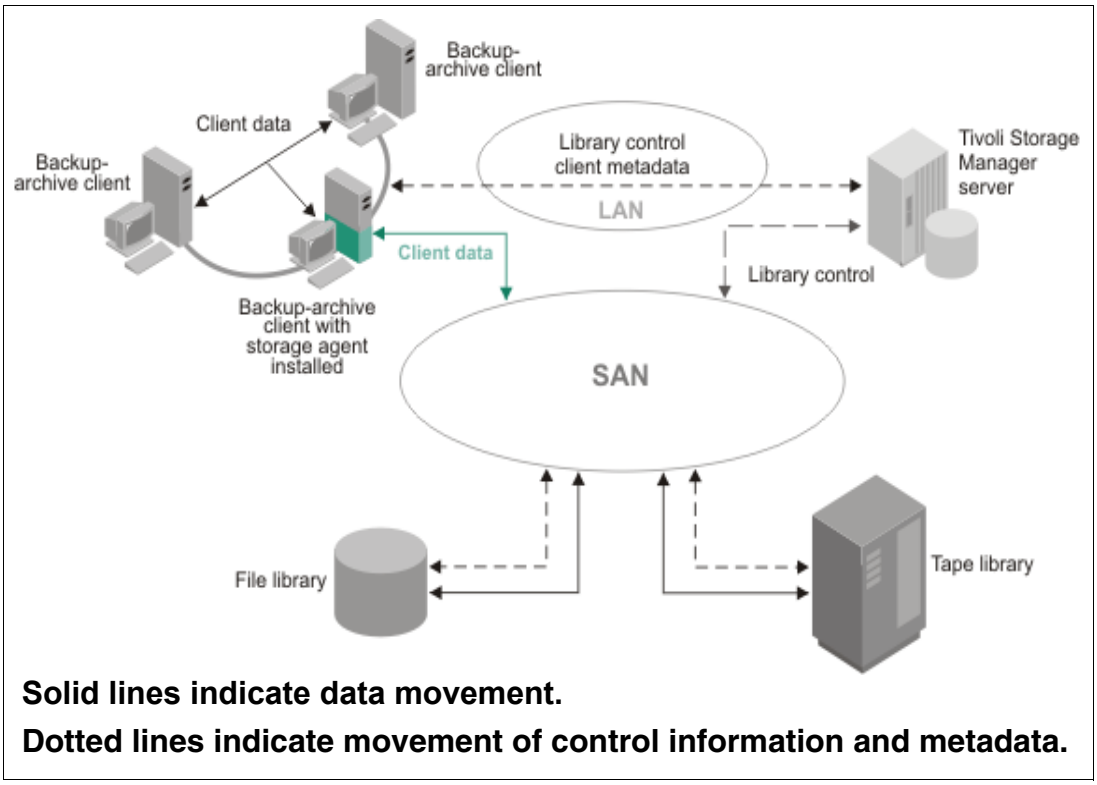

Figure 3-2 SAN data movement with the LANFREECOMMMETHOD option

A Tivoli Storage Manager server, acting as a library manager, controls the storage devices. This server can be the server working in conjunction with the storage agent or another Tivoli Storage Manager server in the enterprise. The Tivoli Storage Manager server keeps track of the metadata that the client has stored. The metadata, such as policy information and file name and size, is passed over the LAN connection between the storage agent and server.

The storage agent communicates with the server to obtain and store database information and to coordinate device and volume access. The server and client coordinate and transfer data access through the SAN. The client uses the storage agent for operations where appropriate. For example, if a SAN path is defined, the client (by means of the storage agent) transfers data on that path. If a failure occurs on the SAN path, failover occurs, and the client uses its LAN connection to the Tivoli Storage Manager server and moves the client data over the LAN.

The storage agent can send the data directly to the server using the LAN control paths between the storage agent and the server. An example is a LAN-free storage pool that is updated to read-only after the client connects to the server and obtains its initial policy information. The storage agent, instead of failing the operation, sends the data to the server. If the storage hierarchy is configured so that the next storage pool destination is available, the server performs the operation.

Tivoli Storage Manager supports SAN-attached device sharing in the following environments:

- Tivoli Storage Manager native library management support consists of an Automated Cartridge System Library Software (ACSLS), SCSI, or 349x library manager and library clients or just a library manager.
- Shared disk storage using a FILE library and the integration of IBM General Parallel File System. IBM General Parallel File System is the preferred option for operating systems on which it is supported.

3.1.4 Microsoft Volume Shadow Copy Service

Volume Shadow Copy Service (VSS) is a set of Microsoft COM APIs that implement a storage management architecture enabling volume-level snapshots to be performed while the applications that contain data on those volumes remain online and continue to write. This architecture can be used for backing up your applications by creating point-in-time snapshots of your data. VSS provides the backup infrastructure as well as the mechanism for creating consistent point-in-time copies of data known as shadow copies. VSS produces consistent shadow copies by coordinating with business applications, file-system services, backup applications, fast-recovery solutions, and storage hardware.

VSS is supported with Data Protection for Exchange and Data Protection for SQL that can be used in conjunction with Tivoli Storage FlashCopy Manager to allow for offload VSS backup operations.

A VSS backup uses Microsoft Volume Shadow Copy Service technology to produce an online snapshot (point-in-time consistent copies) of the data. It allows a snapshot of large amounts of data at once. During a VSS backup, the application server is not in backup mode for an extended period of time, because the length of time to perform the snapshot is usually measured in seconds and not hours.

VSS components

This list defines several VSS components, concepts, and terminology used in this book.

- VSS service

  This is the Microsoft Windows service that directs all of the VSS components that are required to create the point-in-time snapshot or the shadow copy of the volumes. This Windows service is the overall coordinator for all VSS operations.

- Requestor

  This is the software application that commands a shadow copy be created of a specified volume. In a Tivoli Storage Manager VSS environment, the Tivoli Storage Manager Backup-Archive Client is the requestor.

- Writer

  This is the software application that places the persistent information for the shadow copy on the specified volumes. Database applications (such as SQL Server or Exchange Server) or a system service (such as Active Directory) can be a writer.

- Provider

  This is the application that actually produces the shadow copy and also manages their availability. A system provider (such as the one included with the Microsoft Windows operating system), a software provider, or a hardware provider (such as one shipped with a storage system) can be a provider.

- Persistent shadow copy

  This is a shadow copy that remains after the backup application completes its operations. This type of shadow copy also survives system reboots.

- Non-persistent shadow copy

  This is a temporary shadow copy that remains only as long as the backup application needs it in order to copy the data to its backup repository.

- Transportable shadow copy

  This is a shadow copy volume that is accessible from a secondary host so that the backup can be off-loaded. Transportable is a feature of hardware snapshot providers.

- Source volume

  This is the volume that contains the data to be shadow copied. These volumes contain the application data.

- Target or Snapshot volume

  This is the volume that retains the shadow copied storage files. It is an exact copy of the source volume at the time of backup.

Shadow copy methods

The methods used with shadow copy are as follows:

- Clone (full copy and split mirror)

  A clone is a shadow copy volume that is a full copy of the original data as it resides on a volume. The source volume continues to take application changes while the shadow copy volume remains an exact read-only copy of the original data at the point-in-time that it was created. Disk mirroring is the process of writing data to two separate hard disks at the same time. One copy of the data is called a mirror of the other. Splitting a mirror is the process of separating the two copies.

- Copy-on-write (differential copy)

  A copy-on-write shadow copy volume is a differential copy (rather than a full copy) of the original data as it resides on a volume. This method makes a copy of the original data before it is overwritten with new changes. Using the modified blocks and the unchanged blocks in the original volume, a shadow copy can be logically constructed that represents the shadow copy at the point-in-time it was created.

- Redirect-on-write (differential copy)

  A redirect-on-write shadow copy volume is a differential copy (rather than a full copy) of the original data as it resides on a volume. This method is quite similar to copy-on-write, without the double write penalty, and it offers storage space and performance efficient snapshots. New writes to the original volume are redirected to another location set aside for snapshot. The advantage of redirecting the write is that only one write takes place, whereas with copy-on-write, two writes occur (one to copy original data onto the storage space, the other to copy changed data).

The VSS provider implementation determines which type of shadow copy (clone, copy-on-write, or redirect-on-write) is created.

For more information about VSS and Tivoli Storage Manager related product concepts see Appendix B, “Hierarchical storage management (HSM)” on page 323.

3.1.5 Microsoft cluster support wizard

Use the Tivoli Storage Manager cluster wizard to configure the backup-archive client to protect cluster resources. The wizard gathers the information needed so the backup-archive client can protect cluster resources, and log on to the server.

You need to run the wizard only on one node in the cluster; the wizard creates the necessary services on all nodes in the cluster. To ensure that any node can perform backups if any other node fails, the wizard copies the registry data to all nodes in the cluster.

The wizard can configure only one cluster group at a time. If you have multiple cluster groups to protect, run the wizard as many times as needed to configure the client to back up each group.

3.1.6 Virtual environment specifics

Most virtual environments support the concept of virtual machine snapshots - which is to say point in time images of a virtual machine that you can return to at any stage. How Tivoli Storage Manager can take advantage of this snapshot function varies depending on the platform.

Tivoli Storage Manager is able to support the most virtual architectures on the market today. We can protect the data in the virtual environment at different levels depending on the virtual layer used. We divide the protection into three levels:

- Off-host backup of virtual machines

  Tivoli Storage Manager for Virtual Environments provides off-host backup capabilities for VMware environments using the VMware vStorage APIs for Data Protection (VADP). A single backup can be used for both complete virtual machine (VM) recoveries, and file-level recovery. The capabilities of Tivoli Storage Manager for Virtual Environments are further enhanced to provide additional scalability. This includes new capabilities to allow for an incremental-forever backup model, and the ability to use multiple sessions to back up virtual machines in parallel.

- Host-level-based image backup

  Host-level-based image backup means installing the Tivoli Storage Manager backup-archive client on the hypervisor and running the backup of the VM from here. Depending on the options in the hypervisor, we are able to back up the VM while online or offline.

  Tivoli Storage Manager backup-archive client provides integration with the Hyper-V manager to back up and restore snapshots of Hyper-V virtual machines. We exploit the Microsoft Volume Shadow Copy Services (VSS) interface in Hyper-V to make full-vm snapshots. Snapshots can be taken whether the virtual machine is running or stopped. If the virtual machine is running when the snapshot is taken no downtime is involved to create the snapshot.

  It is possible to install the Tivoli Storage Manager backup-archive client on other hypervisors if the operating system platform is supported. For more information see the appropriate documentation from the supplier of the hypervisor.

- In-guest file-level backup

  In all virtual environments where the Tivoli Storage Manager backup-archive client is supported on the operating system level of the VM, we are able to back up the data with the standard progressive incremental forever capability. This is implemented as you would on any other physical host. Disaster recovery of the VM is somewhat more complex with this implementation.

3.1.7 Archive data

This section covers only the backup-archive function available to the Tivoli Storage Manager client.

The Tivoli Storage Manager client **`archive`** command archives a single file, selected files, or all files in a directory and its subdirectories on a server. Archive files that you want to preserve in their current condition. To release storage space on your workstation, delete files as you archive them by using the **`deletefiles`** option. Retrieve the archived files to your workstation when you need them again.

Archiving is a special kind of data protection, sometimes assigned the terms of long time retention and compliant archiving. But in the context of Tivoli Storage Manager, the archive function is sometimes also used to protect a special kind of data such as database log files or SAP application data in a specific way. Archive can also be treated as data protection of important data that must be stored in a secure way and for a long time.

Retention

Backup files are stored and retained based on versioning criteria; archived objects are stored and retained on a number of days basis.

There is a possibility to extend the retention period by using a mechanism called *deletion hold*. If a hold is placed on an object through the client API, the object is not deleted until the hold is released. See the *Backup-Archive Clients Installation and User's Guide* for your operating system platform:

`http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.0/com.ibm.itsm.client.doc/r_pdf_clients.html?lang=en`

There is no limit to how often you alternate holding and releasing an object. An object can have only one hold on it at a time, so if you attempt to hold an object that is already held, you will get an error message.

If an object with event-based policy is on hold, an event can still be signalled. The hold will not extend the retention period for an object. If the retention period that is specified in the **RETVER** and **RETMIN** parameters expires while the object is on hold, the object becomes eligible for deletion when the hold is released.

If an object is held, it is not deleted whether or not data retention protection is active. If an object is not held, it is handled according to existing processing such as normal expiration, data retention protection, or event-based retention. Data that is in deletion hold status can be exported. The hold status will be preserved when the data is imported to another system.

3.1.8 Application or database backup

Application or database data protection is based on backup and archive object protection mechanism, from a retention and storage perspective. The Tivoli Storage Manager family offers modules called Data Protection for Application or Database that allow you to protect the data without downtime on your environment.

See 2.3, “Tivoli Storage Manager recent enhancements” on page 29 for a brief introduction to the features in Tivoli Storage Manager introduced in the Tivoli Storage Manager V6.4 and Tivoli Storage Manager V7.1 releases.

3.1.9 Hierarchical storage management

Hierarchical storage management (HSM) is a data storage system that automatically moves data between high-cost and low-cost storage media. HSM exists because high-speed storage devices, such as hard disk drives, are more expensive per byte stored than slower devices, such as magnetic tape drives. While it would be ideal to have all data available on high-speed devices all the time, this is prohibitively expensive for many organizations. Instead, you can use HSM to store the bulk of your enterprise’s data on slower devices, and then copy data to faster disk drives only when needed.

Additional information about HSM and Space Management is in Appendix B, “Hierarchical storage management (HSM)” on page 323.

3.1.10 Protect data storage devices owning data

Tivoli Storage Manager client and features are also available on platforms to protect the data that is not managed by traditional operating systems, such as Microsoft Windows or UNIX systems.

This section describes available tools to protect data on NAS and Storwize devices. Also included is an explanation on how the Tivoli Storage Manager client protects files stored within GPFS by using the FlashCopy mechanism at the hardware level.

NAS and NDMP backup

For NAS devices, Tivoli Storage Manager Extended Edition uses Network Data Management Protocol (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. Alternatively, NDMP can be used to send data over TCP/IP from the file server to the Tivoli Storage Manager server for storage in the Tivoli Storage Manager storage hierarchy. NDMP enables a full and differential file system image backup and restore of Network Appliance file servers with OS Data ONTAP V7.1 or later, 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.

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

Snapshot Difference API

Snapshot Difference (SnapDiff) Incremental Backup can be used for performing efficient file level backups of NetApp and IBM System Storage N series file servers with very large file systems. It is available on Windows, Linux, and AIX Backup-Archive clients.

The NetApp filer has an embedded operating system. Tivoli Storage Manager backup is done either by allowing the backup-archive client to access files over the network (NFS/CIFS) or through NDMP. This causes major performance problems in the case of big filers with large file systems, as the backup cannot be completed within the backup window.

Snapshot Difference incremental backup attempts to address this problem. Snapshot Difference incremental backup takes the difference between two snapshots and backs up to the Tivoli Storage Manager Server all the files that have been changed between the two. Major performance improvement is achieved as the file system scan over the NAS network is eliminated in addition to the elimination of the scan of the Tivoli Storage Manager server database to determine the list of files to be backed up.

NetApp provides an API (Snapshot Difference API) that takes two snapshots, the base snapshot and the difference snapshot, as input and returns a list of all files that have been added, deleted, modified, and renamed between the two snapshots. The API handles files, directories, ACLs, streams, hard links and symbolic links, and supports both traditional and FlexVol volumes.

SnapMirror to Tape

Tivoli Storage Manager V6 leverages NetApp SnapMirror to Tape technology, enabling data centers to use Tivoli Storage Manager Network Data Management Protocol (NDMP) support to back up FAS and IBM N series volumes directly to tape or across the network to any storage device managed by Tivoli Storage Manager. This can be done without having to scan each volume to find new and changed files, helping reduce the amount of time customers need to back up their FAS and IBM N series systems, which can make the backup process much faster.

You use the NDMP SnapMirror to Tape feature as a disaster recovery option for copying very large NetApp file systems to secondary storage. With a parameter option on the **BACKUP** and **RESTORE NODE** commands, you can back up and restore file systems using SnapMirror to Tape. Consider the following guidelines before you use it as a backup method:

- You cannot initiate a SnapMirror to Tape backup or restore operation from the Tivoli Storage Manager Web client, command-line client, or the Operation Center.
- You cannot perform differential backups of SnapMirror images.
- You cannot perform a directory-level backup using SnapMirror to Tape, because Tivoli Storage Manager does not permit a SMTape backup operation on a server virtual file space.
- You cannot perform an NDMP file-level restore operation from SnapMirror to Tape images. Therefore, a table of contents is never created during SnapMirror to Tape image backups.

SnapMirror to Tape is the only solution, which will preserve the deduplication behavior of the primary storage during backup. The function is primarily a disaster protection solution and because of ONTAP operating system dependencies, not a long time archive option.

Storwize V7000 Unified embedded Tivoli Storage Manager client

The Storwize V7000 Unified system is designed for file and block storage provisioning with the capacity to store a large amount of file data on the file modules. Obviously, working with large amounts of data and a large number of files requires a longer backup duration. Today there is a requirement to reduce backup duration. To meet this requirement, the backup facility in the Storwize V7000 Unified system file modules has been optimized by providing a fully-integrated Tivoli Storage Manager client within file modules.

In a conventional backup method, the backup software can back up data from the shares exported from Storwize V7000 Unified system file modules mounted on the NFS or CIFS clients. Backup software is installed on each NFS or CIFS client and data is backed up over the network to the backup server.

The integrated Tivoli Storage Manager solution with Storwize V7000 Unified system file modules provides the following advantages:

- Tivoli Storage Manager client is pre installed on the file modules of the Storwize V7000 Unified system. Data is directly backed up from file modules to the backup server. Storwize V7000 Unified system file modules have an integrated IBM General Parallel File System (IBM GPFS) policy scan engine to scan the inode table of a file system, which is much faster than the conventional system calls, reducing the overall time required for the backup.
- Storwize V7000 Unified system leverages both its file modules for backing up the data to improve performance and to distribute load across both the file modules.
- Storwize V7000 Unified system file modules can establish several sessions with a Tivoli Storage Manager server.
- An existing Tivoli Storage Manager server environment can be leveraged to back up Storwize V7000 Unified system.
- Multiple Tivoli Storage Manager servers can be configured to back up different GPFS file systems to spread backup workload across multiple backup servers.

IBM General Parallel File System (GPFS)

IBM General Parallel File System (GPFS) is a high-performance and shared-disk file system on AIX, Linux and Windows. It can provide data access from nodes configured in a cluster environment. A GPFS system can hold millions of files, which potentially can challenge the backup processing time. The Tivoli Storage Manager client can interact with GPFS in different ways to support the protection of the data held by the file system. Or GPFS can be used as disk storage for storage pool volumes managed by the Tivoli Storage Manager server.

Note: The use of Tivoli Storage Manager backup-archive client and Space Management client in combination with GPFS is supported only on the AIX and Linux platforms.

For more information visit the GPFS wiki:

http://www.ibm.com/developerworks/wikis/display/hpccentral/General+Parallel+File+System+%28GPFS%29

Client file level backup

The following list describes client file level backup:

- Standard functions in the backup-archive client

  Tivoli Storage Manager backup-archive client works with the GPFS **mmbackup** command to identify files that need to be backed up. See “Solution description using mmbackup” on page 285 for more details.

- Storwize V7000 Unified and SONAS

  In the conventional backup method, the Tivoli Storage Manager client backs up data from the shares exported from the SONAS or the Storwize V7000 system file modules mounted on the NFS or CIFS clients. The client software is installed on each NFS or CIFS client and data is backed up over the network to the Tivoli Storage Manager server. This is not the optimal way to traverse millions of files because the backup time is too long.

  To circumvent this, the Tivoli Storage Manager backup-archive client and Tivoli Storage Manager for Space Management client are embedded in the SONAS and Storwize v7000 systems. The GPFS policy engine is used to scan for candidates to back up or migrate. Data is transferred across the LAN to the Tivoli Storage Manager server. Currently the LAN-free data transfer is not supported for Storwize V7000 Unified or for SONAS. Normally this is not required because the backup data transfer might already be optimized with the embedded client. But if the data volume requires LAN-free data transfer. a host with access to the data must have the Tivoli Storage Manager backup-archive client and Storage Agent installed to perform the backup. Chapter 7, “Protecting your data with Tivoli Storage Manager” on page 141.

Note: A suggestion is to use the same Tivoli Storage Manager Server for the backup-archive client and the Space Management client. In this way, you can take advantage of the integration between the clients for inline copy and stub restore.

Snapshot at the storage hardware layer

Snapshot or FlashCopy is a method by which you can create exact copies of disks (LUNs) in a storage system. The copies may be used for backups, reporting, or disaster recovery purposes. FlashCopy is a term that is used by IBM storage systems. FlashCopy creates a mirror of LUNs at the storage layer and makes that copy available in seconds. This is accomplished by the storage cache, bitmaps, and read/write algorithms that are built into the storage system.

The two main types of storage snapshots are copy-on-write snapshot and split-mirror snapshot. Utilities are available that can automatically generate either type. For more information, see the following website:

http://www.ibm.com/developerworks/tivoli/library/t-snaptsm1/index.html

The FlashCopy is crash-consistent on the operating system level. For application consistency the Tivoli Storage FlashCopy Manager has a built-in capability to support several commonly used applications. On the Microsoft platform, we the Volume Shadow Copy Service to make consistent application snapshots. See 3.1.4, “Microsoft Volume Shadow Copy Service” on page 42 for more information.

Figure 3-3 shows the integration of applications and IBM storage systems with Tivoli Storage Manager. Tivoli Storage FlashCopy Manager ties it all together.

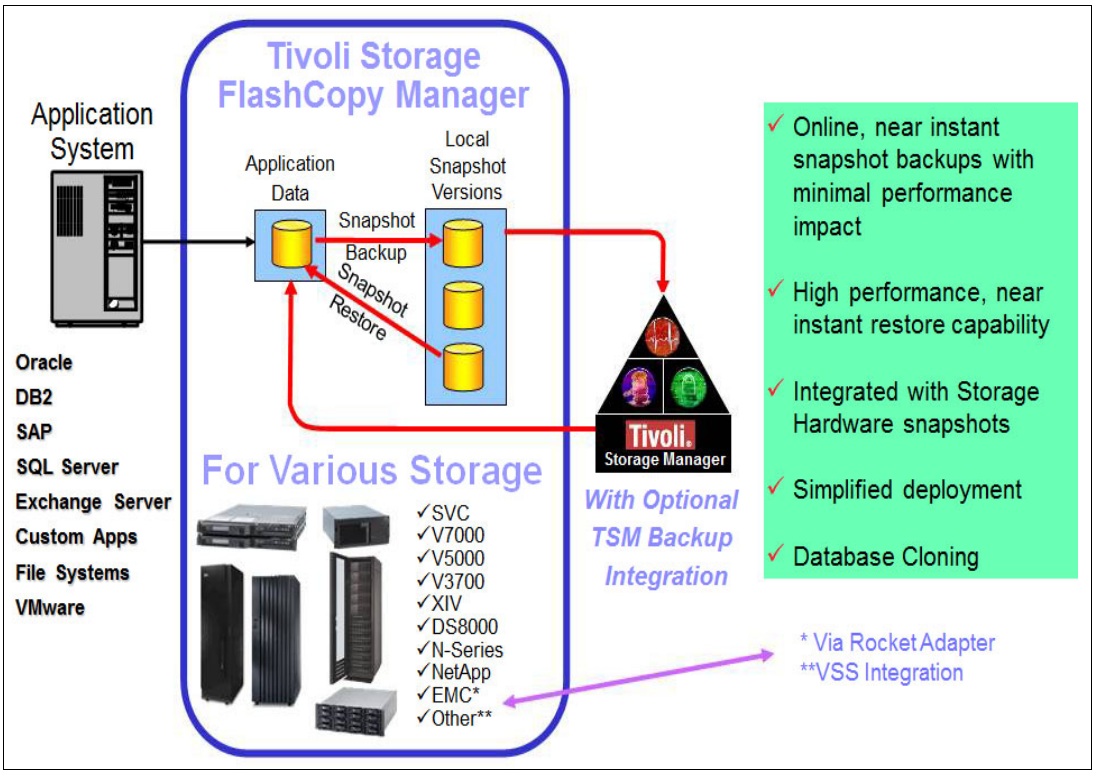

Figure 3-3 Integrate applications and IBM storage: Tivoli Storage FlashCopy Manager

With Tivoli Storage FlashCopy Manager, we can perform near-instant application aware snapshot backups, with minimal performance impact for IBM DB2, Oracle, SAP, Microsoft SQL Server and Exchange, and VMware virtual machines. On the storage system side, it integrates with IBM Storwize family, IBM System Storage DS8000, IBM System Storage SAN Volume Controller, XIV Storage System, IBM N series, and NetApp. We support FlashCopy on AIX, Solaris, HP-UX, Linux, and Microsoft Windows. More information about platform requirements is at the following address:

http://www.ibm.com/support/docview.wss?uid=swg21427692

Tivoli Storage FlashCopy Manager can now be used with other storage vendors such as EMC by leveraging the Rocket Device Adapter for IBM Tivoli Storage FlashCopy Manager. The Rocket Device Adapter is a device agent interface. Rocket Device Adapters for Tivoli Storage FlashCopy Manager extend backup and restore operations seamlessly to the native operations of EMC systems

Tivoli Storage FlashCopy Manager can be configured either as a stand-alone solution where FlashCopies are stored only in the storage system or as a Tivoli Storage Manager integrated solution where the FlashCopy is transferred to the Tivoli Storage Manager server storage hierarchy.

Another implementation is to integrate Tivoli Storage FlashCopy Manager with Metro Mirror and Global Mirror functionality. This is available on AIX, HP-UX, Linux, and Solaris with IBM System Storage SAN Volume Controller, IBM System Storage Storwize V7000, and IBM XIV Storage Systems. This provides application consistent backup and restore at a remote facility for improved disaster recovery.

3.2 Recover data from the client

In this section, we discuss methods to restore data from the client.

3.2.1 Restore data

To restore a file, a directory, or even a whole machine, you need to know two pieces of information:

- What you want to restore (file name, directory)
- Optional: from when (point in time) you want to restore an object 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 gets the location of the objects to restore from its database.

You can restore using the command-line interface (CLI), GUI, or web browser interface. The browser interface enables you to initiate the restore remotely so that you do not have to be physically located at the system being restored.

Point-in-time

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 on the file space, directory, or file level. You must specify a sufficiently long retention period in the management class to enable this to occur. To provide a point-in-time restore capability for, say, up to one month previously, set the following parameters:

- **VEREXISTS** and **VERDELETED** to NOLIMIT
- **RETEXTRA** and **RETONLY** to at least 31 days

This way, the number of times that the files change in the restorable period does not matter because you will always have enough versions stored to be able to perform the restore.

Restore from an active-data pool

Active-data pools are storage pools that contain only active versions of client backup data and can be used to optimize client restore operations. As updated versions of backup data continue to be stored in active-data pools, older versions are deactivated and removed during reclamation processing. Active-data pools associated with a FILE device class are ideal for fast client restores because FILE volumes do not have to be physically mounted and because the server does not have to position past inactive files that do not have to be restored. In addition, client sessions restoring from FILE volumes in an active-data pool can access the volumes concurrently, which also improves restore performance.

Data migrated by hierarchical storage management (HSM) clients and archive data are not permitted in active data pools. As updated versions of backup data continue to be stored in active-data pools, older versions are deactivated and removed during reclamation processing.

Find more information about active data pools, see “Active-data pools” on page 60.

Snapshot restore

Snapshot capabilities are achieved through the use of advanced storage-specific hardware snapshot technology to help create a high performance, low impact application data protection solution. To deliver data protection for business-critical databases restores in minutes instead of hours.

Instant restore

With instant restore, you can restore a volume and almost immediately use the restored volume. To the user or application, the volume appears to contain all the restored data, even though the restore process is in-progress. Less downtime is required before a recovered volume can be used because you can use data on the disk while the restore is in progress.

3.2.2 Retrieve

The retrieve operation obtains copies of archived files from the Tivoli Storage Manager server. You can specify either selected files or whole directories to retrieve archived files. The description option enables you to search for the descriptions assigned to the archive package when it was made; you may decide to put the files into the same directory from which they were archived, or into a different directory.

3.2.3 Recall

This is a brief description of the recall function to bring back a migrated file from IBM Tivoli Storage Manager to its original place on the local file system.

Recall using Tivoli Storage Manager for Space Management

As with the classic HSM solutions, the HSM Client for Windows or the Space Management client for UNIX provides transparent and selective ways for you to bring back a migrated file from IBM Tivoli Storage Manager to its original place on the local file system.

For more information, see Appendix B, “Hierarchical storage management (HSM)” on page 323.

3.3 Data protection at server side

This section lists available tools and features on Tivoli Storage Manager side to help you to protect client’s data.

3.3.1 Disaster recovery with traditional methods

In this section, we look at the methods for disaster recovery with Tivoli Storage Manager.

Disaster Recovery Manager

Disaster recovery is the process of restoring operations of a business or organization in the event of a catastrophe. There may be many aspects related to the restoration, including facilities, equipment, personnel, supplies, customer services, and data. One of the most valuable business assets is the critical data that resides on the computer systems throughout the company. The recovery of this data must be a primary focus of the disaster recovery plan. This topic discusses how Tivoli Storage Manager can help you to avoid losing data in case of disaster.

Tivoli Storage Manager provides disaster recovery of the Tivoli Storage Manager server with its Disaster Recovery Manager function. Disaster Recovery Manager offers various options to configure, control, and automatically generate a disaster recovery plan containing the information, scripts, and procedures to automate recovery of the Tivoli Storage Manager server, and help ensure quick recovery of data after a disaster.

For more information about Disaster Recovery Manager, see “Disaster recovery management” on page 99.

Electronic tape vaulting

Using Tivoli Storage Manager with electronic tape vaulting provides extra data protection capabilities, with backups made to remote tape drives over communication links. Electronic tape vaulting can enable shorter recovery times and reduced data loss if the server is damaged. An electronic tape vaulting solution combined with Tivoli Storage Manager is fundamental to achieving Tier 3 and above RPO and RTOs, that is, less than 24 hours.

With electronic tape vaulting, the Tivoli Storage Manager server will have an alternate location to store primary and copy storage pools as though they are directly attached. The Tivoli Storage Manager server can first write a copy of disk storage pool data to tape pools at the remote site (Datacenter #2), then the data can be migrated to the tape storage pools at the primary site (Datacenter #1). See Figure 3-4.

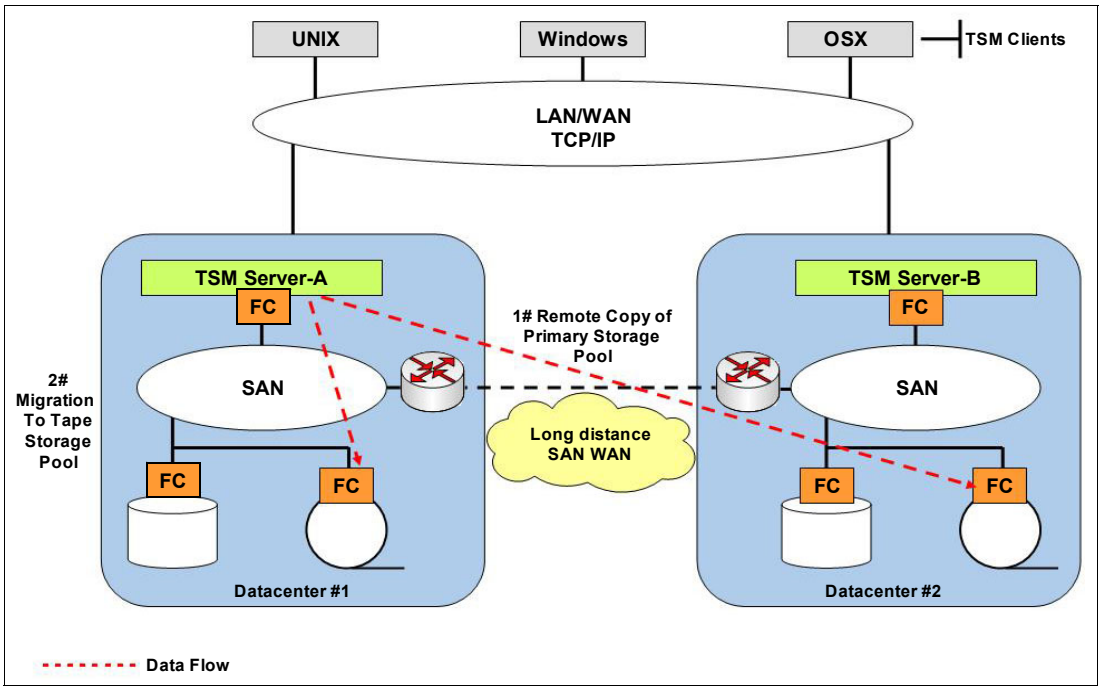

Figure 3-4 Tivoli Storage Manager and electronic tape vaulting

Depending on your configuration (and whether remote disk replication is being used in conjunction with electronic tape vaulting), you might choose to back up the Tivoli Storage Manager database and configuration files by this method. This ensures a copy of the data is stored at both sites and that the Tivoli Storage Manager server can rapidly recover at the remote site. If remote disk replication is used for mirroring of the Tivoli Storage Manager database and storage pools, the Tivoli Storage Manager server can be recovered quickly, without any loss of client data. A peer-to-peer configuration can be used to balance the load of Tivoli Storage Manager services in the enterprise and provide data protection and rapid recovery for a failure at either site. A major consideration with electronic vaulting is synchronization of metadata in the database with data in the storage pool. If the metadata is replicated before the data, invalid database references might occur at the target site.

Using electronic tape vaulting with Tivoli Storage Manager offers these advantages:

- Critical data can be frequently and rapidly vaulted to remote offsite locations.
- Physical tape handling, which can result in damaged tapes, lost tapes, tapes delayed in transit, or data that is sabotaged leading to increased reliability and security, is eliminated.
- Costs associated with couriers and offsite vaulting vendors are eliminated.
- Government offsite vaulting regulations are satisfied.
- Cost of downtime and storage management is lowered.
- Peer solutions eliminate or reduce costs associated with hot-site service providers.

Remote disk mirroring and tape vaulting solutions

Various technologies exist for remote disk mirroring and electronic tape vaulting:

- Long distance SANs
- Dense wavelength division multiplexing (DWDM)
- Fibre extenders
- WAN-based channel extension using telco and IP protocols
- NAS and iSCSI gateways.

Table 3-1 summarizes several of these technologies. The use of such technologies also might depend on a particular vendor's replication or vaulting solution. For example, the IBM DS8000 uses PPRC to achieve data replication. Metro Mirror is supported by ESCON links, which can be further extended by DWDM or WAN channel extension.

Table 3-1 Extended electronic vaulting technologies

Electronic vaulting technology	Common supported distances between sites	Common product vendors	Relative technology costs
Long Distance SAN (shortwave/longwave Fibre Channel)	Up to 30 km	Brocade (OEM) McDATA (OEM)	Low
CWDM, DWDM and fiber extenders	Up to 300 km	Cisco Nortel Ciena Finisar CNT	Medium
WAN, IP Based Routers and Channel Extension	Thousands of kilometers	CNT Cisco (OEM)	Low to high (depending on solution)

Storage area networks (SANs) overcome the distance limitations of other storage channel protocols, such as SCSI. Longwave laser GBICs (available on most SAN hubs), switches, and directors enable a transmission distance of up to 10 kilometers (11 kilometers when you include switch to host connections) when used with 9-micron diameter single-mode optical fiber. Shortwave GBICs use multi-mode fiber and is the ideal choice for shorter distance (less than 500 meters from transmitter to receiver or vice versa).

CWDM or DWDM is a way to open up the conventional optical fiber bandwidth by breaking it up into many channels, each at a particular optical wavelength (a unique color of light). Each wavelength can carry a signal at any bit rate less than an upper limit defined by the electronics, typically up to several gigabits per second. Dense wavelength division multiplexing (DWDMs) are implemented in areas that have dark fiber available through telco and service providers. The DWDM is deployed as part of the physical layer. It is therefore

independent of protocol, simply passing signal information in the format it is received. Examples of the protocols it can support are ATM, Gigabit Ethernet, ESCON, FICON, and Fibre Channel.

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

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

For a detailed overview of technologies, products, costs, and preferred practices with distance solutions, see *Introduction to Storage Area Networks and System Networking*, SG24-5470.

Collocation considerations for offsite vaulting

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

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

- Taking more partially filled volumes offsite, thereby increasing the reclamation activity when the reclamation threshold is lowered or reached
- Leaving these partially filled volumes onsite until they fill and risk not having an offsite copy of the data on these volumes.

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

Reclamation considerations for offsite vaulting

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

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

The **REUSEDELAY** parameter specifies the number of days that must elapse before a volume can be reused or returned to scratch status after all files are expired, deleted, or moved from the volume.

The server reclaims an offsite volume that has reached the reclamation threshold as follows:

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

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

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

If you send copy storage pool volumes offsite, the best approach is to control copy storage pool reclamation by using the default value of 100. This turns off reclamation for the copy

storage pool. You can start reclamation processing at desired times by changing the reclamation threshold for the storage pool.

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

When you perform reclamation for offsite volumes, use the following sequence:

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

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

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

3.3.2 Disaster recovery with node replication

Electronic vaulting is an older way to protect a Tivoli Storage Manager server for disaster recovery over a large distance. Node replication is the newer process of incrementally copying or replicating client node data from one Tivoli Storage Manager server to another for the purpose of disaster recovery.

We discuss this function in detail in 6.3, “Data protection solution using node replication feature” on page 128.

3.3.3 Server-side data deduplication

Server-side data deduplication in Tivoli Storage Manager is a two-phase process:

1. In the first phase, duplicate data is identified.
2. During the second phase, duplicate data is removed by certain server processes, such as reclamation processing of storage-pool volumes.

By default, a duplicate-identification process begins automatically after you define a storage pool for deduplication. (If you specify a duplicate-identification process when you update a storage pool, it also starts automatically.) Because duplication identification requires extra disk I/O and CPU resources, Tivoli Storage Manager lets you control when identification begins and the number and duration of processes.

You can deduplicate any type of data except encrypted data. You can deduplicate client backup and archive data, Tivoli Data Protection data, and so on. Tivoli Storage Manager can deduplicate whole files as well as files that are members of an aggregate. You can deduplicate data that has already been stored. No additional backup, archive, or migration is required.

Figure 3-5 is an overview of the deduplication process, where, in this case, the deduplication is done on the server. See 6.2, “Disk-to-disk data protection solution using deduplication” on page 113.

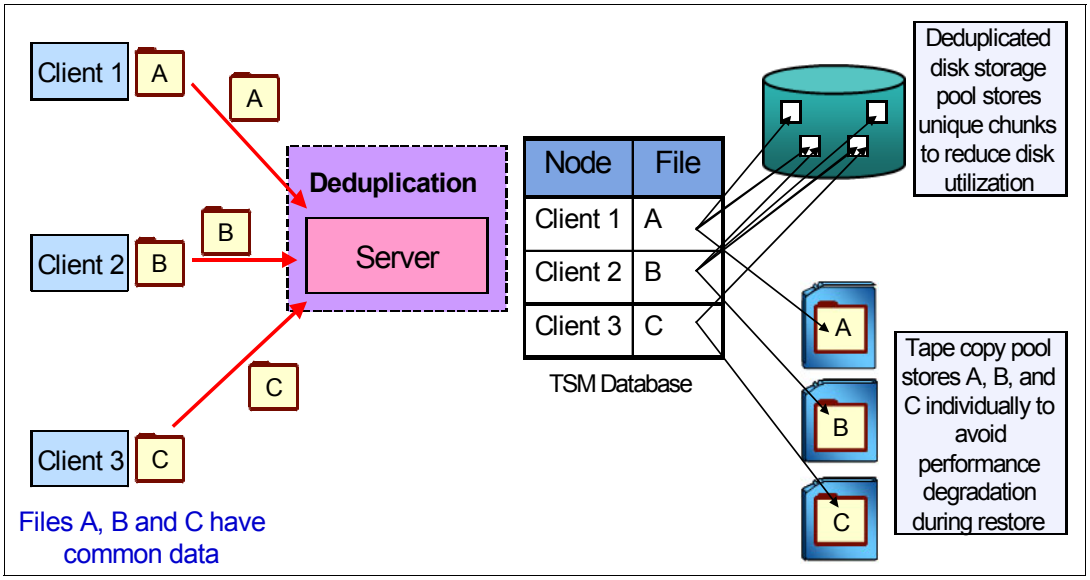

Figure 3-5 Deduplication overview

3.3.4 Storage hierarchy, storage pool features

This section describes how storage hierarchy (storage pools) can help you protect your data. The Tivoli Storage Manager differs between storage pool types:

- Primary
- Copy
- Active storage pool

If the storage pool is configured with a sequential file device class, it is eligible for storage pool deduplication. Data deduplication and its benefits are described in 6.2, “Disk-to-disk data protection solution using deduplication” on page 113.

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 wherever possible. Stored objects are only restored from copy storage pools when the expected primary storage pool volume is unavailable.

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

Copy storage pools

A copy storage pool provides an extra level of protection for client data and is created for the express purpose of backing up a primary storage pool, see Figure 3-6 for sample backup storage pool processes.

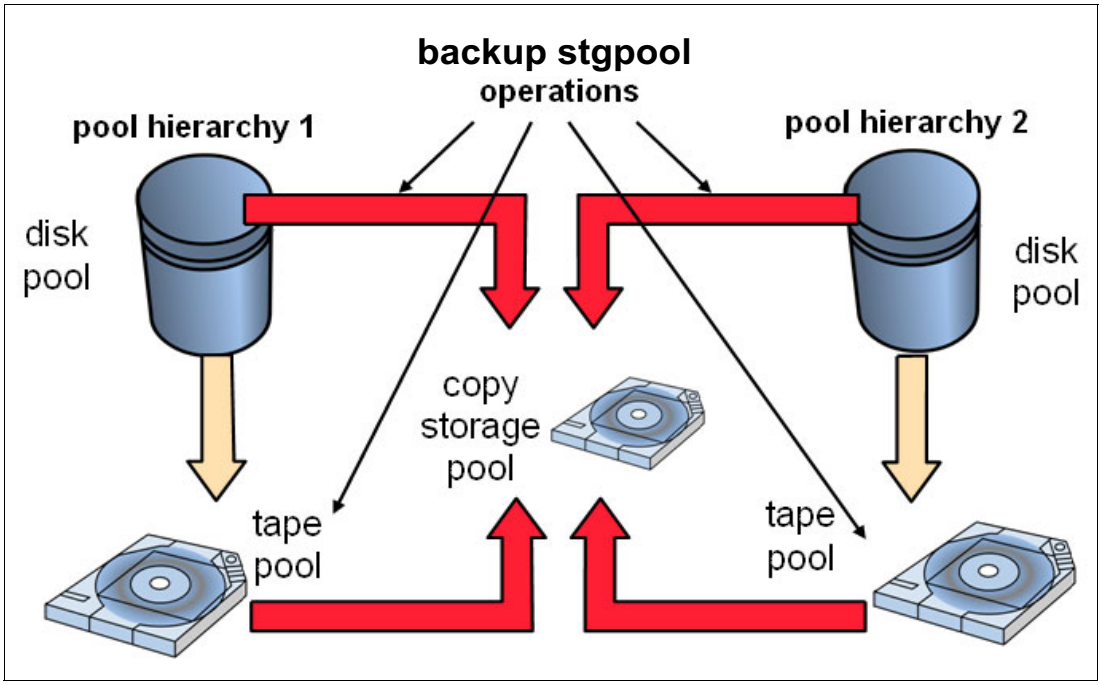

Figure 3-6 Backup storage pool to protect primary storage pools

The copy storage pool may contain copies of all files that appear in the primary storage pool. A copy storage pool provides recovery from partial and complete failures of a primary storage pool.

A partial failure would be, for example, if a single tape in a primary storage pool is damaged by a failing drive, or simply has too many read or write errors.

When backing up your storage pools, you might want to keep your copies local, or remote. Figure 3-7 shows a sample hierarchy.

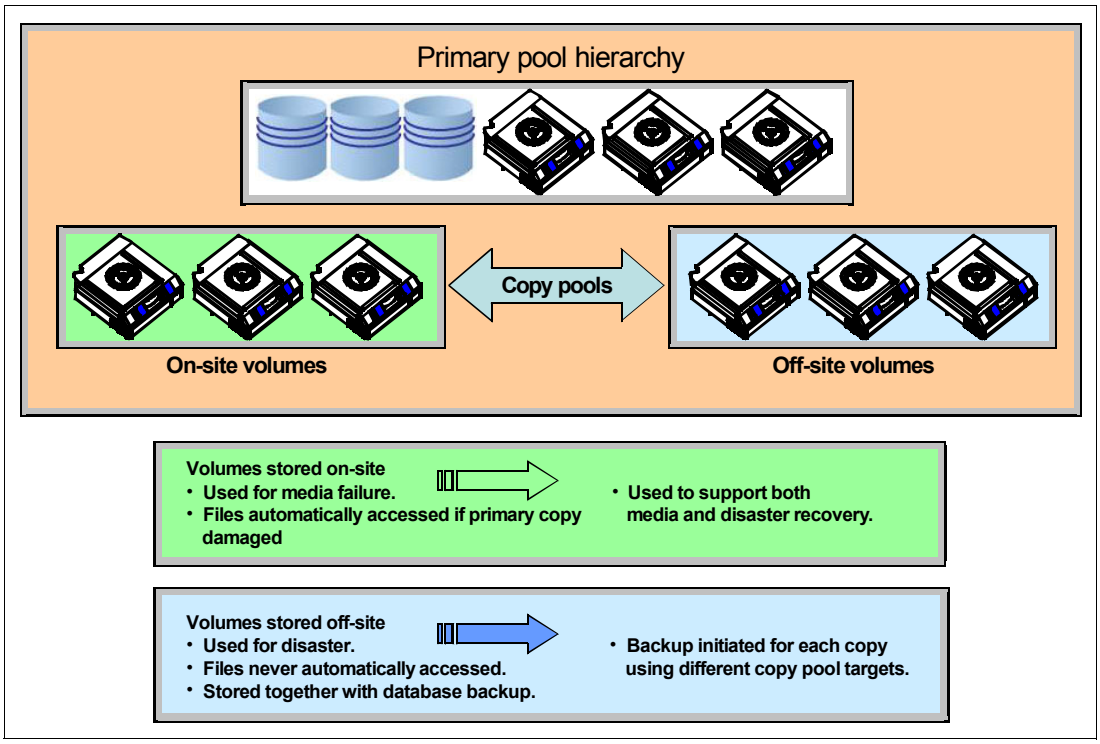

Figure 3-7 Pool hierarchy

Onsite volumes support both, media and disaster recovery. If a media failure is detected for a primary storage pool volume, objects will be automatically accessed using the copy storage pool volume.

Offsite volumes are used for disaster recovery. If a media failure is detected for a primary storage pool volume and no onsite copy volume is available, offsite volumes need to be retrieved at the local site for the object to become accessible again. Typically they are stored at the offsite location together with a database backup to allow for a complete recovery of the server in the event of a disaster.

Active-data pools

With active-data pools (ADP), Tivoli Storage Manager administrators can segregate active and inactive files in the servers storage hierarchy. With ADP, you can organize active data that is likely to be restored in a way that the restore operations do not need to access tapes containing inactive data. If you keep many versions of files, searching tapes for active versions of files that are mixed with inactive versions can take a long time. By storing only the active versions together on tapes, the time spent to position past the inactive files is eliminated. For Tivoli Storage Manager, this means keeping only the active versions of files in a storage pool strictly defined for such a purpose.

Setting up an active-data pool depends on whether your intended use is for faster client restore or onsite and offsite storage of active data. Active-data pools can use any type of sequential-access storage (for example, a tape device class or FILE device class). However, the precise benefits of an active-data pool depends on the specific device type associated with the pool. As mentioned in “Snapshot restore” on page 51, active-data pools associated with a FILE device class are ideal for fast client restores because FILE volumes do not have to be physically mounted and because the server does not have to position past inactive files

that do not have to be restored. In addition, client sessions restoring from FILE volumes in an active-data |pool can access the volumes concurrently, which also improves restore performance.

Active-data pools that use removable media, such as tape or optical, offer similar benefits. Although tapes need to be mounted, the server does not have to position past inactive files. However, the primary benefit of using removable media in active-data pools is the reduction of the number of volumes used for onsite and offsite storage. If you vault data electronically to a remote location, an active-data pool associated with a SERVER device class lets you save bandwidth by copying and restoring only active data.

You can also use active data pool volumes, electronically vaulted, at an offsite location for increased protection of your data in case of a server disaster.

Storage pool caching

Caching is an storage pool option that allows the migration process to leave behind duplicate copies of files on disk after the server migrates these files to the next storage pool in the storage hierarchy. The copies remain in the random disk storage pool, but in a cached state, so that subsequent retrieval requests can be satisfied quickly. However, if space is needed to store new data in the disk storage pool, cached files are erased and the space they occupied is used for the new data. Caching applies only to random-access (disk) pools and not to sequential-access disk (file) pools.

When cache is disabled and migration occurs, the server migrates the files to the next storage pool and erases the files from the disk storage pool. By default, the system disables caching for each disk storage pool because of the potential effects of cache on backup performance.

The advantage of using cache for a disk storage pool is that cache can improve how quickly the server retrieves some files. When you use cache, a copy of the file remains on disk storage after the server migrates the primary file to another storage pool. You may want to consider using a disk storage pool with cache enabled for storing space-managed files that are frequently accessed by clients.

Simultaneous write to primary, copy, and active-data pools

With IBM Tivoli Storage Manager, you can write data simultaneously to a primary storage pool, copy storage pools, and active-data pools. The simultaneous-write function increases your level of data protection and reduces the amount of time required for storage pool backup. You can write data simultaneously during any of the following operations:

- Client store sessions, as in these examples:
  - Backup and archive sessions by Tivoli Storage Manager backup-archive clients.
  - Backup and archive sessions by application clients using Tivoli Storage Manager API.
  - Migration processes by hierarchical storage management (HSM) clients.
  - Migrated data is simultaneously written only to copy storage pools. Migrated data is not permitted in active-data pools.
- Server migration of data within a storage pool hierarchy.
- Server import processes that involve copying exported file data from external media to a primary storage pool that is configured for the simultaneous-write function. Imported data is simultaneously written to copy storage pools. Imported data is not simultaneously written to active-data pools. To store newly imported data into an active-data pool, use the **COPY ACTIVEDATA** command.

The maximum number of copy storage pools and active-data pools to which data can be simultaneously written is three. For example, you can write data simultaneously to three copy storage pools, or you can write data simultaneously to two copy storage pools and one active-data pool.

Figure 3-8 shows how simultaneous writes to primary pools and up to three copy pools and active-data pools occur during the following operations:

- Client backup, archive, and space-management operations
- Server import operations
- Storage pool migration

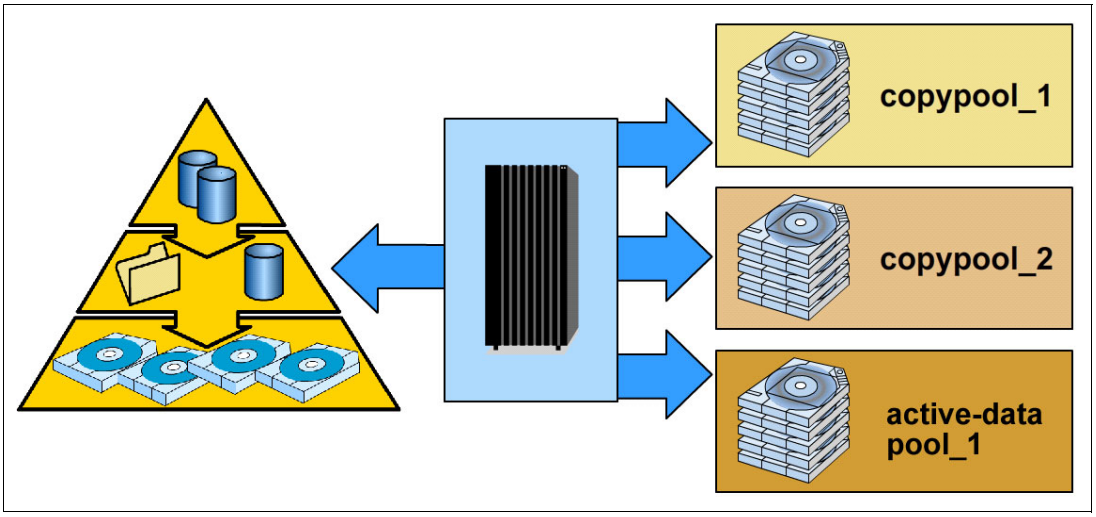

Figure 3-8 Simultaneous write to copy storage pools

With Tivoli Storage Manager, you can now write data simultaneously to copy storage pools and active-data pools during server data-migration processes.

The simultaneous-write function during migration can reduce the amount of time required to back up storage pools or copy active data. Data that is simultaneously written to copy storage pools or active-data pools during migration is not copied again to the copy storage pools or active-data pools. For example, suppose that you migrate all the data in your primary random-access disk storage pool nightly and then back up your primary storage pools. By using the simultaneous-write function during migration, you can significantly reduce the amount of time required for backup operations. You can also use the simultaneous-write function during migration if you have many client nodes and the number of mount points that are required to perform the simultaneous-write function during client store operations is unacceptable. If mounting and demounting tapes when writing data simultaneously during client store operations is taking too much time, consider writing data simultaneously during migration. With Tivoli Storage Manager V6.2, you can specify the simultaneous-write function for a primary storage pool if it is the target for any of the eligible operations (client store sessions, server import processes, and server data-migration processes).

Similar to client-side simultaneous write during data store operations, a copy storage pool list is created at the beginning of a process and is active for the life of the process. While the process is active, the copy storage pools in the list are written to.

Simultaneous write migration is not supported for the server data movements, such as reclamation, move data, move node data, or cloning.

The flow of data during simultaneous write migration is shown in Figure 3-9.

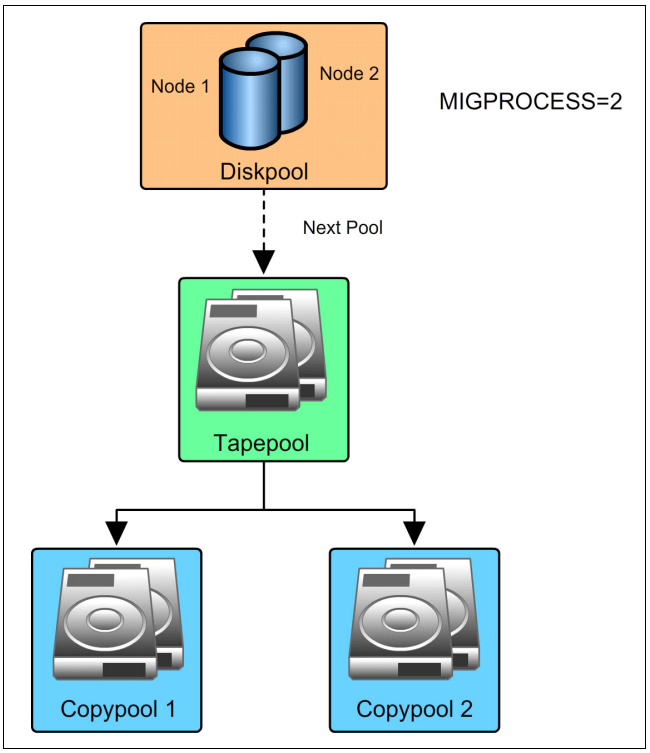

Figure 3-9 Simultaneous write migration overview

Simultaneous writes are not supported for LAN-free backups, or when a NAS backup is writing a TOC file.

Carefully consider the use of simultaneous writes. Because the data is written to the copy storage pool and primary storage pool simultaneously, the backup performance is only as good as the slowest device being used for any of the pools.

Server shared disk storage pool with GPFS

Tivoli Storage Manager supports the IBM GPFS as a shared disk storage pool and with LAN-free access. This is an alternative to sharing the tape library infrastructure among storage agents and servers.

GPFS provides the file and block locking necessary to share a physical device among multiple hosts. The shared disk storage pool allows for storage agents and servers to access FILE devclass storage from multiple physical systems on the LAN. This helps reduce server utilization and enables more concurrent operations by distributing them across multiple storage agents.

Concurrent access to volumes improves restore performance by allowing two or more client sessions, two or more storage agents, or a combination of client sessions and storage agents to access the same volume at the same time. Multiple client sessions and storage agents can read a FILE volume concurrently. In addition, one client session or storage agent can write to the volume while it is being read.

As part of the planning process, decide whether to use LAN-free data movement and whether you want to use client-side data deduplication, server-side deduplication, or both.

If you decide to use LAN-free data movement and both client-side and server-side data deduplication, use one of the following steps:

- For V6.1 or earlier storage agents, store client-side deduplicated data in a separate storage pool. Restore and retrieve deduplicated data from this storage pool over the LAN.
- Use LAN-free data movement to restore and retrieve data from storage pools that contain data that was deduplicated only by the server.

Traditionally, disaster recovery plans include daily offsite tape backups that are picked up from the local site and transported by a courier to a secure facility, which is often a tape vaulting service provider. Vaulting of tapes at offsite locations can provide a secure means to protect data in the event of a disaster at the primary site. To recover from a disaster, you must know the location of offsite recovery media. Tivoli Storage Manager Disaster Recovery Manager (DRM) helps determine which volumes to move offsite and back onsite and tracks the location of the volumes. Disaster Recovery Manager is included in Tivoli Storage Manager Extended Edition.

With tape vaulting, you can back up primary storage pools to a copy storage pool and then send the copy storage pool volumes offsite. You can track these copy storage pool volumes by changing their access mode to offsite, and updating the volume history to identify their location. If an offsite volume becomes expired, the server does not immediately return the volume to the scratch pool. The delay prevents the empty volumes from being deleted from the database, making it easier to determine which volumes must be returned to the onsite location. Disaster Recovery Manager handles all of this automatically.

3.3.5 Tiered storage

You can organize the server's storage pools into one or more hierarchical structures. This storage hierarchy allows flexibility in several ways. For example, you can set policy to have clients send their backup data to disks for faster backup operations, then later have the server automatically migrate the data to tape.

Storage pools will be arranged in a storage hierarchy, which consists of at least one primary storage pool to which a client node backs up, archives, or migrate data. Typically, data is stored initially in a disk storage pool for fast client restores, and then moved to a tape-based storage pool, which is slower to access but which has greater capacity and can be more cost efficient. The location of all data objects is automatically tracked within the server database.

You can set up your devices so that the server automatically moves data from one device to another, or one media type to another. The selection can be based on characteristics such as file size or storage capacity. A typical implementation might have a disk storage pool with a subordinate sequential (tape) storage pool. When a client backs up a file, the server might initially store the file on disk according to the policy for that file. Later, the server might move the file to tape when the disk becomes full. This action by the server is called *migration*. You can also place a size limit on files that are stored on disk, so that large files are stored initially on tape instead of on disk.

You can organize the server's storage pools into one or more hierarchical structures, as Figure 3-10 on page 65 shows.

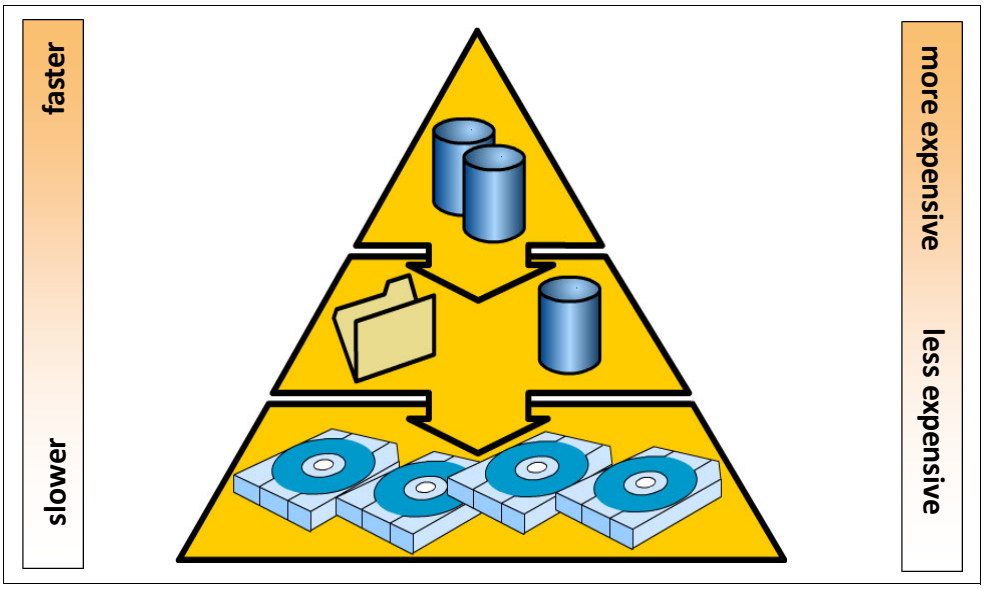

Figure 3-10 Storage Pools can be chained to create a storage hierarchy

Another option to consider for your storage pool hierarchy is IBM tape cartridges and drives, (3592, 1120, 1130 and 1140) which can be configured for an optimal combination of access time and storage capacity.

Migration of files from disk to sequential storage pool volumes is particularly useful because the server migrates all the files for a group of nodes or a single node together. This gives you partial collocation for clients. Migration of files is especially helpful if you decide not to enable collocation for sequential storage pools.

This storage hierarchy allows flexibility in a number of ways. For example, you can set policy to have clients send their backup data to disks for faster backup operations, then later have the server automatically migrate the data to tape.

Data hierarchies can be customized to suit business needs. As Figure 3-11 shows, server Grey's data can go directly from disk to tape. Server Cyan's data might require more complex retention to keep it available longer. This data might need to go from disk to file to tape.

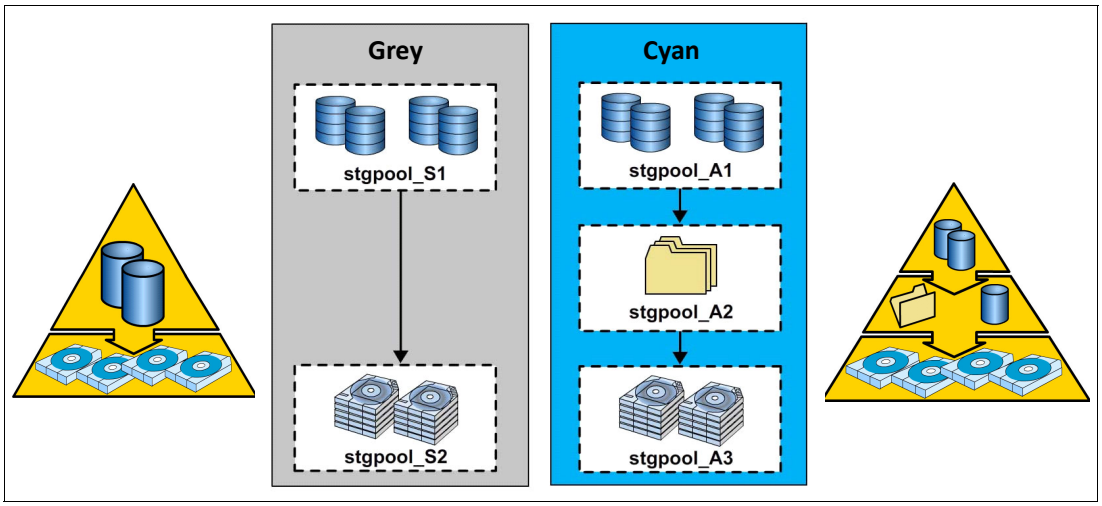

Figure 3-11 Data hierarchies can be customized to suit business needs

3.3.6 Advanced Tiered storage concept with virtual tape libraries

In a modern storage infrastructure, we have a virtualization layer for tape libraries and tape drives. In this case the physical storage represents random disk; the logical storage represents tape. This configuration gives us the opportunity to store data in multistream backups while we can restore them in a fast way from disk. Including the virtual volume storage tier in the traditional layout we are much more flexible to design a valuable storage concept. An extra advantage of virtual tape libraries (VTLs) like IBM TS7650 ProtecTIER is the built-in inline deduplication function, so that much more logical data can be stored than physical space is used.

Using a VTL, you can create variable numbers of drives and volumes because they are only logical entities within the VTL. The ability to create more drives and volumes increases the capability for parallelism, giving you more simultaneous mounts and tape I/O.

VTLs use SCSI and Fibre Channel interfaces to interact with applications. Because VTLs emulate tape drives, libraries, and volumes, an application such as Tivoli Storage Manager cannot distinguish a VTL from real tape hardware unless the library is identified as a VTL.

With enhancements available in Version 6.3, you can define a library as a virtual tape library (VTL) to Tivoli Storage Manager. VTLs primarily use disk subsystems to internally store data. Because they do not use tape media, you can exceed the capabilities of a physical tape library when using VTL storage. Using a VTL, you can define many volumes and drives which provides for greater flexibility in the storage environment and increases productivity by allowing more simultaneous mounts and tape I/O.

There are some considerations for defining a library as a virtual tape library, including enhancements for performance and setup of your hardware. Defining a VTL to the Tivoli Storage Manager server can help improve performance because the server handles mount point processing for VTLs differently than real tape libraries. The physical limitations for real tape hardware are not applicable to a VTL, affording options for better scalability.

You can use a VTL for any virtual tape library when the following conditions are *both true*:

- There is no mixed media involved in the VTL. Only one type and generation of drive and media is emulated in the library.
- Every server and storage agent with access to the VTL has paths that are defined for all drives in the library.

If either of these conditions is not met, any mount performance advantage from defining a VTL library to the Tivoli Storage Manager server can be reduced or negated. VTLs are compatible with earlier versions of both library clients and storage agents. The library client or storage agent is not affected by the type of library that is used for storage. If mixed media and path conditions are true for a SCSI library, it can be defined or updated as LIBTYPE=VTL.

Because VTLs do not have the physical limitations that real tape hardware does, their capacity for storage is more flexible. The concept of storage capacity in a virtual tape library is different from capacity in physical tape hardware. In a physical tape library, each volume has a defined capacity, and the library's capacity is defined in terms of the total number of volumes in the library. The capacity of a VTL, alternatively, is defined in terms of total available disk space. You can increase or decrease the number and size of volumes on disk. This variability affects what it means to run out of space in a VTL. For example, a volume in a VTL can run out of space before reaching its assigned capacity if the total underlying disk runs out of space. In this situation, the server can receive an end-of-volume message without any warning, resulting in backup failures. When out-of-space errors and backup failures occur, disk space is usually still available in the VTL. It is hidden in volumes that are not in

use. For example, volumes that are logically deleted or returned to scratch status in the Tivoli Storage Manager server are only deleted in the server database. The VTL is not notified, and the VTL maintains the full size of the volume as allocated in its capacity considerations.

To help prevent out-of-space errors, ensure that any SCSI library that you update to LIBTYPE=VTL is updated with the RELABELSCRATCH parameter set to YES. The RELABELSCRATCH option enables the server to overwrite the label for any volume that is deleted and to return the volume to scratch status in the library. The RELABELSCRATCH parameter defaults to YES for any library defined as a VTL.

Drive configuration in a VTL is variable, depending on the needs of your environment. Most VTL environments use as many drives as possible to maximize the number of concurrent tape operations. A single tape mount in a VTL environment is typically faster than a physical tape mount. However, using many drives increases the amount of time that the Tivoli Storage Manager server requires when a mount is requested. The selection process takes longer because the number of drives that are defined in a single library object in the server increases. Virtual tape mounts can take as long or longer than physical tape mounts depending on the number of drives in the VTL.

For best results, create VTLs with 300 - 500 drives each. If more drives are required, you can logically partition the VTL into multiple libraries and assign drives to each library. Operating system and SAN hardware configurations could impose limitations on the number of devices that can be utilized within the VTL library.

Figure 3-12 shows the advanced tiering approach with virtual tape libraries. See 6.4, “Tivoli Storage Manager together with ProtecTIER” on page 135.

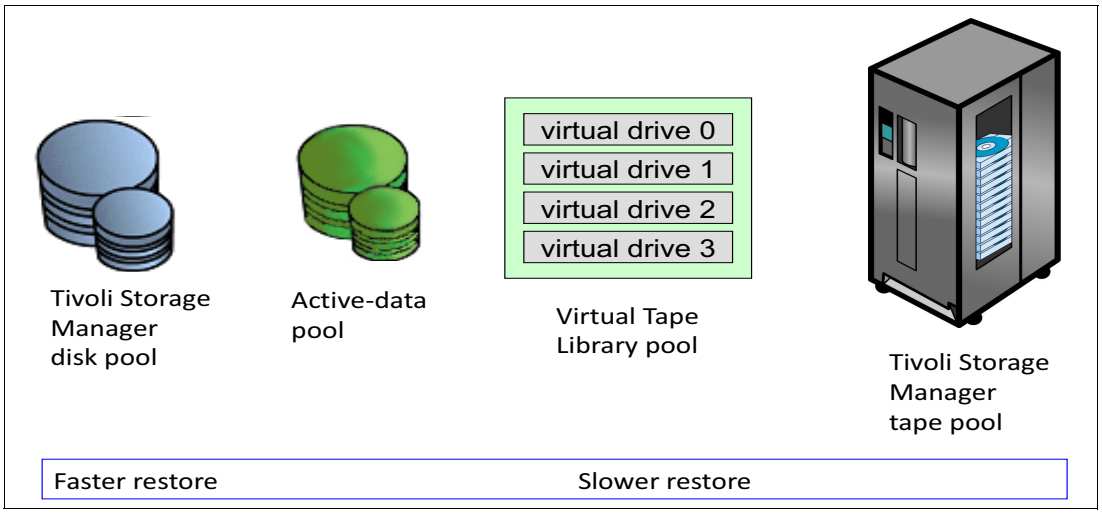

Figure 3-12 Advanced storage tiering with virtual tape library

3.3.7 Policy management

Policy management encompasses all the rules for where data is stored, how many versions can be stored, and for how long it is stored. This is one of the core paradigms of IBM Tivoli Storage Manager that provides the basis of its behavior. You should be familiar with the Tivoli Storage Manager server policy management.

Explaining each of the data storage management components, or the effects of the retention parameters defined in these components on the data stored in the server is far beyond the scope of this book. But, we do provide a brief overview of how client data is stored.

To gain more understanding of policy management, see the “Policy Management” topic in *IBM Tivoli Storage Manager Concepts*, SG24-4877:

http://www.redbooks.ibm.com/abstracts/sg244877.html

Tivoli Storage Manager policies are rules that determine how the client data is stored and managed. The rules include where the data is initially stored, how many backup versions are kept, how long archive copies are kept, and so on.

You can have multiple policies and assign the various policies as needed to specific clients, or even to specific files. Policies Assign a location in server storage where data is initially stored. Server storage is divided into storage pools that are groups of storage volumes.

Server storage can include hard disk, optical, and tape volumes.

Clients use Tivoli Storage Manager to store data for any of the following purposes:

- Backup and restore

  The backup process copies data from client systems to server storage to ensure against loss of data that is regularly changed. A policy includes the number of versions and the retention time for those versions. The server retains versions of a file according to this policy, and replaces older versions of the file with newer versions.

- Archive and retrieve

  The archive process copies data from client systems to server storage for long-term storage. The process can optionally delete the archived files from the client systems. The server retains archive copies according to the policy for archive retention time. A client can retrieve an archived copy of a file.

- Backup set recovery

  Backup set recovery is the creation of a complete set of backed-up files for a client. The set of files is called a *backup set*. A backup set is created on the server from the most recently backed-up files that are already stored in server storage for the client. The policy for the backup set consists of the retention time that you choose when you create the backup set.

  You can copy a backup set onto compatible portable media, which can then be taken directly to the client for rapid recovery without the use of a network and without having to communicate with the Tivoli Storage Manager server.

- Migration and recall

  Migration is a function of the Tivoli Storage Manager for Space Management (for supported UNIX and Linux systems) and Tivoli Storage Manager HSM for Windows programs. It frees client storage space by copying files from client systems to server storage. On the client, the program replaces the original file with a stub file that points to the migrated file in server storage. Files are recalled to the client systems when needed.

  The process of migrating and retrieving data through these programs is transparent to users and applications, other than a possible degradation in performance as compared to data stored on locally attached, tier one disk. Policy determines when files are considered for automatic migration. On the UNIX or Linux systems that support the Tivoli Storage Manager for Space Management program, policies determine whether files must be backed up to the server before being migrated. Space management is also integrated with backup. If the file to be backed up is already migrated to server storage, the file is backed up from there.

Figure 3-13 shows policy as part of the Tivoli Storage Manager process for storing client data.

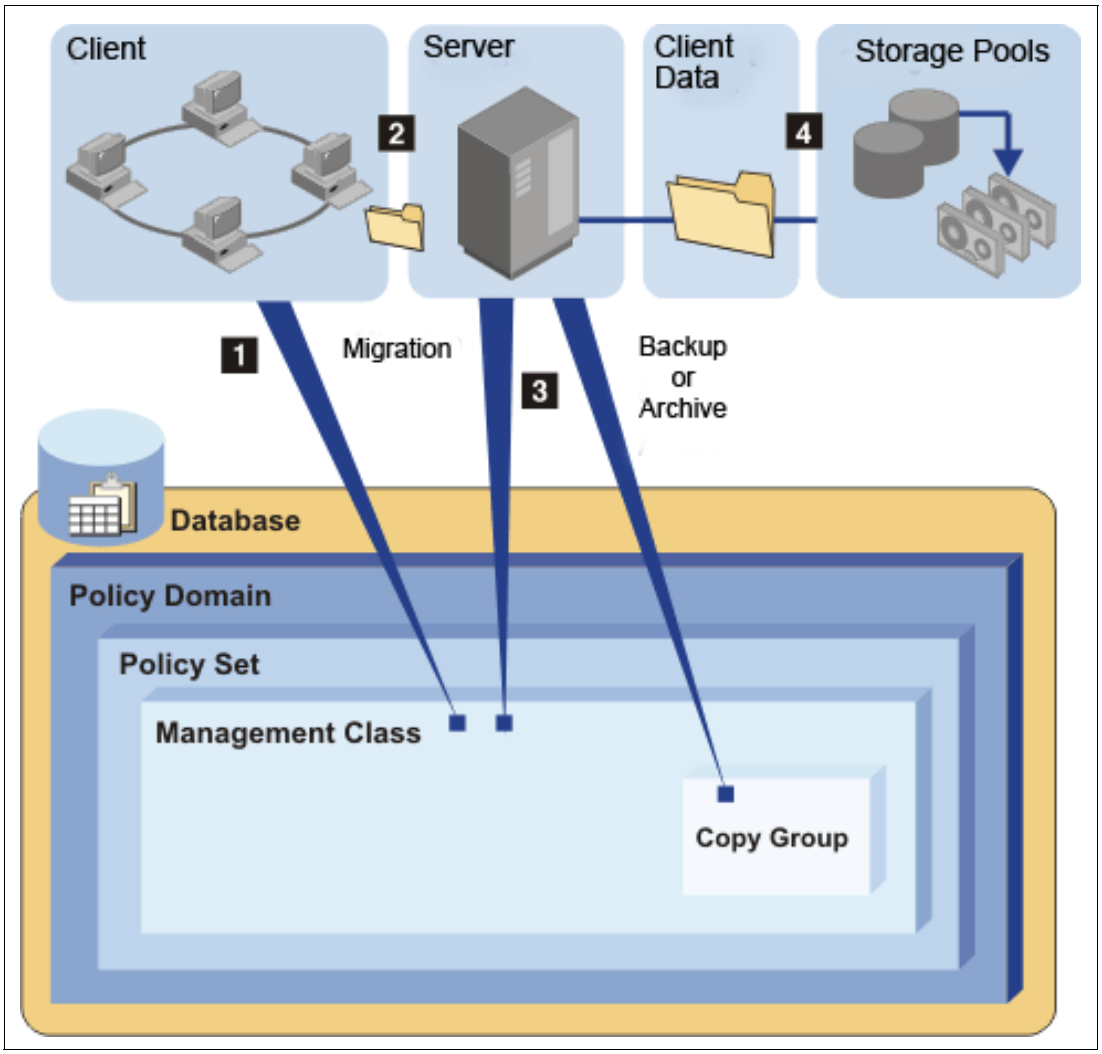

Figure 3-13 Policies as part of the data storage process

The steps in the process are as follows:

1. A client initiates a backup, archive, or migration operation. The file involved in the operation is bound to a management class. The management class is either the default or one specified for the file in client options (the client's include-exclude list).
2. If the file is a candidate for backup, archive, or migration based on information in the management class, the client sends the file and file information to the server.
3. The server checks the management class that is bound to the file to determine the destination, the name of the Tivoli Storage Manager storage pool where the server initially stores the file. For backed-up and archived files, destinations are assigned in the backup and archive copy groups, which are within management classes. For space-managed files, destinations are assigned in the management class itself.
4. The server stores the file in the storage pool that is identified as the storage destination. The Tivoli Storage Manager server saves information in its database about each file that it backs up, archives, or migrates. If you set up server storage in a hierarchy, Tivoli Storage Manager can later migrate the file to a storage pool other than from the one where the file was initially stored. For example, you might set up server storage so that Tivoli Storage

Manager migrates files from a disk storage pool to tape volumes in a tape storage pool after a certain amount of time.

Files remain in server storage until they expire (retention policies are met) and expiration processing occurs, or until they are deleted from server storage. A file expires because of criteria that are set in policy. For example, the criteria include the number of versions that are allowed for a file and the number of days that elapsed since a file was deleted from the client's file system.

3.3.8 Protecting data with Tivoli Storage Manager security

Tivoli Storage Manager offers various levels of security either by protecting data access or by securing the transportation of sensitive information to and from the Tivoli Storage Manager client and server over the network. See 5.3.6, "Protecting the Secure Sockets Layer digital certificate file" on page 109.

LDAP authentication

Instead of the embedded password protection of client data, the Tivoli Storage Manager server can use an external LDAP directory or Microsoft Active Directory to provide login authentication for administrators and nodes. Figure 3-14 shows the communication flow.

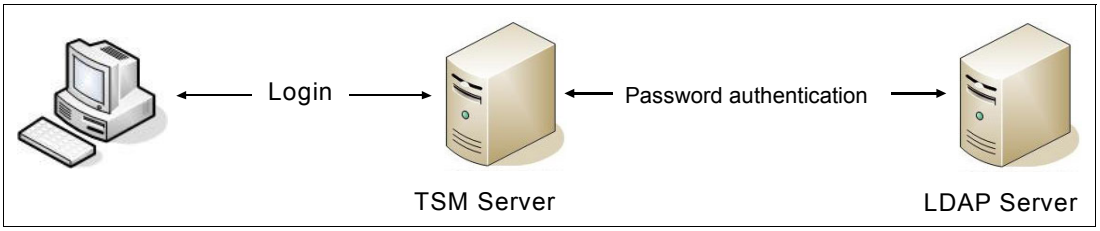

Figure 3-14 LDAP authentication

IBM Tivoli Directory Server and Windows Active Directory are the current supported LDAP servers. LDAP-authenticated passwords give you an extra level of security by being case-sensitive, offering advanced password rule enforcement, and a centralized server on which to authenticate them. You must have an LDAP directory or Microsoft Active Directory on which to authenticate the password.

For more information, go to the following address:

http://ibm.co/1pBawKB

Secure data transportation

You can use Transport Layer Security and Secure Sockets Layer (TLS/SSL) on HP-UX, Linux, Solaris, AIX, and Windows platforms. Figure 3-15 on page 71 shows the file transfer from the client to the server.

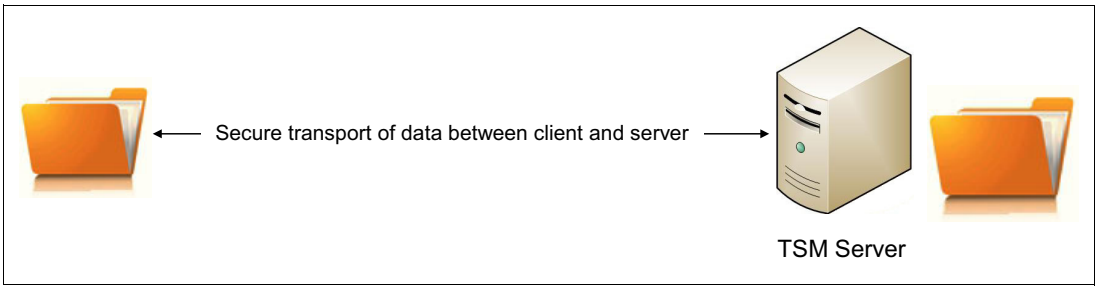

Figure 3-15 Secure data transportation

With TLS/SSL industry-standard communications, you can encrypt all traffic between the backup-archive client, the administrative command-line clients, and the IBM Tivoli Storage Manager server. You can use either self-signed or vendor-acquired SSL certificates. TSL/SSL is available for LAN-free and server-to-server communication too.

Data encryption

To heighten security for Tivoli Storage Manager sessions, data sent can be encrypted. Encryption happens on the client side. The file is password-protected so that it is only the client that can unlock the file during a restore. Encryption exists for both the backup-archive client and API client. Figure 3-16 shows the transport from the client to the server.

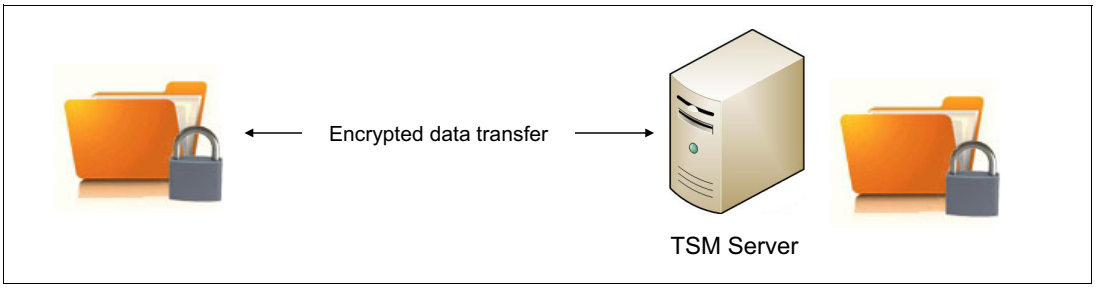

Figure 3-16 Encryption

Data sent to the Tivoli Storage Manager server during backup and archive operations are encrypted with standard 128-bit AES or 56-bit DES encryption. The data that you include is stored in encrypted form, and encryption does not affect the amount of data that is sent or received.

For WAN implementations of Tivoli Storage Manager across public networks, data encryption complements the TLS/SSL communication and completes data security for Tivoli Storage Manager.

Encryption and client-side data deduplication are mutually exclusive. If you enable both encryption and client-side data deduplication, encryption operations complete and client-side data deduplication is ignored. Encryption is incompatible with either client-side or server-side deduplication. Encrypted files, and files that are eligible for client-side data deduplication, can be processed in the same operation, but are done in separate transactions.

Backup over a firewall

IBM Tivoli Storage Manager has enhanced support for environments with firewalls in which communication originating from outside the firewall has to be restricted. Clients normally contact the server but with the firewall support, you can choose to restrict session initiation to the server. Scheduled backup-archive client operations can be restricted to server-initiated sessions.

Tape encryption

Another possibility is to encrypt the data on tape. For more information, see *IBM System Storage Solutions Handbook*, SG24-5250:

http://www.redbooks.ibm.com/abstracts/sg245250.html

3.3.9 Scripting for administration task automation

Server scripting provides integrated mechanisms for serializing server operations. For example, you can use scripts for the automation of database backups. The simplest way of scheduling an automated backup is to write a short script to perform the backup. Then, use Tivoli Storage Manager central scheduler to define a schedule to issue a command to execute the script.

When you install the product, it includes administrative scripts that have an example order of execution for scheduling administrative commands. The sample scripts are in `scripts.smp` and include the following examples:

- Command parameter substitution
- SQL SELECT statements that you specify when the script is processed
- Command execution control, such as PARALLEL and SERIAL processing options
- Conditional logic flow statements:
  - The IF clause: This clause determines how processing should proceed based on the current return code value.
  - The EXIT statement: This statement ends script processing
  - The GOTO and LABEL statement: This statement directs logic flow to continue processing
  - Comments

You can convert the sample scripts into a maintenance script according to your needs.

The Tivoli Storage Manager also supports macros that are called by the administrative client. A macro is a file that contains one or more administrative client commands. You can run a macro only from the administrative client in batch or interactive modes, not from the server.

Macros are stored as a file on the administrative client. Macros are not distributed across servers and cannot be scheduled on the server.

3.3.10 Automatic client deployment

IBM Tivoli Storage Manager can be scheduled to automatically deploy backup-archive client software to all workstations that already have the backup-archive client installed. Automated client deployment helps deliver quicker and less labor intensive software update distribution provided in the Tivoli Storage Manager Administration Center.

Figure 3-17 shows an overview of the automatic client deployment.

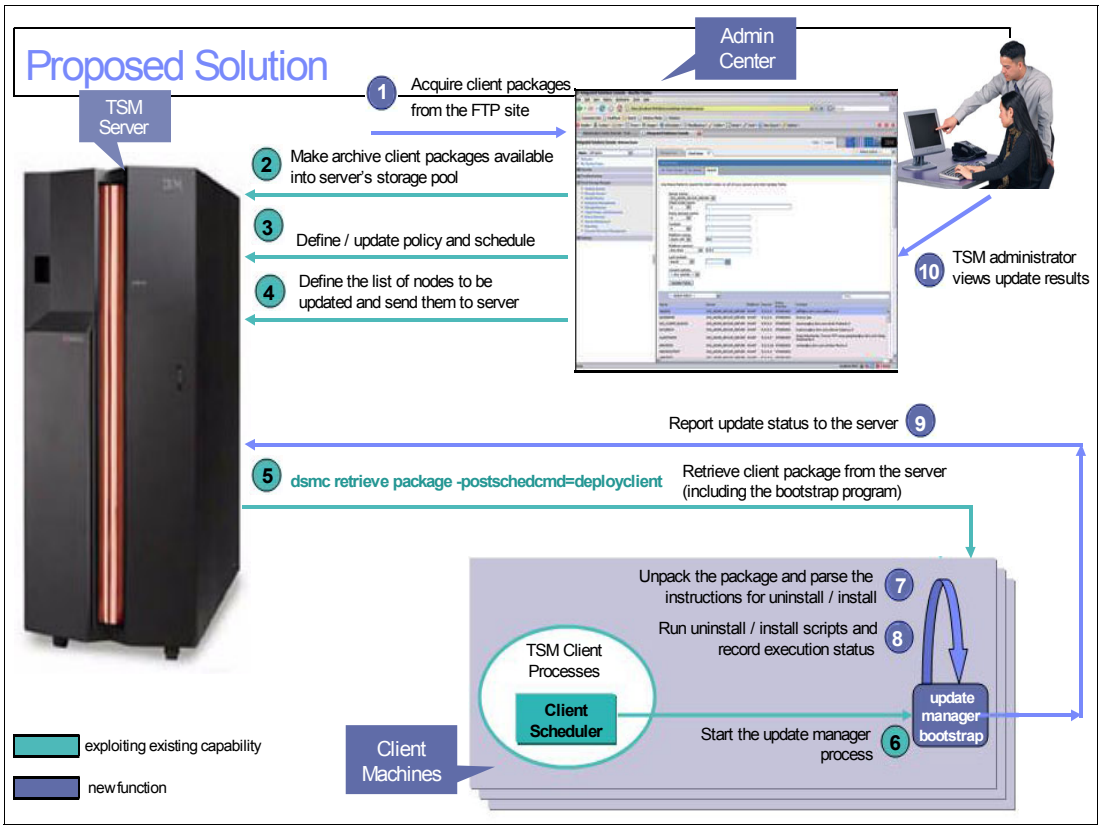

Figure 3-17 Automatic backup-archive client deployment

3.3.11 Reporting analytics

The IBM Tivoli Storage Manager monitoring and reporting feature can be installed on Linux, Solaris, and Microsoft Windows platforms, but can monitor a Tivoli Storage Manager server running on any platform.

You can view the historical reports to determine if any issues or trends need attention, such as uncontrolled growth over time. You can also view workspace that is monitored to see the Tivoli Storage Manager server IDs, database size, agent status, client node status, scheduled events, and so on.

The reporting component, sometimes referred to as Tivoli Common Reporting, reports on the retrieved historical data. IBM Tivoli Monitoring acts as a monitoring application that provides workspace for you to monitor near real-time information.

The Tivoli Storage Manager monitoring agent communicates with the Tivoli Storage Manager reporting and monitoring server to retrieve data from its database and return this data to the Tivoli Monitoring server. The Tivoli Storage Manager monitoring agent communicates with the existing IBM Tivoli Monitoring infrastructure, or Tivoli Storage Manager reporting and monitoring server to retrieve data from its database and return this data to the Tivoli Monitoring server.

Figure 3-18 shows a sample setup.

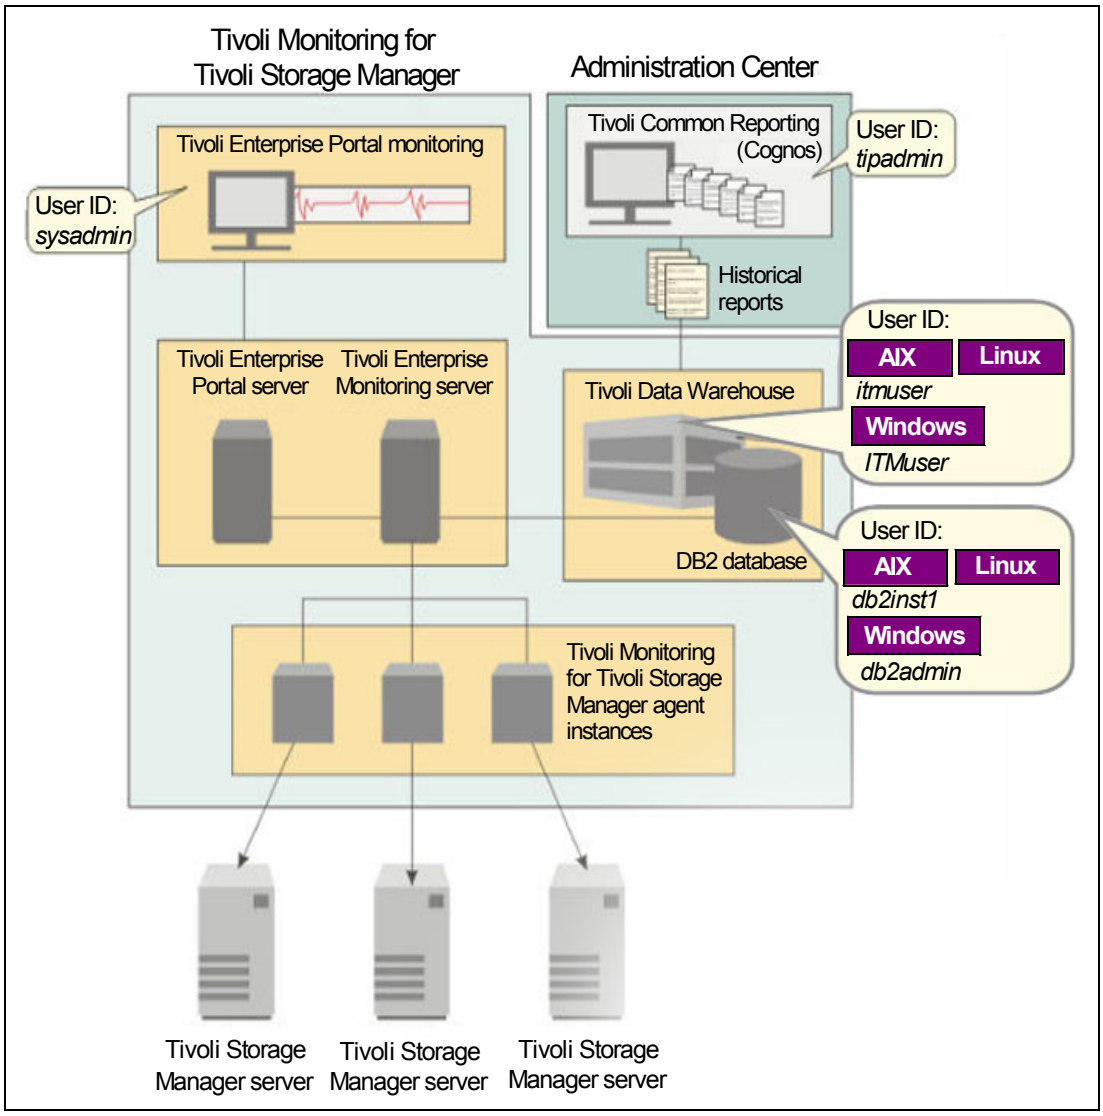

Figure 3-18 Sample setup for monitoring and reporting

3.3.12 Storage area network (SAN) environments

In this section, we look at Tivoli Storage Manager in SAN environments.

Library sharing

When Tivoli Storage Managers share a storage device, a single server, the library manager, controls device operations. These operations include mount, dismount, volume ownership, and library inventory. Other servers and other library clients use server-to-server communications to contact the library manager and request device service. Data moves over the SAN between servers and the storage device. Figure 3-19 on page 75 shows a configuration with a single library client, library server, tape library, the communication paths and data paths.

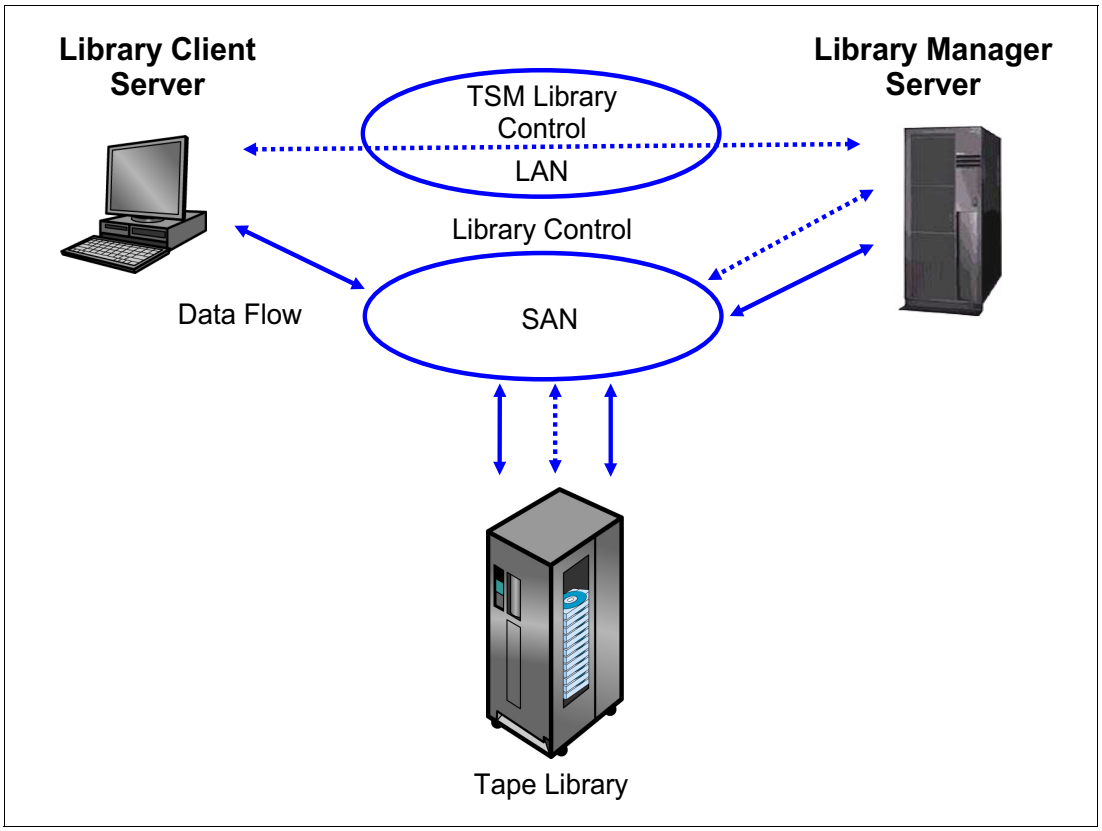

Figure 3-19 Storage area network tape library sharing

The inventory of media volumes in the shared library device is partitioned among servers. Either one Tivoli Storage Manager server owns a particular volume, or the volume is in the global scratch pool. No one server owns the scratch pool at any one time.

Only one Tivoli Storage Manager server accesses each tape drive at a time. Drive access is serialized and controlled so that servers do not dismount other servers' volumes or writes to drives where other servers mount their volumes.

3.4 Tivoli Storage Manager server database management

This section gives a short summary of server database backup and maintenance. A link to more details is provided.

3.4.1 Data Protection with Tivoli Storage Manager server database

The change to use DB2 as a database gives us the indexing capability to exploit the next generation of management, scalability, availability and performance. The evolving world demands a data protection solution with these features, which Tivoli Storage Manager has.

Tivoli Storage Manager DB2 relational database

Historically the protection of the Tivoli Storage Manager database has been easy and we now can exploit the 25 years of innovation of solid reliability and availability in DB2.

You can now scale the database much bigger than was the case with Tivoli Storage Manager versions that used the proprietary database at V5.5 or before. This raises new challenges, like completing the DB backup as quickly as possible, making the application available at all times, and optimizing the utilization of the database. However this is exactly what DB2 helps with.

Database backup

More options to protect the database exist today than previously. The standard way is still valid but as the database grows we need a smarter way to get the backup and restore times decreased. Multiple concurrent data streams while backing up or restoring the database can be used or Tivoli Storage FlashCopy Manager can protect the database.

See 5.2.2, “Database backup and restore” on page 96 for further details.

Replication solution with high availability disaster recovery

High availability disaster recovery (HADR) is for replicating the DB2 database. In combination with remotely attached copy storage pools, the intent is to combine these two to create a “warm standby” solution. The mirror of the Tivoli Storage Manager database eliminates the single point of failure without the need to be recovered if a failure occurs. Using HADR also makes the copy storage pool readily available because of automatic volume switching.

HADR is not intended to replace the database backup, but can increase the availability of Tivoli Storage Manager server application. See “DB2 High Availability Disaster Recovery (HADR)” on page 104.

Clustered solution for the Tivoli Storage Manager database

Depending on which platform you use, various options are available to implement the Tivoli Storage Manager database in a clustered environment. Cluster solutions eliminate single points of failure by making use of redundant components. The OS must deliver good capabilities to do the fail over automatically. It will be a custom configuration, and therefore the implementation varies from platform to platform.

On Microsoft Windows, the Tivoli Storage Manager server can be set up as part of the Microsoft Cluster for Windows resource group. Each Tivoli Storage Manager server instance is configured in its own resource group.

A more detailed discussion of this topic is available in “Clustering services” on page 100.

Database reorganization

The Tivoli Storage Manager server database reorganizes itself automatically in order to enhance performance and usage of the database. The process reconstructs rows to eliminate fragmented data, and compacts information. If the automatic table and index reorganization is affecting server performance, you can manually schedule reorganizations.

See the *Database size, database reorganization, and performance considerations for Tivoli Storage Manager Version 6 servers* technote for details about reorganization:

http://www.ibm.com/support/docview.wss?uid=swg21452146

3.4.2 Tivoli Storage Manager daily operation

The wheel of life (Figure 3-20 on page 77) represents a typical day in the life of a Tivoli Storage Manager server. With the recent enhancements to the server code, new tasks must be implemented into the daily cycle, the identify task, and the node replication task.

