IBM System Storage Business Continuity Solutions Overview

- Anticipating and responding to risks quickly and cost-effectively
- Reviewing of Business Continuity expertise and skills
- Using IBM System Storage products

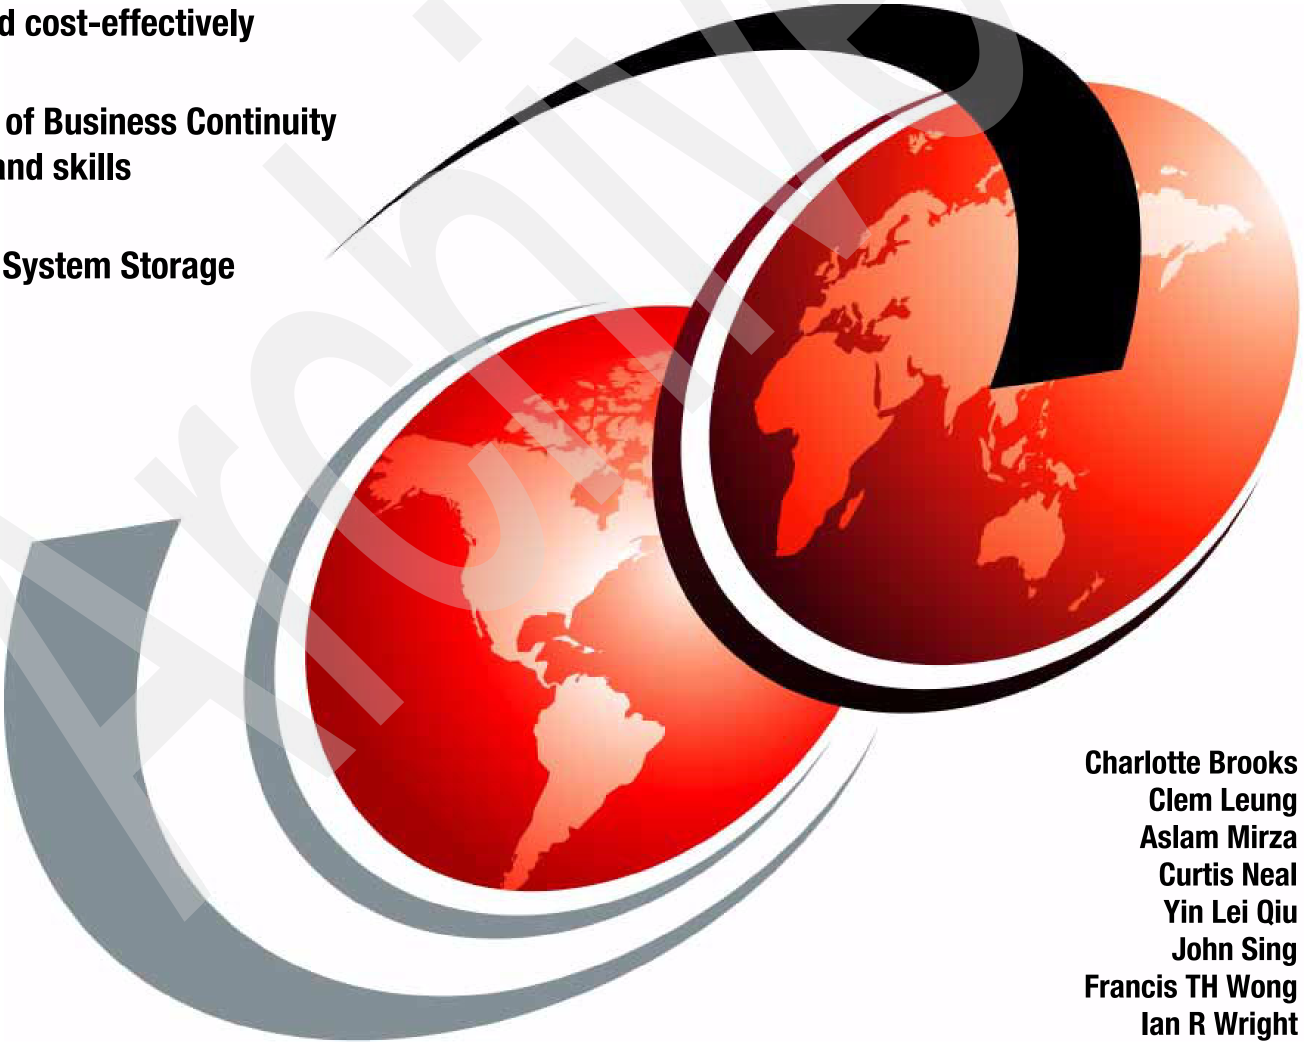

Charlotte Brooks
Clem Leung
Aslam Mirza
Curtis Neal
Yin Lei Qiu
John Sing
Francis TH Wong
Ian R Wright

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM System Storage Business Continuity Solutions Overview

February 2007

SG24-6684-01

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

Second Edition (February 2007)

Contents

Notices vii
Trademarks viii

Preface ix
The team that wrote this redbook. ix
Become a published author xii
Comments welcome. xii

Part 1. Business Continuity principles 1

Chapter 1. Introduction to Business Continuity in an On Demand World 3
1.1 Business Continuity pain points 4
1.2 What is Business Continuity 4
1.2.1 Definition: value of Business Continuity to the entire business. 5
1.2.2 What is IT Business Continuity 5
1.2.3 Definition: high availability. 6
1.2.4 Continuous operations 6
1.2.5 Disaster Recovery. 7
1.2.6 Differentiating Business Continuity from Disaster Recovery. 7
1.3 Business Continuity: Why now 7
1.3.1 If you are down, what do you lose 7
1.3.2 What are the most likely causes of outages 7
1.3.3 Assess where you are today. 8
1.3.4 Determine what your metrics of success should be 8
1.3.5 Final step: Decide where you want to go 8
1.4 Summary. 8

Chapter 2. Selecting Business Continuity solutions 11
2.1 Roadmap to IT Business Continuity 12
2.2 The System Storage Resiliency Portfolio 13
2.2.1 Reliable hardware infrastructure layer 14
2.2.2 Core technologies layer 14
2.2.3 Platform-specific integration layer. 15
2.2.4 Application-specific integration layer. 16
2.2.5 Summary. 16
2.3 Selecting Business Continuity solutions 16
2.4 The tiers of Business Continuity 17
2.5 Blending tiers into an optimized solution. 18
2.5.1 Three Business Continuity solution segments 19
2.6 Select a solution by segment 19
2.6.1 Refine with a detailed evaluation team 20
2.6.2 Value of Business Continuity solution selection via segmentation 20

Part 2. Business Continuity solution offerings 21

Chapter 3. Continuous Availability 23
3.1 Geographically Dispersed Parallel Sysplex (GDPS). 24
3.1.1 GDPS/XRC overview 27
3.1.2 GDPS/Global Mirror (GDPS/GM) overview. 28

3.1.3 GDPS three site support 30
3.1.4 Global Technology Services (GTS) offerings for GDPS overview 31
3.1.5 GDPS summary 32
3.1.6 Additional GDPS information 32
3.2 GDOC 32
3.2.1 Solution components 33
3.2.2 GDOC in greater detail 34
3.2.3 Additional information 35
3.3 HACMP/XD 35
3.4 Continuous Availability for MaxDB and SAP liveCache 38
3.4.1 Solution components 39
3.5 Copy Services for System i 39
3.5.1 Copy Services with IASPs. 40
3.5.2 Copy Services for System i. 41
3.5.3 Metro Cluster for N series overview 42

Chapter 4. Rapid Data Recovery 45
4.1 System Storage Rapid Data Recovery for System z and mixed z+Open platforms (GDPS/PPRC HyperSwap Manager) 46
4.1.1 GDPS/PPRC Open LUN management. 47
4.1.2 Product components 47
4.1.3 Additional information 48
4.2 System Storage Rapid Data Recovery for System z and mixed z+Distributed platforms using TPC for Replication. 48
4.3 TPC for Replication functionality 49
4.3.1 Functionality of TPC for Replication Two Site BC. 49
4.3.2 Environment and supported hardware 49
4.4 IBM System Storage SAN Volume Controller (SVC) 49
4.4.1 SVC FlashCopy 50
4.5 SVC remote mirroring 51
4.6 Rapid Data Recovery with System i 52
4.7 FlashCopy Manager and PPRC Migration Manager. 53

Chapter 5. Backup and Restore. 55
5.1 An overview of Backup and Restore and archive and retrieve 56
5.1.1 What is Backup and Restore 56
5.1.2 What is archive and retrieve 56
5.2 IBM Tivoli Storage Manager overview 56
5.2.1 IBM Tivoli Storage Manager solutions overview. 57
5.2.2 IBM Tivoli Storage Manager solutions in detail. 58
5.2.3 Integrated solution capabilities of IBM Tivoli Storage Manager 71
5.2.4 Further information 72
5.2.5 IBM Tivoli Storage Manager summary 72
5.3 IBM Data Retention 550 73
5.3.1 Retention-managed data 73
5.3.2 Storage and data characteristics 73
5.3.3 IBM strategy and key products 74
5.3.4 IBM System Storage DR550. 76
5.4 Summary. 77
5.5 System z backup and restore software 77
5.5.1 DFSMSdss 77
5.5.2 DFSMShsm. 77
5.5.3 z/VM Utilities 78

Part 3. End-to-end Business Continuity solution components 79

Chapter 6. Networking and inter-site connectivity options 81
6.1 Network topologies 82
6.2 Fiber transport 83
6.2.1 Dedicated fiber 83
6.2.2 SONET 83
6.2.3 Data transport speed, bandwidth, and latency 84
6.2.4 Technology selection 86
6.3 Wavelength Division Multiplexing (WDM) 87
6.3.1 Optical amplification and regenerative repeaters 87
6.3.2 CWDM versus DWDM 88
6.3.3 Subrate multiplexing 88
6.4 Channel extension 89
6.4.1 Methods of channel extension 89
6.5 Testing 90

Chapter 7. High-availability clusters and database applications 93
7.1 High availability 94
7.2 Clustering technologies 95
7.2.1 Shared nothing clusters 95
7.2.2 Common shared cluster 97
7.2.3 Shared nothing application cluster 98
7.2.4 Geographically dispersed clusters 99
7.2.5 Backup and recovery considerations for databases 100

Chapter 8. Business Continuity for small and medium sized business 107
8.1 Small and medium sized business overview 108
8.1.1 SMB company profiles and business continuity needs 108
8.1.2 SMB company IT needs as compared to large enterprises 109
8.1.3 SMB IT data center and staff issues 109
8.2 Business Continuity for SMB companies 110
8.2.1 Major SMB business continuity design components 110
8.2.2 Business Continuity impacts on SMB business 111
8.3 Successful SMB Business Continuity planning and implementation 111
8.3.1 SMB business continuity implementation steps 111
8.3.2 SMB Business Continuity affordability 112
8.4 SMB Business Continuity solution components 113
8.4.1 Typical SMB BC solutions: performance and downtime 113
8.5 Summary 115

Appendix A. Tiers of Business Continuity 117
Further definition of the tiers of Business Continuity 118
BC Tier 0: No off-site data 118
8.5.1 BC Tier 1: Data backup with no hot site 118
BC Tier 2: Data backup with a hot site 118
BC Tier 3: Electronic vaulting 118
BC Tier 4: Point-in-time copies 119
BC Tier 5: Transaction integrity 119
BC Tier 6: Zero or little data loss 120
BC Tier 7: Highly automated, business integrated solution 120