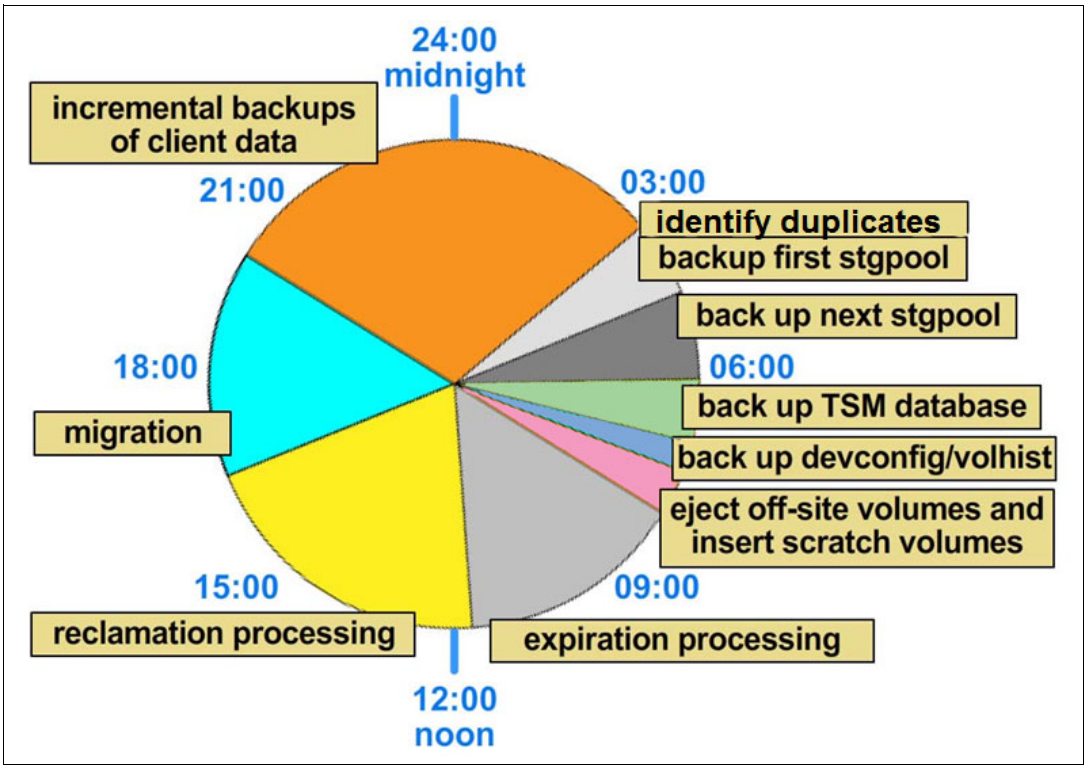

Figure 3-20 Tivoli Storage Manager: Wheel of Life

Keep in mind that this is not an implementation guide and we use the graphic only to explain the concepts. Your real schedules, of course, might look different, and might change depending on how much processing power the server has and the workload. For a detailed introduction to the topic, see the following resources:

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

  http://www.redbooks.ibm.com/abstracts/sg245416.html
- *IBM Tivoli Storage Management Concepts,* SG24-4877

  http://www.redbooks.ibm.com/abstracts/sg244877.html

A general daily lifecycle starts with the client backup operation. Here we do not distinguish between progressive incremental or some data protection agent activity. You usually dedicate some time of the day to the backup or archive operations. The best approach is not to not interrupt the backup task with server maintenance operations.

When the backup window is closed, the server can start its housekeeping. Housekeeping consists of the following tasks:

- Backup storage pools: Typically, the first process after the backup activity completes is to back up the storage pools to be sure valid copies are available for onsite and offsite storage.
- Database backup: A current database backup makes sure you have a copy of your database in case something happens and you need to restore. The backup taken after the storage pool backup does have the correct pointers to the objects managed by the server.

During the backup storage pool and the database backup process, you usually omit client operations so the processes and sessions do not interfere again. When completed, clients can start their normal operations again, for example, restore files if needed. For the server housekeeping, the traditional tasks are as follows:

- Migrate data: Migration frees space by moving objects to the next storage tier, a typical migration operation is the migration from disk to tape.
- Expire inventory: The expiration process identifies objects that are no longer managed by the Tivoli Storage Manager server and updates the database accordingly. Expiration still is a database-intensive task and you might want to make sure there is as little additional load against the server database as possible. When the expiration is complete, the server database is ready to proceed with reclamation.
- Reclamation: The reclamation process will finally free up the resources that have been tied to the objects identified by the expiration process. Usually you do reclamation for offsite volumes first, then for the local tape and sequential file volumes.

After reclamation completes, you are at the end of your housekeeping tasks. You might have scheduled another database backup while continuing with your server operation, waiting for the next backup window to open. With Tivoli Storage Manager V6, additional server housekeeping tasks are introduced.

- Identify duplicates
- Replicate nodes

Now if you go back to the wheel of life, you see how we integrate those processes in the daily operation. We start the identify duplicates process during the ingesting of client data so that we have more time to find the duplicates during the day. Because of the resources required for the identify process and the file expiration process, you might want to be sure that identify and expiration do not run in parallel.

This is why in the sample lifecycle we wait for the expiration to complete before we start the node replication, another job that is database-intensive. Depending on your implementation, you might not take local storage pool copies for your deduplication storage pool and rely on node replication. If that is the case, you can adjust the server process scheduling to your needs, for example copy your non-deduplication storage pools while running the identify duplicate processes in parallel. In any case, be sure that you dedicated the required resources to your server so that the server is able to complete the housekeeping tasks.

There is a recorded Storage Technical Exchange (STE) session that discusses how to perform important, system validations on your Tivoli Storage Manager V6 environment. It explains how the basic steps can be performed manually or automatically. Several components, including the Tivoli Storage Manager Administration Center, Tivoli Common Reporting, Tivoli Enterprise Portal and custom SQL queries are presented, you can find the presentation at the following address:

http://www.ibm.com/support/docview.wss?uid=swg27021836

Also, several videos are available about the Tivoli Storage Manager Operation Center:

https://www.youtube.com/watch?v=5lOwtmCmwkA

Part 3

Solution to challenges using Tivoli Storage Manager

In this part we provide solutions to challenges and how they are resolved with Tivoli Storage Manager.

This part contains the following chapters:

- Chapter 4, “Tivoli Storage Manager challenge matrix” on page 81
- Chapter 5, “Tivoli Storage Manager server protection” on page 93
- Chapter 6, “Tivoli Storage Manager Technologies and Solutions” on page 111
- Chapter 7, “Protecting your data with Tivoli Storage Manager” on page 141

It also contains these appendix topics:

- Appendix A, “Practical approach to creating a Tivoli Storage Manager solution” on page 311
- Appendix B, “Hierarchical storage management (HSM)” on page 323
- Appendix C, “VSS and Tivoli Storage Manager related product concepts” on page 327
- Appendix D, “Additional material” on page 331

4

Tivoli Storage Manager challenge matrix

Tivoli Storage Manager provides a centralized, automated data protection and storage management solution that addresses the business and technical challenges that companies face.

This chapter looks at those challenges, and helps to categorize them, thus allowing you to take advantage of the entire toolkit portfolio to adjust for current and future challenges.

This chapter has an explanation of how the team that wrote this book approached the task to provide an alternate view of the product to you: looking at Tivoli Storage Manager as a data protection solution in addition to a backup product. You will find pointers to solutions the team developed to meet your business requirements.

4.1 Solution matrix

When the team started the project to write this book, there was a lot of brainstorming, mindmapping, and sketching on the whiteboard to structure ourselves. It was our goal to describe the evolving set of data-protection challenges and how capabilities in Tivoli Storage Manager can best be used to address those challenges.

We searched for the best way to explain the transition from what was considered a backup/archive product to a full data protection solution, using well known Tivoli Storage Manager technologies. The goal was to provide possible solutions and the information to allow you to design and take advantage of new solutions combining available technologies.

The initial approach was to introduce something we named *solution tiering*. We tried to bring RTO/RPO objectives into some relation to implementation cost. What we developed on the whiteboard was the chart you see in Figure 4-1, the solution tiering chart.

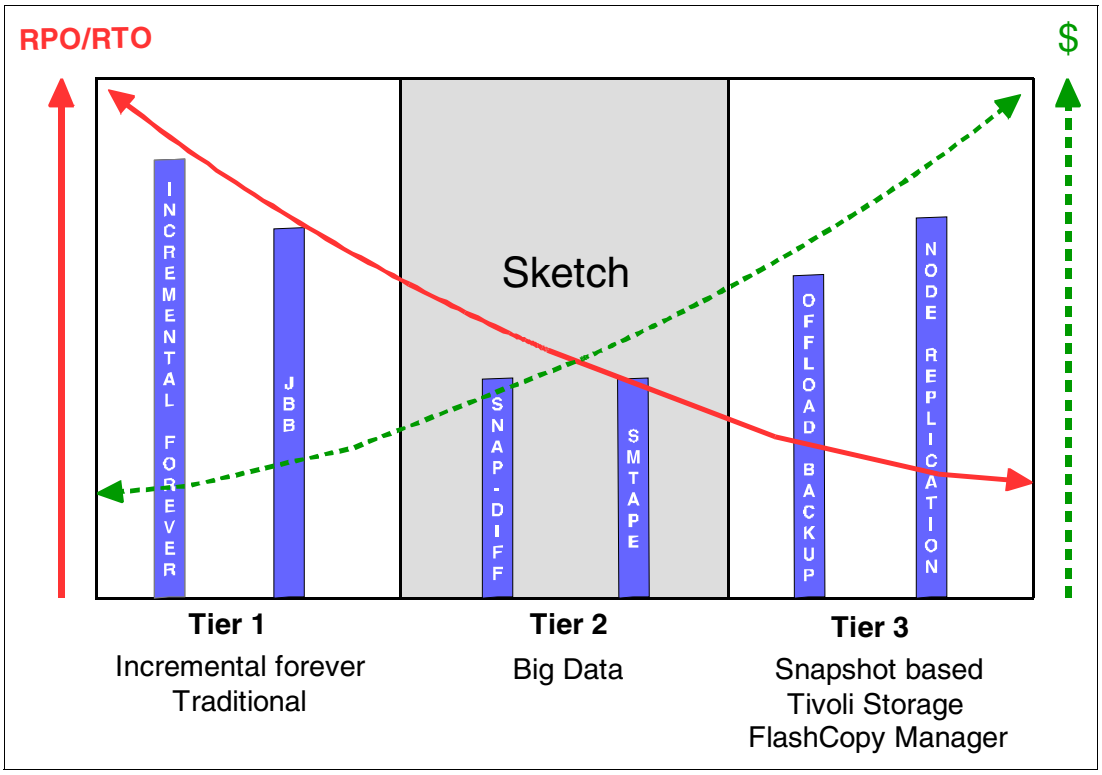

Figure 4-1 Initial solution tiering sketch

The plan was to tweak this sketch to our needs and then have the graphic designer convert this to an impressive graphic. Now, when we started reviewing all types of information to use for the book, we came across the Disaster Recovery Tiers that were defined by the SHARE user group in 1992. This is graphically represented in *Disaster Recovery Strategies with Tivoli Storage Manager*, SG24-6844:

http://www.redbooks.ibm.com/abstracts/sg246844.html

The representation is similar to the one in Figure 4-2.

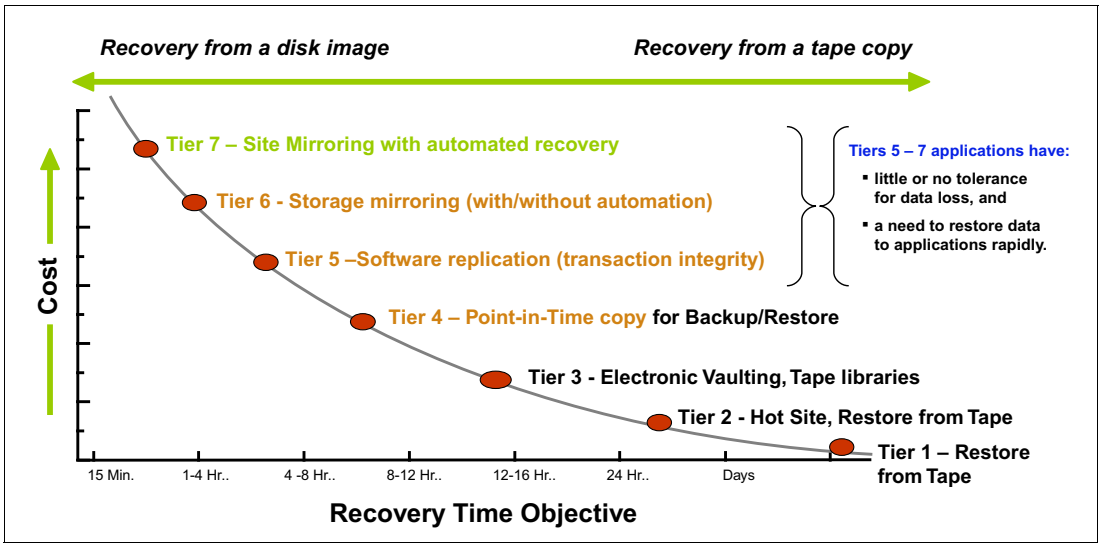

Figure 4-2 Business continuity tiers

The message we learned from this is that the jargon, when discussing data protection, has changed, but the challenges have existed for a long time. The other message was that for all of this time, the Tivoli Storage Manager product was available to address these challenges. For example, although the 1992 concept of “big data” does not look challenging today, the Tivoli Storage Manager product at that time provided the desired level of protection to your data, as it does today. So although you might see further references to that model throughout the book, the team decided to return to the drawing board and devise an alternate approach to introduce the changed mind-set of Tivoli Storage Manager as a data protection solution.

As a next step in our process, we decided to remove challenges such as financial aspects, from our consideration. For example, it is nothing new that full high availability (HA) comes at a price. Our next step was to concentrate on possible data protection architectures.

4.1.1 Setting up the matrix

With all the product knowledge available in the team, the idea arose to introduce a matrix. The intention of the matrix is to assist you, as the reader of this book, to find your way through this publication. The next step in matrix development, we thought, was easy: we listed all the product features and related them to technical and business challenges. Although agreeing on table rows and columns took some time, all of a sudden we had a big matrix (see Figure 4-3 on page 84), and we were impressed by how well this worked.

Work Area							Challenges										
							business					addressed with	technical				
DATA REDUCTION	DATA AVAILABILITY	SECURITY	DISASTER RECOVERY	VIRTUALIZATION		Client or Server or both	reduce recovery time (RTO)	reduce backup time (meet RPO)	Compliance	Business Continuity	security (secure data transfer, secure data store,	Solution (discussed in this book)	Virtual Environments	Significant data growth	automation for business	Remote office and mobile office data	network bandwidth optimization
												Disk to Disk backup					
												Data reduction with TSM					
												Data reduction with ProtecTier					
												Node replication					
												TSM ToolKit					
												Hardware Snapshot technology					
												progressive incremental forever					
												deduplication					
												compression					
												off-load backup					
												Off-Host backup					
												instant restore					
												include / exclude					
												High Availability					
												Online backup of Applications					
												journal based backup					
												image backup					
												LAN-free data transfer					
												GPFS					
												NDMP					
												Snap-Diff					
												Snap-Mirror-to-Tape / SMTAPE					
												archive					
												Space Management / HSM					
												point-in-time restore					
												LDAP					
												SSL					
												data encryption					
												Server maintenance					
												Disaster Recovery Manager					
												Server-to-Server virtual Volumes					
												active datapool					
												caching					
												automatic client deployment					
												centralized monitoring and reporting					
												Tiered Storage Hierarchy					

Figure 4-3 Original matrix

An online version of the matrix is available under the Additional Materials tab at the IBM Redbooks website:

http://www.redbooks.ibm.com/Redbooks.nsf/pages/addmats

Also see Appendix D, “Additional material” on page 331.

Keep in mind that this matrix is something the team agreed upon, and when you review and apply filters, you might want to recategorize the technical and business challenges based on your individual business challenges.

From there, we defined categories that can help you find the right data protection solution for your specific task.

Then we needed to determine how to put the information in the matrix. Initially we thought it would be a perfect tool for what we planned to use it for. We introduced more rows to refer to

possible implementations, and this worked well. When we did this and reviewed the matrix again, we really stopped thinking of Tivoli Storage Manager as a list of product features. Our intention is to assist you to develop the same view of the product and, in the process, help you to enhance your knowledge of Tivoli Storage Manager as a data protection solution.

We next identified the following challenge categories, creating a subset of the original matrix tailored to each challenge:

- Data reduction
- Data availability
- Security
- Disaster recovery
- Virtualization

Again, this was the team's view of how to logically group the challenges, features, and solutions together. In the following sections of this chapter, and in the remainder of this book you will find the matrix subset for each category, guidance for how to interpret it, and references to extra resources, either in this publication or externally.

We hope you find the matrix as useful as we found it during the development of the book.

4.2 Reading the matrix

Each challenge matrix that is described in this chapter is a subset of the original matrix that is shown in Figure 4-3 on page 84.

On the left are the business challenges that relate to the challenge category. On the right are the technical challenges.

In the middle are the related Tivoli Storage Manager toolkit features we identified. You are most likely familiar with them because some have been available for a long time.

The matrix lists sample solutions that allow you to push your data protection limits. From there, you can find references to the sample implementations that demonstrate how creative you can be, implementing the exact data protection solution that fits your needs.

Depending on the specific challenge you have, you might want to look at either of the available solutions to reduce your recovery time, or to help you protect your virtualized environments, or help you with both. Search for an X (as in Figure 4-4 on page 86) in either the solution or toolkit row to match your challenge column. When you find a match, it can help you with the specific challenge category.

4.2.1 Challenge: Data reduction

Figure 4-4 on page 86 shows you the data reduction subset of the challenge matrix.

Here we look at the data reduction challenge, and as you might expect, you see many well known toolkit options to reduce network bandwidth and therefore increase network efficiency. Now if a standard implementation no longer fits your needs, look at the sample solutions we provide; they might be what you are looking for.

You might wonder why you see server maintenance listed under the data reduction challenge. Here it was the author's common understanding that you need a well maintained and healthy server to be able to achieve your goals.

Challenges										
business					addressed with	technical				
reduce recovery time (RTO)	reduce backup time (meet RPO)	Compliance	Business Continuity	security (secure data transfer, secure data store,	Solution (discussed in this book)	Virtual Environments	Significant data growth	automation for business	Remote office and mobile office data	network bandwidth optimization
X	X				Data reduction with TSM	X	X		X	X
X	X		X		Data reduction with ProtecTier	X	X		X	X
X	X	X	X	X	TSM ToolKit	X	X	X	X	X
X	X		X		Hardware Snapshot technology	X	X			
	X				progressive incremental forever	X			X	X
X	X				deduplication	X	X		X	X
	x				compression					X
	X				include / exclude	X	X			X
	X				journal based backup		X			X
X	X		X		image backup	X	X			
	X				Snap-Diff		X			X
					Space Management / HSM		X			
x			X		Server maintenance					

Figure 4-4 Challenge: data reduction

Solution references

Data reduction is a common feature of several market products. One of the most efficient ways to reduce the amount of data that is backed up these days is to use data deduplication. Tivoli Storage Manager has built-in data deduplication features at both the client and server side. In addition, it can also interact with external deduplication systems such as ProtecTIER when Tivoli Storage Manager built-in deduplication does not fit.

We describe two examples of data reduction, based on different technologies; both use Tivoli Storage products. See the following sections:

- 6.3, “Data protection solution using node replication feature” on page 128
- 6.4, “Tivoli Storage Manager together with ProtecTIER” on page 135

Toolkit references

Tivoli Storage Manager has data reduction features such as progressive incremental forever backups, data compression, or subfile backup along with the other features listed in the matrix snapshot. These features are described in Chapter 3, “Data protection with Tivoli Storage Manager” on page 33.

4.2.2 Challenge: Data availability

Figure 4-5 shows the data availability subset of the challenge matrix.

Challenges										
business					addressed with	technical				
reduce recovery time (RTO)	reduce backup time (meet RPO)	Compliance	Business Continuity	security (secure data transfer, secure data store,	Solution (discussed in this book)	Virtual Environments	Significant data growth	automation for business	Remote office and mobile office data	network bandwidth optimization
X	X				Disk to Disk backup	X		X	X	
X	X				Data reduction with TSM	X	X		X	X
X	X		X		Data reduction with ProtecTier	X	X		X	X
X			X		Node replication				X	
X	X	X	X	X	TSM ToolKit	X	X	X	X	X
X	X		X		Hardware Snapshot technology	X	X			
X	X				deduplication	X	X		X	X
X			X		instant restore	X	X			
X	X		X		High Availability	X				
X	X				Online backup of Applications					
X	X		X		image backup	X	X			
X	X				LAN-free data transfer		X			X
	X				GPFS		X		X	
X	X				NDMP		X			
X	X		X		Snap-Mirror-to-Tape / SMTAPE		X			
		X			archive					
					Space Management / HSM		X			
X					point-in-time restore	X				
x			X		Server maintenance					
X			X		Disaster Recovery Manager	X		X		
			X		Server-to-Server virtual Volumes				X	
X			X		active datapool					
X					caching					
					automatic client deployment			X		
					centralized monitoring and reporting			X	X	
X	X				Tiered Storage Hierarchy	X	X			

Figure 4-5 Challenge: data availability

Solution references

All solutions that are described in Chapter 6, “Tivoli Storage Manager Technologies and Solutions” on page 111 meet this data availability category because it is the main purpose of data protection. An aspect that might change regarding data availability is how fast you need the data back to your system. Depending on this recovery time objective criteria, you will probably decide to use a disk-to-disk backup approach or another tape-based solution. Although, consider that the faster your system is, the more expansive the solution is, as shown in Figure 4-2 on page 83.

Toolkit references

As previously stated, the purpose of a data protection solution is to make the data available no matter what happens. The matrix snapshot in Figure 4-5 on page 87 highlights the most relevant tools available in Tivoli Storage Manager based on the experience of the team writing this book. These features are described in Chapter 3, “Data protection with Tivoli Storage Manager” on page 33.

4.2.3 Challenge: Security

Figure 4-6 shows the security subset of the challenge matrix. The Tivoli Storage Manager toolkit options to secure your data are listed.

Challenges										
business					addressed with	technical				
reduce recovery time (RTO)	reduce backup time (meet RPO)	Compliance	Business Continuity	security (secure data transfer, secure data store,	Solution (discussed in this book)	Virtual Environments	Significant data growth	automation for business	Remote office and mobile office data	network bandwidth optimization
X	X	X	X	X	TSM ToolKit	X	X	X	X	X
				X	LDAP					
		X		X	SSL					
		X		X	data encryption					
x			X		Server maintenance					

Figure 4-6 Challenge: security

In this book, we do not describe all security options in detail. *IBM Tivoli Storage Manager: Building a Secure Environment*, SG24-7505 covers the various security features of Tivoli Storage Manager. It shows how to use them, together with preferred practice principles to design, implement, and administer a more secure backup management environment. The book covers passwords, administrative levels of control, the vital role of encryption, and procedures for managing offsite data, among other topics. The book is at this location:

http://www.redbooks.ibm.com/abstracts/sg247505.html

For information about improved user authentication and management by integration with Lightweight Directory Access Protocol (LDAP), go to the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SSGSG7_6.4.1/com.ibm.itsm.srv.doc/c_mgclinod_managepwlogin.html

4.2.4 Challenge: Disaster recovery

Taking advantage of deduplication and replication functionality, you now have many more options to protect your data and your Tivoli Storage Manager server. Figure 4-7 shows the disaster recovery subset of the challenge matrix.

Challenges										
business					addressed with	technical				
reduce recovery time (RTO)	reduce backup time (meet RPO)	Compliance	Business Continuity	security (secure data transfer, secure data store,	Solution (discussed in this book)	Virtual Environments	Significant data growth	automation for business	Remote office and mobile office data	network bandwidth optimization
X	X		X		Data reduction with ProtecTier	X	X		X	X
X			X		Node replication				X	
X	X	X	X	X	TSM ToolKit	X	X	X	X	X
X	X		X		Hardware Snapshot technology	X	X			
X	X		X		High Availability	X				
X	X		X		image backup	X	X			
X	X		X		Snap-Mirror-to-Tape / SMTAPE		X			
x			X		Server maintenance					
X			X		Disaster Recovery Manager	X		X		
			X		Server-to-Server virtual Volumes				X	
X			X		active datapool					

Figure 4-7 Challenge: disaster recovery

Solution references

Disaster recovery has become an entire pane of a data protection solution and must be carefully designed to fulfill the recovery time objectives (RTO) and recovery point objectives (RPO). These are two of the most relevant pieces of information when you build a data protection solution; they are basically the business impact.

Disaster recovery solutions are in examples in these two sections:

- 6.3, “Data protection solution using node replication feature” on page 128
- 6.4, “Tivoli Storage Manager together with ProtecTIER” on page 135

Toolkit references

The first image that comes to mind regarding disaster recovery is to make sure to protect your Tivoli Storage Manager server database, which is key to your protection solution. Tivoli Storage Manager server infrastructure protection is described in 5.2, “Protecting the Tivoli Storage Manager server database” on page 95. It contains (or refers you to other resources that contain) several creative solutions to protect your database.

4.2.5 Challenge: Virtualization

Figure 4-8 shows you the virtualization subset of the challenge matrix. It contains the tools available to protect your data in a virtual environment.

Challenges										
business					addressed with	technical				
reduce recovery time (RTO)	reduce backup time (meet RPO)	Compliance	Business Continuity	security (secure data transfer, secure data store,	Solution (discussed in this book)	Virtual Environments	Significant data growth	automation for business	Remote office and mobile office data	network bandwidth optimization
X	X				Disk to Disk backup	X		X	X	
X	X				Data reduction with TSM	X	X		X	X
X			X		Node replication				X	
X	X	X	X	X	TSM ToolKit	X	X	X	X	X
X	X		X		Hardware Snapshot technology	X	X			
	X				progressive incremental forever	X			X	X
X	X				deduplication	X	X		X	X
	X		X		off-load backup	X	X			
	X				Off-Host backup	X	X			
X			X		instant restore	X	X			
	X				include / exclude	X	X			X
X	X		X		High Availability	X				
	X				journal based backup		X			X
X	X		X		image backup	X	X			
X					point-in-time restore	X				
x			X		Server maintenance					
X			X		Disaster Recovery Manager	X		X		
X	X				Tiered Storage Hierarchy	X	X			

Figure 4-8 Challenge: virtualization

By virtualizing your systems, you separate the physical from logical, you manage and utilize IT resources as a cohesive and holistic unit that is constantly adjusting, reallocating, and responding as fast as the business dictates.

In the first phase of virtualization, when it was a cutting-edge technology, development or test environments were considered good candidates, while real production environments were not, mainly because of performance concerns. Today, with the advent of hypervisor-based virtualization, like IBM PowerVM hypervisor, VMware ESXi, or Microsoft Hyper-V, this perception has radically changed.

Virtualization has evolved and can now handle more data while delivering better performance than in its first phase. Virtualized infrastructures vary in size, therefore they have different requirements for data protection. The Tivoli Storage Manager family of products takes this into consideration and provides a set of scalable solutions to address every need.

Virtualization, from a data protection perspective, means heterogeneous platforms, different configurations, several applications to be handled, and also awareness of the components added by the virtualization layer. Depending on your needs, you might either need to take a picture of the entire virtualized environment (for example, using an off-loaded backup method like the Tivoli Storage FlashCopy Manager), or consider each virtual machine (VM) individually (Tivoli Storage Manager for Virtual Environment product).

Tivoli Storage Manager for Virtual Environments (referred to as Off-host backup in the TSM ToolKit column on Figure 4-8 on page 90) protects VMware vSphere virtual machines by moving the backup workloads to a centralized server and enabling near-instant recovery. It enables you to protect data without the need for a traditional backup window. Now you can protect the massive amounts of information that virtual machines generate without impacting the physical resources of the VMware server.

You might need to use a more conventional approach to protect your data, even in a virtual machine. In this case you can take advantage of the progressive incremental forever method combined with the journal-based backup.

Tivoli Storage Manager is cloud ready

Although business and IT use cloud for different reasons and with different goals, both have a common understanding of cloud's overall value: the ability to deliver IT with fewer boundaries, help improve speed and dexterity, and create new business value.

Therefore we can not talk about virtualization data protection without talking about cloud data protection.

The cloud, beyond virtualization, is the ability to dynamically move resources between different hosts. It adds the ability to create resource pools from virtualized assets, manage a group of VMs as a single entity, and provide self-service capabilities and usage based resource metering.

These requirements were considered when the latest version of Tivoli Storage Manager products were built (for instance, Tivoli Storage Manager for Virtual Environment Data Protection for VMware and Tivoli Storage FlashCopy Manager for VMware). These two products bring the required flexibility to handle the data protection of such a dynamic infrastructure that is the cloud.

Solution references

Details about various sample solution scenarios are in 7.2, "Virtualization: VMware Data Protection solution" on page 144.

Toolkit references

Tivoli Storage Manager toolkit is greatly improved (since V6.2 and even more with V7.1) to handle the virtual environment. Although the traditional features work within a virtual environment, the Tivoli Storage Manager for Virtual Environment product brings a complete set of enhanced features that you can use to handle data protection of virtualized systems from the hypervisor layer.

More information about the traditional tools is in Chapter 3, "Data protection with Tivoli Storage Manager" on page 33. An explanation of how the enhancements work is in 2.2.3, "Tivoli Storage Manager for Virtual Environments: Data Protection for VMware" on page 21 and 2.3, "Tivoli Storage Manager recent enhancements" on page 29.

4.3 Summary

In this chapter we introduce the solution matrix, explain how to use the matrix to find the solution for the data protection challenge you want to meet. We use the matrix to reference you from this chapter to solutions and Tivoli Storage Manager toolkit features that we document in this book.

5

Tivoli Storage Manager server protection

In this chapter, we discuss how to protect the Tivoli Storage Manager configuration and storage data backup from corruption.

5.1 Protecting the Tivoli Storage Manager server infrastructure

When you implement a Tivoli Storage Manager data protection solution, one of the most important issues is server infrastructure protection. In addition to the storage pool volumes, the server database and its infrastructure files must be protected against corruption or loss.

In Figure 5-1, our challenge matrix shows solutions and toolkit options related to server infrastructure protection and disaster recovery (DR). We discuss these in this chapter.

Challenges										
business					addressed with	technical				
reduce recovery time (RTO)	reduce backup time (meet RPO)	Compliance	Business Continuity	security (secure data transfer, secure data store,	Solution (discussed in this book)	Virtual Environments	Significant data growth	automation for business	Remote office and mobile office data	network bandwidth optimization
X	X				Disk to Disk backup	X		X	X	
X	X				Data reduction with TSM	X	X		X	X
X	X	X	X	X	TSM ToolKit	X	X	X	X	X
X	X		X		Hardware Snapshot technology	X	X			
X	X		X		High Availability	X				
x			X		Server maintenance					
X			X		Disaster Recovery Manager	X		X		
			X		Server-to-Server virtual Volumes				X	

Figure 5-1 Protect the server from a disaster

For a long time, performing regular database and storage pool backups to tape volumes, which are then moved to an offsite location, was the most common and implemented method.

Later in this chapter, we describe the traditional server protection techniques and strategies, and how to protect the database, with snapshot technology, which is now possible with the move to the underlying DB2 database.

The Tivoli Storage Manager infrastructure consists of the database and the setup files that are needed to recover the database and client data. The setup files include, for example, the active log and the archive log. Client data includes data that is backed up, archived, and migrated to primary storage pools.

Figure 5-2 on page 95 shows the business continuity tiers and the relation between recovery time objectives and the cost to safeguard against data loss. With Tivoli Storage Manager, you can configure the server infrastructure protection to meet each business continuity objective you require.

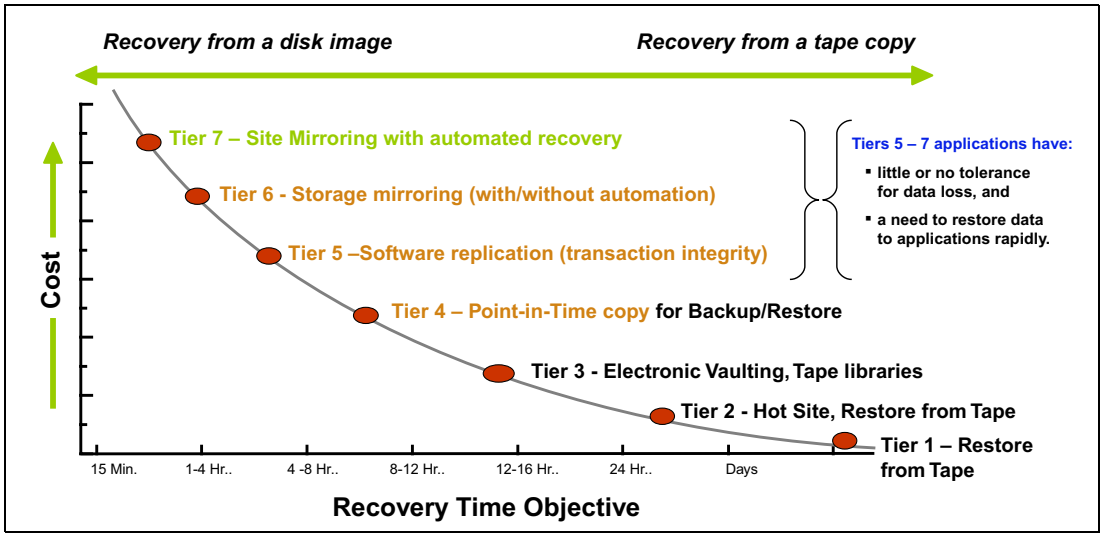

Figure 5-2 Business continuity tiers

Generally, you implement DR solutions on multiple tiers. Critical applications might be protected by a tier 6 solution; other applications might use a tier 3 solution.

Tivoli Storage Manager can be these components:

- The primary component of one of the disaster recovery solution tiers (1, 2, or 3).
- A secondary or complementary component of one of the disaster recovery solution tiers (4, 5, 6, or 7).

The Tivoli Storage Manager application can participate in and be protected by upper tier solutions in a manner similar to other critical applications. Define your current or server recovery requirements as early as possible during your planning phase.

5.2 Protecting the Tivoli Storage Manager server database

In this section, we examine considerations and methods for protecting the Tivoli Storage Manager database.

5.2.1 Database mirroring

To protect from hardware failures for your database and archive log files in Tivoli Storage Manager V6 and V7, you can mirror by using operating system or file system capabilities. You can also use device redundancy, such as RAID capabilities in the storage that is used for the server database.

Further discussion of mirroring techniques to protect the database log files is in 5.3.1, “Protecting the Tivoli Storage Manager database log files” on page 106.

5.2.2 Database backup and restore

To restore a damaged or lost database, you must have a database backup. You must also have copies of the files that are necessary to recover the database and client data. Database backup media and setup files can be stored offsite for protection, either under manual control, assisted by Disaster Recovery Manager, or automated. We describe architectures to protect your server database. We do not describe the restore procedure in detail, because that is implicit with the solution you plan to implement to protect your database.

Standard Tivoli Storage Manager database backup architecture

You can use existing device classes for database backups or you can define new ones to separate the database backup volumes. You can back up the database to sequential media: tape, file, or remote virtual volumes. If you use remote virtual volumes, you can use these volumes for electronic vaulting. Figure 5-3 shows how the standard Tivoli Storage Manager database backup to sequential volumes works.

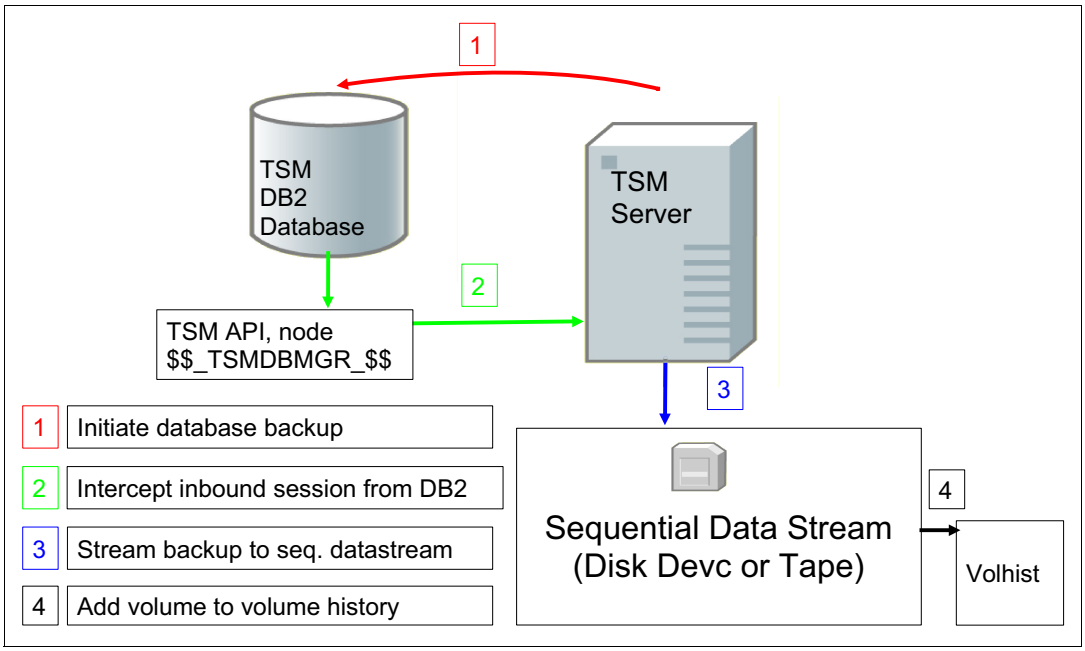

Figure 5-3 Standard Tivoli Storage Manager backup architecture

A special client is embedded within the Tivoli Storage Manager server to perform the backup and the data flows through the server to the media that will store the backup. Backups can be initiated manually by using the **`dsmserv backup db`** command or can be automatic by using database monitoring algorithms. The restore is later performed with the **`dsmserv restore db`** command, which identifies the appropriate backup media using the Tivoli Storage Manager volume history file, and restores the database from it.

Multistream database backup

With V6.3, Tivoli Storage Manager further enhances its database backup performance by providing server database backup using multithreading capabilities. Multithreaded database backups help increase the of Tivoli Storage Manager maintenance activities. You can specify up to four concurrent data streams for automatic or manual database backup operation (see Figure 5-4).

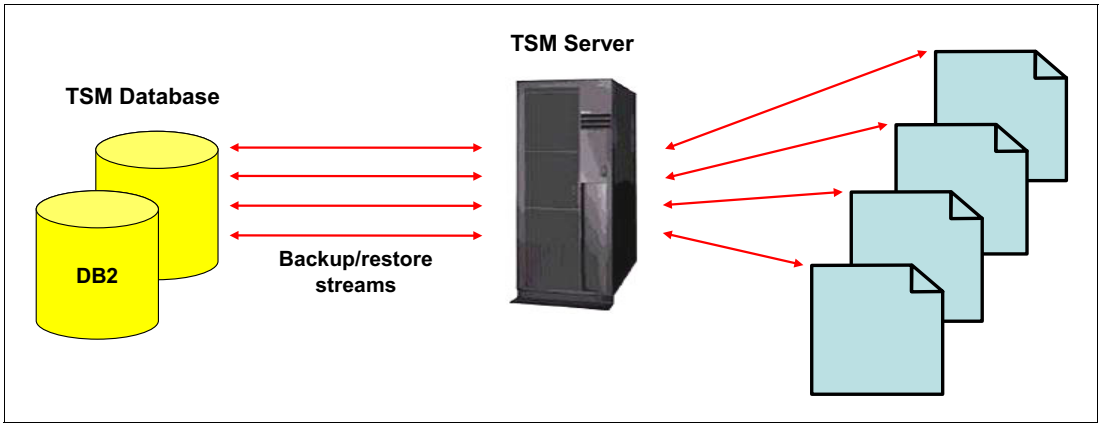

Figure 5-4 Multiple stream database backup and restore

Tivoli Storage FlashCopy Manager for Tivoli Storage Manager database protection

Tivoli Storage FlashCopy Manager software enables fast backup and restore of data by using the copy services functionality of storage hardware. Using hardware-based copy mechanisms rather than traditional backup techniques can significantly reduce backup and restore time windows, particularly where large data volumes are in place.

Although standard Tivoli Storage Manager server backup operations (as described in “Multistream database backup” on page 97) can now be multithreaded, they can still take many hours to backup or restore a multi-terabyte (TB) database. Performance of ongoing server operations can be reduced while the database is being backed up. A larger database takes longer to restore, so the server is offline for longer in the event of a Tivoli Storage Manager database restore.

To solve the backup or restore of a multi-terabyte database problem, you can use the Tivoli Storage FlashCopy Manager product to provide point-in-time snapshot backup and recovery capability for the Tivoli Storage Manager’s underlying database that is functionally equivalent to a standard database backup or restore, or to alleviate the impact of the increased size of Tivoli Storage Manager databases that can be caused when adopting the latest Tivoli Storage Manager capabilities, most notably storage pool data deduplication.

Figure 5-5 on page 98 shows a sample high-level view of an architecture to generate offloaded server database backups with Tivoli Storage FlashCopy Manager technology. The double- headed arrows in the figure represent SAN zoning and mapping requirements.

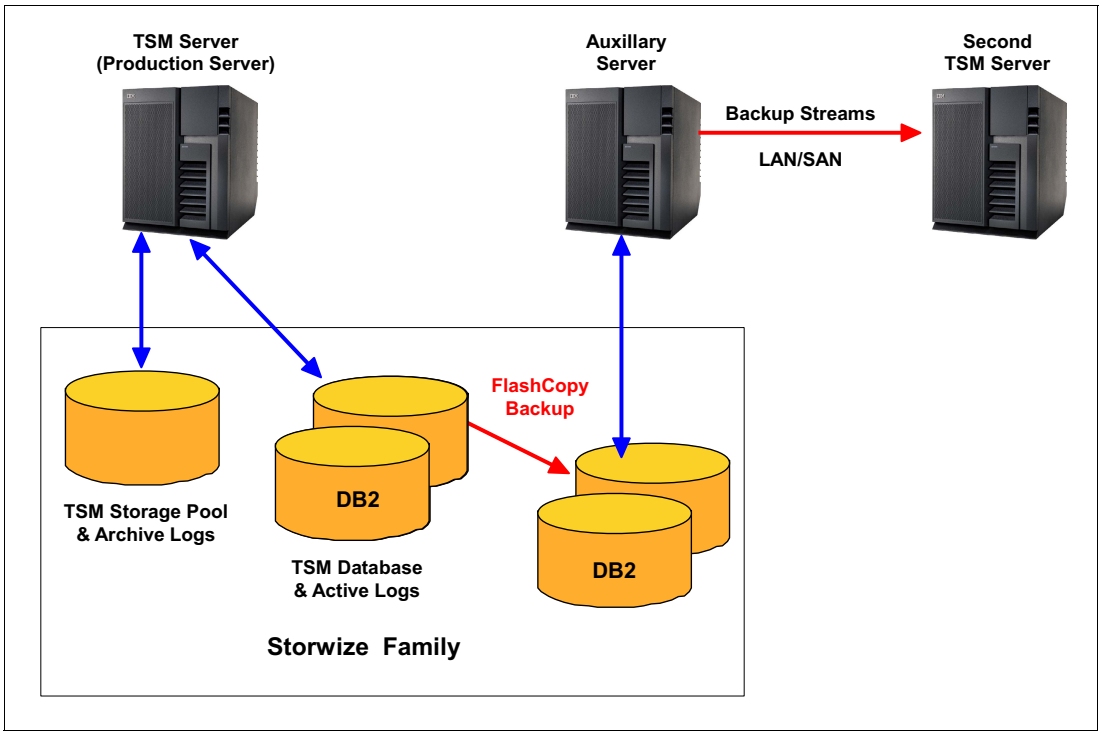

Figure 5-5 Tivoli Storage Manager Server database backup architecture using Tivoli Storage FlashCopy Manager

Tivoli Storage FlashCopy Manager can be used in conjunction with a Tivoli Storage Manager server to “offload” traditional backup processing from the production application servers. Snapshot backups are made by Tivoli Storage FlashCopy Manager on the production server and subsequently offloaded (backed up) to the Tivoli Storage Manager backup infrastructure. In this case backups are performed through an additional server called the backup server (also termed offload server), which performs the actual backup. Because the backup operation is offloaded to the backup server, the production server is free from nearly all the performance impact. The production server’s processor time is dedicated for the actual application tasks, so performance of application users is not affected during backup. When performing offloaded backups, Tivoli Storage FlashCopy Manager is used in conjunction with additional Tivoli Storage Manager modules to provide application consistency and perform the backup from the backup server to Tivoli Storage Manager.

You can extend this approach and as an example perform disk-only backups to an additional set of FlashCopy devices. Disk-only backups can then be used to provide a fast restore option for onsite recovery. For offsite protection, or to protect against failure of the storage device that holds the database they should be used in conjunction with “traditional” Tivoli Storage Manager database backup and restore techniques or with the “offloaded” backups shown above.

The methods for backup to and restore from disk-only backups and offloaded backups are described in the developerWorks Tivoli Storage Manager wiki at the following website:

http://ibm.co/1k614j2

More information about databases (not only Tivoli Storage Manager databases) is in 7.6.3, “Protect very large databases with Tivoli Storage FlashCopy Manager” on page 245.

Disaster recovery management

To store database backup media and setup files offsite, you can use the Tivoli Storage Manager Disaster Recovery Manager functionality, known as Disaster Recovery Manager, with the Tivoli Storage Manager Extended Edition.

Disaster Recovery Manager offers various options to configure, control, and automatically generate a disaster recovery plan containing the information, scripts, and procedures needed to automate recovery of the Tivoli Storage Manager server, and help ensure quick recovery of data after a disaster.

Figure 5-6 illustrates a basic Disaster Recovery Manager strategy.The courier in the figure represents shipping physical tape cartridges to a disaster recovery location. Today the courier can represent the data moving from one location to another, and which can be physical tapes, electronically copied utilizing virtual volumes, and if you think about storage pools you can use replication. There are several ways you can implement this, but fundamentally you ship your data to a remote location to protect from a disaster.

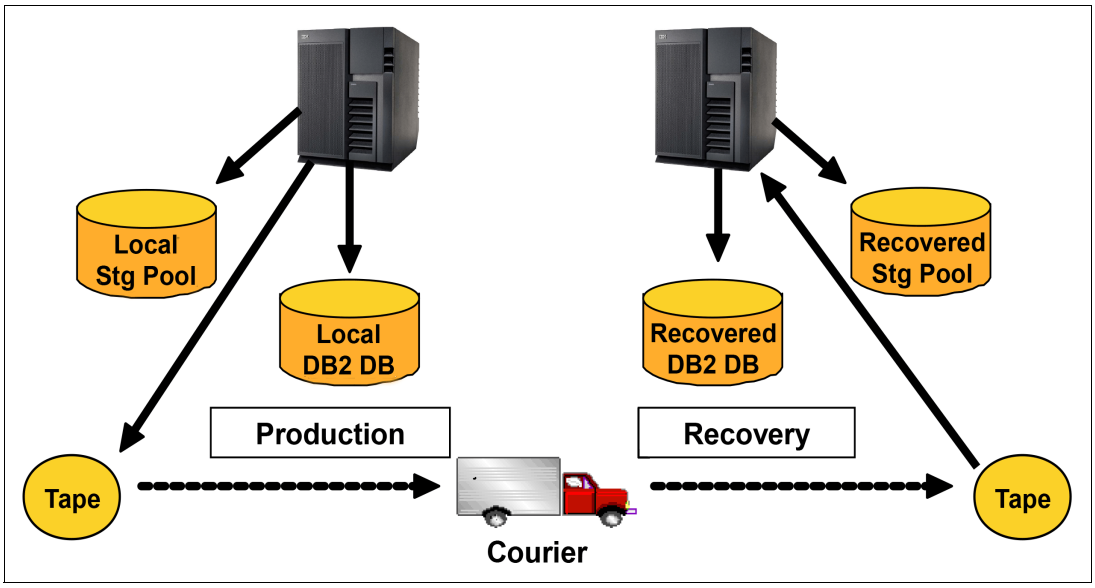

Figure 5-6 Disaster Recovery Manager: basic strategy

The Disaster Recovery Manager media management function manages the movement of database backup and copy storage pool tapes to and from offsite storage, and performs expiration of Tivoli Storage Manager database backup series. Disaster Recovery Manager also manages and tracks the media on which the data is stored, so that data can be easily located if disaster strikes.

The media is tracked in all locations: onsite, in transit, or in a vault as shown in Figure 5-7. Controls are provided to make sure that all previously backed up data can be found and restored in a reasonable amount of time.

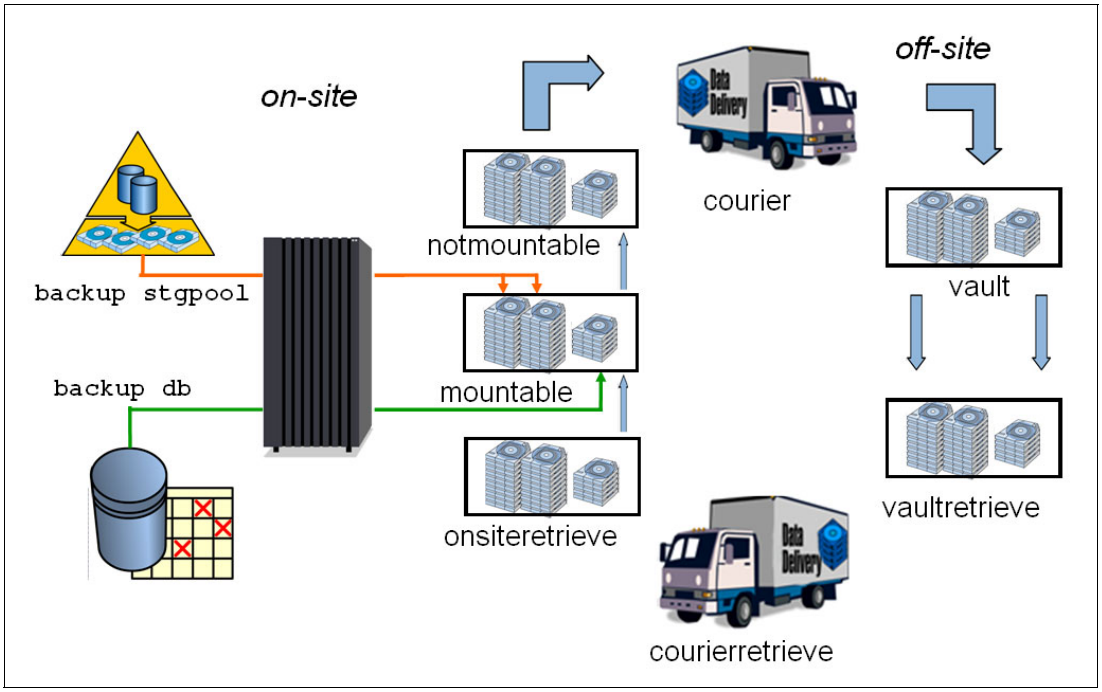

Figure 5-7 Disaster Recovery Manager media state values

High availability server database protection

There are various ways to extend the protection of your server database to meet your requirements. Here we discuss clustering and DB2 HADR. In “Tivoli Storage FlashCopy Manager for Tivoli Storage Manager database protection” on page 97, we document how you can back up your database for high availability.

Clustering services

Clustering is a high-availability solution that minimizes or eliminates many potential sources of downtime. It is the most common technique to achieve high availability (HA), by introducing redundancy in software, hardware, and data. In a failure, the clustering software immediately starts the application on the standby system without requiring administrative intervention.

Table 5-1 details cluster configuration models.

Table 5-1 Cluster configuration models

HA configuration	Second system behavior	Data protection	Recovery time
Cold standby	Second node is a backup, installed/configured only when first node fails. Upon failure, it is activated.	Data from first system can be backed-up and restored on second system as required.	Few hours
Warm standby	Software installed and available on second running node. On failure, second node SW started. Usually automated by cluster manager.	Data is regularly mirrored to second system using disk based replication or shared disk.	Few minutes
Hot standby	Software is installed and running on both nodes. The software on the second system is running but not processing.	Data mirrored near real time and both systems have identical data. Data replication is typically done through software.	Few seconds
Active-Active (load balanced)	Both first and second systems are active and processing requests in parallel.	Data replication happens through software and is bi-directional.	Zero failover time

Figure 5-8 shows a sample cluster configuration: nodes A, B, and C are active, D is passive (standby). Node D can take over from A, B, or C.

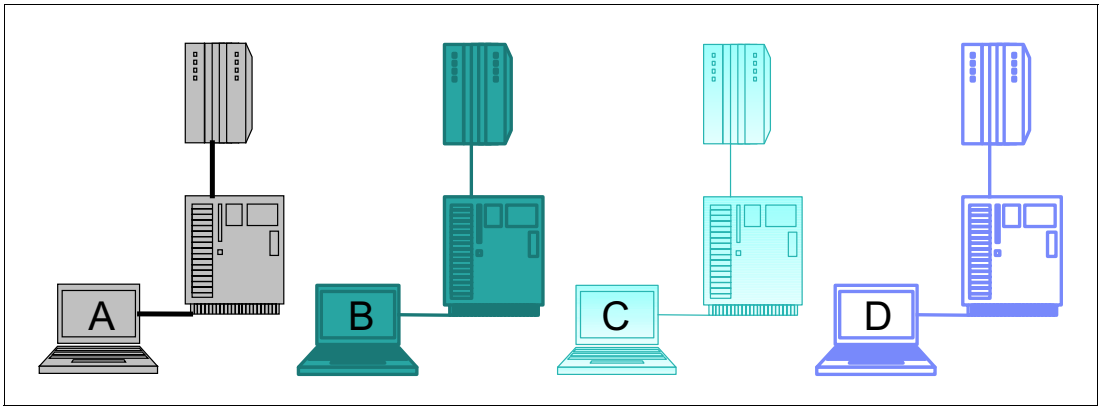

Figure 5-8 Cluster configuration diagram N + 1 and N to 1

If we assume there is a problem with node B, Figure 5-9 show how node D would take over. Now nodes A, C, and D are active.

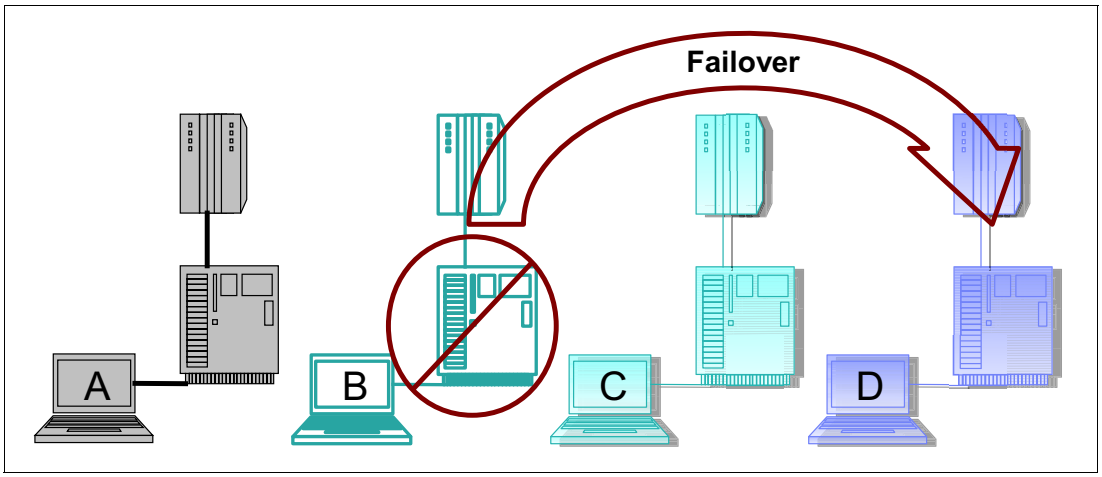

Figure 5-9 Cluster failover

After the problem with node B is fixed, node D can failback to node B as shown in Figure 5-10. Nodes A, B, and C are active again, node D is standby.

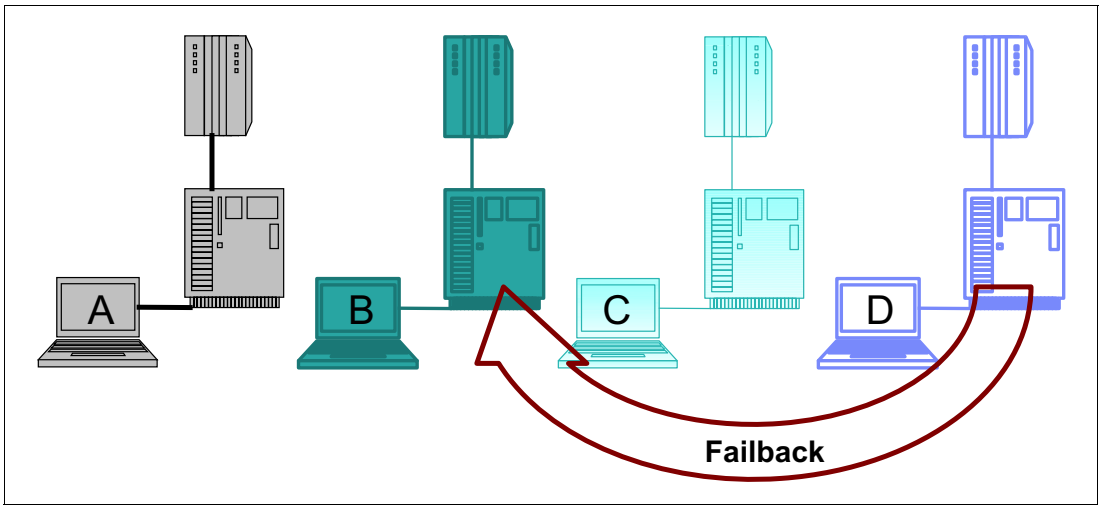

Figure 5-10 Cluster failback

Cluster nodes include the following similarities:

- Every node has access to all cluster configuration data.
- Nodes communicate with other cluster nodes through one or more physically independent networks (interconnects).

Network adapters, referred to in server clusters as network interfaces, attach nodes to networks:

- Every node in the cluster knows when another system joins or leaves.
- Every node in the cluster is aware of the resources that are running locally and also the resources that are running the other cluster nodes.
- All nodes in the cluster are grouped under the cluster name, which is used for accessing and managing the cluster.

When a node starts, it searches for active nodes on the networks designated for internal communication.

- If it finds an active node, it attempts to join the node's cluster.
- If it cannot find an existing cluster, it attempts to form a cluster by taking control of the quorum resource. The quorum resource stores the most current version of the cluster database, which contains cluster configuration and state data.

A server cluster maintains a consistent, updated copy of the cluster database on all active nodes.

The Tivoli Storage Manager server can be configured to work with any HA clustering solution, but is tested and documented for Microsoft Cluster for Windows, IBM HACMP™ and IBM PowerHA® on AIX, and Tivoli System Automation for Multiplatforms (SA MP) on Linux environments.

The following *Cluster support for IBM Tivoli Storage Manager server* technote has details of what to consider if you plan to use the Tivoli Storage Manager server with other HA products:

http://www-01.ibm.com/support/docview.wss?uid=swg21609772

On AIX, you have options to cluster the Tivoli Storage Manager database and log files:

- Logical Volume Manager (LVM) Mirroring
- Live Partition Mobility (LPM)
- PowerHA on AIX (HACMP)

What you choose depends on how automated the fail over of the Tivoli Storage Manager application should be, where LVM is less automated than LPM, which again is less automated than PowerHA.

In Linux environments, there are several cluster software solutions to choose from, such as Red Hat Cluster Suite or SteelEye LifeKeeper.

Oracle Solaris Cluster and HP UNIX Serviceguard are valid cluster software solutions for their respective platforms, Oracle Solaris or HP UNIX.

A solution that applies to clustering of all these platforms is Veritas Cluster Server. It provides application cluster capabilities to systems running other applications, including databases or network file sharing.

Tape fail over support in clustered environments

Because clustering means sharing the same resources among more hosts, the tape library also must be shared in order to be available. This means that the tape library must be a part of the SAN environment or be accessible in another way after a fail over. Proper SAN zoning configurations and SAN discovery should be enabled. Then, Tivoli Storage Manager can update the device relationships automatically if it changes in the environment. Persistent binding of tape drives is way of controlling the devices on the host.

To protect cluster storage resources, a Tivoli Storage Manager client can be used inside or outside the cluster. Consider the following information:

- Outside the cluster, cluster storage resources can be mapped to the client, which will then "follow" the resource to whichever node it may be assigned to.
- The disadvantage to being outside the cluster is that backup of data over network protocols such as NFS or CIFS is slower than the backup of a local file system.
- Placing the client within the cluster allows for faster backups and is supported by Tivoli Storage Manager's incremental forever and other fault tolerant features.

The Tivoli Storage Manager server itself can be included as a cluster resource, thus extending the availability of the data protection functionality.

Additional references

A recorded Storage Technical Exchange session discusses the Tivoli Storage Manager server in a clustered and high availability environment, which has further information, installation examples, and references to more information. You can find the session here:

http://www.ibm.com/support/docview.wss?uid=swg27036541

DB2 High Availability Disaster Recovery (HADR)

With Tivoli Storage Manager server at Version 6.2.2 and later, the use of the DB2 HADR, is supported. This functionality of DB2 provides a replication technique for the server's underlying database, allowing the Tivoli Storage Manager server to immediately restart on the standby host.

HADR is one part of the technique described in *Electronic vaulting using deduplicated remote copy storage pools*, available on the Tivoli Storage Manager wiki pages:

http://ibm.co/1o9YyDv

Also see 6.3, "Data protection solution using node replication feature" on page 128.

The other part of the technique uses remotely attached copy storage pools. Here we concentrate on the HADR side and show how you can use it to prevent delay in restoring the database, in case of a disaster.

When you plan to use HADR, you need to understand that HADR is a DB2 function outside of Tivoli Storage Manager, and requires to be licensed separately. Your Tivoli Storage Manager server is aware of the underlying DB2, but neither DB2 nor HADR is aware of the Tivoli Storage Manager server.

As shown with Figure 5-11 on page 105, HADR establishes a standby copy of the primary database, with the desired synchronous log synchronization mode for the Tivoli Storage Manager database. The HADR standby database is in constant *rollforward* mode to stay synchronized with the primary server.

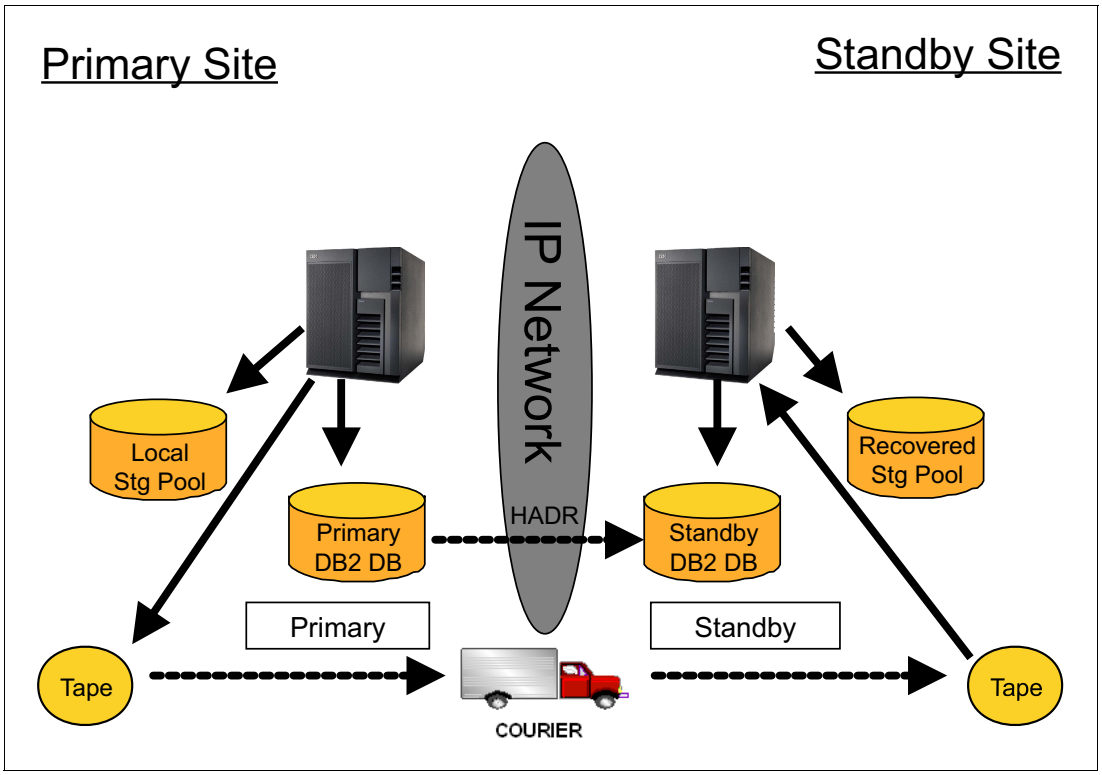

Figure 5-11 DB2 HADR: Synchronize primary and standby using IP

DB2 HADR is a supported and well known feature of DB2. Typically the standby server is initialized with a restore of a remote mounted full offline database backup. After they are started, the primary and standby server databases synchronize using IP to ship database updates.

The activity log[1] messages, from a failover test in the lab, as shown with Figure 5-12 on page 106, document how fast the standby server can restart in an HADR environment after the primary server becomes unavailable.

In our lab environment, the Tivoli Storage Manager server had the following results:

- Starts 36 seconds after shutdown.
- Is unavailable only for 1 minute 22 seconds.

Although these numbers were generated under lab conditions, and your results might vary, they show how fast the server can start by starting the warm standby server.

[1] The *activity* log contains messages that are normally sent to the server console during server options. The *active* log files record transactions that are in progress on the server.

```
02/14/2013 12:27:40 ANR1496I Stopping TSM Server on 4P22 (SESSION: 138)
02/14/2013 12:28:16 ANR4726I The NAS-NDMP support module has been loaded.
02/14/2013 12:28:36 ANR0984I Process 1 for IDENTIFY DUPLICATES started in
the BACKGROUND at 12:28:36. (PROCESS: 1)
02/14/2013 12:28:36 ANR0984I Process 2 for IDENTIFY DUPLICATES started in
the BACKGROUND at 12:28:36. (PROCESS: 2)
02/14/2013 12:28:38 ANR1412W Volume /tsmhastgp/svol002 access mode is
"unavailable".
02/14/2013 12:28:38 ANR1412W Volume /tsmhastgp/svol003 access mode is
"unavailable".
02/14/2013 12:28:38 ANR1412W Volume /tsmhastgp/svol004 access mode is
"unavailable".
02/14/2013 12:28:56 ANR0993I Server initialization complete.
02/14/2013 12:29:02 ANR2552I Server now enabled for Client access.
(SESSION: 5)
```

Figure 5-12 Activity log messages from a failover test in the lab

DB2 HADR helps you to meet your business continuity needs and is part of advanced electronic vaulting implementations.

Summary of database backup and restore

You can tailor the protection of your server database from a courier transporting tapes to an offsite location up to a disk-only database backup solution by using supported and documented technologies and procedures. With these data movement options, you are able to design the protection of your server database and infrastructure to meet your business continuity requirements.

5.3 Additional server infrastructure protection

Infrastructure setup files are prerequisites for recovering the Tivoli Storage Manager database and client data. In most cases, these files cannot be re-created, so you must ensure that copies are up-to-date and easily accessible.

5.3.1 Protecting the Tivoli Storage Manager database log files

You can configure the server to mirror the active log by specifying the `MIRRORLOGDIRECTORY` option, by placing the mirror on a file system that exists on a disk drive that is not on the same drive as the primary active log.

Imagine hardware issues where the write order of database pages to the actual hardware might be affected.

As a result of this, if the active logs are affected (such as a partial write to the storage), this represents a single point of failure where the log data necessary to “reconcile” the database for transaction consistency (crash recovery) is not available.

Having a mirror for the active log provides a higher degree of protection. The protection that is provided results in the active log data being less likely to also be affected by a hardware write issue.

Tips for database log file protection

The following tips and additional information can help protect the Tivoli Storage Manager database log files:

- Consider mirroring the active log and the archive log if retention protection is enabled. If restoring a database is needed, you can restore it to the current point in time with no data loss.
- You can dynamically start or stop mirroring while Tivoli Storage Manager is running.
- Despite its benefits, mirroring does not protect against a disaster or a hardware failure that affects multiple drives or causes the loss of the entire system. In addition, mirroring doubles the amount of disk space that is required for logs. Mirroring also results in decreased performance.

For more information, see the following IBM Knowledge Center topics:

- Protecting the database and infrastructure setup files:

  http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/topic/com.ibm.itsm.srv.doc/c_protect_db-reclog_ovr.html

- Protecting the active, archive, and archive failover logs:

  http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/topic/com.ibm.itsm.srv.doc/t_recovprotect.html

5.3.2 Protecting the volume history file

To restore the database, the server needs the information that is in a volume history file. You can specify duplicate volume history files. When the server updates volume information in the database, it also updates each file.

To specify the file path and name for a volume history file, use multiple VOLUMEHISTORY entries. Tivoli Storage Manager stores duplicate volume histories in all files that are specified with VOLUMEHISTORY options. To find the required volume history information during a database restore operation, the server tries to open volume history files in the order in which the VOLUMEHISTORY entries occur in the server options file.

If configured, Disaster Recovery Manager saves a copy of the volume history file in its disaster recovery plan file.

For more information, see the “Protecting the volume history file” topic:

http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/topic/com.ibm.itsm.srv.doc/t_save_volhist.html

5.3.3 Protecting the device configuration file

The device configuration file contains information that is required to read backup data and restore the database. You can specify duplicate device configuration files. When the server updates device configuration information in the database, it also updates each file. A device configuration file cannot be re-created.

The following device configuration information is stored in the Tivoli Storage Manager database and updated in the device configuration files:

- Devices class definitions
- Library definitions
- Drive definitions

- Path definitions
- Server definitions
- Database manager backup node ID

The device information must match the devices configured on the system where the restore operation can be performed. You might have to edit those commands in an existing file so that they match.

To specify the file path and name for a device configuration file, use the DEVCONFIG server option. To specify more than one path and name, use multiple DEVCONIG entries. Tivoli Storage Manager stores duplicate device configuration information in all the files that are specified with DEVCONFIG options. To find the required device configuration information during a database restore operation, the server tries to open device configuration files in the order in which the DEVCONFIG entries occur in the server options file. If the server cannot read a file, the server tries to open the next device configuration file.

To ensure the availability of device configuration information, use one or more of these steps:

- Store at least one copy of the device configuration file offsite or on a disk separate from the database.
- Store a printout of the file offsite.
- Store a copy of the file offsite with your database backups and volume history file.
- Store a remote copy of the file, for example, on an NFS-mounted file system.

If configured, Disaster Recovery Manager automatically saves a copy of the device configuration file in its disaster recovery plan file.

For more information, see the “Protecting the device configuration file” topic:

http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/topic/com.ibm.itsm.srv.doc/t_save_devconfig.html

5.3.4 Protecting the server options file

To restore the database, you need a copy of the server options file. The server options file includes the file paths of the active log, the archive log, the active log mirror, and the archive failover log. This information is required to restore the database.

To ensure the availability of server options file, use one or more of these steps:

- Store at least one copy of the server options file offsite or on a disk separate from the database.
- Store a printout of the file offsite.
- Store a copy of the file offsite with your database backups and device configuration file.
- Store a remote copy of the file, for example, on an NFS-mounted file system.

If configured, Disaster Recovery Manager automatically saves a copy of the server options file in its disaster recovery plan file.

For more information, see the “Protecting the server options file” topic:

http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/topic/com.ibm.itsm.srv.doc/t_save_opt_db_reclog_info.html

5.3.5 Protecting information about the database and recovery logs

To restore the database, you need detailed information about the database and recovery log. The recovery log includes the active log, the active log mirror, the archive log, and the archive failover log. The recovery log contains records of changes to the database.

You can determine the following information from the recovery log:

- Directory where the recovery log is located
- Amount of disk space required

If you lose the recovery log, you lose the changes that were made since the last database backup. If configured, Disaster Recovery Manager helps you save database and recovery log information.

For more information, see the “Protecting information about the database and recovery logs” topic:

http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/topic/com.ibm.itsm.srv.doc/t_dblog_info.html

5.3.6 Protecting the Secure Sockets Layer digital certificate file

As part of the process of setting up IBM Tivoli Storage Manager to use Secure Sockets Layer (SSL) for client-server authentication, a digital certificate file (`cert.kdb`) is created.

The `cert.kdb` file includes the server’s public key, which allows the client to encrypt data. The digital certificate file cannot be stored in the server database because the Global Security Kit (GSKit) requires a separate file in a certain format. The `cert256.arm` file is generated by the V6.3 server for distribution to the V6.3 clients.

Keep backup copies of the `cert.kdb` and `cert256.arm` files in a secure location. If both of the original files and any copies are lost or corrupted, you can generate a new certificate file.

If client data object encryption is in use and the encryption key is not available, data cannot be restored or retrieved under any circumstance. When using ENABLECLIENTENCRYPTKEY for encryption, the encryption key is stored on the server database. This means that for objects using this method, the server database must exist and have the proper values for the objects for a proper restore operation. Ensure that you back up the server database frequently to prevent data loss.

For more information, see the “Protecting the Secure Sockets Layer digital certificate file” topic:

http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/topic/com.ibm.itsm.srv.doc/c_protect_ssl_certkdb.html

5.3.7 Disaster Recovery Manager: Protecting the disaster recovery plan

The disaster recovery plan file contains the information needed to recover a Tivoli Storage Manager server to the point in time represented by the last database backup operation that is completed before the plan is created.

You can use server-to-server communications to store copies of the recovery plan on a remote target server, in addition to traditional disk-based files.

Storing recovery plan files on a target server provides the following advantages:

- A central repository for recovery plan files
- Automatic expiration of plan files
- Query capabilities for displaying information about plan files and their contents
- Fast retrieval of a recovery plan file if a disaster occurs

You can also store the recovery plan locally, printed, or on a disk.

For more information, see the “Protecting the disaster recovery plan” topic:

http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/topic/com.ibm.itsm.srv.doc/t_protect_drp.html

5.4 Summary

Protecting your Tivoli Storage Manager server infrastructure involves more than a backup copy of your Tivoli storage Manager database. Be sure to have multiple current copies of infrastructure objects necessary for a database restore operation. Because the Tivoli Storage Manager proprietary database has been replaced by a DB2 database, you now have even more options to protect the database. Protecting only the database is not enough, however. You must also have the infrastructure files if you need to restore your DB.

6

Tivoli Storage Manager Technologies and Solutions

Many businesses today have many of the same questions that you do:

- How do I respond to pressures to cut costs and to reduce risk and complexity?
- Data protection and storage architectures are already complex. How can I keep pace with technology?
- How can I react more quickly to take advantage of new business opportunities?
- How do I move my data protection architecture into to the future?

In this chapter, we offer several examples of how to solve these problems with various Tivoli Storage Manager technologies and solutions.

6.1 Providing a set of data protection and restore tools

As mentioned throughout this book, Tivoli Storage Manager provides a great set of products and features to design adaptive and a comprehensive data protection solutions. Whatever your data type and infrastructure size, Tivoli Storage Manager toolkit scales from a small environment, consisting of 10 - 20 machines, to a large environment with thousands of machines to protect.

First, in 6.2, “Disk-to-disk data protection solution using deduplication” on page 113 we explain how to implement a Tivoli Storage Manager solution using data deduplication. We show you how it works, and how important it is for your business.

Second, we describe in 6.3, “Data protection solution using node replication feature” on page 128, how to enlarge the architecture described previously by adding new Tivoli Storage Manager components and functionality.

Third, in 6.5, “Tivoli Storage Manager as a virtual appliance” on page 136 we discuss how you can implement a fully virtualized data protection and restore solution based on Tivoli Storage Manager products.

6.2 Disk-to-disk data protection solution using deduplication

The Tivoli Storage Manager deduplication is more than hype; your business requires it. With all the discussion about deduplication, where is the best place to start and how?

With an increasing data growth and an increased tape media handling, your business requires a change where data deduplication is a technology that removes redundant data to reduce the storage capacity requirement for retaining the data. When deduplication technology is applied to data protection, it can provide a highly effective means for reducing overall cost of a data protection solution.

Deduplication provides additional data reduction capabilities and can be used in addition to the progressive incremental and compression capabilities already available in the product. Figure 6-1 shows the data reduction features in Tivoli Storage Manager.

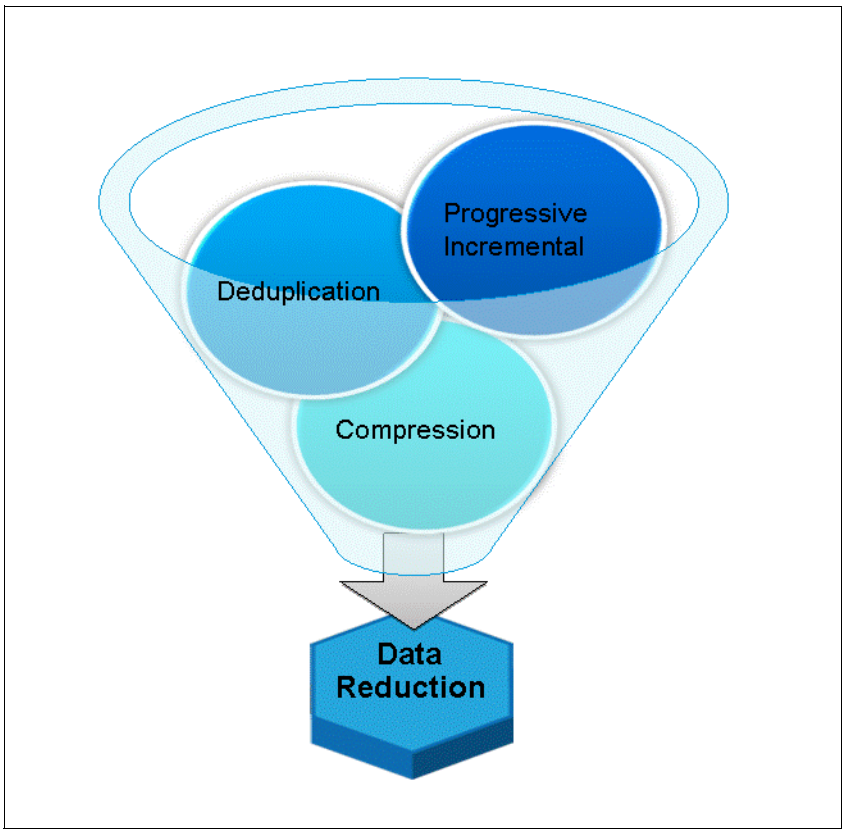

Figure 6-1 Overview of the data reduction features in Tivoli Storage Manager

This section describes the benefits of deduplication and provides guidance for how to make effective use of the Tivoli Storage Manager deduplication feature as part of a well-designed data protection solution. We also provide information for reference purposes and guidance during the planning of other deployments of Tivoli Storage Manager. The architectures are not intended to be a suitable configuration for all situations, and our intention is to describe the benefits of deduplication and with guidance on, how to make effective use of the Tivoli Storage Manager deduplication feature as part of a well-designed data protection solution.

Remember, this document does not replace the need to carefully plan and design the elements of your own implementation of Tivoli Storage Manager.

6.2.1 Benefits deduplication provides

Consider the following key points for using Tivoli Storage Manager deduplication:

- The Tivoli Storage Manager deduplication is an effective tool for reducing overall cost of a backup solution.
- Cost reduction is the result of data reduction where deduplication is just one of several methods that Tivoli Storage Manager provides for data reduction (for example, progressive incremental backup and compression). The goal is overall data reduction when all of the techniques are combined, rather than just on the deduplication ratio.
- Tivoli Storage Manager deduplication is an appropriate data reduction method for many situations. It can also be used as a cost effective option for backing up a subset of an environment that uses a deduplication appliance for the remaining backups.
- Tivoli Storage Manager deduplication can operate on backup, archive, and HSM data. This includes data which is stored via the Tivoli Storage Manager API.
- Additional resources (database capacity, processor, and memory) must be configured for a Tivoli Storage Manager server that is configured with Tivoli Storage Manager deduplication. However, when properly configured, the benefit of storage pool capacity reduction will result in a significant cost reduction benefit.
- Tivoli Storage Manager deduplication is an appropriate data reduction method for many situations. It can also be used as a cost-effective option for backing up a subset of an environment that uses a deduplication appliance for the remaining backups.

Summary

The following list is a short summary of benefits that the Tivoli Storage Manager deduplication solution provides:

- Decreases the amount of storage capacity required to contain data growth. Remember, cost reduction is the result of data reduction; deduplication is just one of several methods that provides for data reduction (such as progressive incremental backup). The goal is overall data reduction when all techniques are combined, rather than just on the deduplication ratio.
- Provides an affordable solution that can be expanded in the future and that delivers a rapid return on investment (ROI) and a lower total cost of ownership (TCO) to justify the investment.
- Network bandwidth optimization (incremental, data deduplication) sends less data over the network.
- Reduces cost in labor with less or no tape media handling at all.

6.2.2 Solution architecture

We describe several examples of Tivoli Storage Manager architectures that can make the most effective use of Tivoli Storage Manager deduplication. Figure 6-2 on page 115 shows Tivoli Storage Manager Server architecture, which provides a disk-based backup solution that retains data in a deduplicated storage pool for its entire retention. Data protection is provided for hundreds of clients in the deduplicated storage pool. Architectural components that are used in this example are the Tivoli Storage Manager Server, Tivoli Storage Manager Client, storage pools, local area network (LAN), and storage area network (SAN.

Figure 6-2 on page 115 illustrates the architecture that includes two primary storage pool hierarchies. The first is a deduplicated sequential-file storage pool. Data remains in this pool for its entire retention and is never allowed to migrate to another storage pool.

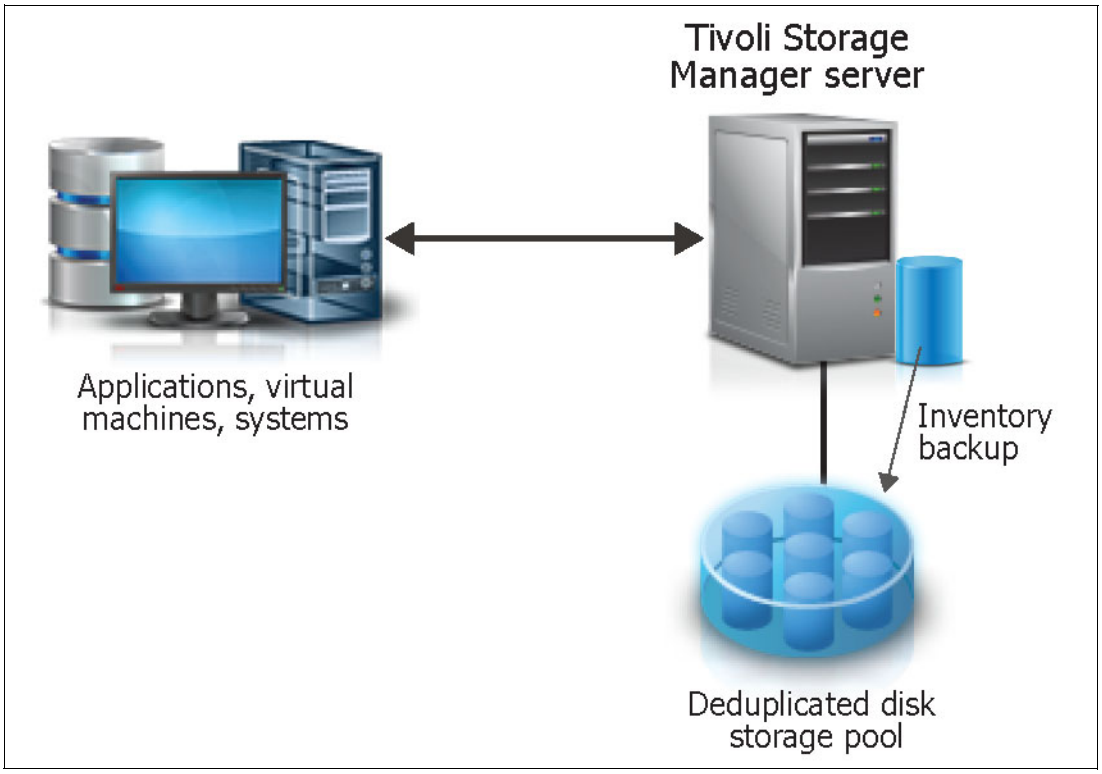

Figure 6-2 Tivoli Storage Manager architecture

The remainder of this section summarizes key aspects of the Tivoli Storage Manager architecture where the client backups are ingested over the LAN to two primary storage pool destinations.

A deduplicated primary file-based disk storage pool is used for a majority of clients with a mix of both client-side and server-side deduplication. Objects are stored in this pool for their entire retention. A random disk-based primary storage pool is used as clients ingest into a random disk storage pool which cannot use deduplication.

Deduplicated primary pool

The deduplicated primary storage pool retains backup objects for their entire retention. A majority of clients ingest backups directly into this storage pool. Of these clients, some use client-side deduplication during backup ingestion (those that are eligible for client-side deduplication, according to considerations listed here). The other clients back up without deduplication, allowing the data to be later deduplicated using server-side deduplication.

The following considerations help to determine which clients use this deduplicated storage pool:

- Fast restore times are needed without the delay that is associated with tape mounts, and data spread across multiple tapes.
- The data responds well to deduplication in terms of the amount of reduction.
- Data is not encrypted on the client.

Random disk storage pool

The remaining clients ingest into a random disk storage pool, which cannot use deduplication. The following factors determine which clients ingest into this storage pool:

- Clients with large objects (greater than 2 TB) that are not suitable for deduplication.
- Lower-priority clients that do not require faster restore times, making the lower-cost tape storage more desirable, or require the use of encryption.

Having a second copy storage pool using disk (that can also be deduplicated) is another option. Server-side deduplication is a two-step process:

1. Duplicate identification
2. Removal of the excess data during a subsequent data movement process such as reclamation or migration

The second step can be done after a storage pool backup copy is created. See the description of the `deduprequiresbackup` option at the IBM Tivoli Storage Manager Version V7.1 page:

http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1/index.jsp

When using server-side deduplication, schedule the storage pool backup process prior to the reclamation processing to ensure that there is minimal overhead when copying the data. After identify duplicate has run, the data is not deduplicated but it is redefined such that it can be reconstructed and dehydrated during the subsequent data movement operation. Sufficient time must be allotted for the scheduled storage pool backup to complete before the start of the schedule for reclamation.

When using client-side deduplication, the storage pool backup processing always occurs after data has been deduplicated. This requires deduplicated data to be reconstructed during the copy (if the copy storage pool is not also deduplicated). The reconstruction processing can result in storage pool backup processing, which is slower when compared with storage pool backup processing of data that has not been deduplicated. For planning purposes, estimate that the duration of storage pool backup will be doubled for data which is already deduplicated.

6.2.3 Disk-to-disk data protection

Disk-to-disk data protection refers to the scenario where the preferred backup storage device is disk-based, as opposed to tape or a virtual tape library (VTL). Disk-based data protection is now more popular because the unit cost of disk storage has fallen. It is also more common as companies distinguish between backup data, which is kept for a relatively short amount of time, and archive data, which has long-term retention.

A secondary copy is suggested, preferably using node replication, but not required. However, with disk-to-disk data protection, the primary storage pool data remains on disk until it expires. You can achieve a significant reduction of disk storage if the primary storage pool is configured for deduplication.

6.2.4 Solution description

Deduplication technology uses a computational technique to detect patterns within data that appear multiple times within the scope of a collection of data. For the purposes of this document, the collection of data consists of Tivoli Storage Manager backup, archive, and HSM data (these types of data are referred to as “backup data” throughout this document). The patterns that are detected are represented as a hash value that is much smaller than the

original pattern, specifically 20 bytes. Except for the original instance of the pattern, subsequent instances of the chunk are referenced by the hash value. As a result, for a pattern that appears many times throughout a collection of data, significant reduction in storage can be achieved.

Unlike compression, deduplication can take advantage of a pattern that occurs multiple times within a collection of data. With compression, a single instance of a pattern is represented by a smaller amount of data that is used to algorithmically re-create the original data pattern. Compression cannot take advantage of data redundancy for patterns that recur throughout the collection of data, and this significantly reduces the potential reduction capability. However, compression can be combined with deduplication to take advantage of both techniques and further reduce the required amount of data storage beyond what would be required by using just one technique or the other.

6.2.5 How Tivoli Storage Manager performs deduplication

Tivoli Storage Manager uses a proprietary algorithm to analyze variable sized, contiguous segments of data, called *chunks*, for patterns that are likely to be duplicated within the same Tivoli Storage Manager storage pool. This process is explained in more detail in various sections of this chapter.

The implementation of Tivoli Storage Manager deduplication applies only to the FILE device class (sequential-access disk) storage pools, and can be used with primary, copy, or active-data pools.

6.2.6 Data reduction and data deduplication

Data deduplication creates substantial opportunity for reduction of storage capacity requirements for backup data. However, deduplication within the context of other data reduction techniques that are available is an important consideration. When you consider the effectiveness of deduplication, the deduplication ratio or percentage of reduction is considered to be the ultimate measurement of effectiveness. However, a more important consideration is the overall effectiveness of data reduction, including deduplication and other techniques that are available, rather than focusing exclusively on deduplication effectiveness.

Unlike other backup products, Tivoli Storage Manager provides a substantial advantage in data reduction through its incremental-forever technology. Combined with deduplication, compression, exclusion of specified objects, and appropriate retention policies, Tivoli Storage Manager provides highly effective data reduction.

Therefore, the business objectives should be clearly defined and understood when you consider how to measure data reduction effectiveness. If reduction of storage and infrastructure costs is the ultimate goal, the focus will be on overall data reduction effectiveness, with data deduplication effectiveness as one component.

6.2.7 Server-side and client-side deduplication

Tivoli Storage Manager provides two methods for performing deduplication: client-side and server-side deduplication. Both methods use the same algorithm to identify redundant data, however the “when” and “where” of deduplication processing differ.

Server-side deduplication

With server-side deduplication, all processing of redundant data occurs on the Tivoli Storage Manager server, after the data is backed up. Server-side deduplication is also called *target-side deduplication*. The key characteristics of server-side deduplication are as follows:

- Duplicate data is identified after backup data has been transferred to the storage pool volume.
- Both server-side and client-side deduplication can be used together within the same storage pool, and reduce data across backups regardless of which type of deduplication is used.
- Duplicate identification processing must run regularly on the server, and consumes Tivoli Storage Manager server processor and Tivoli Storage Manager database resources.
- Storage pool data reduction is not realized until data from the deduplication storage pool is moved to another storage pool volume, usually through a reclamation process, but can also occur during a Tivoli Storage Manager MOVE DATA process.

Client-side deduplication

Client-side deduplication processes the redundant data during the backup process on the host system where the source data is located. The net results of deduplication are virtually the same as with server-side deduplication, except that the storage savings are realized immediately, because only the unique data needs to be sent to the server in its entirety. Data that is duplicate requires only a small signature to be sent to the Tivoli Storage Manager server. Client-side duplication is especially effective when conserving bandwidth between the Tivoli Storage Manager client and server is important.

Client deduplication cache

Although it is necessary for the backup client to “check in” with the server to determine whether a chunk is unique or a duplicate, the amount of data transfer is small. The client must query the server for each chunk of data that is processed. The overhead associated with this query process can be reduced substantially by configuring a cache on the client, which allows previously discovered chunks on the client (during the backup session) to be identified without a query to the Tivoli Storage Manager server.

For the backup-archive client (including VMware backup), a suggestion is to always configure a cache when you use client-side deduplication. For applications that use the Tivoli Storage Manager API, the deduplication cache should not be used because of the potential for backup failures that are caused by the cache being out of sync with the Tivoli Storage Manager server. If multiple, concurrent Tivoli Storage Manager client sessions are configured (such as with a Tivoli Storage Manager for VMware vStorage backup server), a separate cache must be configured for each session.

In some conditions, faster performance is possible when deduplication cache is disabled. When the network between the clients and server has high bandwidth and low latency and the Tivoli Storage Manager server database is on fast storage, deduplication queries that are directly to the Tivoli Storage Manager server can outperform queries to the local cache.

6.2.8 Prerequisites for configuring Tivoli Storage Manager deduplication

Certain prerequisites must be met with Tivoli Storage Manager deduplication. For a complete list of prerequisites, see the Tivoli Storage Manager administrator documentation:

http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1/index.jsp

The prerequisites are as follows:

- Prerequisites that are common to client and server-side deduplication:
  - The destination storage pool must be of type “FILE” (sequential disk)
  - The target storage pool must have the deduplication setting enabled
  - The Tivoli Storage Manager database must be configured according to best practices for high performance
- Prerequisites that are specific to client-side deduplication:
  - The client and server must be at version 6.2.0 or later. Always use the most recent maintenance version.
  - The client must have the client-side deduplication option enabled (`DEDUPLICATION YES`).
  - The server must enable the node for client-side deduplication with the **`DEDUP=CLIENTORSERVER`** parameter using either the **`REGISTER NODE`** or **`UPDATE NODE`** command.
  - The target storage pool must be a deduplication-enabled storage pool.
  - Files must be bound to a correct management class whose destination is a deduplication-enabled storage pool.
  - Files must not be excluded from client-side deduplication processing (by default, all files are included). See the “exclude.dedup” client option on the IBM Tivoli Storage Manager Version 6.4 web page:

    http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/index.jsp
  - Files must be larger than 2 KB, and transactions must be below the value that is specified by the `clientdeduptxnlimit` option.
- Compatibility issues

  The following Tivoli Storage Manager features are incompatible with Tivoli Storage Manager client-side deduplication:
  - Client encryption
  - LAN-free/storage agent
  - UNIX HSM client
  - Subfile backup
  - Simultaneous storage pool write

6.2.9 Exceptions to using Tivoli Storage Manager deduplication

Tivoli Storage Manager deduplication can provide significant benefits and cost savings, but it does not apply to all situations. The following situations are not appropriate for using Tivoli Storage Manager deduplication:

- Primary storage of backup data is on VTL or physical tape

  Movement to tape requires “rehydration” of the deduplicated data. This takes extra time and requires processing resources. If regular migration to tape is required, the benefits of using Tivoli Storage Manager deduplication may be reduced, because the goal is to reduce disk storage as the primary location of the backup data.

- No flexibility with the backup processing window

  Tivoli Storage Manager deduplication processing requires extra resources, which can extend backup windows or server processing times for daily backup activities. For example, a duplicate identification process must run for server-side deduplication. More reclamation activity is required to remove the duplicate data from a storage pool after the duplicate identification processing completes. For client-side deduplication, the client backup speed will generally be reduced for local clients (remote clients may not be impacted if a bandwidth constraint exists).

  If the backup window has already reached the limit for service level agreements, Tivoli Storage Manager deduplication might possibly impact the backup window further unless careful planning is done.

Restore performance considerations

The restore performance from deduplicated storage pools is slower than from a comparable disk storage pool that does not use deduplication. However, restoring from a deduplicated storage pool can compare favorably to restoring from tape devices for certain workloads.

If the fastest restore performance from disk is a high priority, then restore performance benchmarking should be done to determine whether the effects of deduplication can be accommodated. Table 6-1 compares the restore performance of small and large object workloads across several storage scenarios.

Table 6-1 Comparison of restore performance

Storage pool type	Small object workload	Large object workload
Tape	Typically slower due to tape mounts and seeks	Typically faster due to streaming capabilities of modern tape drives
Non-deduplicated disk	Typically faster due to absence of tape mounts and quick seek times	Comparable to or slightly slower than tape
Deduplicated disk	Faster than tape, slower than non-deduplicated disk	Slowest since data must be rehydrated, when compared to tape which is fast for streaming large objects that are not spread across many tapes.

6.2.10 Compare Tivoli Storage Manager deduplication and appliance deduplication

Tivoli Storage Manager deduplication provides the most cost effective solution for reducing backup storage costs, because there is no additional software license charge for it, and it does not require special purpose deduplicating hardware appliances. Deduplication of backup data can also be accomplished by using a hardware deduplicating storage device in the Tivoli Storage Manager storage pool hierarchy. Deduplication appliances, such as IBM ProtecTIER, provide deduplication capability at the storage device level. NAS devices are also available that provide NFS or CIFS mounted storage that removes redundant data through deduplication.

An optimal balance can be made between Tivoli Storage Manager deduplication and storage appliance deduplication. Both techniques can be used in the same environment for separate storage hierarchies or in separate Tivoli Storage Manager server instances. For example, Tivoli Storage Manager client-side deduplication is an ideal choice for backing up remote environments, either to a local Tivoli Storage Manager server or to a central data center. Tivoli

Storage Manager node replication can then take advantage of the deduplicated storage pools to reduce data transfer requirements between Tivoli Storage Manager servers, for disaster recovery purposes.

Alternatively, within a large data center, a separate Tivoli Storage Manager server can be designated for backing up a critical subset of all hosts that use Tivoli Storage Manager deduplication. The remaining hosts can back up to a separate Tivoli Storage Manager server instance that uses a deduplicating appliance, such as ProtecTIER, for its primary storage pool and also supports replication of the deduplicated data.

Tivoli Storage Manager deduplication should not be used in the same storage hierarchy as a deduplicating appliance. For a deduplicating VTL, the Tivoli Storage Manager storage pool data must be “rehydrated” before moving to the VTL (as with any tape device), and there is no data reduction as a result of the Tivoli Storage Manager deduplication; rather it will be re-deduplicated by the VTL. For a deduplicating NAS device, a FILE device type can be created on the NAS. However, because the data is already deduplicated by Tivoli Storage Manager, there is little to no additional data reduction possible by the NAS device.

6.2.11 Factors when deciding between Tivoli Storage Manager and appliance deduplication

Scale, scope, and cost are the three major factors to consider when you are deciding which deduplication technology to use.

Scale

The Tivoli Storage Manager deduplication technology is a scalable solution, which uses software technology that heavily uses Tivoli Storage Manager database transactions. The deduplication processing has an impact on daily server processes such as reclamation and storage pool backup. For a specific Tivoli Storage Manager server hardware configuration (for example, Tivoli Storage Manager database disk speed, processor and memory capability, and storage pool device speeds), a practical guideline can be followed regarding the amount of data that can be backed up using deduplication.

Consider the following two primary points of scalability

- Daily amount of new data that is ingested
- Total amount of data to be protected over time

The practical guideline described is not built in to the product, and can vary based on the capabilities of the hardware which is used. At the time of writing this book, the amount of protected data is presented as a guideline with the purpose of keeping the size of the Tivoli Storage Manager database below the recommended size of 4 TB. A database of this size (4 TB) corresponds roughly to 400 TB of protected data (original data plus all retained versions). There is no harm in occasionally exceeding the limit for daily ingest, which is prescribed with the goal of allowing enough time each day for the Tivoli Storage Manager Server maintenance tasks to run efficiently. Regularly exceeding the practical limit on daily ingest for your specific hardware might impact the ability to achieve the maximum possible amount of data reduction, or cause backup durations to run longer than you want.

Deduplication appliances have dedicated resources for deduplication processing and do not have a direct impact on Tivoli Storage Manager server performance and scalability. If you want to scale up a single Tivoli Storage Manager server instance as much as possible, beyond approximately 400 TB of protected data (original data plus all retained versions), then consider appliance deduplication. However, often a more cost-effective approach is to scale

out with additional Tivoli Storage Manager server instances. Using additional Tivoli Storage Manager server instances can help you to manage many multiples of 400 TB protected data.

The amount of data that can be ingested each day depends on the capability of the hardware that is used for the Tivoli Storage Manager server. A primary consideration is the speed of the disk used for the Tivoli Storage Manager database. To a lesser degree, the speed of the disk used for the storage pool is also important. A high-end Tivoli Storage Manager server that uses solid-state drives (SSD) for the database and serial-attached SCSI (SAS) disk drives for the storage pool can handle up to 20 TB per day with server-side deduplication, and 30 TB per day using client-side deduplication.

Scope

The scope of Tivoli Storage Manager deduplication is limited to a single Tivoli Storage Manager server instance and more precisely within a Tivoli Storage Manager storage pool. A single, shared deduplication appliance can provide deduplication across multiple Tivoli Storage Manager servers. When Tivoli Storage Manager node replication is used in a many-to-one architecture, such as with branch offices, the deduplicated storage pool on the replication target can deduplicate across the data incoming from the multiple source servers.

Cost

Tivoli Storage Manager deduplication functionality is embedded in the product without an additional software license cost. In fact Tivoli Storage Manager software license costs will reduce when capacity-based licensing is in force because the capacity is calculated after deduplication has occurred. An important consideration is that hardware resources must be appropriately sized and configured. Anticipate extra expense when you plan a Tivoli Storage Manager server configuration that will be used with deduplication. However, these additional costs can easily be offset by the savings in disk storage. Also, the software license costs are reduced when capacity-based pricing is in effect.

Deduplication appliances are priced for the performance and capability that they provide, and generally are considered more expensive per gigabyte than the hardware requirements for Tivoli Storage Manager native deduplication. A detailed cost comparison should be done to determine the most cost-effective solution.

6.2.12 Conditions for effective use of Tivoli Storage Manager deduplication

Although Tivoli Storage Manager deduplication provides a cost-effective and convenient method for reducing the amount of disk storage required for backups, specific conditions can provide the most benefit when using Tivoli Storage Manager deduplication. Conversely, conditions exist in which Tivoli Storage Manager deduplication is not effective and in fact might reduce the efficiency of a backup operation.

The following conditions lead to effective use of Tivoli Storage Manager deduplication:

- Need for reduction of the disk space required for backup storage.
- Need for remote backups over limited bandwidth connections.
- Use of Tivoli Storage Manager node replication for disaster recovery across geographically dispersed locations.
- Total amount of backup data and ingested data per day are within the recommended limits of less than 400 TB total and 30 TB per day for each Tivoli Storage Manager server instance (to mitigate this you can deploy a second Tivoli Storage Manager server).
- Either a disk-to-disk backup should be configured (where the final destination of backup data is on a deduplicating disk storage pool) or data should reside in the FILE storage pool

for a significant time (for example, 30 days), or until expiration. The deduplication storage pools should not be used as a temporary staging pool before moving to tape or another non-deduplicating storage pool because this can be highly inefficient.

- Backup data should be a good candidate for data reduction through deduplication. This topic is covered in greater detail in “Resource requirements for Tivoli Storage Manager deduplication” on page 123.
- High-performance disk must be used for the Tivoli Storage Manager database to provide acceptable Tivoli Storage Manager deduplication performance.

6.2.13 Traditional Tivoli Storage Manager architectures compared with deduplication architectures

A traditional Tivoli Storage Manager architecture ingests data into disk storage pools, and moves this data to tape on a frequent basis to maintain adequate free space on disk for continued ingestion. An architecture that includes deduplication changes this model to store the primary copy of data in a sequential file storage pool for its entire life cycle. Deduplication provides enough storage savings to make keeping the primary copy on disk an affordable possibility.

6.2.14 Resource requirements for Tivoli Storage Manager deduplication

Tivoli Storage Manager deduplication provides significant benefits as a result of its data reduction technology, particularly when combined with other data reduction techniques available with Tivoli Storage Manager. However, the use of deduplication in Tivoli Storage Manager adds extra requirements for hardware and database and recovery log storage, which are essential for a successful implementation. When configuring Tivoli Storage Manager to use deduplication, you must ensure that proper resources have been allocated to support the use of the technology. The resources include hardware requirements necessary to meet the additional processing performed during deduplication, additional storage requirements for handling the Tivoli Storage Manager database records used to store the deduplication catalog, and extra storage requirements for the Tivoli Storage Manager server database logs.

The Tivoli Storage Manager internal database plays a central role in enabling deduplication technology. Deduplication requires extra database capacity to be available. In addition, there is a significant increase in the frequency of references to records in the database during many Tivoli Storage Manager operations including backup, restore, duplicate identification, and reclamation. These demands on the database require that the database disk storage be capable of sustaining higher rates of I/O operations than are required without the use of deduplication.

As a result, planning for resources used by the Tivoli Storage Manager database is critical for a successful deduplication deployment.

6.2.15 Database and log size requirements

Use of Tivoli Storage Manager deduplication significantly increases capacity requirements of the Tivoli Storage Manager database. This section provides several guidelines for estimating the capacity requirements of the database. It is important to plan ahead for the database capacity so an adequate amount of higher-performing disk can be reserved for the database (see 6.2.16, “Tivoli Storage Manager database log size estimation” on page 124 for performance requirements).

The use of deduplication in Tivoli Storage Manager requires more storage space in the Tivoli Storage Manager server database than without the use of deduplication. An important factor to consider is that when using deduplication, the Tivoli Storage Manager database grows proportionally to the amount of data that is stored in deduplicated storage pools. This is because each “chunk” of data that is stored in a deduplicated storage pool is referenced by an entry in the database.

Without deduplication, each backed-up object (typically a file) is referenced by a database entry, and the database grows proportionally to the number of objects that are stored. With deduplication, the database grows proportionally to the total amount of data backed up.

The following web page (“Determining the impact of deduplication on Tivoli Storage Manager server database and storage pools”) provides information for estimating the amount of disk storage that will be required for your Tivoli Storage Manager database. It also has formulas for estimating database size, based on the volume of data to be stored.

http://www.ibm.com/support/docview.wss?uid=swg21596944

As a simplified guideline for taking a rough estimate, you can plan for 10 GB of database storage for every 1 TB of data that will be protected in deduplicated storage pools.

The estimation guidelines are approximate, because actual requirements depend on many factors including those that cannot be predicted in advance (for example, a change in the data backup rate, the exact amount of backup data, and other factors).

6.2.16 Tivoli Storage Manager database log size estimation

The use of deduplication adds extra requirements for the Tivoli Storage Manager server database, active log, and archive log storage. Properly sizing the storage capacity for these components is essential for a successful implementation of deduplication.

- Planning active log space requirements

  The database active log stores information about database transactions that are in progress. With deduplication, transactions can run longer, requiring more space to store the active transactions.

  > **Tip:** A suggestion is to begin with an active log size of 120 GB and monitor the space usage and adjust the size of the active log as needed.

- Planning archive log space requirements

  The archive log stores older log files for completed transactions until they are cleaned up as part of the Tivoli Storage Manager server database backup processing. The file system holding the archive log must be given sufficient capacity to avoid running out of space, which can cause the Tivoli Storage Manager server to halt. Space is freed in the archive log every time a full backup is performed of the Tivoli Storage Manager server’s database.

  See the document about sizing the Tivoli Storage Manager archive log for information about how to calculate the space requirements for the Tivoli Storage Manager server archive log:

  http://www.ibm.com/support/docview.wss?uid=swg21389352

  Note that a file system with 500 GB of free space has proven to be more than adequate for a large-scale Tivoli Storage Manager server that ingests several terabytes a day of new data into deduplicated storage pools and performs a full Tivoli Storage Manager database backup once a day.

6.2.17 Estimating capacity for deduplicated storage pools

Tivoli Storage Manager deduplication ratios typically range from 2:1 (50% reduction) to 15:1 (93% reduction), and is data-dependent. Lower ratios are associated with backups of unique data (such as progressive incremental data); higher ratios are associated with backups that are repeated, such as repeated full backups of databases or virtual machine images. Mixtures of unique and repeated data will result in ratios within that range. If you are not sure of what type of data you have and how well it will reduce, use a ratio of 3:1 for planning purposes when comparing with non deduplicated Tivoli Storage Manager storage pool occupancy. This ratio corresponds to an overall data reduction ratio of over 15:1 when factoring in the data reduction benefits of progressive incremental backups.

6.2.18 Estimating storage pool capacity requirements

Consider the following information when you calculate storage disk sizing, so you have sufficient storage pool capacity for ingested data:

- Delayed release of storage pool data

  Due to the latency for deletion of data chunks with multiple references, there is a need for “transient” storage associated with data chunks that must remain in a storage pool volume even though their associated file or object is deleted or expired. As a result of this behavior, storage pool capacity sizing must account for some percentage of data that is retained because of references by other objects. This latency results in the delayed deletion of a storage pool volume if it contains a single chunk that is still being referenced.

- Delayed effect of post-identification processing

  Storage reduction does not always occur immediately with Tivoli Storage Manager deduplication. In the case of server-side deduplication, sufficient storage pool capacity is required to ingest the full amount of daily backup data. With server-side deduplication, removal of redundant data does not occur until after storage pool reclamation completes, which in turn may not complete until after a storage pool backup is done. If client-side deduplication is used, this delay will not apply. Sufficient storage pool free capacity must be maintained to accommodate continued backup ingestion.

- Estimating storage pool capacity requirements

  You can roughly estimate storage pool capacity requirements for a deduplicated storage pool by using the following technique:

  – Estimate the base size of the source data.
  – Estimate the daily backup size, using an estimated change and growth rate.
  – Determine retention requirements.
  – Estimate the total amount of source data by factoring in the base size, daily backup size, and retention requirements.
  – Apply the deduplication ratio factor.
  – Increase the estimate to consider transient storage pool usage.

6.2.19 Database I/O requirements

For optimal performance, fast disk storage is always preferred for the Tivoli Storage Manager database as measured in terms of input/output operations per second (IOPS). Due to the random access I/O patterns of the Tivoli Storage Manager database, minimizing the latency of operations that access the database volumes is critical for optimizing the performance of the Tivoli Storage Manager server. The large tables used for storing deduplication information

in the Tivoli Storage Manager database bring about an even more significant demand for disk storage that can handle a large number of IOPS.

In general, systems based on SSD technology and SAS disk drives, and Fibre Channel provide the best capabilities in terms of increased IOPS. Because the claims of disk manufacturers are not always reliable, we suggest that you measure actual IOPS of a disk system before implementing a new Tivoli Storage Manager database. The highest levels of daily ingest (greater than 8 TB per day) will require the database to use solid-state storage.

Details about how to configure high performing disk storage are beyond the scope of this document. Consider the following key points when you configure disk storage for the Tivoli Storage Manager database:

- The disk used for the Tivoli Storage Manager database should be configured according to best practices for a transactional database.
- Low-latency, enterprise-class disk devices or storage subsystems should be used for the Tivoli Storage Manager database.
- Disk devices or storage systems that are capable of a minimum of approximately 3000 IOPS are suggested for the Tivoli Storage Manager Database disk device. Consider an extra 1000 IOPS per TB of daily ingested data (pre-deduplication). Lower-performing disk devices can be used, but performance might not be optimal. See the Deduplication FAQs for an example configuration:

  https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli+Storage+Manager/page/Deduplication+FAQ
- Disk I/O should be distributed over as many disk devices and controllers as possible.
- Tivoli Storage Manager database and logs should be configured on separate disk volumes (LUNS), and should not share disk volumes with the Tivoli Storage Manager storage pool or any other application or file system.

6.2.20 Hardware requirements for Tivoli Storage Manager client deduplication

Client-side deduplication (and compression if used with deduplication) requires resources on the client system for processing. Prior to deciding to use client-side deduplication you should verify that client systems have adequate resources available during the backup window to perform the deduplication processing.

6.2.21 Processor

The use of deduplication requires additional processor resources on the Tivoli Storage Manager server, particularly for performing the task of duplicate identification. Consider using a minimum of at least eight (2.2 GHz or equivalent) processor cores in any Tivoli Storage Manager server that is configured for deduplication.

A suggested minimum CPU requirement is the equivalent of one 2.2 GHz CPU core per backup process with client-side deduplication. As an example, a system with a single-socket, quad-core, 2.2 Ghz processor that is used 75% or less during the backup window can be a good candidate to use client-side deduplication.

6.2.22 Memory

For the highest performance of a large-scale Tivoli Storage Manager server using deduplication, use additional memory. The memory is used to optimize the frequent lookup of deduplication chunk information stored in the Tivoli Storage Manager database. Consider a minimum of 64 GB of system memory for Tivoli Storage Manager servers that use deduplication. If the retained capacity of backup data grows, the memory requirement might need to be as high as 192 GB. Remember to monitor memory use on a regular basis to determine if more memory is required.

6.2.23 Solution guidelines

A successful implementation of Tivoli Storage Manager deduplication requires careful planning in the following areas:

- Implementing an appropriate architecture suitable for using deduplication
- Properly sizing your Tivoli Storage Manager server hardware and storage
- Configuring Tivoli Storage Manager following best practices for separating data ingestion and data maintenance tasks

6.2.24 Deciding between client and server deduplication

After you decide on an architecture using deduplication for your Tivoli Storage Manager server, decide whether you will perform deduplication on the Tivoli Storage Manager clients, the Tivoli Storage Manager server, or use a combination of the two. The Tivoli Storage Manager deduplication implementation allows storage pools to manage deduplication performed by both clients and the Tivoli Storage Manager server. The server is optimized to only perform deduplication on data that has not been deduplicated by the Tivoli Storage Manager clients. Furthermore, duplicate data can be identified across objects regardless of whether the deduplication is performed on the client or server. These benefits allow for hybrid configurations that efficiently apply client-side deduplication to a subset of clients, and use server-side deduplication for the remaining clients.

Typically a combination of both client-side and server-side data deduplication is the most appropriate. Here are some further points to consider:

- Server-side deduplication is a two-step process of duplicate data identification followed by reclamation to remove the duplicate data. Client-side deduplication stores the data directly in a deduplicated format, eliminating the need for the extra reclamation processing.
- Deduplication on the client can be combined with compression to provide the largest possible storage savings.
- Client-side deduplication processing can increase backup durations. Expect increased backup durations if network bandwidth is not restrictive. A doubling of backup durations is a reasonable estimate when using client-side deduplication in an environment that is not constrained by the network. In addition, if you will be creating a secondary copy using storage pool backup where the copy storage pool is not using deduplication, the data movement will take longer because of the extra processing required to reconstruct the deduplicated data.
- Client-side deduplication can outperform server-side deduplication with a high-performing Tivoli Storage Manager server configuration and a low-latency network connection between the clients and server. In addition, when combining deduplication with node replication, client-side deduplication stores data on the Tivoli Storage Manager server in a deduplicated state that is ready for immediate replication that will take advantage of the

node replication ability to conserve bandwidth by not sending data chunks that have previously been replicated.

- Client-side deduplication can place a significant load on the Tivoli Storage Manager server in cases where a large number of clients are simultaneously driving deduplication processing. The load is a result of the Tivoli Storage Manager server processing duplicate chunk queries from the clients. Server-side deduplication, however, typically has a relatively small number of identification processes running in a controlled fashion.
- Client-side deduplication cannot be combined with LAN-free data movement using the Tivoli Storage Manager for SAN feature. If you are implementing one of Tivoli Storage Manager's supported LAN-free to disk solutions, then you can still consider using server-side deduplication.

Perform deduplication at the client in combination with compression in the following circumstances:

1. Your backup network speed is at a bottleneck.
2. Increased backup durations can be tolerated, and the maximum storage savings is more important than having the fastest possible backup elapsed times.
3. The client does not typically send objects larger than 2 TB in size, or client configuration options can be used to divide large objects into smaller objects.

6.3 Data protection solution using node replication feature

Before the node replication function was available in Tivoli Storage Manager V6.3, we suggested using the electronic vaulting solution in which deduplicated data that is stored in Tivoli Storage Manager storage pools is replicated to a remote site. The server database is replicated to a remote standby server using the DB2 high availability and disaster recovery function. See the following web page ("Electronic vaulting using deduplicated remote copy storage pools"):

https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli+Storage+Manager/page/Electronic+vaulting+using+deduplicated+remote+copy+storage+pools

DB2 high availability and disaster recovery with electronic vaulting solution is still valuable. Now, with node replication and the disaster recovery functionality it brings, it is the future Tivoli Storage Manager disaster recovery solution.

Node replication aims to help simplify disaster recovery from Tivoli Storage Manager with a hot standby Tivoli Storage Manager server. In the ideal world, no requirement exists to recover Tivoli Storage Manager at a remote site, recall the copy pool volumes, and perform time-consuming restores from those tapes. Node replication allows for the backup data to be replicated to a target server, meaning that it is ready and waiting when the disaster recovery begins.

However, like all technologies, you need to be aware of certain aspects. Although there are huge benefits to node replication, it does require careful configuration to ensure that you get the best results from it.

6.3.1 Node replication configuration considerations

Node replication requires bandwidth. If you have a large amount of data to replicate on a daily basis, consider increasing the bandwidth between the two servers. We provide some guidance for calculating how much data you can realistically expect to transmit.

The configuration required to have deduplication benefit replication bandwidth is as follows:

- Both source and target storage pool are enabled for deduplication.
- Data is deduplicated prior to replication, either through client-side deduplication or by allowing identify duplicates to complete on the source server prior to replication.

If you have bandwidth constraints, consider reducing the amount of data sent across the wire. You can work on this in several ways:

- Consider transmitting deduplicated data. While the node replication process is working, after deduplicates are identified, only the deduplicated chunks of data will be replicated, resulting in potentially a huge savings.
- Identify which of your servers has the most pressing recovery point objective (RTO) and recovery time objective (RTO) at a disaster recovery. If bandwidth dictates that you can replicate only a subset of your business data, ensure that you are using that bandwidth intelligently and make sure it is backing up the most important subset. Remember this aspect: Even if bandwidth is not a concern, replicate the critical nodes first.
- Active versions of backup data generally are considered to be more important than the inactive versions. By setting up an initial replication of only active data, you can avoid the overloading of your network or servers had you replicated both active and inactive data. Configure your replication rules to replicate active data first. After the initial replication of the active data completes, configure the replication rules to replicate all versions of the data. During the next scheduled replication, any new active versions, including all inactive versions, will be replicated. The files that were active but are now inactive are not replicated again.
- If you can wait for the initial replication process to complete, finish configuring the nodes and begin or schedule a replication. To ensure that the replication process proceeds as expected, monitor the progress of the replication by issuing a **QUERY PROCESS** command. Summary information in the output includes the amount of data replicated and the process duration. Use the summary information statistics to determine whether the values of the controlled test match the actual replication values. If the values match, continue the replication. If the values do not match and you are concerned about the amount of time that replication is taking, consider one of the replication methods described in the next section (“Planning for incremental replication” on page 129).

6.3.2 Planning for incremental replication

This section describes guidelines for performing replication on a regular basis, after the initial replication has been completed.

Scheduling subsequent incremental replications

After you complete the initial replication, subsequent incremental replications must be frequent, preferably daily, to ensure that the data on the target server is maintained at an acceptable recovery point.

Daily incremental replications typically do not require as much time to complete as the initial replication. However, if the initial replication took four days to complete, and during that time you were ingesting data daily, your first scheduled replication might require more time than

you originally planned for daily incremental replications. However, successive replications will eventually catch up with the data ingested during the initial replication, and you should be able to maintain your replication schedule within the allotted window. If you determine that the replication process is not catching up, consider increasing the number of sessions that are transferring data to the target server at least temporarily until the replications are handling the data ingested each day. Then reduce the number of sessions.

When should you run daily incremental replications? Do not run replication during client backup windows because the data that is replicated is stored at the same time. Wait until after the client backup completes and run replication during the window allotted for server maintenance. You should also avoid running replication at the same time that expiration is running.

Consider replication as a peer process for storage pool backups and schedule replication accordingly. If you are replicating from deduplication enabled storage pools, run processes in the following order (waiting for each to complete before proceeding to the next):

1. Identify duplicates: This process divides files into extents, potentially reducing the amount of data to be sent to the target server when replication occurs.
2. Replication: Only file extents that do not exist on the target server are sent during replication, reducing the required bandwidth and improving performance. Replication performance improves as more deduplicated data is stored on the target server. When more extents are stored on the target server, the probability that a duplicate will be found for an extent increases.
3. Reclamation: Reclamation removes and links duplicated extents. Running reclamation after replication completes improves performance because replication does not need access additional linked extents.

For details about scheduling replication along with other server processes, see the following Technote:

http://www-01.ibm.com/support/docview.wss?uid=swg21419733

Replication mount-point usage

The **REPLICATE NODE** command allows the administrator to specify the maximum allowable number of data sessions to use for sending data to a target replication server by specifying the **MAXSESSIONS** parameter. The default value is 10 and increasing the number of sessions can improve node replication throughput.

When setting the **MAXSESSIONS** value, consider the number of logical and physical drives that can be dedicated to the replication process. To access a sequential-access volume, Tivoli Storage Manager uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on the following factors:

- Other Tivoli Storage Manager and system activity
- The mount limits of the device classes for the sequential access storage pools that are involved
- Ensure that sufficient mount points and drives are available to allow node replication processes to complete. Each replication session might need a mount point on the source and target replication servers for storage pool volumes. If the device type is not FILE, each session might also need a drive on both the source and target replication servers.
- When setting a value for **MAXSESSIONS** on the **REPLICATE NODE** command, consider the available bandwidth and the processor capacity of the source and target replication servers.

Tips:

- The value specified by the **MAXSESSIONS** parameter applies only to data sessions. Data sessions are sessions during which data is sent to a target replication server. However, if you issue a **QUERY SESSION** command, the total number of sessions might exceed the number of data sessions. The difference is due to short control sessions that are used for querying and setting up replication operations.
- Remember that the value of the **MAXSESSIONS** parameter represents the maximum allowable number of sessions and that the number of sessions that are used for replication depends on the amount of data to be replicated. If you are replicating only a small amount of data, there is no benefit in increasing the number of sessions. The total number of sessions might be less than the value that is specified by the **MAXSESSIONS** parameter.

Deduplication and replication mount-point usage

Replicating data using deduplication enabled storage pools can significantly improve the performance by not having to send extents (chunks) of data that already exist on the target server. For chunks that must be sent, the source server may have to mount several volumes to read the file if it contains chunks that are linked to other chunks on different volumes. To reduce the mounting and dismounting of volumes, deduplication enabled storage pools are allowed to keep more than one FILE volume open and mounted. The number of open and mounted FILE volumes is controlled by the NUMOPENVOLSALLOWED server option. The default value for this option is 10. A value of 10 specifies that session or process reading data from FILE volumes in a deduplication enabled storage pool can keep 10 volumes open and mounted. If more volumes are needed to obtain additional extents, the least recently-accessed volume and its mount point are released. The volume with the required extent is mounted on the just-released mount point.

To avoid replication sessions waiting for a mount point in deduplication enabled storage pools, you may need to adjust the number of mount points allowed. A low mount limit might cause replication sessions already holding mount points to needlessly wait for additional mount points. For example, suppose that the mount limit is set to 10 and that all sessions already have a metadata volume mounted and that each session needs to mount another volume. In these circumstances, the sessions will stall while waiting for a mount point that is not available. Other sessions or server processes that are using mount points in the device class can also aggravate this problem.

To avoid having replication sessions wait for a mount point in deduplication-enabled storage pools, use the following formula to set the mount limit in the device class to which the deduplication enabled storage pool is assigned:

```
NUMOPENVOLSALLOWED * MAXSESSIONS
```

The server option NUMOPENVOLSALLOWED specifies the number of input FILE volumes in a deduplicated storage pool that can be open at one time. The default value is 10. You specify the **MAXSESSIONS** parameter on the **REPLICATE NODE** command. The default value for the parameter is 10. If you use the defaults for both values, then update the mount limit in the device class to 100. This value should be sufficient to allow replication and other sessions or processes to acquire mount points. Be aware, however, that running other sessions and processes that use mount points from the same device class as the replication process will use mount points and may affect replication performance.

If your mount limit is already set to a value that is enough to satisfy the number of mount points needed for replication then no action is necessary.

6.3.3 Challenges and benefits the solution addresses

Tivoli Storage Manager replication has challenges and benefits.

With an increasing need for high availability and business continuity, the Tivoli Storage Manager replication is an effective tool for overall cost of a high availability data protection solution. Node replication can offer a huge advantage over the traditional way of recovering data during a disaster recovery, but be sure you understand that it must be configured correctly to ensure you get the results you need and expect.

If a disaster occurs and the source replication server is temporarily unavailable, client nodes can recover their data from the target replication server. If the source replication server cannot be recovered, client nodes can be converted for backup operations on the target replication server. Node replication is the process of incrementally copying or replicating client node data from one Tivoli Storage Manager server to another for the purpose of disaster recovery.

See the “Synchronization of source and target replication server after role change” technote:

http://www.ibm.com/support/docview.wss?uid=swg21628618

Your benefits

Tivoli Storage Manager node replication solution provides the following benefits:

- A backup solution with disaster recovery protection that does not involve tape
- Allows for a hot standby server at a remote location.
- By combining with deduplication, it can reduce bandwidth that is required
- Less tape media handling and decreased labor
- High availability
- Business continuity
- Allows for a many-to-one replication solution where server from multiple branch offices replicate to a common hub server

6.3.4 Solution architecture

The Tivoli Storage Manager node replication capability, which allows for an alternative architecture where deduplicated data is replicated to a second server in an incremental fashion, takes advantage of deduplication and avoids reconstructing the data.

Figure 6-3 on page 133 illustrates the simulated architecture that includes two Tivoli Storage Manager servers at separate sites and their primary storage pool hierarchies. Once a day, these two sites simultaneously replicate their data to each other over a high-speed WAN connection. Each server acts as both a backup target for clients that are local to the site, and a replication target for the Tivoli Storage Manager server from the other site. Each site has roughly the same storage hierarchy and uses the same domain names.

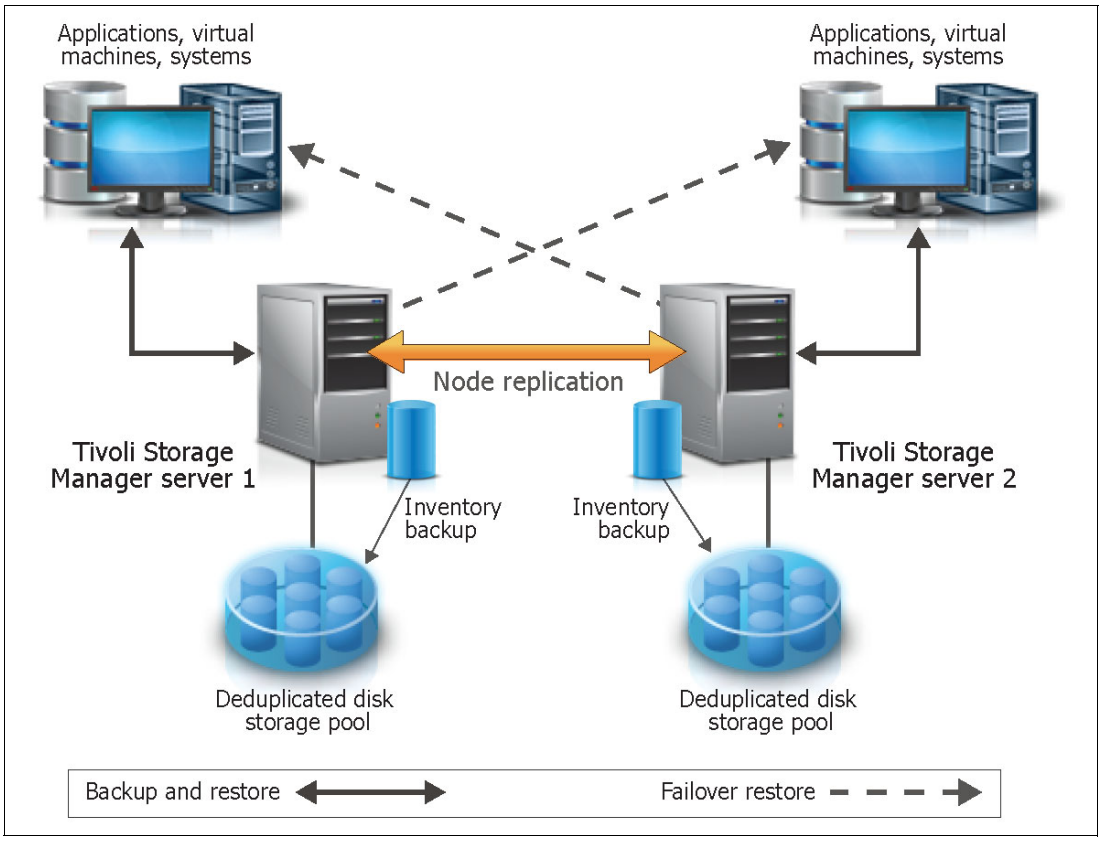

Figure 6-3 Deduplication with node replication copy

Solution description

How does node replication process work? Node replication is incrementally copying, or replicating, data that belongs to backup-archive client nodes. Data is replicated from one Tivoli Storage Manager server to another server. The server from which client node data is replicated is the *source replication server*. The server to which client node data is replicated is the *target replication server*.

Replication provides a tuning option that controls the number of simultaneous replication sessions to use. Tuning this option can increase throughput to take advantage of available bandwidth. Another option is to run more than one replication process, with each process working on a different group of nodes.

A server can function as the source of replicated data for some client nodes and as the target of replicated data for other client nodes. The purpose of replication is to maintain the same level of files on the source and the target replication servers.

As part of replication processing, client node data that was deleted from the source replication server is also deleted from the target replication server. When client node data is replicated, only the data that is not on the target replication server is copied. Replication processing involves the interaction of replication rules, states, and modes:

- Deletes data on the target server that has been deleted on the source server
- Can recover client data directly from the hot standby server (unlike virtual volumes)
- Can be used with or without data deduplication
- Can have multiple servers replicate to one target server

If a disaster occurs and the source replication server is temporarily unavailable, client nodes can recover their data from the target replication server. If the source replication server cannot be recovered, client nodes can be converted for backup operations on the target replication server. The following types of client node data can be replicated:

- Active and inactive backup data together, or only active backup data including application data stored by Tivoli Storage Manager for Data Protection products and Tivoli Storage Manager for Virtual Environments.
- Data that was migrated to a source replication server by Tivoli Storage Manager for Space Management clients.
- Archive data

Only Tivoli Storage Manager V6.3 servers and later can be used for node replication. However, data for client nodes that are running a client level V6.3 or earlier can be replicated. Data that was stored on a Tivoli Storage Manager V6.2 or earlier server before it was upgraded to V6.3 can be replicated.

Summary of key aspects of the architecture

Client backups are ingested over the LAN to one of two primary storage pool destinations. A deduplicated primary file-based disk storage pool is used for a majority of clients with a mix of both client-side and server-side deduplication. Objects are stored in this pool for their entire retention.

All new data ingested each day into the deduplicated storage pool hierarchy is replicated using Tivoli Storage Manager node replication to the other Tivoli Storage Manager server with the following details:

- Node replication occurs after duplicate identification has been completed so that bandwidth savings from deduplication is possible.
- Multiple identify duplicate data sessions are used to increase throughput.

Preferred practices are implemented that separate the activities of data ingestion and the tasks of server data maintenance are divided into distinct windows each day. This includes ordering the data maintenance tasks optimally to avoid resource contention.

Different disk storage is used for holding the Tivoli Storage Manager database and Tivoli Storage Manager storage pools. A faster disk is used for the database; a less-expensive slower disk is used for the storage pools.

The daily processing cycle we prefer for replication combined with deduplication is as follows:

1. Ingest new backup data.
2. Run identify duplicates if necessary (not already deduplicated by the clients). Note that the identify duplicates processing can be overlapped with the backup ingest.
3. Run replicate node (Tivoli Storage Manager database backup can run in parallel).
4. Expire inventory.
5. Start reclamation.

6.3.5 Use scenarios

These scenarios involve integration between node replication and deduplication:

- Using Tivoli Storage Manager for Virtual Environment with node replication and deduplication
- Using Node Replication and Deduplication for Remote Branch Offices data protection
- Using Node Replication for high availability and disaster recovery

6.3.6 Hardware and software requirements

Node replication processing requires extra memory above what is typically recommended for a Tivoli Storage Manager server V6.3 and V7. However, when properly configured, the benefit of node replication results in a significant cost reduction benefit. Running a server with insufficient memory might cause decreased server performance or other operational issues. When running node replication without deduplication, be sure to install a minimum of 64 GB of memory and four processor cores. With deduplication enabled, at least eight processor cores and 128 GB RAM are needed. If the retained capacity of backup data grows, the memory requirement might need to be as high as 192 GB. Remember to monitor memory use on a regular basis to determine if more memory is required.

6.3.7 Database requirements for node replication

Node replication requires additional Tivoli Storage Manager database space to track the files that have been replicated. Be sure your database is properly sized to accommodate the required space. To determine if your database can handle the extra space requirements, t first estimate how much the extra database space node replication will consume. Issue the **QUERY OCCUPANCY** command for each node and data type that you plan to replicate. The output from the command contains the number of files for the node and data type. Use the total number of files from all nodes and data types and multiply it by 300 (the number of additional bytes needed for each replicated file) to determine the additional database space needed. If the additional required space approaches or exceeds the size of your database you must increase the available database space. We suggest that you increase the size of your database by the additional amount needed plus an additional 10% of the current database size. Be sure to examine both replication servers and their databases and then increase the database size if necessary.

6.4 Tivoli Storage Manager together with ProtecTIER

We describe the major factors, such as scale, scope, and cost to consider in 6.2.11, “Factors when deciding between Tivoli Storage Manager and appliance deduplication” on page 121. If you decide to use both techniques in the same environment, you might need help regarding combining your Tivoli Storage Manager Data Protection Solutions together with ProtecTIER solution. *Harnessing the Power of ProtecTIER and Tivoli Storage Manager*, SG24-8209 was released this year to help:

http://www.redbooks.ibm.com/abstracts/sg248209.html

That book describes how you can combine the advanced capabilities and features of Tivoli Storage Manager with the powerful performance-enhancing and cost-reducing capabilities of the ProtecTIER product. The book covers topics such as what to consider (in the planning stage) to correctly implement an IBM ProtecTIER environment that is integrated with IBM Tivoli Storage Manager.

6.5 Tivoli Storage Manager as a virtual appliance

This section describes an architecture using *Tivoli Storage Manager Server and Tivoli Storage Manager for Virtual Environments: Data Protection for VMware*, all running within a virtual machine, in a VMware vSphere environment.

This *appliance-like* deployment of Tivoli Storage Manager Data Protection solution design makes sense to consider for a fully virtualized environment.

We show how Tivoli Storage Manager Server and data mover can run on the same machine, and elaborate on the capability of such installation. This example does not describe the highest performing environment possible; the same performance considerations still apply. Here is the link to the latest optimizing performance for servers and clients publication:

https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/Optimizing%20performance%20for%20servers%20and%20clients

Rather, it describes one example environment to serve as an aid to Tivoli Storage Manager data protection solution deployment in a fully virtualized environment.

The solution discussed here is possible only because Tivoli Storage Manager server and Backup-Archive client are both supported in a virtual environment.

As a reminder, see the following address for the list of Tivoli Storage Manager supported features in a virtual environment (IBM Tivoli Storage Manager guest support for Virtual Machines and Virtualization):

http://www-01.ibm.com/support/docview.wss?uid=swg21239546

Notice that the design here focuses on VMware as a virtualization layer, but can also be done using other virtualization providers as the web page indicates. Again, see that web page to be sure that your virtualization layer is supported.

6.5.1 Challenges the solution addresses

All challenges that Tivoli Storage Manager addresses in a physical environment can also be addressed in a virtual environment, as per the product virtualization support.

Virtualization has several features that Tivoli Storage Manager can use. So, having the data protection solution integrated to the virtual environment makes sense. All challenges that Tivoli Storage Manager addresses in a physical environment can also be addressed in a virtual environment, as per the product virtualization support.

Key points using a virtualized Tivoli Storage Manager Server and data mover

Consider these important points:

- Save hardware installation and maintenance that is required by an external backup server.
- Save LAN and SAN facilities.
- Take advantage of all the virtualization benefits such as these:
  - High availability (VMware HA)
  - Virtual machine move (VMware vMotion)
  - Clone of a virtual machine (ease of deployment, standardization and scalability)

Summary of benefits

Here is a summary of the benefits with a virtual Tivoli Storage Manager Server, in conjunction with Data Protection for VMware:

- Decrease the rate of amount of storage capacity required by using Tivoli Storage Manager deduplication, to contain data growth
- Network bandwidth optimization (incremental, data deduplication)
- Less tape media handling and decreased labor
- Ability to be shared and deduplicate across more than one Tivoli Storage Manager server

6.5.2 Solution architecture

Figure 6-4 shows an installation overview of a virtual Tivoli Storage Manager Server and Data Protection for VMware data mover installed in a unique virtual machine.

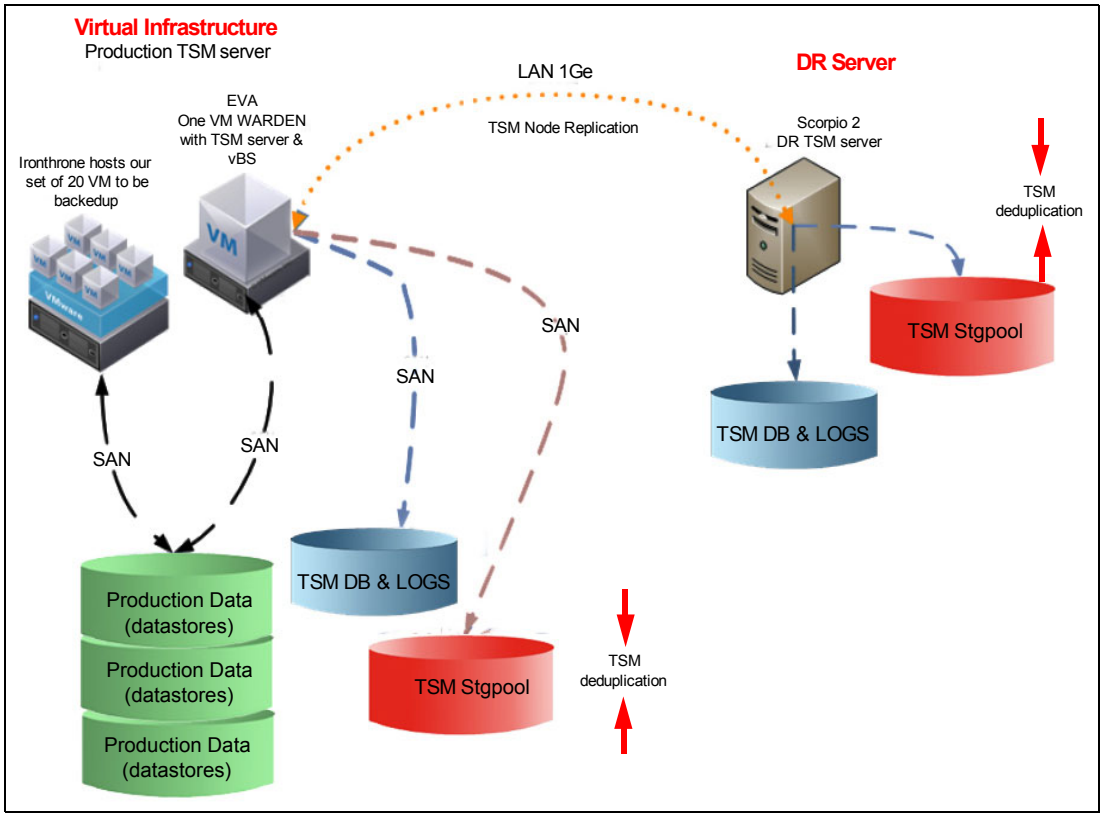

Figure 6-4 Virtualized Tivoli Storage Manager Server and datamover

Notice in this example the possibility to interconnect the virtualized Tivoli Storage Manager server to an external server, using node replication, thus to send the data from the virtualized infrastructure, for disaster recovery purposes, for instance.

To learn more about the IBM Tivoli Storage Manager server and datamover running in a Virtual Environment Sample Solution, see the following document:

http://ibm.co/1kZWbpZ

This document provides a sample architecture describing a consolidation of the Tivoli Storage Manager components such as Tivoli Storage Manager server and Tivoli Storage Manager for Virtual Environments on one virtual machine.

6.5.3 Solution description

The combination of the Tivoli Storage Manager scalability factors and the virtualization benefits allow you to easily increase the capability of such installation.

The virtualization allows you to easily deploy new Tivoli Storage Manager Server and Data Protection for VMware data mover to protect more virtual machines.

By moving the virtual machine that holds the Tivoli Storage Manager server and Data Protection for VMware data mover, you can distribute the backup load and optimize the data transfer within the virtual machine, ESXi hosts and network components. In addition, you can easily move the Tivoli Storage Manager server database volumes, storage pool volumes, to another data store more evenly distribute the load generated by virtual machine data protection.

Tivoli Storage Manager deduplication technology can be also enabled to reduce the amount of data transferred, however it generates an extra load on the ESXi host that is hosting the virtual machine where the Tivoli Storage Manager server and Data Protection for VMware data mover is running. When enabling deduplication, see the following web page to understand the required resources:

https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/Deduplication

Scope

This solution applies to every VMware based virtualization infrastructure.

The number of resources that are required by Tivoli Storage Manager Server and data mover must be calculated to fit the environment.

Cost

When running Tivoli Storage Manager Server and Data Protection for VMware within a virtual machine, no dedicated or separate hardware is required, therefore potential savings can be obtained by avoiding the investment on dedicated and separate hardware.

As you would do in any other Tivoli Storage Manager installation, enabling the data deduplication can save space in storage pools, thus reducing the license costs.

6.5.4 Use scenarios

This use case consists of a single Tivoli Storage Manager Server and Data Protection for VMware data mover agent running within a unique virtual machine to protect several virtual machines within a VMware data center. In addition, backed up data is sent to an external Tivoli Storage Manager Server using node replication, for disaster recovery purposes.

The number of virtual machines that can be protected with this type of solution depends on the virtual machine disk size and the daily changed data rate, which gives the amount of data to be ingested daily by the Tivoli Storage Manager server.

The following list summarizes key aspects of the possible architecture:

- A virtual machine (VM) acts as the data mover component of the vStorage backup server. This capability is implemented in the Tivoli Storage Manager backup-archive client portion of the Data Protection for VMware solution.
- Data is read from VMware data stores across either the LAN or SAN by the vStorage backup server using the VMware Virtual Disk Development Kit (VDDK) for network block device (NBD) or HotAdd transport.
- The VM backups are ingested over either the internal TPC/IP stack or shared memory to the Tivoli Storage Manager server using client-side deduplication and compression, because the data mover agent is located on the same machine as the Tivoli Storage Manager server.
- The Tivoli Storage Manager for Virtual Environments incremental-forever backup capability is used to keep the amount of data process to a minimum before other data reduction technologies are applied.
- The preferred practices implementation of separating the activities of data ingestion and server data maintenance tasks into two distinct windows each day is used on the Tivoli Storage Manager server. This includes ordering the data maintenance tasks optimally to avoid resource contention.
- IBM or any other VMware supported storage can be used for the Tivoli Storage Manager server database and storage pool volumes. These volumes can be assigned as either standard virtual disk or raw device mapping to virtual machine hosting the Tivoli Storage Manager server.

6.5.5 Hardware and software requirements

See Tivoli Storage Manager products virtualization support statements at the “IBM Tivoli Storage Manager guest support for Virtual Machines and Virtualization” web page:

http://www.ibm.com/support/docview.wss?uid=swg21239546

See the Tivoli Storage Manager wiki for documentation about data deduplication:

https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/Deduplication

7

Protecting your data with Tivoli Storage Manager

In this chapter, we provide various data protection solutions based on common challenges in today's data protection world. The solutions are based on the components available in the Tivoli Storage Manager toolkit and reflect the information in the challenges matrix, introduced in Chapter 4, "Tivoli Storage Manager challenge matrix" on page 81.

Systems virtualization is one of the most important transformations today. Tivoli Storage Manager understands this and offers a range of suitable products to assist you in this change.

This chapter explains the various ways of protecting data in a virtualized environment:

- Using in-guest Tivoli Storage Manager agents as you would do for any physical machine
- Using off-host Tivoli Storage Manager agents such as Tivoli Storage Manager backup-archive client and its native "full-vm" backup feature
- Using off-load backups (leverage hardware snapshot technology) such as Tivoli Storage FlashCopy Manager

7.1 Common virtualization challenges

The components that bring value with virtualization become challenges for data protection, which is basically a resource-intensive consumer task.

Regardless of the virtualization platform, and understanding the hypervisor platform, the same technical challenges arise regarding data protection.

- Virtual resource mobility
- Hardware resource sharing

In addition to those technical challenges, business challenges must be addressed too.

- Decrease the amount of data processed for data protection.
- Improve the backup time.
- Improve the restore time.
- Clustered awareness.
- Support of virtual machines with direct attached disks.
- Disaster recovery in a virtual environment.
- Reduce time spent deploying and managing backup agents.

Figure 7-1 on page 143 shows the features that Tivoli Storage Manager provides to protect your data in a virtualized environment, while answering these common challenges.

Challenges										
business					addressed with	technical				
reduce recovery time (RTO)	reduce backup time (meet RPO)	Compliance	Business Continuity	security (secure data transfer, secure data store,	Solution (discussed in this book)	Virtual Environments	Significant data growth	automation for business	Remote office and mobile office data	network bandwidth optimization
X	X				Disk to Disk backup	X		X	X	
X	X				Data reduction with TSM	X	X		X	X
X			X		Node replication				X	
X	X	X	X	X	TSM ToolKit	X	X	X	X	X
X	X		X		Hardware Snapshot technology	X	X			
	X				progressive incremental forever	X			X	X
X	X				deduplication	X	X		X	X
	X		X		off-load backup	X	X			
	X				Off-Host backup	X	X			
X			X		instant restore	X	X			
	X				include / exclude	X	X			X
X	X		X		High Availability	X				
	X				journal based backup		X			X
X	X		X		image backup	X	X			
X					point-in-time restore	X				
x			X		Server maintenance					
X			X		Disaster Recovery Manager	X		X		
X	X				Tiered Storage Hierarchy	X	X			

Figure 7-1 Matrix represents the Tivoli Storage Manager tools to protect data within a virtualized environment

This subset of Tivoli Storage Manager tools can be used when protecting the virtual environment. However, there are differences depending of the hypervisor distributor.

We describe Tivoli Storage Manager toolkit benefits throughout the following sections:

- 7.2, “Virtualization: VMware Data Protection solution” on page 144
- 7.3, “Virtualization: Hyper-V host-based backup” on page 158
- 7.4, “Virtualization: In-guest backup” on page 167

7.2 Virtualization: VMware Data Protection solution

There are various ways to protect the VMware infrastructure, to give you a broader view of what is available within the Tivoli Storage Manager family. In this section, we look at the following available features you can use:

- Tivoli Storage Manager backup-archive client in VMware guest virtual machine (in-guest)

  This option consists of protecting the data in the same way as for a physical machine. It means that you can perform either backup (for example, incremental, selective) or archive within the virtual machine.

- Tivoli Storage Manager backup-archive client off-host backup

  This option consists of making a virtual machine snapshot by a third-party machine, either physical or virtual. This machine is usually named “Backup Proxy.”

  The snapshot capability of Tivoli Storage Manager backup-archive client has been available since version 6.2. It allows the Tivoli Storage Manager backup-archive client to leverage the VMware vStorage API for Data Protection (VADP) to back up the entire virtual machine at block level. This option provides a fast, centralized, off-hosted and scalable backup methodology. Recovery is a single step procedure that re-creates the virtual machine (VM) directly on the VMware ESXi host.

  For a Windows virtual machine, you can perform the file-level backup using the backup-archive client native full-vm feature. By specifying the parameter **`VMbackuptype to file`** in the backup-archive client option file, a file level incremental backup will be performed. This allows you to perform file level recovery without installing a backup-archive client in guest virtual machine.

  If you need to perform file restoration, consider installing Tivoli Storage Manager for Virtual Environment.

  > **Note:** The file-level backup is not supported for Linux virtual machines. Use the in-guest Tivoli Storage Manager agent in addition to the full-vm backup method to have the same level of protection.

- Tivoli Storage Manager for Virtual Environments Data Protection for VMware

  This option is an upgrade of the Tivoli Storage Manager backup-archive client off host backup option described previously. When installing Tivoli Storage Manager for Virtual Environments Data Protection for VMware, an enablement file upgrades the backup-archive client, adding incremental and incremental forever backups capability. It also brings a set of recovery features and a vSphere client plug-in.

  Tivoli Storage Manager for Virtual Environments Data protection for VMware brings recovery agents that allow extracting data to be restored from a unique backup, whatever the type of restore operation. That is, even if the backup is a block level backup, by using the Data Protection for VMware recovery agent you can restore files, volumes, and machines. These operations can be done for both Linux and Windows virtual machines.

- Tivoli Storage FlashCopy Manager for VMware

  This option consists of making the snapshot at the data store level, using Tivoli Storage FlashCopy Manager for VMware. This solution is based on hardware snapshot available at the back-end disk subsystem.

  The recovery possibilities are virtual disk, virtual machine, and data store. File level restores are also provided by attaching a VDisk from a backup to an existing VM.

Tivoli Storage FlashCopy Manager for VMware brings its VMware vSphere plug-in, as Tivoli Storage Manager for Virtual Environments does.

Tivoli Storage FlashCopy Manager for VMware can also be combined with Tivoli Storage Manager for Virtual Environment so the FlashCopy snapshots can be sent onto a Tivoli Storage Manager server. When these products are combined, they share a common and integrated vSphere client plug-in GUI.

- Tivoli Storage FlashCopy Manager within VMware guest

  This option consists of using the FlashCopy Manager within the virtual machine. It allows you to get rid of VMware snapshot limitations (pRDM for instance), providing a powerful backup solution based on hardware snapshots.

Note: You must know your recovery requirements before choosing your backup method. The restore function is strongly influenced by initially choosing the right Tivoli Storage Manager backup option.

7.2.1 Sample solution and benefits

For the purpose of this sample solution, we intentionally combine all these products. You can implement one or more of them, depending on your needs.

For our scenario, we assume the following environment:

- Virtual machines (without applications)
- Virtual machine hosting a Windows large file server
- Virtual machines with a custom applications
- Virtual machines with Active Directory
- Virtual machines with MS-SQL server
- Virtual machines with Exchange Server Cluster, with physical raw device mapping (pRDM)
- Virtual machine with SAP DB2 database (Linux with pRDM disks)
- Virtual machine with Oracle database

Here are the identified recovery requests for our scenario environment:

- Ability to recover single object from the Active directory database
- Ability to recover virtual machine with consistent MS-SQL database
- Ability to recover the SAP DB2 database and Oracle database at a specific point in time
- Ability to recover mailboxes from Exchange Server
- Ability to recover files from file server
- Build a solution that can be part of a disaster recovery scenario

To address these requirements, this solution implements these products:

- Tivoli Storage Manager Backup-Archive Client
- Tivoli Storage Manager for Virtual Environment
- Tivoli Storage FlashCopy Manager for VMware
- Tivoli Storage FlashCopy Manager
- Tivoli Data Protection for Application is a general term that is used in the sample solution and includes these items:
  - Tivoli Storage Manager for Databases
  - Tivoli Storage Manager for Mail
  - Tivoli Storage Manager for Enterprise Resource Planning

Using those Tivoli Storage Manager products, the client benefits are as follows:

- Decrease the amount of data processed for data protection
- Improve the backup time
- Improve the restore time
- Virtual machine clustered system awareness
- Support virtual machine with pRDM disk type
- Facilitate a disaster recovery
- Reduce time spent deploying and managing Tivoli Storage Manager agent

7.2.2 Solution architecture

The sample solution depicted in Figure 7-2 on page 147 takes advantage of the diversity that the Tivoli Storage Manager family of products offers, each providing complementary capabilities. It includes Tivoli Storage Manager Server, Data Protection for Applications, Tivoli Storage FlashCopy Manager for VMware, Tivoli Storage FlashCopy Manager, Tivoli Storage Manager for Virtual Environments, and Tivoli Storage Manager Backup-archive client.

Note: This sample solution depicts only virtual resources because we are talking about virtualization. In some cases, you might use a physical machine such as vStorage Backup Server, backup proxy, or proxied Storage Agent. Then, add Tivoli Storage Manager for SAN to products list, it allows you to transfer data from the back-end disks to Tivoli Storage Manager storage hierarchy using LAN-free transport method. See “LAN-free data movement” on page 41 for more information.

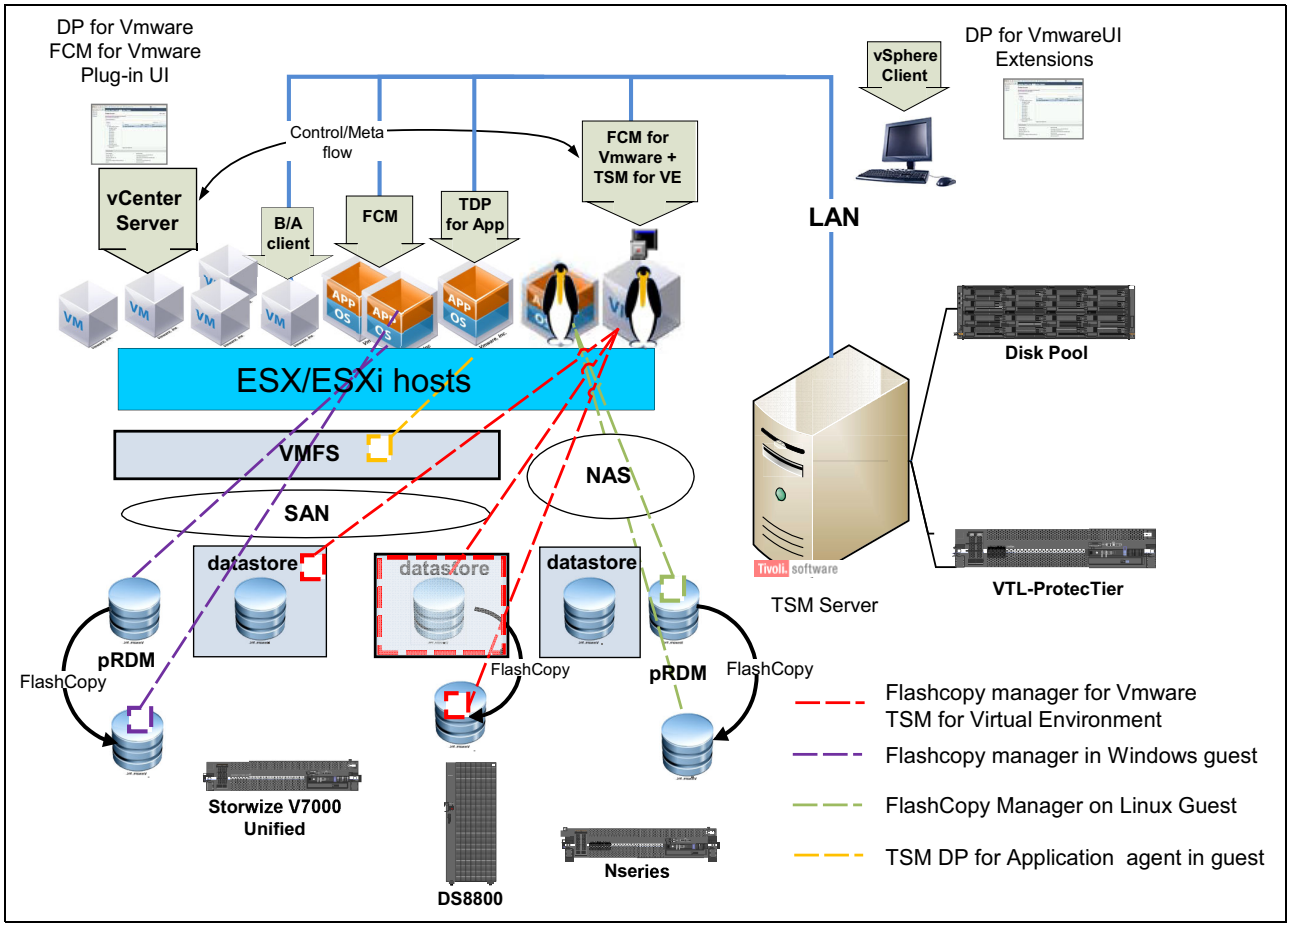

Figure 7-2 Sample architecture of Tivoli Storage Manager products protecting VMware environment

Back-end disk storage

Figure 7-2 shows various types of disk subsystems: DS8000 series, N series, and Storwize V7000 Unified. Those disk subsystems represent what might be used for a VMware environment. It is an opportunity to show that Tivoli Storage Manager products are both SAN and NAS devices compatible. The variety of back-end disk storage allows us to demonstrate the integration of hardware snapshots function into the Tivoli Storage Manager products like for Tivoli Storage FlashCopy Manager, and Tivoli Storage FlashCopy Manager for VMware throughout this scenario.

For more details about FlashCopy and hardware support, see “Snapshot at the storage hardware layer” on page 49.

VMware virtualization layer

In this section we provide a brief introduction to the VMware components.

Virtual machine (VM)

Virtual machine is where the operating system and eventually applications runs. Virtual machines can use two concepts for their virtual hard drives: virtual disks and raw device mappings (RDM).

VMware vSphere

VMware vSphere is a powerful server virtualization solution for x86 platforms. It aggregates server virtualization, network, storage and high availability services, helping you to run an optimized IT environment along with simplified administration and management. VMware vSphere consists of several components. The major component is the ESXi server, a bare metal hypervisor that runs on x86 servers.

ESXi

ESXi provides a virtualization platform for a large variety of open systems such as Windows and Linux. ESXi uses a special file system to store virtual machines called Virtual Machine File System (VMFS). VMFS is a clustered file system that can be accessed concurrently by up to 32 ESXi hosts. A VMFS partition can be spread over 32 logical volumes called *extents*. It can have a maximum size of 64 TB; the maximum file size on a VMFS partition is 2 TB.

On a VMFS partition, ESXi stores all the files of which a virtual machine consists:

- Virtual machine configuration file (`.vmx`)
- Virtual machine BIOS file (`.nvram`)
- Virtual disk files or raw device mappings (`.vmdk`)
- Virtual machine log files (`.log`)

VMware data store

VMware data store represents a storage location for virtual machine files. A data store can be a VMFS volume, a directory on network-attached storage, or a local file system path.

A data store is platform-independent and host-independent. Therefore, data stores do not change when the virtual machines they contain are moved between hosts.

The scope of a data store is a data center. The data store is uniquely named within the data center. It is advisable to have a one-to-one relationship between LUNs and data stores. One LUN should contain only one data store. One data store should not be spread over more than one LUN.

vSphere management

For the management of vSphere, VMware provides two components:

- *vCenter server* is a management suite to manage multiple ESXi hosts and to provide advanced functions (such as on-line migrations, cloning, and high availability).
- *vSphere client* is Windows software to connect to either single ESXi hosts or vCenter Server that provides a graphical user interface.

Tivoli Storage Manager backup-archive client is able to protect the entire virtual machine using vStorage API for Data Protection (VADP). Depending of how the virtual machine is set up, you might need to use an alternate solution to protect the data that resides on this virtual machine.

For instance, when using pRDM or bus sharing, the snapshot can no longer be taken, so you need to use another data protection solution. pRDM and bus sharing can be required when configuring clustered operating systems or clustered application within virtual machines.

Tivoli Storage Manager server

The Tivoli Storage Manager server allows storing the data coming from Tivoli Storage Manager agent, whatever the location of this agent. Therefore, with a set of policies selected wisely to meet the requirements, we can keep several backups for future recovery needs, and even duplicate that data to a remote site, using either copy storage pool operation or node replication.

Because all Tivoli Storage Manager agents used in this example are deduplication-aware, see 3.3.3, “Server-side data deduplication” on page 57 to understand how the deduplication works on both Tivoli Storage Manager client and Tivoli Storage Manager server side.

For more information about Tivoli Storage Manager server protection, see 5.1, “Protecting the Tivoli Storage Manager server infrastructure” on page 94.

Tivoli Storage Manager server storage

Tivoli Storage Manager server provides the storage pools required for the client to be able to store the data when backing up. In this example, we consider disk and virtual tape library (VTL ProtecTier).

The disk storage pool is used to store the control files (CTL) sent by Tivoli Storage Manager for Virtual Environment backups. Its size is roughly estimated to 1% of the amount of VMware data to be protected.

The virtual tape library contains all the backups from the VMware data protection agents and in-guest agents. The ProtecTier brings inline deduplication capabilities, thus allowing us to save storage space without any space overhead. It also provides fast access to virtual machine data blocks in case of file level restore, instant volume restore, and virtual machine restore.

In this scenario, we don’t describe the possibility to replicate the data using node replication. See 6.3, “Data protection solution using node replication feature” on page 128 for a detailed description on how to implement this function.

Tivoli Storage Manager for Virtual Environment, because of its backup format (block based) and strategy (single pass backup, multiple recovery level), requires to store its data on fast access media to enable some of its recovery features. Table 7-1 shows the features available depending on where the data resides. Tivoli Storage Manager for Virtual Environments Data Protection for VMware supports tape and virtual tape library (VTL) devices with the restrictions listed in this document:

http://www.ibm.com/support/docview.wss?uid=swg27021081

Table 7-1 Feature availability based on vStorage backup server and storage type

Storage and features	Disk or file	Virtual tape	Physical tape
LAN-free (supported only for vStorage backup physical server)	Yes (file)	Yes	Yes
Tivoli Storage Manager deduplication	Yes (file)	No	No
FULL VM Backup	Yes	Yes	Yes
INCR VM Backup	Yes	Yes	Yes[1]
FULL VM Restore	Yes	Yes	Yes

Storage and features	Disk or file	Virtual tape	Physical tape
File level restore	Yes	Yes	Yes (lower performance expected)
Instant Volume restore	Yes	Yes	No

1. The incremental VM backup process needs several CTL files to be recalled to determine the differences between previous and current status. This can slow down the backup if CTL files are stored on physical tape. As a suggested practice, always store CTL files on non sequential volume (for example, disk, file). See the tape configuration guidelines document for more information: http://ibm.co/1y6DHa5

Tivoli Storage Manager data protection agents

Here are the Tivoli Storage Manager agents that will be used to protect VMware data in our scenario. Their role and capabilities are described here.

Tivoli Storage Manager Backup-Archive Client

Tivoli Storage Manager Backup-Archive Client allows you to perform a file level backup of the Virtual Machine. It is the base of the other components listed in this section as well, typically as the data mover component. It also provides specific object protection like Windows System State therefore Active Directory objects. It is installed in the virtual machine when Data Protection for Applications agent is to be used.

Tivoli Storage Manager for Virtual Environments: Data Protection for VMware

Tivoli Storage Manager for Virtual Environments Data Protection for VMware provides protection of virtual machines, to provide a crash consistent recovery of the virtual machine whatever the virtual machine content. It provides the full or incremental backup feature, at a block level, based on Changed Block Tracking (CBT) feature that provides VMware through VMware API Data Protection (VADP).

The Microsoft SQL Server and Microsoft Exchange hosted applications can be protected with Tivoli Storage Manager for Virtual Environment. The feature to protect those two applications natively in Tivoli Storage Manager for Virtual Environment is named self-contained application protection. It enables the transaction log management without any in-guest agent installation. The machine, which runs the full virtual machine backup, injects a code to be executed within the virtual machine to clean up the logs after backup completes successfully.

The self-contained application protection must be carefully used, based on the recovery granularity you want to have. If you need to recover to a point in time through transaction logs (for Microsoft SQL Server), individual database (for Microsoft SQL Server), or individual mailboxes (Exchange), plan to use the in-guest data protection agent. The same advice applies when your applications are clustered.

Tivoli Data Protection for applications

Tivoli Data Protection for applications installed within the virtual machine offers the same level of protection as if you run it on a physical machine. One advantage to use it in-guest is the restore granularity it provides. If you need point-in-time, rollforward, or single object restore, from an application perspective, then plan to install it.

There is another case where you can install it, typically when the virtual machine has pRDM attached. These disks are skipped by Tivoli Storage Manager for Virtual Environment full-virtual machine backups.

Tivoli Storage FlashCopy Manager

Tivoli Storage FlashCopy Manager is the alternative of Tivoli Data Protection for Applications within the virtual machine. This component fits when want to leverage the hardware snapshot capabilities to protect in-guest applications, providing the ability to keep the snapshots on disk for readily instant recovery. You can offload those snapshots to a Tivoli Storage Manager server as well.

Tivoli Storage FlashCopy Manager for VMware

Tivoli Storage FlashCopy Manager for VMware is a data management solution that you can use to streamline storage management in a VMware vSphere environment. This application can back up VMware environments from Linux-based backup servers by integrating with VMware vSphere APIs and hardware snapshot mechanisms.

Tivoli Storage FlashCopy Manager for VMware optionally integrates with Tivoli Storage Manager for Virtual Environments to store VMware image backups on Tivoli Storage Manager server storage. With Tivoli Storage FlashCopy Manager for VMware, you can create off-host backups for VMware virtual machines in a vSphere environment. Backups are generated at the VMware data store level and restores can be done at the data store, virtual machine, virtual disk or file level. A command-line interface (CLI), the Data Protection for VMware command-line interface, and a vCenter GUI plug-in, the Data Protection for VMware vCenter plug-in, are provided. Tivoli Storage FlashCopy Manager for VMware protects the virtual infrastructure through automated data protection and recovery of your virtual machines. It offers an easy-to-use interface to help you manage backup and recovery of virtual machines in a multiple VMware ESX server environment. With Tivoli Storage FlashCopy Manager for VMware, you can create off-host storage hardware snapshot backups from VMware VMs. The following features are provided by Tivoli Storage FlashCopy Manager for VMware:

- Backup, restore, and disaster recovery operations for virtual machines are streamlined and simplified.
- File system consistent backups are provided and the backup window of the virtual machine is reduced by using hardware snapshots of complete data stores in combination with offloaded backups to Tivoli Storage Manager.

The combination of Tivoli FlashCopy Manager for VMware and Tivoli Storage Manager for Virtual Environment fits when planning for a disaster recovery scenario.

7.2.3 Solution description

In this section we explain how Tivoli Storage Manager data protection agents address each challenge in 7.1, “Common virtualization challenges” on page 142.

For the challenges, we describe how the data is backed up, how you can recover, and explain their interaction with the disaster recovery plan when required.

To protect the virtual machines without applications or recovery requirements, we trigger snapshots using Tivoli Storage Manager for Virtual Environments, Data Protection for VMware, running from a vStorage Backup server. It allows us to protect these virtual machines using block level progressive Incremental backup using the vStorage API for Data Protection. This provides an easy, fast and space efficient way to protect the virtual machine, keeping the possibility to recover either file, volumes or entire virtual machine from a single pass backup.

We apply the same strategy to protect the virtual machine hosting the File Server. This allow us to save backup time, the amount of data to be processed, and Virtual machine (ESXi hosts) computing resources as compared to a traditional in-guest backup-archive client

backup. Using the Data Protection for VMware Recovery Agent, we can restore a single file or directory. In case of severe damage on a volume of this file server, we will take advantage of the instant restore volume restore.

Tivoli Storage Manager for Virtual Environments Data Protection for VMware is also the solution we use to protect the virtual machines with custom applications. In addition, we use a custom pre-freeze script to quiesce the I/O before the snapshot occurs. This ensures the application consistency. We take advantage of the instant volume recovery feature from the Data Protection for VMware recovery agent, and in the meantime provide the ability restore the entire machine in case of disaster recovery. For such a custom application, having the entire virtual machine snapshot makes sense so that you keep all possible (and most likely poorly identified) dependencies between the operating system and this custom application.

For a virtual machine that runs Microsoft Active Directory, we use Tivoli Storage Manager backup-archive client because it allows us to recover individual active directory objects from accidental corruption or deletion. This operation is then online, without impacting any other Active Directory activities. To learn more about restoring active directory object, see the "Restore Windows individual Active Directory objects" topic in the *IBM Tivoli Storage Manager for Windows Backup-Archive Clients Installation and User's Guide*.

In our scenario, based on the assumption saying that Microsoft SQL Server recovery includes the entire virtual machine restore, we use Tivoli Storage Manager for Virtual Environments Data Protection for VMware snapshots, with the self-contained application protection feature enabled. We assume in this case that running version are among the supported ones: Microsoft SQL Server 2008, Microsoft SQL Server 2008 R2, Microsoft SQL Server 2012. Self-contained application protection has a mechanism to truncate the logs once the backup ends, therefore the application never runs out of log space. This avoids installing and maintaining in guest data protection for the application agent. However, in the recovery of individual databases, log files are not possible without restoring the entire virtual machine.

You might have a recovery strategy stating that you need to be able to restore partial Microsoft SQL databases, or you must provide a point-in-time recovery capability. If so, you might choose to implement Tivoli Data Protection for Microsoft SQL server instead. If you want to recover individual Microsoft SQL databases from a VM backup, you must use both Data Protection for VMware and the IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server V7.1.

These statements apply to Microsoft Exchange. See *IBM Tivoli Storage Manager for Virtual Environments: Data Protection for VMware User's Guide* for more details about self-contained application protection:

http://www.ibm.com/support/knowledgecenter/api/content/SSGSG7_7.1.0/com.ibm.itsm.ve.doc/b_ve_user_guide.pdf

Next, we discuss Microsoft Exchange Cluster defined over multiple virtual machines. In our case these virtual machines have physical raw device mapping (pRDM) disks. To protect this cluster, we implement Tivoli Storage FlashCopy Manager within the virtual machines. Because we know that we have a disaster recovery strategy, we set up Tivoli Storage FlashCopy Manager with the Tivoli Storage Manager support instead of stand-alone mode. Indeed stand-alone mode creates VSS-based snapshots and manages them on the server that is performing the backup. The VSS backup, using Microsoft Volume Shadow Copy Service, produces an online snapshot (point-in-time consistent copy) of Exchange. The Tivoli Storage Manager support mode means that backups can be stored either locally or on the Tivoli Storage Manager server. The Tivoli Storage FlashCopy Manager software protects and manages Exchange by using a Tivoli Storage Manager server. With Tivoli Storage Manager, you can also offload your backups to another computer and move the data to the Tivoli Storage Manager server.

Notice that the same FlashCopy statements apply for MS-SQL, custom applications, or file systems running on Windows platform.

By installing and using Tivoli Storage FlashCopy Manager in the virtual machines to protect our Microsoft Exchange cluster, we are able to protect the data stored on pRDM disks, which are not supported by VMware snapshot operations. We also meet the mailbox's recovery requirement. Even better, we can restore individual mail using the mailbox restore browser utility, this either using local snapshot copy (from FlashCopy) or using offloaded copies stored on the Tivoli Storage Manager server.

The Linux virtual machine hosting the SAP DB2 database uses IBM System Storage N series disks, which are assigned to ESXi hosts as Fibre Channel LUNs. These LUNs are dedicated to the Linux in a pRDM mode to fulfill performance requirements. We use Tivoli Storage FlashCopy Manager for the SAP DB2 database protection. DB2 backups to Tivoli Storage Manager can be done with either Tivoli Storage Manager for Enterprise Resource Planning environment or the Tivoli Storage Manager agent that is included with DB2.

Tivoli Storage FlashCopy Manager either initiates a tape backup from the snapshot target set when the snapshot completes or backs up a previously generated snapshot. We use the second mode to offload snapshots onto Tivoli Storage Manager server, on a regular basis for the disaster recovery purpose, relying on IBM System Storage N series to restore SAP DB2 database from snapshot when needed. If an older backup is needed, we restore the backup from the Tivoli Storage Manager server.

To protect the Oracle database, we use Oracle Recovery Manager (RMAN); Data Protection for Oracle provides an interface between Oracle Media Management API calls and Tivoli Storage Manager API routines. RMAN provides consistent and secure backup, restore, and recovery performance for Oracle databases. While the Oracle RMAN initiates a backup or restore, Data Protection for Oracle acts as the interface to the Tivoli Storage Manager server. The Tivoli Storage Manager server then applies our storage management policies to the data. Data Protection for Oracle implements the Oracle defined Media Management application program interface (SBTAPI) 2.0.

Thereby we meet the restore requirements and send the data out of the production environment to the Tivoli Storage Manager server, so we meet the disaster recovery objectives.

On top of all these backup methodologies we implement Tivoli Storage FlashCopy Manager for VMware, to protect all the data that resides on VMware data stores in an easy and efficient way, to optimize as much as possible the Disaster Recovery. Tivoli Storage FlashCopy Manager for VMware allows backing up all data stored on the VMware data stores, and excludes those on pRDM disks (toleration mode only, explicitly excluded). This allows a fast and easy recovery.

Using Tivoli Storage FlashCopy Manager backup methodology, restoring all virtual machines with their configuration and layout is streamlined, except for those with pRDM.

For the pRDM content, because we are using Tivoli Storage FlashCopy Manager in guest as a data protection solution, we need to separately restore that data from the hardware snapshots, either from the back-end disks or from the Tivoli Storage Manager server because it must be offloaded to a Tivoli Storage Manager managed storage.

On the Tivoli Storage Manager server side, all data from our backup strategy is duplicated using Tivoli Storage Manager Server node replication to another Tivoli Storage Manager server, to make them available in case of disaster. In addition, the data is also available for any purpose other than disaster recovery. For more information about node replication, see 6.3, “Data protection solution using node replication feature” on page 128.

7.2.4 Use scenarios

This scenario examines how all the components interrelate and the backup strategy that is associated to each of them.

Backup

In these scenarios, we examine various components and data protection methods that fit to the data type to be protected. Data protection methods hereafter consist of backup, recovery, and backup frequency (controlled by schedules).

- Virtual Environment backups

  Incremental forever backup uses Tivoli Storage Manager grouping technology to create synthetic-full recovery points by combining required blocks from previous backups with the changes from daily incremental backups. An incremental forever backup strategy minimizes backup windows while providing faster recovery of your data. Data Protection for VMware provides a backup strategy called incremental forever. Rather than scheduling weekly (periodic) full backups, this backup solution requires only one initial full backup. After, an ongoing (forever) sequence of incremental backups occurs.

  The filtering options available (**`exclude.vmdisk`**) are used to exclude the virtual machine or virtual machine disks that Tivoli Storage Manager for virtual environment is not supposed to protect, because they are protected by another in-guest solution. We also use the performance options (**`vmmaxparallel`**, **`vmlimitperdatastore`**, and **`vmlimitperhost`**) to decrease the backup window. To save network bandwidth, we enable both compression and deduplication on Tivoli Storage Manager client side.

- Virtual machine running Microsoft Active Directory

  Active directory objects are part of the systemstate backup. Every day we run an incremental backup of the machine, including the systemstate.

- DB2 redologs and datafiles backup

  Tivoli Storage Manager agent will be installed in the virtual machine to allow DB2 database to send its DB2 logs by using the Tivoli Storage Manager Backup-Archive Client API. Those logs will be sent directly on the Tivoli Storage Manager server.

  Using Tivoli Storage FlashCopy Manager with IBM System Storage N series and NetApp storage system, we protect the DB2 databases.

  Because the IBM System Storage N series and NetApp storage systems can be attached to hosts using SAN, the following tasks can be completed:

  - Use Tivoli Storage FlashCopy Manager to complete snapshots.
  - Offload the Tivoli Storage Manager backup of the FlashCopy data to an auxiliary host.
  - Use the FlashCopy data that backs up DB2, SAP with DB2, to restore data.
  - Create database clones.

  We set up policies within the IBM System Storage N series storage systems that delete snapshots created with Tivoli Storage FlashCopy Manager. For this purpose, Tivoli Storage FlashCopy Manager periodically checks whether backups on the storage system remain valid. This checking process is referred to as *reconciliation*.

- Exchange backup

  In our case the Exchange server databases are in a Database Availability Group (DAG) environment, and we back up the databases to a common node, therefore we set up a Database Availability Group node name (DAGNODE).

  For Microsoft Exchange Server databases in a Database Availability Group (DAG) environment, several online copies of a database are maintained for high availability. To reduce the number of database backups that are made, we set up to back up the copies of

a database from different Database Availability Group members under a single Database Availability Group node. DAG is available for Exchange Server 2010 and 2013.

For more information about Data Protection for Exchange see 2.2.1, “Tivoli Storage Manager for Mail” on page 20.

- Oracle archive logs and datafile backups

  According to customer’s service level objectives, the Oracle database must remain online all the time, which is why we use RMAN to perform online database backups. Therefore, to maintain the consistency and allow the point-in-time recovery, we need to protect the Oracle application logs and the database files. Using Data Protection for Oracle, we use Oracle Recovery Manager (RMAN) tools and backup models, such as Level0 (full backup), Level1 (incremental or differential backup), and archive log backup.

Restore

The restore scenarios depend on how the backup is done.

Using Tivoli Storage Manager for Virtual Environment, we can recover data from the full virtual machine backup and, when necessary, extract either files or an entire volume. For the entire volume restore we can use the instant restore volume method if we need to bring up the service as fast as possible. In that case, the volume is accessible for everyone after a few minutes, while a background process runs to restore all the content. Or, we can restore only the specific disk by using the **`dsmc restore vm`** command, and the following option:

```
:vmdk=disklabel
```

Depending of the task you want to initiate, you can chose among vStorage Backup Server, virtual machine, or the Data Protection for VMware vCenter plug-in as Table 7-2 describes.

Table 7-2 Availability of Data Protection for VMware operation depending of the location

Location task	vStorage Backup Server	Virtual machine	vCenter plug-in
Full/Incr VM Backup	Using backup-archive client GUI or CLI	Not available	Data Protection for VMware plug-in interface, backup Tab
Full VM Restore	Using backup-archive client GUI or CLI	Not available	Data Protection for VMware plug-in interface, restore Tab
Instant Volume restore	Not available	Using Data Protection for VMware recovery agent	Not available
File Level restore	Using Data Protection for VMware recovery agent	Using Data Protection for VMware recovery agent	Not available

Using Tivoli Storage FlashCopy Manager, we take advantage of the near-instant restore feature, that allows you to recover the data from the local snapshot as well as use the backups stored on the Tivoli Storage Manager server.

With the Data Protection for Oracle and RMAN, the table, table space, logs, or any other point-in-time recovery is satisfied. When needed, the data is sent back from Tivoli Storage Manager to the virtual machine using RMAN commands, via the Data Protection for Oracle API.

The Tivoli Storage FlashCopy Manager for VMware is used in case of severe incident or disaster. With Tivoli Storage FlashCopy Manager for VMware, VMware data stores are backed up. With the data that is backed up, we can restore data to virtual machines and virtual disks. The data can be restored to both original and alternative data store locations. Only virtual machines that we explicitly select for backup at backup time can be restored.

When restoring data with Tivoli Storage FlashCopy Manager for VMware, the following options are available:

- Restore to either the original data store or to a different data store at the virtual machine level.
- Restore a single virtual disk to the original location or a different virtual machine. This restore occurs by attaching a virtual disk from within a backup to a target virtual machine.
- Attach single virtual disks in the backup to the original or a different virtual machine to enable file level restore operations.
- Restore one or more data stores by using the **`instant_restore`** command or from the GUI. You can select which data store to restore and check the dependencies of virtual machines that are stored in this data store to other data stores in the environment. If a distributed virtual machine is present, a list of extra data stores are identified and you can select more data stores to restore for consistency.

Schedules

To have a central management of the schedules, we use the Tivoli Storage Manager server central scheduler. All the Tivoli Storage Manager agents deployed in our solution can be triggered by the Tivoli Storage Manager server scheduler, then take advantage of the built-in event status and other reporting features.

For Tivoli Storage FlashCopy Manager tasks (Microsoft Exchange and SAP DB2 in our case) and Data Protection for Oracle, we create Tivoli Storage Manager schedules, managed by the Tivoli Storage Manager server and not using local scheduler agent, like Windows Task Scheduler, even if it is supported for Tivoli Storage FlashCopy Manager. Thus we keep control and track all events from a single source, the Tivoli Storage Manager server.

For Tivoli Storage Manager for VMware and in-guest backup-archive client, the schedule plan is simplified. We need to create only one type of schedule since we use an incremental forever backup strategy.

For the Tivoli Storage Manager for VMware features, like Tivoli Storage Manager for Virtual Environments Data Protection for VMware and Tivoli Storage FlashCopy Manager for VMware, you can control, monitor, and restart the schedules using the vCenter plug-in as shown in Figure 7-3 on page 157. However, notice that only those created with the plug-in can be managed by the plug-in. If some tasks are created directly on the Tivoli Storage Manager server, they will not be available within the vCenter plug-in.

In our case, we always control the event, whatever the type of backup, from the Tivoli Storage Manager server.

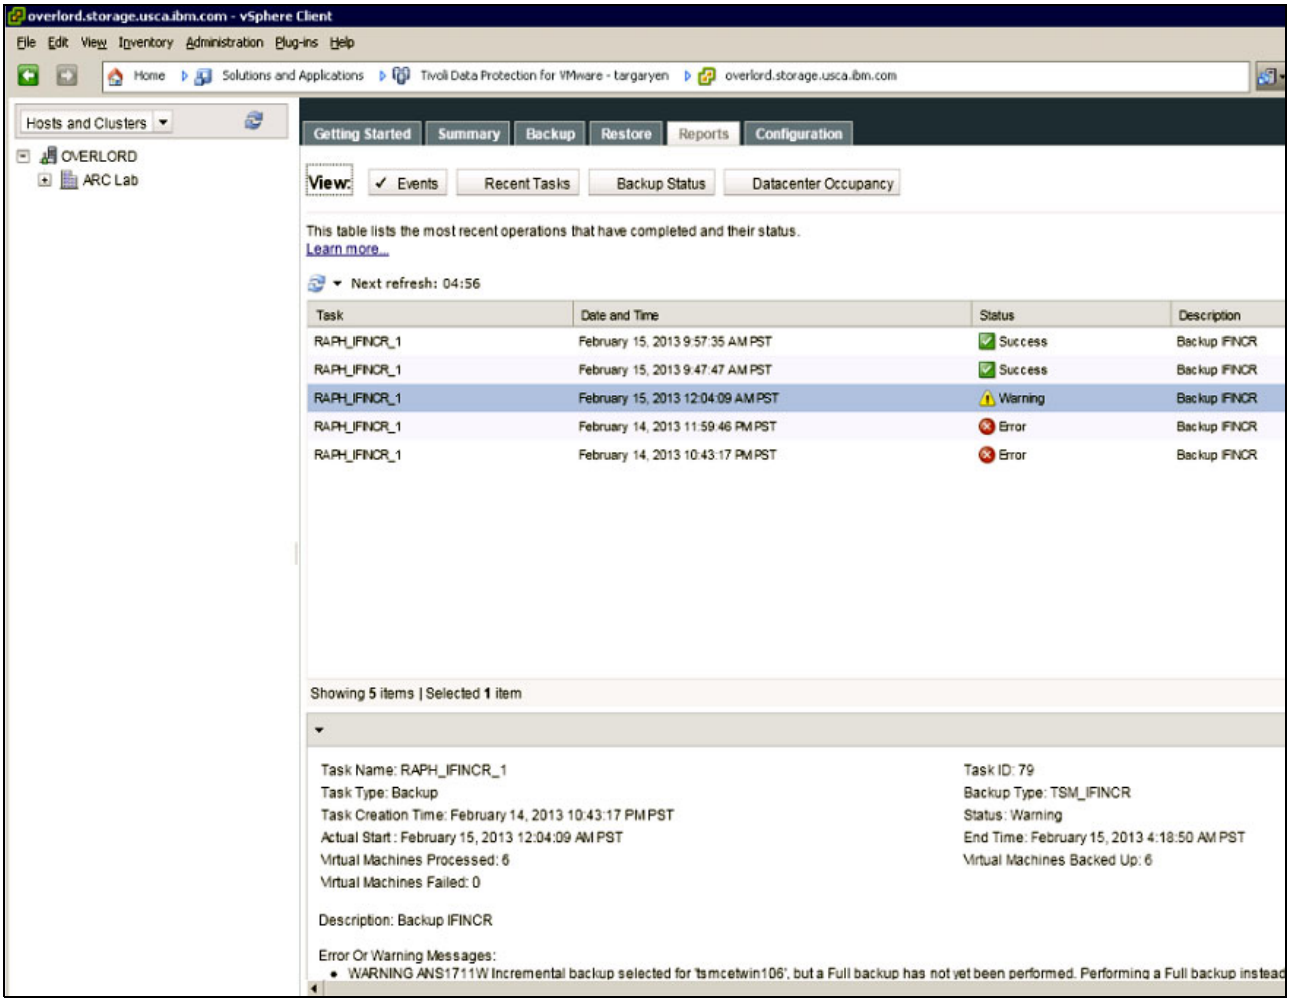

Figure 7-3 Data Protection for VMware vCenter plug-in to monitor and control your backups

The following schedules are defined to the Tivoli Storage Manager server:

- Oracle archivelogs backups, period=2 perunits=hours dayofweek=any weekofmonth=any action=command obj="<backup archivelog>"
- Oracle incremental backups, dayofweek=mon,tue,wed,thur,fri,sat weekofmonth=any action=command obj="<backup incremental level 1>"
- Oracle full backups, dayofweek=sun weekofmonth=any action=command obj="<backup incremental level 0>"
- FlashCopy incremental backups, dayofweek=mon,tue,wed,thur,fri,sat action=command obj="<flashcopy incremental>"
- FlashCopy full backups, dayofweek=sun weekofmonth=sun action=command obj="<flashcopy full>"
- Tivoli Storage Manager for Virtual Environment backups; period=1 perunits=day dayofweek=any weekofmonth=any action=backup subaction=vm option="-mode=incremental"
- In guest backup-archive client backup, period=1 perunits=day dayofweek=any weekofmonth=any action=incremental

7.2.5 Hardware and software requirements

The backup-archive client is supported within a virtual machine as described in the following IBM technote:

http://www.ibm.com/support/docview.wss?&uid=swg21239546

Tivoli Storage FlashCopy Manager prerequisites are described here:

http://www.ibm.com/support/docview.wss?uid=swg21612732

Tivoli Storage FlashCopy Manager for VMware requires the installation of a Linux vStorage Backup Server.

7.2.6 Additional information

See the following resources for more information:

- Tivoli Storage Manager for Virtual Environments Data Protection for VMware wiki page:

  https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/Tivoli%20Storage%20Manager/page/Tivoli%20Storage%20Manager%20for%20Virtual%20Environments

- Tivoli Storage FlashCopy Manager IBM wiki page:

  https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/Tivoli%20Storage%20FlashCopy%20Manager/page/Home

- Tivoli Storage Manager IBM wiki page:

  https://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli%20Storage%20Manager?lang=en

7.3 Virtualization: Hyper-V host-based backup

In this section, we describe the backup of data to the Tivoli Storage Manager server from the Hyper-V infrastructure. We show that combining the capabilities of Tivoli Storage Manager can protect your complete Hyper-V environment.

7.3.1 Summary of the client benefits

The client benefits that the Tivoli Storage Manager backup in a Hyper-V solution provides are listed here:

- Facilitate the disaster recovery of VMs.
- Supports Virtualized Clustered System at the host-level.
- Make consistent VSS snapshots of applications supported by the appropriate VSS writer.
- Improve recovery time.
- FlashCopy integration of pass-through disks or iSCSI disks.

7.3.2 Solution architecture

Hyper-V introduces the concept of virtual machine (VM) snapshots, which is to say point-in-time images of a virtual machine that you can return to at any stage. This functionality can be used in the total disaster recovery planning for your Hyper-V environment. We refer to this as *host-based backup*.

The latest addition to the Tivoli Storage Manager virtual machine backup capabilities is the backup and restore of Microsoft's Hyper-V virtual machines. This is a full virtual machine backup, using Volume Shadow Copy Service (VSS) to create a snapshot. See 3.1.4, "Microsoft Volume Shadow Copy Service" on page 42 for information about VSS mechanism.

Figure 7-4 illustrates the solution design.

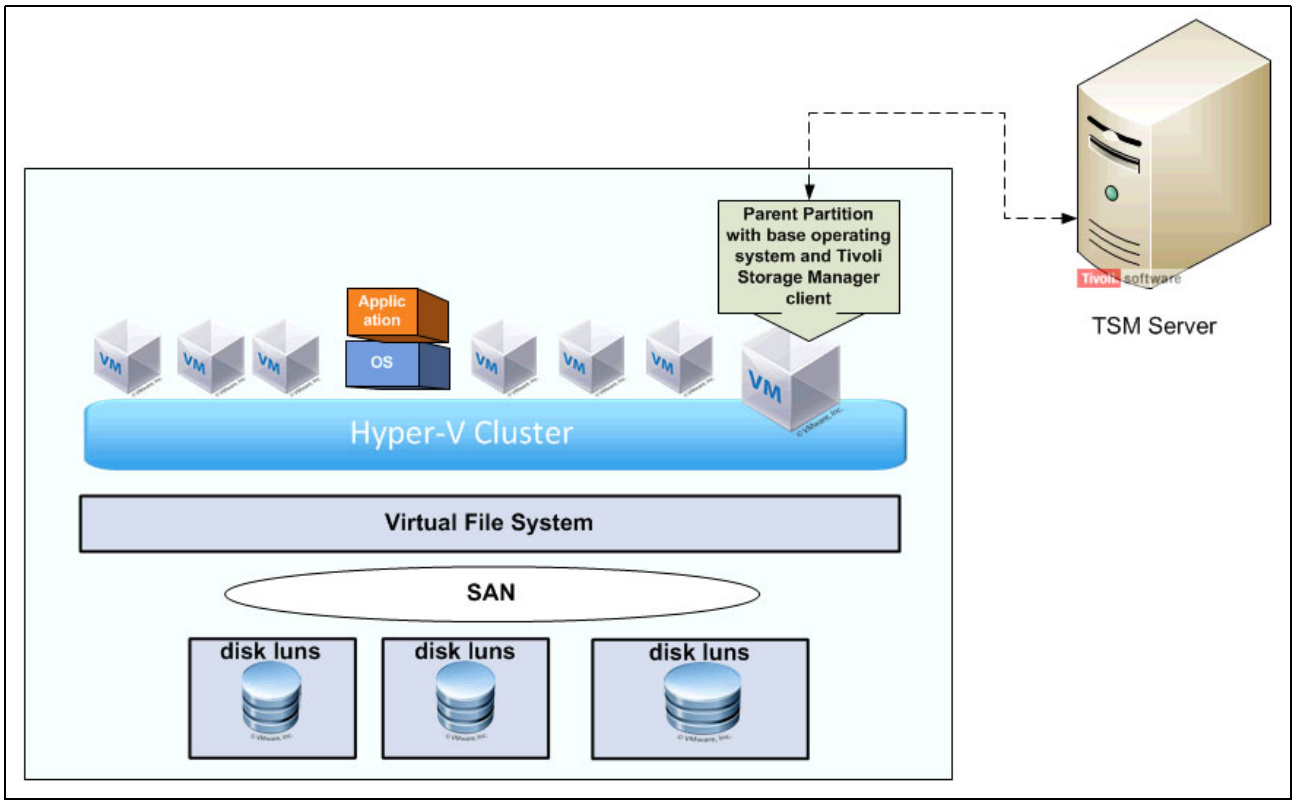

Figure 7-4 Infrastructure of the Hyper-V system

The Hyper-V cluster can service different kinds of VM operating systems and applications. The data that need protection may require more levels of protection in order to satisfy both disaster recovery capabilities and to make sure that the data backup is consistent across the infrastructure. Also consider the level of restore granularity.

The hypervisor layer is an interface to the hardware. It is isolated from all other layers, such as VMs and base operating systems. Hypervisor is done by Windows operating system, starting with Windows 2008, which have a Hyper-V role installed (at this time Windows 2008R2, Windows 2012). This is known as the parent partition. The guest VM operating systems, known as the child partitions, can be Windows operating system running 64 bit, 32 bit, or Linux.

Figure 7-4 illustrates a Hyper-V clustered environment running different kinds of VMs. It could be Microsoft SQL servers running on each Hyper-V node or a file server and print services.

The data is stored on a shared storage system. You must have only one Tivoli Storage Manager client installed on each physical machine, that is the base Hyper-V host system, not on the virtual machines.

7.3.3 Solution description

The snapshot process is implemented in the virtualization layer and can be taken at any time with any guest operating system. Snapshots can be taken whether the virtual machine is running or stopped. If the virtual machine is running when the snapshot is taken there is no downtime involved to create the snapshot. Guest virtual machine is basically a single entity, one or more .VHD files with a few configuration files. It can be backed up from the host operating system if the guest virtual machine is shut down.

Hyper-V has its own VSS writer which means that you have the ability to get application consistent backups of virtual machines from the host level. Some conditions must be met to ensure this works as expected. Look up the appropriate documentation on the VSS writer of the application used.

The basic operation is that when the Hyper-V VSS writer is triggered by Tivoli Storage Manager, it propagates the requests to the internal Hyper-V VSS requestor of the VM. The VM then in turn notifies all of its registered VSS writers (Exchange, SQL-Server, and more) that a backup is taking place.

The most common need for application-consistent backups is the use of database servers. Even though Microsoft SQL and Exchange server have their own VSS writer, this VSS mechanism will provide a crash consistent backup only. For proper database protection and advanced recovery possibilities, Tivoli Data Protection for Application within the virtual machine should be used.

If no VSS writer is available for the application you are running the backup will be crash consistent from which the application might or might not recover. In this situation, the advice is to install the Tivoli Storage Manager backup-archive client on the VM and maybe run a pre- and post-script to ensure application consistency.

Not all situations require an application-consistent backup. File and print servers are fine with crash consistent and possibly inconsistent backups, therefore full virtual machine backup is fine, with or without VSS enabled.

For the VMs that do not support VSS, such as Linux and Windows 2000, the saved state is used. Hyper-V will make a crash-consistent snapshot. The saved state is also used if the backup integration service is not installed on the VM. If VSS writer in the hypervisor is unable to communicate with VSS through integration components, its default behavior is to put the VM into a saved state before a snapshot of host volumes are taken for backup.

Note: Pass-through and iSCSI disks are not visible to the base operating system configured with the Hyper-V role. Therefore, the data on these type of disks must be backed up by the Tivoli Storage Manager that is installed on the VM.

Client deduplication can be enabled to reduce the data transferred to the Tivoli Storage Manager server. The default value (CLIENTDEDUPTXNLIMIT) for the maximum file size to be deduplicated is 300 GB, which should be greater than the size of the files sent to the Tivoli Storage Manager server. The policy used for the data must have a destination storage pool that is deduplication-enabled. See “Client-side data deduplication” on page 40.

7.3.4 Use scenarios

The use scenarios describe backup and restore of the virtual machines illustrated in the Figure 7-4 on page 159.

Command-line client interface

The command line is available for backup and restore and can be used as appropriate.

Graphical user interface (GUI)

When the Tivoli Storage Manager GUI client detects that it is running on a Hyper-V server, the GUI displays a list of Hyper-V guest virtual machines that can be backed up or restored.

Backup

Tivoli Storage Manager backs up or restores all files that are associated with a guest virtual machine. Tivoli Storage Manager server uses the file grouping to keep all files that comprise a virtual machine together as one virtual entity. Normal versioning in the central configured policies are used. Figure 7-5 shows the GUI for backup the VMs.

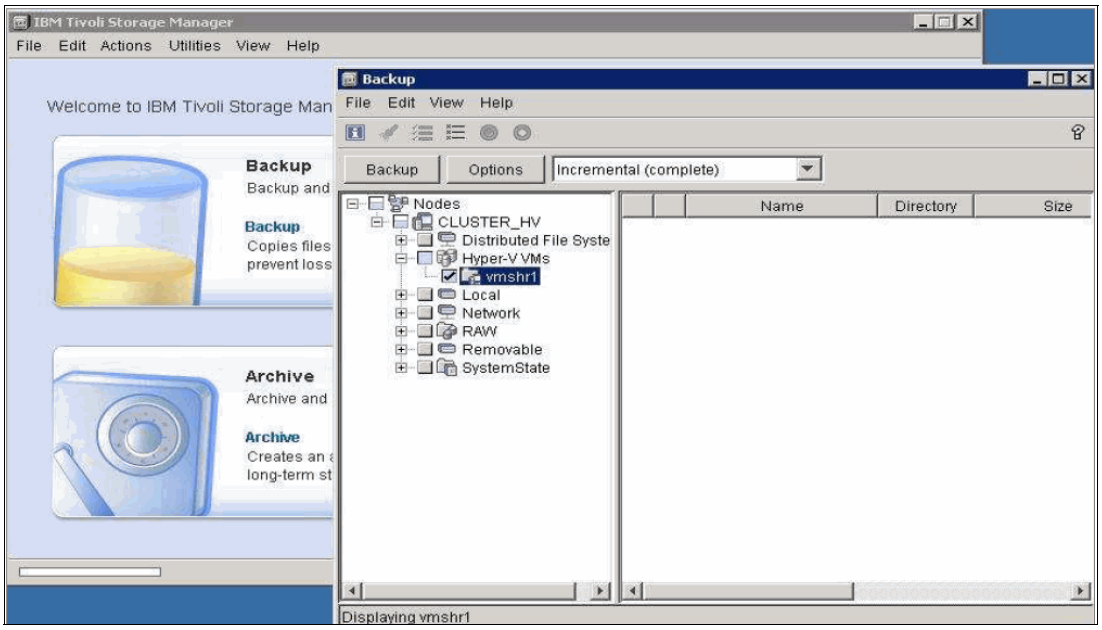

Figure 7-5 Tivoli Storage Manager client backup GUI interface with Hyper-V integration

The virtual machines must be backed up to the same client node in Tivoli Storage Manager using the **ASNODE** parameter. In that way, all VM backups are kept under the same node name and it is easier to manage if the VMs move to another physical Hyper-V node in the cluster.

Backing up multiple virtual machines onto the same Tivoli Storage Manager target node streamlines the management of backup and recovery operation. It simplifies the scheduling tasks because only one event must be created. That simplifies the recovery and the backup operations because virtual machines are stored within a target node.

Figure 7-6 Shows how the data is backed up to the same node name in Tivoli Storage Manager.

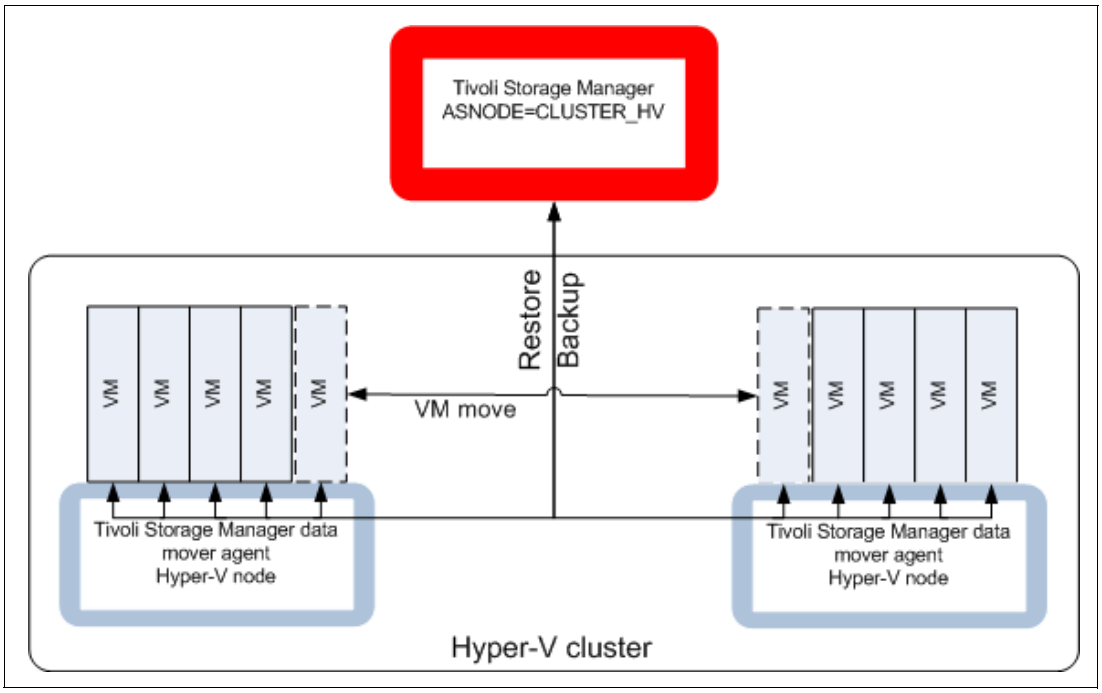

Figure 7-6 Host based backup and restore of virtual machines

Backing up a file server is straight forward because the data is consistent with a normal VSS snapshot. Because this is a full backup of the VM data, the size of the disks and the backup time must be considered.