Appendix B. IBM System Storage Components and Copy Services 121
IBM System Storage DS Family 122

IBM System Storage DS8000 series. 122
IBM System Storage DS6000 122
The IBM TotalStorage ESS. 122
Advanced Copy Services for DS6000 and DS8000 122
IBM System Storage DS4000 Family 130
Advanced Copy Services for DS4000. 130
IBM System Storage N series 136
Advanced Copy Services for N series. 136
IBM System Storage SAN Volume Controller 140
Advanced Copy Services for SVC 141
IBM System Storage Virtualization Engine TS7500 and TS7700 144

Appendix C. Tivoli Storage Manager family overview 151
IBM Tivoli Storage Manager. 152
Disaster Recovery Manager (DRM) 153
IBM Tivoli Storage Manager for Databases 153
IBM Tivoli Storage Manager for Mail 154
IBM Tivoli Storage Manager for Enterprise Resource Planning. 154
IBM Tivoli Storage Manager for Copy Services 154
IBM Tivoli Storage Manager for Advanced Copy Services 155
IBM Tivoli Storage Manager for System Backup and Recovery 155

Appendix D. Services and planning 157
Services and services providers 158
IBM Global Services services families and life cycle 158
On Demand services 160
IBM Global Services solutions for resilient infrastructures. 162
Network Consulting and Integration services. 163
Optical/Storage Networking 163
Network Deployment. 165
Network Management. 165
Network Cabling 166

Related publications 167
IBM Redbooks 167
How to get IBM Redbooks 168
Help from IBM 168

Index 169

Notices

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Trademarks

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

AIX 5L™
AIX®
Domino®
DB2 Universal Database™
DB2®
DFSMS™
DFSMS/VM™
DFSMSdfp™
DFSMSdss™
DFSMShsm™
DS4000™
DS6000™
DS8000™
Enterprise Storage Server®
ESCON®
FlashCopy®
FICON®
Geographically Dispersed Parallel Sysplex™
GDPS®
HyperSwap™
HACMP™
Redbooks (logo) ™
i5/OS®
IBM®
Informix®
Lotus®
NetView®
OS/400®
Parallel Sysplex®
POWER5™
Redbooks™
S/390®
System i™
System p™
System p5™
System z™
System Storage™
System Storage DS™
SysBack™
Tivoli®
TotalStorage®
Virtualization Engine™
VSE/ESA™
WebSphere®
z/OS®
z/VM®

The following terms are trademarks of other companies:

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

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

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

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

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

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

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

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

Preface

This IBM® Redbook presents an overview and descriptions of IBM System Storage™ Business Continuity Solutions for Backup / Restore, Rapid Data Recovery, and Continuous Availability.

IT Business Continuity concepts are discussed; advice, tips, and a roadmap for selecting the optimum IT Business Continuity solution for your organization is provided.

IBM System Storage products are described that can fulfill your IT Business Continuity solution design. This IBM Redbook is a summary of the detailed information contained in *IBM System Storage Business Continuity: Part 1 Planning Guide*, SG24-6547 and *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548.

The team that wrote this redbook

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

Figure 1 The team: Curtis, Aslam, Yin Lei, Ian, John, Charlotte, Clem, and Francis

Charlotte Brooks is an IBM Certified IT Specialist and Project Leader for Storage Solutions at the International Technical Support Organization, San Jose Center. She has 15 years of experience with IBM in storage hardware and software support, deployment, and management. She has written many IBM Redbooks™, and has developed and taught IBM classes in all areas of storage and storage management. Before joining the ITSO in 2000, she was the Technical Support Manager for Tivoli® Storage Manager in the Asia Pacific Region.

Clem Leung is an Executive IT Architect with the IBM Global Small and Medium Business sector, supporting emerging and competitive customers. He specializes in IT infrastructure simplification and Business Continuity technologies and solutions. Previously, he was in worldwide technical sales support for IBM storage and storage networking solutions and products. Clem has worked for IBM for 25 years in various technical sales capacities, including networking, distributed computing, data center design and more. He was an author of a previous edition of this IBM Redbook.

Aslam Mirza is a Certified Senior Consulting Storage Specialist in New York, working as a pre-sales advisor for enterprise storage topics. He has more than 30 years of experience with IBM large systems, storage systems, tape systems, and system storage resiliency portfolio. His area of expertise is strategy and design of storage solutions.

Curtis Neal is a Senior IT Specialist working for the System Storage Group in San Jose, California. He has over 25 years of experience in various technical capacities, including mainframe and open system test, design, and implementation. For the past six years, he has led the Open Storage Competency Center, which helps customers and IBM Business Partners with the planning, demonstration, and integration of IBM System Storage Solutions.

Yin Lei Qiu is a senior IT specialist working for the Storage Systems Group in Shanghai, China. He is the leader of the storage technical team in East China and a pre-sales advisor, and provides technical support storage solutions to IBM professionals, IBM Business Partners, and clients. He has more than six years of solution design experience with IBM Enterprise Disk Storage Systems, Midrange Disk Storage Systems, NAS Storage Systems, Tape Storage Systems, Storage Virtualization Systems, and the System Storage Resiliency Portfolio.

John Sing is a Senior Consultant with IBM Systems and Technology Group, Business Continuity Strategy and Planning. He helps with planning and integrating IBM System Storage products into the overall IBM Business Continuity strategy and product portfolio. He started in the Business Continuity arena in 1994 while on assignment to IBM China. In 1998, John joined the IBM ESS planning team for PPRC, XRC, and FlashCopy®, and then in 2000, became the Marketing Manager for the ESS Copy Services. In 2002, he joined the Systems Group. John has been with IBM for 23 years. He was an author of a previous edition of this IBM Redbook.