For VMs that do not have applications that are installed, such as a print server service, using the VSS snapshot might be the only solution that you need. In fact, you might consider running a new full backup only when the server is software-updated. Otherwise, the crash consistent backup is adequate.

For application servers like Microsoft SQL or Microsoft Exchange we are able to take an application consistent VSS snapshot. It might be adequate to only backup the database at the host level once a day. If you need to make several backups during the day to lower the RPO you probably need to back up the database and its transaction logs on the VM using Tivoli Storage Manager for databases software.

To use this backup combination, see “In-guest backup capability” on page 164.

Restore

Multiple guest virtual machines can be restored. For example, all the guest virtual machines for a particular host are stored under one node name. Figure 7-7 shows the restore GUI for VMs.

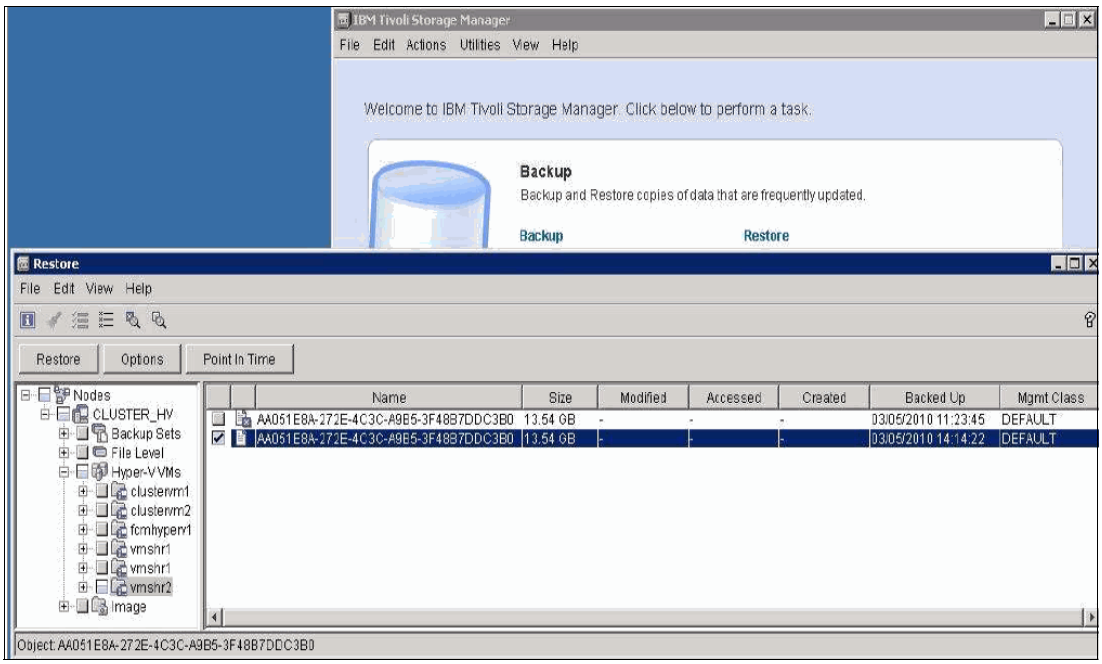

Figure 7-7 Tivoli Storage Manager client restore GUI interface with Hyper-V integration

The VM is restored to the cluster node that is listed as the current owner. It should be restored only by a backup-archive client installed on the cluster node that is the current owner of the VM. To determine the VM configurations of the cluster, use Windows and select **Server Manager →Features →Failover Cluster Manager →Cluster name →Services and Applications**. Use the Services and Applications window to change which node is the owner of the VM before you restore the VM.

When using the command line to restore all the Hyper-V, backup objects are displayed on the first page. After selecting the objects that you want to restore, the second page lists the backup details for the active and inactive objects.

When a restore operation is complete, the existing virtual machine is stopped (if it is running) and all existing files (for example, VHD) are deleted. When the backup snapshot is restored, the virtual machine that existed when the files were backed up is re-created. This includes any Hyper-V snapshots that existed when the files were backed up.

File level restore granularity is possible only if you have backed up the files from the VM with the Tivoli Storage Manager backup-archive client installed. This might be important for file servers where you need this kind of restore granularity.

For servers like print servers where a crash-consistent backup is adequate, the restore is quick to get the VM running again from the host-based snapshot backup. In this case, you would not install a Tivoli Storage Manager backup-archive client because you would not need to restore single files or directories.

With application servers like Microsoft SQL servers, more options need to be considered. These applications tend to have many transactions running during time. It is crucial that the state of the data is consistent and can be regenerated to a specific point in time. For these

data types we suggest to make a host-based FullVM once a week for crash-consistency. This must be combined with the Tivoli Storage Manager for databases or even Tivoli Storage FlashCopy Manager (see “Tivoli Storage FlashCopy Manager and Hyper-V VMs” on page 165) to make sure you can recover the transactions you need. See the next section (“In-guest backup capability” on page 164) for details about the in-guest solution.

In-guest backup capability

In-guest backups are started in the virtual machine and provide the same capability as for a regular environment (not virtualized); the following data protection solutions are possible:

- Virtual machines (without applications)
- Virtual machines with MS-SQL server and Exchange Server
- Virtual machines with applications such as Oracle or DB2
- Virtual machines with physical raw disks
- Virtual machines with any custom applications

If the guest OS of the VM is supported by Tivoli Storage Manager, we are able to back up the data on a file level basis. If the VM is running an application that is supported by a Tivoli Storage Manager Data Protection client, the data can be backed up also with this client.

Backup data is sent across the LAN from the VM to the Tivoli Storage Manager server storage hierarchy, as shown in Figure 7-8.

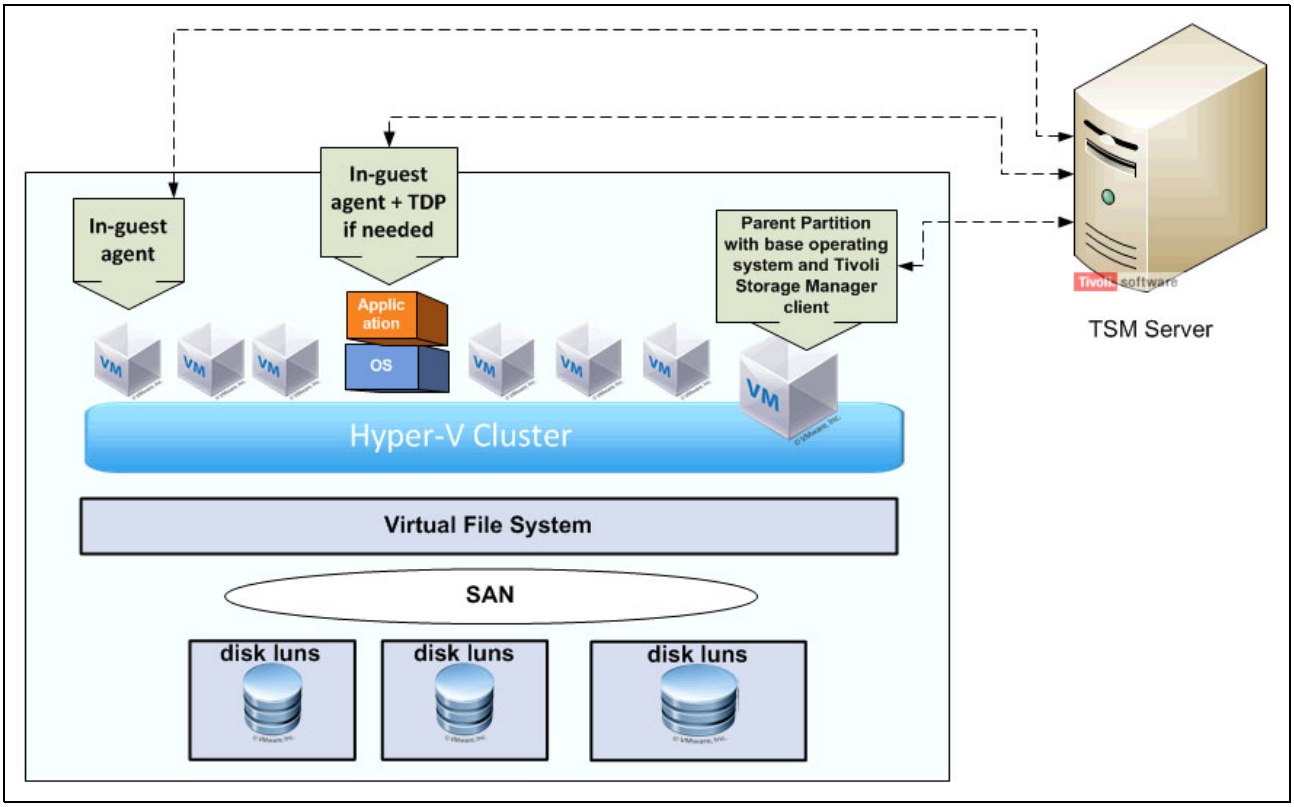

Figure 7-8 Combination of in-guest and host-based backup

If you are able to complete a full backup of all VMs every day and the RPO is adequate, do only the host-based VMFull backup. In most cases, this is not an option and we suggest that you combine the two options (host-based and in-guest backup) to get as much flexibility into the solution as you need.

The in-guest backup can protect the data that are stored on pass-through or iSCSI disks. Read about the in-guest backup solution in 7.4, “Virtualization: In-guest backup” on page 167

Tivoli Storage FlashCopy Manager and Hyper-V VMs

Tivoli Storage FlashCopy Manager meets this challenge to back up business critical applications, such as Microsoft SQL Server or Microsoft Exchange running on virtual machines (VMs).

Hardware-assisted snapshots within Hyper-V guest machines is supported to provide the application-aware VSS backup capability for Hyper-V guests, with the VSS Hardware Provider and Tivoli Storage FlashCopy Manager running on the guest machine.

Figure 7-9 illustrates the solution design. Tivoli Storage FlashCopy Manager can be used as a stand-alone product or integrate with a Tivoli Storage Manager server for advanced archiving, versioning, and scheduling capabilities. This specific solution focuses on the stand-alone installation for managing local snapshots.

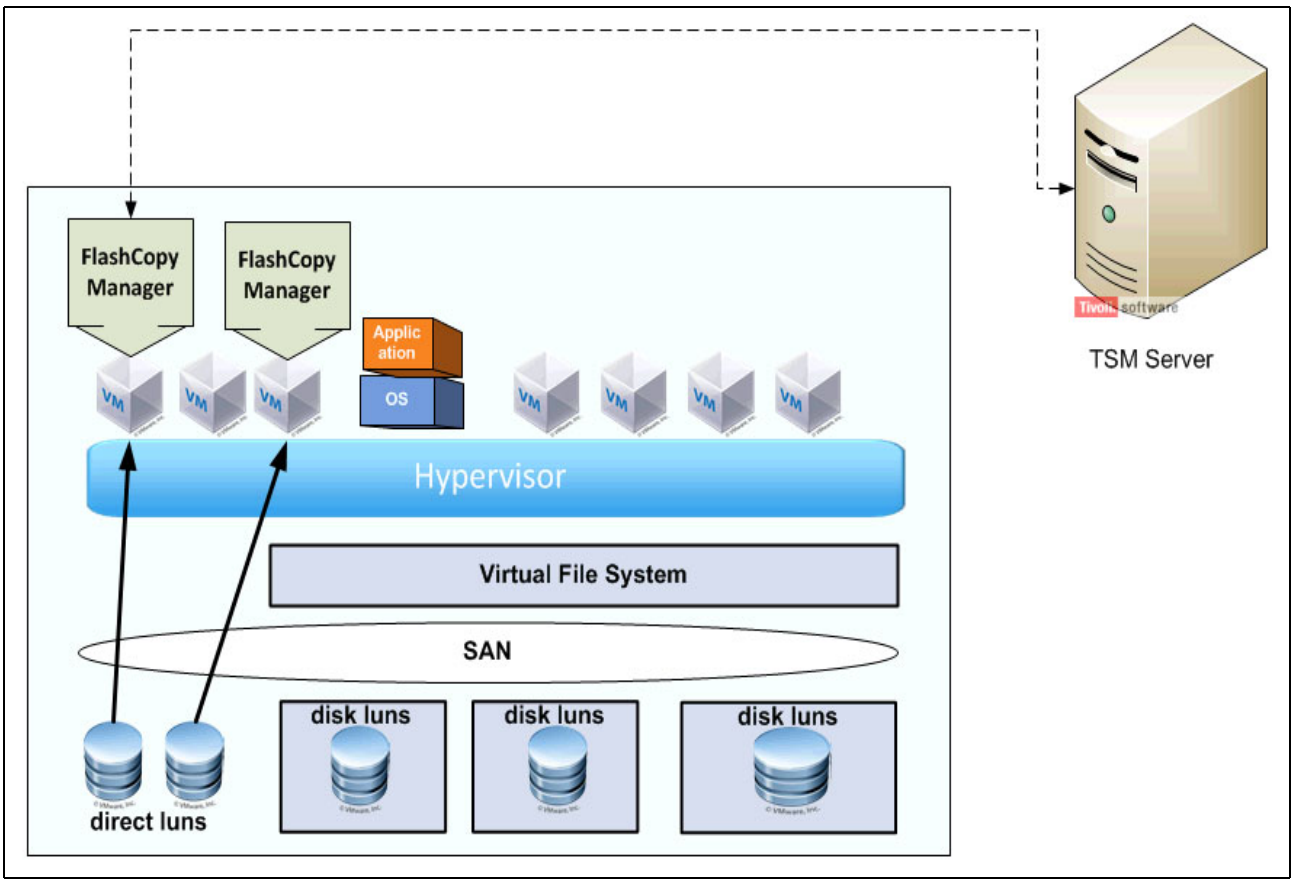

Figure 7-9 Using Tivoli Storage FlashCopy Manager in a Hyper-V environment

The IBM Storwize family products and SAN Volume Controller systems provide internal storage as well as external storage virtualization to make it possible to integrate with and manage existing heterogeneous storage, along with the internal storage from a single interface.

The virtual Fibre Channel switch for Hyper-V, combined with the virtual FC adapter on the virtual machine, provides direct FC connections from SAN storage to the guest virtual machine making use of the NPIV (N_Port ID Virtualization) technology. Data used by the VM is stored on a separate LUN than managed by Hyper-V. These disks are known as SCSI pass-through disks and iSCSI attached disks. We support FlashCopy backup of these disks through Tivoli Storage FlashCopy Manager, which the host-based backup of Hyper-V does not.

Combining the capabilities of Tivoli Storage Manager can protect your complete Hyper-V environment.

Scheduling: Automating the backup

Because we have different use case scenarios with Hyper-V, we also have different backup options in terms of scheduling:

- Host-based scheduling

  The Tivoli Storage Manager client scheduler is configured on the parent partition to run the scheduled host-based backup. It should run on each node to periodically backup the VMs owned by that node. The schedules must be defined with non-overlapping start times to avoid contention for common shared volumes. If the scheduler on each node backs up with the parameter **-vmlist="v,1.v,2"**, there is no need for the scheduler service to fail over because the scheduler service on the other node automatically picks up nodes, which move over on the next backup. Consider using ASNODE to allow for backup from multiple nodes and restore to fail over nodes if needed.

  The scheduler will run a command file that you build or a command scheduled directly from the Tivoli Storage Manager server. The backup command is shown in Example 7-1.

  Example 7-1 Backup command for VM based backup in Hyper-V

  ```
  dsmc backup vm -vmbackuptype=hypervfull
  ```

  The **-vmlist="vm1,vm2"** option can filter which VMs to include in the backup.

- In-guest scheduling

  The Tivoli Storage Manager client scheduler is configured on the VM. All schedules that the VM should run (file-level or database backup) is configured on the Tivoli Storage Manager server.

- Scheduling Tivoli Storage FlashCopy Manager copies

  Tivoli Storage FlashCopy Manager has its own scheduler that you can use. Or if you want to centralize the initiation and reporting on the schedules, you can use the Tivoli Storage Manager client scheduler to automate the copies.

7.3.5 Hardware and software requirements

The hardware and software requirements for Tivoli Storage Manager backup in a Hyper-V solution are as follows:

- Tivoli Storage Manager requirements

  Tivoli Storage Manager Windows Backup-Archive client Version 6.2 and later

- Hyper-V parent partition requirements

  Windows Server 2008 (or R2) 64-bit versions with Hyper-V role added

7.3.6 References

For more information and implementation references see the following sources:

- Tivoli Storage Manager guest support for virtual machines and virtualization

  http://www.ibm.com/support/docview.wss?rs=663&context=SSGSG7&uid=swg21239546#Microsoft%20Hyper-V%20Virtual%20Guest

- Protecting Microsoft SQL Server 2012 with IBM Tivoli Storage FlashCopy Manager

  https://www-304.ibm.com/partnerworld/wps/servlet/ContentHandler/stg_ast_sto_wp_protecting_sql_server_with_flashcopy

7.4 Virtualization: In-guest backup

In this section, we describe the backup of data to the Tivoli Storage Manager server when the virtualized infrastructure is based on platforms other than VMware and Hyper-V.

7.4.1 Summary of the client benefits

The client benefits of the Tivoli Storage Manager in-guest backup solution are as follows:

- Decrease the amount of data processed for data protection
- Improve the backup time
- Support Virtualized Clustered System
- Support Virtual Machine with raw disk type
- Facilitate the Disaster Recovery of VMs
- File level backup granularity

7.4.2 Solution architecture

The In-guest backup method can be used in all types of virtualized environments where the operating system of the VM is supported by a Tivoli Storage Manager backup-archive client, Tivoli Data Protection, or FlashCopy Manager. In this way, Tivoli Storage Manager supports a multitude of hypervisors like kernel-based virtual machine (KVM) or IBM PowerVM Workload Partitions Manager™. In Figure 7-10 on page 168, the in-guest Tivoli Storage Manager client is shown as the data protection solution.

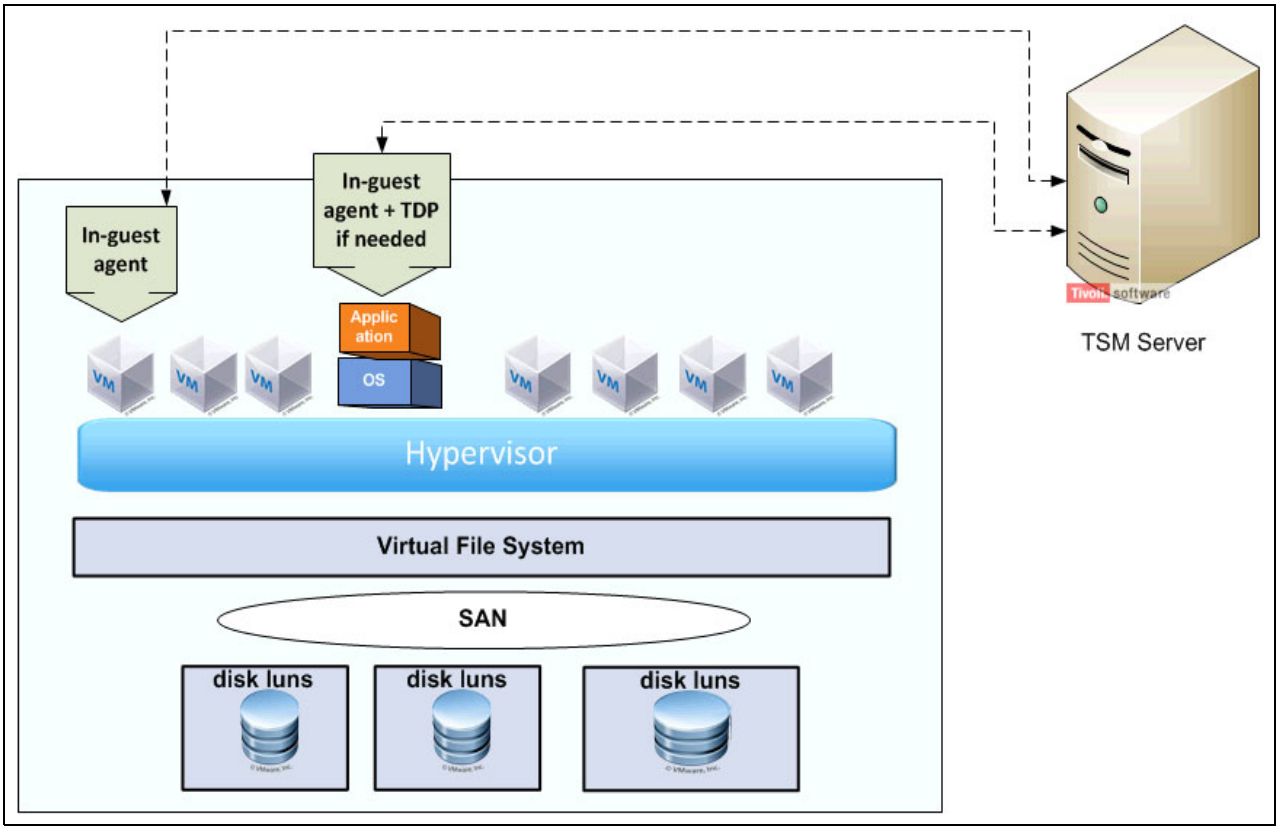

Figure 7-10 In-guest backup solution

Solution description

In-guest backups are started in the virtual machine and provide the following features:

- Virtual machines (without applications)
- Virtual machines with MS-SQL server and Exchange Server
- Virtual machines with applications such as Oracle or DB2
- Virtual machines with physical raw disks
- Virtual machines with any custom applications

If the guest OS of the VM is supported by Tivoli Storage Manager, we are able to back up the data on a file level basis. If the VM is running an application that is supported by a Tivoli Storage Manager Data Protection client, the data can be backed up with this client also.

Backup data is send across the LAN from the VM to the Tivoli Storage Manager server storage hierarchy.

7.4.3 Use scenarios

The use scenarios describe backup and restore of the virtual machines illustrated in Figure 7-10.

Backup

The same product backup capabilities that apply to a physical host, also apply to a virtual machine. All the standard function of the Tivoli Storage Manager backup-archive client can be used, such as progressive incremental forever, journal-based, and image backup. The VSS support of Microsoft operating systems is also supported.

For more information, see these sections:

- “Progressive incremental backup” on page 36
- “Journal-based backup” on page 37
- “Image backup” on page 36

If the virtual machine is part of a Microsoft Shadow Copy Services solution, we are able to support the disks in resource groups for backup.

You can use a server-side deduplication storage pool as a destination for your backup data to get the data reduction in the Tivoli Storage Manager storage pool.

Restore granularity

We can restore single files, directories, system state, or image level backups. The standard Tivoli Storage Manager backup-archive client GUI or command line can be used.

Restore virtual machine

We are doing a backup at the file level or image level, which means a partition or file system. Therefore if the virtual machine needs to be recovered completely, other mechanisms must be used, combined with the Tivoli Storage Manager file level backup. The basic steps to recover the virtual machine are as follows:

1. Provision a new base virtual machine with the used applications.
2. Restore the file level data that are not included in the standard image provisioned.
3. Restore system state data (if Windows).
4. Restore application data if they exist.

Although these steps seem easy, several factors must be considered to be sure the process is working on every type of virtual machine you have in your environment.

System backup and recovery solutions can be used to make the recovery process easier and more robust against change. One example is the Network Installation Manager (NIM) server in AIX. The NIM server can keep **mksysb** backups of the AIX system of a VM and make the recovery procedure easier. For the given hypervisor, other solutions might be considered to incorporate it to a complete disaster recovery plan of the IT infrastructure.

Tivoli Storage FlashCopy Manager

The FlashCopy integration can work in any type of virtual environment if the Tivoli Storage FlashCopy Manager software requirements are met, hardware and software. See “Snapshot at the storage hardware layer” on page 49 to understand the concept of Tivoli Storage FlashCopy Manager.

You can also review the virtualization support:

http://www.ibm.com/support/docview.wss?uid=swg21433737

Figure 7-11 on page 170 shows how data is kept on its own set of LUNs separate from the LUNs that the hypervisor manages, otherwise FlashCopy Manager will not be recognized as the FlashCopy enabled disks.

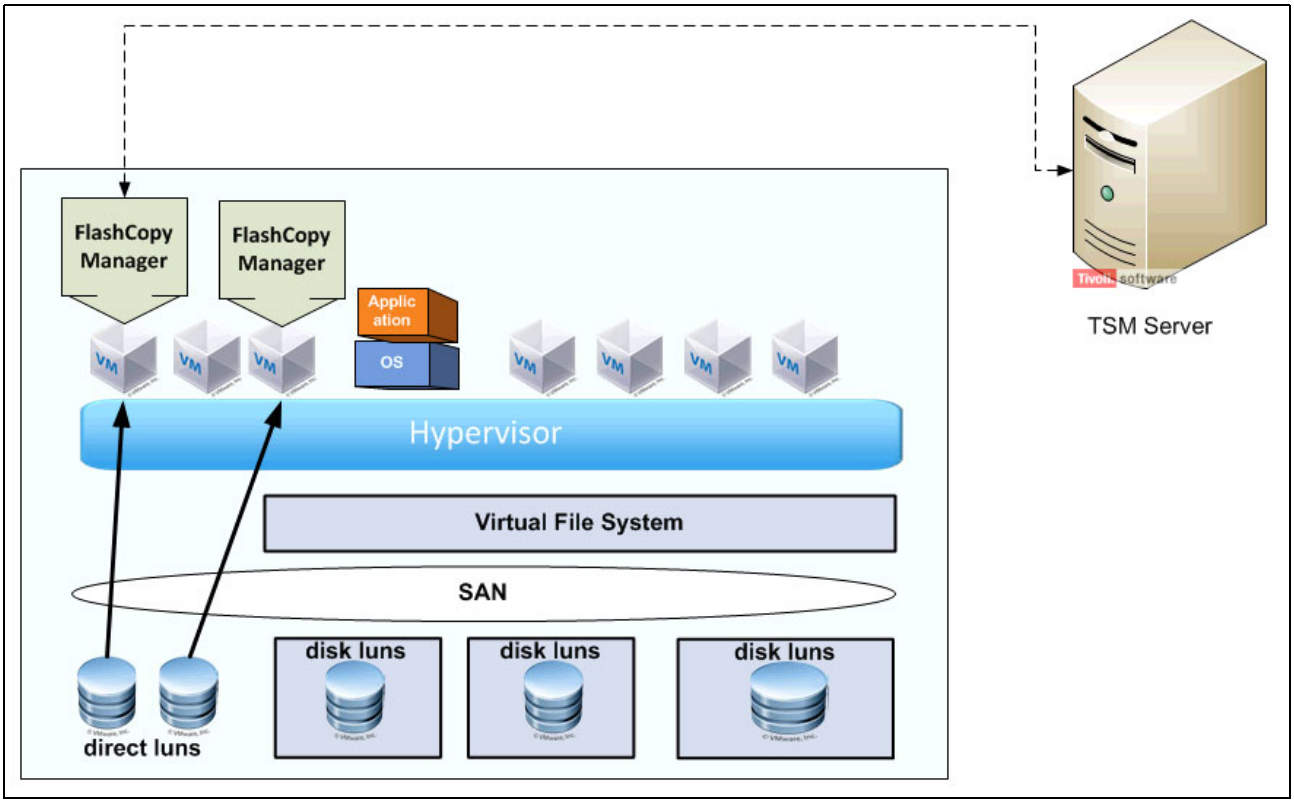

Figure 7-11 Tivoli Storage FlashCopy Manager integration in virtualized environments

Scheduling

In a virtual environment, you can easily have 10 - 50 VMs running on a single physical host. Each VM will have its own Tivoli Storage Manager client scheduler running to trigger the daily incremental backup job to back up the file system, services database, and applications.

The backup window and randomization used by Tivoli Storage Manager scheduler can help you distribute the backup load on the physical host to minimize the resources used at a specific time.

Another add-on option is to use journal-based backup. It uses a change journal maintained by the Tivoli Storage Manager journal service process running in the Tivoli Storage Manager backup-archive client. Because we do not need to use resources to scan the entire file system of the VM, it runs much quicker and fewer less CPU cycles to complete the backup.

7.4.4 Software requirements

On Windows, journal-based backup is supported on 2003, 2003 R2 32-bit and 64-bit, Vista, Windows 7 32-bit and 64-bit, 2008 32-bit, 2008, 2008 R2 64-bit. It does not use the journaling facility inherent in Windows NTFS file systems or any other journaled file system.

On AIX, journal-based backup is supported on JFS and JFS2 file systems.

On Linux, journal-based backup is supported on Ext2, Ext3, Ext4; XFS, ReiserFS, JFS, VxFS, and NSS, and for a local file system shared through NFS. GPFS is not supported for journal-based backups.

7.4.5 Tivoli Storage Manager as a comprehensive solution for a virtualized system

As this chapter shows, Tivoli Storage Manager has adapted to virtualization. Even more, it anticipates the future by providing a cloud-ready solution and hardware snapshot based solution. Tivoli Storage Manager no longer needs to rely on the system virtualization layer capability to protect the data. Therefore Tivoli Storage Manager products can help you with future system virtualization challenges.

7.5 Application and email servers

The most important factor in a backup and recovery strategy is to meet the recovery business requirements. A database or application can be used for a variety of roles in an organization, from a development database or application with a few megabytes of information being accessed by a few developers to a non-stop production database/application with several terabytes of information being accessed by thousands of users around the world. Likewise, the recovery requirements for these particular databases/applications might differ significantly: maybe a weekly backup will suffice, or a comprehensive strategy should be designed to achieve the recovery requirements.

As you see from the matrix in Chapter 4, “Tivoli Storage Manager challenge matrix” on page 81, the online backups can help you meet the recovery time objective (RTO) and recovery point objective (RPO). To identify the recovery requirements, answer the following questions:

- How critical are your databases and applications to your business?
- Is it acceptable to lose any database or application activity? How much time and changes are acceptable to lose?
- What is the maximum acceptable downtime for your databases and applications?
- Is it necessary to perform a restore right up to a point of failure?
- What is the backup window? How many resources are available for the backup?
- Must the database or application be available only during commercial hours or working days, or does it have to operate on a 24x7x365 basis?
- Are there peak periods of database and application utilization? How frequently does the data change during peak and nonpeak periods?
- Do you have two or more databases or applications that must maintain a logical consistency?
- What are the legal requirements for your backup routines? Are you required to retain backups for a long period of time (more than your normal retention period)?

The answers to these questions can help you determine a backup strategy.

Tivoli Storage Manager for Mail and Tivoli Storage Manager for Databases enable data protection of your mail and databases while they are online. They automate data protection, enable hot backups without shutting down the application and improve data restore performance.

Offline backups

Sometimes you need to run offline backups of your database or application. At this time, you closed your database/application for maintenance and will use Tivoli Storage Manager Backup-Archive client to perform the backup. This will copy the files to the Tivoli Storage Manager server at the file system level, without any application reference and file dependency.

7.5.1 Tivoli Storage Manager for Databases: Data Protection for Oracle

The Tivoli Data Protection for Oracle application client provides an integrated solution for performing backup and restore operations on Oracle databases. It is a client application that provides backup of online databases and restore of databases to the original or different location.

Challenges and benefits the solution addresses

Tivoli Data Protection for Oracle is not intended as a substitute for the standard Tivoli Storage Manager backup-archive client. Tivoli Data Protection for Oracle cannot be used to back up or restore any non-database data, such as history files or any other system configuration files. Those files must be backed up by the Tivoli Storage Manager backup-archive client. Therefore, the two client types work together to provide full data protection for your Oracle environment. The Tivoli Data Protection for Oracle and the Tivoli Storage Manager backup-archive client can run simultaneously on the same Oracle server, however, they are totally separate clients where the Tivoli Storage Manager server is concerned.

IBM Tivoli Storage Manager for Databases helps protect Oracle Databases data no matter where it is stored. You can continue running primary applications on your database servers, while they back up and restore data to and from offline storage using automated tasks, utilities and interfaces. This software performs online, consistent, and centralized backups to help you avoid downtime, protect vital enterprise data, and minimize operational costs.

Tivoli Storage Manager for Databases provides the following functions:

- Protects a wide range of application data by securing the underlying database management systems that contain that data.
- Takes advantage of Oracle Server backup-certified utilities and interfaces.
- Applies IBM Tivoli Storage Manager automated data protection tasks to running database servers.
- Supports the Oracle Automated Storage Management (ASM) feature available in Oracle 10g and later releases.

Tivoli Data Protection for Oracle and RMAN

Oracle Recovery Manager (RMAN) provides consistent and secure backup, restore, and recovery performance for Oracle databases. The Oracle RMAN initiates a backup or restore; Data Protection for Oracle acts as the interface to the Tivoli Storage Manager server. The Tivoli Storage Manager server then applies administrator-defined storage management policies to the data. Data Protection for Oracle implements the Oracle defined Media Management application program interface (SBTAPI) 2.0. This SBTAPI interfaces with RMAN and translates Oracle commands into Tivoli Storage Manager API calls to the Tivoli Storage Manager server.

With the use of RMAN, Data Protection for Oracle allows you to do these functions:

- Full and incremental backup function for the following items while online or offline:
  - Databases
  - Table spaces
  - Data files
  - Archive log files
  - Control files
- Full database restores while offline
- Table space and data file restore while online or offline

RMAN provides the interface to the Oracle database and the functions for backup, restore, and recovery. It does not provide any storage management capabilities and must be integrated with other storage management products such as Tivoli Storage Manager to provide a complete enterprise wide storage management solution.

Solution architecture

This section describes the solution components for using Tivoli Data Protection for Oracle to protect your Oracle database.

Fundamentals of Relational Database Management Systems (RDBMS)

RDBMS products share a common set of principles. The purpose of this section is to explain the basic principles in order to design a backup and recovery system for a relational database. Although all RDBMS products are based on the same set of principles, not all use the same terminology or structures.

Databases

A database consists of one or more logical units called *table spaces*. On the physical layer, the database has data files, control files, and optionally a password file.

Table spaces

When a database is created, a data dictionary is created in the system table space. Although It is not required to create additional table space, additional table spaces are suggested for user data. The data dictionary is critical to the operation of the database because it records, verifies, and conducts ongoing work. The system table space is always online and cannot be taken offline because the data dictionary must always be available to Oracle.

Data files

A table space is a logical grouping of data storage called *data files*. A data file can be a file or a raw device. A table space can have a mixture of both files and raw devices as data files. Backup can be performed on a logical level (database and table spaces) or on a physical level (data files).

Control files

Control files keep information about the physical structure of the database and log files. They are commonly multiplexed and are defined in the initialization parameter files. Keep at least three copies on separate disks. Just like the data dictionary, it is important to make backups of your control files regularly. Losing the control files makes recovery much more complicated.

Initialization parameter files

The initialization parameter files are text files that contain instance configuration parameters, such as how much memory to use and what to do with filled online redo logs. Back up the initialization parameter files whenever configuration parameters change.

Password file

A database can optionally have a password file. The password file is used during remote administration of a database server. The password file must be backed up when there changes or additions of administrative users.

RMAN system components

RMAN consists of several components that interact during the backup and recovery process. The system components involved with RMAN, and the flow of an RMAN operation to Tivoli Storage Manager, are illustrated in Figure 7-12 on page 176.

RMAN command

The **RMAN** command is the database administrator's interface to RMAN. It invokes a command-line interface that provides a scripting language (operating system independent) for performing backup and recovery operations. RMAN can be executed either interactively, where a command prompt is displayed and additional RMAN commands entered, or in batch mode, where an RMAN command file is executed.

Target database

The target database is the Oracle database instance on which RMAN executes specified backup, restore, and recovery actions. When the **RMAN** command is run, it connects to the target database. The target database is specified by using RMAN parameters.

Communication channel

RMAN can perform backup and restore functions to either local disk or to external media management products such as an Tivoli Storage Manager server through the `libobk.a` library provided by Tivoli Data Protection for Oracle. These I/O operations are performed over a communication channel that defines the device to be used for the operation. The channel is used by RMAN to send or receive backup data to and from the I/O device. For backup and restore operations, you must allocate a channel before the operation is performed. A channel corresponds to a single device. With the Tivoli Data Protection for Oracle, a channel is a single session to a Tivoli Storage Manager server using the SBT API. Multiple channels can be allocated. RMAN provides a multiplexing feature that enables parallel data streams to be sent over multiple allocated channels to maximize backup and recovery performance.

Recovery catalog

The recovery catalog is the repository for information about backup objects created by RMAN. It is an Oracle database instance, separate from the target databases, and can contain information for multiple target databases. The data stored in the recovery catalog comprises structural information about the target databases to back up and restore. The recovery catalog contains the following information:

- Physical schema of a target database

  You must register the target database at the recovery catalog to define the physical schema of the target database. RMAN needs to know about any structural change of the target database, and obtains this information from the target database control file.

- Database backup history

  RMAN backs up databases, table spaces, data files, control files, and archive logs to the Tivoli Storage Manager server. Details of these backup objects held on Tivoli Storage Manager is stored in the recovery catalog.

- Backup and recovery history

  RMAN stores backup, restore, and recovery information to maintain a history of previously performed operations. When backup and restore operations are done, this information enables RMAN to determine the following information:

  - Database files that require backing up
  - Old backup files that can be deleted
  - Files that are not recoverable

- Stored RMAN scripts

  RMAN commands can be stored in the recovery catalog as stored scripts. Scripts can be created to automate the execution of a several RMAN operations.

Tivoll Storage Manager components

Tivoli Storage Manager components interact with RMAN during backup and restore operations.

Tivoli Storage Manager server

This is a computer where the Tivoli Storage Manager server program has been installed. The Tivoli Storage Manager server is responsible for managing data objects sent by a Tivoli Storage Manager client. A Tivoli Storage Manager Administrator allocates disk and tape storage to the Tivoli Storage Manager Server.

Tivoli Storage Manager client

This is a computer where a Tivoli Storage Manager client program has been installed. A Tivoli Storage Manager client can be a backup-archive client, a Tivoli Data Protection (Tivoli Data Protection client, or a third-party product that uses the Tivoli Storage Manager Application Program Interface (API).

Tivoli Storage Manager backup-archive client

This is a Tivoli software program that allows data objects to be sent to a Tivoli Storage Manager server. It typically consists of the files and directories on a computer.

Tivoli Storage Manager Application Program Interface (API)

This is the Tivoli software that programmers use to interface with the Tivoli Storage Manager server. The Tivoli Storage Manager API is used by both Tivoli Data Protection products and other vendor products to send data objects to the Tivoli Storage Manager server.

Flow of RMAN operation

Figure 7-12 on page 176 shows the operation flow as follows:

1. The operation is started by invoking the `RMAN` command and entering the appropriate commands directly or by submitting a command file containing the commands.
2. RMAN connects to the recovery catalog.
3. RMAN also connects to the target database.
4. Before any backup or restore operations can be performed, RMAN allocates a channel to Tivoli Storage Manager, using the SBT API.
5. RMAN then creates a server process on the target database instance that performs the operation.

For restore operations, RMAN queries the recovery catalog to determine which files to restore from Tivoli Storage Manager. For a backup operation, RMAN backs up the objects specified in the command to Tivoli Storage Manager. In both cases, the data is transferred on the previously defined channel.

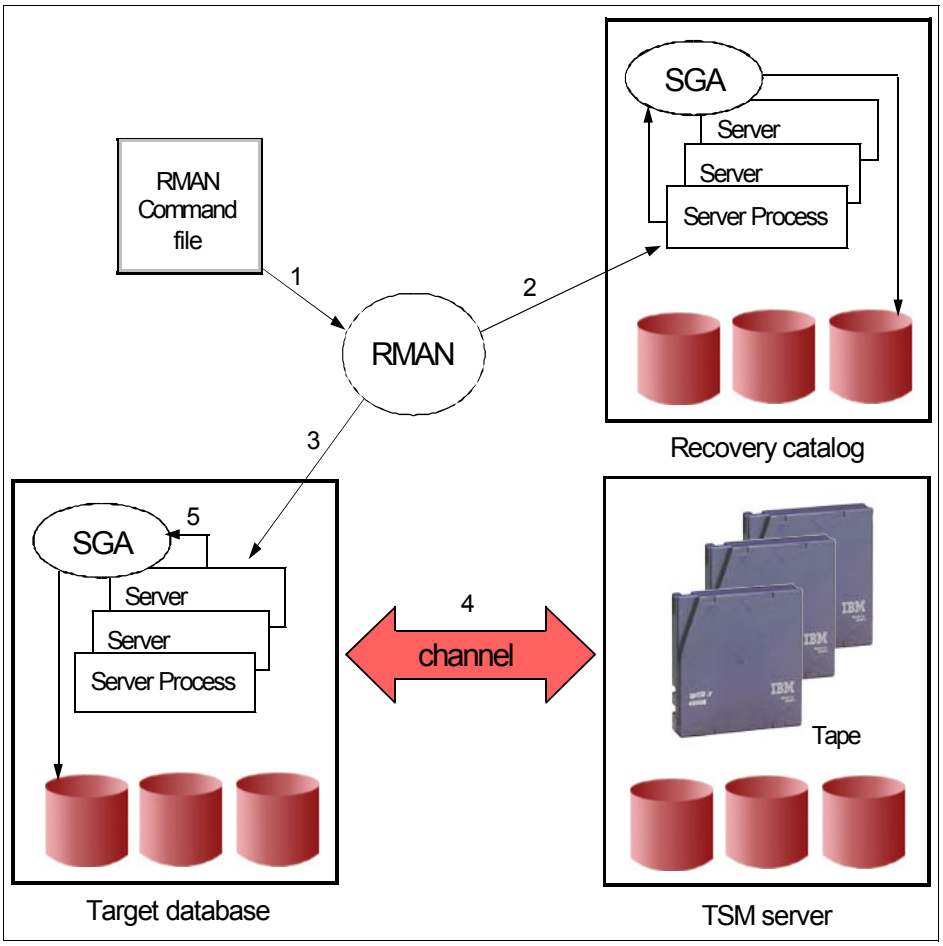

Figure 7-12 Flow of an RMAN operation to Tivoli Storage Manager

Solution description

Planning is one of the most important areas for consideration before beginning to use Tivoli Storage Manager for database. It is important that the database administrator and the Tivoli Storage Manager administrator work together to anticipate the circumstances in which recovery will be required, and also the resource and configuration requirements. These ideas apply to all types of databases.

Backup requirements

A backup strategy is only one part of your overall data management plan. You must consider how important your data is to the function (or even existence) of your organization. The less time that your organization can function without its data, the more important that data is to you. Your system must be designed in such a way to keep important data available when a failure occurs. Reliance on backups is not necessarily sufficient. Before you design a backup strategy, define the requirements that the strategy must satisfy. These are some factors to consider when you define the requirements for your backup strategy:

- Types of events (the categories of incidents that may occur)
- Speed of recovery (how quickly you need to be able to recover)
- Backup windows (the periods of time at which backups can be performed)
- Recovery points (to which points in time you need to be able to recover)
- Units of recovery (which other tables and files need to be recovered to the same point in time)

Also consider the following factors:

- Redundant Arrays of Independent Disks (RAID) devices
- Dual access paths
- Dual I/O controllers
- Dual power supplies
- Backup or standby processors
- Uninterruptable power supplies

None of these on their own can guarantee the availability of your data, but in combination they can reduce the impact of a failure.

Tivoli Storage Manager server considerations for Oracle backups

Whenever you use the Tivoli Storage Manager server to back up and restore data objects, considering which management class the data objects will be bound to is especially important. This is true of both API and backup-archive clients. Failure to do so will result in storing the data objects in one of three situations:

- Just right: It is highly unlikely that you will manage the objects “just right” if you do not take the time to define your storage requirements, configure the Tivoli Storage Manager server appropriately, and configure the Tivoli Storage Manager client to use the correct management classes.
- Too long: If you store the data objects for too long, you waste space and storage resources on the Tivoli Storage Manager server.
- Too short: If you store the data objects for too short a time, you do not have the required files when you need them.

The Tivoli Storage Manager server cannot and does not know what length of time a client program needs to keep the data objects. This must be done by the client.

How Tivoli Data Protection for Oracle stores data objects

Database objects stored on the Tivoli Storage Manager server by Tivoli Data Protection for Oracle are stored as backup objects. Each Oracle backup is stored as a unique object by generating a random character string as the low level qualifier (LL_NAME). The RMAN script can control what the LL_NAME or backup piece name is by using the RMAN format command. If using the format command, you should generate unique backup piece names by either random character string (%U option) or timestamp (%s and %t). This is documented in the Oracle RMAN manual. Failure to do so will cause inconsistency between the RMAN catalog and the Tivoli Storage Manager server. Using unique names means that the Oracle backups must be manually inactivated. This is done by allocating a channel for deletion using the same node name and file space name that was used to perform the initial backup. This also means that the management class to which the backup objects are bound should have retention settings that change the inactivated backup objects to be expired immediately. The following retention settings for a backup copy group can facilitate this:

- `RETONLY=0`
- `VERDELETED=0`

Policy management considerations

When you decide how to set up your Tivoli Storage Manager server to store Oracle backup objects, you must decide how you will organize your nodes into domains. You must also correctly define and configure your management classes.

Domain considerations

You can use one domain for all your client data or you can specify multiple Tivoli Storage Manager domains to group nodes with the same backup characteristics together. Using multiple domains helps ensure that the data objects get bound to the appropriate management class. Consider the following items if you want to use multiple domains:

- Backup-archive client, API client
- Platforms separated, such as AIX, Oracle, Windows
- Critical data, non-critical data
- Group nodes with the same data characteristics together

Because the policy requirements for Oracle backups differ from the settings you want for regular Tivoli Storage Manager backup clients, a different management class must be defined within Tivoli Storage Manager for managing these Oracle backups. There are two ways to implement this different management class setup:

- Define a new management class within an existing policy domain.
- Define a separate policy domain where the default management class contains the required settings.

Backup copy group considerations.

Normally Tivoli Storage Manager backup copy groups are designed to hold multiple versions of files and directories in order to restore not only the latest (active) but also older versions that were changed or deleted (inactive). Tivoli Data Protection for Oracle sends backups through the Tivoli Storage Manager client API directly to the backup copy group of the default management class to which the node is assigned. Oracle assigns unique names to every database backup. The settings that pertain to multiple versions do not apply. Use the following retention settings for the management class that will be bound to the Oracle backups.

- VEREXISTS=1

  Keeps only one version of the backup file because the name of the backup is unique. (There will not be a newer version of the backup image with the same name.)

- VERDELETED=0

  If the backup file was deleted, then Tivoli Storage Manager should not keep an inactive version of this file.

- RETEXTRA=0 (the same value as RETONLY)

  This parameter will never be used because you will never have more than one version of the backup file. To prevent confusion set this parameter to the same value as RETONLY.

- RETONLY=0

  When a backup image file becomes inactive it will be purged from the Tivoli Storage Manager server at the next expiration.

Crosscheck utility

The crosscheck utility verifies whether backups still exist on disk or Tivoli Storage Manager. RMAN does not delete backup entries that it could not find, but instead marks them as expired. If the backup was erroneously marked expired, for example, because Tivoli Storage Manager was unavailable or misconfigured, the crosscheck utility will mark it available the next time it is run if the backup still exists.

Life cycle of Tivoli Storage Manager data objects

Backup data objects and archive data objects are managed differently. We use the term *life cycle* to describe how these data objects exist on Tivoli Storage Manager storage from initial creation to when they are purged.

Life cycle of backup data objects

A backup object exists in three states, *active*, *inactive*, and *expired*, before being purged from the Tivoli Storage Manager server. The four steps involved in the life cycle of a backup data object are listed here.

1. A copy of the client data is sent to the Tivoli Storage Manager server as a backup object. When a backup object is sent to the Tivoli Storage Manager server, it becomes the active version.
2. It remains in an active state until the Tivoli Storage Manager client program deletes the backup object manually, or a newer version of the backup object is sent. The backup object changes state from active to inactive.
3. The backup object remains inactive until it exceeds its retention settings. A backup object can exceed retention settings by either time or number of versions. The backup object changes state from inactive to expired.
4. The backup object remains in the expired state until expiration processing runs on the Tivoli Storage Manager server. This process is invoked by a Tivoli Storage Manager administrator with the expire inventory command. When expiration processing encounters a backup object in the expired state, it purges that object from the Tivoli Storage Manager database and frees up the storage space where the backup object resided.

A backup object that is the active version or in the active state will never be purged from Tivoli Storage Manager storage, that is, it never expires. It must first be inactivated by the Tivoli Storage Manager client program. The Tivoli Storage Manager client program can do this by manually deleting the backup object or sending a new version of the backup object. When a backup object becomes inactive or moves into the inactive state, it is still accessible by the Tivoli Storage Manager client. A main difference between active and inactive is that an active object becomes inactive due to a client operation. An inactive object becomes expired automatically by the Tivoli Storage Manager server as soon as it exceeds its retention criteria. Changing from inactive to expired does not require a client operation. There is no way for a backup object to change back to the active state once it has become inactive. When a backup archive object moves into the expired state, it is no longer accessible by the Tivoli Storage Manager client. Additionally, there is no way for the backup object to change back to the inactive state once it has become expired. If the retention for the backup object is set to retain zero inactive objects (retextra=1,verdel=0) or to retain inactive copies for zero days (retextra=0, retonly=0), the active backup object will change to the expired state as soon as the active backup is inactivated

Automatic deletion of old backups

Automating deletion of Tivoli Data Protection for Oracle for Windows should be planned and tested carefully before implementation in a production environment. There is no simple command that can be run to easily deactivate backups based on redundancy or time. One way to get around this is to use command files for automation. Another way is to use tags as part of your backup process and use the same tags for deletion.

Use scenarios

This section describes additional scenarios for using Tivoli Data Protection for Oracle.

Back up the database using Tivoli Data Protection for Oracle and RMAN

When you execute the backup command, you create one or more backup sets. A *backup set,* as shown in Figure 7-13 on page 180, which is a logical construction, contains one or more physical *backup pieces*. Backup pieces are operating system files that contain the backed up data files, control files, or archived redo logs. You cannot split a file across different backup sets or mix archived redo logs and data files into a single backup set.

A backup set is a complete set of backup pieces that constitute a full or incremental backup of the objects specified in the backup command. Backup sets are in an RMAN-specific format; image copies, in contrast, are available for use without additional processing. Each backup piece contains control and checksum information that allows the Oracle server process to validate the backup piece during a restore. A backup set is created by the backup command. A restore command is required to extract files from a backup set.

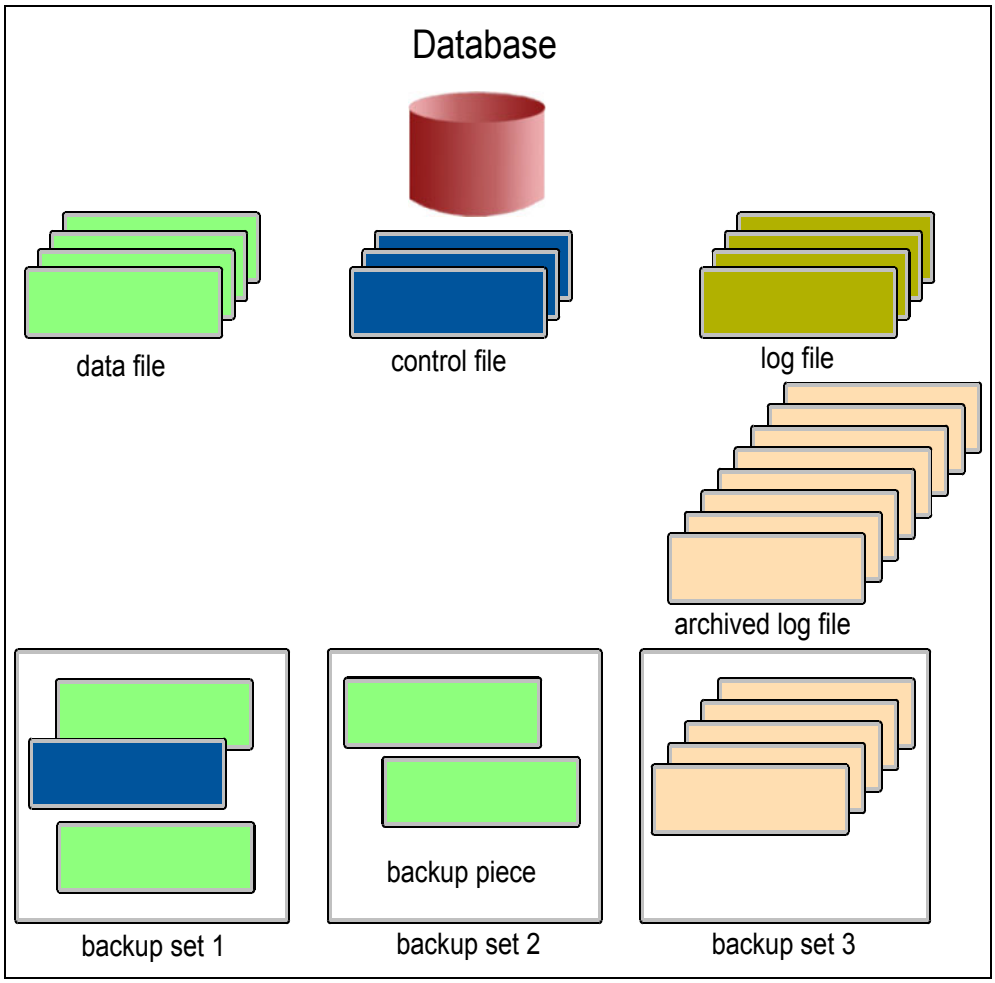

Figure 7-13 Backup sets

Full backup

A full backup is a non-incremental backup of one or more data files. A full backup has no effect on incremental backups and is not considered to be part of the incremental strategy.

If the database is in ARCHIVELOG mode, you can choose to do full backup while the database is online or offline. If the database is in NOARCHIVELOG mode, the database must be closed by a clean shutdown. Full backups can be taken of the following items:

- Data files
- Table spaces
- Databases
- Control files
- Archive logs

Whole database backup

A whole database backup set contains the control files and all database files that belong to that database. Whole database backups do not require the database to be operated in a specific archiving mode. They can be taken whether a database is operating in ARCHIVELOG or NOARCHIVELOG mode. If the database is in ARCHIVELOG mode, you can choose to back up the database while it is open or closed. If running in NOARCHIVELOG mode, the database must be shut down first. There are two types of whole database backups:

- Consistent whole database backup

  A consistent whole database backup is a backup set where all files within it are consistent to the same point in time. A consistent whole database is the only valid backup for databases running in NOARCHIVELOG mode. The only way to take a consistent whole database backup is to shut down the database cleanly and take a backup while the database is offline.

- Inconsistent whole database backup

  An inconsistent whole database backup is a backup of an online database. It is inconsistent because portions of the databases may have been modified and written to disk during the backup process. The database must be in ARCHIVELOG mode to run an inconsistent backup.

Incremental backup

RMAN can incrementally back up databases at the individual block level. An incremental backup is a backup of one or more data files and contains only those blocks that have been modified since a previous backup at the same or lower level.

The *multilevel incremental backup* feature allows you to create various levels of incremental backups. Each level is denoted by an integer, with 0 being the lowest backup level. An incremental backup performed at a given level backs up only those blocks that have been modified since the last backup at the same or lower level.

An incremental backup can be performed on these items:

- Individual data files
- Table spaces
- The entire database

Incremental backup of control files or archived logs is not supported. There are two types of incremental backups: *non-cumulative* and *cumulative*.

Non-cumulative incremental backup

A non-cumulative incremental backup backs up only those blocks that have changed since the previous incremental backup at the same or lower level. This is the default mode of operation for incremental backups. A level 0 backup backs up all blocks that contain data. It performs the same backup as a full backup. A level 0 backup is required for subsequent incremental backups at other levels. An incremental backup at a level greater than level 0 backs up only those blocks that have changed since a previous incremental backup at the same or lower level. The size of the backup depends on the number of blocks modified.

Figure 7-14 on page 182 illustrates part of a monthly cycle of non-cumulative incremental backups. The cycle is based on backup levels 0, 1, and 2. A weekly backup is performed at level 0 on Sunday, incremental backups level 2 are performed on Monday through Wednesday and on Friday and Saturday, and weekly incremental backups at level 1 on each Thursday.

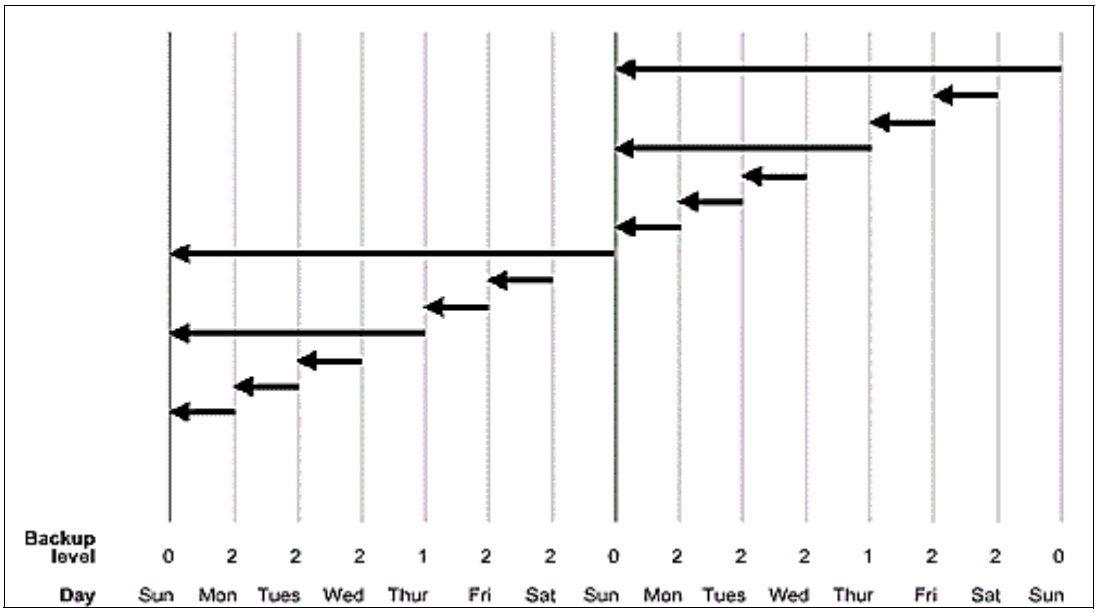

Figure 7-14 Non Cumulative Incremental Backups

Cumulative incremental backup

A cumulative incremental backup at level *n* contains only blocks that have been changed since the most recent backup at level *n* - 1 (minus 1) or lower. Cumulative backups require more storage space than differential backups, but they are preferable during a restore operation because only one backup for a given level is needed. See Figure 7-15.

Note that the first incremental backup must be a level 0 backup that contains all used blocks. A cumulative backup at level 2 will contain all blocks changed since the most recent level 1 backup, copying all blocks changed since the base level 0 backup only if a previous level 1 is unavailable. In contrast to a cumulative backup, a differential backup at level 2 determines which level 1 *or* level 2 backup occurred most recently and copies all blocks changed since that backup.

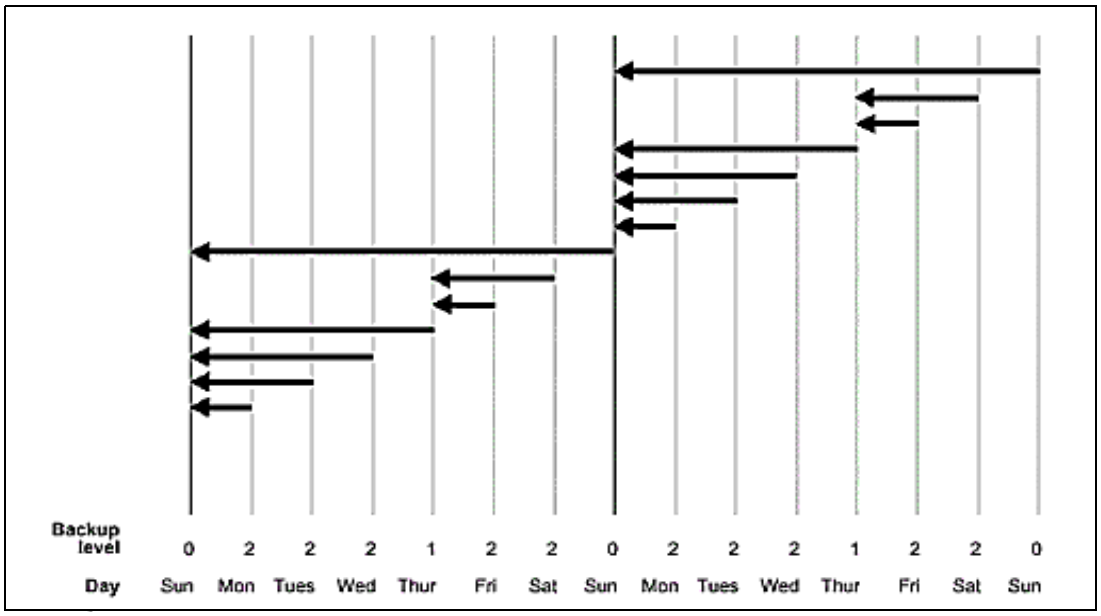

Figure 7-15 Cumulative Incremental Backups

Image copies

An image copy is a single file (data file, archive log, or control file) that can be used as-is to perform a recovery. It is similar to an operating system copy of a single file, except that it is produced by an Oracle server process which performs additional tasks such as validating the blocks in the file and registering the copy in the control files. An image copy can be done only to disk. Image copies are not discussed in detail, because RMAN does not send image copies to Tivoli Storage Manager (they are always stored locally on disk).

Restore operations

Recovering a database, table space, or data file is a two-stage process. The object is first restored and then it is recovered. The restore process restores the necessary full or incremental level 0 backups. Incremental backups at levels greater than 0 are not restored. These are restored during the subsequent recovery process. By default, the objects are restored to their original location as specified in the recovery catalog. An alternative location can also be specified if required. RMAN uses the recovery catalog to select the most current backup sets for use in the restore. The mode in which the database is running determines whether a *consistent* or *inconsistent* restore operation can be performed.

Consistent database recovery

If the database is running in NOARCHIVELOG mode, it can be restored only from a whole database backup. The control files and all data files are restored from a consistent backup. The database must be mounted but not open during the restore operation. After a consistent restore, the database can be opened without performing any recovery. Any updates to the database after the backup are lost.

Inconsistent database recovery

If the database is running in ARCHIVELOG mode, a subset of the database such as a data file, table space, or the entire database can be restored to the most current state or to a specific point in time. This type of restore is inconsistent because the database cannot be started after the restore. The restore operation must be followed by a recovery operation. After the recovery is finished, the database can then be opened.

The following objects can be restored:

- Database
- Table space
- Data file
- Control file
- Archive log

If the control files are lost, they must be restored before other restore operations can be performed. Only after restoring the control files can the target database be mounted and the other restore operations started. A restore operation is typically followed by a recovery operation.

Recovery Operation

Recovery is a process whereby a restored file is made available, either to the most current state or to a specific point in time. Once the data files are restored, they have to be made consistent with each other. The recover command is used to perform media recovery and to apply incremental backups. During the recovery process, RMAN automatically restores the archived redo logs as required. If RMAN has a choice between restoring incremental backup sets or applying redo logs, it always uses the incremental backups. If overlapping levels of incremental backup are available, the lowest level of incremental backup, the one covering the longest period of time, is chosen automatically.

The following objects can be recovered:

- Individual data files

  One or more data files can be recovered. The target database must be started and mounted. If the target database is open, the data file must be offline.

- Table spaces

  All data files in a table space can be recovered or a table space can be recovered to a previous point in time. The target database must be started and mounted. If the target database is open, the table space must be offline.

- Entire database

  The entire database can be recovered under the following circumstances:

  – A media failure has damaged the entire database.
  – The entire database must be recovered to a previous point in time.
  – The control files have been lost.

For database recovery, the database must be started but not open. Archive log backup sets are restored as needed to perform a recovery. They are restored to the current archive log destination as specified in the `init.ora` file

LAN-free data transfer

Data Protection for Oracle supports backup and restore operations in a LAN-free environment. This environment shifts the movement of data from the communications network to a storage area network (SAN). Data moves over the SAN to a SAN-attached storage device by the Tivoli Storage Manager Storage Agent. Running Data Protection for Oracle in a LAN-free environment avoids constraints of the network and decreases the load on the Tivoli Storage Manager server, allowing the server to support a greater number of simultaneous connections.

Before enabling LAN-free support, you must install the Tivoli Storage Manager Managed System for SAN Storage Agent on the same system as Data Protection for Oracle. If your database is considered very large, see 7.6, “Big data: Structured, very large databases” on page 231.

See the IBM Tivoli Storage Manager for SAN for your operating environment for more information about LAN-free requirements.

Using data deduplication with Data Protection for Oracle

You can use data deduplication with Data Protection for Oracle to reduce the amount of redundant data that is backed up to the Tivoli Storage Manager server.

Considerations

Consider the following information:

- The `deduplication` and `enablelanfree` options are mutually exclusive. Therefore, you can use either one option or the other, but not both options together.
- The `deduplication` and `enableclientencryptkey` options are also mutually exclusive. Therefore, you can use either one option or the other, but not both options together.

- A local deduplication cache is an optimization that can reduce network traffic between the Tivoli Storage Manager server and the client. Client-side data deduplication can occur with or without it. Do not use the deduplication cache with Data Protection for Oracle for the following reasons:
  - The cache cannot be used when multiple processes, such as concurrent backups or Tivoli Storage Manager API applications, transfer content concurrently. Data Protection for Oracle backup operations that use multiple channels use multiple processes.
  - It is possible that the client deduplication cache can become out of sync with the server-deduplicated disk storage pool. This state can be the result of object expiration, file space deletion, and overflow to an associated tape storage pool. When the client cache contains entries that are no longer in the Tivoli Storage Manager server deduplicated pool, the cache is reset and the backup operations fails. The Tivoli Storage Manager API does not attempt the backup again.
  - When Tivoli Storage Manager server expiration or a similar process that removes deduplicated data extents runs concurrently with a deduplicated backup, the backup might fail. Backup operations with client-side deduplication enabled fails with the following message:

    ```
    Return code=254
    Error message: ANS7899E The client referenced a deduplicated extent that
    does not exist on the Tivoli Storage Manager server.
    ```

More information about deduplication is in 3.3.3, “Server-side data deduplication” on page 57.

Additional information

A description of the offering is provided, along with links to installation resources and additional informations about the product.

System requirements

The detailed system requirements for Tivoli Storage Manager for Databases and other maintenance updates are available from the All Requirement Documents web page:

http://www.ibm.com/support/docview.wss?uid=swg21218747

Installing Data Protection for Oracle

For information about where to download Data Protection for Oracle, see the download web page:

http://www.ibm.com/support/docview.wss?uid=swg24033228

Documentation updates

For information that was unavailable at the time of publication, see the following web page:

http://www.ibm.com/support/docview.wss?uid=swg27022321

7.5.2 Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL

Tivoli Data Protection for SQL Server enables you to perform online backups and restores of Microsoft SQL Server databases to Tivoli Storage Manager Server storage using either command-line or graphical interfaces on Windows servers, in a stand-alone or clustered environment. Use Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server to protect business-critical databases with automated tasks, utilities, and interfaces

Challenges and benefits the solution addresses

Tivoli Data Protection for SQL Server helps protect and manage Microsoft SQL Server data to more easily do these tasks:

- Back up any SQL Server database to any Tivoli Storage Manager server.
- Perform full and transaction log backups and restores of SQL Server databases.
- Perform backups with an expanded range of options such as differential, file, and group operations.
- Perform operations from multiple Microsoft SQL Server instances on the same machine as Data Protection for SQL Server. You can access only one SQL server per execution of Data Protection for SQL Server from either the command line or GUI.
- Schedule automated backups.
- Perform expanded restore operations on backup objects such as relocating, restoring to named marks, and partially restoring full backups.
- Restore database backups to a different SQL server. Data Protection for SQL Server can restore database backups that were performed on either 32-bit or 64-bit versions of Microsoft SQL Server. See the Microsoft documentation regarding which combinations are supported by Microsoft.
- Retain with a backup the information needed to re-create or move SQL Server databases or files - such as sort order, code page, and Unicode information - and filegroup and logical and physical file names. The meta object information is retained on the IBM Tivoli Storage Manager separately from the backup data objects.
- Use the restorefiles command to restore VSS backups to flat files without involving the SQL Server.
- Enhanced Statistics.
- Back up and restore AlwaysOn Availability Databases (SQL Server 2012 environment) to provide high availability and disaster recovery at the SQL server database level.
- Back up and restore any SQL Server database by using any node in an AlwaysOn Failover Cluster instance ((SQL Server 2012 environment)) to provide high availability and disaster recovery at the SQL server instance level.

Solution architecture

Data Protection for Microsoft SQL Server performs online backups and restores of Microsoft SQL databases to Tivoli Storage Manager storage or local shadow volumes. You can perform backups and restores using a command-line or graphical user interface (GUI).

Data Protection for Microsoft SQL operations use the Tivoli Storage Manager application programming interface (API) to communicate with the Tivoli Storage Manager server, and use the SQL API to communicate with SQL Server. In addition to using these APIs, Data Protection for Microsoft SQL VSS operations use the Tivoli Storage Manager backup-archive client (VSS Requestor) and Microsoft Volume Shadow Copy Service to produce an online snapshot (point-in-time consistent copy) of SQL data that can be stored on local shadow volumes or on Tivoli Storage Manager server storage.

Data transfer between client and server can be done via LAN (TCP/IP) or SAN (LAN-free backups).

Solution description

Data Protection for SQL provides two methods of backing up SQL Server data:

- Legacy backup
- VSS backup

Legacy backup overview

A Legacy backup is a specialized API backup that functions with the Microsoft SQL Server storage engine, as show in Figure 7-16. This is the type of backup provided by previous releases of Data Protection for SQL Server.

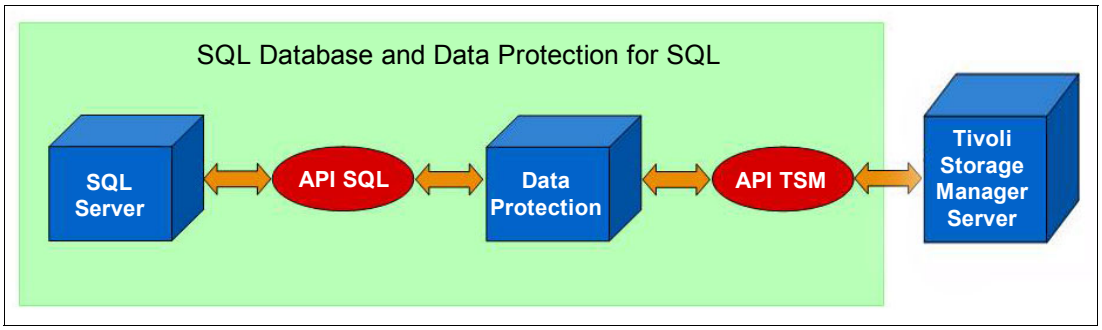

Figure 7-16 Data Protection for SQL Legacy backup communications

A Legacy backup creates a copy of all or part of a SQL database or logs on Tivoli Storage Manager storage media.

Data Protection for SQL provides selection mechanisms and the logic that are required to back up and restore SQL data. When you initiate a Legacy backup operation, Data Protection for SQL completes the following actions:

1. Begins a session with a Tivoli Storage Manager server using the Tivoli Storage Manager API and information contained in a client options file.
2. Starts a session with the SQL Server using the SQL-SMO interface.
3. Instructs the SQL Server using the SQL VDI interface to begin a backup of the selected database objects.
4. Receives data from the SQL Server and sends it to the Tivoli Storage Manager server.
5. Informs the SQL Server that the backup is complete.
6. Ends the Tivoli Storage Manager server and SQL Server sessions.

When a backup is performed, Tivoli Storage Manager server retains information about the SQL Server and database. This information is available for query and restore operations after the backup is completed. The information about the names and sizes of the database file groups and files is stored along with the database data, as a sub-object. This sub-object is referred to as *metadata*.

The following characteristics are true of Legacy backups:

- Full, copy, incremental, differential, and database copy types are supported.
- Backup granularity is at the database level.
- Backups are stored on IBM Tivoli Storage Manager server storage.
- Backups are managed through IBM Tivoli Storage Manager server policy.
- Backups can be performed in a Microsoft Cluster Server (MSCS) environment.
- Backups provide Microsoft SQL Server database integrity check functionality

VSS backup

A VSS backup uses Microsoft Volume Shadow Copy Service technology to produce an online snapshot (point-in-time consistent copy) of SQL data. A VSS backup means that the SQL Server is not in backup mode for an extended period of time. The length of time to perform the snapshot is usually measured in seconds, not hours. In addition, a VSS backup allows a snapshot of large amounts of data at one time because the snapshot works at the volume level. VSS backups can be stored on Tivoli Storage Manager server storage or local VSS shadow volumes. VSS backups stored locally on VSS shadow volumes are directly accessible by the SQL system (as long as sufficient space is available for the snapshot).

For local VSS backups you must have a licensed version of IBM Tivoli Storage FlashCopy Manager installed on your system. We cover this option in 7.5.6, “Tivoli Storage FlashCopy Manager” on page 222.

When performing VSS backups and moving data to Tivoli Storage Manager server storage, sufficient space on local snapshot volumes is still temporarily required to hold the snapshot. For SQL data backed up to Tivoli Storage Manager server storage, the SQL data on the snapshot volume is sent to the Tivoli Storage Manager server. After the data transfer to the server is complete, the snapshot volume is made available for reuse.

The Tivoli Storage Manager backup-archive client serves as the VSS requestor that communicates with VSS to access the SQL data to create shadow copies of SQL databases. Data Protection for SQL serves as a front end for VSS backup operations. Because of the role that the backup-archive client performs as the VSS requestor, features such as LAN-free backup, client-side deduplication, data encryption, and data compression require that options related to these features be specified in the backup-archive client options file for VSS operations.

The VSS architecture is depicted in Figure 7-17.

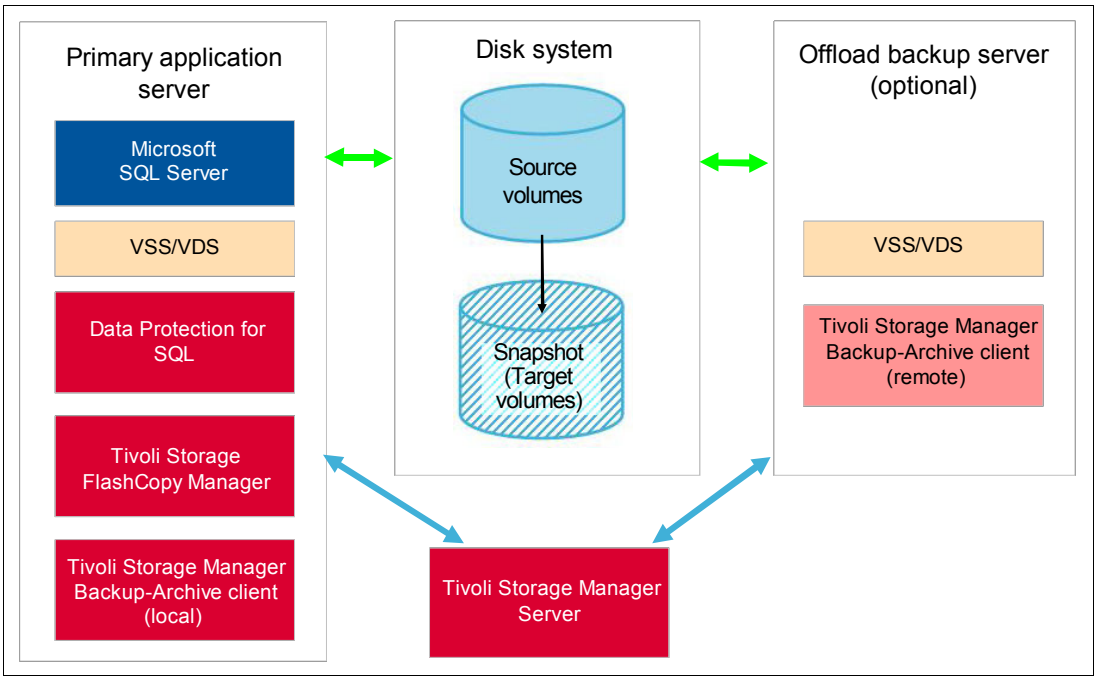

Figure 7-17 General VSS architecture

The following characteristics are true of VSS backups:

- Full and copy-only full (COPYFull) VSS snapshot backups and full and copy-only full VSS offloaded snapshot backups are supported. Incremental, differential, and transaction log backup types are not supported.
- Backup granularity is at the database level only.
- Backups are managed through Tivoli Storage Manager policy.
- Backups can be stored on local shadow volumes, Tivoli Storage Manager server storage, or both locations.
- Different policy settings can be defined for each storage location and backup method.
- Backups to Tivoli Storage Manager server storage can be offloaded to an alternate machine, to reduce the workload on the production servers.

Offloaded VSS backups

An offloaded backup uses another machine to move the data to the Tivoli Storage Manager server. This type of backup shifts the backup load from the production machine to another machine. An offloaded VSS backup requires a VSS hardware provider that supports transportable shadow copy volumes is installed on the production and secondary machines. Offloaded VSS backups require a Tivoli Storage FlashCopy Manager license.

Use scenarios

Different backup strategies are available depending on specific requirements regarding network traffic, backup window and acceptable restore times. Table 7-3 on page 190 summarizes these strategies.

Table 7-3 Strategies defined by backup type

Backup type	Use	Legacy	VSS
Full backup only	This approach is best for SQL databases that are relatively small because it implies that the entire database is backed up each time. Each full backup takes longer to perform, but the restore process is most efficient because only the most recent (or other appropriate) full backup need be restored. This is the appropriate strategy for system databases such as master, model, and msdb due to their normally small size.	X	X
Full plus log backup	A full plus transaction log backup strategy is commonly used when the normal backup window or network capacity cannot support a full backup each time. In such cases, a periodic full backup followed by a series of log backups allows the backup window and network traffic to be minimized. The full backups can be done during low usage times when a larger backup window and increased network traffic can be tolerated. The restore process becomes more complex, however, because a full backup, as well as subsequent log backups, must be restored. It is also possible to do a point-in-time restore to restore a transaction log to a specified date and time. You can apply Legacy log backups after a full VSS backup has been restored. To do this, you must leave the database in a recovering state by specifying `/recovery=no` in the CLI or by making sure that the Recovery option in the GUI Restore Databases or Restore Groups/Files is *not* selected when restoring the VSS backup.	X	X
Full plus differential backup	This strategy can be used between full backups. A differential database backup can save both time and space. Space is saved because the backup consists of only the changed portions of a database since the last full backup (it is cumulative). Time is saved because you can avoid applying all individual log backups within that time to the operation. This applies to restore operations too; only the last differential backup (latest version) must be restored. Although VSS supports full backups only, Legacy differential backups can be applied to the VSS full backup. To do this, you must leave the database in a recovering state by specifying `/recovery=no` in the CLI or by making sure that the Recovery option is *not* selected when restoring the VSS backup.	X	X
Full plus differential plus log backup	This strategy allows for a faster restore scenario by reducing the number of transactions that may need to be restored and applied. If, for example, a full Legacy or VSS backup is done weekly, a differential nightly, and a log backup every four hours, the restore would involve the full backup, a differential, and at most five log backups. However, simply a full plus log backup scheme on the same cycle could require a full plus up to 41 log backups to be restored (six days times six log backups per day plus up to five backups on the day the full backup was done). Although VSS supports full backups only, Legacy log backups and Legacy differential backups can be applied to the VSS full backup	X	X
File or group backups	When a group is created on the SQL Server, database files are identified with that group. The group used for the group backup is dependant on the group to which the database files are defined. Use a file backup strategy when it is impractical to backup an entire database because of size and accompanying time and performance issues. When performing restore operations for a file or file group, provide a separate backup of the transaction log. File or group options can also save both backup and restore time in cases when certain tables or indexes have more updates than others and need to be backed up more often. It is time-effective to place such data in their own file group or files and then back up only those items.	X	

Using VSS backups and Legacy backups together

Using VSS backups and Legacy backups together can implement a highly-effective backup solution for Data Protection for SQL data. Although VSS supports only full backups, Legacy differential and Legacy log backups can be applied after a full VSS backup has been restored.

Use the following preferred practices:

- Legacy and VSS backups to Tivoli Storage Manager server storage are usually dictated by time, not versions.
- Backups to local shadow volumes are usually dictated by versions because of space limitations and provisioning of VSS storage.
- When running VSS operations, have at least 200 MB of free disk space on your Windows system folder. You can set this location with the **`vssaltstagingdir`** parameter.
- Continue to schedule and perform Legacy backups in your strategy.

Backups of availability databases

Data Protection for SQL backs up and restores each availability database as a single object, regardless of which availability replica is used for backup and restore operations.

An AlwaysOn Availability Group can contain a set of primary databases and one to four copies of the set of primary databases, called secondary databases.

Databases in an availability group are called *availability databases*, and they fail over together as a group.

An AlwaysOn Availability Group requires SQL Server instances on Windows Failover Cluster nodes. An availability group can have several replicas. For example, availability group 1 might have replicas node1, node2, and node3.

A cluster node might be an availability replica for one or more availability groups. For example, node1 might be a replica for availability group 1 and another availability group. For the secondary replica, read-only is an option that can be set at the availability group level.

The AlwaysOn Node is used to manage backups of availability databases. When working in a Tivoli Storage Manager environment, the AlwaysOn Node should be common in a Windows Failover Cluster. This presence enables the management of backups of an availability database in a single location, regardless of the replica that is used to perform the backup.

The following types of VSS backup operations are supported:

- Full VSS backups of the primary availability replica
- VSS copy-only full backups of availability replicas

The following types of Legacy backup operations are supported:

- On the primary replica, legacy full, differential, file, set, group, and log backups are supported.
- On the secondary replica, legacy full, file, set, group, and log backups are supported.
- VSS and legacy copy-only full backups, legacy copy-only file, set, or group backups, and legacy copy-only and normal log backups are supported.

For all backup operations of secondary availability replicas, the secondary replicas must be in the synchronized or synchronizing state.

Restoration of availability databases

Depending on how availability databases are backed up, Legacy restore and VSS restore operations are available to restore the availability databases on primary or secondary availability replicas.

- Legacy restore

  You can restore an availability database on either a primary or secondary replica.

  During the restore process, the restored database is removed from the availability group. When a database is removed from the availability group, the database becomes a local database on that replica. The database is restored as a local database. After this restore is complete, manually add the database back to the availability group. However, before adding the database to the availability group, verify that the data on all replicas is transactionally consistent. For example, the database and any log files must be restored so the files are all at the same level. After verifying the data is transactionally consistent, the database can be added to the availability group.

- VSS restore

  Because of a SQL Server limitation, you cannot restore a VSS backup to an alternative SQL server instance. Therefore, VSS backups must be restored to the same SQL server instance where the snapshot was taken.

LAN-free environment (Legacy and VSS)

Running Data Protection for SQL in a LAN-free environment if you are equipped to do so avoids network constraints.

- For Legacy backups, specify `enablelanfree yes` in the Data Protection for SQL options file.
- For VSS backups, specify `enablelanfree yes` in the DSMAGENT (VSS Requestor) `dsm.opt` file only.

For information about setting up a LAN-free environment, see *IBM Tivoli Storage Manager for SAN for Windows Storage Agent User's Guide*:

http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1/topic/com.ibm.itsm.sta.doc/welcome.html

Hardware and software requirements

Details of the hardware and software requirements change over time because of maintenance updates and the addition of operating system, application, and other software currency support. For the most current requirements, see the Hardware and Software Requirements technote that is associated with your level of software. This technote is available from the following website:

http://www.ibm.com/support/docview.wss?uid=swg21218747

Before installing Data Protection for SQL, ensure that your system meets the minimum software and operating requirements. Details of the software and operating system requirements for Data Protection for SQL can change over time. For current software requirements, see the Tivoli Storage Manager for Databases - All Requirements Documents website:

http://www.ibm.com/support/docview.wss?uid=swg21218747

7.5.3 Tivoli Storage Manager for Mail: Data Protection for Domino

The Tivoli Data Protection for Lotus Domino application client provides an integrated solution for performing full backup and restore operations on Lotus Domino databases and database templates. It is a client application that provides full backup of online databases and restore of full databases to the original or different location. Tivoli Data Protection for Lotus Domino also archives the transaction log extents of a Domino server and retrieves the appropriate transaction log extents for the recovery of databases if archive transaction logging is enabled on the Domino server.

Tivoli Data Protection for Lotus Domino is not intended as a substitute for the standard Tivoli Storage Manager backup-archive client. Tivoli Data Protection for Domino cannot be used to back up or restore any non-database data, such as Notes ID files, or `notes.ini`, or any other system configuration files. Those files need to be backed up by the Tivoli Storage Manager backup-archive client. Therefore, the two client types work together to provide full data protection for your Notes environment. The Tivoli Data Protection for Lotus Domino application client and the Tivoli Storage Manager backup-archive client can run simultaneously on the same Domino server, however, they are totally separate clients as far as the Tivoli Storage Manager server is concerned.

Client benefits

IBM Tivoli Storage Manager for Mail protects data on email servers running IBM Lotus Domino. This software module for IBM Tivoli Storage Manager enables data protection of your mail databases while they are online. It automates data protection, enables hot backups without shutting down the application and improves data restore performance. Tivoli Storage Manager for Mail does these tasks:

- Uses APIs provided by mail application vendors to perform online hot backups and improve restores without shutting down the mail application.
- Helps protect the growing amount of new and changing data that should be securely backed up to help maintain continuous availability.
- Exploits the Lotus Domino transaction logging feature, which enables the capture and logging of database changes, resulting in less-frequent full backups.
- Backs up Lotus Domino NSF databases.
- Maintains multiple backup versions of Domino databases.
- Archives Lotus Domino transaction log files when archival logging is in effect.
- Restores backup versions of a Lotus Domino database and apply changes since the last backup from the transaction log.
- Restores Domino databases to a specific point in time.
- Restores one or more archived transaction log files.
- Expires database backups automatically based on version limit and retention period.
- Expires archived transaction log files when no longer needed.
- Automates scheduled backups.
- Restores Domino databases to an alternate server or partition.
- Accesses Data Protection for Domino remotely using the Tivoli Storage Manager Web client.
- Accesses Data Protection for Domino using the client GUI based on Oracle Java.
- Accesses Data Protection for Domino using the command-line interface.

Solution architecture

Data Protection for Domino component communicates with a Tivoli Storage Manager server using the Tivoli Storage Manager API. Tivoli Storage Manager communicates with a Domino Server using the Domino API, as shown in Figure 7-18.

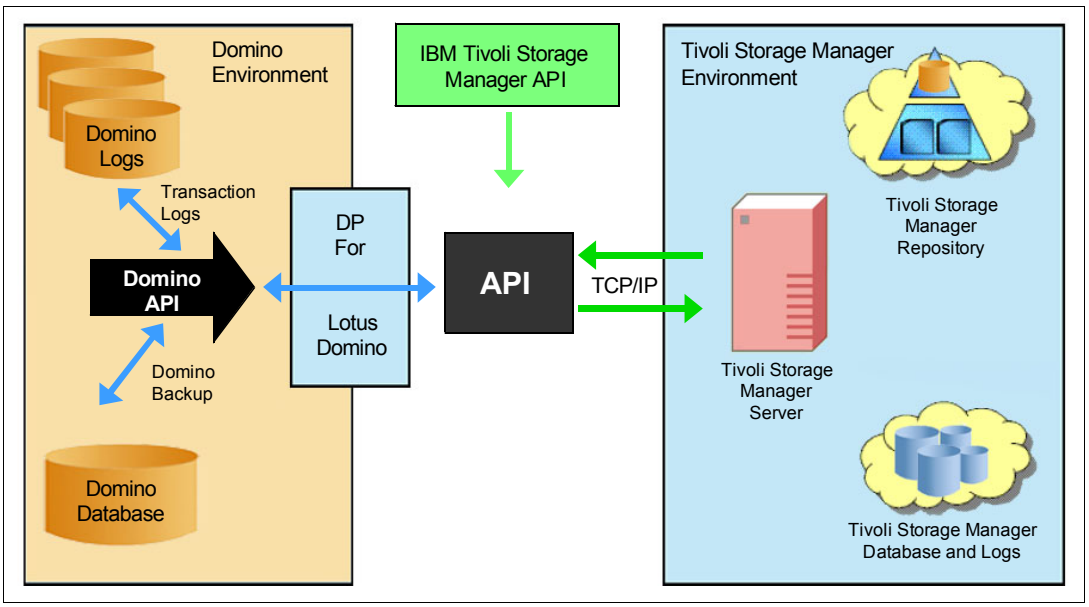

Figure 7-18 Tivoli Storage Manager for Mail - Data Protection for Domino

The backup and recovery API in Domino provides the capability to perform online full backups of individual databases and archives of the transaction log when archival logging is in effect. When archival logging is used on the Domino server, it archives the transaction log files and retrieves them as required for a database recovery. Database backups are archived. The Tivoli Data Protection for Lotus Domino application client provides a command line interface for performing backups and restores. The application client commands are issued from a command prompt. On Windows, Tivoli Data Protection for Lotus Domino also provides a GUI, which supports most of the functions of the application client; transaction log files are stored in Tivoli Storage Manager storage. A transaction log captures database changes for logged databases, so full database backups are not required as frequently.

Data Protection for Domino provides support for two types of Domino databases: NSF and DB2.

Domino NSF database backup and transaction log archive solution

This section describes the concepts associated with Data Protection for Domino back ups of Domino databases and transaction logs. The backup and recovery API in Domino provides the capability to perform these tasks:

- Online full backups of individual databases
- Archives of the transaction log when archival logging is in effect

Domino server transaction log

Updates to a logged database are recorded in the Domino server transaction log, so full database backups are not required as frequently. Changes to a database since the last full backup can be applied from the transaction log after the backup is restored from the last full backup. Enabling transaction logging for all databases on a Domino server is not required, so the backup process must handle both logged and non-logged databases. Domino allows the active transaction log to be backed up too.

Transactions recorded in the transaction log are keyed by a Database Instance Identifier (DBIID), which is unique for each database on a Domino server. The DBIID must match that of a restored database for transactions in the log to be applied to the database. The most common reason for a DBIID to change is compaction of the database to reduce file size. Therefore, whenever the DBIID of a database changes, a full backup must be taken so that subsequent updates (which are recorded in the transaction log) can be applied to a restored backup of that database. Transactions recorded since the DBIID change cannot be applied to prior backups of that database because the DBIID will not match. See your Domino server documentation for more information about the DBIID and when it can change.

Types of NSF backup and archive logs

Data Protection for Domino provides two types of database backups and an archive log function:

- Incremental backup

  An incremental backup provides a conditional backup function that performs a full online backup of Domino databases under the following conditions:

  - The database is not excluded in the Tivoli Storage Manager include-exclude options file (standard include and exclude processing is supported).
  - The database is not logged and was modified since the last active backup image for that database. Both data and non-data modification dates are checked. If either is different from that of the active backup, the database is backed up.
  - Archival logging is in effect and the DBIID of a logged database changed. If the DBIID has not changed, then logged databases are not backed up (the changes are captured in the transaction log backups). In this case, periodic selective backups of all logged databases should be done to refresh the active backup images. This reduces the number of transaction logs to be applied during a recovery.

    When circular logging is used on the Domino server or when logging is disabled on the Domino server, transaction log files are not archived.
  - The database is new or newly included in the backup and an active backup image does not exist on the Tivoli Storage Manager server.

  The incremental command includes a function that determines if active backup database copies exist on the Tivoli Storage Manager server that are deleted from the Domino server or excluded from backup. If so, they are marked inactive so that automatic expiration of these backup copies can occur according to defined Tivoli Storage Manager management class parameters for backup files.

- Selective backup

  A selective backup unconditionally performs a full online backup of the specified Domino databases, unless they are excluded from backup through exclude statements within the Data Protection for Domino options file dsm.opt.

- Archive log

  An archive log stores filled transaction log files on the Tivoli Storage Manager server so that space allocated to these files can be reused by the Domino logger. The archivelog command is available when transaction logging on the Domino server is enabled in archival mode. Filled transaction log files must be archived frequently enough to ensure the transaction log never fills completely and stops the Domino server. Transaction log files stored on the Tivoli Storage Manager server are automatically restored as needed for a database recovery. Archived transaction log files are retained on the Tivoli Storage Manager server as long as a database backup exists that needs these log files for a complete recovery.

When circular or linear loop logging is used on the Domino server (or when logging is disabled on the Domino Server), transaction log files are not archived.

Expiration of NSF archived transaction log files

This section describes concepts associated with expiring archived transaction log files. The **`inactivatelogs`** command expires transaction log files from backup storage. There is a single shared transaction log for all logged databases on a Domino server. Thus log files cannot be deactivated (and allowed to expire) until all databases that require that log file for recovery are inactive. This command queries the database backups on the Tivoli Storage Manager server to determine which log files are required by any active database backup. This command also deactivates log files that are no longer required (because the database backups were deactivated). Run the **`inactivatelogs`** command after a full database backup is completed to deactivate the transaction logs as the database backups requiring them are deactivated.

Restore NSF databases

A Domino database recovery can involve restoring several transaction log files in addition to the database backup file from the Tivoli Storage Manager server, depending on the backup strategy you choose. The function to restore database files is separate from the function that applies updates from the transaction log. This allows you to restore database files separately while transaction logs are processed for all restored databases; it avoids restoring the same transaction log files multiple times. Restoring and updating a database with current changes from the transaction log is a two-step process, implemented by the **`restore`** and **`activatedbs`** commands.

Domino database restore: Step 1 of 2

Restore is the first step of a two-stage recovery process. This function restores a single database or group of databases from Tivoli Storage Manager storage to the Domino server. You can restore the database to a different database file name or to a different Domino server. You can also restore a group of databases to a different directory and preserve existing file names. In addition, if you specify a point in time on the restore command, the most recent backup version prior to that time is restored. To restore a database without applying updates from the transaction log, the two steps can be combined into one step by specifying **`/activate=yes`** during the **`restore`** command.

Domino database activation: Step 2 of 2

This is the second step of the two stage recovery process. This function brings restored databases online for use by the Domino server. You can optionally apply transactions from the transaction log to update the database. Transactions can be applied up to a specific point in time or up through the most recent changes recorded in the transaction log. If archival logging is in effect, Data Protection for Domino automatically restores archived transaction log files as needed. The Domino server provides an alternate restore path feature that allows you to specify the directory where transaction logs are restored. You can use this feature with the **`activatedbs`** command.

Restore of archived transaction logs

This function allows a single, archived transaction log file to be restored independently of a routine database restore. Restoring a single, archived transaction log file assists with disaster recovery operations. By retrieving the most recent archived log file, it is possible to rebuild the Domino transaction log control file. This allows archived transaction log files to be used to recover restored database backups to a more current state, even after a loss of the active transaction log. More than one archived transaction log file can be restored at a time.

Restore at document level

Data Protection for Domino restores Domino databases at the database level. To restore a document in a database, the entire database must first be restored and the document copied.

A database can be restored to the production server under a temporary name and the desired document can be copied to the appropriate database. If, for performance reasons, the production server cannot be used in the restore process, the database can be restored to an alternate server and copied to the production server. You should perform alternate server restores when possible to reduce demands on the production Domino server. Alternate server restores can be performed to an alternate partition or to a separate Domino server.

DB2 Tivoli Storage Manager Agent

DB2 provides a Tivoli Storage Manager Agent and a utility program (**db2adutl**) that interfaces with the DB2 Recovery API to manage Tivoli Storage Manager objects created on the DB2 server. Data Protection for Domino uses the DB2 Tivoli Storage Manager Agent through the DB2 Recovery API to back up and restore the Domino DB2 database and DB2 Groups (table space). These Tivoli Storage Manager objects associated with DB2 backups are unique in that only one Tivoli Storage Manager object is created for each backup operation per session. The **db2adutl** program, for example, can be used to expire these objects.

Data Protection for Domino and the DB2 API

Data Protection for Domino uses the DB2 Recovery API to communicate with the DB2 Tivoli Storage Manager Agent to back up DB2 data to the Tivoli Storage Manager server. Configure the DB2 Tivoli Storage Manager Agent to use the same Tivoli Storage Manager node name and to access the same Tivoli Storage Manager server as Data Protection for Domino. This enables the Tivoli Storage Manager objects (created by the DB2 Recovery API) to belong to the same Tivoli Storage Manager node as the objects created by Data Protection for Domino for NSF databases. Specify the desired options file with the DSMI_CONFIG environment variable.

Types of DB2 backups

Data Protection for Domino provides three types of database backups:

- DB2 database backup

  Data Protection for Domino DB2 database backups create a selective backup image that can be used for disaster recovery of the Domino DB2 database or for restoring individual DB2 enabled Notes databases. Only selective backup (db2selective) is provided for DB2 enabled Notes databases.

- DB2 group (table space) backup

  Data Protection for Domino DB2 group backups create a selective table space backup image. This type of backup can only be performed after the DB2 database is enabled for rollforward recovery.

- Full DB2 database and NSF database backup

  Data Protection for Domino can perform a selective NSF database backup and a full Domino DB2 database backup in a single operation.

Expiration of DB2 backups and transaction log objects

This section describes concepts associated with expiring DB2 backup objects and DB2 transaction log files. Data Protection for Domino uses the DB2 Recovery API to access the DB2 Tivoli Storage Manager Agent. When Data Protection for Domino performs a backup, it informs the DB2 Tivoli Storage Manager Agent to back up the DB2 data to the Tivoli Storage Manager server. During backup processing, Data Protection for Domino creates a group of Tivoli Storage Manager objects that describes the contents of each Tivoli Storage Manager

object created by the DB2 Tivoli Storage Manager Agent. Each object describes the type of backup performed and the name of the DB2 enabled Notes databases contained in the backup. The Tivoli Storage Manager group object has a reference to the object created by the DB2 Tivoli Storage Manager Agent. Policy settings are applied to the Tivoli Storage Manager group object. As a result, when a backup version is no longer needed, the objects that are referenced by the Tivoli Storage Manager group object must also be inactivated. These Tivoli Storage Manager group objects can be inactivated by using the **db2inactivateobjs** command. This command displays how to issue the DB2 Tivoli Storage Manager Agent **db2adutl** utility to inactivate these objects. The **db2adutl** utility ensures that information on the DB2 server remains consistent after objects have been inactivated. The Domino DB2 database transaction logs are archived automatically by the DB2 server (using the DB2 Tivoli Storage Manager Agent) to the Tivoli Storage Manager server. The `db2archivelog` command forces a backup of the Domino DB2 database transaction log file. This command can be used to guarantee that the latest updates are available during an alternate DB2 database rollforward to the current time operation.