Francis TH Wong is a storage solution architect for Asia Pacific, where he provides training and technical support to the regional storage team, as well as designing customer storage solutions. He has 20 years IT experience in various positions with IBM in both Australia and China, including data center operations and S/390® storage support, as well as customer sales, technical support, and services. His areas of expertise include Business Continuity solutions for mainframe and open systems, disk, tape, and virtualization.

Ian R Wright is a Senior IT Specialist with Advanced Technical Support, in Gaithersburg, and is part of the Business Continuity Center of Competence. He holds a Bachelor of Science in Business Administration degree from Shippensburg University of Pennsylvania. He has seven years of IT experience, encompassing Advanced Business Continuity Solutions, network connectivity, and GDPS® for the S/390 division. He has written educational material on Business Continuity and taught at the Business Continuity Top Gun. He was an author of a previous edition of this IBM Redbook.

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

Gustavo Castets, Bertrand Dufrasne, Babette Haeusser, Emma Jacobs, Mary Lovelace, Alex Osuna, Jon Tate, Wade Wallace
International Technical Support Organization, San Jose Center

Michael Stanek
IBM Atlanta

Steven Cook, Douglas Hilken, Bob Kern
IBM Beaverton

Tony Abete, David Sacks
IBM Chicago

Shawn Bodily, Dan Braden, Mike Herrera, Eric Hess, Judy Ruby-Brown, Dan Sunday
IBM Dallas

Bill Wiegand
IBM Fort Wayne

Craig Gordon, Rosemary McCutchen, David Petersen,
IBM Gaithersburg

Thomas Luther
IBM Germany

Manny Cabezas
IBM Miami

Nick Clayton
IBM Portsmouth

Noshir Dhondy, Scott Epter, David Raften
IBM Poughkeepsie

John Foley, Harold Pike
IBM Raleigh

Selwyn Dickey
IBM Rochester

Jeff Barckley, Charlie Burger, Don Chesarek, Pete Danforth, Scott Drummond, John Hulsey, Tricia Jiang, Sathees Kodi, Vic Peltz, John Power, Peter Thurston
IBM San Jose

Greg Gendron
IBM San Ramon

Chooi Ling Lee
IBM Singapore

Thomas Maher
IBM Southfield

Mark S. Putnam
IBM St Louis

Matthias Werner
IBM Switzerland

Bob Bartfai, Ken Boyd, James Bridges, Ken Day, Brad Johns, Carl Jones, Greg McBride, JD Metzger, Jon Peake, Tony Pearson, Gail Spear, Paul Suddath, Steve West
IBM Tucson

Patrick Keyes
IBM UK

Tom and Jenny Chang and their staff
Garden Inn, Los Gatos

Become a published author

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

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an e-mail to:

  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

Part 1

Business Continuity principles

In this part, we discuss the terms and definitions of Business Continuity, as well as an overview of the Business Continuity Solution Selection Methodology.

1

Introduction to Business Continuity in an On Demand World

In today's online, highly connected, fast-paced world, we all expect that today's information technology (IT) systems will provide high availability, continuous operations, and can be quickly recovered in the event of a disaster.

Yet today's IT environment also features an ever growing time to market pressure, with more projects to complete, more IT problems to solve, and a steep rise in time and resource limitations. In light of these realities, what is to be done?

Thankfully, today's IT technology also features unprecedented levels of functionality, features, and lowered cost. In many ways, it is easier than ever before to find IT technology that can address today's business concerns.

This IBM Redbook discusses the thought processes, methods, solutions, and product concepts that can be applied to today's Business Continuity requirements. This chapter begins by examining Business Continuity from a *business* perspective.

By gaining an increased ability to understand, evaluate, select, and implement solutions that successfully answer today's Business Continuity requirements, today's enterprises can continue to maintain marketplace readiness, competitive advantage, and sustainable growth.

1.1 Business Continuity pain points

There is little doubt that today's highly competitive business marketplace leaves little room for error in terms of availability, continuous operations, or recovery in the event of an unplanned outage.

In today's world, an outage could be a power outage, a network outage, a logical corruption outage via unintentional or intentional corruption of files, and many other events. An outage can be triggered by the loss of access to a call center, or by loss of access to an IBM Business Partner's application. In today's connected world, an event that makes the business data *unavailable, even for relatively short periods of time,* has the potential to be of major impact.

In light of these possible impacts, the questions that we need to ask ourselves include:

- Do we have adequate control over prevention of business process or IT infrastructure downtime?
- Do we have adequate IT capabilities to insure continuous operations?
- Do we have adequate regulatory compliance auditability in place? Do we have adequate procedures and planning that can support maintenance of regulatory compliance?
- Can we afford the necessary IT Business Continuity and regulatory compliance with our current cash flow and budgetary posture?

If you answered "No" to any of the above questions, this IBM Redbook can be of help.

As budget pressures tighten, you may also be asking yourself:

- What business problem(s) will IT Business Continuity solve, especially if you do not experience an unplanned IT outage?
- How much will IT Business Continuity cost?
- How do we discover how much we can afford?
- In fact, how should we define what level of IT Business Continuity we really need anyway?

We will answer these questions as well.

1.2 What is Business Continuity

Today's senior management, with ever increasing time to market and resource constraint pressures, will need the answers to two major questions when considering the value of any proposed investment in improved Business Continuity IT infrastructure. These questions are:

1. What is the overall *business* value of the proposed Business Continuity project to the business? What will this function do for our daily competitiveness, responsiveness, expense posture, and profit?
2. What is the relationship between these business benefits and the associated IT infrastructure and IT operations requirements?

Good questions. In this IBM Redbook, we will define what Business Continuity means to the *business* as follows:

Business Continuity:

The ability to adapt and respond to risks, as well as opportunities, in order to maintain continuous business operations, be a more trusted partner, and enable growth.

1.2.1 Definition: value of Business Continuity to the entire business

Here we document many reasons that Business Continuity has been raised to the higher level of a *business tool* for being responsive, flexible, and competitive. In effect, what business problem(s) does a Business Continuity solution solve?

- Continuity of business operations: Helps the business become more anticipatory, adaptive, and robust, from IT to all business processes (for example, taking orders, shipping, manufacturing, and so on).
- Regulatory compliance: Helps the business comply with new government rules and regulations, more quickly and cost effectively.
- Reduce the cost of risk management: Helps the business stay competitive by managing risk more efficiently and cost effectively.
- Security, privacy, and data protection: Help the business protect against threats, both internal and external, and develop a critical information management strategy.
- Expertise and skills (outsourcing or training): Help the business obtain and program manage the expertise and skills necessary to maintain continuous business operations.
- Maintaining market readiness: Helps the business stay competitive by anticipating and quickly responding/adapting to changes in market requirements, and accelerating project completion, to ensure delivery of the right products, at the right place, at the right time, with the right infrastructure.
- Become a more attractive partner: Helps the IBM Business Partner quickly and effectively within their industry by becoming a trusted and reliable IBM Business Partner in their supply chain or value-net.

Strategically, the factors above should allow IT infrastructure Business Continuity to be an important strategic factor in sustainable growth, better profitability, and increased shareholder value.

Now that we have reviewed Business Continuity as viewed by the business, in the remainder of this IBM Redbook, we will focus on the *preservation and recovery of the IT infrastructure* portions of the business.

1.2.2 What is IT Business Continuity

Next, let us define what *IT* Business Continuity is.

There are three primary aspects to Business Continuity; they are related, yet each is different from the others. These aspects are:

- High availability
- Continuous operations
- Disaster Recovery

Figure 1-1 gives an overview of these aspects.

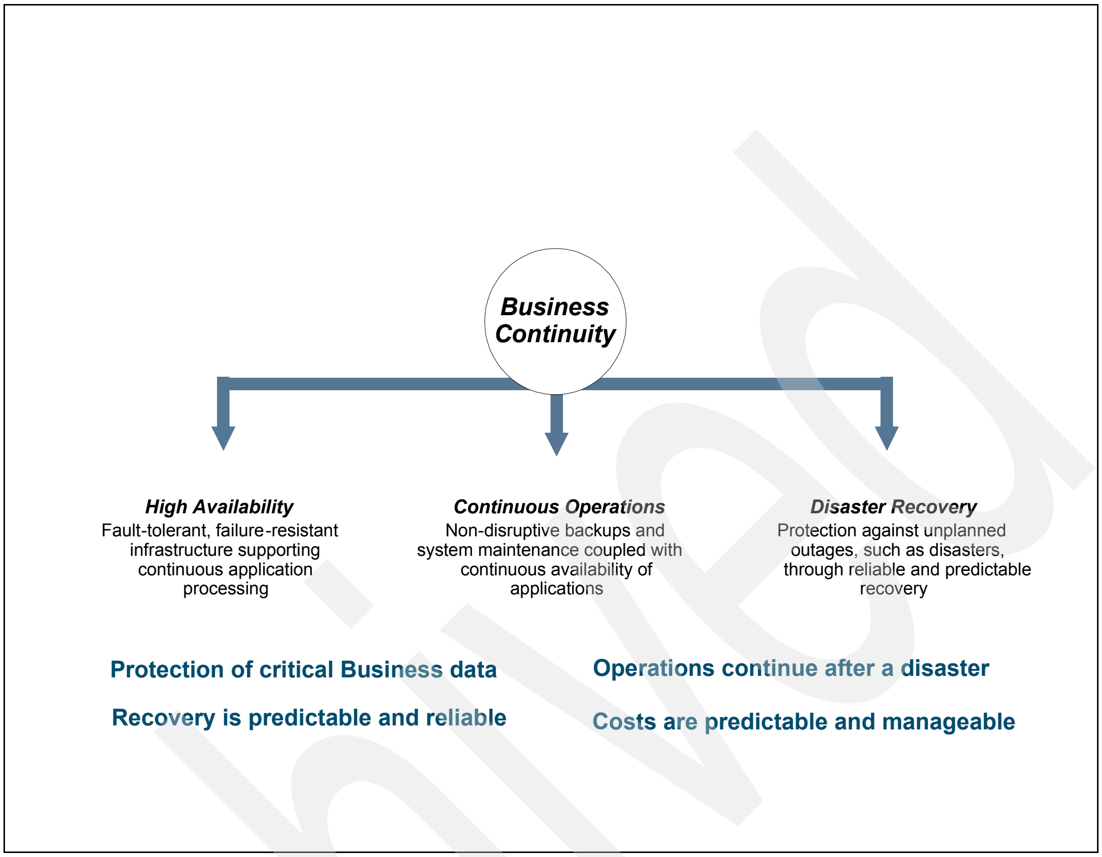
Business Continuity
High Availability
Fault-tolerant, failure-resistant infrastructure supporting continuous application processing
Continuous Operations
Non-disruptive backups and system maintenance coupled with continuous availability of applications
Disaster Recovery
Protection against unplanned outages, such as disasters, through reliable and predictable recovery
Protection of critical Business data
Recovery is predictable and reliable
Operations continue after a disaster
Costs are predictable and manageable

Figure 1-1 Three aspects of Business Continuity: high availability, continuous operations, and Disaster Recovery

1.2.3 Definition: high availability

High availability is the ability and processes to provide access to applications regardless of *local* failures, whether they be in the business processes, in the physical facilities, or in the IT hardware or software.

From an IT standpoint, high availability is often provided by reliable, redundant hardware and software, often running on server clustering solutions that work within the operating systems, and coupled with the hardware infrastructure to remove single points of failure.

1.2.4 Continuous operations

Continuous operations is the ability to keep things running *when everything is working properly*, that is, where you do not have to take applications down merely to do scheduled backups or planned maintenance. Continuous operations technologies provide the ability to perform repeating, ongoing, and necessary infrastructure actions, while still maintaining high availability.