Note: Because transaction log file names are unique, they will not expire because of version limit.

Archived transaction log files are retained on the Tivoli Storage Manager server as long as a database backup exists that needs these log files for a complete recovery.

Domino DB2 enabled Notes database restore, rollforward, and activation

Concepts associated with a Domino DB2 enabled Notes database three-stage recovery process (restore, rollforward, and activation) are as follows:

1. Restore

   Data Protection for Domino provides the ability to restore a single DB2 enabled Notes database or a group of DB2 enabled Notes databases. A Domino DB2 Group (table space) can be restored from either a full DB2 database backup image or a DB2 table space backup image. Only one DB2 Group can be restored at a time If the DB2 Group is being restored from a DB2 Group backup. The DB2 Group is restored to an alternate DB2 database within the same DB2 instance. If more than one DB2 Group is restored, each DB2 Group must be restored to a different DB2 database. Otherwise, restoring more than one DB2 Group to the same alternate DB2 database will overwrite the previously restored DB2 Group. If the DB2 Group is being restored from a full DB2 backup image, then more than one DB2 Group can be restored to the same alternate DB2 database. A Domino DB2 database can be restored from a full DB2 database backup image to an alternate DB2 database. This makes the individual DB2 enabled Notes databases available for restore. The DB2 database can also be restored directly to the Domino DB2 database. This type of *in-place* restore operation is useful for disaster recovery purposes.

2. Rollforward

   Rollforward is an intermediate step that is required when the Domino DB2 database is enabled for rollforward recovery. This task rolls the Domino DB2 database forward to the specified point in time and marks the rollforward as complete. The DB2 database can be an alternate DB2 database or the Domino DB2 database.

3. Activation

   This is the last step of the three stage recovery process. This function brings DB2 enabled Notes databases online for use by the Domino server. DB2 enabled Notes databases that are restored from a DB2 table space backup image can be activated after first rolling the alternate DB2 database forward to the desired point-in-time. The DB2 enabled Notes database can be restored to a time later than the backup time by applying necessary transaction log files by specifying the /applylogs parameter during the rollforward

operation. The logs are then applied to the alternate DB2 database or to the Domino DB2 database if it is an *in-place* restore. Although the DB2 application automatically archives transaction log files when they become full, the active transaction log files should be archived before starting the rollforward operation to ensure that the latest transactions are available. The necessary logs (from those archived) are automatically restored during the rollforward operation. The DB2 enabled Notes databases are then copied into the Domino DB2 database to their original filename location or to a new filename location. DB2 enabled Notes databases that are restored from a full DB2 database backup image are activated in the same manner as described for activating DB2 enabled Notes databases restored from a DB2 table space backup image. However, DB2 enabled Notes databases that reside on different table spaces can be rolled forward simultaneously if more than one table space is restored from the full backup image.

Use scenarios

Different backup strategies are available depending on the database used with Domino Server.

NSF backup strategy considerations

You can choose different backup strategies depending on your specific requirements regarding network traffic, backup window, and acceptable restore times. Your choice of strategy includes selecting the type of backup commands to use and the type of transaction logging to be done on the Domino server. Data Protection for Domino can back up transaction logs only from a Domino server that has archival logging in effect. Transaction logs cannot be backed up from a Domino server that has circular or linear loop logging in effect.

Archival logging allows transaction log data to be archived on the Tivoli Storage Manager server so that changes to logged databases can be stored on the Tivoli Storage Manager server without having to perform a full backup. This allows a strategy with less frequent full database backups because changes to logged databases are available for restore in the archived transaction log files.

The **`archivelog`** command backs up Domino transaction log files when archival logging is in effect on the Domino server. The command queries the Domino server to determine if any log extents are ready for archiving. If so, the log files are backed up to Tivoli Storage Manager server storage, and the Domino server is notified of their availability for reuse.

In addition, high and low threshold values can be specified as a percentage of the log capacity to control whether log files should be archived when the command is run. This allows the command to be scheduled regularly to protect against a log full condition but to actually do the archive only if the log is getting close to being full.

Consider the following information when choosing a backup strategy:

- When using archival transaction logging, the frequency of **`archivelog`** command use depends on the size of your log and the rate of change for logged databases. Perform archival transaction logging several times per day if you generate a large volume of changes at a rapid rate.
- When a DBIID for a logged database changes, the database cannot be recovered until another backup of that database is performed. The **`incremental`** command detects the changed DBIID. Any changes recorded in the log between the DBIID change and backup are not restored if the original database is lost. The Domino server sends a message to the server console when a DBIID change occurs. It is useful to monitor the server console and perform a backup when a DBIID change occurs.
- When restoring a group of logged databases for which transactions need to be applied, activate them together when possible. This avoids restoring the same transaction log files

multiple times. Restored transaction log files are deleted during a database recover by the Domino server. Activating and applying logs to the database separately requires retransmitting log files for each database.

- Data Protection for Domino provides backup and restore functions for the Domino databases (including template files) and associated transaction logs. However, Data Protection for Domino does not provide a complete disaster recovery solution for a Domino server by itself. There are many other files that are part of the Domino server installation, such as executable files and configuration files. For example, database link files have an nsf extension but are not considered databases and are not backed up by Data Protection for Domino. These files must be recovered in a disaster recovery situation. A comprehensive disaster recovery plan can be achieved using the normal Tivoli Storage Manager backup-archive client for your server platform together with Data Protection for Domino.
- Personal copies (replicas) of Domino databases that are stored on Notes clients (not on the Domino server) are not protected by Data Protection for Domino. You can use the Tivoli Storage Manager backup-archive client on the Notes client platform to back up and restore these files or rely on Domino server replication if you need to recover them.
- To restore an individual Notes document, you must restore the entire database to an alternate name. Choose a time when the document existed for both the **`restore /pit`** and **`activate /applylogs`** commands but before the document was deleted, and then copy the desired document using the Notes client.
- The Tivoli Storage Manager encryption, deduplication and compression functions can be used with Data Protection for Domino. For more information, read the documentation about using the application programming interface documentation, which is available at the IBM Knowledge Center for Tivoli Storage Manager.

Full backups only

The following backup option can be implemented if your network capacity and backup window support regular full database backups:

- Select circular transaction logging.
- Perform regular selective backups.
- Perform occasional Incremental backups to deactivate backup copies of databases that have been deleted from the Domino server.

Each backup takes longer to perform, but the restore process is most efficient because only the most recent (or other appropriate) full backup needs to be restored.

> **Note:** You can apply updates to the restored database from the transaction log if the log has not wrapped since the backup was performed. If the log has wrapped, the attempt to apply logs fails

Full backup plus transaction log archives

It is often not practical to back up entire databases with each regular backup for large Domino installations. Archival logging captures changes to all logged databases in the archived transaction log files. This enables you to perform full database backups less frequently and reduce burdens on network and storage resources. To implement this strategy, use these steps:

- Select archival transaction logging.
- Perform regular log archives using the archivelog command. This ensures the log does not fill and captures changes to logged databases.

- Perform regular Incremental backups. This does not back up logged databases unless the DBIID has changed.
- Perform occasional Selective backups of all logged databases. This reduces the number of transaction log files to be processed during a restore.
- Issue the archivelog command (following Selective backups) to allow nonessential transaction log files to expire.

The **`archivelog`** command captures changes to all logged databases in between full backups of selected databases. To restore a database to its most recent state, restore the most recent database backup and specify **`/applylogs`** when activating the restored database. This automatically restores the necessary archived transaction log files so that updates for the database can be applied.

DB2 enabled Notes database backup strategy considerations

You can choose different backup strategies depending on your specific requirements regarding network traffic, backup window, and acceptable restore times. Your choice of strategy will include selecting the type of DB2 backup commands to use.

> **Note:** The DB2 commands do not return information about whether a backup to a Tivoli Storage Manager server was compressed, encrypted, sent LAN-free or deduplicated.

Some strategies you can employ are described next.

Full DB2 database backups only

This backup strategy can be followed when the Domino DB2 database is enabled for rollforward recovery:

- Perform full DB2 database backups on a regular basis.
- Routinely inactivate (and delete) DB2 objects from the Tivoli Storage Manager server that are no longer needed.

  A full DB2 database backup completes quicker and requires less storage space than DB2 Group backups. However, DB2 enabled Notes databases cannot be restored to a specific point in time because the database is not enabled for rollforward recovery and requires less storage space than backing up all the DB2 Groups individually.

Full DB2 database backups plus DB2 Group backups

This backup strategy can be followed when the Domino DB2 database is enabled for rollforward recovery:

- Perform full DB2 database backups on a regular basis.
- Perform DB2 Group backups on a regular basis in between full DB2 database backups. Note that only those DB2 Groups with the strictest restore time requirements should be backed up.
- Maintain a complete set of transaction log files to a specified point-in-time. DB2 automatically archives the transaction logs when the DB2 database is enabled for rollforward recovery.
- Routinely inactivate (and delete) DB2 objects from the Tivoli Storage Manager server that are no longer needed.

To restore a DB2 enabled Notes database to its most recent time, first select the most recent backup from a DB2 Group backup or from a full DB2 database backup that contains the DB2 enabled Notes database (if available). If the most recent DB2 Group backup is not available, restore the DB2 Group from the most recent full DB2 database backup. This type of restore is

to an alternate DB2 database. Rollforward the DB2 Group (or full DB2 database backup) and activate (copy) the DB2 enabled Notes database that you want to the alternate Domino DB2 database.

Environments that contain both NSF and DB2 enabled Notes databases

Domino environments that contain both NSF and DB2 enabled Notes databases can implement the following backup strategy:

- Perform full DB2 database backups and NSF selective backups on a regular basis.
- Perform routine incremental backups of NSF databases to inactivate backup copies that have been deleted from the Domino server.
- Perform regular DB2 Group backups if the DB2 database is enabled for rollforward recovery.
- Perform routine archiving of the transaction log files if archival transaction logging is enabled on the Domino server.
- Routinely inactivate the Domino server log file and routinely inactivate (and delete) DB2 objects from the Tivoli Storage Manager server.

Additional information

A description of the offering is provided, with links to installation resources and additional informations about the product.

System requirements

The detailed system requirements for this release, and other V6.4 maintenance updates, are available from the Tivoli Storage Manager for Mail - All Requirement Documents:

http://www.ibm.com/support/docview.wss?uid=swg21219345

Installing Data Protection for Domino

For information about where to download Data Protection for Domino, go to this address:

http://www.ibm.com/support/docview.wss?uid=swg24033227

7.5.4 Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server

Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server provides online backups and restores of Microsoft Exchange Server components to Tivoli Storage Manager storage. Data Protection for Microsoft Exchange Server provides a connection between an Exchange Server and a Tivoli Storage Manager server which allows Exchange data to be protected and managed by Tivoli Storage Manager. Data Protection for Microsoft Exchange Server protects Exchange Server data and improves the availability of Exchange databases.

Support: Starting with Exchange Server 2010, Microsoft no longer supports legacy-style (streaming) backups. Only VSS-style backups are supported. Tivoli Storage Manager 7.1 supports only Microsoft Exchange Servers 2010 and 2013. For these Exchange servers only VSS backups are supported. References to Legacy backups in this book are supported by Tivoli Storage Manager 6.4.1 for Microsoft Exchange Server 2007.

Challenges and benefits the solution addresses

See the appropriate matrices we created to describe the challenges this solution addresses.

Today, email systems play a key role in making an enterprise successful. Businesses are often severely impacted when email service is down, even if other production services are running. Consequently, ensuring email server availability is a critical business concern.

In the face of such constraints, certain requirements must be addressed, such as these:

- Fast recovery
- Fast backups
- *Zero impact*, high performance backups
- Intelligent management of these backups

Data Protection for Microsoft Exchange helps protect and manage Exchange Server environments by facilitating the backup, restore, and recovery of Exchange Server data.

Data Protection for Microsoft Exchange provides these key features:

- Perform individual mailbox recovery and item-level recovery from Data Protection for Microsoft Exchange backups.
- Back up Exchange Server 2007 databases using the Exchange Server streaming backup and restore API.
- Back up Exchange Server databases using Microsoft Volume Shadow Copy Service (VSS).
- Back up Exchange Server 2007 continuous replica copies (CCR or LCR) using VSS technology.
- Back up Exchange Server Database Availability Group databases to a common node to manage all DAG members using a single policy.
- Perform a VSS backup to the Tivoli Storage Manager server using an alternate machine instead of a production machine.
- Restore VSS backups that reside on Tivoli Storage Manager server storage to their original location.
- Restore VSS backups that reside on local shadow volumes using file-level copy mechanisms.
- Restore VSS backups that reside on local shadow volumes using hardware-assisted volume-level copy mechanisms.
- Restore a VSS backup of Exchange Server 2007 data into a Recovery Storage Group, alternate storage group, or relocated storage group.
- Restore a VSS backup of Exchange Server 2010 data into a Recovery database, alternate database, or relocated database.
- Query the managed capacity for VSS backups that reside on local shadow volumes.
- Delete a VSS backup of an Exchange Server storage group or database.
- Manage policy for VSS backups that reside on local shadow volumes.
- Integrate with Tivoli Storage FlashCopy Manager.
- Tivoli Storage Manager policy-based management of VSS snapshot backups.
- Use the **`restorefiles`** command to restore VSS backups to flat files without involving the Exchange Server.

Certain Data Protection for Microsoft Exchange functions vary based on the version of Exchange Server in your environment. Exchange Server 2010 and 2013 have functions that differ from functions available with Exchange Server 2007:

- Exchange Server 2010 and 2013 provide Database Availability Groups (DAG). A DAG consists of mailbox servers that provide recovery from database, server, or network failures. They provide continuous replication and continuous mailbox availability; replaces Exchange Server 2007 local continuous replication (LCR), cluster continuous replication (CCR), and standby continuous replication (SCR) replication.
- Exchange databases replace Exchange storage groups.
- Exchange Management Shell commands have been changed to support the new Exchange features and storage configuration.
- The Recovery Database (RDB) replaces the Recovery Storage Group (RSG).
- The number of databases allowed for each Exchange server increases from 50 to 100.
- Single Copy Clustering (SCC) is not available starting with Exchange Server 2010.
- Only VSS-style backups are supported. Starting with Exchange Server 2010, Microsoft no longer supports legacy-style (streaming) backups.

Solution architecture

Data Protection for Microsoft Exchange performs online backup and restore operations of Microsoft Exchange Server databases (Exchange Server 2010 and 2013) to Tivoli Storage Manager storage or local shadow volumes. You can perform backups and restores using a command-line interface (CLI) or graphical user interface (GUI). See your Exchange Server documentation for complete, detailed information about the backup and restore process of Microsoft Exchange Servers.

Beginning with Exchange Server 2010 Microsoft no longer supports the Microsoft Legacy API (streaming) for backup and restore operations. It only supports the use of VSS for the backup and restore.

Data Protection for Microsoft Exchange operations use the Tivoli Storage Manager application programming interface (API) to communicate with the Tivoli Storage Manager server, and use the Exchange API to communicate with Exchange Server. In addition to using these APIs, Data Protection for Microsoft Exchange VSS operations use the Tivoli Storage Manager backup-archive client (VSS Requestor) and Microsoft Volume Shadow Copy Service to produce an online snapshot (point-in-time consistent copy) of Exchange data that can be stored on local shadow volumes or on Tivoli Storage Manager server storage.

Solution description

Data Protection for Exchange provides a Legacy method and a VSS method for backing up data.

Data Protection for Exchange tracks and stores mailbox location history, which is used to automate mailbox restore operations. This causes a delay before each backup. In a small or centralized Active Directory environment, the delay might be a few seconds or minutes. In large or geographically disperse Active Directory environments, the delay might be more time. If you do not plan to use mailbox restore, you can safely disable mailbox history.

Legacy backup overview

A Legacy backup creates a copy of an Exchange Server 2007 storage group on Tivoli Storage Manager storage media. The backup includes any associated transaction logs.

When you initiate a Legacy backup operation, Data Protection for Exchange completes the following actions:

1. Begins a session with a Tivoli Storage Manager server using the Tivoli Storage Manager API and information contained in a client options file.
2. Informs the Exchange Server that a backup is ready to begin.
3. Receives data from the Exchange Server and sends it to the Tivoli Storage Manager server.
4. Informs the Exchange Server that the backup is complete.
5. Ends the Tivoli Storage Manager server session.

The following characteristics are true of Legacy backups:

- You can use full, copy, incremental, differential, and database copy backup types.
- Backup granularity is at the database and storage-group level.
- Backups are stored on Tivoli Storage Manager server storage.
- Backups are managed through the Tivoli Storage Manager server policy.
- Backups can be performed in a Microsoft Windows Failover Clustering (previously MSCS) or Veritas Cluster Server (VCS) environments.
- Backups provide Exchange Server database zeroing function.
- Backups provide Exchange Server database integrity check function.

The following restrictions apply:

- Microsoft does not support Legacy backups on cluster continuous replication (CCR) or local continuous replication (LCR).
- Microsoft does not support performing legacy (or VSS) backups on standby continuous replication (SCR) databases. You cannot run Legacy backups on CCR, LCR, or SCR databases.
- You can run VSS backups on CCR and LCR databases.
- You can use Legacy backups to backup databases that have CCR, LCR, or SCR replicas, but you must back up the primary database, not the replica.

Data Protection for Microsoft Exchange provides backup and restore functions for the Exchange storage groups and associated transaction logs. Data Protection for Microsoft Exchange does not provide a complete disaster recovery solution for an Exchange Server. In a disaster recovery situation, Data Protection for Microsoft Exchange restores local continuous replication (LCR) storage groups. Other files need to be restored in a disaster recovery situation. See your Microsoft Exchange Server documentation for disaster recovery considerations.

Personal folders and personal address books that are stored on Microsoft Outlook clients are not protected by Data Protection for Microsoft Exchange. The Tivoli Storage Manager backup-archive client can be used on the Outlook client platform to back up and restore these files. Because the Outlook client normally keeps these files locked when running, you should stop the Outlook client before backing up or restoring these files. Because Tivoli Storage Manager Backup-Archive client provides open file support, you might be able to back up and restore these files while the Outlook client is running.

VSS backup

A VSS backup uses Microsoft Volume Shadow Copy Service technology to produce an online snapshot (point-in-time consistent copy) of Exchange data.

Figure 7-19 summarizes the components of a Tivoli Storage Manager with Data Protection for Exchange solution providing VSS backup restore services.

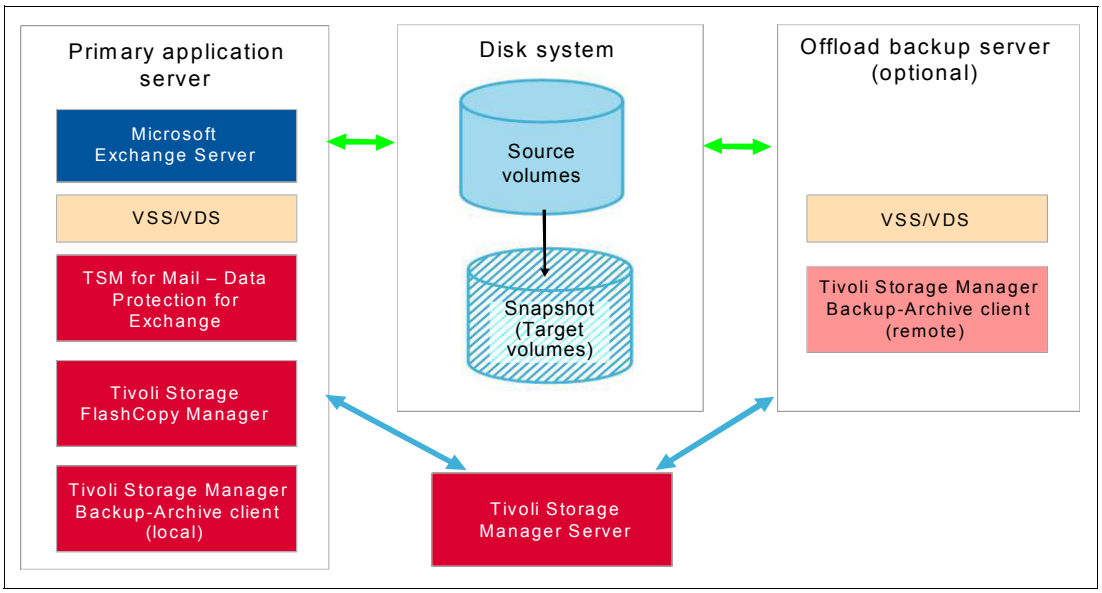

Figure 7-19 Data Protection for Exchange with Tivoli Storage Manager integration

A VSS backup means that the Exchange Server is not in backup mode for an extended period of time. The length of time to perform the snapshot is usually measured in seconds, not hours. In addition, a VSS backup allows a snapshot of large amounts of data at one time because the snapshot works at the volume level. VSS backups can be stored on Tivoli Storage Manager server storage or local VSS shadow volumes. Both of these storage destinations require that sufficient space be available for the snapshot. VSS backups stored locally on VSS shadow volumes are directly accessible by the Exchange system (as long as sufficient space is available for the snapshot).

These types of local VSS backups are faster for a couple of reasons:

- Because of the way the snapshots are managed
- Because the data is not placed into Tivoli Storage Manager server storage

Restoring these backups is also fast because the Exchange data is not transferred from Tivoli Storage Manager server storage over the network.

For local VSS backups you must have a licensed version of IBM Tivoli Storage FlashCopy Manager installed on your system.

When performing VSS backups and moving data to Tivoli Storage Manager server storage, sufficient space on local snapshot volumes is still temporarily required to hold the snapshot. For Exchange data backed up to Tivoli Storage Manager server storage, the Exchange data on the snapshot volume is sent to the Tivoli Storage Manager server. After the data transfer to the server is complete, the snapshot volume is made available for reuse. If you are storing VSS backups locally and the maximum number of local backup versions to be maintained (as specified by the Tivoli Storage Manager policy) is reached, the oldest backup version is expired to create the snapshot for the backup to Tivoli Storage Manager server storage.

For Exchange data backed up to local VSS shadow volumes, the snapshot backup resides on the shadow copy volume. For Exchange data backed up to both destinations, a local snapshot backup is performed and the Exchange data on the local snapshot volume is sent to the Tivoli Storage Manager server. The local snapshot volume is retained as a local backup.

The Tivoli Storage Manager backup-archive client serves as the VSS requestor that communicates with VSS to access the Exchange data to create shadow copies of Exchange storage groups. Data Protection for Exchange serves as a front end for VSS backup operations. Because of the role that the backup-archive client performs as the VSS requestor, features such as LAN-free backup, client-side deduplication, data encryption, and data compression require that options related to these features be specified in the backup-archive client options file for VSS operations.

VSS backup operation steps

When a VSS backup operation is initiated the following actions are performed:

1. Data Protection for Exchange validates the state of Exchange server objects.
2. Data Protection for Exchange begins a session with a Tivoli Storage Manager server.
3. Data Protection for Exchange verifies that the VSS service is running and that the Exchange writer is available.
4. The Tivoli Storage Manager VSS requestor lists the backup components through the VSS Writer.
5. The Tivoli Storage Manager VSS requestor performs the VSS snapshot backup preparation stage.
6. The Tivoli Storage Manager VSS requestor performs the actual VSS backup.
7. The Tivoli Storage Manager VSS requestor performs an integrity check on the VSS backup.
8. Optional: The integrity check can be offloaded to an alternative machine that has the Tivoli Storage Manager VSS requestor installed and configured.
9. The Tivoli Storage Manager VSS requestor backs up the data, including metadata, to a Tivoli Storage Manager server. Optionally, the movement of data to a Tivoli Storage Manager server can be offloaded to an alternate machine that has the Tivoli Storage Manager VSS Requestor installed and configured.
10. The Tivoli Storage Manager VSS requestor marks the backup as complete in VSS.
11. Data Protection for Exchange ends the Tivoli Storage Manager server session.

Some VSS backup characteristics differ from Legacy backup characteristics. Examples of these differences are the backup characteristics for types supported, the backup granularity, and the backup storage location options.

VSS backup characteristics

The following characteristics are true of VSS backups:

- Full, copy, incremental, and differential backup types are supported. Database copy backup types are not supported.
- Backup granularity is at the storage group or database level only.
- Backups are managed through Tivoli Storage Manager server policy.
- Backups can be stored on local shadow volumes, Tivoli Storage Manager server storage, or both locations.
- Different policy settings can be defined for each storage location, backup method, and backup type (FULL or COPY).

- Backups to Tivoli Storage Manager server storage can be offloaded to an alternate machine as resource relief for production servers.
- Backups can be run in a Microsoft Windows Failover Clustering or Veritas Cluster Server (VCS) environments.
- Backups do not provide an Exchange Server database zeroing function.
- Backups provide an Exchange Server database integrity check function.
- Restores storage groups or database backups to the Recovery Storage Group or Recovery Database, or to an alternative location.
- VSS backups can be restored to flat files without the involvement of the Exchange Server. See the restorefiles command for more details.
- VSS is the only available backup method on Exchange Server 2010 and 2013.
- For databases in an Exchange Server DAG that have two or more healthy copies, the database integrity check can be skipped.
- Exchange Server Database Availability Group (DAG) databases can be backed up under a common DAG node name, regardless of which DAG member runs the backup. When backing up data to a common node, the backups are managed by a common policy, and the user can restore to any Exchange Server. Setting the minimum interval can prevent too frequent backups.

Legacy restore

When a legacy restore is initiated from the Data Protection for Exchange GUI, you are prompted to either dismount the databases or cancel the restore operation. If the restore is run from the CLI, the Exchange Administrator must first dismount the necessary databases.

Data Protection for Exchange performs the following actions to restore an Exchange database or storage group.

1. It starts a session with the Tivoli Storage Manager Server.
2. It informs the Exchange Server that a restore is about to start.
3. It restores the specified database and log files from the Tivoli Storage Manager server. The logs are restored to a temporary location as specified by the Exchange Administrator.
4. It informs the Exchange Server that the restore has completed. You have the option of either starting the recovery or mounting the database (when the recovery completes).
5. It ends the Tivoli Storage Manager server session.

Depending on what backup strategy you choose, restoring an Exchange Storage Group can involve restoring multiple backup objects from the Tivoli Storage Manager server. For details about restoring Legacy backups, see *Data Protection for Microsoft Exchange Server Installation and User's Guide*, SC32-9058.

VSS restore

Only backups that are made through VSS can be restored through VSS. Therefore, incremental or differential Legacy backups cannot be restored with this method, and neither can full or copy Legacy backups. Because of current Microsoft restrictions, a Recovery Storage Group (RSG) restore from VSS snapshot backups is not supported. Site Replication Service (SRS) and Key Management Service (KMS) also cannot be restored with VSS. A VSS restore will be directed to the same drive letters and paths as when the backup was run.

When you initiate a restore operation, the following actions are performed:

1. Data Protection for Exchange validates the state of Exchange server objects.
2. When using the Data Protection for Microsoft Exchange GUI, you are prompted whether to dismount the databases within the selected storage group you are restoring into.
3. Data Protection for Exchange begins a session with a Tivoli Storage Manager server.
4. Data Protection for Exchange verifies that the VSS service is running and that the Exchange writer is available.
5. The Tivoli Storage Manager VSS Requestor performs the VSS snapshot restore preparation stage.
6. The Tivoli Storage Manager VSS Requestor restores the backup data.
7. The Tivoli Storage Manager VSS Requestor marks the restore as complete in VSS.
8. Optionally, Data Protection for Exchange mounts databases to run recovery.

Use scenarios

Depending on your specific requirements regarding network traffic, backup window, and acceptable restore times, you might choose to follow different backup strategies. It is important to understand all aspects of Exchange Server disaster recovery, and backup considerations recommended by Microsoft.

Data Protection for Exchange provides five types of backup:

- These four can be performed with legacy (Exchange Server 2007 only) and VSS operations:
  - Full backup
  - Copy backup
  - Incremental backup
  - Differential backup
- The database copy backup type, DBCOPY, can be performed with legacy operations on Exchange Server 2007 only.

Table 7-4 summarizes these methods and can help you plan your backup strategy.

Table 7-4 Backup strategies with Data Protection for Microsoft Exchange Server

Backup type	Use	Legacy	VSS
Full backup	This type backs up the specified storage group or database, and associated transaction logs. The Exchange Server deletes the committed log files after the storage group or database, and logs are successfully checked for integrity and backed up. If the storage group or database is not mounted, the backup fails and the transaction logs are not truncated	X	X
Copy backup	This type is similar to a full backup except that transaction log files are not deleted after the backup. A copy backup is used to make a full backup of the Exchange Server storage group without disrupting any backup procedures that use a incremental or differential backup	X	X
Database copy backup	This type backs up only the specified database and its associated transaction logs. The transaction log files are not deleted after the backup. A database copy backup is used to make a special full backup of the database without disrupting any backup procedures that use incremental or differential backups	X	

Backup type	Use	Legacy	VSS
Incremental backup	This type backs up only transaction logs. The Exchange Server deletes the committed log files after they are successfully backed up. These log files are not deleted if the backup fails. Restoration of an Exchange Server storage group or database from an incremental backup requires the following tasks: 1. Restore of the last full backup 2. Restore of any other incremental backups that are performed between the full backup and this incremental backup 3. Restore of this incremental backup The log files are not deleted if storage groups or databases are not mounted.	X	X
Differential backup	This type backs up transaction logs. The log files are not deleted. When you perform a full backup followed by only differential backups, the last full backup and the latest differential backup has all data required to bring the storage group back to the most recent state. This type of backup is also called a cumulative incremental backup. Restoring an Exchange Server storage group from a differential backup requires the following tasks: 1. Restore of the last full backup 2. Restore of this differential backup, but no other differential backups	X	X

Circular logging: When circular logging is enabled, you cannot use differential or incremental backups. Data loss might occur if the log wraps before an incremental or differential backup is finished. If you choose a backup strategy that involves incremental or differential backups, you must disable circular logging for the Exchange storage group or database from the Exchange Administrator program. See your Microsoft Exchange Server documentation for more information about circular logging.

When you are considering backup strategies, use the following guidelines:

- Do not use incremental and differential backups together.
- If you choose a strategy that involves incremental or differential backups, circular logging must be disabled on the storage groups or databases of the Exchange Server.
- For Exchange Server 2010 and 2013, consider using DAG database replication technologies. See your Microsoft documentation for details regarding this technology. For an additional Database Availability Group (DAG) backup strategy, set up all DAG members to back up all of the database copies. Use the /MIN and /PREFERDAGPAS flags.

Possible backup strategies

The following backup strategies are based on the type of backup being performed.

- Full backup only

  This approach is best for Exchange Servers that are relatively small because each backup contains enough data to restore the entire storage group. Each backup takes longer to perform, but the restore process is the most efficient because only the most recent (or other appropriate) full backup needs to be restored.

- Full backup plus incremental backups

  This strategy is commonly used when the normal backup window or network capacity cannot support a full backup each time.

In such cases, a periodic full backup followed by a series of incremental backups allows the backup window and network traffic to be minimized during peak usage times. For example, you can perform full backups on the weekend and incremental backups during the week. The full backups can be done during low usage times when a larger backup window and increased network traffic can be tolerated. The restore process becomes more complex, however, because a full backup, as well as subsequent incremental backups, must be restored. In addition, transactions within the logs must be applied which increases process time. As a result, the more transactions applied, the longer the recovery process.

If you use this backup strategy, you must decide whether the Tivoli Storage Manager storage management policies are modified, to ensure all incremental backups are stored together on the Tivoli Storage Manager server (collocated). This helps improve restore process performance by reducing the number of media mounts necessary for restoring a series of incremental backups.

- Full backup plus differentials

  This strategy provides easier restoring than the full plus incremental backup strategy.

  This approach might be useful if your backup window and network capacity can process the backup of all transaction logs that accumulate between full backups. This is because it requires the transfer of only one differential plus the last full backup to accomplish a restore. However, the same amount of data must be transferred in the differential image, as in the series of incremental backups.

  Therefore, a full backup plus differential backup policy increases network traffic and Tivoli Storage Manager storage usage. This assumes that the differential backups are done with the same frequency as the incremental backups.

  Carefully consider whether there is sufficient advantage to justify the extra resources necessary to resend all prior transaction logs with each subsequent differential backup.

Using VSS and Legacy backups together

With Exchange Server 2007, you can use both VSS and Legacy backups in your complete backup strategy, but you cannot mix the two types of backups. Use the following guidelines:

- Legacy and VSS backups to Tivoli Storage Manager server storage are usually dictated by time, not versions.
- Backups to local shadow volumes are usually dictated by versions because of space limitations and provisioning of VSS storage

Hardware and software requirements

You must install Data Protection for Microsoft Exchange on the same system as the Exchange Server. Data Protection for Microsoft Exchange also supports operations in a Microsoft Windows Failover Clustering (previously MSCS), Veritas Cluster Server (VCS), and Database Availability Group (DAG) environment.

Details of the hardware and software requirements change over time because of maintenance updates and the addition of operating system, application, and other software currency support.

For the most current requirements, see the hardware and software requirements technote that is associated with your level of software. This technote is available from the following site:

http://www.ibm.com/support/docview.wss?uid=swg21219345

7.5.5 Tivoli Storage Manager for Enterprise Resource Planning

Data Protection for SAP is fully integrated to the SAP environment. The communication with the backup-archive server is performed using an API called ProLE, as shown in Figure 7-20. This API is shared with other Tivoli Data Protection products. ProLE runs as a background process for communicating with the Tivoli Storage Manager server. The administration assistant is available for the SAP or Database Administrator.

This optional component provides the following features:

- Configuration of Data Protection for SAP instances
- Performance monitoring of data transfer between Database and Tivoli
- Storage Manager Server
- Analysis and simulation
- Monitoring of backup and restore operations
- Administration all of your Data Protection for SAP instances remotely

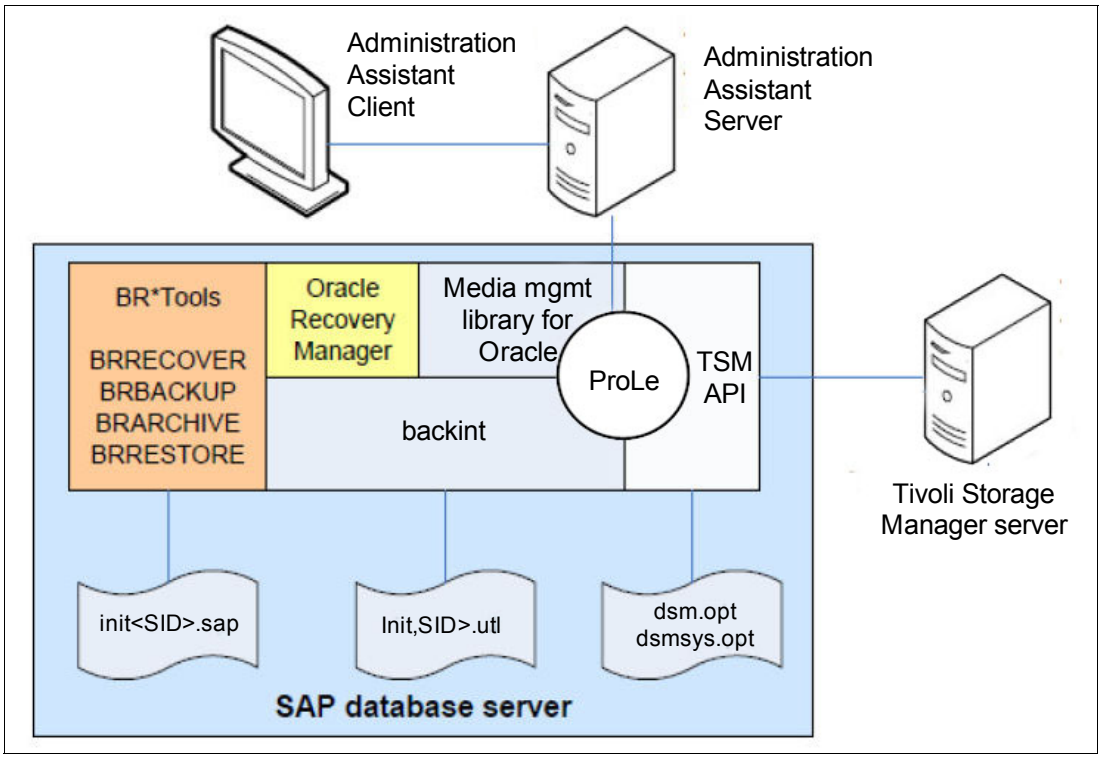

Figure 7-20 Tivoli Data Protection for SAP sample scenario

Client benefits

IBM Tivoli Storage Manager for Enterprise Resource Planning protects your vital SAP system data. It provides automated data protection for mySAP and SAP R/3 environments. You can improve availability of your SAP database servers and reduce your administration workload.

Tivoli Storage Manager for Enterprise Resource Planning offers these benefits:

- Delivers business value by protecting SAP system data efficiently and consistently.
- Helps increase productivity by reducing repetitive administrative tasks.
- Is SAP certified for heterogeneous environments.
- Provides more efficient backup of very large SAP databases.
- Integrates with database-specific utilities of IBM DB2 for Linux, UNIX and Windows, Oracle and SAP BR*Tools.

Solution architecture

Data Protection for SAP and Tivoli Storage Manager provide a reliable, high performance, and production-oriented solution that enables back up and restore of DB2-based Oracle-based SAP systems. It is integrated with DB2 backup and recovery facilities SAP backup and recovery utilities BRBACKUP, BRARCHIVE, BRRESTORE, and BRRECOVER, and applies SAP backup and recovery procedures. Data Protection for SAP is optimized for SAP databases and therefore provides efficient management of large data volumes.

As demonstrated in Figure 7-21, SAP backup and recovery utilities center on database objects where more than 90% of the data resides on an SAP database server. As a result, Data Protection for SAP backs up and restores data files, control files, and online or offline redo logs.

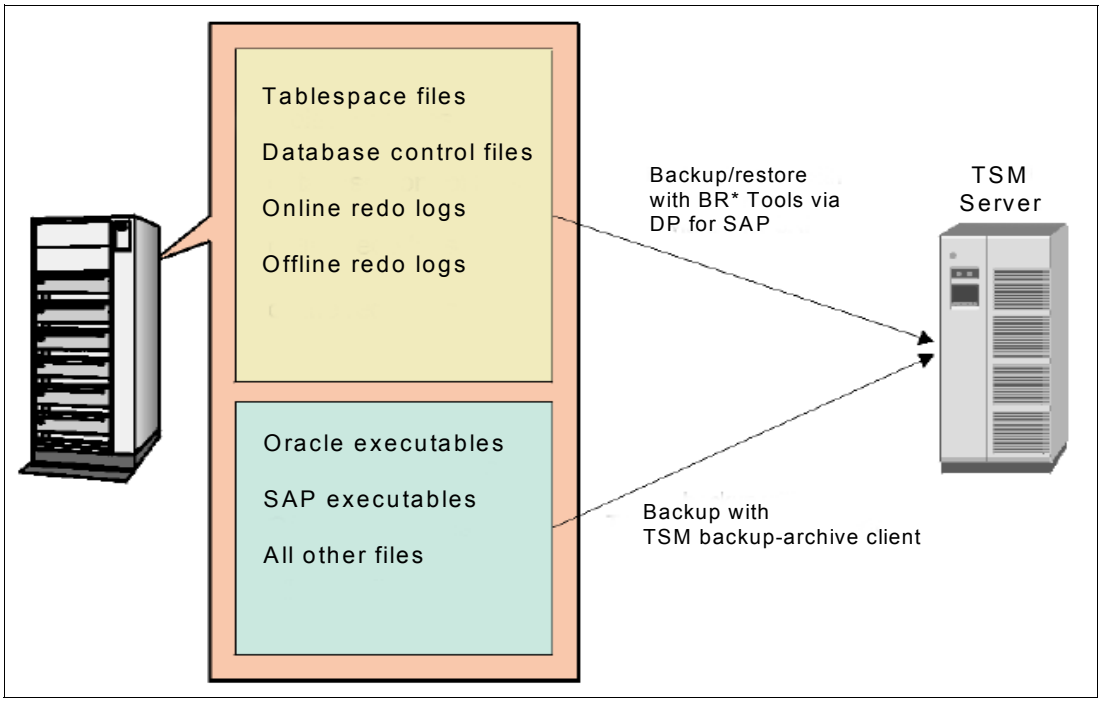

Figure 7-21 Scope of Data Protection for SAP for Oracle

Figure 7-22 shows that SAP backup and recovery utilities center on database objects where more than 90% of the data resides on an SAP database server. As a result, Data Protection for SAP backs up and restores database contents, database specific control files, for example the database configuration, the history and the log file header, and offline DB2 log files.

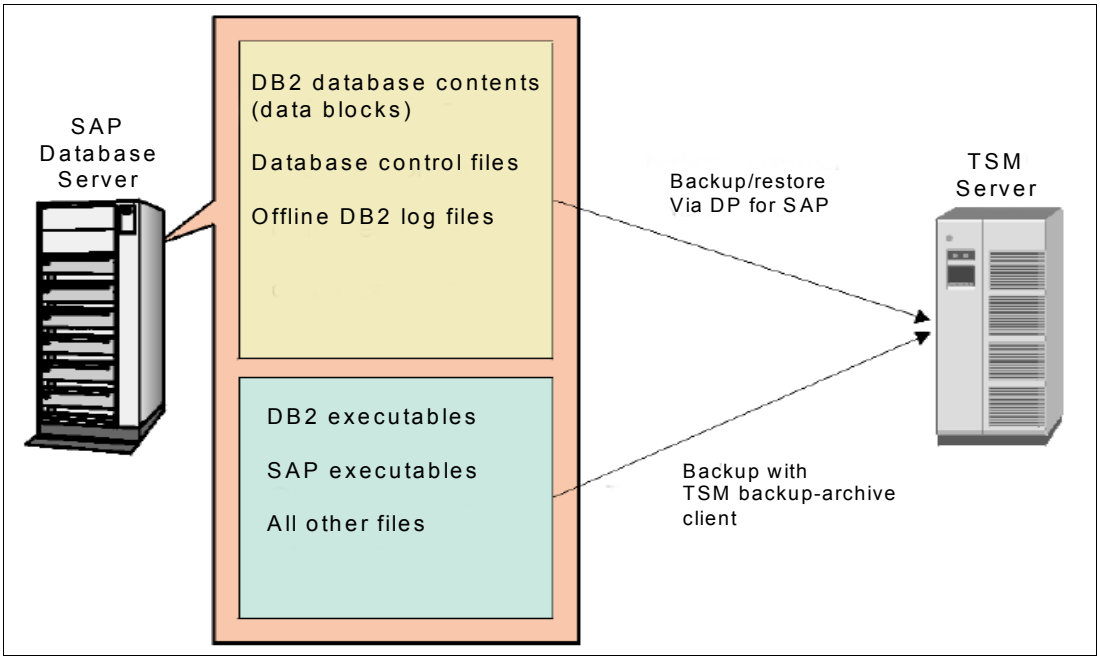

Figure 7-22 Scope of Data Protection for SAP for DB2

Other files (such as SAP and Oracle or DB2 executable files) can be backed up using the IBM Tivoli Storage Manager Backup-Archive Client. This is important for disaster recovery purposes, because all SAP and Oracle or DB2 executable files must be available before using Data Protection for SAP to restore and recover the database.

Solution for Data Protection for SAP for Oracle database

Data Protection for SAP is a client/server program that manage backups and restores in connection with Tivoli Storage Manager. With Data Protection for SAP, handling SAP database backups is possible, and it includes the ability to manage backup storage and processing independently from normal SAP operations. Furthermore, Data Protection for SAP in combination with Tivoli Storage Manager provides reliable, high performance, and repeatable backup and restore processes to manage large volumes of data more efficiently. For Oracle databases, two options implement a backup using Tivoli Storage Manager:

- Data Protector for SAP using the BACKINT interface
- Data Protector for SAP using Oracle Recovery Manager (RMAN)

With the integration, it is possible to follow ERP backup/restore procedures and to use the integrated SAP database utilities BRBACKUP, BRARCHIVE, BRRESTORE and SAPDBA for backup and restore. Other SAP related files (executables) are backed up by using Tivoli Storage Manager standard techniques for file backup and restore, for example, incremental backup, file filtering, and point-in-time recovery.

Figure 7-23 shows a diagram of the backup and restore process in connection with Data Protection for SAP.

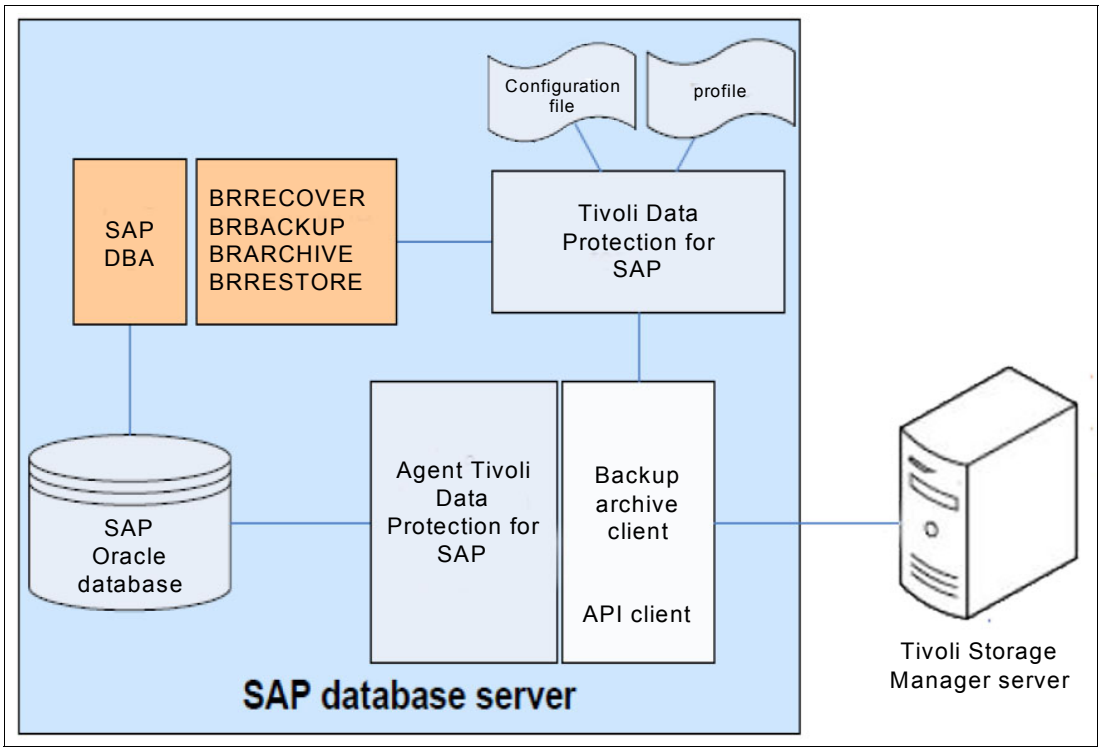

Figure 7-23 Data Protection for SAP architecture for Oracle

Figure 7-24 shows the Tivoli Data Protection for SAP for Oracle function and interfaces.

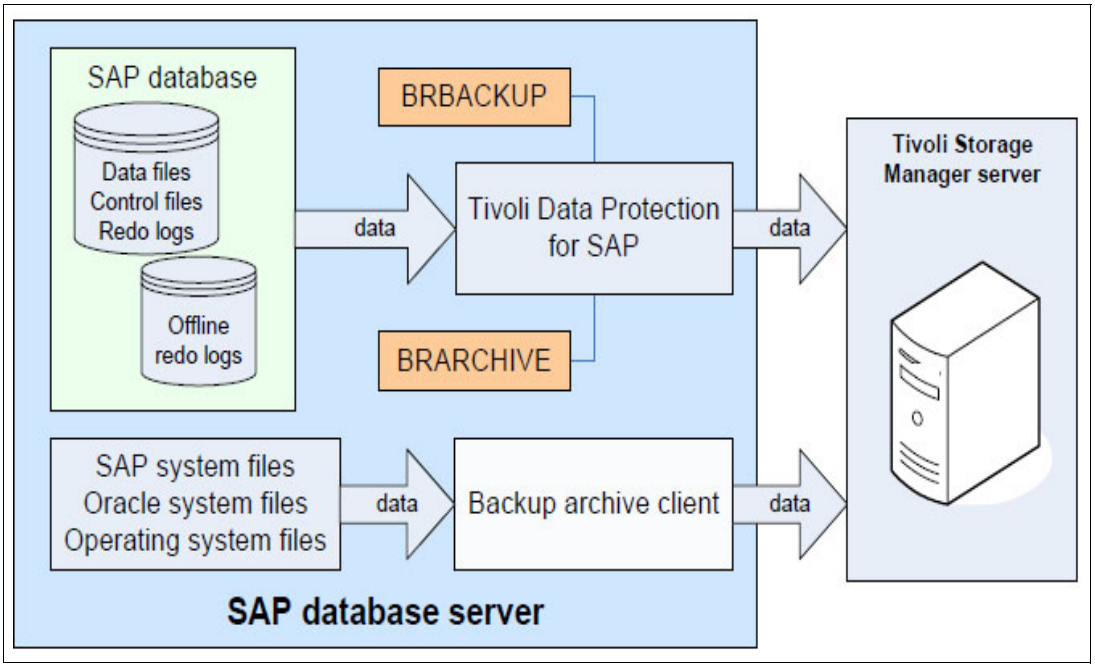

Figure 7-24 Tivoli Data Protection for SAP for Oracle overview

Data Protection for SAP using BACKINT

Using this feature, you are able to perform the traditional Oracle online backup with automation provided by BACKINT.

Figure 7-25 shows the data interface between Oracle Databases and Tivoli Storage Manager using Data Protection for SAP for Oracle using backint interface.

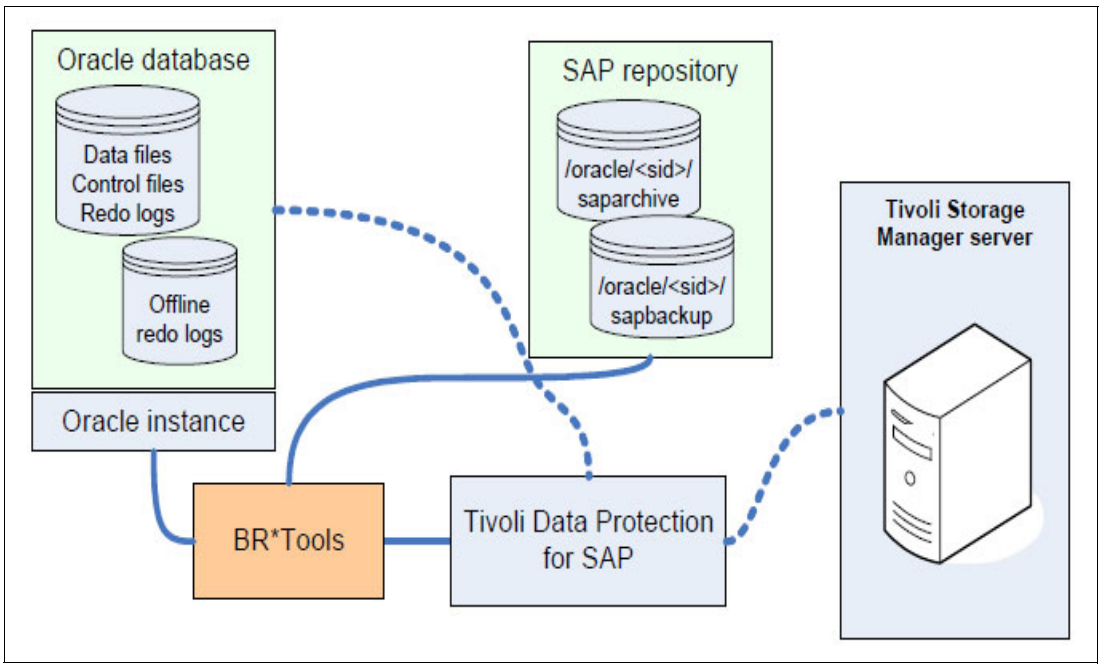

Figure 7-25 Data Protection for SAP for Oracle using Backint

The backup steps are as follows:

1. BR*Tools takes control.
2. BRBACKUP calls the Data Protection for SAP using Backint.
3. Backint changes the table spaces to backup mode by using the following command:

   ```
   alter tablespace <tablespace name> begin backup
   ```

4. Backint using Tivoli Data protector for SAP reads all the data files and saves them to Tivoli Storage Manager Server.
5. BR*Tools updates the catalog with information regarding the backed up data file.

Notes:

- Using this method, the chosen data files are sent to Tivoli Storage Manager one by one. No compression or block checking are performed at this level.
- When a database is in backup mode, the amount of redo logs written to disk increases. This is because Oracle writes the entire dirty block to the disk, not just the updated data.
- In some cases, when the backup routine fails for any reason, the data file remains in active backup mode. This can cause some performance impact and additional I/O to the disk.

Data Protection for SAP using RMAN

Using this feature, you are able to take all the advantages and facilities provided by RMAN. In general, RMAN is able to perform backup in less time compared to the traditional backup using Backint. This is possible because RMAN sends only used data blocks (in an Oracle data file) to Tivoli Storage Manager. The other interesting feature is block checking, which discovers bad blocks as soon as they occur. In addition, you can use the Oracle RMAN (Recovery Manager) utility to execute some tasks not provided by BR*Tools, such as incremental backups, releasing backup versions, and catalog maintenance.

Figure 7-26 shows the data interface between Oracle Database and Tivoli Storage Manager using Data Protection for Oracle for SAP using RMAN.

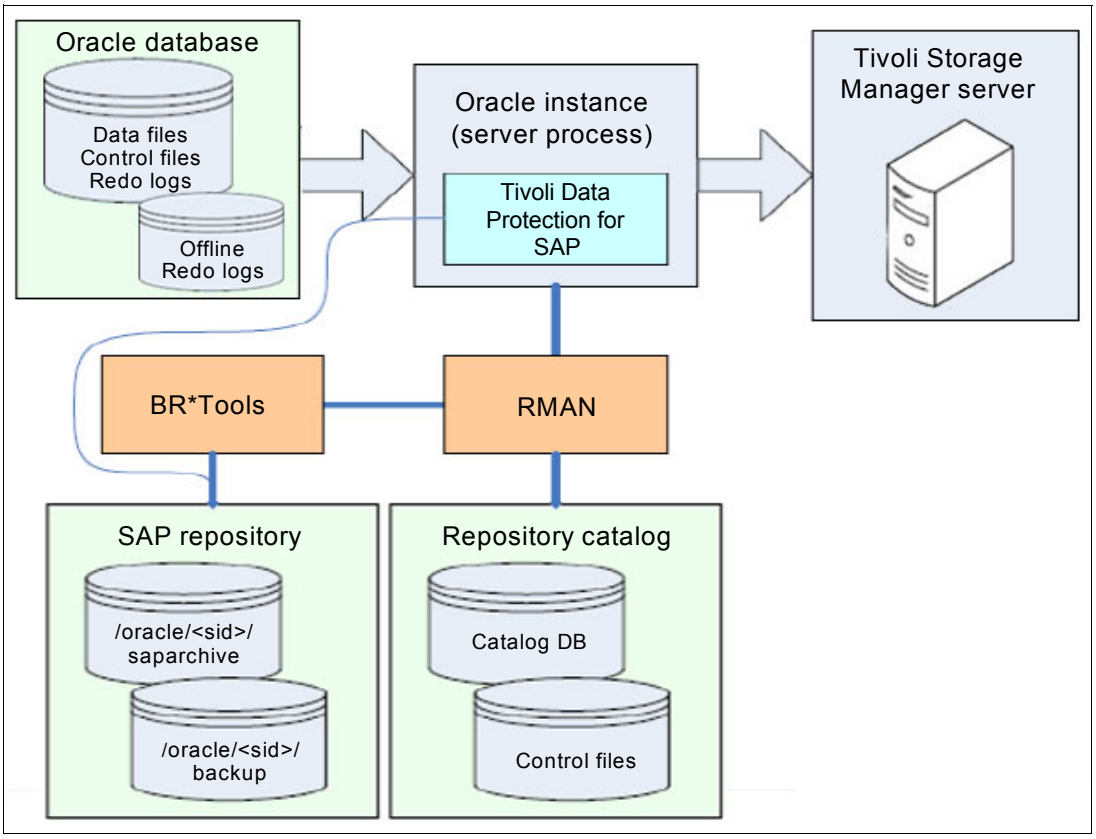

Figure 7-26 Data Protection for SAP for Oracle using RMAN

These extra features are available only to RMAN:

- Data Recovery Advisor (Version 11g and later). This feature helps you in analyzing and deciding what is the best recovery plan.
- Fast Backup Compression (Version 10g and later). This feature helps reduce the amount of data sent to the tapes.
- Network-enabled database duplication (Version 11g and later). Cloning databases in the network is easy. There are no requirements to have an existing backup.
- Integration with Windows Volume Shadow Copy Services - VSS (Version 11g and later). This feature allows the Oracle Database to participate in the VSS infrastructure on Windows platforms.

- Backup Set encryption (Version 10g and later). With this feature, only the backup creator is able to restore.
- Unused Block Compression (Version 10g and above). This feature is enabled by default. It reduces the amount of data sent to the tape by reducing the backup time.

Tivoli Data Protection for SAP for DB2

Tivoli Data Protection for SAP for DB2 was created to provide an intelligent interface to manage backup and restore using Tivoli Storage Manager. It is fully integrated in the SAP environment. The backup command (DB2 BACKUP DATABASE) and the restore command (DB2 RESTORE DATABASE) are initiated by the DB2 command line (CLI), which calls the Tivoli Data Protection for SAP for DBA module. The backup and restore of the DB2 log files is provided by the BR*Tools commands BRARCHIVE and BRRESTORE. In addition, you can use the Tivoli Data Protection for SAP for DB2 Tools BackOM and the built-in Log Manager.

Figure 7-27 shows the data interface between DB2 Databases and Tivoli Storage Manager using Data Protection for SAP for DB2.

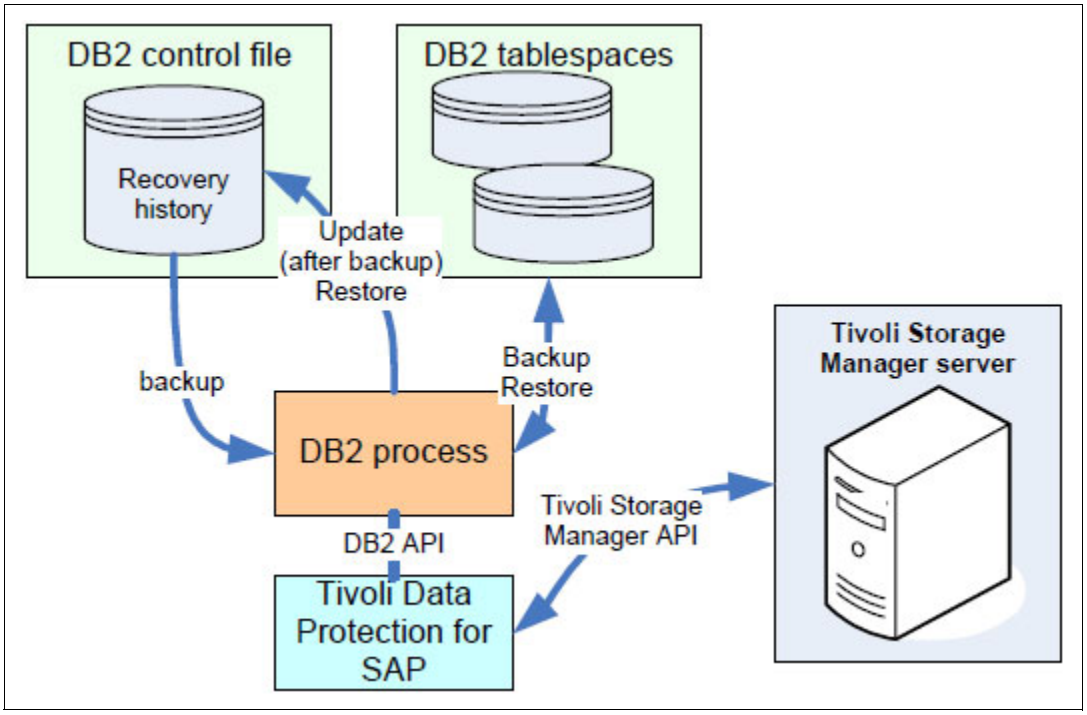

Figure 7-27 Data Protection for SAP for DB2

The archiving of DB2 offline log files is provided by the SAP tool BRARCHIVE. The retrieval of DB2 offline log files is provided by the SAP tool BRRESTORE and by the Data Protection for SAP tool BackOM. As of DB2 Version 9.X, offline log files can be archived and retrieved with the DB2 built-in Log Manager. The DB2 Command-Line Processor (CLP) interprets commands for the DB2 database and passes control to a DB2 Server Process. In the case of Data Protection for SAP, the `LOAD <libraryname>` option causes DB2 to invoke the Data Protection for SAP shared library. This process kicks off the backup or restore, loads the library dynamically, and communicates with it through the TivoliStorage Manager API. For starting a backup or restore, the DB2 CLP communicates with the DB2 Server Process, providing the Server Process with the relevant information for processing the database

Use scenarios

Database backup and recovery tools (backup managers) provide the functions to control backup and recovery of databases in case of media failure, logical failure, or a disaster. In addition to database recoverability, backup and restore tools can assist in other administration tasks, such as database cloning and data migration.

Backup and recovery tools are either an integrated component of a relational database system, such as Oracle Recovery Manager (RMAN), DB2 backup/restore functions, or may be installed as an independent software package developed by a third party (such as SAP BR*Tools for Oracle).

Aside from performing backup and recovery operations, backup managers integrated with media management systems can manage and enforce backup storage policies and backup retention policies. The backup retention policies control retention and deletion of backup versions stored on the backup media. Backup storage policies determine the destination media to be used for the particular backup objects, such as backup images of table spaces, and backups of archived logs can be stored on different media, as dictated by the storage policy.

The basic functions supported by the backup management tools allow you to do these tasks:

- Automate data protection tasks and allow for backup of database objects using various methods and techniques (such as online or offline backup and full or partial backup).
- Provide automatic or semi-automatic restore and recovery functions specific to a particular RDBMS. A backup repository is used for the decision making about which backup image to restore (for example, backup manager picks the last version of full backup and determines which archive redo log backups will be needed for rollforward recovery).
- Use a backup repository to keep track of the location of backup images and their time stamps. The information from the backup repository is retrieved whenever the user sends a list of the history of backup operations. The backup repository records are also used whenever backup manager has to determine which backup objects will be retrieved during the automatic database restore or recovery.
- Control the retention and expiration of backup versions according to retention policies. The retention policy may be defined by either the number of days the backup objects are to be retained or by a number of backup versions to be kept.

Database managers provide interfaces to transfer backup data to backup media or third-party media management systems such as IBM Tivoli Storage Manager. The media interfaces either are based on open standards (such as DB2 XBSA) or can be specific for the particular RDBMS (such as BR*Tools BACKINT interface or Oracle RMAN SBT interface). Interface adapters for particular backup management systems are developed by the media management systems vendors. This conceptual relationship is shown in Figure 7-28.

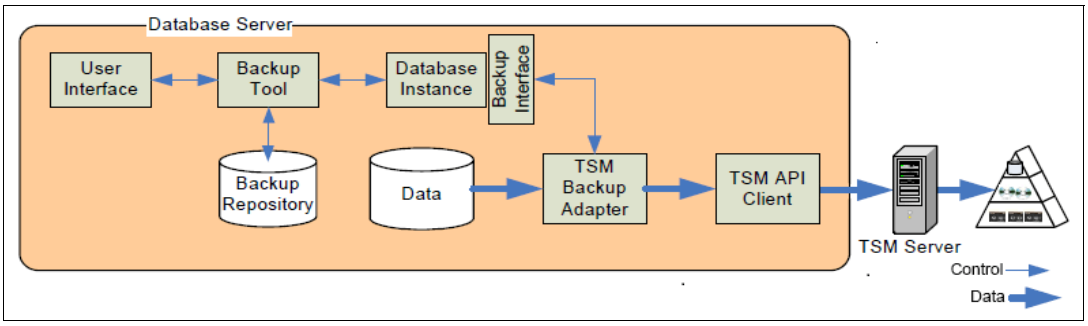

Figure 7-28 System architecture: database backup adapters for Tivoli Storage Manager

Backup adapters for IBM Tivoli Storage Manager are installed as part of the software packages Tivoli Storage Manager for Databases or Tivoli Storage Manager for Enterprise Resource Planning (formerly known as Tivoli Data Protection modules). Those Tivoli Storage Manager backup solutions include backup adapters based on the shared libraries. The backup adapters are compatible with the interfaces of database backup tools. The Tivoli Storage Manager backup adapter is called (or dynamically linked) by the backup management component of the database using the backup interface. The backup adapter communicates with the Tivoli Storage Manager server through the Tivoli Storage Manager client application program interface (API). The basic functions provided by backup adapters are transfer the backup objects to and from the Tivoli Storage Manager server, instruct the Tivoli Storage Manager server to delete the selected backup objects, and retrieve information about the backup objects.

Note: IBM DB2 UDB provides integrated Tivoli Storage Manager backup support, so it can directly call the Tivoli Storage Manager API client to transfer data to Tivoli Storage Manager server.

Backup adapters

The following Tivoli Storage Manager clients provide backup adapters for Oracle, MS-SQL, IBM DB2 UDB, and SAP MaxDB databases:

- General purpose Tivoli Storage Manager backup adapters for databases:
  - IBM Tivoli Storage Manager for Databases - Oracle:

    Tivoli Storage Manager backup adapter for Oracle RMAN implements Oracle Media Management API, Secure Backup to Tape (SBT) API V2.0.
  - IBM Tivoli Storage Manager for Databases

    Microsoft SQL Server: Tivoli Storage Manager adapter for Microsoft SQL Server implements Microsoft SQL Server Virtual Device Interface (VDI), which enables the Microsoft SQL backups to be transferred to Tivoli Storage Manager server. It supports both earlier and VSS backups.
  - IBM DB2 built-in Tivoli Storage Manager support

    DB2 UDB provides an integrated support for Tivoli Storage Manager. Thus, the Tivoli Storage Manager API client is the only component that is required to enable transfers of DB2 backup data to the Tivoli Storage Manager server.
  - IBM ADINT/TSM

    The Tivoli Storage Manager adapter is similar to Tivoli Storage Manager for ERP and provides seamless integration with SAP MaxDB backup, restore, and recover utilities. ADINT/TSM is a service offering with central hotline support sold by Tivoli Services organizations in the countries. For more information see the following site:

    `http://www-05.ibm.com/de/entwicklung/adint_tsm/`
- SAP-oriented Tivoli Storage Manager backup adapters for databases:
  - Tivoli Storage Manager for ERP - Oracle

    This software package implements the Tivoli Storage Manager backup adapters for Oracle RMAN (SBTAPI) and for SAP BR*Tools (BACKINT). Tivoli Storage Manager for ERP (Oracle) can serve as a backup adapter either between SAP BR*Tools and Tivoli Storage Manager server or between Oracle RMAN and Tivoli Storage Manager server.
  - Tivoli Storage Manager for ERP - DB2

    This implements shared libraries that can integrate DB2 UDB backup tools with Tivoli Storage Manager server.

SAP-oriented Tivoli Storage Manager adapters support additional functions specific to the SAP environment (see Table 7-5). You can integrate Tivoli Storage Manager for ERP clients with the Administration Center, which serves as a solution for the central administration of database backup/recovery tasks. The Administration Center includes productivity tools such as a configuration manager, a performance monitor, operation monitoring, backup recovery simulation, and bottleneck analysis for tuning, cloning support, and reporting. These functions are intended to relieve the administrator of repetitive tasks and to help manage and optimize the overall backup and recovery process.

Table 7-5 Summary of Tivoli Storage Manager adapters for DB2 UDB, Oracle, and MaxDB

Adapter or backup tool	**DB2 UDB**	**Oracle RMAN**	**SAP BR*Tools for Oracle**	**SAP MaxDB**
Built-in Tivoli Storage Manager Support	Yes (API)	Not available	Not available	Not available
Tivoli Storage Manager for Databases (Oracle)	Not available	Yes (SBTAPI)	Not available	Not available
Tivoli Storage Manager for ERP (DB2)	Yes (API)	Not available	Not available	Not available
Tivoli Storage Manager for ERP (Oracle)	Not available	Yes (SBTAPI)	Yes (BACKINT)	Not available
ADINT/TSM	Not available	Not available	Not available	Yes (ADINT)

All the backup solutions mentioned can be integrated with the advanced backup techniques such as LAN-free backup, parallel transfer of backup data to and from Tivoli Storage Manager server, or multiplexing. Implementation of these techniques can significantly reduce backup and restore times and eliminate the impact of backup data transfers on LAN throughput

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 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.

A storage area network (SAN) is a high-speed network that connects different kinds of data storage devices, such as disks subsystems, tape libraries, and juke boxes with associated data servers. The primary purpose of the SAN is to transfer data between systems and storage elements. For more information about SAN, see *Introduction to Storage Area Networks and System Networking*, SG24-5470. The SAN feature with the storage agent provides a great opportunity for lowering the backup window, reducing the traffic on the LAN, and reducing the use of the IBM Tivoli Storage Manager server. You can install the storage agent on a client machine that shares storage resources with the Tivoli Storage Manager server or on a client machine that does not share storage resources but is connected to a client machine that does share storage resources with the Tivoli Storage Manager server.

For more information, see *IBM Tivoli Storage Manager for SAN for AIX: Storage Agent User's Guide*, SC32-0129.

SAP HANA

New and changed features for Version 6.4 are available only in the Data Protection for SAP HANA component. The Data Protection for SAP for Oracle and Data Protection for SAP for DB2 components remain at the Version 6.3 level.

IBM Tivoli Storage Manager for Enterprise Resource Planning is a simple, scalable data protection solution for SAP HANA and SAP ERP. SAP HANA's memory-to-disk backups create export files that must be preserved in case recovery is required. Tivoli Storage Manager for ERP includes a one-step command that automates SAP HANA backup and Tivoli Storage Manager data protection. This is a simple process that preserves SAP HANA data safely in Tivoli Storage Manager. Full and incremental (log only) backups can be performed, so you can make frequent backups throughout the day to help reduce the risk of data loss. Parallel multiplexed sessions can be configured for faster backup processing as the SAP HANA environment grows.

Recovery is also simplified. The most recent Tivoli Storage Manager managed SAP HANA backup files can be recovered to their home directory with one command. Alternately, administrators can select from a list of available backups or restore to alternate systems. Tivoli Storage Manager for ERP reduces complexity by enabling all files associated with a SAP HANA backup to be restored as a single unit. Tivoli Storage Manager data protection solutions can help clients implement consistent, integrated backup policies for SAP HANA, SAP ERP, GPFS and other critical components.

For information about the Data Protection for SAP HANA component, see this location:

http://www.ibm.com/support/knowledgecenter/SSZHVN_6.4.1/com.ibm.itsm.erp.doc/welcome.html

Additional information

For more specific information, see the following Tivoli Storage Management web pages:

- Tivoli Storage Manager for Enterprise Resource Planning:

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

- IBM ADINT/TSM

  http://www-05.ibm.com/de/entwicklung/adint_tsm/index.html

- Tivoli Storage Manager for Databases

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

7.5.6 Tivoli Storage FlashCopy Manager

IBM Tivoli Storage FlashCopy Manager delivers high levels of data protection for business-critical databases and applications through integrated application snapshot backup and restore capabilities.

The applications supported by Tivoli Storage FlashCopy Manager are IBM DB2, Oracle, Microsoft Exchange, and Microsoft SQL Server, and the VMware vSphere platform. In addition, IBM DB2 and Oracle databases are supported for use either with or without SAP environments. Other applications can be supported on IBM AIX, HP-UX, Linux, Solaris, and Microsoft Windows platforms with script customization.

These capabilities are achieved through the utilization of advanced storage specific hardware snapshot technology to help create a high-performance, low-impact application data protection solution. Tivoli Storage FlashCopy Manager is easy to install, configure, and

deploy. It seamlessly integrates with various storage systems such as IBM Storwize V7000, IBM System Storage SAN Volume Controller, IBM System Storage DS8000, and IBM XIV Storage System products.

In addition to the above devices, Tivoli Storage FlashCopy Manager on Windows supports any storage system that is VSS capable, by using the Microsoft Volume Shadow Copy Services (VSS) system provider or a VSS Hardware Provider that strictly adheres to the Microsoft VSS Provider interface.

Challenges the solution addresses

Large databases and applications bring a series of challenges that are mostly related to RTO and RPO. Backup takes too long and backup window is not sufficient anymore. As a consequence, recovery takes too long, risking service level agreements (SLAs) and service availability. Large databases are often core applications, which are used by too many people within and outside of the organization and cannot have a downtime.

Tivoli Storage FlashCopy Manager performs and manages frequent application-aware snapshot backups, using advanced technologies in IBM storage systems. Application-consistent backups are performed in minutes without taking the application off-line, and restores can be performed in just a few minutes, instead of hours.

Client benefits

Tivoli Storage FlashCopy Manager can help with these challenges in the following ways:

- Performing and managing frequent, near-instant, non-disruptive, application-aware backups and restores.
- Using advanced snapshot technologies in IBM storage systems.
- Improving availability of IBM DB2, SAP, Oracle, Microsoft Exchange, Microsoft SQL and other application data
- Supporting application development, data mining, and so on through database cloning
- Integrating seamlessly with Tivoli Storage Manager to enable a full range of data lifecycle and recovery management.

Solution architecture

Tivoli Storage FlashCopy Manager is available for three platforms:

- Tivoli Storage FlashCopy Manager for UNIX and Linux:

  Tivoli Storage FlashCopy Manager uses the copy services capabilities of intelligent disk subsystems to create point-in-time copies. These are application-aware copies (FlashCopy or snapshot) of the production data.

  This copy is then retained on disk as backup allowing for a fast restore operation (Flashback). Tivoli Storage FlashCopy Manager also allows mounting the copy on an auxiliary server (backup server) as a logical copy. This copy (instead of the original data on the production server) is made accessible for further processing. This processing includes creating a tape backup or performing backup verification functions (for example, the Database Verify Utility).
- Tivoli Storage FlashCopy Manager for Windows:

  Tivoli Storage FlashCopy Manager provides the tools and information needed to create and manage volume-level snapshots of Microsoft SQL Server and Microsoft Exchange server data. Snapshots are created while the applications remain online.

- Tivoli Storage FlashCopy Manager for VMWare:

  Tivoli Storage FlashCopy Manager is a data management solution that you can use to streamline storage management in a VMware vSphere environment. This application can back up VMware environments from Linux-based backup servers, by integrating with VMware vSphere APIs and hardware snapshot mechanisms. Tivoli Storage FlashCopy Manager for VMware optionally integrates with Tivoli Storage Manager for Virtual Environments to store VMware image backups on Tivoli Storage Manager server storage

Figure 7-29 shows the components of the Tivoli Storage FlashCopy Manager. Tivoli Storage FlashCopy Manager can now be used with other storage vendors such as EMC by leveraging the Rocket Device Adapter for IBM Tivoli Storage FlashCopy Manager. For more information about the Rocket Device Adapter see 2.2.5, “Tivoli Storage FlashCopy Manager” on page 23.

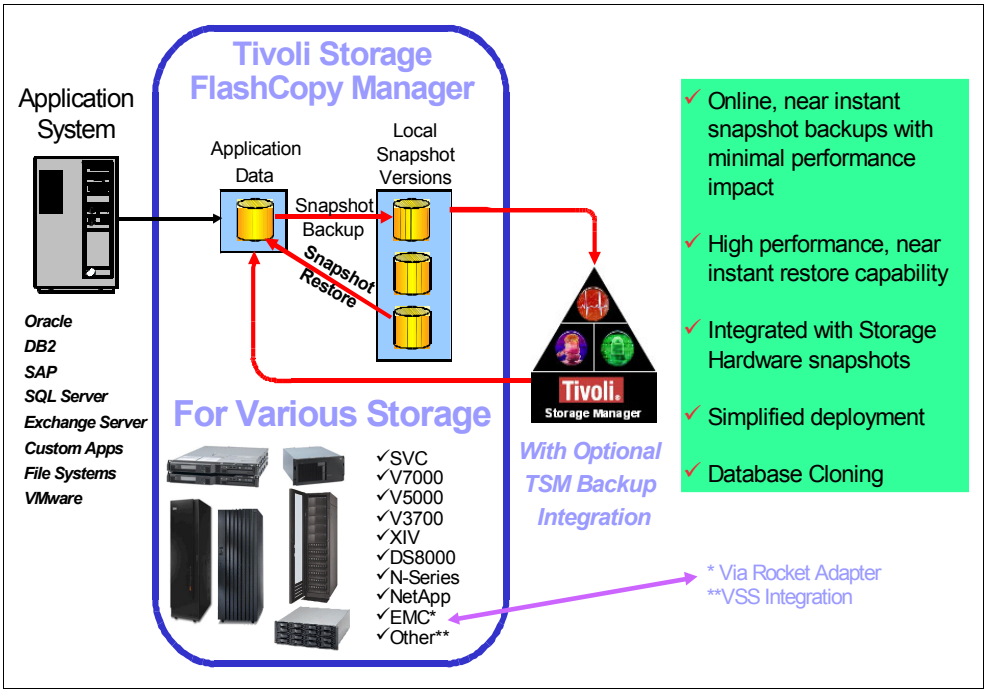

Figure 7-29 Tivoli Storage FlashCopy Manager solution overview

Solution description

This section provides a description about Tivoli Storage FlashCopy Manager for Windows and Tivoli Storage FlashCopy Manager for UNIX. For information about Tivoli Storage FlashCopy Manager for VMware see 7.1, “Common virtualization challenges” on page 142.
Tivoli Storage FlashCopy Manager for Windows

Tivoli Storage FlashCopy Manager for Windows supports five VSS backup types for Microsoft Exchange Servers: full, copy, incremental, differential, and copy without integrity check. With each backup type, there are options to specify the recovery preferences handled by Exchange after the restore. Depending on the backup type selection, Exchange performs tasks such as deleting committed log files, integrity checking (ESEUTIL), log replay, and mounting the databases.

In Local Continuous Replication (LCR), Cluster Continuous Replication (CCR), Standby Continuous Replication (SCR) environments, and Exchange Server 2010 and 2013 DAGs, log truncation is delayed until the Exchange Server knows the changes have been sent to the replica or database copy location.

Tivoli Storage FlashCopy Manager supports only full VSS backup types for Microsoft SQL Server.

Tivoli Storage FlashCopy Manager for Windows provides four types of VSS restores:

- Both Exchange and SQL (three types):
  - VSS Restore: This restores VSS backups from a remote Tivoli Storage Manager server.
  - VSS Fast Restore: This restores VSS backups from local shadow volumes using file level copy mechanisms. In general, restores conclude within minutes instead of hours.
  - VSS Instant Restore: This restores by copying snapshots from IBM Tivoli Storage FlashCopy Manager volumes back to the original source volumes with hardware-assisted volume-level copy mechanisms. The application returns to normal operation as soon as the volume-level copy starts and the log replay is complete.
- Exchange only (one type):
  - Mailbox Restore from VSS backups: This provides granular recovery of any mailbox or item from IBM Tivoli Storage FlashCopy Manager for Windows VSS-based snapshot backups. The ability to restore individual mailbox and granular mailbox items is a new feature for Microsoft Exchange Server 2007 or later using IBM Tivoli Storage FlashCopy Manager for Windows. IBM Tivoli Storage FlashCopy Manager for Windows maintains mailbox location history. For backups taken with prior versions to Data Protection for Exchange 6.1, no mailbox location history is available. When restoring from these prior version backups, if the mailbox to be restored from has been moved or deleted since the time of the backup, the /mailboxoriglocation parameter is necessary.

VSS restore considerations

Be aware of the following considerations when performing VSS restores. Unless otherwise specified, *VSS restores* refers to all restore types that use VSS (VSS Restore, VSS Fast Restore, and VSS Instant Restore):

- If restoring a CCR database, the cluster database is mounted successfully. However, due to a Microsoft Exchange Server 2007 limitation, the database resources are not brought online. Bring the database resources online using the Microsoft Cluster Administrator interface. See the following Microsoft article for details regarding this limitation:

  http://support.microsoft.com/kb/938442/en-us

- When performing a VSS Instant Restore in a CCR or DAG environment, stop the Microsoft Exchange Replication Service on both the active node and the passive node before running the restore operation.
- Performing any type of Restore Into function automatically disables VSS Instant Restore. If performing a VSS restore of a storage group that has been relocated (system file path, log file path, or database file path), use the Restore Into function and specify the same storage group name as the one being restored. The restore will fail if the same storage group name is not specified.

  Important: An attempt to perform a VSS restore into a Recovery Storage Group on Exchange Server 2003 will ignore the Recovery Storage Group and be placed directly into the production database.

- A VSS Instant Restore overwrites the entire contents of the source volumes. However, overwriting the source volumes can be avoided by selecting the Disable VSS Instant Restore option. This option bypasses volume-level copy and uses file-level copy instead to restore the files from a VSS backup that resides on local shadow volumes.

- Unlike Legacy restores (which only dismount the database being restored), VSS restores dismount all databases in the storage group being restored. This is a Microsoft requirement.
- When a VSS restore from local shadow volumes is performed, the bytes transferred will display zero (0). That is because no data (0) is restored from the Tivoli Storage Manager server.
- When performing a VSS Instant Restore, restore all databases within the specified storage group. A partial restore (/partial) cannot be performed while using VSS Instant Restore. Although Data Protection for Exchange allows this operation to begin, it will either fail or complete with undesirable consequences. If only one database from a VSS backup that resides on local VSS shadow volumes is needed for restore, make sure to select the Disable VSS Instant Restore option in the Data Protection for Exchange GUI Restore Window. If VSS Instant Restore capability is needed for single databases, make sure to place these databases in their own storage group.

Tivoli Storage FlashCopy Manager for UNIX

The applications shown in Figure 7-30 are the key components of the Tivoli Storage FlashCopy Manager for UNIX components.

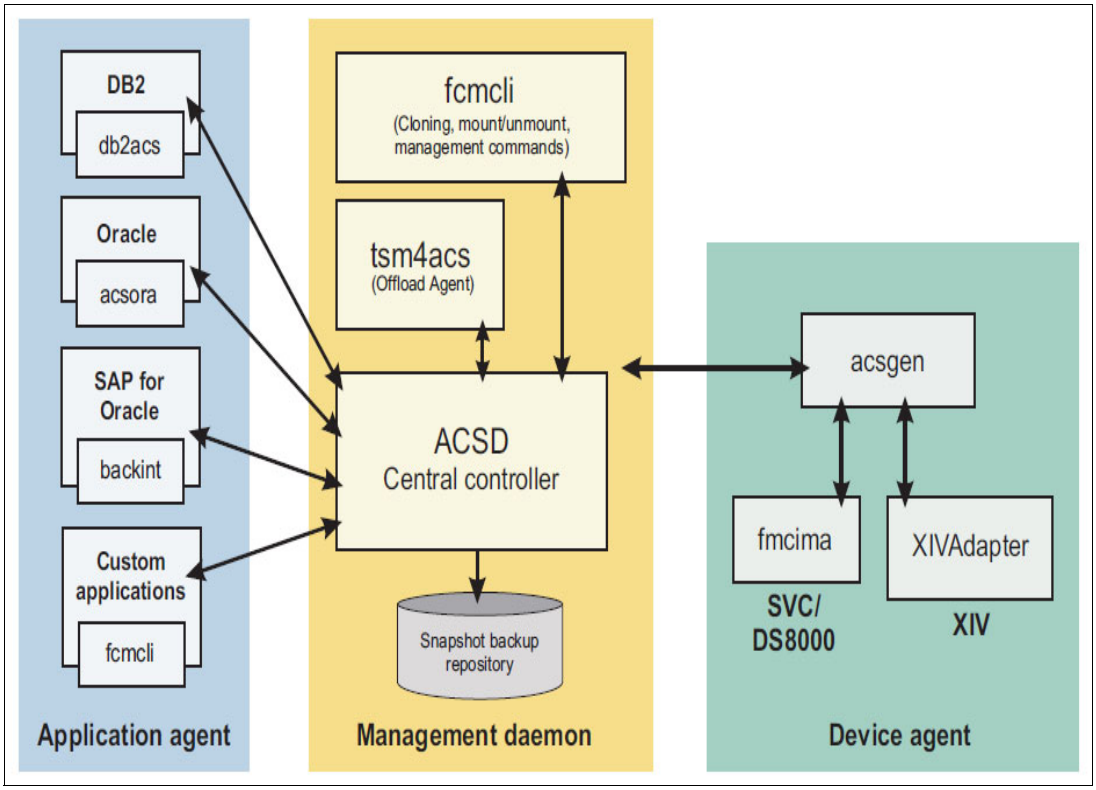

Figure 7-30 Tivoli Storage FlashCopy Manager for UNIX components

Application agent

The Application client provides the necessary support for implementing snapshot-based backup and restore operations.

- (DB2) The client is implemented as the Snapshot Backup Library (referred to as a vendor library in DB2 terms). The library is also a component of Tivoli Storage FlashCopy Manager and is invoked by using the *use snapshot* phrase in the **db2 backup database** or **db2 restore database** commands.
- (Oracle, SAP with Oracle) The client functions are **acsora** or **backint**.
- (Custom applications) Tivoli Storage FlashCopy Manager for Custom Applications provides custom application support with the Tivoli Storage FlashCopy Manager command line interface, **fcmcli**.

Management Agent (acsd)

The Management Agent (**acsd**) coordinates the backup operation. It controls the backup flow and mediates between the application and device agents. The Management Agent also provides access to the snapshot backup repository which contains information about the valid snapshot backups and their relationships to snapshot-capable storage devices.

Device Agent for Generic Devices (acsgen)

The Device Agent for Generic Devices (**acsgen**) is an operating system independent and storage device independent software layer that interacts with operating system specific and storage device-specific adapters. This agent is also used to send and request updates of the progress and usability information that is stored in the local snapshot backup repository.

CIM Adapter (fmcima)

The CIM Adapter (**fmcima**) is used with the Generic Device Agent (acsgen). It is the component that invokes a snapshot command on a FlashCopy device (such as DS8000, Storwize V7000, and SAN Volume Controller) using the CIM interface.

XIV Adapter Oracle Java Archive (XivAdapter.jar)

The XIV Adapter (`XivAdapter.jar`) is used with the Generic Device Agent (acsgen). It communicates with acsgen and issues commands to the XIV command-line interface (XCLI).

Query Capacity (fmquery)

The Query Capacity (**fmquery**) command lists all backups (FlashCopy or snapshot backups) that are registered in a particular repository. Use this command to periodically check the amount of storage space used for backups and to verify compliance with the licensed capacity amount.

Volume Group Takeover script (acsvg.sh) for AIX only

The Volume Group Takeover utility (**acsvg.sh**) is a shell script. It is required only in special high-availability scenarios where enhanced concurrent capable volume groups are used on production systems. In these situations, this script exports and reimports the volume groups on an HACMP takeover system after a snapshot restore is performed. This process is necessary in order to synchronize the AIX Object Data Manager (ODM) on the production and HACMP takeover systems.

Offload Agent (tsm4acs)

The primary role of the Offload Agent is to provide a single user interface for backing up an existing snapshot to Tivoli Storage Manager. Tivoli Storage FlashCopy Manager includes a license file that enables the use of the enhanced functions of the Offload Agent. The Offload Agent also calls the generic device agent for mount and unmount operations on the backup systems.

Tivoli Storage FlashCopy Manager command line interface (fcmcli)

The Tivoli Storage FlashCopy Manager manager executable file, `tsm4acs`, is also the cloning interface for Oracle, DB2, DB2 DPF databases in SAP and non-SAP environments on AIX, Solaris, xLinux and HP-UX. For further details on Tivoli Storage FlashCopy Manager see the Tivoli Storage FlashCopy Manager wiki:

http://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli%20Storage%20FlashCopy%20Manager/page/Home

Tivoli Storage FlashCopy Manager use scenarios

Table 7-6 shows the various uses of Tivoli Storage FlashCopy Manager with databases, mail and application products.

Table 7-6 Tivoli Storage FlashCopy Manager features

Feature	Benefit
Provide advanced, granular restoration of Microsoft Exchange data	► Snapshot restore of Microsoft Exchange storage groups ► File copy restore of a storage group or database from a mounted snapshot image ► Restore into a recovery storage group, alternate storage group, or relocated storage group ► File copy restore of transaction logs from incremental or differential backups ► Single or multiple user mailbox support ► Restore selectable based on user name and date/time specifications ► Recover any mailbox items, including Inbox, Deleted Items, Drafts, Outbox, Sent Items, Journal, Calendar, Contacts, Notes, Tasks or User Folders ► Limit scope of restore based on a range of variables, such as Sender Name, Subject Texts, Attachment Name, Folder Name, Message Delivery Date and Time Range, Message Body and all content (searches Subject Text, Message Body, and Attachment names) ► Recover data into the original mailbox or alternate mailbox and folder on the production Exchange server, or into a .PST file
Provide advanced protection and restoration of Microsoft SQL databases	► Snapshot restore of a full database backup ► File copy restore of a full database to an alternate database or location
Provide advanced protection and restoration of DB2 databases on UNIX platforms	► Snapshot backup and restore of a full database including and excluding log files ► Backup and restore of individual partitions for multi-partition DB2 databases ► Support of databases that are mirrored (between sites) using Logical Volume Manager (LVM) mirroring technology
Provide advanced protection and restoration of Oracle databases on UNIX platforms	► Snapshot backup and restore of a full database ► Support for databases that are mirrored (between sites) using LVM mirroring technology ► Support for Oracle ASM configurations and failure groups
Provide snapshot backup and restore of a full SAP database running on DB2 or Oracle on UNIX platforms	► Support of databases that are mirrored (between sites) using LVM mirroring technology

Feature	Benefit
Make copies of active databases (cloning) on UNIX platforms	► Facilitate database application development and testing ► Provide database copies for other purposes, such as data mining, migrations, quality assurance and education
Provide customization capabilities on UNIX platforms	► Enable snapshot backup and restore of custom applications that are not natively supported by Tivoli Storage FlashCopy Manager, for example MySQL
Support for VMware vSphere 4.1 and later	► Leverage hardware snapshots in virtualized server environments ► Integrated with vCenter client interface for easy management
Integration with Tivoli Storage Manager	► Extend data protection and retention off the source storage system for long-term data management and offsite disaster recovery capabilities ► Take advantage of automated data reduction and data management policies that can reduce overall costs

Integration with Tivoli Storage Manager

Tivoli Storage FlashCopy Manager can back up data from a remote system (backup server) to Tivoli Storage Manager. These components must be installed and configured on the backup server in order to back up to Tivoli Storage Manager:

- IBM Tivoli Storage Manager for Enterprise Resource Planning (SAP with DB2, SAP with Oracle)
- The DB2 native Tivoli Storage Manager agent (DB2 in non-SAP environments)
- Tivoli Storage Manager for Databases (Oracle in non-SAP environments)

Tivoli Storage Manager Backup-Archive Client (custom application environments) Tivoli Storage FlashCopy Manager provides these functions with Tivoli Storage Manager; see Figure 7-31 on page 230:

- Back up to Tivoli Storage Manager immediately after the Tivoli Storage FlashCopy Manager backup completes successfully.
- Perform the Tivoli Storage Manager backup with a separate schedule. This function allows delaying the backup to Tivoli Storage Manager to a time when the availability of tape drives is at its best.
- Manually restart a backup to Tivoli Storage Manager after an error. In this situation, data that has already been committed on the Tivoli Storage Manager server is not sent again.

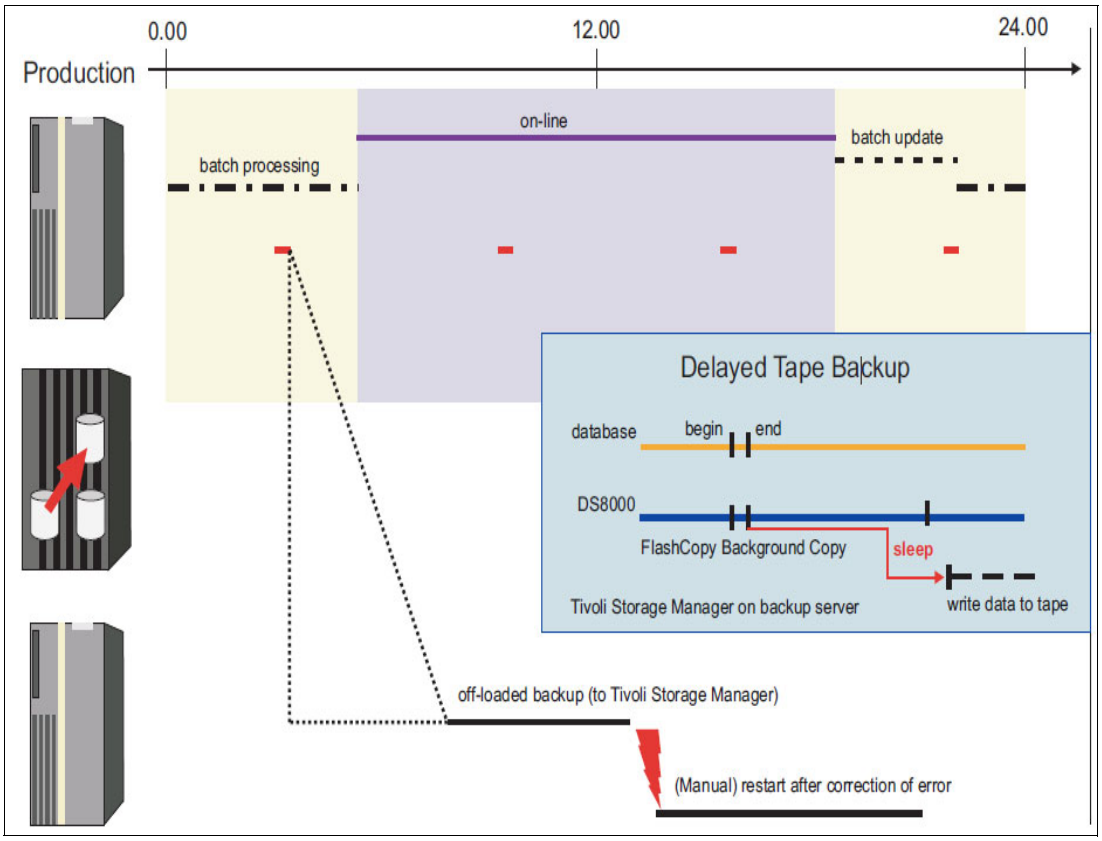

Figure 7-31 Tivoli Storage FlashCopy Manager environment when integrated with Tivoli Storage Manager

Hardware and software requirements

The following document contains links to all of the Tivoli Storage FlashCopy Manager hardware and software requirement documents:

http://www.ibm.com/support/docview.wss?uid=swg21427692

Additional information

These resources have examples of Tivoli Storage FlashCopy Manager implementation.

- *IBM Tivoli Storage FlashCopy Manager 2.2 for Windows Backup & Recovery Solution for Microsoft SQL Server 2008 and Exchange Server 2010 on the IBM XIV Storage System*:

  http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101840

- *SAP with IBM Tivoli Storage FlashCopy Manager for VMware and IBM XIV and IBM Storwize V7000 storage systems*:

  http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102290

7.6 Big data: Structured, very large databases

Through the years the definition of very large database (VLDB) has changed. A 50 GB database was considered very large 20 years ago. These days, a new database deployment in a large company might start at 1 TB. It is not hard to find databases running more than 10 - 100 TB and even PBs for data warehousing or online transaction processing (OLTP) activities. The big challenge is to protect large volume of data sets in a timeframe that is ever shrinking due to stringent service level requirements

So how do you back up very large databases? The answer is by careful planning. We need to find ways to complete the backup and the restore within a shortened time frame. As the data grows the backup window and recovery time object can become too long and new technologies need to be put into play to satisfy business needs.

So the challenge is to have a solution to back up and restore data as quickly as possible, with lowest impact on the production data. This should be done at the lowest cost even when you experience significant data growth.

The matrix Figure 7-32 shows which Tivoli Storage Manager solutions that can address these challenges when protecting very large databases.

Challenges										
business					addressed with	technical				
reduce recovery time (RTO)	reduce backup time (meet RPO)	Compliance	Business Continuity	security (secure data transfer, secure data store,	Solution (discussed in this book)	Virtual Environments	Significant data growth	automation for business	Remote office and mobile office data	network bandwidth optimization
X	X				Disk to Disk backup	X		X	X	
X	X		X		Data reduction with ProtecTier	X	X		X	X
X	X	X	X	X	TSM ToolKit	X	X	X	X	X
X	X		X		Hardware Snapshot technology	X	X			
X	X				deduplication	X	X		X	X
	X		X		off-load backup	X	X			
	X				Off-Host backup	X	X			
X	X				Online backup of Applications					
X	X				LAN-free data transfer		X			X
	X				GPFS		X		X	
X	X				Tiered Storage Hierarchy	X	X			

Figure 7-32 Challenge matrix for big structured data

To meet the challenges we need to look at the drivers to move the data from the primary production storage to the Tivoli Storage Manager back-end storage. For this particular type of data we look at the tired storage hierarchy which can be a disk to disk, disk to tape, or a disk to VTL solution. In our Tivoli Storage Manager toolkit, we will make use of hardware-based snapshot technology and LAN-free data transfer. Because VLDBs do not change the data structure, we can lower the footprint by using different ways of deduplication functions.

Figure 7-33 shows the difference in RTO and RPO when using various technologies: FlashCopy to disk or software-based backup to tape or VTL. For instance a FlashCopy backup of the database can be done in minutes but would take hours to tape. And a FlashCopy backup is not dependent of the size of the data. The restore process is performing the same way.

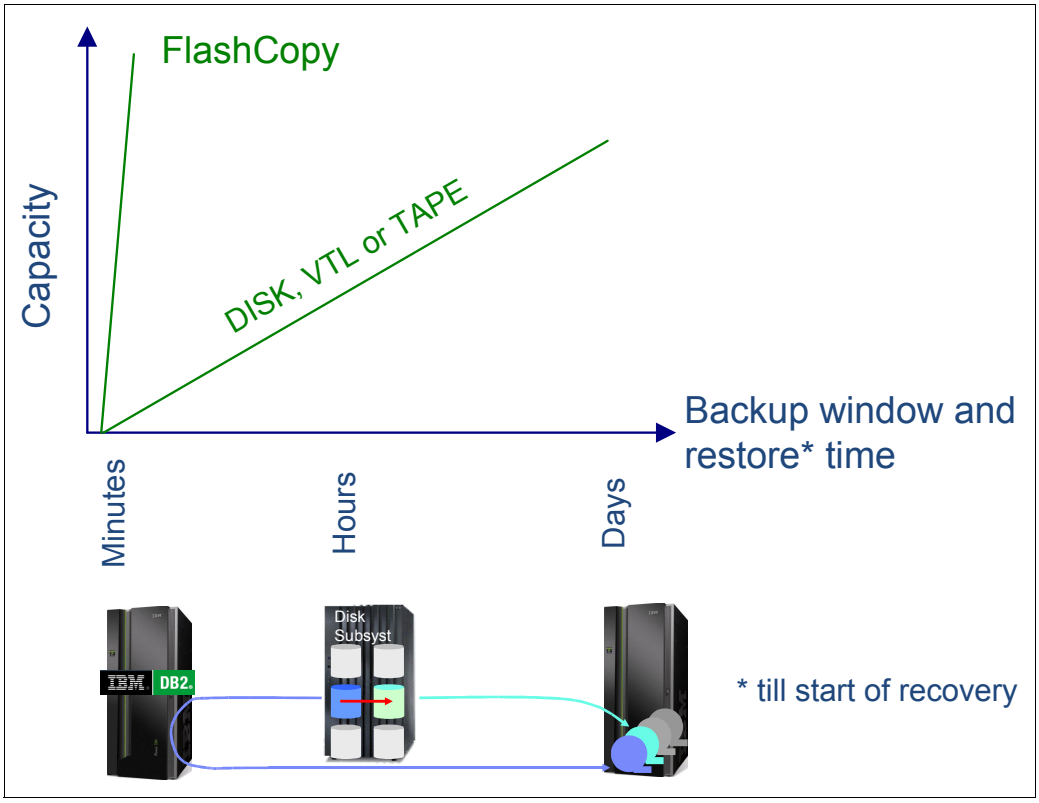

Figure 7-33 Comparing FlashCopy to disk, VTL and tape

In this section, we describe what options exist to protect large amounts of data stored in structured databases with a Tivoli Storage Manager solution.

There are discussions about which solution is best for backup data in general, whether it should be disk-based or tape-based. Some reports might indicate VTL as being the most cost efficient and others indicate tape. In reality, this must be considered case by case. Maybe we need not choose at all. A combination of disk and tape might be the optimal strategy for many storage administrators to address various needs. It all comes down to good planning and a fundamental understanding of the data to be protected and also what the backup window and recovery time objectives are.

A VTL or disk-based backup can provide the performance that is needed to prioritize the recall of files for high risk applications. As data backed up to disk becomes infrequently or never accessed, it should be moved to tape for long-term retention. You must define what long term means to your company. Tape technology can provide data security, compliance, and offline protection (against viruses, hackers, system errors, and so on) and a long term, low-cost archive repository.

7.6.1 LAN-free backups to tape

LAN-free is generally used for backup direct to tape because the tape can be much faster than disk when streaming. Today we have LTO6 that natively can go 160MBps per backup thread and much more with compression.

The LAN-free data transfer concepts are the same for all devices supported for that purpose. We give a short description in this section about LAN-free backups to tape. The concepts will also apply to disk and VTL, which we describe later.

LAN-free backup should be the standard way to backup big data with Tivoli Storage Manager. To achieve the best overall performance, you should consider the type of data on the client. LAN-free backup makes substantial use of the LAN for control and other information exchange. When trying to decide if LAN-free is appropriate for your environment, consider the following common factors:

- A congested network (LAN): This factor includes overall network congestion as well as any network limitations between the client and the server machines.
- Constrained server: Streaming data to tape over the SAN can be faster than using the network, if the client systems have access to the SAN storage resources.
- Existing SAN storage resources
- Type of data to be backed up
- Number of mount points available

Because the LAN-free path is used for sending the actual data, and not the metadata, a client workload, which has proportionately more metadata than data to send, will in general see less benefit from using the LAN-free path. Conversely, a client workload that spends most of its time sending data will see a greater benefit from using LAN-free backup. From this information, you can extrapolate that large file sizes are better suited for LAN-free backup.

Important: In this topic we discuss LAN-free backups to tape in very large database environments. For more information about protecting databases, see 7.5, “Application and email servers” on page 171.

Solution benefits

Using LAN-free to tape should give some immediate benefits:

- Improved performance reducing backup time.
- Decreased recovery time of data compared to transferring data across the LAN.
- LAN-free data movement makes LAN bandwidth available for other uses and decreases the load on the Tivoli Storage Manager server, allowing it to support a greater number of concurrent client connections.
- Exploit SAN infrastructure to relieve LAN.
- Control of data as it grows.
- Support for open system connectivity.
- Heterogeneous operating system platforms can share the same storage devices, which allows for better use of storage and reductions in storage costs.

Solution architecture

This section presents the Tivoli Storage Manager protection for a very large database. The concepts on that architecture can be used for all supported databases and applications. To see supported databases and applications, see 7.5, “Application and email servers” on page 171.

LAN-free architecture

Figure 7-34 shows a typical LAN-free architecture.

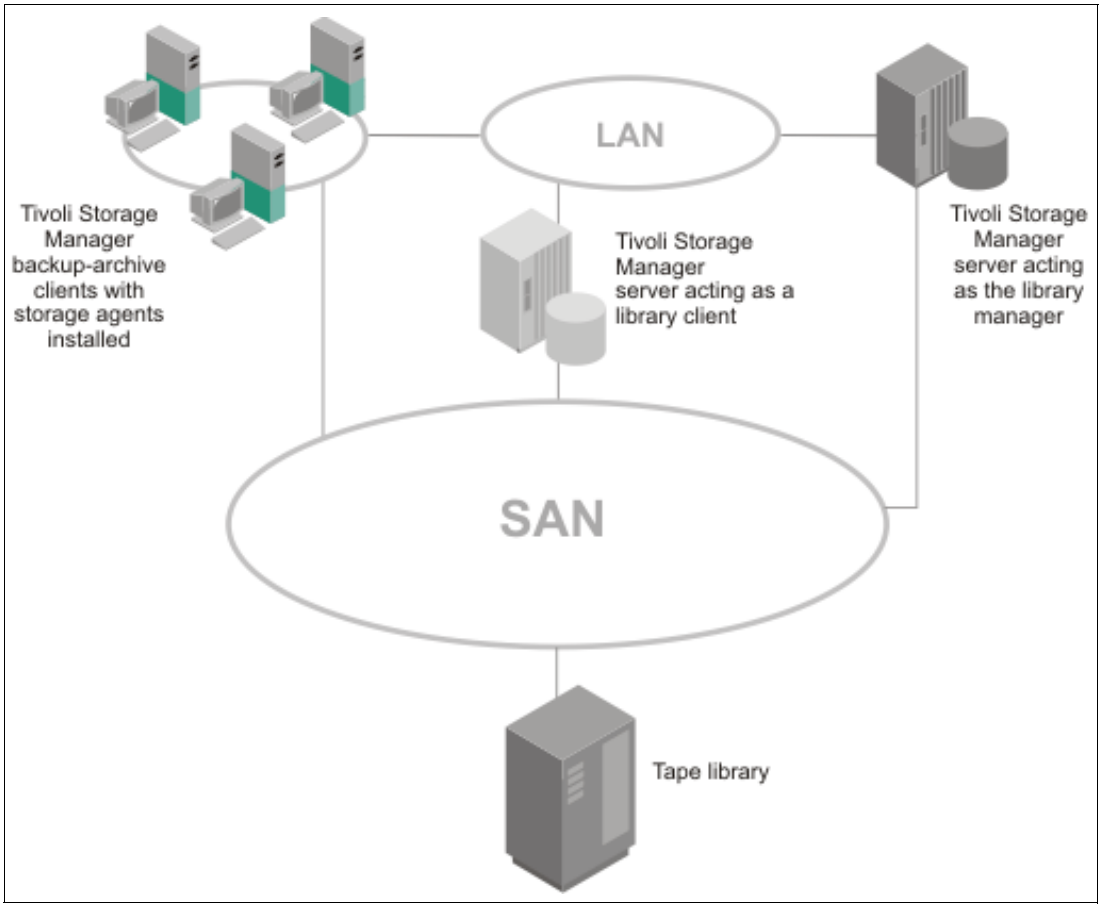

Figure 7-34 LAN-free data movement

The components of the solution are as follows:

- Tivoli Storage Manager Database: This is one of the principal architectural components of the Tivoli Storage Manager solution. All policy information, logging, authentication and security, media management, and object inventory is managed through this database. The database does not store client data; it points to the locations of the client files in the storage pools. The Tivoli Storage Manager Database contains information about the Tivoli Storage Manager server.
- Tivoli Storage Manager 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.
- Tivoli Storage Manager Client: Backup-archive clients are implemented as multi-session clients, which exploit the multithreading capabilities of modern operating systems. Backup and archive operations can run in parallel to maximize the throughput to the server, and parallel restore operations allow faster restore time to be achieved.

- Tivoli Storage Manager API Client: This is used to implement application clients to integrate popular business applications, such as databases or groupware applications, into the Tivoli Storage Management solution. These IBM solutions are named with the prefix of *IBM Tivoli Storage Manager for* (as in these examples: IBM Tivoli Storage Manager for Mail, IBM Tivoli Storage Manager for Databases). The API is also published, which allows customers or vendors to implement specialist clients for special data management needs or non-standard computing environments.
- Tivoli Data Protection Client: The *IBM Tivoli Storage Manager for* products are separate program products that connect business applications to the Tivoli Storage Manager data management API. Such applications (for example, Oracle, IBM Lotus Notes® and Domino, Microsoft Exchange, Microsoft SQL Server, and mySAP) have their own storage management interfaces that are used to interface to Tivoli Storage Manager.
- SAN/ Tape Library: The Tivoli Storage Manager SAN tape resource sharing capability delivers immediate benefits by reducing the traffic on the IP network and enabling shared utilization of resources over a SAN. SANs remove the overhead commonly found with slow, overworked communication networks and facilitate quicker access time. Tape library and drive resources are used more efficiently because they can be shared by multiple Tivoli Storage Manager servers across the SAN.

Configuration preparations

Before you set up a LAN-free configuration, you should have a good understanding of the essential pieces of information and components. This section discusses these pieces.

To assist you in setting up LAN-free data movement, identify or be aware of these items:

- The client machine on which LAN-free is to be set up

  After you decide to use LAN-free backup, consider where to install the Storage Agent. The Storage Agent is installed most usually on the actual client machine, having the Storage Agent running on a separate system is also possible. For informations about having the Storage Agent run on a separate machine, see product documentation:

  http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/topic/com.ibm.itsm.sta.doc/t_tsmsan.html

- The Tivoli Storage Manager server to be used

  The version of the Tivoli Storage Manager server and the Storage Agent must match. You can find the compatibility matrix at the following location:

  http://www-01.ibm.com/support/docview.wss?uid=swg21302789

- The type of library sharing method to be used

  The SAN-attached storage devices that you have (or plan to have) in your environment most likely determines the type of library sharing method you will use. IBM tape devices or non IBM storage devices that are supported by the Tivoli Storage Manager device driver can be used for sharing. These device types include SCSI, ACSLS, external, 349*x* tape libraries and virtual tape libraries (VTL). For more information about LAN-free to VTL, see 7.6.2, “LAN-free backups to VTL” on page 238.

- The device names

  The device names for the SAN-attached storage devices are required when defining the path on the server. SAN-attached devices might appear with different device names on each host that is attached to the same SAN.

- A proper management class destination

For a client node to use the Storage Agent to perform LAN-free data movements, the client node must be bound to a management class that uses the SAN-attached storage device. This management class can be set as the default to allow any client node that is created within the policy domain to use it automatically for all backup operations. Alternatively you can explicitly specify to use the management class by specifying INCLUDE statements within the client options file or client option set that point to the LAN-free management class.

For more information about planning and configuring LAN-free, see the following location:

http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.0/com.ibm.itsm.ic.doc/welcome.html?lang=en

Considerations for very large databases environments

Special considerations for backing up very large databases LAN-free are as follows:

- Tape Drive allocation must be analyzed to prevent Tivoli Storage Manager from exceeding the number of mount points. Your backup window cannot use more drives than you have in you environment. You probably want to parallelize the backup of the database, which will require a tape drive for each backup or restore thread.
- Collocate to improve restore times by reducing the number of tapes needed. Enabling collocation also improves restore times because it allows you to run more restores concurrently. The reason is because the data that is required to restore a particular very large database is on its own media. See “Collocation” on page 241.
- The use of performance parameters for tuning Tivoli Storage Manager and Storage Agent communication is a consideration. For more information go to the following location:

  http://pic.dhe.ibm.com/infocenter/tsminfo/v6r3/topic/com.ibm.itsm.nav.doc/c_performance.html
- Large database objects are ideal for LAN-free data transfer. However the transaction logs might not be adequate for LAN-free transfer because of the metadata sent on the LAN. See the steps for LAN-free communication in Figure 7-35 on page 237. You might also want to make these files available from a disk-based storage pool to be sure the rollforward process of the recovery process does not require tape mounts.

Solution description

The solution in Figure 7-34 on page 234 shows how the data goes across the SAN instead of LAN. SANs provide an alternative path for data movement between the IBM Tivoli Storage Manager client and server. LAN-free data transfer exploits this SAN path by enabling the IBM Tivoli Storage Manager client to back up and restore data directly to and from SAN-attached storage. Storage is shared between the Tivoli Storage Manager library server and client. The Tivoli Storage Manager library manager server is responsible for managing the tape library.

The following steps are involved in a Tivoli Storage Manager LAN-free backup:

1. The Backup-Archive Client begins a backup operation. The Tivoli Storage Manager server reports policy information to the client, including whether a destination is LAN-free. As the client assigns policy settings for files during backup processing, it uses the Storage Agent to send the data via LAN-free when the destination for that policy is LAN-free enabled.
2. The Storage Agent receives data for those files backed up by the client and assigned to policy settings that use a LAN-free enabled storage pool. The Storage Agent sends a request for a volume mount to the Library Manager server.
3. A request is sent from the Library Manager to the storage device to mount the appropriate media.

4. The Library Manager notifies the Storage Agent of the location where the mounted media resides.
5. The client, by means of the Storage Agent, writes the backup data directly to the device over the SAN.
6. The Storage Agent sends metadata information to the Tivoli Storage Manager server over the LAN, and the server stores the information in its database.

Note: LAN-free data movement takes precedence over client-side data deduplication. If LAN-free data movement occurs during client-side data deduplication, client-side data deduplication is turned off, and a message is issued in the error log.

All other tasks, which are metadata-related, use the LAN path. Therefore, depending on the profile of the backup to be performed, the proportion of time spent transferring over the LAN versus the SAN differs.

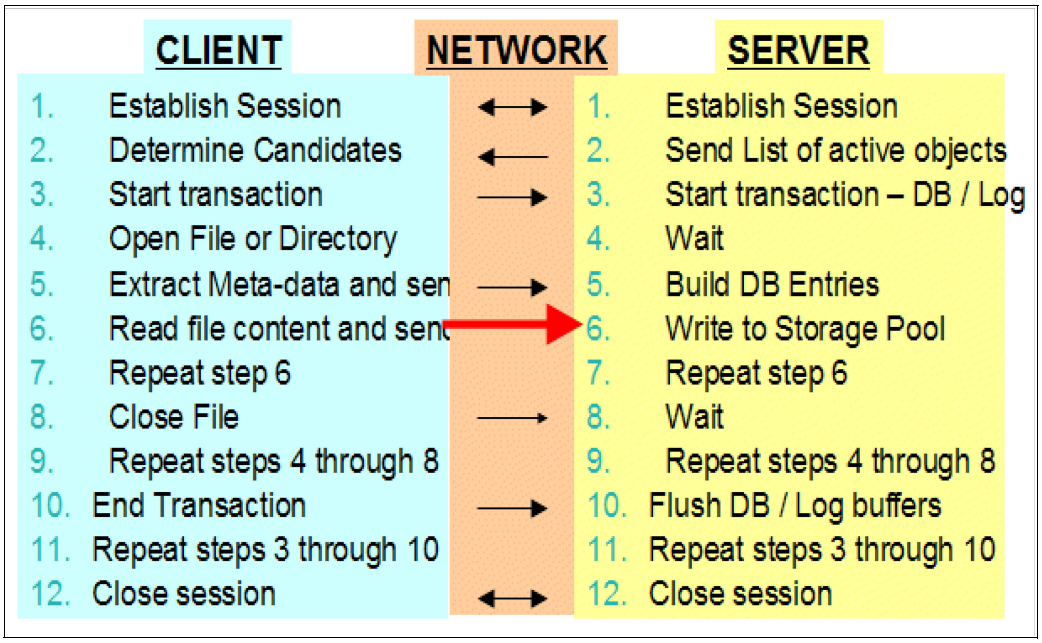

Figure 7-35 Steps for a LAN-free backup

LAN-free is typically ideal for transferring large amounts of data, such as when you backup a database. However, LAN-free data movement is not optimal for short bursts of data or for small files, such as archived log files. For that reason, do not use LAN-free data movement for archived transaction log files.

Hardware and software requirements

The hardware and software requirements for LAN-free support are shown in Table 7-7.

Table 7-7 Hardware and software requirements for LAN-free client support

Servers	Backup-archive clients	Tivoli Data Protection
AIX	AIX	R3
Solaris	Solaris	Oracle
Windows	Windows	Domino
HP-UX	HP-UX	Informix
Linux	Linux	Microsoft Exchange
z/OS		MS-SQL

See the Tivoli Storage Manager Server-Client Compatibility and Upgrade Considerations technote:

http://www.ibm.com/support/docview.wss?uid=swg21053218

7.6.2 LAN-free backups to VTL

With disk prices decreasing, organizations have started adopting disk as an excellent medium for storing backup data. It can be used as a supplement to tape, which is still cost effective but has the limitation on physical media mount points that might give administrative challenges in a LAN-free backup environment.

Although Tivoli Storage Manager has supported disk storage pools for many years, there are some obvious advantages to using a virtual tape library in place of, or as a supplement to, your disk storage pools. For example, you do not need to assign a dedicated disk pool for every server.

Backup and recovery from disk can solve all of the listed challenges in Figure 7-32 on page 231. An effective disk-based backup and recovery environment requires the application to know how to take advantage of the random access capabilities of disk like the Tivoli Storage Manager application does.

You can use any supported virtual tape library when the following conditions are true:

- There is no mixed media involved in the VTL.
- Only one type and generation of drive and media is emulated in the library.
- Every server and storage agent with access to the VTL has paths that are defined for all drives in the library.

If any of these conditions are not met, any mount performance advantage, from defining a VTL library to the Tivoli Storage Manager server, can be reduced or negated.

VTLs are compatible with earlier versions of both library clients and storage agents. The library client or storage agent is not affected by the type of library that is used for storage. If mixed media and path conditions are true for a SCSI library, it can be defined or updated as LIBTYPE=VTL.

Solution benefits

Using LAN-free to VTL offers several immediate benefits:

- You can create small cartridge sizes to allow as much parallelism as possible.
- Recovery time of data is decreased compared to transferring data across the LAN or LAN-free to physical tapes because of tape mount time.
- Performance is improved because the server handles mount point processing for VTLs differently than real tape libraries.
- LAN-free data movement makes LAN bandwidth available for other uses and decreases the load on the Tivoli Storage Manager server, allowing it to support a greater number of concurrent client connections.
- The physical limitations for real tape hardware are not applicable to a VTL, affording options for better scalability.
- You can create more slots than needed for future growth. Adding cartridges is an online procedure, while changing library dimensions is an offline procedure.
- Open system connectivity is supported.
- Integration with an existing Fibre Channel or tape-based infrastructure is possible.
- Use drawer based deduplication (ProtecTIER) to save space and decrease the load on Tivoli Storage Manager Server.
- Hardware replication allows more options on the DR planning.
- Reliability is improved compared to tapes because you are not dependent on the mechanical operations to be working.

Solution architecture

By emulating an actual tape library, the VTL enables Tivoli Storage Manager to function identically with virtual tape as it does with physical tape, except that it does it much faster and much more reliable.

We are in this solution using the IBM TS7650G ProtecTIER as an example of a LAN-free to VTL environment. Figure 7-36 on page 240 shows the backup servers that are connected to FC Switch and ProtecTIER as a data destination. Backup servers can be either Tivoli Storage Manager servers or Storage Agents.

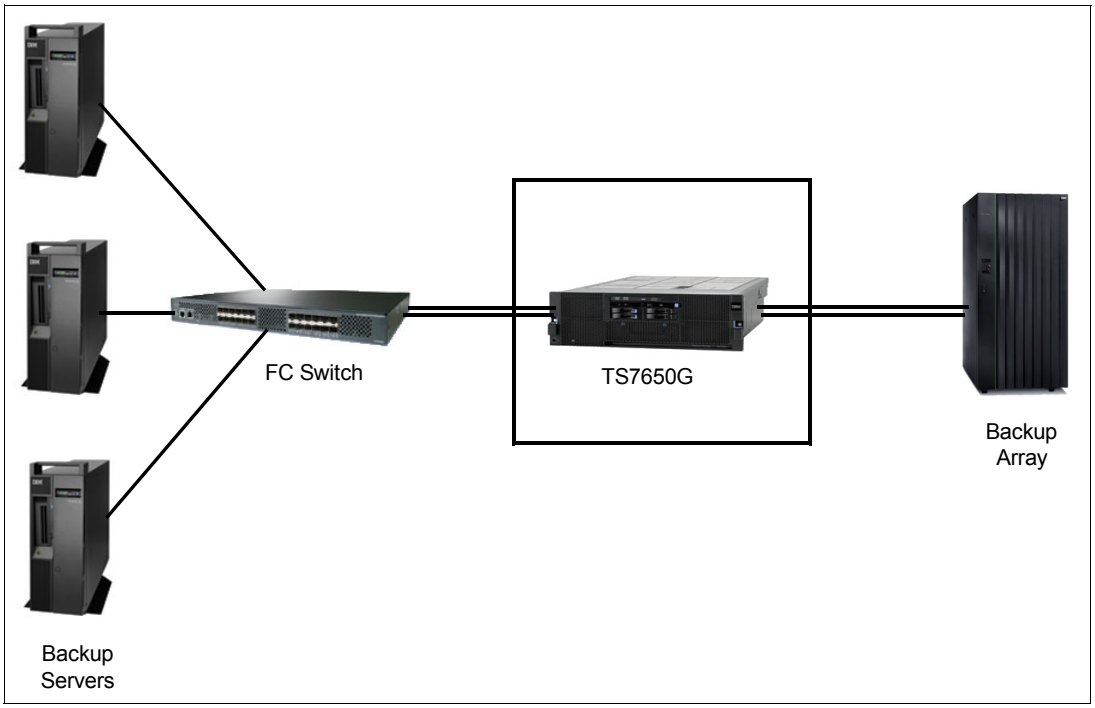

Figure 7-36 Tivoli Storage Manager with ProtecTIER Solution

With the IBM TS7650G, you can either share a virtual tape library between many servers and storage agents or you can create several virtual tape libraries, one or more for each server. This means you are not required to dedicate storage to every server, because the VTL that uses ProtecTIER technology can assign capacity as required.

More important, the TS7650G's unique IBM HyperFactor® data deduplication technology provides up to 25:1 reduction of storage space required, delivering the capacity to retain months of backup images at a total cost of ownership (TCO), less than that of actual tape.

The TS7650G ProtecTIER deduplication technology was designed to work seamlessly with Tivoli Storage Manager to offer remarkable improvements in speed and reliability, and to greatly reduce the TCO of a disk-based backup and recovery environment.

Solution description

The solution provided uses ProtecTIER as the VTL with deduplication enabled. Combining the advanced capabilities and features of Tivoli Storage Manager with the powerful capabilities of the ProtecTIER product provide IT organizations a cost-effective way to improve the performance, reliability, and scalability of data protection.

LAN-free backups are simpler with the ProtecTIER product because there are increased tape resources and fewer hardware restrictions. Because ProtecTIER is a VTL, it has the advantage of presenting greatly increased tape resources to the backup server. This feature is available with only Tivoli Storage Manager V6.3 or later. So, you are able to perform LAN-free backups to the ProtecTIER server without the limitations normally applied to these backups, such as tape drive availability. If you have many LAN-free clients already, then it is possible that your LAN-free backup windows are dictated not entirely by business needs but also by hardware availability. With the ProtecTIER product and its maximum of 256 virtual tape drives per ProtecTIER node, you can virtually eliminate any previous hardware restrictions, and schedule your backups as and when they are required by your business needs.

LAN-free backup to VTL sends the data over SAN in the same way as LAN-free backup to tape does. The existing LAN connection is used to exchange control information such as policy information and meta data about the objects being backed up, but the data movement utilizes the SAN to write directly to the storage media. If there is a failure when using the SAN path, the LAN-free client recovers from these errors by failing-over to a LAN connection to the IBM Tivoli Storage Manager server and proceeds to move data over the LAN network.

The backup servers (Tivoli Storage Manager Server or Tivoli Storage Manager Storage Agent) must have a dedicated host bus adapter (HBA) port or ports for the ProtecTIER VTL. This port or ports can be shared with a physical tape library. However, the physical tape library must not be in the same storage area network (SAN) zone as the VTL. When it is not possible to dedicate HBA ports for VTL and physical tape library, you should have different zones to separate the traffic.

Port sharing: Although sharing a Fibre Channel (FC) port between physical and virtual tape is possible when they are in different SAN zones, you should avoid sharing a port with a SAN-attached disk. For more information see the following technote:

http://www.ibm.com/support/docview.wss?uid=swg21194590

Increasing backup performance and flexibility

In a tape-based data protection environment, storage administrators must carefully tune Tivoli Storage Manager to stream just the right amount of data to a tape drive. Streaming too little data can result in “shoe shining,” an effect of starting, stopping, and repositioning the tape that results in poor drive performance. Multiplexing and parallel sessions are used to make sure enough data is being sent to maximize tape drive utilization, compensating for slow clients, slow networks, or small backup jobs.

The ProtecTIER frees storage administrators from shuffling backup jobs and tuning tape libraries and operations for optimal performance. Regardless of the amount of data being streamed, virtual tape libraries back up data at the speed of the target disk.

In addition, although virtual tape libraries accept and process all tape operation commands, they are not hit with the delays associated with allocating tape cartridges, robotic movements, media mounting, media positioning, or media eject operations. In a tape-based environment, these delays can last from a few seconds to a few minutes per tape.

Increased backup performance allows IT organizations to complete backups within allowed backup windows and meet service level requirements of mission-critical applications.

You might be able to reduce your current backup window by taking full advantage of the throughput performance capabilities of the ProtecTIER product. It is easy to define a greater number of virtual drives and to schedule backups to run at the same time to maximize the number of allowable parallel tape operations on ProtecTIER servers.

Collocation

Collocation means that all of the data for a node or node group is contained on the same set of virtual cartridges. Because you do not have any of the restrictions as you do with physical cartridges normally associated with this feature (such as media and slot consumption), you can enable the option with benefit.

Preferably, collocate data with similar expiration characteristics. This practice can help with the following tasks:

- Minimizing reclamation processes
- Reducing the Tivoli Storage Manager workload
- Reducing the risk of replicated cartridges being out of synchronization because of the timing of the reclamation activity
- Increasing the likelihood of a cartridge being directly available for restore tasks

Increasing restore speed

Multiplexing is used to ensure a constant stream of data flow to each tape drive in order to reach optimal drive performance. Multiplexing is often the only way to back up all data in a backup window with the tape drive resources provided.

Taking advantage of Tivoli Storage Manager multiplexing capabilities and parallel sessions increases backup performance but has a devastating impact of slowing down restore performance significantly. This is primarily due to the extra time required to read through backup images and skipping over unwanted data from other backup jobs combined into that image.

The VTL eliminates the delays associated with moving, loading, and mounting of cartridges and the seek times inherent to physical tape.

By recovering data from VTLs at the speed of disk, Tivoli Storage Manager can deliver an order of magnitude improvement in recovery performance and, more important, reduce downtime and its associated costs.

Avoiding mount conflicts

To avoid a mount conflict, increase the number of drives (according to your need) up to 512 per dual-node cluster (256 per node). Depending on your Tivoli Storage Manager version or operating system, these maximum values might change.

Accommodating increased sessions

Ensure that the MAXSESSIONS setting on the Tivoli Storage Manager server can accommodate the increased sessions. Also, update the NODE definition on the Tivoli Storage Manager server to allow more than one mount point (MAXNUMMP).

For more information, see the following website:

http://pic.dhe.ibm.com/infocenter/tsminfo/v6r3/topic/com.ibm.itsm.srv.ref.doc/r_cmd_node_update.html

Migrating data: Do not expect an effective deduplication when you migrate your existing data from physical tape to the ProtecTIER repository if the data was originally backed up without suggested practices in place. Use the most current version of the ProtecTIER product so that you implement the appropriate Tivoli Storage Manager parser, which maximizes your overall deduplication factoring ratio.

Reducing costs

Most TCO studies compare only the cost of automated tape libraries and their associated media to the cost of large disk storage systems with virtual tape library servers and software. These hard costs are only part of the total equation. These TCO studies are weak on operational cost and resort to making generalized assumptions. Strategic Research Corporation recently completed a study that includes hard costs as well as the cost to administer failed backups (that first-time backup error rate problem that tape struggles with),

labor cost of recovery, vaulting, tape management, etc. The research data shows an average 30 % reduction in the operational costs of administrating backup and recovery across the 10 - 100 TB environment size just by moving from tape to a disk-based virtual tape library.

Combining the data protection capabilities of Tivoli Storage Manager and Virtual Tape Library Solutions with the ProtecTIER reduces the time and cost associated with managing backup and recovery processes, and eliminates cumbersome tape media management and maintenance activities like tape head cleaning and tape re-tensioning.

There is no doubt that the need to maintain data on physical tape media will continue for the foreseeable future. However, supplementing Tivoli Storage Manager with the ProtecTIER technology can enable IT managers to consolidate tape library resources, lowering hardware replacement costs, implementation and training costs, remote vaulting services costs, and significant annual drive and library maintenance costs.

Improving reliability

If you are looking for the best way to increase the reliability of your Tivoli Storage Manager environment, begin by reducing the cause of most failed backup jobs, which is tape. Tapes, tape drives, and the robotic tape libraries are mechanical devices and are prone to failures, including these:

- Tape drive failure
- Media errors
- Tape getting stuck in a drive
- Cannot read tape due to aging or normal wear and tear
- Tape not loaded
- Wrong tapes loaded
- Tape demagnetized
- Robotic arm jam

The reliability of tape is a big issue and frequently causes backup jobs to fail, requiring manual intervention and time spent restarting jobs to maintain the necessary level of data protection.

Use scenarios

The LAN-free data movement in itself is fairly simple in its usage. However, when data arrives at the destination VTL, options can be used for data reduction and disaster recovery purposes.

Data reduction by deduplication

Deduplication aims to eliminate redundant data. Currently it is mainly used in backups and archiving. Although it can be employed for any data, this is not recommended because of performance considerations.

Deduplication performance is the result of two processes: identification of duplicate data, which requires a database or index look-up, and a store operation in the repository. Because ProtecTIER uses inline processing, the performance is increased. The reason is that is that the deduplication appliance keeps the index in its memory during processing. Tivoli Storage Manager server use an after-processing deduplication method unless client-side deduplication is used. In a very large database configuration, using client-side deduplication will impact the server performance. If you need to lower the backup footprint on large amount of data ProtecTIER is the obvious choice.

Replication of data for disaster recovery

The ProtecTIER Replication Manager is a server that remotely manages the replication grids within an organization. The ProtecTIER Replication Manager can be installed on a ProtecTIER node, and can manage up to one grid with 24 repositories.

In most Tivoli Storage Manager environments, backup data is written to tape and physically transported to offsite storage or a hot site location for disaster recovery protection. This manually intensive operation is required because electronically transporting the huge amount of backup data generated daily is cost prohibitive.

With the ProtecTIER advanced HyperFactor data de-duplication capabilities, moving data to and from remote locations can be done quickly and affordably. By significantly shrinking the amount of data sent, the ProtecTIER technology reduces the bandwidth that is required to electronically vault data to remote locations. This enables IT organizations to replicate selected mission-critical backup data from a primary location over a WAN to a secure offsite location.

You can use ProtecTIER replication to offload tape cloning to your secondary site. Many users are replicating their data from the primary site to the secondary (DR) site, and then moving it from the disk-based repository onto physical tape cartridges for long-term retention. One of the advantages of this practice at the secondary site is that it shifts the burden of cloning to physical tape from the production environment to the DR site location. The DR site cloning operation uses the cartridge replicas at the ProtecTIER VTL shelf of the destination. The process imitates the commonly used physical process for the transportation of physical cartridges from the primary site to a DR site. This feature is effective in single domain backup deployments because in these environments the backup application servers at both sites share the catalog and can be concurrently connected to the ProtecTIER systems. The replication visibility switch control feature is used in these environments. The cartridges to be cloned are moved from the primary repository to the secondary repository and then cloned to physical tapes.

A typical site-to-site disaster recovery model is shown in Figure 7-37 on page 245. A backup server can either be Tivoli Storage Manager server or Storage Agent.

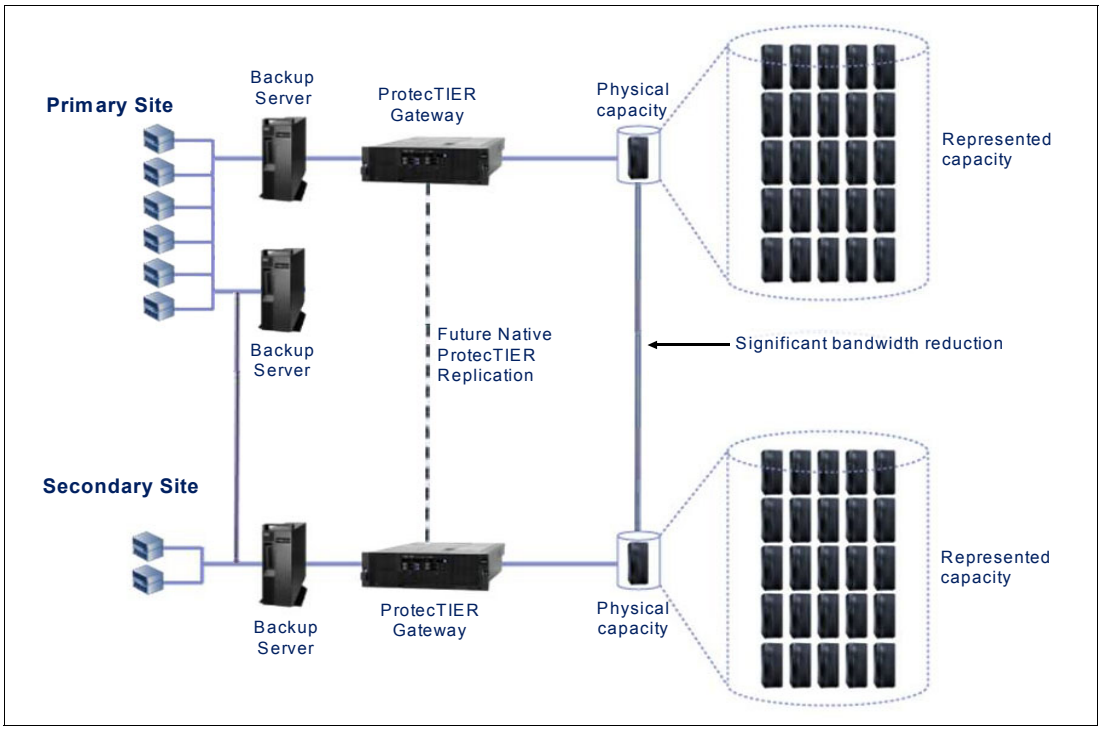

Figure 7-37 Site-to-site disaster recovery model

Cloning of Data to Physical Tape for Archiving

Regular daily backups can be directed to the VTL to exploit disk technology advantages, but when the data on a virtual tape needs to be archived, the backup application or the VTL control program copies the data from the virtual tape to a physical tape as a cloning operation.

7.6.3 Protect very large databases with Tivoli Storage FlashCopy Manager

This section describes how to protect very large databases with Tivoli Storage FlashCopy Manager integrated with Tivoli Storage Manager. The solution is based on an example but the problem description applies to most other environments of this type, and the same concepts can be used for that particular database system you have. We point out what to consider and which functionalities are available to protect these types of environments, both on the client side and on the Tivoli Storage Manager server side.

Solution benefits

Several obvious benefits to this solution should be considered. Some of them that are listed result from the exploitation of LAN-free and ProtecTIER, based on the integration with Tivoli Storage Manager:

- RTO is improved.
- Backup window is shorter.
- Off-loaded backup to Tivoli Storage Manager server so that the production server is not involved in the backup and therefore zero load on production server.
- ProtecTIER provides deduplication to backup data to lower storage costs.
- LAN-free to VTL removes the backup off the LAN to the SAN.
- LAN-free takes backup loads away from the Tivoli Storage Manager server CPU and I/O.

- Cloning capabilities of production database for testing purposes.
- Quiesce time of database is short.

The backup and recovery time is almost instantaneous with FlashCopy and FlashBack. The restore time can vary depending on how many transactions need to be rolled forward after the FlashBack has occurred.

Solution architecture

Big database environments are based on an enterprise server and storage infrastructure. The solution in Figure 7-38 covers the protection of an environment based on the AIX platform with SAP running on a DB2 database system. The concepts also apply to an Oracle environment, so if you are using Oracle by itself or with SAP you have the same options as described here.

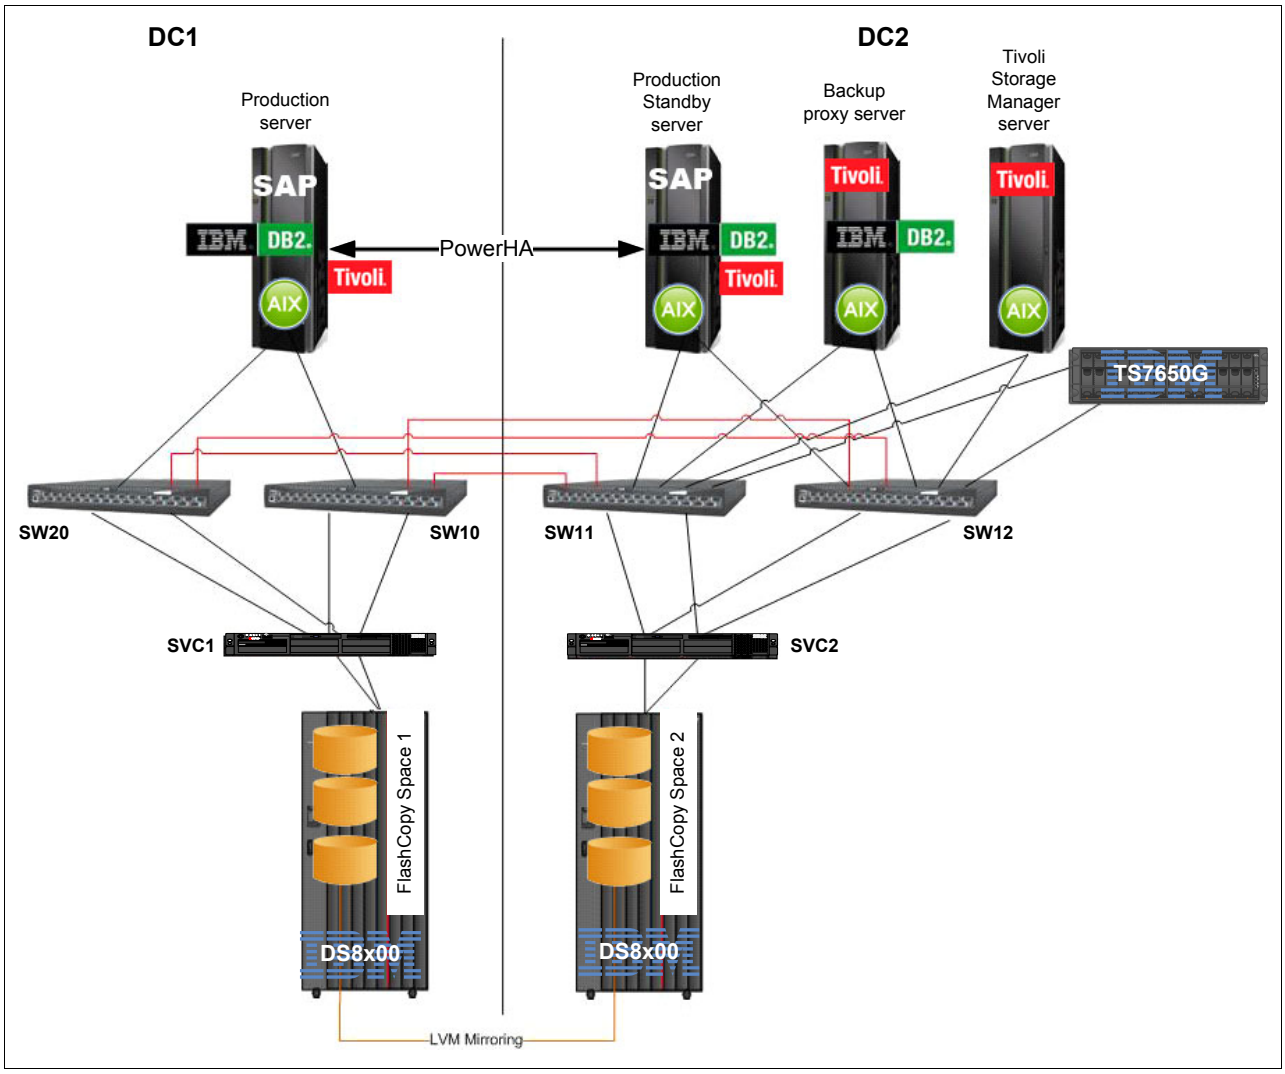

Figure 7-38 Architectural concept in a big data concept of structured data.

In this type of environment, you normally keep the production data in two data centers with some sort of replication or mirroring between them. In this case, we use AIX LVM mirroring to have a synchronized copy of the SAP database data on both sites. Another option is to exploit storage replication (Metro Mirror or Global Mirror) to replicate the database volumes to the remote site. The two servers are set up in a PowerHA where the standby server can take over

the production automatically and run the database in DC2 if something happens to DC1. It could even be for hardware maintenance of the production server in DC1. We do not distinguish between whether the database is running in an AIX logical partition (LPAR) or on a physical machine.

As a back-end solution for Tivoli Storage Manager server storage, we chose IBM ProtecTIER gateway. This is the most scalable VTL solution with hardware deduplication capabilities that IBM has today. Read 7.6.2, “LAN-free backups to VTL” on page 238 for more information about LAN-free to VTL solutions.

Considerations before implementing the solution

Consider the following information before implementing big database environments:

- Spread the FlashCopy source and target volumes on all arrays. As a result, the DS8x00 will give full performance to production when no backup is running.
- TS7650G emulates IBM Ultrium TD3. The model DD4 can perform with 7.2 TB per hour sustained inline deduplication backup. Performance tests show that the number of available Fibre Channel (FC) adapters on the Tivoli Storage Manager backup proxy server limits the actual bandwidth, not the ProtecTIER. For instance, if the backup window is 6 hours for 20 TB, you must use 4 VTL drives = ~3.4 TB/hr, ~880 GB/hr/stream. If we can manage to feed more FC adapters on the Tivoli Storage Manager backup proxy server, we can decrease the backup and restore time even further.
- Review the number of ports available in the storage system to perform the backup. It should at least fit the bandwidth required to move the data to the TS7650G within the given time frame.
- Calculate the total I/O capacity in the SAN Volume Controller and storage system to make sure it is able to deliver the bandwidth required to do the backup and the restore.
- Use both ProtecTIER deduplicated and non-deduplicated as back-end storage pools.

  Database logs are known to have a high change rate (100%). As database logs track all changes within the database, they are never identical. Consider multiple storage pool types in the back-end Tivoli Storage Manager storage tier using both ProtecTIER with deduplication enabled and other disk without deduplication enabled to store data.

  Using ProtecTIER with deduplication is described in *IBM ProtecTIER Implementation and Best Practices Guide*, SG24-8025:

  `http://www.redbooks.ibm.com/redpieces/abstracts/sg248025.html`
- Make sure you are guaranteed to have the mount points available during backup otherwise the whole backup operation might fail.
- Scheduling of backups to the ProtecTIER must be planned so that the load is spread equally. Because we use off-loaded backup, you are not as limited on the backup start windows as you normally would be.
- Use replication of ProtecTIER file system to other ProtecTIER instead of Tivoli Storage Manager copypool for DR

  ProtecTIER allows the configuration of replication policies to replicate file system’s directories and all objects contained in these directories recursively to remote ProtecTIER repositories. This is done without any disruption to the operation of the file system as a target for backup.
- Use migration delay (**`migdelay`**) on non-deduplicated storage pool to make sure that transaction log files are available for restore from disk and not from tape. Preferably keep this data on disk until it expires.

- Tivoli Storage Manager passwords must be the same at all times. Use ASNODE to access data and authenticate from each host.
- Establish a test environment to test the restore procedure at least four times a year, or at least when any changes are made in the environment (hardware or software).

Solution description

We installed Tivoli Storage FlashCopy Manager on the targeted SAP source servers to provide the FlashCopy backup to disk. We use the technology of the advanced storage subsystems and the SAP integration provided by Tivoli Storage FlashCopy Manager. When the FlashCopy has been completed the disk is then mounted to a backup proxy server to either verify the data or to copy the data to ProtecTIER so it can be managed by the Tivoli Storage Manager server policies.

All copy services functions used by IBM Tivoli Storage FlashCopy Manager are at the volume or LUN level. In addition, multiple volumes that are organized into volume groups require IBM Tivoli Storage FlashCopy Manager to process these volume groups consistently. As a result, non-application data residing on a volume group that is processed by IBM Tivoli Storage FlashCopy Manager is included in the backup. Similarly, all data that resides on a volume group that is being restored is overwritten. This means that the disk configuration must be planned to be sure that the FlashCopy will include the correct data.

Because SAP environments are fully integrated with DB2, the DB2 backup command is used. DB2 notifies Tivoli Storage FlashCopy Manager of the current environment in order to enable Tivoli Storage FlashCopy Manager to implement the appropriate workflow. Tivoli Storage FlashCopy Manager supports single partition databases, and logically or physically partitioned databases on journaled file systems (JFS and JFS2). Supported DB2 backup options are documented in the DB2 user publications. Tivoli Storage FlashCopy Manager supports these DB2 backup functions:

- Full database backups, both online and offline
- Backups of selected database partitions
- Backups of database partitions including or excluding database logs

Consider the following guidelines:

- In a multi-partition database environment for SAP workloads, all partitions are suspended in parallel (parallel mode).
- The **`db2 backup`** command is available in the DB2 Control Center. For SAP environments, this command is also available in the Computing Center Management System (CCMS).

Snapshots are created using space efficient FlashCopy. This minimizes the space requirement on the DS8x00 to store the FlashCopy snapshots. However we are then also restricted to have access to both the source volume and the FlashCopy volume to make a restore. If we do a full copy of the source volume to the target FlashCopy volume we do not need the source volume for restore. But it would require the same space in both the source and target volumes and we need to copy the data before it will be available for restore. The copy frequency would then be limited to how long time the copy will take before making another FlashCopy.

The FlashCopy frequency is another aspect to consider. One suggestion is to take a FlashCopy four times a day: three in DC1 and one in DC2. The daily backup timeline in Figure 7-39 on page 249 shows when the scheduled operations are running and when DB2 sends the database log files to Tivoli Storage Manager.

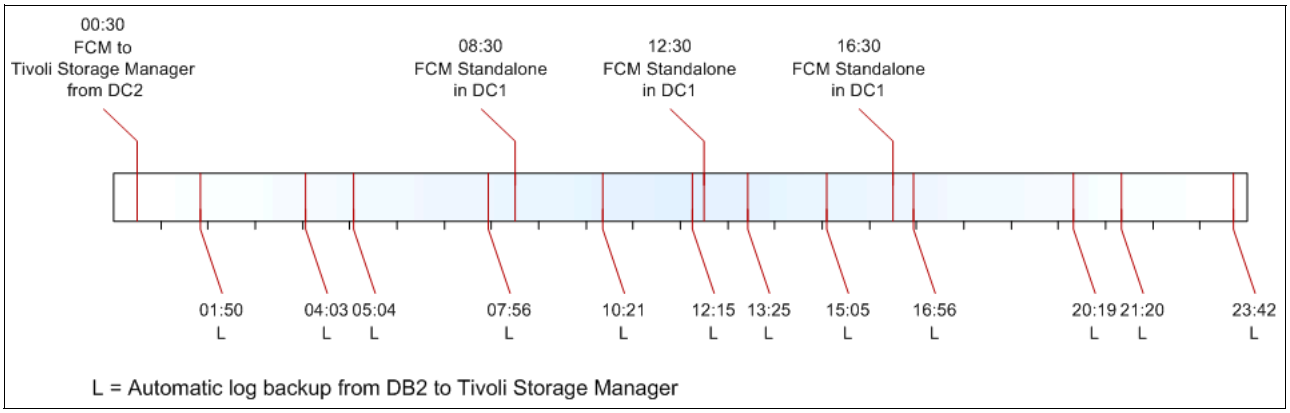

Figure 7-39 Time line of data protection schedules.

The schedules can run at 00:30 in DC2 for DR; and at 08:30, 12:30, and 16:30 a schedule can run in DC1. If the working hours of the company is 08:30 to 17:30, then at a maximum, you roll forward 4 hours of logs from the production day. We suggest that you combine the FlashCopy snapshots with the Tivoli Data Protection for ERP to Tivoli Storage Manager server at DC2, which is closest to the Tivoli Storage Manager server storage hierarchy.

Table 7-8 shows source and target volumes and which location is used at the scheduled time.

Table 7-8 FlashCopy schedules and location

Scheduled time	Source volume location	Target volume location	Disk only or to Tivoli Storage Manager
00:30	DC2	DC2	Copy to Tivoli Storage Manager after FlashCopy
08:30	DC1	DC1	FlashCopy to disk
12:30	DC1	DC1	FlashCopy to disk
16:30	DC1	DC1	FlashCopy to disk

Scenarios

Several components are involved in the backup and restore process of data when using Tivoli Storage FlashCopy Manager:

- Application and database layer (SAP DB2)

  DB2 initiates the process. The database is put into backup mode to make a consistent FlashCopy.

- Backup and restore initiator

  Tivoli Storage FlashCopy Manager is responsible for triggering the FlashCopy process and for managing the retention policies of the FlashCopy snapshots.

- Virtualization layer

  The CIMOM agent resides on the SAN Volume Controller clustered system and is the interface used to communicate with the SAN Volume Controller and trigger the FlashCopy command.

Figure 7-40 shows the components and how they communicate.

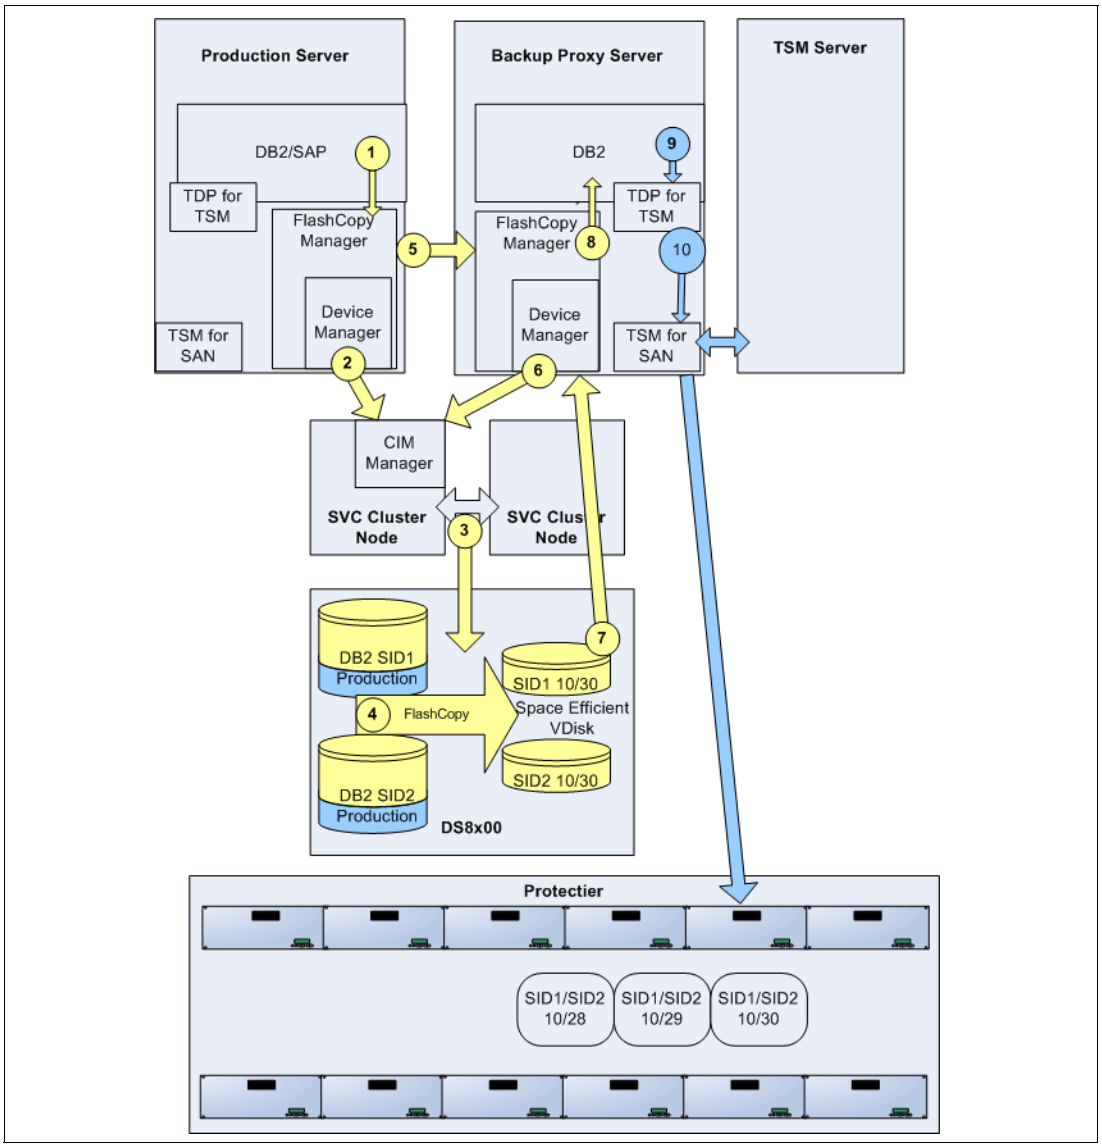

Figure 7-40 Backup flow

Backup using Tivoli Storage FlashCopy Manager stand-alone solution

The following steps show the backup process to make a FlashCopy in the storage system as a stand-alone solution. The process is shown in Figure 7-40 and is indicated in yellow.

1. Backup is initiated as scheduled using DB2, SAP CCMS, or the Tivoli Storage Manager scheduler. DB2 communicates to Tivoli Storage FlashCopy Manager to start the process by issuing the following backup command:

   ```
   backup database…use snapshot
   ```

2. The Device Manager, which is part of the Tivoli Storage FlashCopy Manager component, on the production server notifies the CIMOM interface on the SAN Volume Controller to initiate a FlashCopy snapshot of the volumes.
3. SAN Volume Controller sets up needed pointers for FlashCopy to occur.

4. FlashCopy recovery point starts. Any changes since FlashCopy recovery point are tracked in the space efficient volume (pre-change copies).
5. Tivoli Storage FlashCopy Manager on the production server notifies Tivoli Storage FlashCopy Manager on the backup proxy server that the FlashCopy is complete.
6. The Device Manager on the backup proxy server notifies the CIMOM interface on the SAN Volume Controller to mount the space efficient VDisks to the backup proxy server.
7. The SAN Volume Controller present the space efficient VDisks to the backup proxy server as if they were full VDisk. Unchanged tracks, on the source production server are shown and tracks preserved before change which are kept on the space efficient VDisk.

The process stops here if you do not want to make a backup to the Tivoli Storage Manager server.

Backup using Tivoli Storage FlashCopy Manager and Tivoli Storage Manager server

If you want to make a copy of the database to Tivoli Storage Manager you can proceed with the following steps. The process is shown in Figure 7-40 on page 250 indicated in blue.

1. Tivoli Storage FlashCopy Manager on the backup proxy server notifies DB2 that a backup can occur.
2. DB2 initiates a backup using standard Tivoli Storage Manager methods. Tivoli Storage Manager thinks the backup is occurring on the production server.
3. The DB2 backup uses the Tivoli Storage Manager for SAN agent to run a LAN-free backup to VTL.

Tivoli Storage FlashCopy restore from Tivoli Storage Manager

The FlashCopy restore will also be LAN-free and at this point, the backup proxy server is not involved. The Tivoli Storage Manager server works directly with the production server and the Tivoli Storage Manager for SAN agent to mount the needed tape-to-tape drives attached to the production server. This is basically a standard Tivoli Storage Manager restore at this point. See Figure 7-41 on page 252.

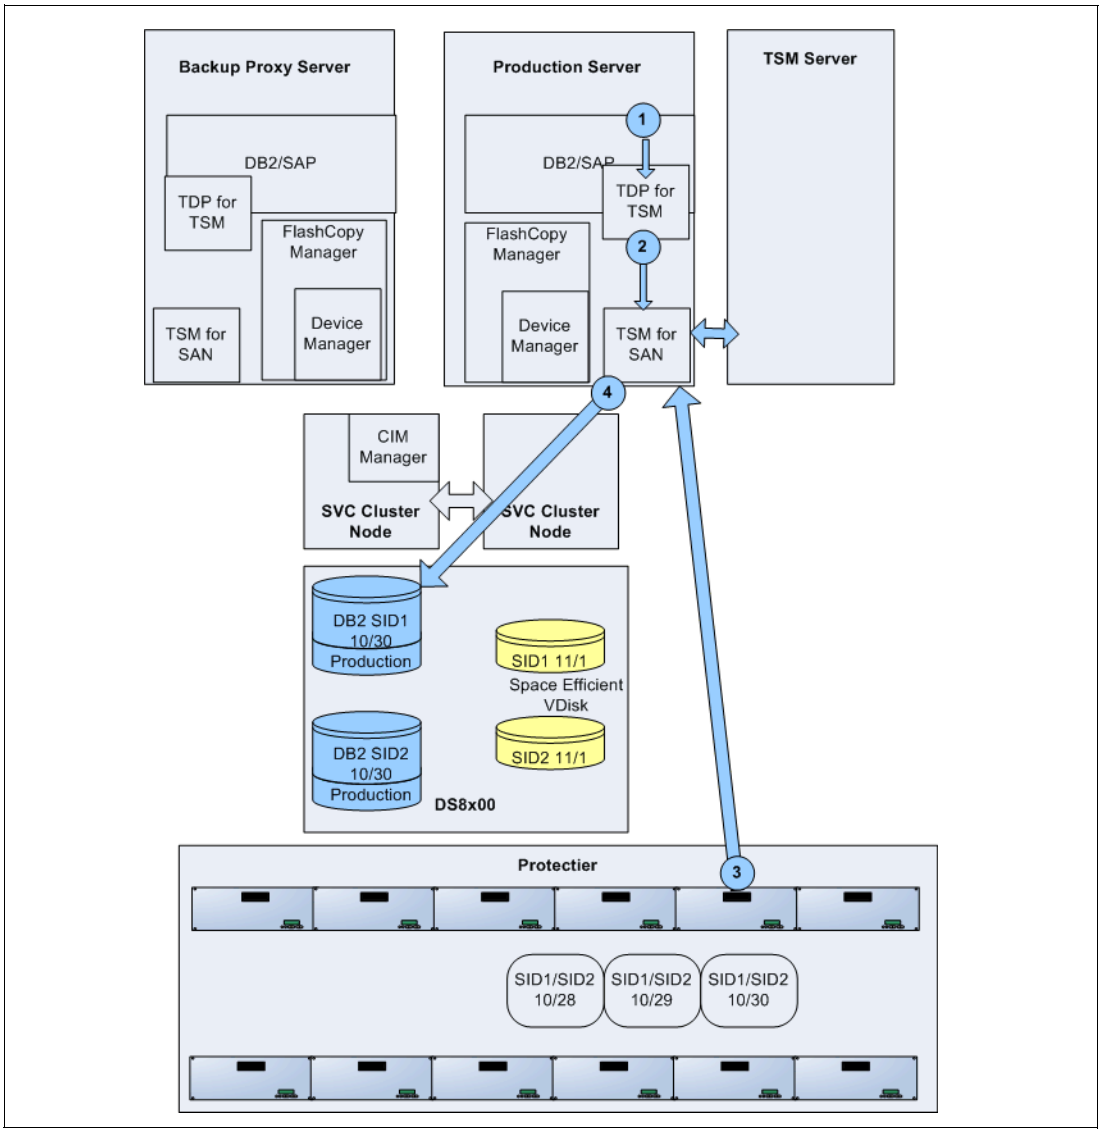

Figure 7-41 Recovery process with Tivoli Data Protection

Steps for the recovery process as shown in Figure 7-41 are as follows:

1. DB2 issues standard recovery commands of **`db2 restore`** or **`db2 recover`** and interfaces with Tivoli Data Protection for Tivoli Storage Manager to invoke Tivoli Storage Manager for the restore.
2. Tivoli Data Protection for Tivoli Storage Manager interfaces with Tivoli Storage Manager for SAN to perform a LAN-free restore from the tape drive connected to the production server.
3. Tape data is transferred to the production server.
4. Production DB2 is recovered according to the recovery request.

Tivoli Storage FlashBack restore

The FlashBack restore does not involve Tivoli Storage Manager. The production server communicates directly with Tivoli Storage FlashCopy Manager and the SAN Volume Controller to recover the database using the copies of tracks at the time of the FlashCopy recovery point to overlay the changed tracks in production.

Note: This cannot be used to recover lost LUNs on the production server database because space efficient FlashCopy.

Figure 7-42 shows the communication flow to recover data from a FlashCopy.

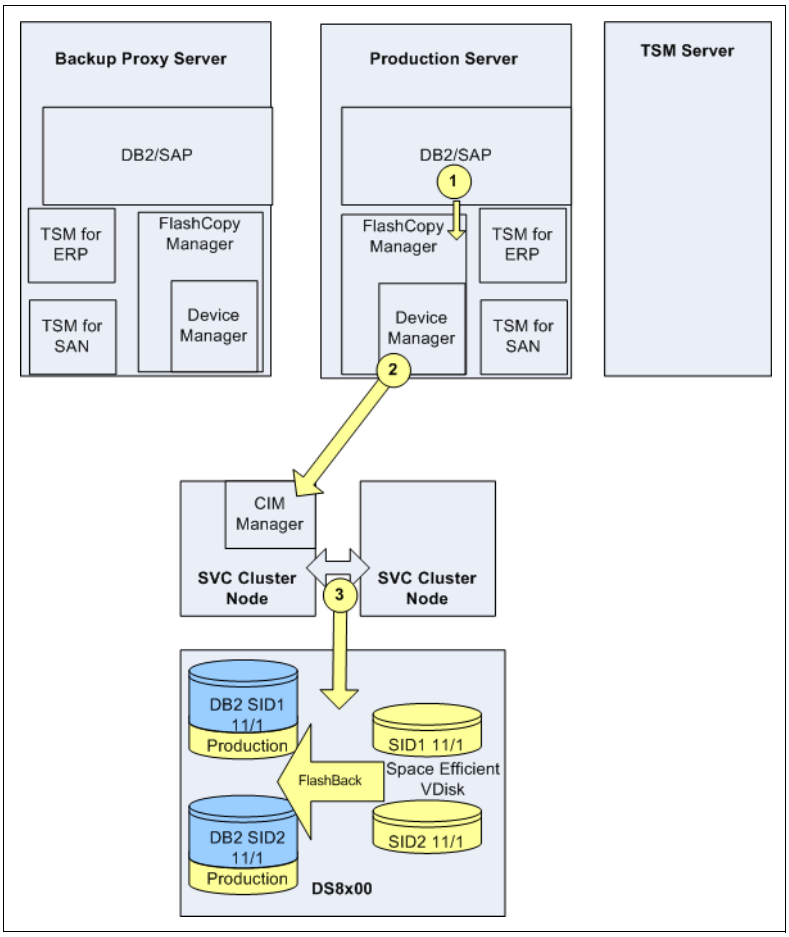

Figure 7-42 FlashBack flow

Steps for the recovery process as shown in Figure 7-42 are as follows:

1. DB2 issues special recovery commands to invoke Tivoli Storage FlashCopy Manager.
2. The device manager on the production server communicates with the CIMOM on the SAN Volume Controller to request the FlashBack operation
3. The FlashBack is performed so that the tracks unchanged before the FlashCopy recovery point are used to overwrite the changed tracks in the production database.

When you perform a restore in an LVM mirrored environment, the mirror is re-created automatically. But the resynchronization of the volume groups is not performed. This operation consumes numerous CPU and I/O resources. This leads to a degraded operation, for example, in case of log recovery of the database instance was restored. Thus the re-creation and the resynchronization of the LVM mirror must be done separately when you decide it can run.

Tivoli Storage FlashBack and DB2 recovery of log files to last transaction

According to the default, the `include logs` option is used for FlashCopy backups. The transaction logs must reside on at least one volume group, distinct from the table space containers.

Forward recovery can start immediately after a FlashCopy relationship is established to apply the transaction logs including those from the Tivoli Data Protection based backup and rolling the database forward to a point in time.

The database at the primary server is now fully accessible, and all applications on the primary server can start.

Using Tivoli Storage FlashCopy Manager to clone databases

Cloning is a different kind of operation from replication and backups. The cloned environment is fully functional and separate. Possible reasons to perform system copies are as follows:

- To create test and quality assurance systems that are re-created regularly from the production systems to test new developments with the most actual production data
- To create migration or upgrade systems from a production system prior to phasing in a new release or functions into production
- To create education systems from a master training system to reset before starting a new course
- To create dedicated reporting system to offload a workload from production

The SAP System Copy guidelines describe several extra actions to be performed in the copied SAP system, as in the following examples:

- Disable Remote Function Call (RFC) destinations.
- Disable batch-job processing.

Some of these actions defuse the cloned system that is they anticipate the running of SAP tasks that are planned in the production system, but that must not be performed or even repeated in the cloned system (for example, data transfer to or from other applications, batch jobs or spool jobs). Tivoli Storage FlashCopy Manager provides the ability to automatically run scripts before and after clone creation and before the cloned SAP system is started.

Tivoli Storage FlashCopy Manager uses the FlashCopy function of the SAN Volume Controller for database cloning. This method eliminates downtime and minimizes the impact on the production database.

For FlashCopy backup, the physical volume identification numbers (PVIDs) are not changed. For FlashCopy cloning, the PVIDs of the FlashCopy target disks are automatically changed by the Tivoli Storage FlashCopy Manager software. You can have several cloned databases of one source database running on one host.

With Tivoli Storage FlashCopy Manager, a cloning process can be started with an online or offline source database. For online Tivoli Storage FlashCopy Manager cloning, the source database is suspended for a short time. The suspension occurs when the storage system creates its FlashCopy or snapshot of the source database.

The cloned database (target database) can have the same database name as the source database. The cloned database can also be renamed to any valid database name during the Tivoli Storage FlashCopy Manager cloning process. Tivoli Storage FlashCopy Manager requires the cloned database to be created on a different database server than the source database server, regardless of whether the clone database name is changed.

The basic cloning steps are shown in Figure 7-43.

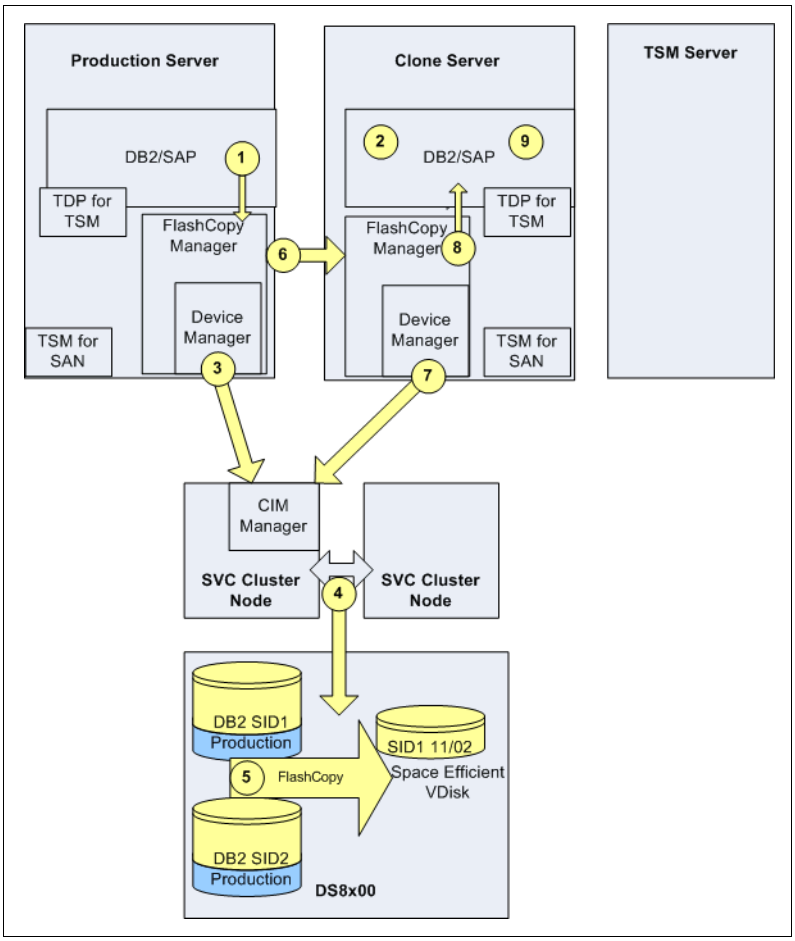

Figure 7-43 Cloning flow

The cloning function is started from the command line on the production system. As shown in Figure 7-43, the steps are as follows:

1. Start cloning on the production system.
2. The preprocessing scripts run against the clone database. This task is optional and depends on available preprocessing scripts on the clone. The scripts are not part of the Tivoli Storage FlashCopy Manager software.
3. The Device Manager on the production server notifies the CIMOM interface on the SAN Volume Controller to initiate a FlashCopy of the volumes.
4. SAN Volume Controller sets up needed pointers for FlashCopy to occur.
5. A consistent FlashCopy or snapshot backup, including database logs, is created on the storage server.
6. Tivoli Storage FlashCopy Manager on the production server notifies Tivoli Storage FlashCopy Manager on the clone server that the FlashCopy is complete.
7. Mount the cloned production database on the clone system and rename the PVIDs, logical volumes and volume groups.

8. The database on the clone system is recovered. The database is renamed to match the name (SID) of the clone database.
9. The postprocessing scripts run against the clone database. This task is optional and depends on the available postprocessing scripts on the clone. The scripts are not part of the Tivoli Storage FlashCopy Manager software.

When you clone databases and use SAN Volume Controller, the space-efficient disks can be used as a target for FlashCopy cloning operations, but there are restrictions on the FlashCopy backups. You can complete cloning operations from the cloning source volumes. If you want to complete FlashCopy backup and FlashCopy cloning from the same source disks, use full target disks.

Tivoli Storage FlashCopy Manager can also create a database clone on a remote storage system. See “Metro Mirror for backup or cloning to remote site” (the next section).

Note: Cloning a FlashCopy Backup image is not supported. If multiple clones are required, these are always created from the same running production or source system, not from an offline backup image.

Metro Mirror for backup or cloning to remote site

As an alternative to using AIX LVM mirroring you can use Metro Mirror and Global Mirror with SAN Volume Controller. Today this functionality is available on AIX, HP-UX, Linux, and Solaris with IBM System Storage SAN Volume Controller and IBM XIV Storage Systems.

Figure 7-44 illustrates the hosts and volumes involved in remote mirroring that uses Metro and Global mirrors and where FlashCopy Manager is installed to make FlashCopy snapshots on both sites.

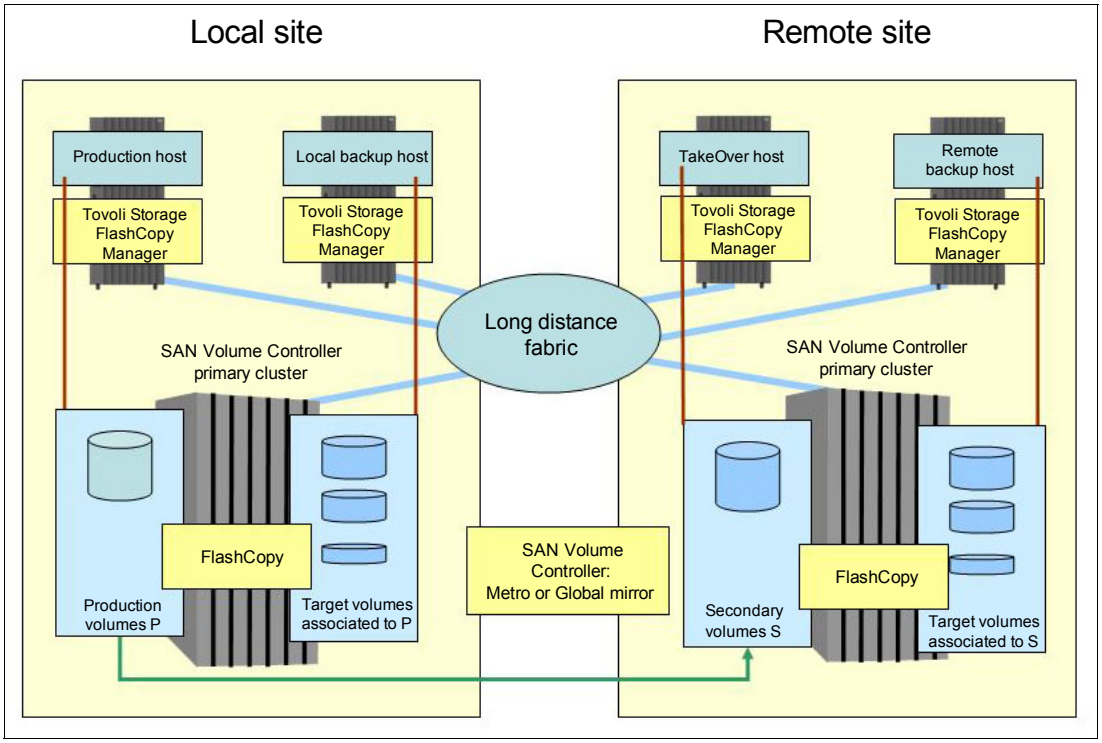

Figure 7-44 Remote mirroring using Metro Mirror and Global Mirror sources

When implementing the solution, Tivoli Storage FlashCopy Manager can create FlashCopy backups on the secondary storage system, not on the primary system that is currently used by the application. This provides more flexibility for backup scenarios and a fail-safe database copy in case the primary location is lost.

Figure 7-45 shows the data movement using Metro Mirror or Global Mirror and the data transfer to the Tivoli Storage Manager server storage hierarchy.

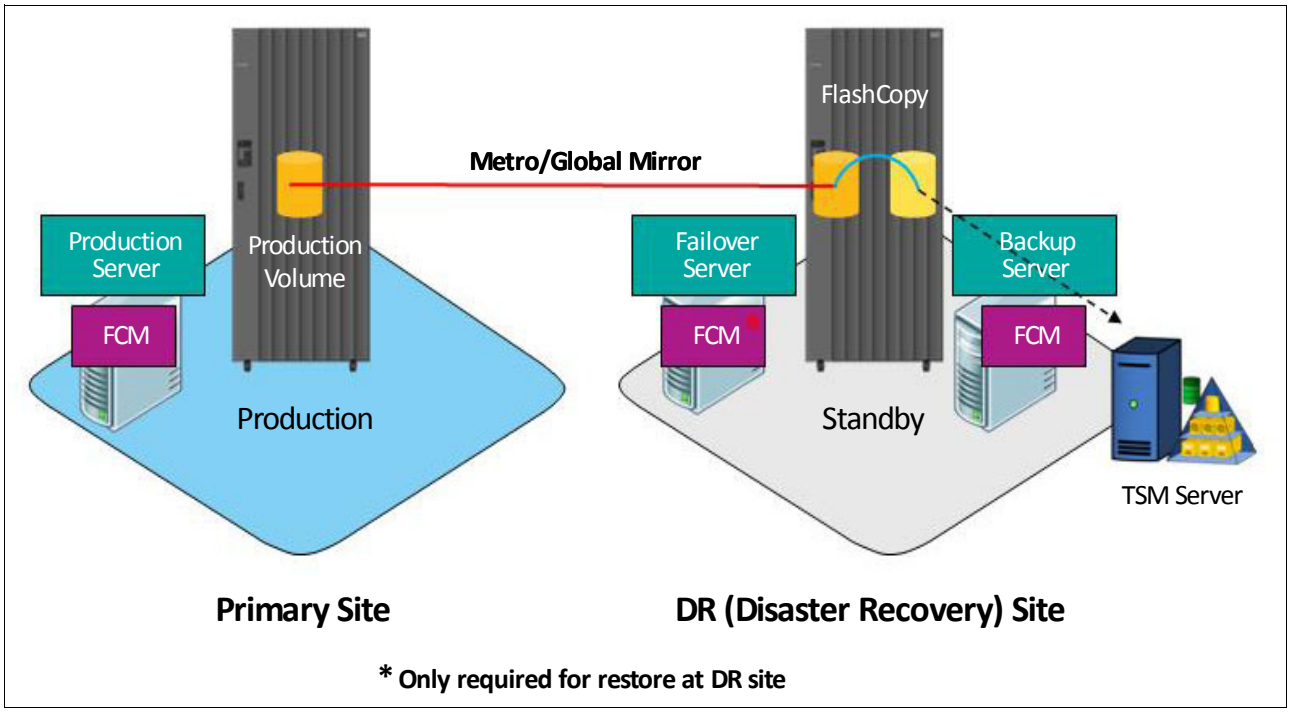

Figure 7-45 Metro Mirror and Global Mirror architecture

The benefit to keep the FlashCopy snapshots on a DR site is evident:

- Offload the workload from production storage control unit
- Consolidate multiple backups to the DR site
- The backups on the DR site can be used for restore in a disaster situation and to offload the backup to the Tivoli Storage Manager server
- When fail over to DR site occurs, the backup continues to work

A use scenario might be to keep FlashCopy snapshots on the primary site and the secondary site. Different configurations of the target volumes provide more flexibility in how the storage capacity is being used:

- FlashCopy target volumes that are smaller than the source volumes that have been created for the database to enable space efficient FlashCopy snapshots on the primary system.
- FlashCopy target volumes that have the same size as the source volumes to enable full and incremental FlashCopy snapshots on the secondary SAN Volume Controller cluster. These copies can also be used for cloning the database.

If you want to create FlashCopy backups or clones on both the primary and the secondary storage system, the Tivoli Storage FlashCopy Manager configuration must provide the information of when to perform which kind of backup or clone when running the schedule.

If you consider the recovery scenarios, also consider why you would implement a Metro Mirror or Global Mirror solution. It will primarily be to increase the availability of the application and be able to run it on other hardware in a remote site.

The notion is that restoring from the FlashCopy snapshots will only occur if the replicated data to remote site include the error that you want to overcome. Or it could be that you do not want to make a fail over to the remote site and need to recover the data on the primary site.

So when would you make use of the FlashCopy snapshots in the disaster site? Definitely for backup to Tivoli Storage Manager server or cloning processes. The restore of the remote FlashCopy snapshots will be used if:

- The primary source LUN is destroyed or corrupted and the replication of the data is corrupted as well.

  Since the data cannot be recovered from the primary target FlashCopy snapshots because we used Space-efficient FlashCopy volumes we need to recover it from the disaster site
- The primary site is down and a fail over to the disaster site did not work and you need to recover the data on the disaster site.

Note: For environments with a SAN Volume Controller version 6.1 and earlier, Tivoli Storage FlashCopy Manager must stop and deactivate the Global Mirror and Metro Mirror before running a restore operation.

Additional information

For more information, see the following resources:

- Tivoli Storage FlashCopy Manager; system requirements and supported environments:

  http://www.ibm.com/software/tivoli/products/storage-flashcopy-mgr/
- IBM System Storage SAN Volume Controller:

  http://www.ibm.com/systems/storage/software/virtualization/svc/index.html
- *Implementing the IBM System Storage SAN Volume Controller*, SG24-7933 (description and FlashCopy):

  http://publib-b.boulder.ibm.com/abstracts/sg247933.html
- *Best Practices for Tivoli Storage FlashCopy Manager in a SAN Volume Controller Metro Mirror environment*:

  http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102196
- *Infrastructure Solutions: Design, Manage, and Optimize a 20 TB SAP NetWeaver Business Intelligence Data Warehouse*, SG24-7289:

  http://www.redbooks.ibm.com/abstracts/SG247289.html
- *IBM ProtecTIER Implementation and Best Practices Guide*, SG24-8025:

  http://www.redbooks.ibm.com/redpieces/abstracts/sg248025.html
- Tivoli Storage FlashCopy Manager reconciliation:

  http://www-01.ibm.com/support/docview.wss?uid=swg21618086

7.7 Big data: Unstructured

In this section, we look at big data from the perspective of the growing use of unstructured data and the resulting continuous increase of the number of files in a file system.

The big challenge is to find the best way to protect this environment, which sees millions of new files being written every day to single file systems, making sure backups complete successfully in a timely manner and that data is restorable also in a timely manner. New technologies should be considered when you want to back up files systems that have millions of files.

In this section, we explain what components you can use with Tivoli Storage Manager to back up unstructured data, both on local file systems and in remote NAS devices, based on the matrix shown in Figure 7-46. Starting at the most traditional backup method of progressive incremental backups, we describe other alternatives for the backup of large file systems.

Challenges										
business					addressed with	technical				
reduce recovery time (RTO)	reduce backup time (meet RPO)	Compliance	Business Continuity	security (secure data transfer,	Solution (discussed in this book)	Virtual Environments	Significant data growth	automation for business	Remote office and mobile	network bandwidth
X	X				Disk to Disk backup	X		X	X	
X	X				Data reduction with TSM	X	X		X	X
X	X	X	X	X	TSM ToolKit	X	X	X	X	X
	X				progressive incremental forever	X			X	X
	X		X		off-load backup	X	X			
X			X		instant restore	X	X			
	X				journal based backup		X			X
X	X				image backup	X	X			
	X				GPFS		X		X	
X	X				NDMP		X			
	X				Snap-Diff		X			X
X	X		X		Snap-Mirror-to-Tape / SMTAPE		X			
X					point-in-time restore	X				

Figure 7-46 Challenges with big data unstructured

For several reasons, you might have use different procedures to do a successful and reliable backup, always with the goal to be prepared for a fast and effective restore. The various functions are described in the next topics. Figure 7-47 on page 260 shows the procedures and functions that are available and can be combined (analogy is a gear shifter knob). Notice that progressive incremental and GPFS policy-driven backups are grouped together. The Tivoli Storage Manager backup-archive client is integrated with the GPFS command **mmbackup** to provide a scalable and performant backup solution for GPFS file systems. Therefore the progressive incremental backup process as provided from Tivoli Storage Manager backup-archive client can be replaced by **mmbackup**, provided from GPFS.

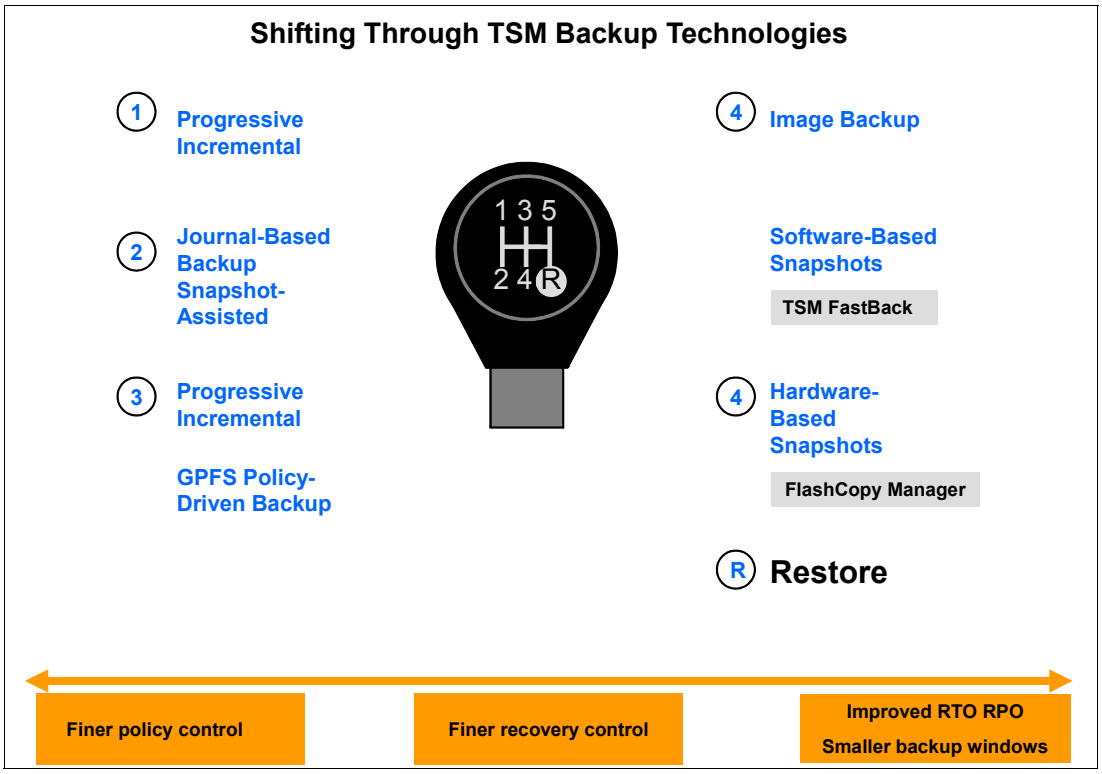

Figure 7-47 Like a gear shift, use the methodologies for backup

7.7.1 Progressive incremental backups

One of the key differentiators between Tivoli Storage Manager and other data protection products is the progressive incremental backup methodology. Tivoli Storage Manager backs up only new or changed files. It tracks all of the backups at a file level. It has no concept of a full backup with dependent incrementals or differentials. Because of Tivoli Storage Manager's powerful relational database, it does not require periodic full backups. This methodology reduces network and storage resource consumption and lowers the overall cost of storage management. Tivoli Storage Manager file-level progressive backup methodology is far superior to other traditional backup methods such as full-plus-incremental or full-plus-differential, because progressive incremental backups are never redundant.

Challenges and benefits the solution addresses

Backing up large file servers has always been a challenge for backup administrators because of the increasing numbers of files that need protection on either local file systems on servers of any operating system, or using Network File System (NFS) for UNIX and Linux operating systems and Common Internet File System (CIFS) for Windows.

Progressive incremental backup is Tivoli Storage Manager's most traditional component. It offers the advantage of not needing to run weekly full backups, which reduce the backup window and the amount of backup storage, either on tape or disk. Only files that have changed are backed up.

Progressive incremental backup relies on a local file system scan and a query of active entries in the Tivoli Storage Manager server database.

Benefits of the solution

A summary of the benefits of progressive incremental backups is as follows:

- Single-instance store of data

  Saves time and storage space by backing up only new files and modified files. The progressive backup feature uses the Tivoli Storage Manager relational database to track data where 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.

- File deletion support

  Files deleted from local machine will result in file expiration on Tivoli Storage Manager server. Active file deletion on Tivoli Storage Manager server will result in a new backup of local file during the next backup.

- Policy support

  If management class differs from the current backup and the file has not changed, the file will be rebound to a new management class (new includes rule or policy change). It honors copy mode (that means the action to take if file changes during backup) and copy frequency (number of days that pass between subsequent backups of a file).

Solution architecture

Preferred storage pool destination is disk or file, because the small files that are sent directly to tape can cause tape drive contention. This is also true for the journal-based backup solution in 7.7.2, “Journal-based backups” on page 264.

Running an incremental backup, the backup-archive client first inspects the file on the local systems and compares it with files that are already backed up on the server and have not been changed on the client. If it finds a file on the client that has been changed, it backs up that file.

Figure 7-48 shows how progressive incremental backup works.

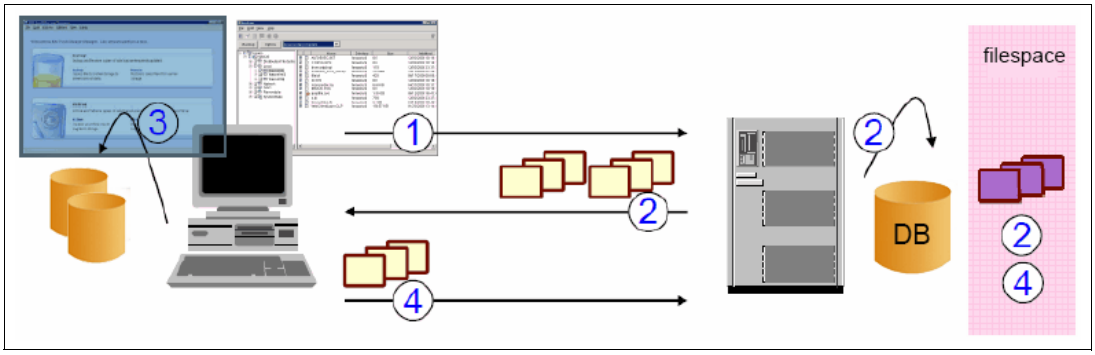

Figure 7-48 Progressive incremental backups with large file servers

The progressive incremental backup shown in Figure 7-48 has these steps:

1. Client queries the server for a current view of file system (**`dsmc incremental`**).
2. Server returns a list of files for the entire file system (**`db query`**).
3. Client scans and compares the local file system.
4. Client creates transactions of files for backup.

During the inspection phase (while the client is checking if files are changed), the session remains idle. Because the object inspected number is high and the session is idle during that time, your connection is closed by the Tivoli Storage Manager server based on the value of

IDLETIMEOUT. In this case, you must increase or adjust the value of this parameter to best fit your needs.

For very large file systems, the time spent to inspect the file system might add so much time to the overall time of the backup that the available backup window is exceeded. You might need to implement other options, as discussed in this book. The time required to enumerate the files is done at relatively the same speed that the **dir** command would list all the files on the computer. Another reason is that if the number of files is very large, more memory will be required. Tivoli Storage Manager requires approximately 300 bytes of memory per file to inspect. When the number of files is very large, more memory will be used and can affect performance.

Progressive incremental optimizations

These are features that, when used together with progressive incremental backups, can further help the backups of file systems with large number of files:

- Memory efficient backup

  Use the memory efficient backups with the incremental command when your client is memory constrained. You can also use this option as a parameter to the `include.fs` option. Use this option only if memory on the client is a bottleneck, otherwise its implementation might affect backup performance. With the `MEMORYEFFICIENTBACKUP` option, Tivoli Storage Manager backs up one directory at a time.

- Memory efficient disk caching

  Use the `memoryefficientbackup=diskcache` method for any file space that has too many files for Tivoli Storage Manager to complete the incremental backup with either the default setting, `memoryefficientbackup=no`, or with `memoryefficientbackup=yes`. In this case, the use of RAM will be decreased, but the use of disk for caching will be increased. The actual amount of disk space required for the disk cache file created by disk cache incremental backups depends on the number of files and directories included in the backup and on the average path length of the files and directories. It can require up to 5 GB of disk space for each million files or directories being backed up.

- Virtual mount points

  If you have a large file system you want to force **resourceutilization** logic to treat differently, **virtualmountpoint** will tell Tivoli Storage Manager to treat a directory like a discrete mount point. Using collocation by file space can help ensure that restores are multi-session also. If you are an authorized user and you want to back up files beginning with a specific directory within a file system, you can define that directory as a virtual mount point.

  Defining a virtual mount point within a file system provides a direct path to the files you want to back up, saving processing time. It is more efficient than defining the file system with the domain option and then using an exclude option to exclude the files you do not want to back up. It also allows you to store backups and archives for specific directories in separate storage file spaces.

- Incremental by date

  This checks fewer attributes of the files to be backed up, and much fewer lookups on the Tivoli Storage Manager server (just last file space backup date). It takes less time to process than a full incremental backup and requires less memory.

Note: When using incremental by date, periodic standard incremental should be also run because incremental-by-date backups do not expire deleted files or rebind backup versions to a new management class if you change the management class. And incremental-by-date backup also does not back up files with attributes that change, unless the modification dates and times also change.

Use scenarios

You can use progressive incremental to backup and recover CIFS/NFS file system data on NAS devices with progressive incremental as shown in Figure 7-49.

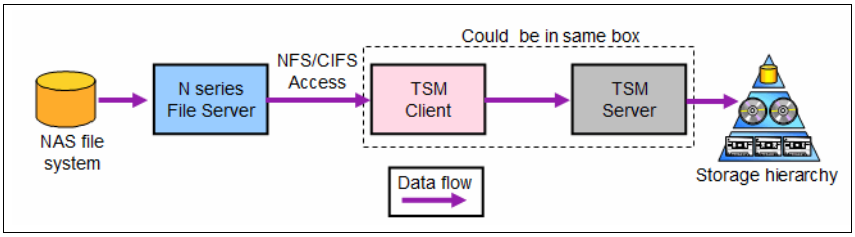

Figure 7-49 Progressive incremental and NFS/CIFS

There are two ways to set up progressive incrementals with CIFS/NFS file systems:

- Use a Tivoli Storage Manager backup-archive client to back up and restore data, by using CIFS or NFS to access files from the backup-archive client. Data can be stored on the Tivoli Storage Manager server with file-level granularity, by using the progressive incremental backup method. The data is stored in the Tivoli Storage Manager storage hierarchy and can be migrated, reclaimed, and backed up to a copy storage pool.

  This method increases processor usage when the Tivoli Storage Manager client accesses individual files. Data flows through the Tivoli Storage Manager client and through the Tivoli Storage Manager server, unless a LAN-free configuration is used.

- Use a Tivoli Storage Manager backup-archive client running on the NAS device, if you can use external programs with the NAS operating system. Data can be stored on the Tivoli Storage Manager server with file-level granularity using progressive-incremental backup. The data is stored in the Tivoli Storage Manager storage hierarchy and can be migrated, reclaimed, and backed up to a copy storage pool.

  This method decreases processor usage of CIFS or NFS. Data flows through the Tivoli Storage Manager client and through the Tivoli Storage Manager server, unless a LAN-free configuration is used.

The Tivoli Storage Manager backup-archive client can be configured to protect files which are accessed with the Network File System (NFS) protocol.

Backup performance is better when you install the backup-archive client where the file system physically resides, but sometimes it is necessary to access file systems using NFS for purposes of backup and recovery. The Tivoli Storage Manager UNIX and Linux backup-archive client can back up, archive, restore and retrieve file data using an NFS mount. This includes all versions of the NFS protocol.

The Windows client backs up the CIFS share definition under the root directory, the mapped CIFS share, or the UNC name. This support requires that the NAS file server is running DATA ONTAP software which presents CIFS shares to remote clients as ordinary remote NTFS shares.

The root directory of a CIFS share is backed up with a full progressive incremental backup of the mapped drive or UNC name. See Example 7-2.

Example 7-2 Backing up CIFS share

```
net use x: \\NetAppFiler\CifsShareName
dsmc incr x:
or
dsmc incr \\NetAppFiler\CifsShareName
```

The output in Example 7-3 is displayed when the root directory (and share definition) is backed up.

Example 7-3 Output from dsmc incr command

```
Directory-->                    0 \\NetAppFiler\CifsShare\ [Sent]
```

7.7.2 Journal-based backups

Journal-based backup is an alternate method of backup that uses a change journal maintained by the Tivoli Storage Manager journal service process and improves incremental backup performance in most environments.

Challenges and benefits the solution addresses

When a file system has too many objects, the scanning of changed files alone can take too much time, increasing the backup window. With journal-based backups, because this initial scanning is not performed, the time to determine which files have changed is greatly reduced.

Network traffic between the client and the server is also reduced because less pre-backup communication happens. The backup-archive client will still send data as usual to the server to be stored and at this time the network traffic is the same without journaling.

Because the backup-archive does not carry out the initial metadata conversation, it does not sit idle, waiting for the list of files to be backed up. The backup-archive client begins to send the files to the Tivoli Storage Manager as soon as the journal-based backup is initiated. This means faster backup times and less backup-archive client idle time.

Journaling is best used for these situations:

- Small number of files (< 1,000,000) and small number of changes between backups
- Large number of objects (< 10,000,000 but > 1,000,000) with 10 - 15% change rate

Solution description

With journal-based backups, the file system and Tivoli Storage Manager database scans are not performed to determine which files to backup. Instead, a local database is created with all changes that occur in the file system before backups. Figure 7-50 on page 265 explains how journal-based backups work.

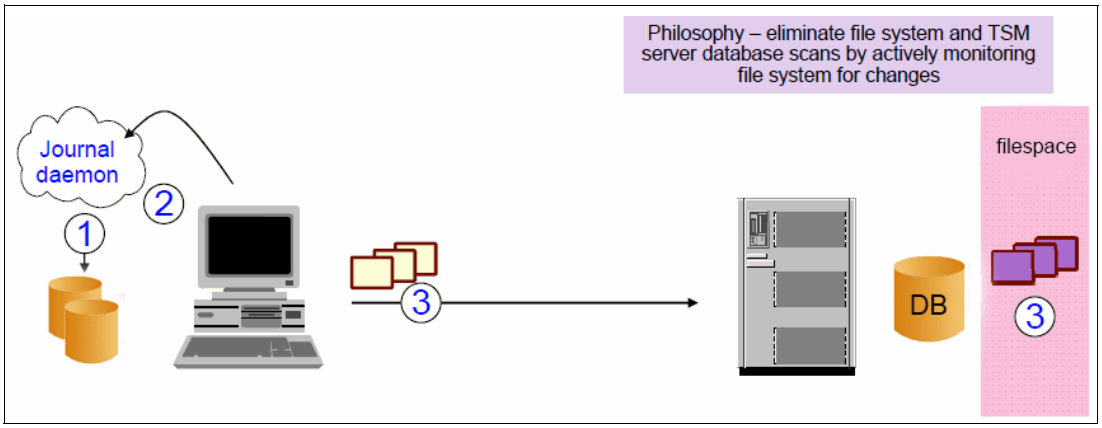

Figure 7-50 Journal-based backup architecture

The figure explains the following steps:

1. Journal daemon monitors the file system for changes.
2. When backup is executed, the client queries the journal for file system changes.
3. Client translates the journal changes into an internal list of files and creates transactions of files for backup.

A *change journal* is a database of file system change information. The Tivoli Storage Manager Journal daemon creates and maintains a change journal for each file system that is being journaled, as specified in the Tivoli Storage Manager Journal Service `tsmjbbd.ini` configuration file.

A backup of a particular file system will be journal-based when the Tivoli Storage Manager journal daemon has been installed and configured to journal the particular file system, and a valid journal has been established for the file system. Journal-based backup is enabled by successfully completing a full incremental backup.

The primary difference between traditional incremental backup and journal-based backup is the method used for backup and expiration candidates.

A journal database will be created in the client server and sometimes its settings need to be refined for better performance or troubleshooting. See the following technote to estimate the size of a journal database:

http://www.ibm.com/support/docview.wss?uid=swg21213920

Use scenarios

You will use journal-based backup when backing up file systems with small or moderate amounts of change activity between backup cycles. If you have many file changes between backup cycles, you will have very large change journals. Large change journals might create memory and performance problems that can negate the benefits of journal-based backup. For example, creating, deleting, renaming, or moving very large directory trees can also negate the benefit of using journal-based backup instead of normal incremental backup.

Journal-based backup is not intended to be a complete replacement for traditional incremental backup. You should supplement journal-based backup with a full progressive incremental backup on a regular basis.

Initial successful progressive incremental backup of entire file system must be performed to enable the journal.

Progressive incremental backup will be executed in these situations:

- The file space was deleted on the server since the last backup.
- The node policy set was updated since the last backup.
- Any error occurs in the journal database.
- The velocity of changing files is high enough to cause a notification buffer overrun (Windows).
- The journal is taken offline or the journal service is stopped and restarted (and the PreserveDbOnExit setting is not specified).

7.7.3 Image backups

An image backup is a block-by-block copy, single object backup of a volume (typically a UNIX file system or raw logical volume, or Windows drive) on a Tivoli Storage Manager client. Being able to restore an entire volume as one object can lead to faster recoveries. Image backup is available at the time of writing on AIX, HP-UX, Oracle Solaris, Linux, and Windows client platforms. Review the specific image backup requirements for each platform in their associated installation guides.