Normally, all the components that make up continuous operations are situated in the same computer room. The building, therefore, becomes the single point-of-failure. Thus, a continuous operation technology does not necessarily mean that it is also a disaster recovery solution.

1.2.5 Disaster Recovery

Finally, Disaster Recovery is the ability to recover a datacenter *at a different site, on different hardware,* if a disaster destroys the primary site or renders it inoperable. A non-disaster problem, such as a corruption of a key client database, is not a disaster in this sense of the term, unless processing must be resumed at a different location and on different hardware.

1.2.6 Differentiating Business Continuity from Disaster Recovery

Strictly speaking, Disaster Recovery is the ability to recover a data center at a different site if a disaster destroys the primary site or otherwise renders it inoperable. It is only one component of an *overall* Business Continuity plan.

The Business Continuity plan has a much larger focus and includes business processes, such as a crisis management plan, and human resources management, in addition to IT recovery.

1.3 Business Continuity: Why now

The fundamental question about any Business Continuity requirement is: *Why now?* With all of the budget and time pressures of the modern world, what business issues determine the urgency of implementing Business Continuity?

We suggest the questions to consider in the following sections.

1.3.1 If you are down, what do you lose

The first step is to identify what your business stands to lose in the event of an outage. A partial list of possible impacts to the business in the event of an unplanned outage follows:

- Lost revenue, loss of cash flow, and loss of profits
- Loss of clients (lifetime value of each) and market share
- Fines, penalties, and liability claims for failure to meet regulatory compliance
- Lost ability to respond to subsequent marketplace opportunities
- Cost of re-creation and recovery of lost data
- Salaries paid to staff unable to undertake billable work
- Salaries paid to staff to recover work backlog and maintain deadlines
- Employee idleness, labor cost, and overtime compensation
- Lost market share, loss of share value, and loss of brand image

1.3.2 What are the most likely causes of outages

Next, assess what are the most likely causes of outages for your organization. Organize them into a priority list and a risk assessment. Which of these should or can you protect against?

These can comprise a large variety of components, both business and IT. Some components (of many) that could cause business outages are:

- Servers, storage (either disk or tape), and software (either database or application)
- Network components, telecom providers, and access to call centers
- Power grid and physical infrastructure damage (water and fire)

- Logical data corruption (unintentional, intentional due to error, or virus)

1.3.3 Assess where you are today

Third, do an assessment of where you stand today regarding your current processes, procedures, and management infrastructure. In particular, pay attention to:

- What tool sets for IT Business Continuity do you already have? What software, information flows, management, and automation processes? How effective are they?
- What are your problem and change management procedures? Are they rigorous enough, or will they need to be upgraded? How adaptable and responsive are they?
- What are your current monitoring tools for the status of your IT and business systems?
- What is the physical infrastructure that you have today? Where are the important points, and where will you need to recover them? What is the current size of your business and infrastructure? What is the growth rate of each of the components?
- What are the current and future regulations that you will need to comply with? What is your audit trail and audit proof capability?
- What is your recovery time today, if you had an outage? How would you accomplish this goal?
- How often have you successfully tested your recovery? How long ago was that?

1.3.4 Determine what your metrics of success should be

Fourth, define what the metrics for success in Business Continuity would be:

- What is the desired uptime that you need on a weekly, monthly, and annual basis?
- What is the desired system response time that you will need, even during a recovery?
- What is your desired ability during recovery for IT transactions per minute, hour, and day?
- What is your Recovery Time Objective? What amount of data can you afford to recreate?
- What is your desired budget? Who pays for that budget? Is it part of a larger project?

1.3.5 Final step: Decide where you want to go

Finally, once we have the answers above, we can compare our answers from where we are with where we want to go. We can then see what the gap is, and decide where we need to go.

1.4 Summary

Each organization has unique business processes necessary for its survival.

The primary leverage to optimize a cost-effective, best Business Continuity solution is in recognizing the business justifications that were discussed in this chapter, balanced with the current status, the desired future status, and affordability.

From a cost optimization standpoint, we will leverage the fact that applications and business processes are not homogeneous across the entire enterprise. This fundamental reality is a central premise of this IBM Redbook: *One size or one solution definitely does not fit all*, either within an organization or within their data.

In following chapters, we will discuss information leading to selecting an optimized IT Business Continuity solution. We introduce a roadmap to Business Continuity (the IBM System Storage Resiliency Portfolio is designed to address this roadmap). We also review a Business Continuity Solution Selection Methodology, by which we place applications and data in tiers.

These approaches can yield an optimized, tiered solution architecture that can maximize Business Continuity coverage, for a given amount of budget and investment.

2

Selecting Business Continuity solutions

In this chapter, we review the underlying principles for the selection of an *IT infrastructure* Business Continuity solution. With this information, you will be able to better evaluate and compare what is needed for your desired level of IT Business Continuity.

With those concepts in place, we will then discuss a roadmap to IT Business Continuity.

We will introduce the System Storage Resiliency Portfolio, which contains the IBM strategy, products, and services that can address the roadmap to IT Business Continuity, and which provides a full solution set to address a wide spectrum of Business Continuity solution requirements.

We will describe the fundamental principles related to defining the IT infrastructure requirements for the most cost-effective solution at various levels of recovery, by defining key Business Continuity terms, such as Recovery Time Objective, Recovery Point Objective, and the tiers of Business Continuity.

2.1 Roadmap to IT Business Continuity

To understand how to optimize a Business Continuity solution, we begin by walking through a roadmap to IT Business Continuity. We relate each of the necessary components in the IT infrastructure to each other, in a logical sequence, for building a Business Continuity solution. The roadmap is built from the bottom up, as shown in Figure 2-1.