Challenges and benefits that the solution addresses

Ask yourself this question: What is more important, having a quick restore of the data as a whole or being able to restore individual files? If the probability of needing a file level restore is low, image backup is one solution.

Advantages of image backup are as follows:

- Faster backups and restores because they are done at the logical volume level, especially for volumes with a large number of files
- Online image backup using snapshot (Windows, Linux LVM2, AIX JFS2).
- Volume is left online.
- Windows and AIX: Only used blocks will make the image; non-used blocks are not transferred.
- Provides a point-in-time picture of the logical volume.
- No scan time (local file system or backup server) to determine what has changed.
- Faster overall data movement.
- Fewer resources on Tivoli Storage Manager server because only one database entry will be written for the image.
- Ability to take image plus incremental backups.

Solution description

You can back up a logical volume as a single object (image backup) on your system. The traditional static image backup prevents write access to the volume by other system applications during the operation.

You must be a root/administrative user to perform this task on AIX, HP-UX, Linux, Solaris and Windows operating systems, and image backup does not apply to Mac OS X. You can perform the backup only on formatted volumes; not on raw logical volumes. Windows uses VSS to create snapshot of volume.

To restore an image backup of a volume, the Tivoli Storage Manager client must be able to obtain an exclusive lock on the volume being restored.

Use scenarios

We describe two backup methods to be used with image backups that allow you to perform a point-in-time restore of your file systems and improve backup and restore performance. You can always use image backups alone but then you do not have the possibility of restoring single files and it will generate larger backup volumes.

Method 1: Using image backups with file system incremental backups

Follow these steps to perform image backups with file system incremental backup:

1. Perform a full incremental backup of the file system. This establishes a baseline for future incremental backups.
2. Perform an image backup of the same file system to make image restores possible.
3. Perform incremental backups of the file system periodically to ensure that the server records additions and deletions accurately.
4. Perform an image backup periodically to ensure faster restore.

Restore your data by performing an incremental restore. Before you start the restore, select the image plus incremental directories and files, and delete inactive files from local options in the Restore Options window. During the restore, the client does the following steps:

1. Restores the most recent image on the server.
2. Deletes all of the files restored in the previous step that are inactive on the server. These are files that existed at the time of the image backup, but were subsequently deleted and recorded by a later incremental backup.
3. Restores new and changed files from the incremental backups.

Method 2: Using image backups with incremental-by-date image backups

Follow these steps to perform image backups with incremental by date image backup:

1. Perform an image backup of the file system.
2. Perform an incremental-by-date image backup of the file system. This sends only those files that were added or changed since the last image backup to the server.
3. Periodically, perform full image backups.

Restore your volume by performing an incremental restore. Before you start the restore, select the Image plus incremental directories and files option in the Restore Options window. This first restores the most recent image and then restores all of the incremental backups performed since that date.

Perform full image backups periodically, which improves restore time. The file system can have no previous full incremental backups.

Incremental-by-date image backup does not inactivate files on the server. When you restore an image with the incremental option, files deleted after the original image backup are present after the restore.

Comparing method 1 and method 2

To help you decide which method is appropriate for your environment, see the comparison in Table 7-9.

Table 7-9 Comparison of method 1 and method 2

Method 1: Using image backup with file system incremental	**Method 2: Using image backups with incremental-by-date image backups**
Files are expired on the server when they are deleted from the file system. On restore, you have an option to delete files that are expired on the server from the image.	Files are not expired on the server. After the image incremental restore completes, all files that are deleted on the file system after the image backup are present after the restore. If file systems are running at or near capacity, an out-of-space condition might result.
Incremental backup time is the same as regular incremental backups.	Incremental image backup is faster because the client does not query the server for each file that is copied.
Restore is much faster compared to a full incremental file system restore.	Restore is much faster compared to a full incremental file system restore.
Directories deleted from the file system after the last image backup are not expired.	Directories and files deleted from the file system after the last full image backup are not expired.

7.7.4 Traditional file system backups and RTO

The solutions presented in this section can be combined together in order to meet specific service level agreements (SLAs). Table 7-10 lists the use of each of solution in relation to the recovery time objective (RTO).

Note: These are only suggested numbers, which can be used as a baseline.

Table 7-10 RTO table showing greater than (>) and less than (<) values

Client data			
Recovery time objective (RTO)	**< 100 GB < 1,000,000 files**	**100 - 500 GB < 10,000,000 files**	**> 500 GB > 10,000,000 files**
Low: > 24 hours	Progressive backup	Progressive backup + journal	Progressive backup + journal + image backup
Medium: > 4 hours but < 24 hours	Progressive backup + journal	Progressive backup + journal + image backup	Progressive backup + journal + local snapshot / non Tivoli Storage Manager
High: < 4 hours	Progressive backup + journal + image backup	Progressive backup + journal + local snapshot / non Tivoli Storage Manager	Progressive backup + journal + local snapshot / non Tivoli Storage Manager

7.7.5 NDMP backups

Network-attached storage (NAS) file servers are dedicated storage machines whose operating systems are optimized for file-serving functions. NAS file servers typically do not run third-party software. Instead, they interact with programs like Tivoli Storage Manager through industry-standard network protocols, such as network data management protocol (NDMP).Tivoli Storage Manager uses NDMP to perform high-performance, scalable backups and restores.

Challenges and benefits the solution addresses

Data residing on N series and NAS devices have unique backup needs. Tivoli Storage Manager provides enhancements to address the big filer challenge. Companies that are using very large file systems often realize where it is not possible to continue with standard incremental backups, even with the enhancements we presented.

Benefits of NDMP backups are as follows:

- Minimizes network traffic and outboard transfer data outboard of the Tivoli Storage Manager client and server
  - Backup/restore direct to library
  - No LAN data traffic
- Does high performance, scalable backups and restores with full and differential images.
- File-level restores use table of contents (TOC) and direct access recovery (DAR).
- Supports dynamic drive sharing.
- Allows centralization of tape resource.
- Exploits full capability of Tivoli Storage Manager storage hierarchy.
- Works with many NAS devices:
  - IBM N series, Network Appliance
  - EMC Celerra
  - NAS devices certified for NDMP with Tivoli Storage Manager. For a list of NAS file servers that are certified through the Ready for IBM Tivoli software, visit the following site and enter search word `NDMP`.

    http://www.ibm.com/software/brandcatalog/opal

Solution architecture

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 7-51 on page 270.

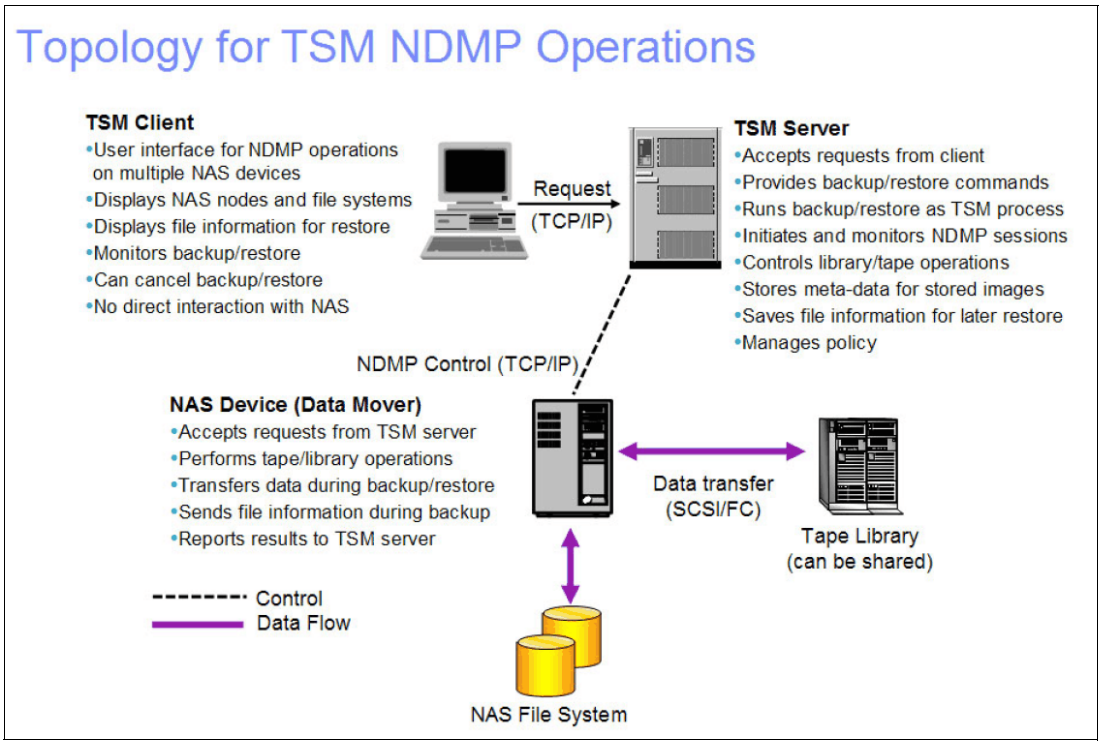

Figure 7-51 Topology for NDMP using IBM Tivoli Storage Manager

During backup and restore operations, data flows directly between the NAS appliance and the tape drive. 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 through the NAS file server.

Solution description

Tivoli Storage Manager combines NDMP methodology with existing Tivoli Storage Manager infrastructure and capabilities allowing the following items:

- Administrative authority and remote access
- Policy management
- Library and drive sharing
- Tape management
- System and operation management (scheduling, reporting, logging)

NDMP backup methods are either full or differential. Full backup includes all files within a NAS file system or directory tree. Differential backup includes all files that have changed since the most recent full backup. If differential backup is specified and full backup not found, a full backup is performed.

A restore of differential image is automatically preceded by a restore of corresponding full backup. NDMP restore behavior of full plus differentials is dependant on the NAS device.

Deactivation, versioning, or expiration is done at the file system level.

Tivoli Storage Manager tracks file-system image backups and can perform NDMP file-level restores using Direct Access Recovery (DAR).

During backup, the NAS device sends metadata, including position information, for each file backed up. Tivoli Storage Manager stores position information and other metadata in the TOC for this backup.

During a restore of individual files (as opposed to directory tree or file system), Tivoli Storage Manager accesses TOC to get position information for each file to be restored. Tivoli Storage Manager initiates a DAR operation, providing position information for each file. NAS device positions directly to each file, avoiding scan of the entire image.

If the full path of the file is known, use the server's **RESTORE NODE** command. The difference is that the restore will not be a DAR restore, just a scan.

Figure 7-52 shows how backup and restore work with the table of contents.

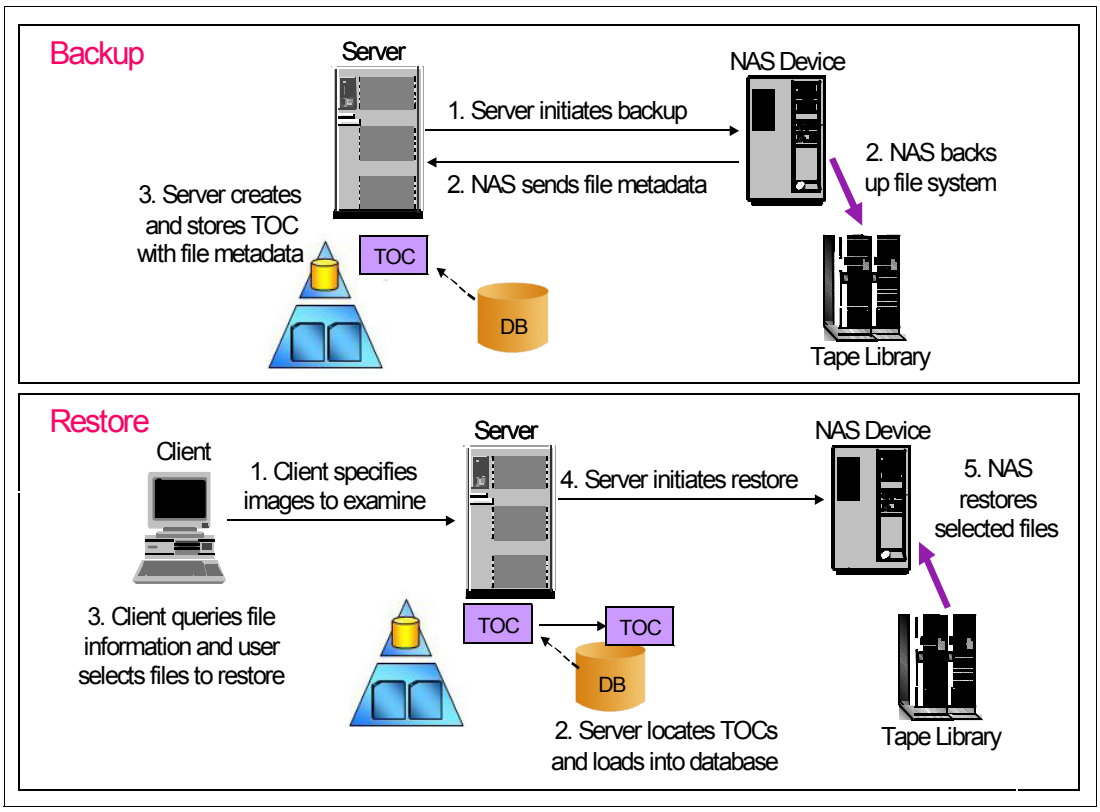

Figure 7-52 File-level restore using TOC

Storage location of TOC is specified in Tivoli Storage Manager policy and we suggest that it stays on disk storage for fast access.

When performing a file-level restore for a NAS file server, the table of contents (TOC) is loaded into temporary database tables in order to choose files to restore. The amount of space required depends on the average length of the file and directory names and the average depth of the directory structure. The amount of space also depends on the file server vendor and whether non-English characters are used in file and directory names.

In general, the amount of space required is about 280 bytes for each file or directory in the TOC when using English file and directory names. If the NAS file server is Network Appliance and contains non-English file or directory names, the amount of space required is about 340 bytes for each non-English file or directory in the TOC.

Each independent restore operation will require storage. For example, a file-level restore of a NAS file server that includes 10 million directories and files with English names will require about 2.8 GB of temporary space in the database. Restoring Network Appliance file servers, whose directories and files contain non-English characters, requires additional space for the TOC. After the TOC has been unreferenced for a specified period of time, the temporary database space is released. Use the **SET TOCLOADRETENTION** command to specify the approximate number of minutes that unreferenced TOC data will remain loaded in the server database.

Policy management with NDMP

Full and differential NAS backups must use the same management class, otherwise re-binding will happen during the backup, causing backups to take different retention values depending on the management class. The following link provides more information about expiration with NDMP.

http://www.ibm.com/support/docview.wss?uid=swg21200154

Tivoli Storage Manager versioning applies to the complete NDMP dump. Tivoli Storage Manager server is not aware of the single objects included in the NDMP dump (except when reading the TOC).

Copygroups determine the destination of the TOC.

Offsite protection

The following operations are supported with NDMP:

- Back up storage pool.
- Restore storage pool or volume.
- Move data:
  - Intra-pool for space recovery
  - Inter-pool for migration to new device type
- Disaster Recovery Manager has support for NDMP data.
- Restore node uses copy pool if primary data is not accessible.

Use scenarios

Tivoli Storage Manager provides two of the most common configurations that use NDMP for backing up and managing NAS file servers.

Configuration 1: Library connected to the Tivoli Storage Manager server

In this configuration, we have a library whose drives are attached through Fibre Channel or SCSI to the NAS server, and whose robotics are attached to the Tivoli Storage Manager server. Note that for configuration 1 to be used, the library must have separate ports for library robotics control and drive access. It further requires that both the Tivoli Storage Manager server and NAS server are within either Fibre Channel or SCSI range of the tape library. Figure 7-53 on page 273 shows the details of this solution.

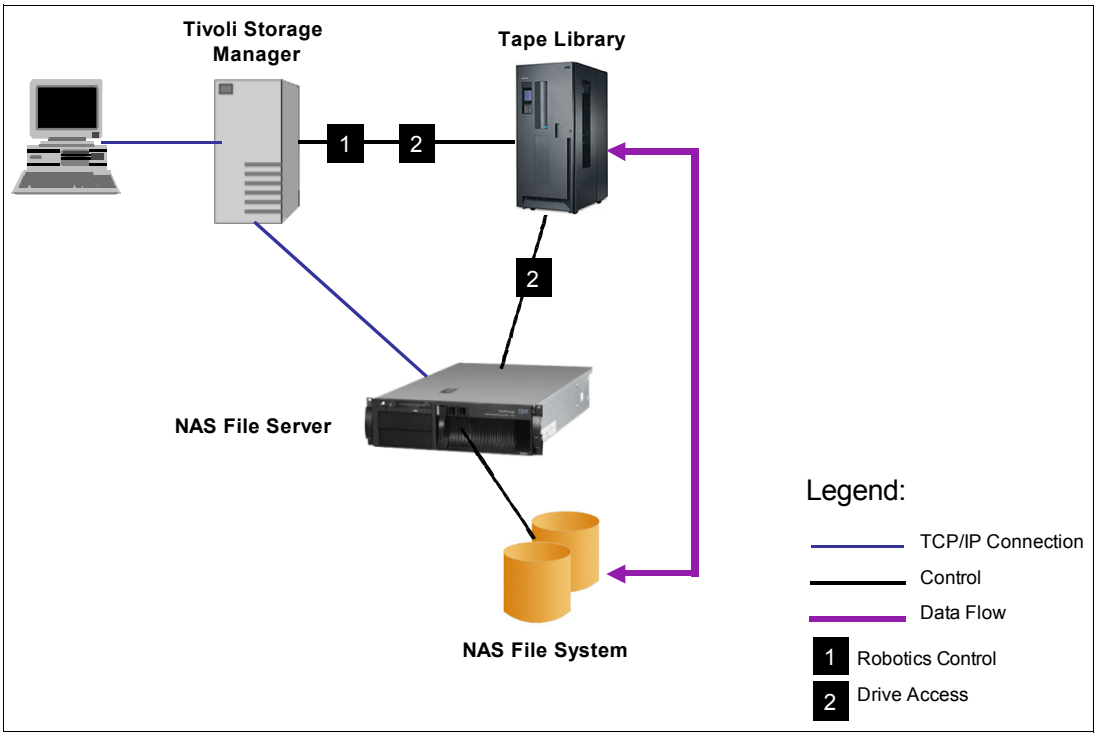

Figure 7-53 Library device attached to Tivoli Storage Manager server

NDMP filer-to-server

Another variation is to use Tivoli Storage Manager hierarchy to send backups from NAS devices using Network Data Management Protocol (NDMP).

With the filer-to-server configuration, the library is attached to the DMA (Tivoli Storage Manager server). The NAS device does not have access to the library, as shown in Figure 7-54 on page 274. This configuration is also known as *3-way NDMP backup*. The backup data from the NAS device is transferred over the network (TCP/IP) to the Tivoli Storage Manager server.

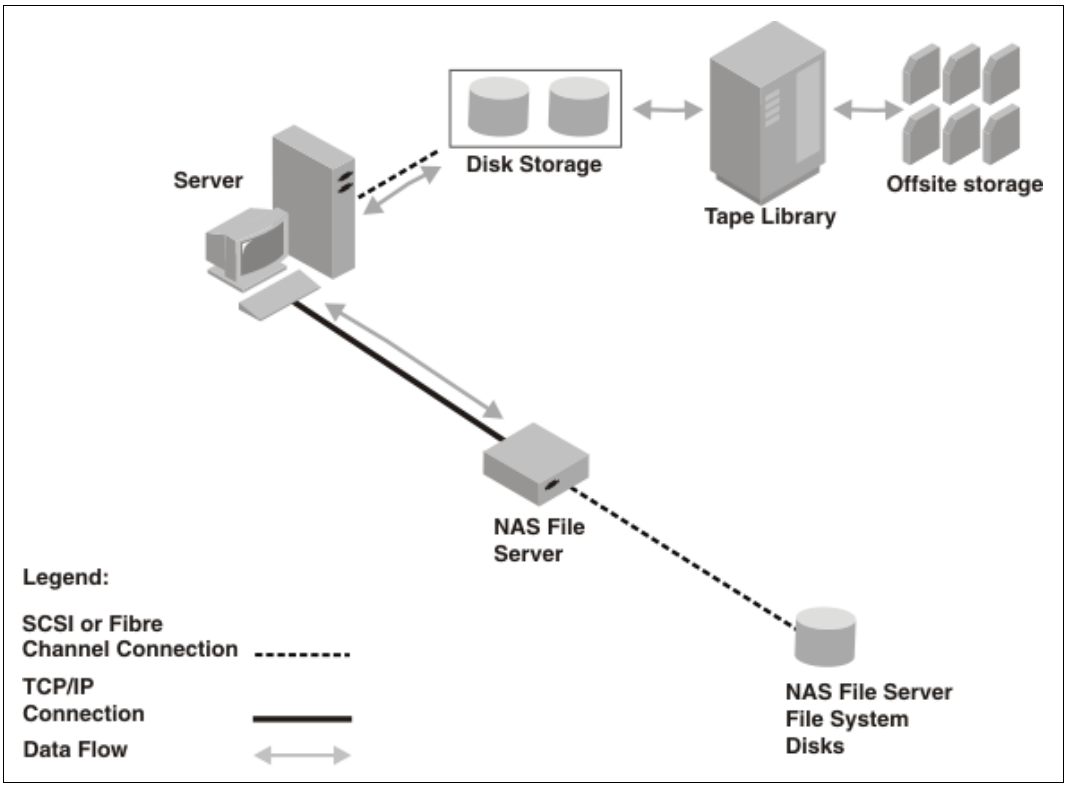

Figure 7-54 Filer-to-server configuration

For this configuration, several server options are established.

The `NDMPCONTROLPORT` option specifies the port number to be used for internal communications for certain NDMP operations. The Tivoli Storage Manager server does not function as a general purpose NDMP tape server.

Specify the port number to be used for internal communications for certain NDMP operations. The port number must be in the range of 1024 - 32767. The default is 10000.

Firewall considerations are more stringent than they are for filer-to-attached-library because communications can be initiated by either the Tivoli Storage Manager server or the NAS file server. NDMP tape servers run as threads within the Tivoli Storage Manager server and the tape server accepts connections on port of 10001.

This port number can be changed through the options file (Example 7-4) in the Tivoli Storage Manager server options file shown in “Hardware and software requirements” on page 192.

Example 7-4 NDMPPORTRANGE options file

```
NDMPPORTRANGE port-number-low,port-number-high
```

During NDMP filer-to-server backup operations, you can use the `NDMPPREFDATAINTERFACE` option to specify which network interface the Tivoli Storage Manager server uses to receive NDMP backup data. The value for this option is a host name or IPv4 address that is associated with one of the active network interfaces of the system on which the Tivoli Storage Manager server is running. This interface must be IPv4-enabled.

Before using this option, verify that your NAS device supports NDMP operations that use a different network interface for NDMP control and NDMP data connections. NDMP control

connections are used by Tivoli Storage Manager to authenticate with an NDMP server and monitor an NDMP operation; NDMP data connections are used to transmit and receive backup data during NDMP operations. You must still configure your NAS device to route NDMP backup and restore data to the appropriate network interface.

When enabled, the `NDMPPREFDATAINTERFACE` option affects all subsequent NDMP filer-to-server operations. It does not affect NDMP control connections because they use the system's default network interface. You can update this server option without stopping and restarting the server by using the **`SETOPT`** command (set a server option for dynamic update).

NetApp file servers provide an NDMP option (`ndmpd.preferred_interface`) to change the interface used for NDMP data connections. For more information, see the documentation that is included with your NAS device.

Configuration 2: Library connected to the NAS server

In this type of configuration, Tivoli Storage Manager uses NDMP to back up a NAS file server to a library device directly attached to the NAS file server. The NAS file server, which can be distant from the Tivoli Storage Manager server, transfers backup data directly to a drive in a SCSI-attached tape library. The Tivoli Storage Manager server controls library robotics by sending library commands across the network to the NAS file server. The NAS file server passes the commands to the tape library. Any responses generated by the library are sent to the NAS file server, and passed back across the network to the Tivoli Storage Manager server. This configuration supports a physically distant Tivoli Storage Manager server and NAS file server. For example, the Tivoli Storage Manager server might be in one city, and the NAS file server and tape library are in another city.

Figure 7-55 shows the NAS file server with control of the library.

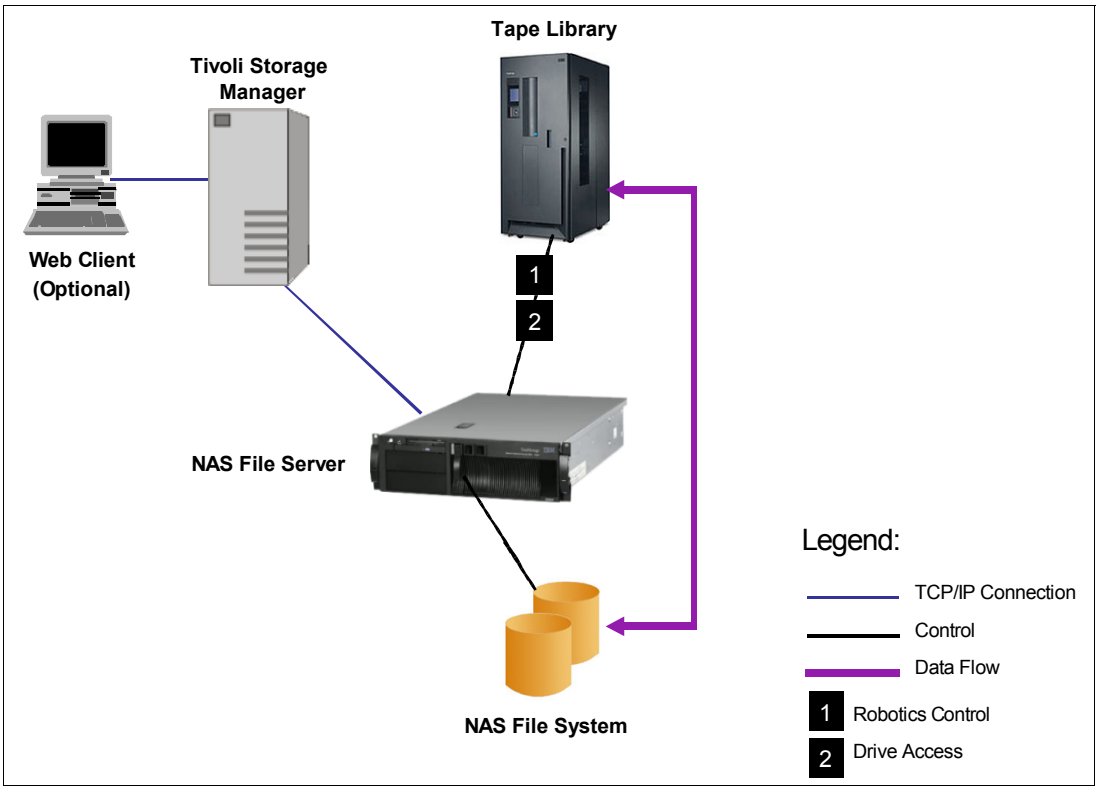

Figure 7-55 NAS file server controls the library

Configuring Tivoli Storage Manager to use NDMP

Assuming that the tape library has been correctly attached to the NAS filer and (if required) the Tivoli Storage Manager server (configuration 1 and 2), and the device configured for use, these are the basic steps to be completed in the Tivoli Storage Manager server:

1. Define the library.
2. Define the drives and their associated paths.
3. Define a device class for NDMP operations.
4. Define the storage pool for backups performed by using NDMP operations.
5. Optionally select or define a storage pool for storing tables of contents for the backups.
6. Configure Tivoli Storage Manager policy for NDMP operations.
7. Register the NAS nodes with the server.
8. Define a data mover for the NAS file server.
9. Label and check in media to the library.

For the procedure, see the *Quick NDMP Setup for Tivoli Storage Manager* white paper:

https://www-304.ibm.com/software/brandcatalog/ismlibrary/details?catalog.label=1TW10SM36#

NDMP operations require a specialized device class that is not set up in the default configuration. Therefore, you must define it yourself. Indicate this information in the device class definition:

- Specify NAS as the value for the `DEVTYPE` parameter.
- Specify the `MOUNTRETENTION=0` setting, which is required for NDMP operations.
- Specify a value for the `ESTCAPACITY` parameter. This should correspond to the nominal capacity of the tape media being used for the tests.

NDMP media requires its own distinctly formatted storage pool. Only one data format is available for general use, NDMPDUMP. Use `dataformat=ndmpdump` when you define the storage pools for NDMP backups.

For NDMP backups, rather than use the standard backup client we use a data mover, which provides the Tivoli Storage Manager server with a communication path to the NAS filer. Each data mover must be associated with a node, which is the first parameter after the **`define datamover`** command. The data format must be specified as `ndmpdump`. The usual low level address is 10000.

When it is time to define library and drive paths, whether you use configuration 1 or 2 will determine how you set them up:

- For configuration 1 the Tivoli Storage Manager server will control the library robotics, so a path must be defined between the server instance and the robotics port.
- For configuration 2, robotics control is done by the NAS filer, so a path must be defined between the NAS node and the library.

Note: The device name to use in this case is the device parameter used by the NAS server to identify the library robotics.

Hardware and software requirements

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 through the same Tivoli Storage Manager server. Drives can be also shared with LAN-free backup and restore operations, provided that the library is controlled directly by the Tivoli Storage Manager server.

Tape drives can be dynamically shared in the Tivoli Storage Manager with the use of library manager. See Figure 7-56 for details.

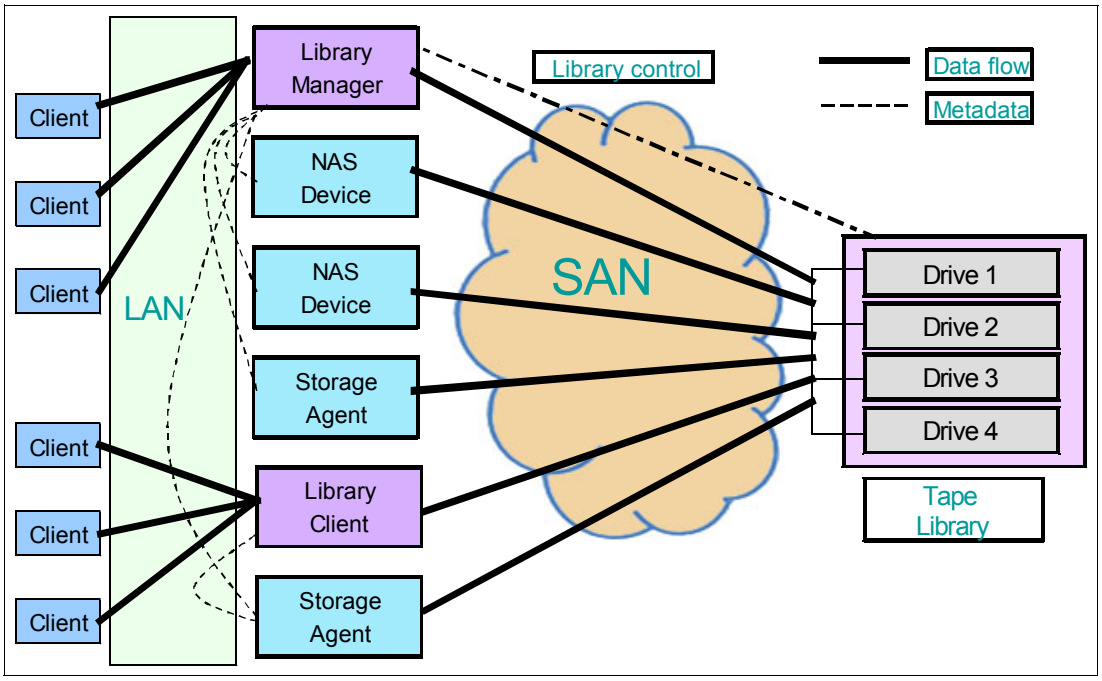

Figure 7-56 Dynamically sharing tape drives

7.7.6 Off-host backup using snapshot differencing

When large file servers are moved to network-attached storage (NAS), such as NetApp or IBM System Storage N series, backup can be taken from a network-attached Tivoli Storage Manager backup-archive client. To avoid long file systems scan during incremental backup, use snapshot differencing API feature, included in the ONTAP operating system of NetApp and N series.

For NAS and IBM System Storage N series file servers that are running ONTAP 7.3.0, or later, you can use the `snapdiff` option to start the snapshot differencing backup from NetApp when you run a full-volume incremental backup. Using this option reduces memory usage and is faster.

Consider the following restrictions when running a full-volume incremental backup using the `snapdiff` option, to ensure that data is backed up when it should be.

- A file is excluded because of an exclude rule in the include-exclude file. Tivoli Storage Manager runs a backup of the current snapshot with that exclude rule in effect. This happens when you have not made changes to the file, but you have removed the rule that excluded the file. NetApp will not detect this include-exclude change because it detects only file changes between two snapshots.

- If you added an include statement to the option file, that include option does not take effect unless NetApp detects that the file has changed. Tivoli Storage Manager does not inspect every file on the volume during backup.
- If you used the **`dsmc delete backup`** command to explicitly delete a file from the Tivoli Storage Manager inventory, NetApp cannot detect that a file was manually deleted from Tivoli Storage Manager. Therefore, the file remains unprotected in Tivoli Storage Manager storage until it is changed on the volume and the change is detected by NetApp, signaling Tivoli Storage Manager to back it up again.
- Policy changes such as changing the policy from `mode=modified` to `mode=absolute` are not detected.
- The entire file space is deleted from the Tivoli Storage Manager inventory. This action causes the `snapdiff` option to create a new snapshot to use as the source, and a full incremental backup to be run.

The NetApp software (not Tivoli Storage Manager) determines what is a changed object. If you run a full volume backup of an NFS-mounted or a CIFS-mapped NetApp or N series volume, all the snapshots under the snapshot directory might also be backed up. To avoid this situation, you can do one of the following actions:

- Run NDMP backups.
- Run backups using the `snapshotroot` option.
- Run incremental backups using the `snapdiff` option.

> **Tip:** If you run an incremental backup using the `snapdiff` option and you schedule periodic incremental backups, use the `createnewbase=yes` option with the `snapdiff` option to create a base snapshot and use it as a source to run an incremental backup.

- Exclude the snapshot directory from backups.
  - On Windows systems, the snapshot directory is in `~snapshot`.
  - On AIX and Linux systems, the snapshot directory is in `.snapshot`.

> **Note:** The `.snapshot` directory is not backed up for some versions of Red Hat Linux, so you are not required to exclude it.

Challenges the solution addresses

If you must back up large file systems that are network-attached through CIFS or NFS and are located on a NetApp or N series system, a solution described here can help you solve this problem.

The solution provides the following benefits:

- Significant data growth
- Reduced backup time
- Network bandwidth optimization (incremental, data deduplication)
- Reduced RPO, increased the backup frequency
- Reduced data management cost: total cost of ownership (TCO), operational expense (OPEX)

Solution architecture

To increase the probability of completing nightly backups for NetApp filers or N series systems and to continue to provide the highest level of Tivoli Storage Manager data protection, we consider the use of snapshot differencing for our daily backups.

Snapshot differencing is a method that Tivoli Storage Manager provides in combination with the application interface that NetApp supports to use the snapshot functionality to determine what data changed and now must be backed up.

With Tivoli Storage Manager Version 6.4 we have a full function solution package for NetApp and N series file servers, which can use snapshot functions to provide an optimized backup strategy. This includes the following items:

- File Access Protocol
  - Data OnTAP 7.3.3 and 8.1 have a feature named *File Access Protocol* on the snapshot differencing API, but not on OnTAP 8.0.
  - Enables snapshot differencing to handle file names with 7-bit ASCII characters and also non 7-bit ASCII character correctly either through NFS or CIFS.
- Supports NetApp SnapDiff vFiler in ONTAP 8.1.1.
- Snapshot differential backup of SnapMirror-restricted (read-only) mirror volumes.

This gives you overall flexibility to back up the network-attached file systems that are hosted on a NetApp or N series system.

Also know that you can use your standardized backup methodology which is incremental forever on a file level basis.

Solution description

As mentioned in Chapter 3, “Data protection with Tivoli Storage Manager” on page 33, the NetApp snapshot difference API can be used to remove the long scan time Tivoli Storage Manager incremental backup needs to find the data that needs to be backed up.

The N series or NetApp volumes should be connected through NFS or CIFS over a NAS network to Linux, AIX, or Windows Backup-Archive client hosts. On Linux or AIX, files are accessed through NFS using UNIX file semanticsm and with secure access using Kerberos authentication. On Windows, files are accessed through CIFS shares using Windows file semantics with NTFS Security style.

Data OnTAP 7.3 or later is supported. Only entire volumes can be specified on the incremental command. Read-only volumes are not supported, unless in a SnapMirror relationship. Qtrees and subdirectories are not supported. vFilers are now also supported.

MultiStore (vFiler) support

Snapshot differencing backups can now be done on NetApp vFiler volumes, if the targeted NetApp vFiler is running ONTAP 8.1.1, or a newer version. The software changes that allow the client to protect vFiler volumes are all in ONTAP. There are no changes to the backup-archive client to use this new capability. You use same client commands and options to protect vFilers that you use for physical filers. Like a physical filer, a vFiler has unique credentials, including an IP address or host name, and a user name and password. The backup-archive client Set Password command must be used to initially store these credentials before a snapshot differencing backup can be run.

Figure 7-57 explains how snapshot differencing backup is implemented with vFilers.

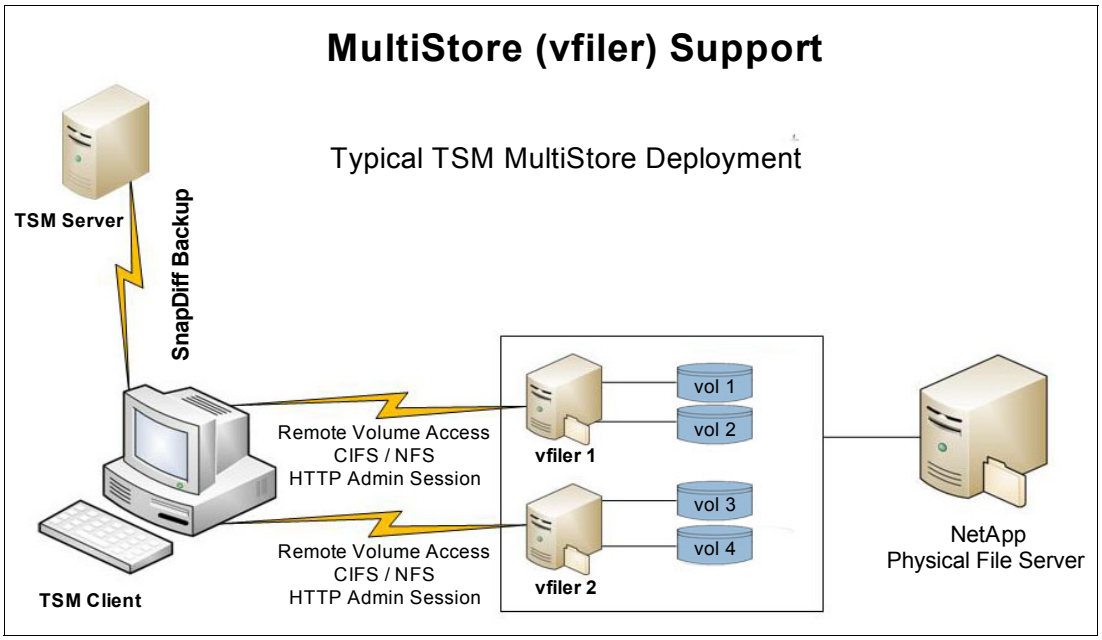

Figure 7-57 Snapshot differencing backup for vFilers support

The details about the new functionality are as follows:

- NetApp snapshot differencing backup vFiler is supported in ONTAP 8.1.1.
- Tivoli Storage Manager 6.40 client is enabled for snapshot differencing backup support of vFilers.
- Client recognizes if a targeted filer has snapdiff vFiler support.
- To Tivoli Storage Manager, vFilers appear the same as physical filers.
- Each targeted vFiler requires these items:
  - Unique credentials that must be set through the **SET PASSWORD** command; these are Host/IP address, user, and password
  - CIFS share definitions and NFS exported volume
- Volume access is mutually exclusive to a particular vFiler including the physical filer (default vfiler0).

SnapMirror support

You can use NetApp SnapDiff backup processing in conjunction with NetApp SnapMirror replication to back up NetApp or N series production or remote filer volumes.

The overview in Figure 7-58 on page 281 shows how the components relate.

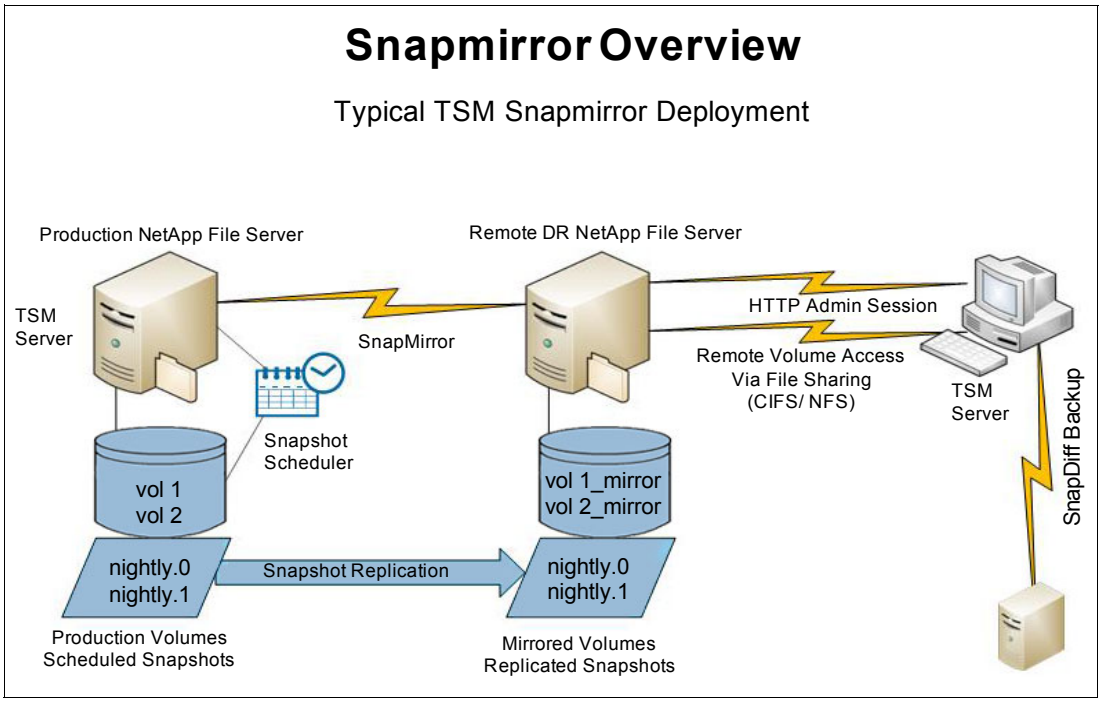

Figure 7-58 Snapshot differential backup in a SnapMirror relationship

The new support allows snapshot differencing backup of SnapMirror restricted (read-only) mirror volumes. By default, Tivoli Storage Manager creates base and difference snapshots on volumes being backed up, which is not possible on read-only mirror volumes.

A typical configuration is to offload the backups from the source filer by creating backups of the source volumes by using the replicated volume snapshots stored on the destination filer. Ordinarily, backing up a destination filer presents a problem because creating a snapshot differencing backup requires that a new snapshot be created on the volume that you are backing up. The destination filer volumes that mirror the contents of the source volumes are read-only volumes, so snapshots cannot be created on them.

To overcome this read-only restriction, Tivoli Storage Manager provides client configuration options so you can use the existing base and differential snapshots on the read-only destination volume to back up changes to the Tivoli Storage Manager server.

For more Information about the N series functions in Tivoli Storage Manager, see *IBM Tivoli Storage Manager for Windows Backup-Archive Clients Installation and User's Guide,* SC23-9792.

Use scenarios

NetApp Snapshot Differencing was implemented for environments with “big fat filers” that house millions of files, with unacceptable backup windows. Using incremental backup by NetApp, the Tivoli Storage Manager Client does not need to crawl the file space looking for changed files, but instead queries the ONTAP OS on a filer for files that have changed since the last **`-snapdiff`** or **`-diffsnapshot`** backup.

This method specifically increases the speed of incremental backups of filers with many files, and with a small percentage of changed files. In a classic incremental backup, the Tivoli Storage Manager Client can spend hours on a large file system, searching for changed files. In both extremes, NetApp Snapshot Differencing performs better than classic incremental, but especially so when relatively few files have been changed. Testing was also performed on a

small file system to ensure that no performance degradation occurred when using NetApp or IBM System Storage N series on such a system.

The solution is combined with your traditional incremental forever backup strategy. For the backup window, this means that scan time is reduced to find what data has changed. You can compare it with journal-based backup, where the compare-time is also reduced. To determine what time-slice in your overall backup time the compare time is, query the summary table, as shown in Example 7-5, to find the idle time amount.

Example 7-5 Query summary table

```
select * from summary where entity=“BACKUP”
```

The first incremental backup taken with the **-snapdiff** option creates the base snapshot, to which the next incremental backup will reference. After this point, the idle time in the summary table should tend to zero.

When you use the incremental backup with snapdiff option, remember the side effects that can occur. A file change or metadata change results in file backup. Mass directory changes are handled on the directory level so that only an incremental backup of that directory takes place. If a file is deleted from the local system, the active version is expired on the Tivoli Storage Manager server. That means the full version and retention control on an object basis will be considered. A file that is unexpectedly removed from the Tivoli Storage Manager server results in a new file backup. Because the Tivoli Storage Manager server is not returning information about current version, events such as file deletion from Tivoli Storage Manager server, new include or exclude rules, copy mode, and frequency are not detected and honored.

If you have an environment with NetApp or IBM System Storage N series filers in a SnapMirror relationship, perform the incremental backup with the snapdiff option from the system to which the snapshot is synchronized. That means, refer to the snapshot, created on the primary system and copied to the SnapMirror system, with these options: -useexistingbase, -diffsnapshot, -basesnasphotname, and -diffsnapshotname.

Example 7-6 shows how these options are used in the command to back up the latest nightly snapshot, created by the default snapshot scheduler.

Example 7-6 Sample SnapMirror incremental backup

```
dsmc incr x: -snapdiff –useexistingbase –diffsnapshot=latest
–basesnapshotname=nightly.? –diffsnapshotname=nightly.?
```

For information about using this command see these resources:

- I*BM Tivoli Storage Manager for Windows Backup-Archive Clients Version Installation and User's Guide,* SC23-9792
- *IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Version Installation and User's Guide,* SC23-9791

With this solution you are able to manage your files and objects on a single object level. The interface for backup and restore is the standard backup-archive client. Command line, GUI client, and web client can all be used. Each version of each file will be stored as a single object in the database. It fits in the standard configuration, so that no special knowledge for the backup administrator is necessary.

Software requirements

The solution, to back up your data that resides on NetApp or N series file servers with the snapdiff function has several software levels requirements:

- Tivoli Storage Manager Server Version 6.1 or later
- Tivoli Storage Manager Backup-Archive Client Version 6.4 or later
- Ontap Version 8.1.1 7 mode or later

Table 7-11 lists which levels of the 6.2, 6.3, 6.4, and 7.1 IBM Tivoli Storage Manager Windows, AIX, and Linux x86/x86_64 Backup-Archive clients are supported with which levels of NetApp Data ONTAP for the NetApp snapshot-assisted progressive incremental functionality.

Table 7-11 Software requirements

Tivoli Storage Manager Backup-Archive Client levels	Platforms supported	NetApp Data ONTAP levels[1]
6.2.0 and 6.2.1	Windows, AIX 64-bit	7.3.0, 7.3.1, and 7.3.2
6.2.x (where x is 2 or higher)	Windows, AIX 64-bit, and Linux x86	7.3.y (where y is 0 or higher)[2] and 8.0.1
6.3	Windows, AIX 64-bit, and Linux x86_64	7.3.y (where y is 0 or higher)[2] and 8.0.1
6.4	Windows, AIX 64-bit, and Linux x86_64	7.3.y (where y is 0 or higher)[2] and 8.1.1[3]
7.1	Windows, AIX 64-bit, and Linux x86_64	7.3.y (where y is 0 or higher)[2] and 8.1.1[3]

1. Snapshot-assisted backup is supported only for supported versions of ONTAP running in 7-mode. C-mode is not supported.
2. With 6.2.2 or higher clients, Tivoli Storage Manager recommends using NetApp Data ONTAP level 7.3.3 or higher levels of 7.3 (7.3.y where y is 3 or later), due to the Unicode support at those levels of Data ONTAP.
3. Support for backing up vFiler volumes requires ONTAP version 8.1.1 or later.

Additional information

For further information about implementing the snapshot differencing solution, see the following resources:

- *Using the IBM System Storage N series with IBM Tivoli Storage Manager,* SG24-7243:

  http://www.redbooks.ibm.com/abstracts/sg247243.html

- Snapshot Differencing Caveats page in IBM developerWorks highlights peculiarities in the way the NetApp snapshot differencing works as compared to standard backups:

  https://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli%20Storage%20Manager/page/Snapshot%20Differencing%20Caveats?lang=en

7.7.7 GPFS, SONAS and V7000 Unified

IBM General Parallel File System (GPFS) is a high performance shared-disk file management solution for Windows, AIX, and Linux and is also used in IBM Scale Out Network Attached Storage (SONAS).

The following solutions apply to GPFS file systems:

- File level backup with **mmbackup** command and Tivoli Storage Manager client (preinstalled in SONAS and V7000 Unified)
- Integration with information lifecycle management (ILM) with hierarchical storage manager (HSM)

Challenges and benefits that the solution addresses

GPFS stores petabytes of data and with customers dealing with billions of files and large amounts of data. Obviously, working with large amounts of data and large numbers of files leads to longer duration of time for backup in today's industry, where there is a demand to reduce backup duration. To meet this requirement, the backup facility in GPFS provides a fully-integrated Tivoli Storage Manager client.

In case of restoring data, selective restore operations can be performed without the need to restore a full backup.

With the Tivoli Storage Manager server, administrators of SONAS can do these tasks:

- Perform a scheduled full-incremental backup of a file system.
- Perform an individual full-incremental backup of a file system started manually
- Restore several single file system objects, objects according to file pattern, whole directories, and all subdirectories or a whole file systems from Tivoli Storage Manager server.
- Query previous or active backup sessions.
- Show the log files generated by backup and restore sessions.

SONAS provides clustered serviceability to its users. In case of a failure of interface nodes, access to data is provided through the alternate interface nodes.

- If helper backup nodes fail during the active backup session, all backup processes of that interface node are redistributed to the remaining backup nodes and backup continues to the Tivoli Storage Manager server.
- If the primary backup node fails during the active backup session, the entire backup fails. You will need to restart the backup again after the interface node becomes healthy.
- Because the restore occurs only on a single interface node, if that particular interface node fails during the active restore session, the entire restore process fails. You must restart the restore process after the interface node becomes healthy.
- If the management node fails during the backup or restore session, the backup or restore session continues. After the management node returns to a good state, the status of backup or restore request along with the corresponding result is reported to the management node.

Solution architecture

The Tivoli Storage Manager client is directly installed on the SONAS interface nodes; data is backed up directly from the interface nodes to the Tivoli Storage Manager server as shown in Figure 7-59 on page 285. SONAS integrated the GPFS policy scan engine internally to scan the inode table of a file system. This is much faster than the conventional system calls, reducing the overall time required for the backup.

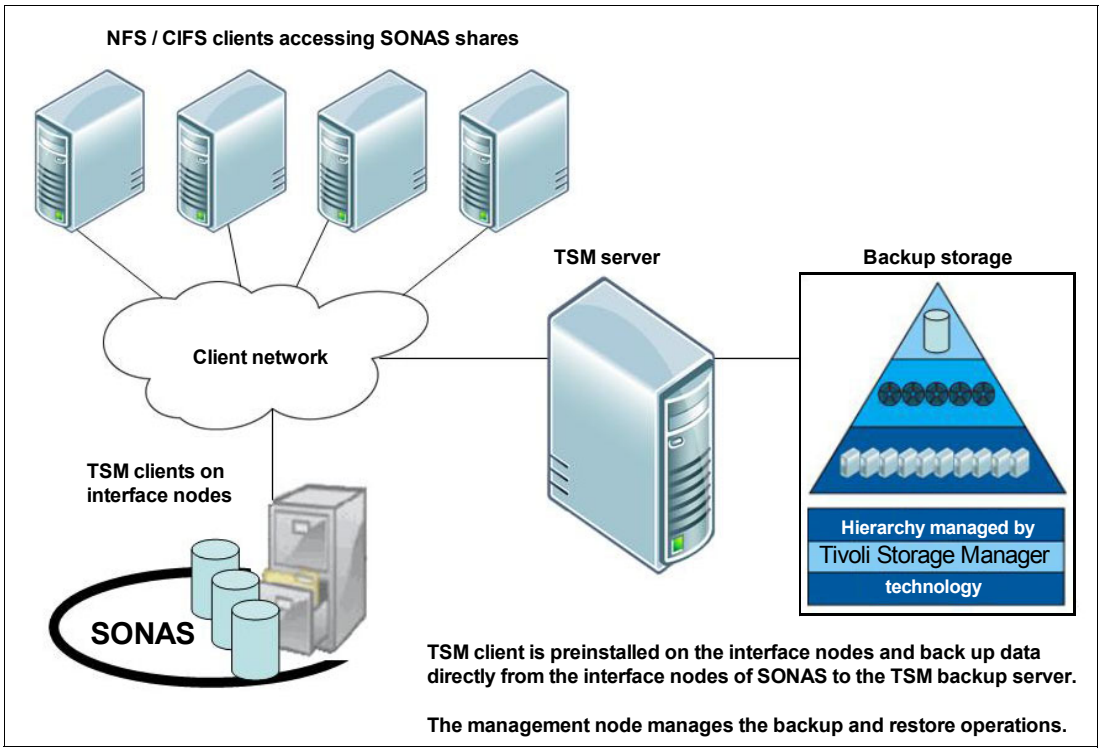

Figure 7-59 SONAS with Tivoli Storage Manager

Multiple Tivoli Storage Manager servers can back up a SONAS system. All Tivoli Storage Manager servers are connected to the network and reachable from the interface nodes of SONAS. It is possible to have one Tivoli Storage Manager server configured per file system but it is not possible to have multiple Tivoli Storage Manager servers backing up the same file system. All Tivoli Storage Manager servers do not have to be on the same release. However, verify that they all are compatible with the Tivoli Storage Manager client version on SONAS.

Solution description using mmbackup

Tivoli Storage Manager backup does not use the classical backup process to traverse the file system, query the server, and identify changes. Instead the GPFS command **mmbackup** is run, which does these tasks:

- Identifies candidates using the GPFS inode scanner (reduced time to identify such candidates).
- Generates a list of files that must be expired, and a list of files that must be backed up.
- Distributes parts of the file list as backup jobs across the file modules.

The **mmbackup** command does not need to check each file against the Tivoli Storage Manager server to identify changes. It uses the policy engine and a shadow database of Tivoli Storage Manager inventory to identify changes in the file system. File changes are tracked by GPFS.

Each File Module then operates on its own file-lists, by default spawning one thread. The **mmbackup** command calls the **/usr/lpp/mmfs/bin/mmexectsmcmd** script, which uses the Tivoli Storage Manager commands shown in Example 7-7.

Example 7-7 Backup commands for SONAS

```
dsmc expire -filelist=,name>
dsmc selective -filelist=<name>
```

Each **dsmc** process can establish several sessions to the Tivoli Storage manager server. The restore is basically calling only the Tivoli Storage Manager command shown in Example 7-8.

Example 7-8 Restore command syntax

```
dsmc restore <pattern> <targetdir>...
```

Several scripts are provided to define the File Modules in the backup, the relationship of which file system needs to be backed up to which Tivoli Storage Manager server, and to start and stop backups or restores.

Use scenarios

The integrated Tivoli Storage Manager client solution enables enterprises to back up and restore data in a minimum time period. It increases efficiency, performance, and reliability of the backup and restore operations. Taking a full backup every time ensures reliability of backups. However, not all the data changes every time a new backup is invoked, and the time required for a full backup increases the backup window. To reduce duplication of backed up data, you need to back up only those files and objects that have been changed or created since the last backup. This identification can be done with a full file system traversal using standard system calls, but the time increases linearly with the number of objects in the file system. SONAS integrated IBM General Parallel File System (IBM GPFS) policy engine scans the *inode* table of a file system, which is much faster than the conventional system calls reducing the overall time required for the backup. SONAS also leverages its multiple interface nodes in the backup process to further reduce the time needed for backups.

Tivoli Storage Manager client software is preinstalled on SONAS on each of the interface nodes. This is done during the SONAS software installation. Not all interface nodes need to participate in the backup and restore activities. All interface nodes that are configured for backup purposes are called *backup nodes*. The first backup node that is configured with Tivoli Storage Manager Server is called the *primary backup node*. Other backup nodes are called as *helper backup nodes*.

See the *Integrating IBM SONAS with IBM Tivoli Storage Manager* white paper for more details:

http://public.dhe.ibm.com/partnerworld/pub/whitepaper/19356.pdf

GPFS and HSM

GPFS has a built-in disk-based information lifecycle management (ILM) implementation with the storage pool concept. Tivoli Storage Manager for Space Management is a hierarchical storage management (HSM) solution. Tivoli Storage Manager for Space Management integrates tape and external storage pools into the GPFS ILM solution (see Figure 7-60 on page 287).

GPFS provides a single name space across all pools. Files in the same directory can be in different pools. Files are placed in storage pools at creation time using placement policies. Files can be moved between pools based on migration policies and files can be removed based on given policies.

GPFS implements an engine for fast file system scans and file-system monitoring, which is also known as policy engine. The policy engine implements a user programmable interface which is called the GPFS policy rule language. This section focuses on the use of this engine in conjunction with the Tivoli Storage Manager for Space Management client.

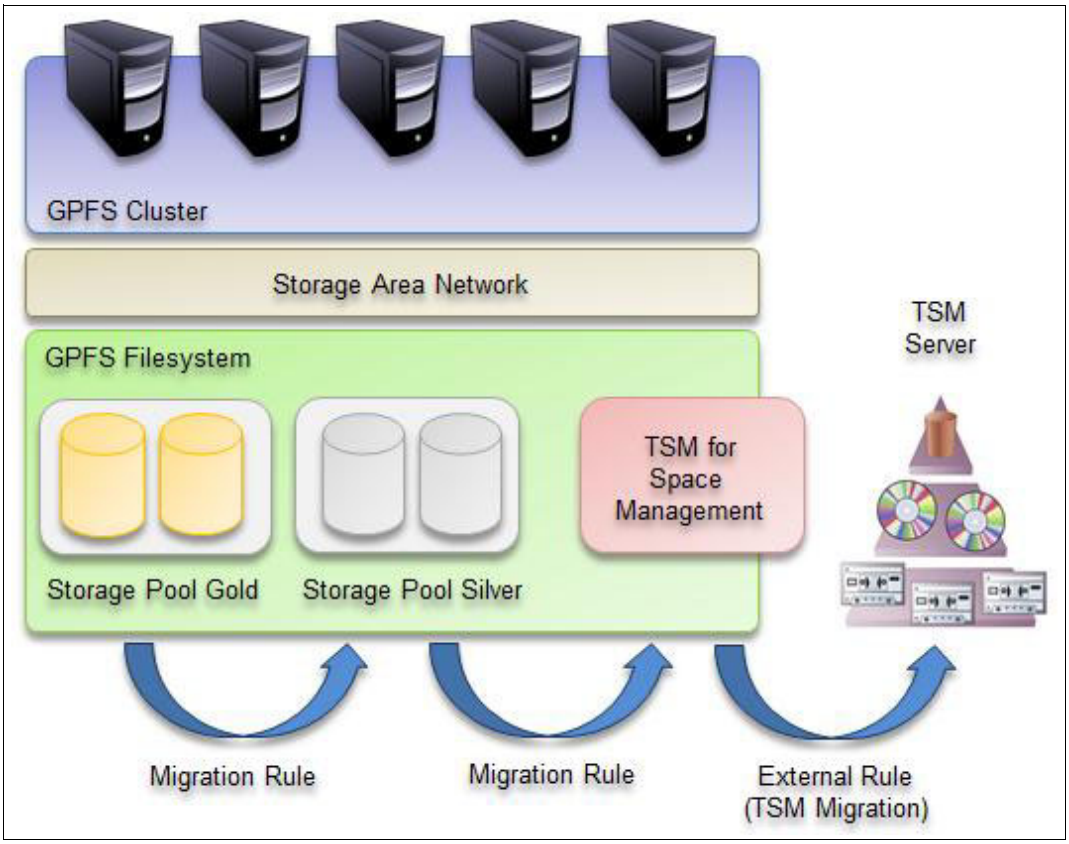

Figure 7-60 Tivoli Storage Manager for Space Management integration with GPFS

GPFS 3.2 introduces a new feature called external storage pools. You can set up external storage pools and GPFS policies allowing the GPFS policy manager to coordinate file migrations from a native GPFS online pool to external pools on the Tivoli Storage Manager server. The GPFS policy manager invokes the migration through the HSM client command line interface.

The migration candidate selection is identical to the GPFS native pool-to-pool migration rule. The Policy Engine uses scripts to call the **`dsmmigrate -filelist`** Tivoli Storage Manager command for the migration of files from a native storage pool to the Tivoli Storage Manager server.

An HSM solution typically moves the file's data to back-end storage and leaves a small stub file back in the local storage. The stub file consumes minimal space but leaves all metadata information on the local storage in such a way that for a user or a program, the file looks like a normal local stored file. When the user or a program accesses the file, the HSM solution automatically recalls (moves back) the file's data from the back-end storage and gives the reading application access to the file when all the data is back online. The back-end storage for Tivoli Storage Manager HSM is the Tivoli Storage Manager server, which handles tier2 (disk storage) and tier3 (tape storage) of the storage hierarchy. That means Tivoli Storage Manager HSM virtually extends the managed file system with the space that is provided by the Tivoli Storage Manager server. Migrating files to the Tivoli Storage Manager server frees space for new data on your local file system and takes advantage of lower-cost storage resources that are available in your network environment.

The whole functionality to perform the virtual file space expansion is implemented in the product. In GPFS environments, the preference is to use the GPFS and HSM integration because of scalability and performance reasons. For more information, see *TSM for Space Management for UNIX – GPFS Integration*, which you can download from the following address:

http://www-01.ibm.com/support/docview.wss?uid=swg27018848&aid=1

Setting up external storage pools and GPFS policies allows GPFS policy manager to coordinate file migrations from a native GPFS online pool to external pools on the Tivoli Storage Manager server. The GPFS policy manager invokes the migration through the HSM client CLI.

HSM client on AIX and Linux GPFS, and AIX JFS2 support LAN-free data transfer. With assistance from GPFS, the need for a file system scan with HSM is eliminated, allowing Tivoli Storage Manager for Space Management to migrate data more efficiently.

V7000 Unified

IBM Storwize V7000 Unified is a unified storage product, providing file storage through network-attached storage (NAS) protocols, and block storage through block storage protocols such as Fibre Channel and iSCSI. The NAS part of V7000 Unified allows storing files through standardized protocols such as NFS and CIFS.

V7000 Unified includes backup module with Tivoli Storage Manager backup client (V 6.3). It can be configured to back up data to external Tivoli Storage Manager server.

Incremental backups are performed per file system via LAN interfaces, and exploit IBM Active Cloud Engine® for fast identification of files. Tivoli Storage Manager client architecture is shown in Figure 7-61.

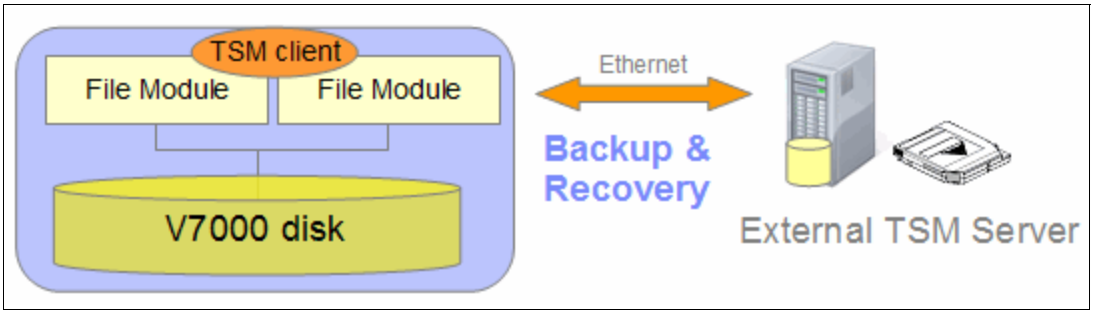

Figure 7-61 V7000 Unified and Tivoli Storage Manager client

The backup is scheduled on V7000 Unified that starts the **mmbackup** utility. That utility invokes Tivoli Storage Manager backup client to copy files to the Tivoli Storage Manager server. The backup client can be configured to run on both file modules in parallel.

Restoring files, directories, or file systems can be done through the GUI or CLI, but only one restore process at a time.

V7000 Unified can also be configured to automatically move a file between different tiers of storage. The integrated information lifecycle management (ILM) function allows the moving of files between disk-tiers, which are managed by the V7000 system. Furthermore the HSM capability allows moving files from the V7000 managed disk storage to an external Tivoli Storage Manager server while keeping the access to the file transparent.

The V7000 Unified includes Tivoli Storage Manager HSM client for migrating files from V7000 Unified managed disk to tape, which is attached to an external Tivoli Storage Manager server.

Figure 7-62 shows the general architecture of HSM in the V7000 Unified and the external Tivoli Storage Manager server.

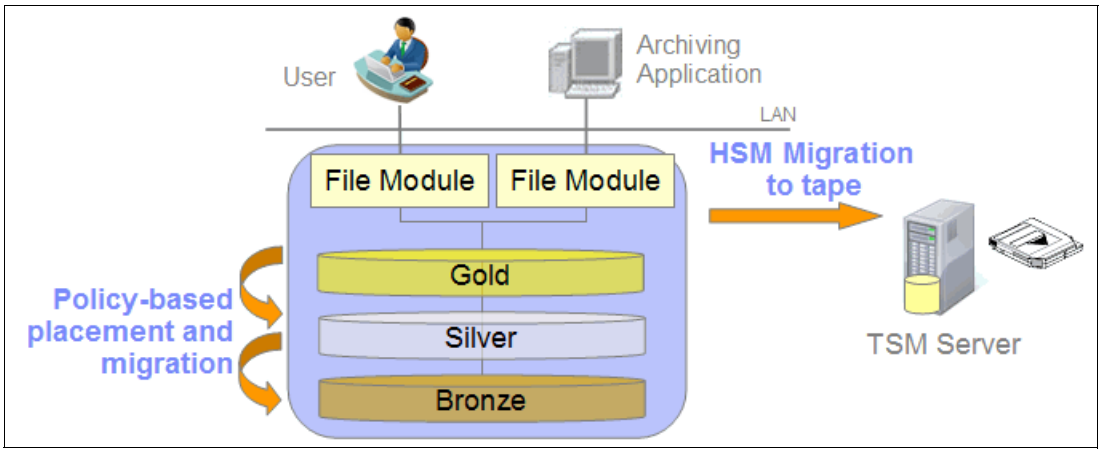

Figure 7-62 V7000 Unified and HSM

The local Active Cloud Engine, which is an integral part of the V7000 Unified, in combination with the Tivoli Storage Manager HSM client and Tivoli Storage Manager server is used to perform policy-based migration of files from V7000 managed disk to external tape, managed by an external Tivoli Storage Manager server. For example, policies can be based on age or last-access time of files, whereby files older than 1 year or files that have not been accessed for the last 30 days can be migrated to an external Tivoli Storage Manager server with attached tape storage.

The Tivoli Storage Manager HSM client can use the same Tivoli Storage Manager server as the Tivoli Storage Manager backup client, which is configured for the file system. Using the same Tivoli Storage Manager server for file system backup and migration enables the inline backup function (when a file that was migrated is backed up, the file will not be recalled and copied again to the Tivoli Storage Manager server). Instead the Tivoli Storage Manager server performs an inline copy of the file from the HSM storage pool to the backup storage pool of the Tivoli Storage Manager server. This improves operational efficiency and reduces the amount of data being transferred over the network.

HSM is enabled on file-system basis and operates on the entire file system. Policies and rules can be configured to operate on a file set basis. If multiple file systems are configured for HSM, one or more Tivoli Storage Manager servers can be used as a migration target, however a single file system can have only one migration destination server.

HSM can be configured to run on one or both file modules of the V7000 Unified system. The recommendation is to configure Tivoli Storage Manager HSM to run on all file modules for all file systems, because the workload (file migration) is distributed among the file modules (approximately 100 files per file module for each chunk).

When files are migrated from the disk storage to an external Tivoli Storage Manager server, these files remain represented in the name space of the file share, facilitating transparent access from a user perspective. If the user accesses a migrated file, the HSM client performs a recall of the file from the Tivoli Storage Manager server.

See the *IBM Storwize V7000 Unified - TSM HSM Overview and Implementation* white paper, which shows how to implement HSM in V7000 Unified:

http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102116

Additional information

See the following sources for more information:

- *IBM Scale Out Network Attached Storage Concepts*, SG24-7874:

  http://ibm.com/redbooks/redpieces/abstracts/sg247874.html

- *IBM SONAS Introduction and Planning Guide*, GA32-0716:

  http://publib.boulder.ibm.com/infocenter/sonasic/sonas1ic/topic/com.ibm.sonas.doc/sonas_ipg.pdf

7.7.8 SnapMirror to Tape

You can back up very large NetApp file systems using the NetApp SnapMirror to Tape feature. Using a block-level copy of data for backup, the SnapMirror to Tape method is faster than a traditional Network Data Management Protocol (NDMP) full backup and can be used when NDMP full backups are impractical.

Use the NDMP SnapMirror to Tape feature as a disaster recovery option for copying very large NetApp file systems to secondary storage. For most NetApp file systems, use the standard NDMP full or differential backup method.

Challenges the solution addresses

Tivoli Storage Manager as a data protection solution can help solve the challenges you have to protect big data in an unstructured format on NetApp or N series filers.

In this way, the solution provides benefits as follows:

- Significant data growth
- Reduce backup time
- Fast recovery
- Reduce data management cost (TCO, OPEX)
- Reduce RPO, increase the backup frequency
- Business continuity

Solution architecture

Tivoli Storage Manager NetApp and N series integration enhancements address the backup challenges associated with large filers. NDMP SnapMirror to Tape feature provides an efficient disaster recovery or vaulting solution.

NDMP backups of very large NAS devices can be slow. SnapMirror to Tape, also called SMTAPE at ONTAP 8, is a block-level movement of a flexible volume or volume. The function is similar to backup-archive client backup image, which is covered in 7.7.3, “Image backups” on page 266. The data written to tape is a block by block copy of the file system. No time is wasted to search through the file system. The block structure is preserved so that deduplicated data savings are kept. Another advantage is that the whole snapshot structure is saved and will be available after restore process. Because the disk geometry is used, restoration of traditional volume to unlike geometry can be very slow. The implementation is provided by extending NDMP functions.

The transport layer for the backup and restore with SnapMirror to Tape uses the standard NDMP protocol. You have the choice to control the data flow. Either, filer-to-tape, which means LAN-free data transfer, or filer-to-server through LAN can be used, controlled by the destination storage pool definition. Only the additional option `type=snapm` is necessary to switch between traditional NDMP backup and SnapMirror to Tape. Every backup process will create a separate backup version of the volume on the Tivoli Storage Manager server. Even if

you use traditional NDMP backup in combination with SnapMirror to Tape, the data in the file space will be separated.

To restore a volume, and it is always a whole volume, it must be prepared and put in restricted mode. After the restore, it will automatically switch to normal mode. The destination of the retrieval must use the same or later version of Data ONTAP. There is a different file system format between traditional volume and the Flexvol volume; the FlexVol volume cannot be restored to a traditional volume and vice versa. Restoring to the same disk geometry as backup should be done for best performance results.

Information about SnapMirror copies can be displayed with the following command:

```
Query NASbackup Type=SNAPM
```

SnapMirror to Tape is a preferred solution for disaster recovery or offsite vaulting. It is not a replacement for traditional file-based backup like incremental forever or full or differential Tivoli Storage Manager NDMP backup.

Consider the following factors:

- Always use full backups of SnapMirror images.
- Backups are taken on a volume level.
- No table of contents is created or used.
- At the start of a SnapMirror to Tape copy operation, the file server generates a snapshot of the file system. NetApp provides an NDMP environment variable to control whether this snapshot should be removed at the end of the SnapMirror to Tape operation. Tivoli Storage Manager always sets this variable to remove the snapshot.
- After a SnapMirror to Tape image is retrieved and copied to a NetApp file system, the target file system is left configured as a SnapMirror partner. NetApp provides an NDMP environment variable to control whether this SnapMirror relationship should be broken. Tivoli Storage Manager always breaks the SnapMirror relationship during the retrieval. After the restore operation is complete, the target file system is in the same state as that of the original file system at the point-in-time of backup.

Use scenarios

If you need a fast disaster protection of your NetApp or N series filers, SnapMirror to Tape is a valid solution. You can do a quick protection of whole volumes for fast restores. The big advantage is that all snapshots including all pointers will be available after restore.

Consider that this kind of data protection is almost an addition to any of the standard protection solutions you will use. For example, you will do a daily NDMP based file level backup and at the weekend an additional backup with SnapMirror to Tape will be scheduled. Or you can take a daily incremental file level backup with the **snapdiff** option instead. Even if different node names are used and the data resides on different storage pool targets, the data will fit together after the restore process. In the case of a disaster, you will restore the broken volumes completely from the SnapMirror to Tape image and later on the files, which are changed after the time that the SnapMirror to Tape backup was taken. You can either use the files saved with NDMP or the incremental backup. If restoring the files from the incremental backup, the option **-ifnewe**r can help you select the involved objects.

The **-ifnewe**r option replaces an existing file with the latest backup version only if the backup version is newer than the existing file.

The only disadvantage is that files deleted between the SnapMirror to Tape backup and the incremental backup are restored and available after the process. This must be manually cleaned up.

Figure 7-63 shows how the solutions work together. While the standard backup method for all filers is NDMP and can also be used for archiving, SnapMirror to Tape should be used for disaster protection.

NDMP Engine and Ontap-Version→	DUMP	SM2T Ontap 7.x	SMTape 8.0.x
recommended backup use case	Tape DP Archival	Tape Disasterprotection	
File- or block level backup ?	File level	Block level	
Backupspeed Volumes with medium/large files	high	high	
Backupspeed Volumes with millions of small files	limited (file-level)	high	
License costs NetApp?	no	no (no SnapMirror License)	
PVR for DP usage at NetApp?	no	yes	No
Tape format	Dump	SM2T	SMTape
supported Volume-Modes	7-Mode and Clustered ONTAP	7-Mode	currently 7-Mode only
Tape format backward support	Longtime (10 Years)	2 Ontap-Versions only (~3 Years)	
backup modes	Full / filelevel Incr/Diff	Full; with ONTAP 8.1.1: +blocklevel Incr / Diff	
single file / directory restore?	yes	No	
# included Snapshots (backup & restore)	1 (typical "last" used)	All (good for DR & easy Snapshot Integration)	
Keeping dedupe/compression/cloning benefits	No, rehydrate	Yes, densing benefit are kept (=faster)	
Protocol (NDMP DMA / CLI-Command)	NDMP-Standard	NDMP-Extension type=snapm separate CLI-Commands each	

Figure 7-63 Comparison of different backup solutions

Hardware and software requirements

The solution to backup your data that resides on NetApp or N series file servers with the SnapMirror to Tape function requires several hardware and software requires.

Hardware requirements

The hardware requirements are as follows:

- Tivoli Storage Manager Server on Windows, UNIX, or Linux
- NetApp FAS or IBM System Storage N series

Software requirements

The software requirements are as follows:

- Tivoli Storage Manager Server Version 6.1 or later
- Tivoli Storage Manager Backup-Archive Client Version 6.4 or later
- Data ONTAP Version 7.4 or later

References

For more information about implementing the SnapMirror to Tape solution see *Using the IBM System Storage N series with IBM Tivoli Storage Manager*, SG24-7243:

http://www.redbooks.ibm.com/abstracts/sg247243.html

7.8 Remote office backup

Remote or branch offices are increasingly at the front lines of business; they have the closest contact with customers and partners and therefore have a dramatic impact on the success of the business. Many, if not most, of these offices run autonomously from headquarters and are responsible for managing their own operations, including protecting and retaining the electronic information that they generate. Ignoring the protection and recovery needs of this distributed data is simply not an option.

As companies expand operations into new markets, the percentage of total corporate data in remote offices is increasing. However, many companies are not adequately protecting these assets. Obviously this is a very complicated problem and can be extremely expensive to solve completely, considering all that can go wrong with remote office data:

- Accidentally or maliciously deleted files: You need a way to easily and quickly restore these assets at the file level from a local backup copy of the data.
- Virus, worm, or hacker attack: You need to restore your systems and data to a point in time before the attack.
- Corrupted database: You need to ensure that your backup data is “application-consistent,” meaning that all parts of recent, complex database transactions are included in the backup.
- Disk or server crash: You need to be able to quickly restore at the volume level, and to restore operating systems and applications, even on different hardware.
- Local or regional disaster: You need a recent copy of your data in another location that is outside the potential disaster area, and ability to quickly restore the affected workloads.
- Combine these challenges with the lack of skilled and dedicated IT staff in most remote offices and you can see that much business risk is associated with data in these offices.

We put these challenges together in our solution matrix, as described in Chapter 4, “Tivoli Storage Manager challenge matrix” on page 81.

Consider the following challenges:

- Low network bandwidth
- Short backup window
- Secure data store and transport
- Appropriate restore time
- Re-creation from scratch, when the device is broken or stolen

See Figure 7-64 on page 294 to find the combination.

Challenges										
business					addressed with	technical				
reduce recovery time (RTO)	reduce backup time (meet RPO)	Compliance	Business Continuity	security (secure data transfer, secure data store,	Solution (discussed in this book)	Virtual Environments	Significant data growth	automation for business	Remote office and mobile office data	network bandwidth optimization
X	X	X	X	X	TSM ToolKit	X	X	X	X	X
X	X		X		Hardware Snapshot technology	X	X			
	X				progressive incremental forever	X			X	X
X	X				deduplication	X	X		X	X
	x				compression					X
X			X		instant restore	X	X			
	X				include / exclude	X	X			X
X	X		X		High Availability	X				
X					point-in-time restore	X				
X			X		Disaster Recovery Manager	X		X		
			X		Server-to-Server virtual Volumes				X	

Figure 7-64 Solution matrix for remote office data protection

IBM offers solutions for remote office data protection and recovery that are automated, easy to use, and cost-effective. These solutions offer enterprise-class data protection at small business prices, and integrate seamless with core data storage management solutions in the data center, such as Tivoli Storage Manager.

7.8.1 Remote office client attached to data center server

For smaller offices, a Tivoli Storage Manager backup-archive agent can be deployed on the local servers, and send deduplicated, incremental backup data to a central Tivoli Storage Manager Server.

Depending on what kind of data is stored in the remote office, the corresponding data protections solution should be used. If you have file servers, big or small, and many or few objects, then incremental backup or journal-based backup might be the best choice. See 7.7.1, “Progressive incremental backups” on page 260.

If you have applications, databases, and mail servers to protect, you can use the application protection like Tivoli Storage Manager for ERP and others. If you have a virtual environment, you can use the Tivoli Storage Manager for Virtual Environments to protect this kind of data. For all kinds of data protection consider that the bandwidth to you data center might be limited and you need any kind of data reduction for the data transfer.

To reduce the amount of data sent over the network you can use these solutions:

- Client-side deduplication
- Compression
- Subfile backup (only Windows based client)
- Hardware WAN accelerator like Cisco or Riverbed

Recovery considerations

In all cases, be aware of what happens in a recovery situation. If much data must be transferred over a WAN connection with low bandwidth, your service-level agreements (SLAs) for recovery time (RTO) might be violated. In this case, a recovery plan can help you be prepared so that you can recover data in a reasonable time. Possible solutions are as follows:

- Move your hardware to the data center and connect it to the Tivoli Storage Manager server over a high-speed network.
- Use local backup sets.

7.8.2 Remote office server connected to data center server

Most companies rely on tape backups for data protection and recovery, and enterprise-class tape drives and media are very reliable products when used in automated and controlled environments such as corporate data centers. But in remote offices, the manual processes used in operating the tape backup system might not be as reliable. Tape backup usually involves several manual processes to label, load, unload, tension, ship offsite, reuse and erase tapes, in addition to running backup and restore operations. These processes are often performed by nontechnical office personnel who have other responsibilities and who might have little or no training in backup and recovery procedures. In addition, tape backups can also require application downtime for backup windows, a requirement that can impact productivity in some offices, or require staff to be available to run (or re-run) backups after normal business hours and on weekends.

Trying to recover data in remote offices from tape backups can also be problematic and often requires the help of central IT staff or outside contractors. If you use incremental backups to reduce the time needed to run the backup job, you must restore the last full backup and then each sequential incremental tape, all of which can increase your recovery time while increasing complexity and risk. In many cases, only the most essential data can be recovered. Successful recovery from tape backups also assumes that all backup operations were completed successfully, with no tape errors, labeling errors, or tape loss. Tape is a point-in-time technology, so it can work well for recovering from a variety of data losses. However, the time needed to perform a recovery makes tape a poor choice for remote offices. And given the infrequency of backup jobs, tape probably does not address your more business-critical recovery needs.

For this reason, we show a solution, how you can protect your remote office clients locally, and have also a valid disaster protection of your whole remote office in your data center.

Figure 7-65 shows a sample environment.

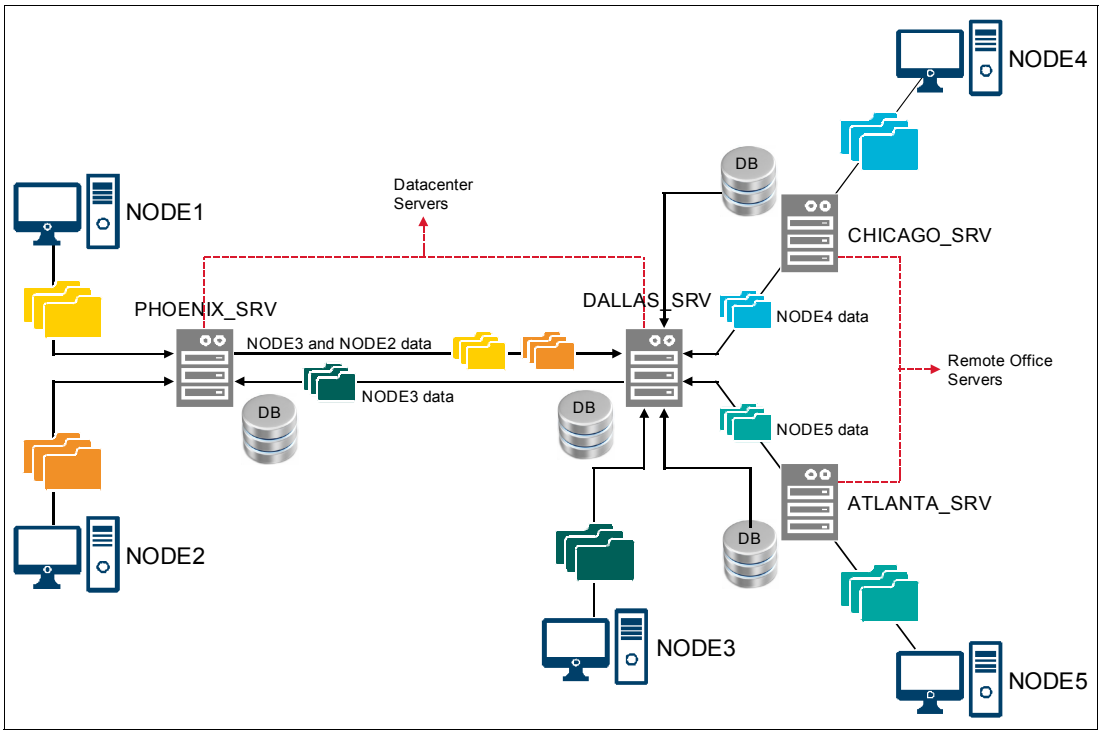

Figure 7-65 Remote office infrastructure connected to data center server

Install a Tivoli Storage Manager instance on your remote site. Use a disk-only solution and keep the total amount of data in disk storage pools. For deduplication purposes, they must be FILE disk storage pools. All your local clients are attached to this Tivoli Storage Manager server. The disk hardware should be reliable storage and RAID-protected to prevent failures and outages. The database and recovery logs should also reside on reliable midrange storage systems like Storwize family. To protect the remote Tivoli Storage Manager server against disaster outages, we do a copy of the database and storage pools with server-to-server virtual volumes to the central Tivoli Storage Manager server in the data center.

Now we have the helpful and improved function of node replication. This will help us to protect the data in a way that the objects and the corresponding metadata will be stored on a target server. There is no need to send the database backup as virtual volumes over WAN to an on-premises data center or to IBM SoftLayer, however, for increasing the level of protection, this is still possible. If you decide to do that, keep in mind that you should also prepare the disaster recovery manager plan file and send it as a virtual volume to the data center. In the case of a disaster outage, you can connect the remote clients directly to the Tivoli Storage Manager server in the data center, without recovery of the server database and storage pools in the remote location. The limitations and challenges in the case of recovery are the same as described in “Recovery considerations” on page 295.

Sometimes, important clients need a very fast recovery during a disaster situation. In this case, it is possible to generate a backup set from the data, replicate data to the target server in the data center, and store it on a storage device, which will be accessed locally on the client. The local backup set restore will be used to recover the data on the client in the remote office. Figure 7-66 on page 297 shows how this is done.

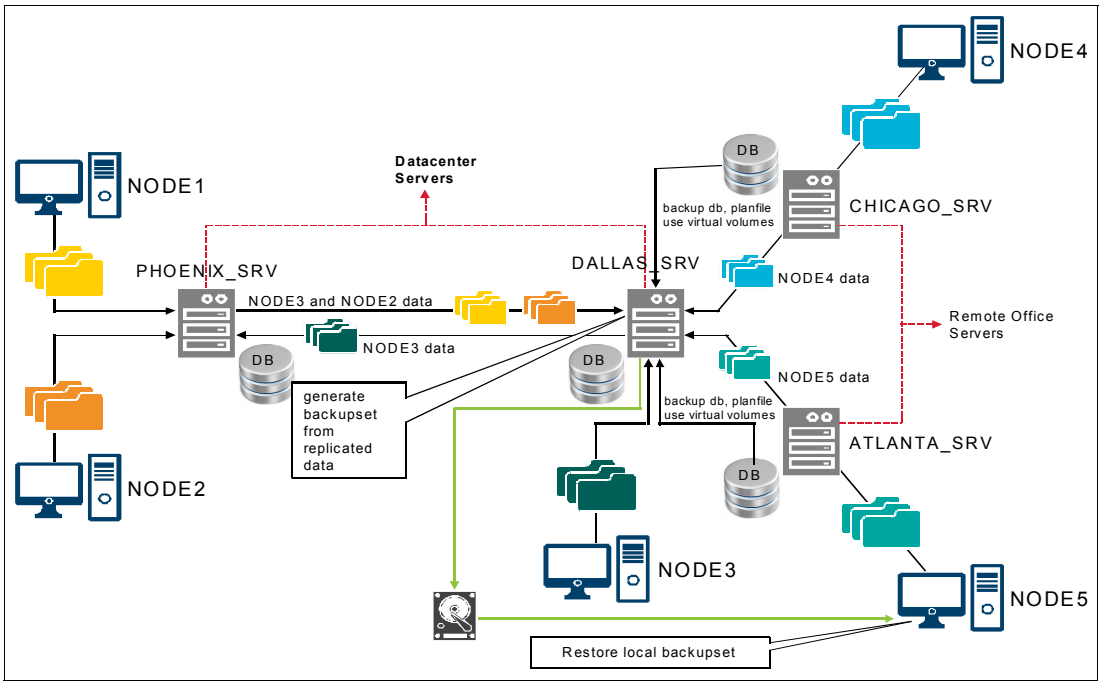

Figure 7-66 Restore local backup set from replicated node data

If you decide to send deduplicated data with node replication to your data center, see *Guidelines for node replication* at the following web page:

https://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli%20Storage%20Manager/page/Guidelines%20for%20node%20replication?lang=en

7.8.3 Remote office server connected to data center server with 3-site copy

For special requirements, you should have three copies of the backup data and they must be stored at different places for disaster protection reasons. As described in the previous section, we have only a two-site data protection, even if we copy the data with node replication to the data center. How can we achieve this challenge? Maybe you consider a solution where you replicate the already-replicated data to the next level, in this case to the disaster protection site. However, this is not possible today. With node replication you can copy the data to only one target.

How can we solve this problem? The Tivoli Storage Manager architecture can easily solve this problem. We use the standard copy pool mechanism, which gives the choice to store the copy storage pool on the primary site using online attached devices such as sequential disk, virtual tape, or physical tape. But we can also use the disaster recover functionality to send the copy pool data offsite.

Figure 7-67 on page 298 shows an overview for this solution.

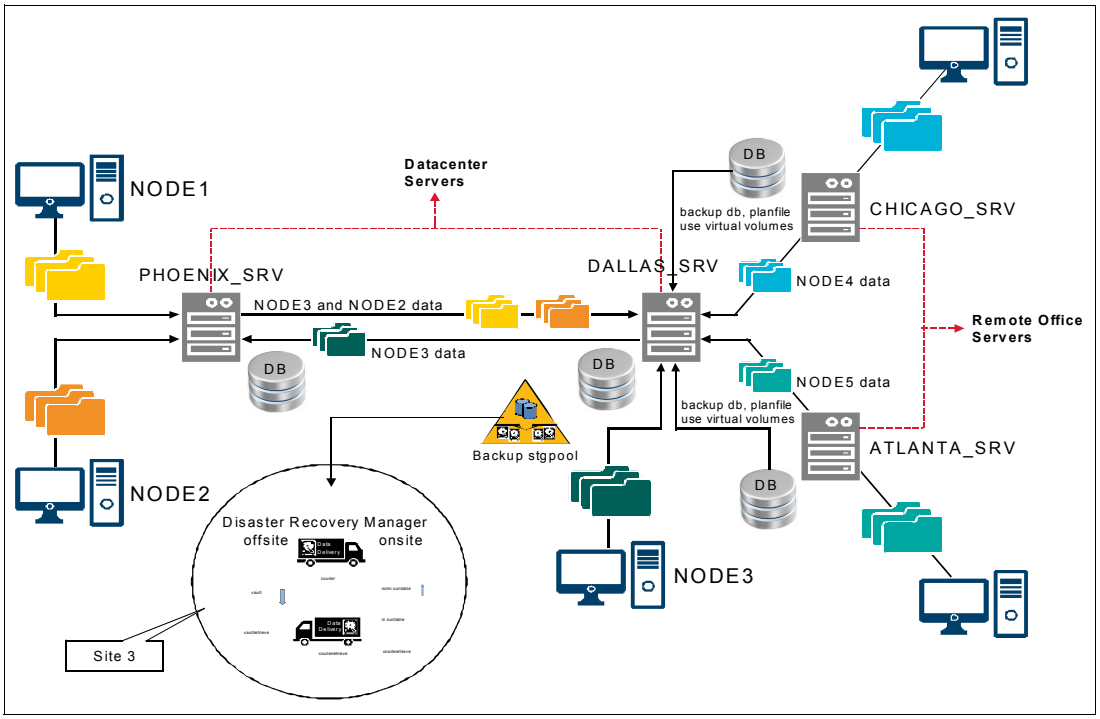

Figure 7-67 3-site data protection with node replication and disaster recovery manager

If a server outage occurs, you have different places, where the data is stored for recovery. If it is only a local incident, you can restore from the Tivoli Storage Manager server on the remote site. If the server on the remote site is unavailable, you can recover from the data center site, by using either the replicated data or the local backup set. If the data center is also affected, you must first recover the data center from the disaster recovery site. In this case, it is a multi-step procedure, which is well-planned, during the setup of the disaster protection procedures.

7.9 Employee workstation backup

In many companies, servers and big data bases must be protected, and valuable information is also stored on distributed systems directly used by the user. For IT, the relationship between private devices and business devices shrinks and the boundary is not often clear. If you are responsilbe for protecting such devices, doing your job in the expected manner can sometimes be a challenge.

This section shows solutions to successfully protect the data and its value.

The challenges are as follows:

- Low network bandwidth
- Short backup window
- Secure data store and transport
- Appropriate restore time
- Re-create from scratch, when the device is broken or stolen

In our solutions matrix, we put those challenges together. See Figure 7-69 on page 300 to see, the combination.

Challenges										
business					addressed with	technical				
reduce recovery time (RTO)	reduce backup time (meet RPO)	Compliance	Business Continuity	security (secure data transfer, secure data store,	Solution (discussed in this book)	Virtual Environments	Significant data growth	automation for business	Remote office and mobile office data	network bandwidth optimization
X	X				Disk to Disk backup	X		X	X	
X	X				Data reduction with TSM	X	X		X	X
X			X		Node replication				X	
X	X	X	X	X	TSM ToolKit	X	X	X	X	X
	X				progressive incremental forever	X			X	X
X	X				deduplication	X	X		X	X
	x				compression					X
	X				include / exclude	X	X			X
	X				journal based backup		X			X
X					point-in-time restore	X				
		X		X	data encryption					
					centralized monitoring and reporting			X	X	

Figure 7-68 Solution matrix for workstation backup

Client benefit

The solution offers these benefits:

- Secure data, secure data transfer, secure data access, secure data store
- Reduced backup time
- Fast recovery
- Remote office and mobile data management
- Reduced RPO, increased backup frequency
- Business continuity

Solution architecture

This solution describes how you can build a secure environment to protect your valuable business data, even when you are in a remote office or traveling to client but have low network bandwidth. We show available options to use the traditional way with a daily based backup or the more efficient version to protect your data continuously.

Figure 7-69 on page 300 shows the challenge and how to solve it. In this environment, you have remote locations and mobile devices, which must be protected by the centrally located Tivoli Storage Manager infrastructure.

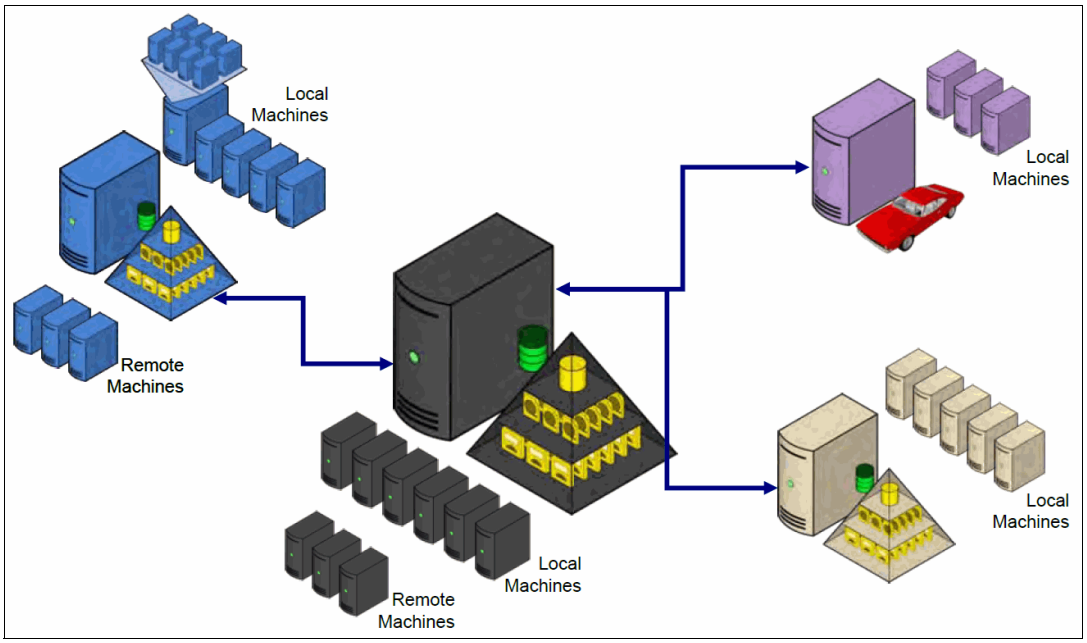

Figure 7-69 Overview of a data center structure with remote office and mobile clients

7.9.1 Tivoli Storage Manager Backup-Archive Client

Before discussing backup-archive client, consider several well-known terms.

Backup and archive

Backup means creating a copy of a data object to be used for recovery. This data object can be a file, a part of a file, a volume image, a directory, or a user-defined data object like a database table. The backup version of this data object is stored separately in the Tivoli Storage Manager server storage hierarchy. Tivoli Storage Manager policy tools allow great flexibility for the way data is managed for each client. Backup frequency, retention, migration and copy policies are easily implemented on the Tivoli Storage Manager client.

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

Therefore, the difference between backup and archive is that backup creates and controls multiple backup versions; archive creates an additional file that is retained for a specific period of time.

Security and encryption

Security is a vital aspect for enterprise storage management. Data must be protected, available, and secure with *encryption*. From the moment data is backed up from the client, IBM Tivoli Storage Manager provides a secure storage management environment.

Transport encryption

While sending data to the Tivoli Storage Manager Server, transport encryption is a key point and is necessary.

To improve the security of stored data, the backup-archive client implements an optional encryption function, which allows for encrypting 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 is encrypted and thus is unreadable by any malicious intruders. The user can choose which files are subject to encryption through include/exclude processing. The encryption uses a 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 backup and for archive.

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

To heighten security for Tivoli Storage Manager sessions, data sent to the Tivoli Storage Manager server during backup and archive operations can be encrypted with standard DES 56-bit encryption. The solution is backup-archive client, built-in and integrated with Tivoli Storage Manager server. Alternatively, you can set the ENCRYPTIONTYPE option to AES128 or DES56 in the client options file (`dsm.opt`).

Compression

Tivoli Storage Manager client *compression* option helps to instruct the Tivoli Storage Manager client to compress files before sending them to the Tivoli Storage Manager server. This industry-standard compression gives us effective bandwidth use.

Compressing files reduces the file data storage space and can improve throughput over slow networks with a powerful client. The throughput, however, can be degraded when running on a slow client system using a fast network, because software compression uses significant client CPU resources and costs additional elapsed time.

If the compression option is set to YES, the compression processing can be controlled in the following ways:

- Use the `include.compression` option to include files within a broad group of excluded files for compression processing.
- Use the `exclude.compression` option to exclude specific files or groups of files from compression processing, especially for the objects that are already compressed or encrypted, such as GIF, JPG, ZIP, MP3, and others.

For compression, use these suggestions:

- For fast network and fast server, set compression to NO.
- For LAN-free with tape, set compression to NO.
- For slow network or slow server, set compression to YES.
- Normally, set compressalways to YES.

If the Tivoli Storage Manager client compression and encryption is used for the same file during backup, the file is first compressed and then encrypted, which results in a smaller file. On restore, the file is decrypted first and then decompressed.

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 also with Tivoli Storage Manager family of products. These products work together with the IBM Tivoli Storage Manager server base product to ensure that any data stored is managed as defined.

The IBM Tivoli Storage Manager Backup-Archive Client, included with the server, provides the operational backup and archive functions. The client implements the patented progressive incremental backup methodology, deduplication, adaptive subfile 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.