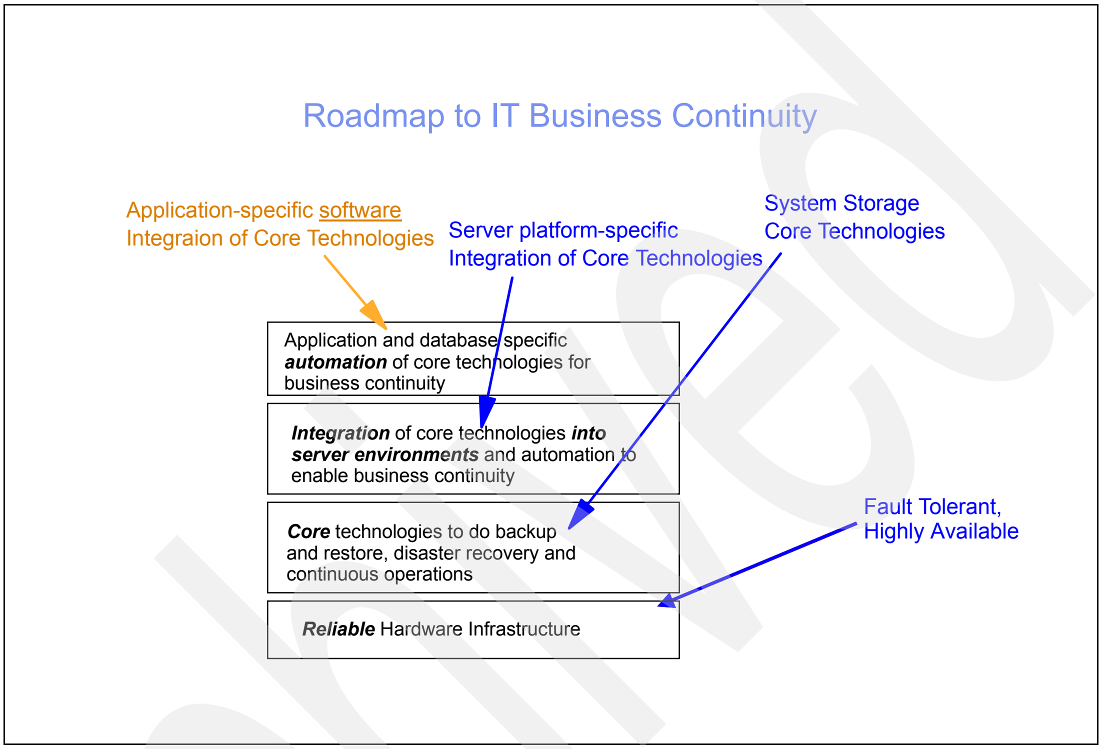

Roadmap to IT Business Continuity

Application-specific software
Integraion of Core Technologies

Server platform-specific
Integration of Core Technologies

System Storage
Core Technologies

Fault Tolerant,
Highly Available

Application and database specific ***automation*** of core technologies for business continuity

Integration of core technologies ***into server environments*** and automation to enable business continuity

Core technologies to do backup and restore, disaster recovery and continuous operations

Reliable Hardware Infrastructure

Figure 2-1 Roadmap to IT Business Continuity

Starting from the bottom and working up, we build a solution, layer upon layer, step by step.

We first start with:

1. Reliable hardware infrastructure: Storage devices, servers, and SANs. The base IT component(s) of hardware and microcode reliability must be adequate, with the objective being no single-points-of-failure. This layer can be considered prevention of outages by assuring that the base components are reliable.

Next, we would proceed to selecting and implementing:

2. Core Technologies: Advanced copy technologies that enable reliability, redundancy, and Business Continuity, at the *solution and system operations* level (rather than at the component level).

   Core technologies can be resident in server operating systems, storage systems, storage software, file systems, and other components. These core technologies typically include advanced copy functions for making point-in-time copies, remote replicated copies, file backup and restore, and volume images.

   They also include management software and services for enterprise-wide policy-based management of backup copies, administration of the storage replication services, and other necessary infrastructure management functions.

Upon the base of core technologies, we add:

3. Server Integration: Each operating system platform has its own requirements for how to access and use the data, volume copies, and replication recovery copies. The next layer provides operating system-specific tasks and integration that are needed to make use of the copied data. Administrators can take care of these platform-specific tasks by sitting at a console, but it is preferable to have automation that will take care of these steps without administrator intervention. Administrators expect the automation to be in a format familiar to them, in other words, platform-specific interfaces and commands for Windows®, AIX®, z/OS®, Linux®, or whatever platform is being used.

Finally, after server integration, we add:

4. Application integration: Applications need to have specific integration into the previous layers to take full advantage of the underlying functions. This takes the form of integrated application commands, familiar to the application users, that can transparently exploit advanced server or storage capabilities. This functionality often includes coordinating core technology operations across a large number of multiple LUNs, datasets, and objects, disk systems, and servers that may make up an application system. Application-aware automation should know the specific steps that are required for these databases to perform necessary steps in the right order, check for successful completion, and provide automation and logic for these multiple steps.

2.2 The System Storage Resiliency Portfolio

Designed to match this roadmap to IT Business Continuity is an architected portfolio of IBM products called the System Storage Resiliency Portfolio. This architecture maps the portfolio of products according to their function, as shown in Figure 2-2.