Deduplication

Deduplication is also another bandwidth and data reduction technique. Tivoli Storage Manager helps to decrease the rate of amount of storage capacity required to contain data growth with a built-in data deduplication feature that helps eliminate redundant data. Deduplication eliminates redundant data chunks when client-side deduplication is used. This can enable significantly less data to be moved over the network, and significantly more backup data to be stored on disk.

Client-side data deduplication restrictions

Client-side data deduplication is mutually exclusive with several Tivoli Storage Manager Server features. When these features and client-side data deduplication are enabled, the features take precedence over client-side data deduplication.

- When LAN-free data movement operations are run, client-side data deduplication is ignored.
- Files that are included for encryption are excluded from data deduplication.
- Files that are in encrypted file systems are automatically excluded from data deduplication.
- To protect encrypted data that is inflight, use Secure Sockets Layer (SSL) with client-side data deduplication.
- When backing up subfiles, client-side data deduplication is ignored.
- A simultaneous write on the server takes precedence over client-side data deduplication.
- The destination storage pool must be of type FILE (sequential disk); the target storage pool must be a deduplication-enabled storage pool.
- The client and server must be at version 6.2.0 or later. Always use the most recent maintenance version.
- The client must have the client-side deduplication option enabled (`DEDUPLICATION YES`).
- The server must enable the node for client-side deduplication with the `DEDUP=CLIENTORSERVER` parameter using either the **`REGISTER NODE`** or **`UPDATE NODE`** commands.
- Files must not be excluded from client-side deduplication processing (by default all files are included).

- Files must be larger than 2 KB, and transactions must be below the value that is specified by the `clientdeduptxnlimit` option.
- The following Tivoli Storage Manager features are incompatible with Tivoli Storage Manager client-side deduplication:
  - Client encryption
  - LAN-free and storage agent
  - UNIX HSM client
  - Subfile backup
  - Simultaneous storage pool write

Figure 7-70 shows our solution approach to get benefits from Tivoli Storage Manager client deduplication.

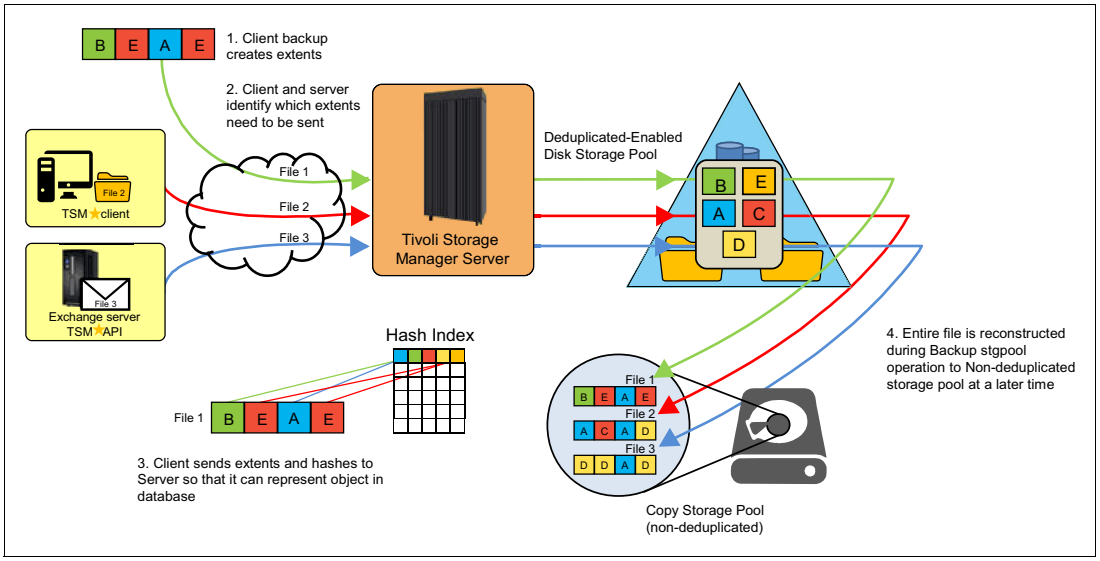

Figure 7-70 Client deduplication data flow

You can get more benefit from deduplication if you combine it with Tivoli Storage Manager Node replication for Disaster Recovery, which is discussed as another solution in this book.

Subfile backup

Mobile and remote computer branches have limited access to the infrastructure that serves the rest of the organization. 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 subfile backup feature which reduces the amount of data transferred while backing up changed files.

Figure 7-71 shows this feature. The backup-archive client (web client, CLI, or GUI) backs up only the changed portion of a file, either on a byte level or block level, instead of transferring the whole file to the server every time.

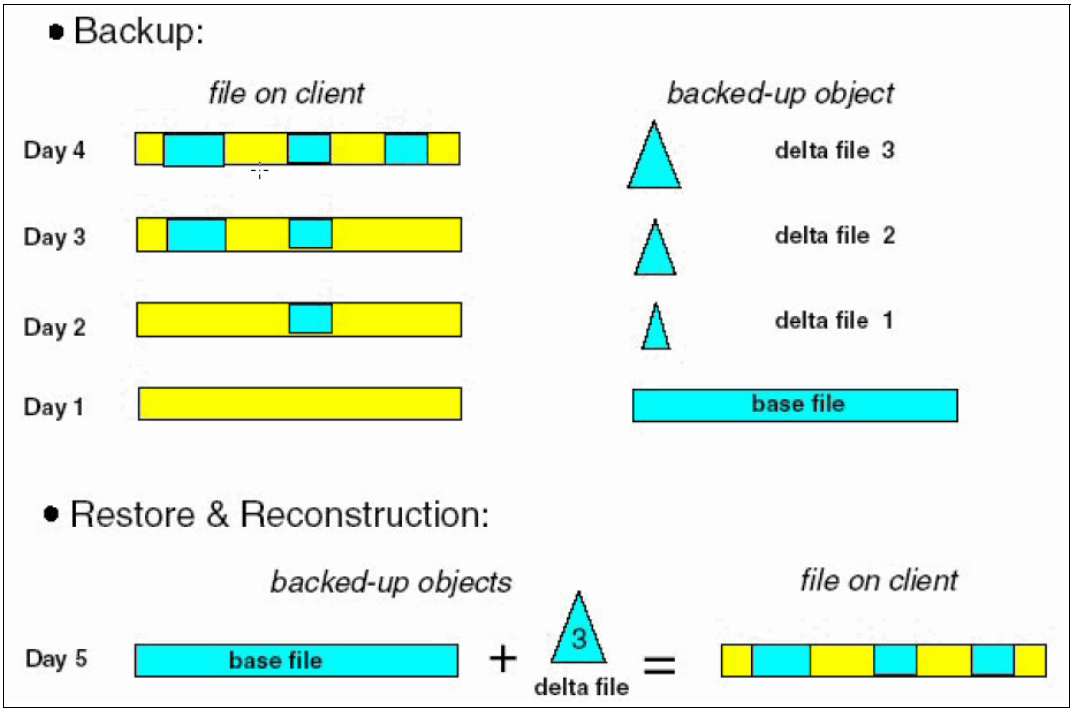

Figure 7-71 Adaptive subfile backup and restore

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 operation, this feature allows for restoring the whole file by restoring only two subfile components, one delta file, and the last complete backup of the whole file.

Bare machine recovery

Several solutions exist for recovering a system from scratch. But all of them should be fire-proof and heavily tested, because in a situation where you need to recover a system from scratch, you have mostly only one attempt.

For further information, and best practices for recovering Windows Server 2008, Windows Server 2008 R2, Windows 7, and Windows Vista, see the following document:

http://ibm.co/1ptFRi7

Another solution is the product from Cristie Software Limited, which also addresses this challenge.

TBMR is a software package from Cristie Software that offers these benefits:

- Leverages a customer's investment in Tivoli Storage Manager to provide a fully automated method of recovering or cloning a system that is running Microsoft Windows operating system, Linux, Oracle Solaris, and AIX to the same or new hardware.
- Complements your Tivoli Storage Manager backup and recovery strategy by offering a solution to recover, migrate, or clone your operating systems without the need to perform a separate backup. TBMR helps to maximize your recovery flexibility by supporting recovery to dissimilar hardware or virtual machines, which in turn minimizes business disruption.

- Provides fast automated recovery for Tivoli Storage Manager users. In the event of a machine failure, TBMR can recover the operating system and applications directly from the Tivoli Storage Manager backup. The recovery can be to the original or to new dissimilar hardware or to a virtual machine.
- Is unique in its ability to recover directly from Tivoli Storage Manager and without requiring any extra backup. This means that (unlike other BMR software products) TBMR consumes no extra storage or network bandwidth. TBMR also has the advantage of not requiring a separate backup application to be installed, managed and monitored.

Tivoli Storage Manager is an IBM enterprise class backup and archiving software. TBMR works exclusively with Tivoli Storage Manager and is by IBM and its channel partners worldwide as the preferred BMR solution for Tivoli Storage Manager.

Tivoli recovery is available for several operating systems:

- Windows
- Linux
- Solaris
- AIX

The main features of the BMR solution for Tivoli are as follows:

- Fully integrated with Tivoli Storage Manager
- Rebuild a server OS in under ten minutes
- Uses the Tivoli Storage Manager backups for disaster recovery
- Multiple servers can be recovered simultaneously
- TBMR is the *ONLY* software that can recover direct from the Tivoli Storage Manager server without doing a separate backup

7.9.2 Fastback for Workstations

IBM Tivoli Storage Manager FastBack for Workstations 7.1 is a file protection system for workstations and notebook computers. Your most important files can be continuously protected. Your less important files can be protected at scheduled intervals to save time and storage space. You can prevent any changes (including deletions) to files in folders that you designate as vaults.

Continuously protected files can be backed up to a local drive. This means that backup copies are created even when network conditions prevent storing backup copies on remote storage locations. Continuously protected files can also be stored on remote storage locations, when network connections allow. If a remote location is not available when you change a continuously protected file, the FastBack for Workstations client makes a backup copy on that device as soon as the device becomes available. Scheduled backup copies are created on the interval that you configure (hourly, weekly, daily, or monthly). If the remote device for scheduled backups is not available at the time of the backup, the FastBack for Workstations client makes backup copies on the remote location as soon as that device becomes available.

Figure 7-72 on page 306 shows a backup paradigm using a hybrid approach.

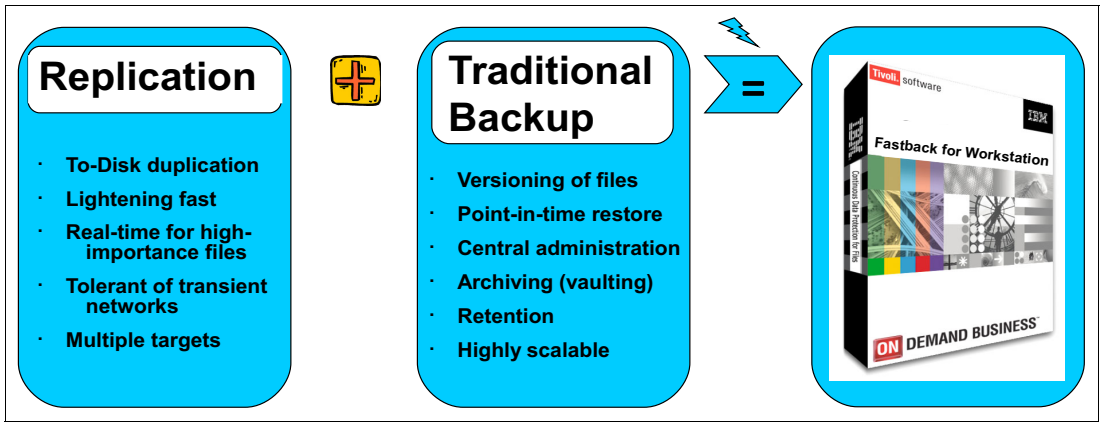

Figure 7-72 Flashback for Workstations backup paradigm using a hybrid approach

IBM Tivoli FastBack for Workstations is designed to comprehensively protect end users workstations. It hooks into the file system, so that new or changed files, applications, or directories are detected and are then copied to a local storage as shown in Figure 7-73. Files can also be replicated continuously or on a scheduled basis to a secondary media type or target. This can be a file-share, a target from a WebDev server, a Tivoli Storage Manager server, or a removable disk device (external hard drive, USB flash drive, and others). A user can maintain the configurations of the software or this can be controlled from a central management console. Although control of the configuration can be delegated to an administrator, users can still recover their own data. This can reduce the number of restoration requests for help desks and allow IT administrators to focus on other tasks.

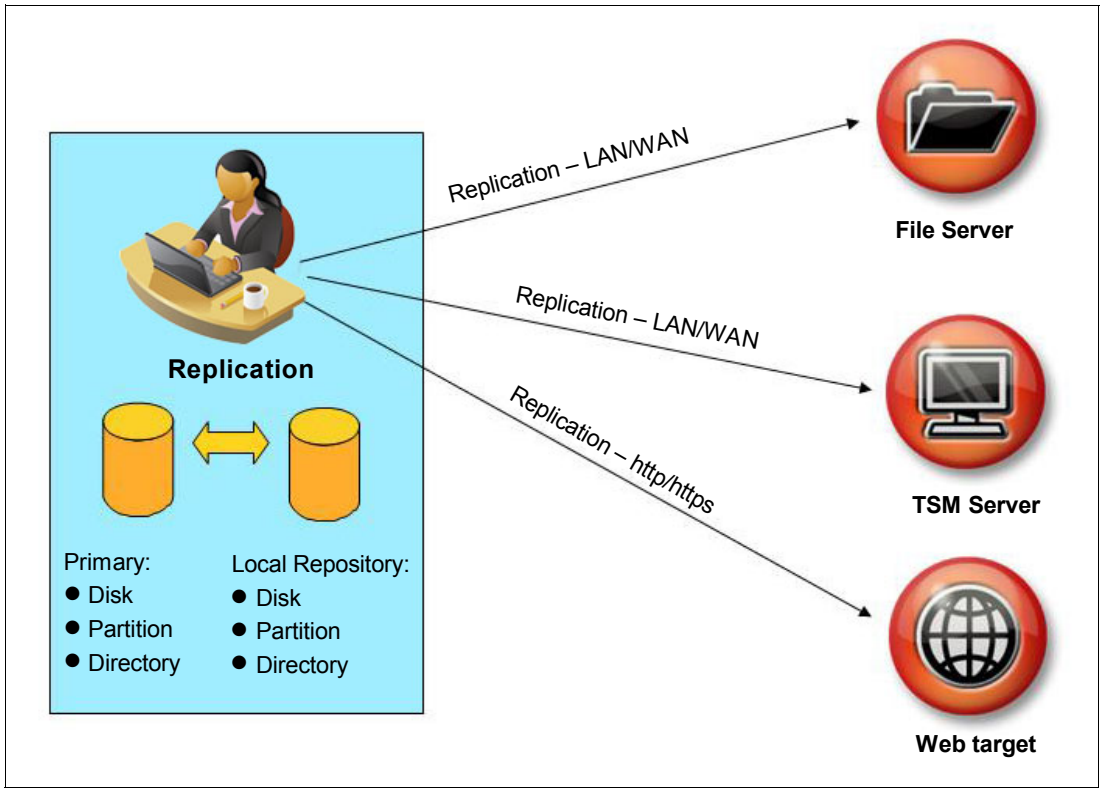

Figure 7-73 Fastback for Workstation standard configuration

Lightweight solution

Tivoli FastBack for Workstations is a lightweight data protection solution, combined with all of the premium data protection features. It offers these benefits:

- Easy to set up:
  - Select the files and mailboxes (Outlook and Lotus Notes) that you want to protect.
  - Select the number of local disk to allocate for backups.
  - Select the targets for off-machine copies of data.
- Easy to manage:

  Drop-down status windows for reports on last backup, capacity usage, network, or connectivity.
- Easy to restore:

  Simply select the files and directories you want to restore, and the point in time for which you want to restore.

Easy integration with Tivoli Storage Manager server

Tivoli Storage Manager FastBack for Workstations 7.1 can use Tivoli Storage Manager as a target system for disaster recovery purposes. When sending data to a Tivoli Storage Manager server, users can leverage some of the advanced data reduction and security capabilities that are built into the Tivoli Storage Manager clients and application programming interfaces (APIs) such as WAN deduplication, compression algorithms, and data encryption.

Deployments with Tivoli Storage Manager Server

If you are deploying clients to back up to a Tivoli Storage Manager Server, follow these steps. In this example, we use `tsmserver.ibm.com` as the Tivoli Storage Manager server and `client1` as the computer name for the client computer.

1. Register the client as a Tivoli Storage Manager node with delete backup authority. The node name can be the computer name of the client or a given name. If a given name is used, you must update `dsm.opt` to specify the node name option. Using the Tivoli Storage Manager Administrative command-line client, issue these commands:

   ```
   tsm: TSMSERVER>reg node client1 client1 backdel=y or some given name
   tsm: TSMSERVER>reg node jtn jtn backdel=y
   ```

2. Using the central administration console, create a group configuration called TSMGroup and set the remote storage to the Tivoli Storage Manager server. Make sure the Continuous Protection Level is set to `remote` storage only or `local and remote` storage. On the Remote Storage panel, select **Tivoli Storage Manager** from the Back up to menu and enter the following text in the Location field:

   ```
   tsm://tsmserver.ibm.com
   ```

3. Define an administration folder named `TSMAdminFolder` at the following address:

   ```
   \\AdminServer\TSMAdminFolder\RealTimeBackup
   ```

   Assign TSMGroup to this folder and then create the configuration file. You now have the file `fpa.txt` file.

4. If the computer name is used as the node name, you can skip this step. The installation creates the default `dsm.opt` for you. If node name is not the computer name, create the `dsm.opt` file with the following content:

   ```
   COMMMETHOD TCPIP PASSWORDACCESS generate NODENAME jtn
   ```

5. Copy `fpa.txt` and `dsm.opt` if needed to a temporary folder named `C:\DefaultConfigFiles` on the client computer.

6. Copy the FastBack for Workstations installation package to the client computer.
7. On the client computer, make sure you have connection to the AdminServer computer and the computer hosting tsmserver.
8. Open a command prompt and issue the following silent installation command. Make sure you open the command prompt with **Run** as administrator if you are running on Windows 7 or Windows Vista.

   ```
   6.3.0.0-TIV-FB4WKSTNS-x64_windows.exe /S /v"/qn /l*v c:\temp\msi.log
   REBOOT=ReallySuppress CUSTOM_CONFIG_FILES_PATH=C:\DefaultConfigFiles "
   ```

9. When the installation is finished, the client runs in the background. If you installed FastBack for Workstations 6.3 and later, the configuration files are backed up to the server.

Format of backup copies

Tivoli Storage Manager FastBack for Workstations keeps most backup copies in the same format as the original file. Tivoli Storage Manager FastBack for Workstations provides tools and views to see the backup copies and to restore them. However, in many cases it is not necessary to use Tivoli Storage Manager FastBack for Workstations to restore those backup copies. These files have content exactly like the originals, in a directory tree structure that simulates the original tree. Some backup copies are not in the same format as the original files, and must be restored using Tivoli Storage Manager FastBack for Workstations:

- Backup copies stored on Tivoli Storage Manager server
- Backup copies that were encrypted
- Backup copies that were compressed
- Large files that were backed up with subfile copy. In the storage area, the subfile copies have the `-FPdelta` file name suffix.
- Versions of bitmap backups. In the storage area, these backup copies have the `-TPdelta` file name suffix.

Bandwidth throttling

By specifying throttle settings and network rules for Tivoli Storage Manager FastBack for Workstations, you can manage the bandwidth usage in the networks that you specified.

Use the Network Rules settings to manage bandwidth usage for each network. When a network is accessed, Tivoli Storage Manager FastBack for Workstations uses the first rule in the list that matches the network. As a result, the throttle setting does not require a manual update every time Tivoli Storage Manager FastBack for Workstations accesses a different network. When a new network is detected, a default network rule is created. This default rule is added to the end of the network rule list.

This feature helps you control the data transfer to your remote system and not to overload the connection, which might be a tight WAN connection.

For information about configuring bandwidth throttling, see *IBM Tivoli Storage Manager FastBack for Workstations Version 6.3 Client Installation and User's Guide,* SC27-2809.

Central management

Key to providing a unified data protection infrastructure is the ability to centrally control and manage data protection operations. Tivoli Storage Manager FastBack for Workstations V6.3 provides a centralized management interface that can help administrators manage thousands of laptops and desktops. This central management interface runs in the same integrated solutions console in which the Tivoli Storage Manager Administration Center runs, enabling the management of data protection policies across the data center, remote office, and mobile

users in a centralized way. This centralized management capability allows various operations such as these:

- Automatic discovery of Tivoli Storage Manager FastBack for Workstations V6.3 clients
- Viewing various types of client information such as these items:
  - Deployment information (operating system version, Tivoli Storage Manager FastBack for Workstations version, and so on)
  - Amount of storage that a client is using
  - Client activity
  - Current client configurations for potential editing
  - Information about a client's storage target
  - Log files
  - Alerts
- Taking client actions, as in these examples:
  - Initiate an incremental backup.
  - Push a client configuration.
  - Lock a client configuration so it can not be changed.
  - Deploy updates.
  - Configure email alerts for administrators based on your own criteria.

Figure 7-74 shows how the central administration is implemented.

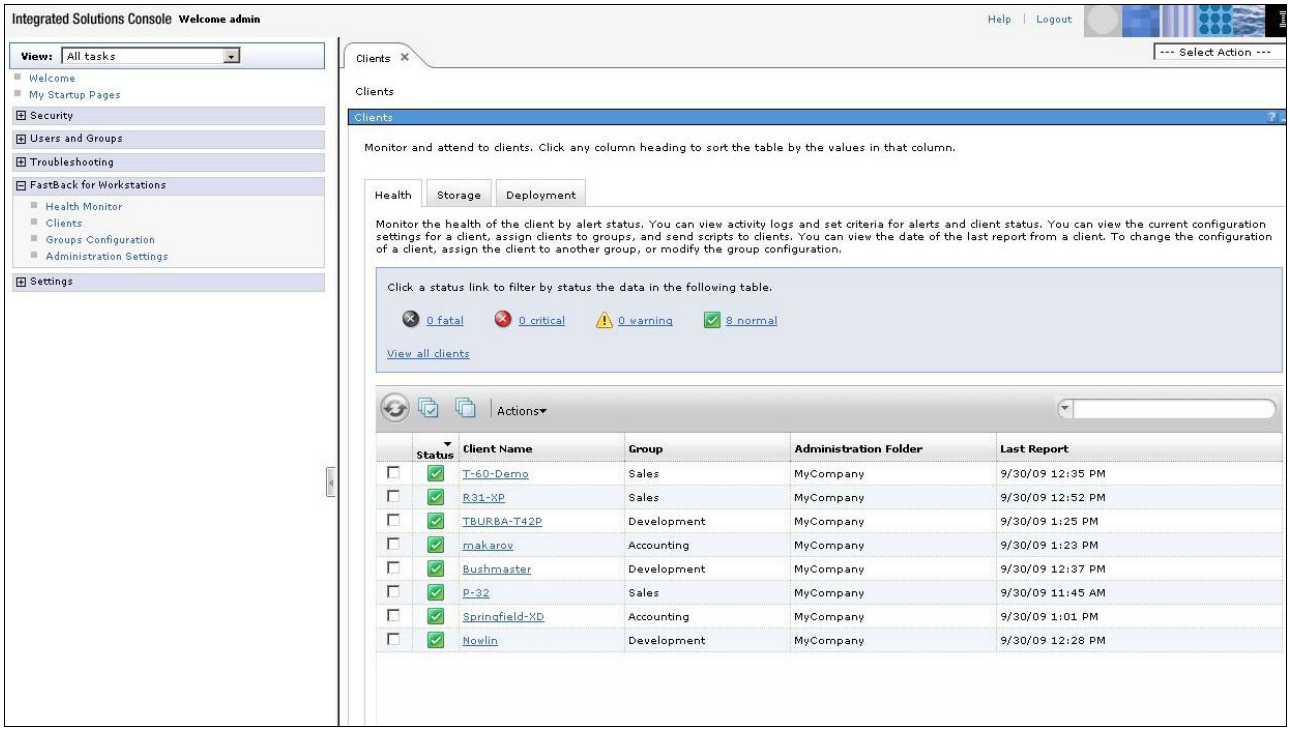

Figure 7-74 Central administration in one instance of the Tivoli Integrated Solutions Console

Tivoli Storage Manager FastBack for Workstations

Table 7-12 shows details of Tivoli Storage Manager FastBack for Workstation as a data protection solution in the context of Tivoli Storage Manager.

Table 7-12 Features, advantages, benefits

Features	Advantages	Benefits
Continuos data protection	Real-time data protection	Simplified storage management can save IT and user labor.
Restore to point-in-time	Multiple versions of the files retained	Continuous data protection provides data integrity when viruses and corruption attack systems.
Backups up on periodic basis to a remote file server	Faster backup	Reduces or eliminates backup windows and optimizes integration to network and enterprise.
Backs up only changed files	Less data is transferred	Optimizes bandwidth and network transfer of data.
Backs up your files the moment they change	Real-time data protection	Continuously protects versions of the files to allow customers choice of recovery points.
Backs up files to local cache	Stand-alone protection	Ability to write-protect data locally even when not connected in case of virus, corruptions, logical error or user error.
Remote and local disk support	Choice of backup devices	Ability to send data to mix of backup devices: disk, NAS, USB, local partition, LUN from SAN.

Additional information

For hints and tips, see *Deploying FastBack for Workstations In an Enterprise Environment*:

http://www-01.ibm.com/support/docview.wss?uid=swg21638113

A

Practical approach to creating a Tivoli Storage Manager solution

In this appendix, we draw a conclusion from the possible solutions we outline in the previous chapters and introduce you to the term *Tivoli Storage Manager policy-based dynamic tiering*.

We address this question:

How can the Tivoli Storage Manager product help you with your data protection challenges?

How to put solutions together

The volume, variety, and velocity of data is growing, and the value that the data represents to its owner is increasing. Consolidation of data centers to meet decreasing budgets for administration and operating require more flexibility for the infrastructure and entire data centers. The challenges for the responsible managers are the same worldwide.

To help you to meet your requirements for business-valuable processes and applications that are related to recovery time objective (RTO) and recovery point objective (RPO), Tivoli Storage Manager assists with intelligent storage architecture methods. Tivoli Storage Manager can act as an interface between intraday and interday protection, as shown in Figure A-1.

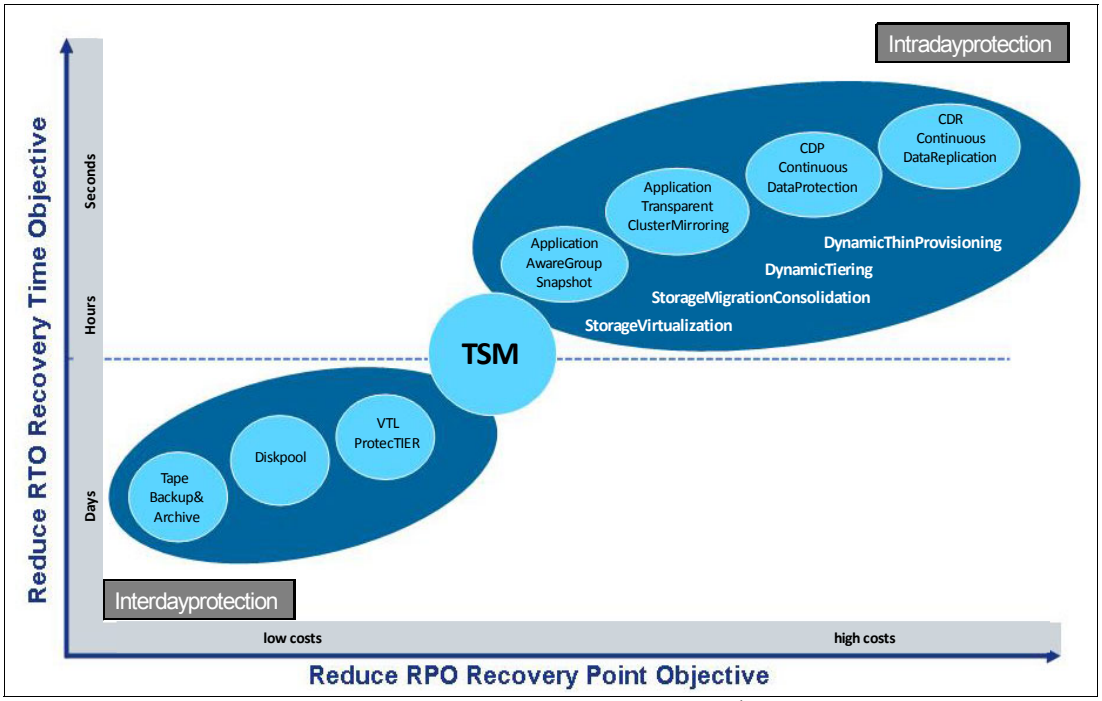

Figure A-1 Tivoli Storage Manager Intraday and Interday protection [1]

The Tivoli Storage Management layer interfaces between the application layer and the storage layer for the application and data-specific backup and recovery. You can take advantage of this and configure a disk and tape based solution dependent on the particular application.

Approaching the solution

How do you best approach a data protection solution using the Tivoli Storage Manager product? Of course there is more than one way to design the solution to the challenges you want to fulfill. The following example shows a possible way that might help you to design your exact custom data protection solution.

[1] Courtesy of Stephane Criachi, Concat AG

Start by collecting the data to identify your data protection task. We assume you try to consolidate three existing Tivoli Storage Manager servers. With the sample analysis table from Figure A-4 on page 317 the following data is collected for analysis:

- TSM SERVER

  The current Tivoli Storage Manager server in use
- NODE

  The node name for the node backing up to that server
- CLIENT VERSION

  The current client version for the node
- DOMAIN_NAME

  The domain the node is assigned to
- PLATFORM NAME

  The client node operating system platform
- SESSIONS

  Number of sessions that are configured
- DBSIZE GB

  Size (in gigabytes) of nodes database, for example, for Tivoli Data Protection applications
- max LOG /h GB

  How much database log amount to expect within an hour
- OCCSIZE GB

  Occupancy in gigabytes
- NUM OBJECTS

  Number of objects for that node
- AVG OBJSIZE GB

  The OCCSIZE/ NUM OBJECTS (occupancy size divided by the number of objects gives the average object size; it can be an indicator of whether the node can take advantage of LAN-free backups)
- MB PER SECOND LAN

  LAN backup throughput in megabytes
- GB LAN

  LAN backup volume in gigabytes
- MB PER SECOND LAN-free

  SAN backup throughput in megabytes
- GB LAN-free

  SAN backup volume in gigabytes
- AVG MB per second per session

  Average throughput in megabytes per second per session
- ACCUMULATED MB per second per session

  Overall session throughput in megabytes

- AVG GB DAILY

  Average daily backup volume in gigabytes

- MAX GB DAILY (*1,3)

  Add some buffer to the maximum amount (in gigabytes) of data to be transferred

Gather the information in a spreadsheet using information from the activity log, the summary table and query commands (see Figure A-4 on page 317).

To collect the data use commands similar to Example A-1.

Example A-1 Sample select statement to gather data for analysis

```
select * from summary where activity='BACKUP'
```

The output is presented in a comma separated format, where it can be used for further calculations. Example A-2 shows output.

Example A-2 Comma separated output from previous select statement

```
2013-09-18 00:00:27.000000,2013-09-18
00:13:06.000000,BACKUP,47470,SRV6015,Tcp/Ip,SRV6015:32780,,67899,125,0,990431,721,0,2,YES,,,,,0,
2013-09-18 00:00:29.000000,2013-09-18
00:13:02.000000,BACKUP,47477,SRV6014,Tcp/Ip,SRV6014:34190,,67434,54,0,810086,697,0,2,YES,,,,,0,
2013-09-18 00:00:30.000000,2013-09-18
00:12:35.000000,BACKUP,47482,SRV6012,Tcp/Ip,SRV6012:33301,,69320,48,0,6066848,682,0,2,YES,,,,,2,
2013-09-18 00:00:30.000000,2013-09-18
00:13:15.000000,BACKUP,47481,SRV6013,Tcp/Ip,SRV6013:32979,,67519,55,0,5191665,715,0,2,YES,,,,,1,
2013-09-18 00:00:31.000000,2013-09-18
00:13:00.000000,BACKUP,47487,SRV6011,Tcp/Ip,SRV6011:33519,,69904,49,0,702472,725,0,2,YES,,,,,0,
2013-09-18 00:00:33.000000,2013-09-18
00:12:45.000000,BACKUP,47494,SRV6009,Tcp/Ip,SRV6009:33054,,70335,59,0,822834,703,0,2,YES,,,,,0,
2013-09-18 00:00:34.000000,2013-09-18
00:00:58.000000,BACKUP,47500,SRV6004,Tcp/Ip,SRV6004:42442,,5027,16,0,1388107,0,0,2,YES,,,,,0,
2013-09-18 00:00:34.000000,2013-09-18
00:12:31.000000,BACKUP,47499,SRV6008,Tcp/Ip,SRV6008:32962,,68954,63,0,2704098,692,0,2,YES,,,,,3,
2013-09-18 00:00:36.000000,2013-09-18
00:05:58.000000,BACKUP,47505,SRV6003,Tcp/Ip,SRV6003:63456,,229517,147,0,109754055,243,0,2,YES,,,,,1,
2013-09-18 00:00:37.000000,2013-09-18
00:14:21.000000,BACKUP,47507,SRV6005,Tcp/Ip,SRV6005:35012,,99919,99,1,32585270,760,0,2,YES,,,,,13,
2013-09-18 00:00:38.000000,2013-09-18
00:04:54.000000,BACKUP,47508,SRV6002,Tcp/Ip,SRV6002:56424,,127134,299,0,27178390,222,0,2,YES,,,,,4,
2013-09-18 00:00:38.000000,2013-09-18
00:52:26.000000,BACKUP,47509,SRV6001,Tcp/Ip,SRV6001:32801,,3028255,833,0,1980221601,2998,0,2,YES,,,,,34,
2013-09-18 12:00:05.000000,2013-09-18
12:12:40.000000,BACKUP,47913,SRV6016,Tcp/Ip,SRV6016:32851,,70081,51,0,2235003,732,1,2,YES,,,,,0,
2013-09-18 12:00:06.000000,2013-09-18
12:12:39.000000,BACKUP,47919,SRV6007,Tcp/Ip,SRV6007:33171,,69192,86,0,4834289,718,0,2,YES,,,,,2,
2013-09-18 12:00:09.000000,2013-09-18
12:12:54.000000,BACKUP,47927,SRV6006,Tcp/Ip,SRV6006:32931,,68406,55,0,6707630,708,0,2,YES,,,,,1,
2013-09-18 12:00:10.000000,2013-09-18
12:15:47.000000,BACKUP,47932,SRV6017,Tcp/Ip,SRV6017:33069,,94296,44,0,5761771,889,0,2,YES,,,,,2,
2013-09-18 12:00:12.000000,2013-09-18
12:13:30.000000,BACKUP,47940,SRV6030,Tcp/Ip,SRV6030:32822,,66112,49,0,873626,752,0,2,YES,,,,,0,
2013-09-18 12:00:12.000000,2013-09-18
12:13:57.000000,BACKUP,47941,SRV6015,Tcp/Ip,SRV6015:32791,,67899,50,0,1622416,758,0,2,YES,,,,,1,
2013-09-18 12:00:12.000000,2013-09-18
12:15:37.000000,BACKUP,47937,SRV6010,Tcp/Ip,SRV6010:33359,,100677,86,0,3123415,893,0,2,YES,,,,,4,
2013-09-18 12:00:15.000000,2013-09-18
12:13:32.000000,BACKUP,47952,SRV6013,Tcp/Ip,SRV6013:33107,,67515,79,0,4185352,725,1,2,YES,,,,,2,
2013-09-18 12:00:16.000000,2013-09-18
12:13:07.000000,BACKUP,47955,SRV6012,Tcp/Ip,SRV6012:33350,,69328,93,1,12938566,697,0,2,YES,,,,,3,
2013-09-18 12:00:17.000000,2013-09-18
12:13:31.000000,BACKUP,47962,SRV6009,Tcp/Ip,SRV6009:33065,,70334,52,0,778410,768,0,2,YES,,,,,0,
2013-09-18 12:00:17.000000,2013-09-18
12:13:43.000000,BACKUP,47959,SRV6014,Tcp/Ip,SRV6014:34457,,67433,79,0,2895454,736,0,2,YES,,,,,3,
2013-09-18 12:00:17.000000,2013-09-18
12:13:52.000000,BACKUP,47960,SRV6011,Tcp/Ip,SRV6011:33783,,69902,83,0,6325644,787,1,2,YES,,,,,5,
2013-09-18 12:00:19.000000,2013-09-18
12:01:07.000000,BACKUP,47970,SRV6004,Tcp/Ip,SRV6004:43046,,5026,9,0,703355,0,0,2,YES,,,,,0,
2013-09-18 12:00:20.000000,2013-09-18
12:13:32.000000,BACKUP,47972,SRV6008,Tcp/Ip,SRV6008:33088,,68953,125,0,4485908,762,0,2,YES,,,,,6,
2013-09-18 12:00:21.000000,2013-09-18
12:05:44.000000,BACKUP,47975,SRV6003,Tcp/Ip,SRV6003:56633,,229517,36,0,80935682,259,0,2,YES,,,,,0,
```

```
2013-09-18 12:00:22.000000,2013-09-18
12:15:55.000000,BACKUP,47978,SRV6005,Tcp/Ip,SRV6005:35641,,99924,121,0,63463354,847,1,2,YES,,,,22,
2013-09-18 12:00:23.000000,2013-09-18
12:05:00.000000,BACKUP,47980,SRV6002,Tcp/Ip,SRV6002:58209,,127138,279,0,48835219,242,0,2,YES,,,,4,
2013-09-18 12:00:23.000000,2013-09-18
12:51:32.000000,BACKUP,47981,SRV6001,Tcp/Ip,SRV6001:33846,,3028866,967,0,818242468,2987,0,2,YES,,,,13,
2013-09-18 12:00:33.000000,2013-09-18
12:54:13.000000,BACKUP,47988,SRVDM02,Tcp/Ip,192.168.111..146:65529,,819033,195,0,549029732,318,0,2,YES,,,,9,
2013-09-18 19:28:29.000000,2013-09-18
```

You can also create a batch query from an existing Tivoli Storage Manager instance, as shown in Example A-3.

Example A-3 Batch query to gather data

```
dsmadmc -userid=<admin> -pass=<password> -outfile=<filename> -commadelimited query node f=d
```

The data can now be imported in a spreadsheet program for analysis as shown in Figure A-2.

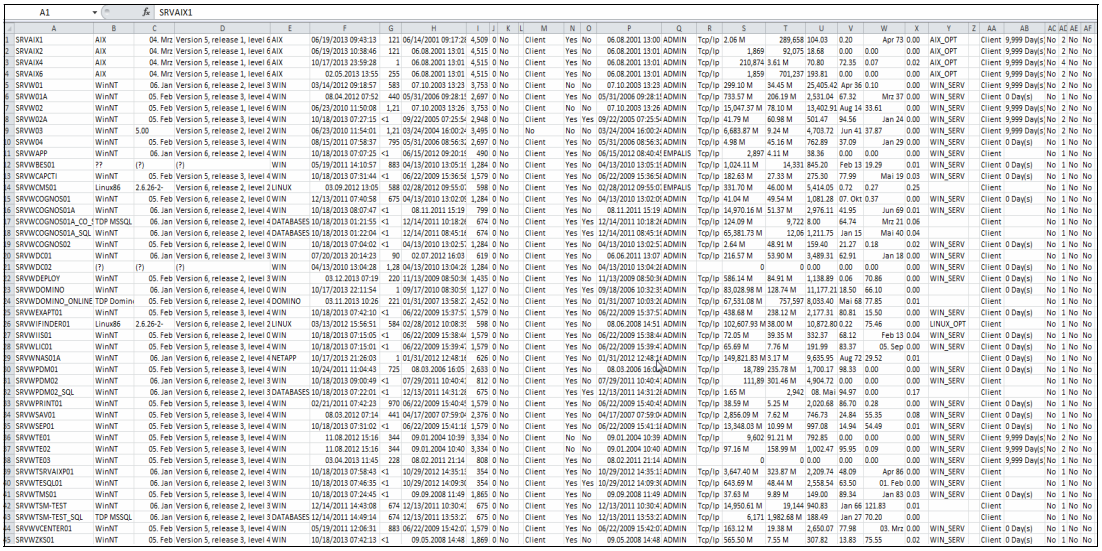

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	AA	AB	AC	AD	AE	AF
1	SRVAIX1	AIX	04. Mrz	Version 5, release 1, level 6	AIX	06/19/2013 09:43:13	121	06/14/2001 09:17:2	4,509	0	No		Client	Yes	No	06.08.2001 13:00	ADMIN	Tcp/Ip	2.06 M	289,658	104.03	0.20	Apr 73	0.00	AIX_OPT		Client	9,999 Day(s)	No	2	No	No
2	SRVAIX2	AIX	04. Mrz	Version 5, release 1, level 6	AIX	06/19/2013 10:38:46	121	06.08.2001 13:01	4,515	0	No		Client	Yes	No	06.08.2001 13:01	ADMIN	Tcp/Ip	1,869	92,075	18.68	0.00	0.00	0.00	AIX_OPT		Client	9,999 Day(s)	No	2	No	No
3	SRVAIX4	AIX	04. Mrz	Version 5, release 1, level 6	AIX	10/17/2013 23:59:28	1	06.08.2001 13:01	4,515	0	No		Client	Yes	No	06.08.2001 13:01	ADMIN	Tcp/Ip	210,874	3.61 M	70.80	72.35	0.07	0.02	AIX_OPT		Client	9,999 Day(s)	No	4	No	No
4	SRVAIX6	AIX	04. Mrz	Version 5, release 1, level 6	AIX	02.05.2013 13:55	255	06.08.2001 13:01	4,515	0	No		Client	Yes	No	06.08.2001 13:01	ADMIN	Tcp/Ip	1,859	701,237	193.81	0.00	0.00	0.00	AIX_OPT		Client	9,999 Day(s)	No	2	No	No
5	SRVW01	WinNT	06. Jan	Version 6, release 2, level 3	WIN	03/14/2012 09:18:57	583	07.10.2003 13:23	3,753	0	No		Client	No	No	07.10.2003 13:23	ADMIN	Tcp/Ip	299.10 M	34.45 M	25,405.42	Apr 36	0.10	0.00	WIN_SERV		Client	9,999 Day(s)	No	2	No	No
6	SRVW01A	WinNT	05. Feb	Version 5, release 3, level 4	WIN	08.04.2012 07:52	440	05/31/2006 09:28:1	2,697	0	No		Client	Yes	No	05/31/2006 09:28:1	ADMIN	Tcp/Ip	733.57 M	206.19 M	2,531.04	47.32	Mrz 37	0.00	WIN_SERV		Client	0 Day(s)	No	1	No	No
7	SRVW02	WinNT	05. Feb	Version 5, release 1, level 6	WIN	06/23/2010 11:50:08	1,21	07.10.2003 13:26	3,753	0	No		Client	No	No	07.10.2003 13:26	ADMIN	Tcp/Ip	15,047.37 M	78.10 M	13,402.91	Aug 14	33.61	0.00	WIN_SERV		Client	9,999 Day(s)	No	2	No	No
8	SRVW02A	WinNT	05. Feb	Version 5, release 3, level 4	WIN	10/18/2013 07:27:15	<1	09/22/2005 07:25:5	2,948	0	No		Client	Yes	Yes	09/22/2005 07:25:5	ADMIN	Tcp/Ip	41.79 M	60.98 M	501.47	94.56	Jan 24	0.00	WIN_SERV		Client	9,999 Day(s)	No	2	No	No
9	SRVW03	WinNT	5.00	Version 5, release 2, level 2	WIN	06/23/2010 11:54:01	1,21	03/24/2004 16:00:2	3,495	0	No		No	No	No	03/24/2004 16:00:2	ADMIN	Tcp/Ip	6,683.87 M	9.24 M	4,703.72	Jun 41	37.87	0.00	WIN_SERV		Client	9,999 Day(s)	No	2	No	No
10	SRVW04	WinNT	05. Feb	Version 5, release 3, level 4	WIN	08/15/2011 07:58:37	795	05/31/2006 08:56:3	2,697	0	No		Client	Yes	No	05/31/2006 08:56:3	ADMIN	Tcp/Ip	4.98 M	45.16 M	762.89	37.09	Jan 29	0.00	WIN_SERV		Client	0 Day(s)	No	1	No	No
11	SRVWAPP	WinNT	06. Jan	Version 6, release 2, level 4	WIN	10/18/2013 07:07:25	<1	06/15/2012 09:20:1	490	0	No		Client	Yes	No	06/15/2012 08:40:4	EMPALIS	Tcp/Ip	2,897	4.11 M	38.36	0.00	0.00	0.00	WIN_SERV		Client		No	1	No	No
12	SRVWBE501	??	(?)	(?)	WIN	05/19/2011 14:10:57	883	04/13/2010 13:05:1	1,284	0	No		Client	Yes	No	04/13/2010 13:05:1	ADMIN	Tcp/Ip	1,024.11 M	14,331	845.20	Feb 13	19.29	0.01	WIN_SERV		Client	0 Day(s)	No	1	No	No
13	SRVWCAPCTI	WinNT	05. Feb	Version 5, release 3, level 4	WIN	10/18/2013 07:51:44	<1	06/22/2009 15:36:5	1,579	0	No		Client	Yes	No	06/22/2009 15:36:5	ADMIN	Tcp/Ip	182.63 M	27.33 M	275.30	77.99	Mai 19	0.03	WIN_SERV		Client	0 Day(s)	No	1	No	No
14	SRVWCM501	Linux86	2.6.26-2-	Version 6, release 2, level 2	LINUX	03.09.2012 13:05	588	02/28/2012 09:55:0	598	0	No		Client	Yes	No	02/28/2012 09:55:0	EMPALIS	Tcp/Ip	331.70 M	46.00 M	5,414.05	0.72	0.27	0.25			Client		No	1	No	No
15	SRVWCOGNOS01	WinNT	05. Feb	Version 5, release 1, level 6	WIN	12/13/2011 07:40:58	675	04/13/2010 13:02:0	1,284	0	No		Client	Yes	No	04/13/2010 13:02:0	ADMIN	Tcp/Ip	41.04 M	49.54 M	1,081.28	07. Okt	0.37	0.00	WIN_SERV		Client	0 Day(s)	No	1	No	No
16	SRVWCOGNOS01A	WinNT	06. Jan	Version 6, release 2, level 4	WIN	10/18/2013 08:07:47	<1	08.11.2011 15:19	799	0	No		Client	Yes	No	08.11.2011 15:19	ADMIN	Tcp/Ip	14,970.16 M	51.37 M	2,976.11	41.95	Jun 69	0.01	WIN_SERV		Client		No	1	No	No
17	SRVWCOGNOS01A_CO_T	TDP MSSQL	06. Jan	Version 6, release 2, level 4	DATABASES	10/18/2013 01:21:55	<1	12/14/2011 10:18:2	674	0	No		Client	Yes	Yes	12/14/2011 10:18:2	ADMIN	Tcp/Ip	124.09 M	9,722	8.00	64.74	Mrz 21	0.06			Client		No	1	No	No
18	SRVWCOGNOS01A_SQL	WinNT	06. Jan	Version 6, release 2, level 4	DATABASES	10/18/2013 01:22:04	<1	12/14/2011 08:45:1	674	0	No		Client	Yes	Yes	12/14/2011 08:45:1	ADMIN	Tcp/Ip	65,381.73 M	12,06	1,211.75	Jan 15	Mai 40	0.04			Client		No	1	No	No
19	SRVWCOGNOS02	WinNT	05. Feb	Version 6, release 2, level 0	WIN	10/18/2013 07:04:02	<1	04/13/2010 13:02:5	1,284	0	No		Client	Yes	No	04/13/2010 13:02:5	ADMIN	Tcp/Ip	2.64 M	48.91 M	159.40	21.27	0.18	0.02	WIN_SERV		Client	0 Day(s)	No	1	No	No
20	SRVWDC01	WinNT	06. Jan	Version 6, release 2, level 3	WIN	07/20/2013 20:14:23	90	02.07.2012 16:03	619	0	No		Client	Yes	No	06.06.2011 13:07	ADMIN	Tcp/Ip	216.57 M	53.90 M	3,489.31	62.91	Jan 18	0.00	WIN_SERV		Client		No	1	No	No
21	SRVWDC02	(?)	(?)	(?)	WIN	04/13/2010 13:04:28	1,28	04/13/2010 13:04:2	1,284	0	No		Client	Yes	No	04/13/2010 13:04:2	ADMIN			0	0.00	0.00	0.00	0.00	WIN_SERV		Client	0 Day(s)	No	1	No	No
22	SRVWDEPLOY	WinNT	05. Feb	Version 6, release 2, level 3	WIN	03.12.2013 07:19	220	11/13/2009 08:50:3	1,435	0	No		Client	Yes	No	11/13/2009 08:50:3	ADMIN	Tcp/Ip	586.14 M	84.91 M	1,138.89	0.06	70.86	0.00	WIN_SERV		Client	0 Day(s)	No	1	No	No
23	SRVWDOMINO	WinNT	06. Jan	Version 6, release 4, level 0	WIN	10/17/2013 22:11:54	1	09/17/2010 08:30:5	1,127	0	No		Client	Yes	Yes	09/18/2006 10:32:3	ADMIN	Tcp/Ip	83,028.98 M	128.74 M	11,177.21	18.50	66.10	0.00			Client	0 Day(s)	No	1	No	No
24	SRVWDOMINO_ONLINE	TDP Domino	05. Feb	Version 6, release 2, level 4	DOMINO	03.11.2013 10:26	221	01/31/2007 13:58:2	2,452	0	No		Client	Yes	No	01/31/2007 10:03:2	ADMIN	Tcp/Ip	67,531.08 M	757,597	8,033.40	Mai 68	77.85	0.01			Client		No	1	No	No
25	SRVWEXAPT01	WinNT	05. Feb	Version 5, release 3, level 4	WIN	10/18/2013 07:42:10	<1	06/22/2009 15:37:5	1,579	0	No		Client	Yes	No	06/22/2009 15:37:5	ADMIN	Tcp/Ip	438.68 M	238.12 M	2,177.31	80.81	15.50	0.00	WIN_SERV		Client	0 Day(s)	No	1	No	No
26	SRVWIFINDER01	Linux86	2.6.26-2-	Version 6, release 2, level 2	LINUX	03/13/2012 15:56:51	584	02/28/2012 10:08:3	598	0	No		Client	Yes	No	08.06.2008 14:51	ADMIN	Tcp/Ip	102,607.93 M	38.00 M	10,872.80	0.22	75.46	0.00	LINUX_OPT		Client		No	1	No	No
27	SRVWIIS01	WinNT	05. Feb	Version 6, release 2, level 0	WIN	10/18/2013 07:15:05	<1	06/22/2009 15:38:4	1,579	0	No		Client	Yes	No	06/22/2009 15:38:4	ADMIN	Tcp/Ip	72.05 M	39.35 M	332.37	68.12	Feb 13	0.04	WIN_SERV		Client	0 Day(s)	No	1	No	No
28	SRVWLIC01	WinNT	05. Feb	Version 5, release 3, level 4	WIN	10/18/2013 07:15:01	<1	06/22/2009 15:39:4	1,579	0	No		Client	Yes	No	06/22/2009 15:39:4	ADMIN	Tcp/Ip	65.69 M	7.76 M	191.99	83.37	05. Sep	0.00	WIN_SERV		Client	0 Day(s)	No	1	No	No
29	SRVWNAS01A	WinNT	06. Jan	Version 6, release 2, level 4	NETAPP	10/17/2013 21:26:03	1	01/31/2012 12:48:1	626	0	No		Client	Yes	No	01/31/2012 12:48:1	ADMIN	Tcp/Ip	149,821.83 M	3.17 M	9,635.95	Aug 72	29.52	0.01			Client	0 Day(s)	No	1	No	No
30	SRVWPDM01	WinNT	05. Feb	Version 5, release 3, level 4	WIN	10/24/2011 11:04:43	725	08.03.2006 16:05	2,633	0	No		Client	Yes	No	08.03.2006 16:05	ADMIN	Tcp/Ip	18,789	235.78 M	1,700.17	98.33	0.00	0.00	WIN_SERV		Client	0 Day(s)	No	1	No	No
31	SRVWPDM02	WinNT	06. Jan	Version 6, release 2, level 3	WIN	10/18/2013 09:00:49	<1	07/29/2011 10:40:4	812	0	No		Client	Yes	No	07/29/2011 10:40:4	ADMIN	Tcp/Ip	111,89	301.46 M	4,904.72	0.00	0.00	0.00	WIN_SERV		Client		No	1	No	No
32	SRVWPDM02_SQL	WinNT	06. Jan	Version 6, release 2, level 3	DATABASES	10/18/2013 07:22:01	<1	12/13/2011 14:31:2	675	0	No		Client	Yes	Yes	12/13/2011 14:31:2	ADMIN	Tcp/Ip	1.65 M	2,942	08. Mai	94.97	0.00	0.17			Client		No	1	No	No
33	SRVWPRINT01	WinNT	05. Feb	Version 5, release 3, level 4	WIN	02/21/2011 07:42:23	970	06/22/2009 15:40:4	1,579	0	No		Client	Yes	No	06/22/2009 15:40:4	ADMIN	Tcp/Ip	38.59 M	5.25 M	2,020.68	86.70	0.28	0.00	WIN_SERV		Client	0 Day(s)	No	1	No	No
34	SRVWSAV01	WinNT	05. Feb	Version 5, release 3, level 4	WIN	08.03.2012 07:14	441	04/17/2007 07:59:0	2,376	0	No		Client	Yes	No	04/17/2007 07:59:0	ADMIN	Tcp/Ip	2,856.09 M	7.62 M	746.73	24.84	55.35	0.08	WIN_SERV		Client	0 Day(s)	No	1	No	No
35	SRVWSEP01	WinNT	05. Feb	Version 5, release 3, level 4	WIN	10/18/2013 07:31:02	<1	06/22/2009 15:41:1	1,579	0	No		Client	Yes	No	06/22/2009 15:41:1	ADMIN	Tcp/Ip	13,348.03 M	10.99 M	997.08	14.94	54.49	0.01	WIN_SERV		Client	0 Day(s)	No	1	No	No
36	SRVWTE01	WinNT	05. Feb	Version 5, release 3, level 4	WIN	11.08.2012 15:16	344	09.01.2004 10:39	3,334	0	No		Client	No	No	09.01.2004 10:39	ADMIN	Tcp/Ip	9,602	91.21 M	792.85	0.00	0.00	0.00	WIN_SERV		Client	9,999 Day(s)	No	2	No	No
37	SRVWTE02	WinNT	05. Feb	Version 5, release 3, level 4	WIN	11.08.2012 15:16	344	09.01.2004 10:40	3,334	0	No		Client	No	No	09.01.2004 10:40	ADMIN	Tcp/Ip	97.16 M	158.99 M	1,002.47	95.95	0.09	0.00	WIN_SERV		Client	9,999 Day(s)	No	2	No	No
38	SRVWTE03	WinNT	05. Feb	Version 5, release 3, level 4	WIN	03.04.2013 11:45	228	08.02.2011 21:14	808	0	No		Client	Yes	No	08.02.2011 21:14	ADMIN			0	0.00	0.00	0.00	0.00	WIN_SERV		Client	9,999 Day(s)	No	1	No	No
39	SRVWTSRVAIXP01	WinNT	06. Jan	Version 6, release 2, level 4	WIN	10/18/2013 07:58:43	<1	10/29/2012 14:35:1	354	0	No		Client	Yes	No	10/29/2012 14:35:1	ADMIN	Tcp/Ip	3,647.40 M	323.87 M	2,209.74	48.09	Apr 86	0.00	WIN_SERV		Client		No	1	No	No
40	SRVWTESQL01	WinNT	06. Jan	Version 6, release 2, level 3	WIN	10/18/2013 07:46:35	<1	10/29/2012 14:09:3	354	0	No		Client	Yes	Yes	10/29/2012 14:09:3	ADMIN	Tcp/Ip	643.69 M	48.44 M	2,558.54	63.50	01. Feb	0.00	WIN_SERV		Client		No	1	No	No
41	SRVWTMS01	WinNT	05. Feb	Version 5, release 3, level 4	WIN	10/18/2013 07:24:45	<1	09.09.2008 11:49	1,865	0	No		Client	Yes	No	09.09.2008 11:49	ADMIN	Tcp/Ip	37.63 M	9.89 M	149.00	89.34	Jan 83	0.03	WIN_SERV		Client	0 Day(s)	No	1	No	No
42	SRVWTSM-TEST	WinNT	06. Jan	Version 6, release 2, level 3	WIN	12/14/2011 14:43:08	674	12/13/2011 10:30:4	675	0	No		Client	Yes	No	12/13/2011 10:30:4	ADMIN	Tcp/Ip	14,950.61 M	19,144	940.83	Jan 66	121.83	0.01			Client		No	1	No	No
43	SRVWTSM-TEST_SQL	TDP MSSQL	06. Jan	Version 6, release 2, level 3	DATABASES	12/14/2011 14:49:14	674	12/13/2011 13:53:2	675	0	No		Client	Yes	No	12/13/2011 13:53:2	ADMIN	Tcp/Ip	6,171	1,982.68 M	188.49	Jan 27	70.20	0.00			Client		No	1	No	No
44	SRVWVCENTER01	WinNT	05. Feb	Version 5, release 3, level 4	WIN	05/19/2011 12:06:31	883	06/22/2009 15:42:0	1,579	0	No		Client	Yes	No	06/22/2009 15:42:0	ADMIN	Tcp/Ip	163.12 M	19.38 M	2,650.07	77.98	03. Mrz	0.00	WIN_SERV		Client	0 Day(s)	No	1	No	No
45	SRVWZKS01	WinNT	05. Feb	Version 5, release 3, level 4	WIN	10/18/2013 07:42:13	<1	09.05.2008 14:48	1,869	0	No		Client	Yes	No	09.05.2008 14:48	ADMIN	Tcp/Ip	565.50 M	7.55 M	307.82	13.83	75.55	0.02	WIN_SERV		Client	0 Day(s)	No	1	No	No

Figure A-2 sample spreadsheet for analysis of the output from query node

Valuable information is also in the occupancy table of the Tivoli Storage Manager Server database. To query this information, use the administrative **query** command (Example A-4).

Example A-4 Sample query command to get information about the occupancy for backup and archive

```
query auditocc
```

The tab-separated output is shown in Figure A-3.

```
SRV6002 13,272  0       0       13,272
SRV6001 990,979 0       0       990,979
SRV6003 208,867 0       0       208,867
SRV6004 5,17    0       0       5,17
SRV6005 3,184   0       0       3,184
SRV6008 2,384   0       0       2,384
SRV6009 1,844   0       0       1,844
SRV6011 4,191   0       0       4,191
SRV6012 2,86    0       0       2,86
SRV6013 1,992   0       0       1,992
SRV6014 2,473   0       0       2,473
SRV6015 1,969   0       0       1,969
SRV1    259     0       0       259
SRV2    48      0       0       48
SRV4    7,558   0       0       7,558
SRV6    690     0       0       690
SRV6030 4,958   0       0       4,958
SRVW-3D-6       7,654   0       0       7,654
SRVW-3D-7       5,383   0       0       5,383
SRVW-3D-4       37,85   0       0       37,85
SRVW-3D-9       5,48    0       0       5,48
SRVW-3D-8       5,726   0       0       5,726
SRV01   75,338  0       0       75,338
SRV02   355     0       0       355
SRV6010 2,785   0       0       2,785
SRV6017 2,4     0       0       2,4
SRVW-3D-5       12,818  0       0       12,818
SRV03   1,138   0       0       1,138
SRVTE01 215,969 0       0       215,969
SRVTE02 208,552 0       0       208,552
SRV02A  10,32   0       0       10,32
SRVW-3D-26      5,854   0       0       5,854
SRVW-3D-27      6,749   0       0       6,749
SRVW-3D-28      5,234   0       0       5,234
SRVW-3D-29      15,153  0       0       15,153
SRV04   23,638  0       0       23,638
SRV01A  16,204  0       0       16,204
SRVW-3D-30      6,95    0       0       6,95
SRVPDM01        240,081 0       0       240,081
SRVDOMINO       717,595 0       0       717,595
SRVW-3D-31      14,742  0       0       14,742
```

Figure A-3 Output of the select statement

You also can collect data from the accounting log or from the schedule log of the clients.

To evaluate the value of the information about client backups, check the schdule log. Example A-5 shows a `dsmsched.log` file.

Example A-5 Sample schedule log

```
10/19/2013 20:51:40 Backup of object 'SystemState' component 'System State' finished successfully.
10/19/2013 20:51:40 Successful incremental backup of 'GBECKERT520\SystemState\NULL\System State\SystemState'

10/19/2013 20:51:43 --- SCHEDULEREC STATUS BEGIN
10/19/2013 20:51:43 Total number of objects inspected:      276,767
10/19/2013 20:51:43 Total number of objects assigned:        76,463
10/19/2013 20:51:43 Total number of objects backed up:        6,019
10/19/2013 20:51:43 Total number of objects updated:              0
10/19/2013 20:51:43 Total number of objects rebound:              0
10/19/2013 20:51:43 Total number of objects deleted:              0
10/19/2013 20:51:43 Total number of objects expired:              2
10/19/2013 20:51:43 Total number of objects failed:               0
10/19/2013 20:51:43 Total number of subfile objects:              0
10/19/2013 20:51:43 Total number of bytes inspected:          53.49 GB
10/19/2013 20:51:43 Total number of bytes transferred:         4.81 GB
10/19/2013 20:51:43 Data transfer time:                      675.68 sec
10/19/2013 20:51:43 Network data transfer rate:            7,479.27 KB/sec
10/19/2013 20:51:43 Aggregate data transfer rate:          3,012.69 KB/sec
10/19/2013 20:51:43 Objects compressed by:                        0%
10/19/2013 20:51:43 Total data reduction ratio:               91,00%
10/19/2013 20:51:43 Subfile objects reduced by:                   0%
10/19/2013 20:51:43 Elapsed processing time:               00:27:57
10/19/2013 20:51:43 --- SCHEDULEREC STATUS END
10/19/2013 20:51:43 --- SCHEDULEREC OBJECT END TAEGLICHES_BACKUP 10/19/2013 18:00:00
```

```
10/19/2013 20:51:43
Executing Operating System command or script:
   C:\Dasi_ende.cmd
10/19/2013 20:51:43 Scheduled event 'TAEGLICHES_BACKUP' completed successfully.
10/19/2013 20:51:43 Sending results for scheduled event 'TAEGLICHES_BACKUP'.
10/19/2013 20:51:43 Results sent to server for scheduled event 'TAEGLICHES_BACKUP'.

10/19/2013 20:51:43 ANS1483I Schedule log pruning started.
10/19/2013 20:51:43 ANS1484I Schedule log pruning finished successfully.
10/19/2013 20:51:43 TSM Backup-Archive Client Version 6, Release 4, Level 0.38 120723B
10/19/2013 20:51:43 Querying server for next scheduled event.
10/19/2013 20:51:43 Node Name: GBECKERT520
10/19/2013 20:51:43 Session established with server CHEF2_SERVER1: Windows
10/19/2013 20:51:43   Server Version 5, Release 5, Level 7.0
10/19/2013 20:51:43   Server date/time: 10/19/2013 20:51:43  Last access: 10/19/2013 20:45:35
```

For further data analytics, review the data that is collected in the accounting log. A sample of the comma separated content of the accounting log is shown in Example A-6.

Example A-6 Sample accounting log

```
5,0,ADSM,10/18/2013,00:51:03,SRV6001,,AIX,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,556385,3024,1632,0,0,4,0,0,0,0,5,6
5,0,ADSM,10/18/2013,01:00:16,SRVPDM02,,WinNT,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,1,2,2,0,0,4,0,0,0,0,5,6
5,0,ADSM,10/18/2013,01:00:17,SRVPDM02,,WinNT,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,308805,3518,2447,0,0,4,0,0,0,0,5,6
5,0,ADSM,10/18/2013,01:00:18,SRVPDM02_SQL,,WinNT,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,1,3550,3550,0,0,5,0,0,0,0,5,6
5,0,ADSM,10/18/2013,01:04:13,SRVPDM02_SQL,,TDP MSSQL
Win64,1,Tcp/Ip,1,0,0,0,0,20,8178903,0,0,8179992,233,7,76,0,4,0,0,0,0,5,6
5,0,ADSM,10/18/2013,01:04:14,SRVPDM02_SQL,,WinNT,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,1,236,236,0,0,5,0,0,0,0,5,6
5,0,ADSM,10/18/2013,01:04:14,SRVPDM02_SQL,,WinNT,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,1,0,0,0,0,5,0,0,0,0,5,6
5,0,ADSM,10/18/2013,01:21:53,SRVCOGNOS01A_SQL,,TDP MSSQL
Win64,1,Tcp/Ip,1,0,0,0,0,22,66942438,0,0,66950900,1212,14,65,1,4,0,0,0,0,5,6
5,0,ADSM,10/18/2013,01:22:03,SRVCOGNOS01A_CO_SQL,,TDP MSSQL
Win64,1,Tcp/Ip,1,0,0,0,0,16,127033,0,0,127077,8,5,0,0,4,0,0,0,0,5,6
5,0,ADSM,10/18/2013,01:22:03,SRVCOGNOS01A_SQL,,WinNT,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,1,1298,1298,0,0,5,0,0,0,0,5,6
5,0,ADSM,10/18/2013,01:22:04,SRVCOGNOS01A_SQL,,WinNT,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,1,0,0,0,0,5,0,0,0,0,5,6
5,0,ADSM,10/18/2013,06:00:14,SRVPDM02_SQL,,TDP MSSQL Win64,1,Tcp/Ip,1,0,0,0,0,16,1669,0,0,1688,8,8,0,0,4,0,0,0,0,5,6
5,0,ADSM,10/18/2013,06:03:00,SRV6017,,AIX,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,1,0,0,0,0,5,0,0,0,0,5,6
5,0,ADSM,10/18/2013,07:01:03,SRVPDM02_SQL,,WinNT,1,Tcp/Ip,0,0,0,0,0,0,0,0,0,1,3660,1,0,0,5,0,0,0,0,5,6
5,0,ADSM,10/18/2013,07:04:02,SRVCOGNOS02,,WinNT,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,1,0,0,0,0,5,0,0,0,0,5,6
5,0,ADSM,10/18/2013,07:04:04,SRV-3D-47,,WinNT,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,2,0,0,0,0,4,0,0,0,0,5,6
5,0,ADSM,10/18/2013,07:04:04,SRV-3D-47,,WinNT,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,1,0,0,0,0,5,0,0,0,0,5,6
5,0,ADSM,10/18/2013,07:07:25,SRVAPP,,WinNT,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,1,0,0,0,0,5,0,0,0,0,5,6
5,0,ADSM,10/18/2013,07:15:01,SRVLIC01,,WinNT,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,1,0,0,0,0,5,0,0,0,0,5,6
5,0,ADSM,10/18/2013,07:15:05,SRVIIS01,,WinNT,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,1,0,0,0,0,5,0,0,0,0,5,6
5,0,ADSM,10/18/2013,07:21:58,SRVPDM02,,WinNT,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,1,2,2,0,0,4,0,0,0,0,5,6
5,0,ADSM,10/18/2013,07:22:00,SRVPDM02,,WinNT,1,Tcp/Ip,1,0,0,0,0,0,0,0,0,308805,4905,4428,0,0,4,0,0,0,0,5,6
5,0,ADSM,10/18/2013,07:22:01,SRVPDM02
```

In our case, we also collected information about the application and databases, such as size of the database from the application system, and stored it in a separate database such as mysql. We can use this daily collected data for our analysis. If Tivoli Storage Manager is not running, you must find a way to collect this data manually.

The spreadsheet can look like the table in Figure A-4. You can expand this information list based on your requirements. For example, you might decide to add some numbers for typical restore requirements.

Indicator 3	Indicator 2	Indicator 1		TSM SERVER	NODE	CLIENT VERSION	DOMAIN NAME	PLATFORM NAME	SESSIONS	DBSIZE GB	max. LOG / h GB	OCCSIZE GB	NUM OBJECTS	AVG OBJSIZE GB	MB PER SECOND LAN	GB LAN	MB PER SECOND LAN-free	GB LAN-free	AVG MB per second per session	AC MB (p
			1	Server1	Node 1	5,5	DOM_A	AIX	2	0,000	0,000	5.000,00	100.000,00	2,00	20,00	0,00	0,00	0,00	10,00	
			2	Server1	Node 2	5,5	DOM_A	AIX	1	0,000	0,000	3.000,00	5.000,00	50,00	30,00	0,00	0,00	0,00	20,00	
			3	Server2	Node1	5,5	DOM_A	DB2	4	500,000	5,000	30.000,00	2.000,00	2.000,00	50,00	0,00	100,00	0,00	100,00	
			3	Server2	Node2	6,2	DOM_A	Linux86	1	0,000	0,000	200,00	200.000,00	1,00	20,00	0,00	0,00	0,00	20,00	
			2	Server3	Node1	6,2	DOM_A	Linux86	2	0,000	0,000	300,00	500.000,00	0,50	30,00	0,00	0,00	0,00	30,00	
			1	Server3	Node 2	6,2	DOM_A	Windows	1	0,000	0,000	200,00	50.000,00	2,00	10,00	0,00	0,00	0,00	10,00	
			1	Server3	Node3	6,2	DOM_A	TDP SQL	3	1.000,000	10,000	30.000,00	10.000,00	3.000,00	5,00	0,00	200,00	0,00	300,00	
		Summary			8				7			68.700,00			165,00	0,00	300,00	0,00	490,00	

Figure A-4 Sample analysis table

After you transfer the numbers to the spreadsheet, you can use the it to categorize the nodes and assign to domains that you define for the categories that you identify. We use indicator fields to categorize our environment and requirements. Indicator fields can define a category such as database size, type of LAN connection, type of SAN environment, and backup or restore requirements. In the next step we assign values to the indicator fields.

Indicators are helpful numbers or characters that can help you find categories and structures of the data, and can help you sort and organize the data. The indicators are described in a legend. Figure A-5 shows sample indicators in a legend.

Legend:	Indicator Fields / Values		Tier 1: Disk - Tape	Tier 2: Disk - VTL	Tier 3: LAN - VTL	Tier 4: LAN-free - VTL	Tier 5: LAN-free Tape	special case
1	**Indikator total backup amount per day in TB**							
	0	Amount = 0 GB, Occupancy < 1 GB, MB/s LAN=0, MB/s LANfree=0, GB total=0						x
	1	data amount < 1 GB	x					
	2	data amount > 1 GB	x					
	3	data amount > 100 GB		x				
	4	data amount > 200 GB		x				
	5	data amount > 300 GB				x		
	6	data amount > 400 GB					x	
	7	data amount > 500 GB						
2	**indicator average object size**							
	0	average objectsize = 0						x
	1	average objectsize < 1 GB	x					
	2	average objectsize > 1 GB < 10 GB			x			
	3	average objectsize > 10 GB					x	
3	**indicator database size**							
	0	db size = 0	x					
	1	db size < 1 GB	x					
	2	db size > 1 GB	x					
	3	db size > 10 GB		x				
	4	db size > 20 GB			x			
	5	db size > 500 GB				x		
	6	db size > 3000 GB					x	
4	**Indikator throughput LAN, MB/S**							
	1	throughput cumulated < 20 MB/sec	x					
	2	throughput cumulated < 30 MB /sec		x				
	3	throughput cumulated < 40 MB /sec			x			
	4	throughput cumulated < 50 MB /sec				x		
	5	throughput cumulated >= 50 MB/sec					x	
5	**Indikator number parallel sessions**							
	1	number of sessions < 4					x	
	2	number of sessions > = 4	x	x	x	x		

Figure A-5 Legend of sample indicator fields and assigned Indicator Values

The value of the indicators to the data can be done manually or by intelligent formulas. In our example, the value of the indicator is done manually, because during the process we were able to decide the value to assign.

In our example, we use database size as an indicator field and assign these values:

- 0 for file systems without databases
- 1 for small databases up to 10 GB
- 2 for databases between 10 and 20 GB
- 3 for databases up to 50 GB
- 4 for databases up to 500 GB
- 5 for large databases up to 3 TB
- 6 for huge databases above the limit

For network categories, we have another indicator field with these assigned values:

- 1 for 100 Mb connection
- 2 for 1 Gb
- 3 for 10 Gb

The SAN Indicator is defined with these values:

- 0 for no SAN connection available
- 1 for 4 Gb
- 2 for 8 Gb
- 3 for 16 Gb

Another indicator field is the restore requirement, with these values:

- 1 for low priority
- 2 for medium priority
- 3 for high priority

With the summary of all indicators, you can identify your service requirements. With these results you can build your storage categories. In our case (Figure A-6 on page 320), these are the values:

- Values 0, 1, and 2 represent a category that uses a traditional storage hierarchy.
- Value 3 sends the data to disk and migrates to virtual tape library (VTL).
- Value 4 are the nodes to do the backup through LAN direct to the VTL.
- Value 5 will do it LAN-free.
- Value 6 is for huge database systems that will use LAN-free data transfer direct to physical tape.

Another way might be to implement a disk-to-disk backup strategy data replication to another server if the data protection requirements will be identified to a high level of security for special systems that are represented by Tivoli Storage Manager nodes.

The indicators may represent your service requirements. Identify the factors that are important for your backup and restore strategy:

- If you want to meet fast restore requirements for a huge amount of data in few objects, LAN-free data transfer to tape might be your solution.
- If you have millions of small files to protect, incremental backup to random disk pools is a possible implementation.
- If you want to take advantage of data deduplication, put the data on sequential file storage pools with data deduplication and leave it there for a disk only solution.
- If you want to have extra protection for your deduplicated data, take advantage of node replication.
- If the amount is too big, migrate the data from the random disk storage pool to physical tape.

After you categorize the tiers of protection you want to implement, you can use Tivoli Storage Manager to implement the data protection solution to fit your needs. And of course you can configure any combination of them, even in a single server, taking advantage of the powerful Tivoli Storage Manager policy management. Depending on the values of the spreadsheet, and mostly of the indicators, now you can categorize your nodes and assign them to Tivoli Storage Manager objects such as domains. You can use the *Indicator* fields to summarize as in these examples:

- `Indicator value 1 in IndicatorFiled 1` (database size) could represent a category that uses a traditional storage hierarchy
- `Indicator value 2 in Indicator Filed1` (database size) if you want to implement a disk-to-disk backup strategy with data replication to another server.

This is why we refer to this as *Tivoli Storage Manager policy-based dynamic tiering*.

Our goal is to find a way to determine what data ends in what storage tier or protection tier. For our project, we have many file system backups, many small databases, some medium databases, and a few huge databases.

- The file system backups are assigned indicator value 1 in the indicator field 1 (database size), which represents in the storage tiering the way of traditional disk (disk storage pool) migrated to tape.
- Small databases are assigned indicator value of 2 and use for the first storage tier disk, but for the second tier VTL, because we need a good RTO value for recovery and the database backups are good candidates for deduplication.
- Mid-size databases are assigned indicator value 3 and will be stored direct-to-VTL because there are not so many objects and the VTL has enough mount points.
- In an additional indicator value, we can decide that this backup will be done LAN-free. The indicator value 4 describes the LAN-free data transfer directly to physical tape because the daily amount of data is so much that we cannot move it through the network. In addition, the occupancy is also high so that we cannot store it on the expensive VTL.

Figure A-6 shows our project's data flow to solve the tiered storage approach.

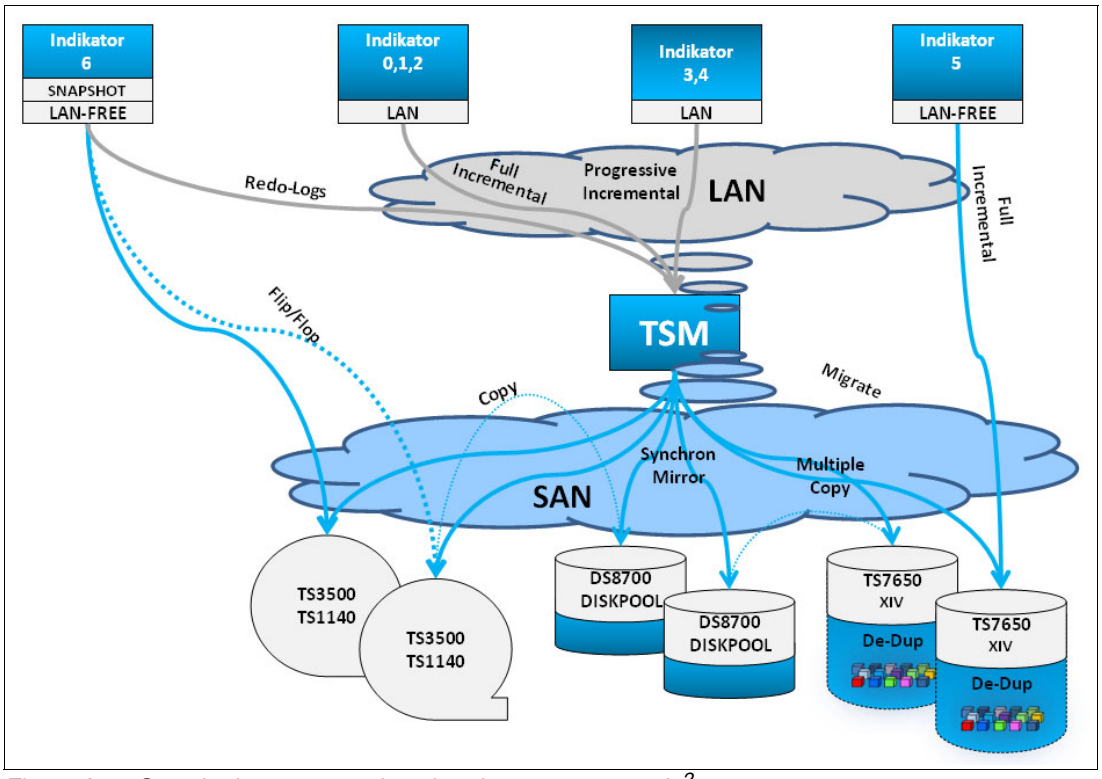

Figure A-6 Sample data structure in a tiered storage approach.[2]

The message here is that you can be creative and innovative when you use the Tivoli Storage Manager toolbox for your data protection solution.

[2] Courtesy of Stephane Criachi - Concat AG

Tivoli Storage Manager policy-based dynamic tiering

Tivoli Storage Manager is the interface between your data and application protection and the data flow through the storage tiers of your available hardware. Figure A-7 shows the layered model.

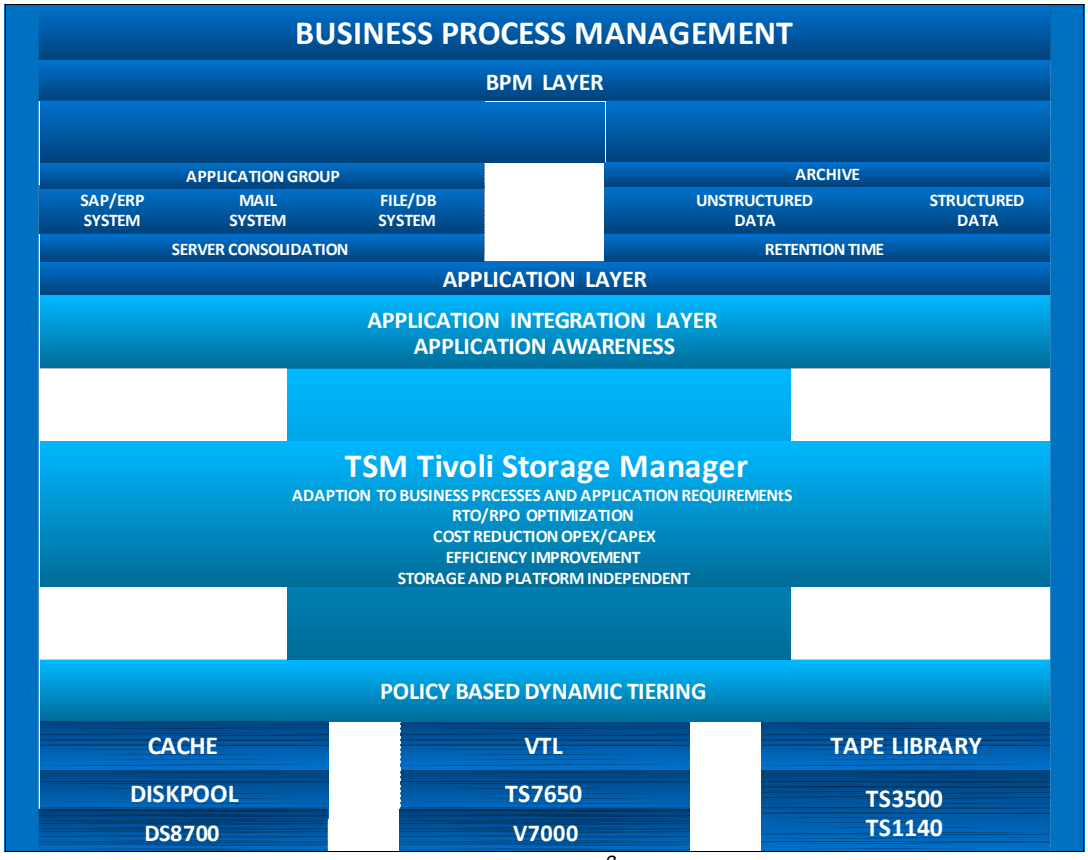

Figure A-7 Layer model to build a flexible infrastructure [3]

Business processes change frequently. As a result, you want to adjust the service level requirements for RTO and RPO dynamically. To assist with this, you can take advantage of the centralized functions for the policy and tiered storage management that is available with the Tivoli Storage Manager product.

You can use the existing mechanisms for data migration, data copy, data deduplication, and data replication. The copy and replication functions are used to create redundant data for disaster protection and offsite vaulting. Migration is used to change the storage tier to another level in the storage hierarchy, for example to move the data to new storage technology or, in the case of lifecycle management, for long-term retention to high capacity, low cost, and lower RTO physical tape. Successful dynamic tiering requires that the defined storage classes for diskpool, filepool, virtual tape library, and physical tape are transparently integrated and controlled in the Tivoli Storage Manager server. Avoid special functions of the storage hardware, for example, VTL tape caching or post deduplication. Be sure that the server controls the data management, and that the service level requirements are guaranteed at recovery time. Special functions of the storage hardware might influence the effectiveness of the Tivoli Storage Manager system or the operating of it.

[3] Courtesy of Stephane Criachi - Concat AG

With dynamic tiering, as explained in this appendix, using all your storage tiers in parallel is possible. This reduces the risk of bottlenecks for backup and restore processes because a central instance controls the data flow. Defined data classes are assigned to the storage tiers and backups will be done, for example, reflecting the available network backbone bandwidth. Tivoli Storage Manager always controls the backup and restore processes.

The storage tiers can handle the calculated amount of data that must be moved in and out and must be stored in the data pools. Each tier is a collection of storage pools and storage pool volumes that represent an amount of available space. They are stretched over all layers in the hierarchy, including disk, VTL, and physical tape and can be adjusted to your needs. You can scale to the calculated data amount from the top to the bottom of the storage hierarchy without a change in the basic technical concept.

This provides you with the necessary investment protection of the overall solution for the life of the data.

Parallelism and distribution are the keys to independent modular storage tiers, giving you enhanced redundancy and linear scalability.

B

Hierarchical storage management (HSM)

This appendix describes Tivoli Storage Manager HSM for Windows and Tivoli Storage Manager for Space management, two important Tivoli Storage Manager family products that benefit from Tivoli Storage Manager data protection.

The hierarchical storage management (HSM) function is implemented differently in Tivoli Storage Manager for Windows and UNIX. In the Windows environment, it is called HSM for Windows; in the UNIX part it is called Tivoli Storage Manager for Space Management. There are major differences in the way they are implemented in their respective supported platforms.

Tivoli Storage Manager HSM for Windows

IBM Tivoli Storage Manager HSM for Windows, referred to as HSM for Windows client, provides Hierarchical Storage Management for Windows file systems. Individual files and files from parts of supported file systems are migrated to remote storage in IBM Tivoli Storage Manager.

In effect, HSM turns the fast disk drives into caches for the slower mass storage devices. The HSM for Windows client monitors the way files are used and lets you automate policies as to which files can safely be moved (migrated) to slower devices and which files should stay on the hard disks.

The HSM for Windows client manages the migration of individual files, files from parts of file systems, or complete file systems, to remote storage in Tivoli Storage Manager. Migrated files can be accessed, opened, and updated by the Windows application corresponding to the file extension.

The HSM for Windows client includes a graphical user interface (HSM for Windows client GUI) that you use to define and run migration jobs, threshold migration, reconciliation, searches and file retrieval, and to define general settings. You can also do many of these tasks using HSM for Windows client commands from a Command Prompt window.

The HSM for Windows client supports local, fixed NTFS, and Resilient File System (ReFS) file systems. This includes Microsoft Cluster Server (MSCS) cluster volumes, if they are formatted in NTFS or ReFS. Windows File Allocation Table (FAT) partitions, Common Internet File System (CIFS) shared folders, network-attached storage (NAS) drives, and other file systems are not supported.

Transparent

Again, from a user or running process perspective, all files in the local file system are actually available. Directory listings and other commands that do not require access to the entire file will appear exactly as they would without the HSM client. When a migrated file is needed by an application or command, the operating system initiates a transparent recall for the file to the Tivoli Storage Manager server. This helps to more easily locate and access offloaded data without administrator interaction. The transparency is not affected by the capability of keeping separate versions of migrated files because a user-initiated recall always gets the latest version. Retrieval is transparently invoked by actions resulting in an open call to the file:

1. Double-click the file from an Explorer window.
2. Select **File** → **Open** on the files icon from the appropriate program, like Windows Explorer or any other file navigator.

If a file is recalled from the server, in the next scheduled archiving run, the retrieved document is replaced by a shortcut ("re-stubbed"). No additional version is stored on the server. An MD5 key is computed for the retrieved document. This MD5 key is compared with the MD5 key stored in the migrated document. If the two MD5 keys match, the file is only replaced with a stub, otherwise a new version of the file is stored in the repository. In this way, the file is migrated again only when necessary (for example, if it changed on the client).

Files migrated to the Tivoli Storage Manager server using the HSM Client for Windows are retained on the server for the length of time defined in the Retain Version field of the archive copy group. You should set this field according to your needs and the space available. This field can be set to NOLIMIT, which means the migrated files are kept on the server indefinitely, regardless of whether the original is deleted from the client. If you set this field to a lesser value, be careful of the possibility that the stub file still exists on the client, when the migrated file on the server has expired. Upon backup of a stub file, the stub file becomes the active copy of the data, marking the original copy inactive. Depending on your policy settings, the original file can be processed by expiration, making it impossible to restore from a backup.

Selective

Selective recalls of migrated files are performed by the HSM Client GUI or command line. These interfaces provide powerful controls of what to retrieve and where. They provide the ability to recall different versions of the files that were sent to the Tivoli Storage Manager archive pool, or files where the stub file has been deleted from the client system. A selective recall can be directed to the same or different directory from the one where the file was originally migrated. Selective recalls can be submitted only by the Tivoli Storage Manager HSM for Windows administrator.

Tivoli Storage Manager for Space Management

The IBM Tivoli Storage Manager for Space Management client for UNIX and Linux (the HSM client) migrates files from your local file system to distributed storage and can then recall the files automatically or selectively. Migrating files to storage frees space for new data on your local file system and takes advantage of lower-cost storage resources that are available in your network environment.

When a file is migrated from your local system to Tivoli Storage Manager storage, a placeholder, or stub file, is created in place of the original file. Stub files contain the necessary

information to recall your migrated files and remain on your local file system so that the files appear to reside locally. This process contrasts with archiving, where you usually delete files from your local file system after archiving them.

The HSM client provides space management services for locally mounted file systems, and it migrates regular files only. It does not migrate character-special files, block-special files, named pipe files, or directories.

File migration, unlike file backup, does not protect against accidental file deletion, file corruption, or disk failure. Continue to back up your files regardless of whether they reside on your local file system or are migrated to Tivoli Storage Manager storage. The IBM Tivoli Storage Manager backup-archive client is used to back up and restore migrated files in the same manner as you would back up and restore files that reside on your local file system. If you accidentally delete stub files from your local file system, or if you lose your local file system, you can restore the stub files or the complete files.

For planned processes, such as storing a large group of files in storage and returning them to your local file system for processing, use the archive and retrieve processes. The backup-archive client is used to archive and retrieve copies of migrated files in the same manner as you would archive and retrieve copies of files that reside on your local file system.

The HSM client functions for threshold migration, demand migration, selective migration, and selective and transparent recall include processing GPFS file systems containing multiple space-managed storage pools.

The HSM client has both a graphical user interface (the HSM GUI) and commands you can run from a shell. You can also use the commands in scripts and cron jobs.

C

VSS and Tivoli Storage Manager related product concepts

Two important Tivoli Storage Manager product concepts must be understood before you can comprehend how the software application components interact with each other.

- Proxy node

  This is the name given to a Tivoli Storage Manager node (agent) that can act on behalf of another node (target). The proxy node functionality was created to facilitate backup and restore of data to a single namespace by multiple Tivoli Storage Manager client nodes. In a Tivoli Storage Manager VSS environment, the proxy node function is utilized to enable the Tivoli Storage Manager backup-archive client to act as an agent node on behalf of the Data Protection client target node. The Tivoli Storage Manager backup-archive client is able to store data under the Data Protection client node at backup time and restore data stored under the Data Protection Client node at restore time.

- Client-to-client communication

  This is the method used for one Tivoli Storage Manager client node to contact and communicate with another Tivoli Storage Manager client node. The nodes can reside on the same or on different computer hosts. The contacting node sends commands to the receiving node to perform backup, query, or restore tasks. The receiving node is always considered a DSM Agent node. Authentication and coordination between these two nodes is tightly coupled with Tivoli Storage Manager security; the two communicating Tivoli Storage Manager client nodes are in a target-agent relationship as defined on the Tivoli Storage Manager server.

Figure C-1 shows the Tivoli Storage Manager VSS environment.

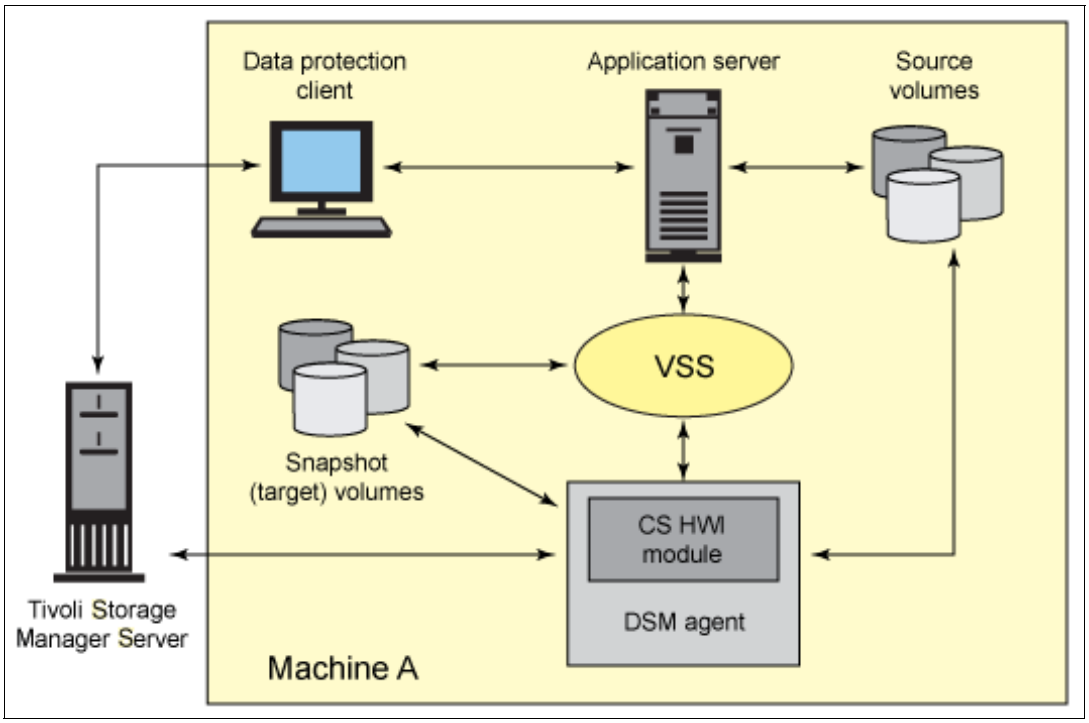

Figure C-1 Tivoli Storage Manager VSS environment

Because the Data Protection Client and the DSM agent are in a proxy relationship, they are able to communicate using client-to-client communication. As a result, during the backup operation, the Data Protection Client communicates with the DSM agent, which in turn communicates with the Volume Shadow Copy Service. The Volume Shadow Copy Service then communicates with the Application Server (such as Microsoft Exchange Server or Microsoft SQL Server) via the application VSS Writer to start the snapshot.

The Data Protection client supports the following types of VSS backups:

- Local

  This backup type creates a persistent shadow copy that resides on a snapshot volume. Although this type of backup is managed by Tivoli Storage Manager policy, the actual data resides on volumes local to the VSS Provider and not on Tivoli Storage Manager controlled storage. The VSS Provider software controls how that shadow copy is created and maintained, and the Tivoli Storage Manager policy manages its lifecycle.

- Tivoli Storage Manager

  This backup type creates a copy of application data on Tivoli Storage Manager server storage. This data from an application server VSS snapshot is copied to the Tivoli Storage Manager server as backup data. This data is also managed by Tivoli Storage Manager policy. When this backup is complete, the snapshot volumes are released.

- Both

  This backup type creates two copies of application server data. One copy resides on local volumes as a snapshot and another copy (taken from the snapshot copy) is sent to Tivoli Storage Manager server storage.

- Off-loaded backup

  This backup type refers to a Tivoli Storage Manager backup that is done from a system other than the application server. To accomplish this, a shadow copy volume is imported to

a secondary host for the purpose of backing up the application data from the snapshot copy to Tivoli Storage Manager server storage. The benefit is that most of the resource load to back up the files is shifted from the production machine to a secondary host machine. As a result, there is little or no resource impact on the production machine during the VSS backup.

Relating Tivoli Storage Manager backup types with VSS terms

The backup types described previously have a close relationship with the VSS snapshot types in terms of where the data is stored. This section describes how the VSS snapshot types relate to the Tivoli Storage Manager backup types.

- Tivoli Storage Manager

  This backup type means that the VSS snapshot is non-persistent. The snapshot is created locally and resides locally and exists only long enough to complete the backup to Tivoli Storage Manager server storage.

- Local

  When the VSS snapshot data resides locally to the application server, the VSS snapshot is persistent.

- Both

  When the VSS snapshot data resides both locally to the VSS Provider and on Tivoli Storage Manager server storage, the VSS snapshot is persistent.

- Off-loaded backup

This backup type means the VSS snapshot is both transportable and non-persistent.

D

Additional material

This book refers to additional material that can be downloaded from the Internet as described in the following sections.

Locating the Web material

The Web material associated with this book is available in softcopy on the Internet from the IBM Redbooks Web server. Point your Web browser at:

ftp://www.redbooks.ibm.com/redbooks/SG248134

Alternatively, you can go to the IBM Redbooks website:

ibm.com/redbooks

Select the **Additional materials** and open the directory that corresponds with the IBM Redbooks form number, SG248134.

Using the Web material

The additional Web material that accompanies this book includes the following files:

File name	*Description*
solution_matrix.zip	Matrix hardcopy spreadsheet source

System requirements for downloading the Web material

The Web material requires the following system configuration:

Hard disk space: 205 KB

Downloading and extracting the Web material

Create a subdirectory (folder) on your workstation, and extract the contents of the web material `.zip` file into this folder.

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

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *IBM ProtecTIER Implementation and Best Practices Guide,* SG24-8025
- *IBM Tivoli Storage Manager: Building a Secure Environment*, SG24-7505
- *Infrastructure Solutions: Design, Manage, and Optimize a 20 TB SAP NetWeaver Business Intelligence Data Warehouse*, SG24-7289
- *SONAS Concepts, Architecture, and Planning Guide*, SG24-7963

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources:

- *Tivoli Storage Manager for Enterprise Resource Planning - Data Protection for SAP, Installation and User's Guide for DB2*, SC33-6341
- *Tivoli Storage Manager for Enterprise Resource Planning - Data Protection for SAP, Installation and User's Guide for Oracle*, SC33-6340
- *IBM Tivoli Storage FlashCopy Manager V3.2: Installation and User's Guide on UNIX and Linux,* SC27-4005
- *IBM SONAS Introduction and Planning Guide,* GA32-0716

Online resources

These websites are also relevant as further information sources:

- Tivoli Storage Manager:

  http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/index.jsp

- Tivoli Storage Manager Data Deduplication FAQ:

  https://www.ibm.com/developerworks/community/wikis/home?lang=en#/wiki/Tivoli%20Storage%20Manager/page/Deduplication%20FAQ

- Guidelines for node replication (on Tivoli Storage Manager Wiki):

  https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/Guidelines%20for%20node%20replication

- Sample Architectures (on Tivoli Storage Manager Wiki):

  https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/Sample%20Architectures

- Deployment recommendations for Tivoli Storage Manager V6 servers:

  http://www.ibm.com/support/docview.wss?uid=swg21421060

- Optimizing performance for servers and clients:

  https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/Optimizing%20performance%20for%20servers%20and%20clients

- Tivoli Storage Manager Deduplication Sample Architecture on AIX:

  https://www.ibm.com/developerworks/community/wikis/home?lang=en#/wiki/Tivoli%20Storage%20Manager/page/Tivoli%20Storage%20Manager%20Deduplication%20Sample%20Architecture%20on%20AIX

- Effective Planning and Use of IBM Tivoli Storage Manager V6 Deduplication:

  https://www.ibm.com/developerworks/community/wikis/home?lang=en#/wiki/Tivoli%20Storage%20Manager/page/Effective%20Planning%20and%20Use%20of%20IBM%20Tivoli%20Storage%20Manager%20V6%20Deduplication

- ProtecTIER and Tivoli Storage Manager Performance Tuning:

  http://www.ibm.com/support/docview.wss?uid=tss1wp102008

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

IBM Tivoli Storage Manager as a Data Protection Solution

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

IBM Tivoli Storage Manager as a Data Protection Solution

Provides solutions for common data protection challenges

Includes challenges and solution matrices

Effectively use the Tivoli Storage Manager Toolkit

When you hear IBM Tivoli Storage Manager, the first thing that you typically think of is data backup. Tivoli Storage Manager is the premier storage management solution for mixed platform environments.

Businesses face a tidal wave of information and data that seems to increase daily. The ability to successfully and efficiently manage information and data has become imperative. The Tivoli Storage Manager family of products helps businesses successfully gain better control and efficiently manage the information tidal wave through significant enhancements in multiple facets of data protection.

Tivoli Storage Manager is a highly scalable and available data protection solution. It takes data protection scalability to the next level with a relational database, which is based on IBM DB2 technology. Greater availability is delivered through enhancements such as online, automated database reorganization.

This IBM Redbooks publication describes the evolving set of data-protection challenges and how capabilities in Tivoli Storage Manager can best be used to address those challenges. This book is more than merely a description of new and changed functions in Tivoli Storage Manager; it is a guide to use for your overall data protection solution.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-8134-00 ISBN 0738439983