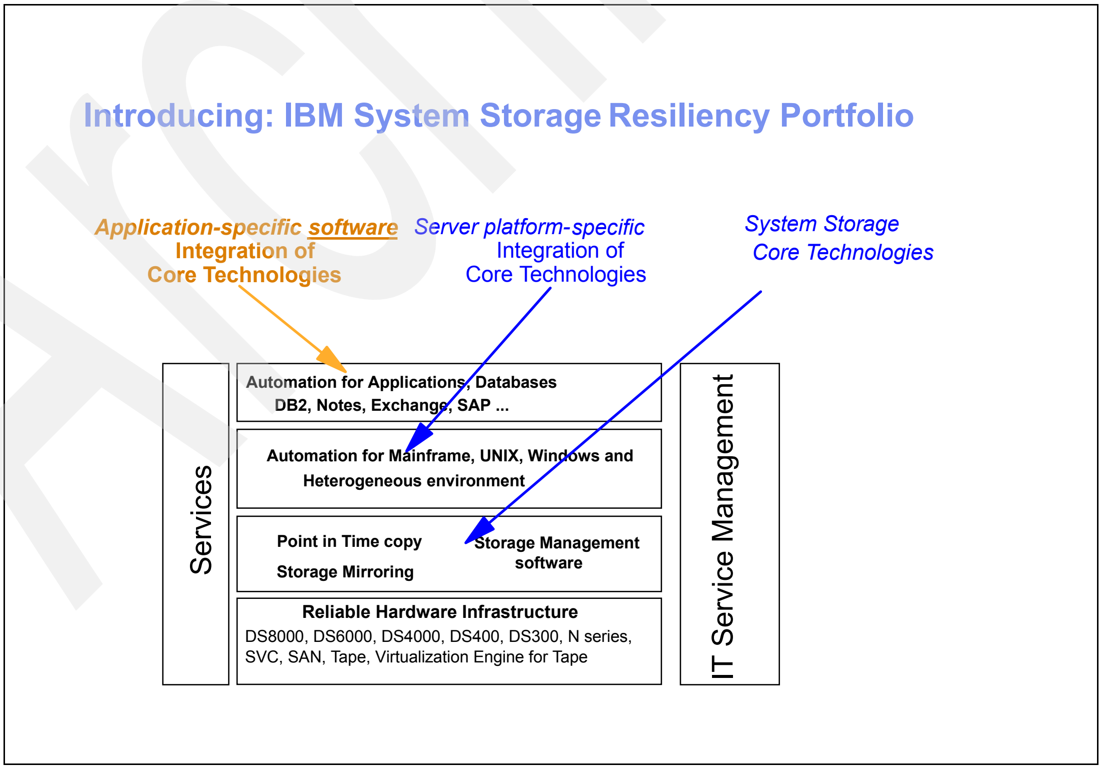

Introducing: IBM System Storage Resiliency Portfolio

Application-specific software Integration of Core Technologies

Server platform-specific Integration of Core Technologies

System Storage Core Technologies

Services

Automation for Applications, Databases DB2, Notes, Exchange, SAP ...

Automation for Mainframe, UNIX, Windows and Heterogeneous environment

Point in Time copy

Storage Mirroring

Storage Management software

Reliable Hardware Infrastructure

DS8000, DS6000, DS4000, DS400, DS300, N series, SVC, SAN, Tape, Virtualization Engine for Tape

IT Service Management

Figure 2-2 System Storage Resiliency Portfolio

In addition to the products being mapped to the roadmap, *Services and Skills* are added, which are available through IBM Business Partners or IBM Global Services. Services and Skills provide definition and implementation of processes, such as component failure impact analysis, business impact analysis, definition of the recovery time objectives, implementation of problem management, change management, system management, monitoring, integration, and testing.

The System Storage Resiliency Portfolio architecture also has a strategic purpose. The architecture is designed to provide the *autonomic* foundation for future On Demand Business Continuity solutions. The System Storage Resiliency Portfolio architecture provides for an open standards-based set of *consistent interfaces*, between the Resiliency Family layers, providing a pathway for On Demand dynamic discovery and dynamic exploitation of new function.

Let us see what is in each of the product layers of the System Storage Resiliency Portfolio.

2.2.1 Reliable hardware infrastructure layer

As discussed in the roadmap to IT Business Continuity, this layer provides the fault-tolerant and highly-available infrastructure. Functionalities such as storage RAID, dual power supplies, dual internal controllers, redundant internal component failover, and so on, all reside in this layer.

IBM System Storage products, designed to provide robust reliability, that reside in this layer include:

- IBM DS Family disk: DS8000™, DS6000™, DS4000™, DS400, and DS300
- IBM NAS: N series
- IBM Storage Virtualization: SAN Volume Controller
- IBM Storage Area Network: switches and directors
- IBM Tape: IBM tape libraries and virtual tape products, including Virtualization Engine™ for Tape

2.2.2 Core technologies layer

Core technologies provide advanced copy functions for making point-in-time copies, remote replicated copies, file backup and restore, volume images, and other advanced replication services. Core technology functionality examples include (but are not limited to):

- Non-disruptive Point-in-Time copies: FlashCopy or SnapShot
- Synchronous storage mirroring at metropolitan distances: Metro Mirror or SyncMirror®
- Asynchronous storage mirroring at long distances: Global Mirror or SnapMirror®

Specific product examples of core technologies on storage devices include (but are not limited to):

- FlashCopy: DS8000, DS6000, DS4000, ESS, DS400, DS300, and SAN Volume Controller
- Metro Mirror: DS8000, DS6000, DS4000, SAN Volume Controller, and Virtual Tape Server
- Global Mirror: DS8000, DS6000, DS4000, SAN Volume Controller, and Virtual Tape Server
- z/OS Global Mirror (Extended Remote Copy XRC): DS8000

> **Note:** On IBM N series, the point-in-time copy function is called SnapShot. Synchronous local mirroring is called SyncMirror, and remote mirroring is called SnapMirror.
>
> The implementations of copy functions across different hardware products can differ; for example, FlashCopy on a DS8000 or DS6000 works slightly differently than FlashCopy on a SAN Volume Controller, or on a DS4000.

Also within this layer is management software and services for disaster recovery planning, enterprise-wide policy-based management of backup copies, and administration and management of the storage-provided replication services:

- Tivoli Storage Manager for application-aware, online, backups of popular applications, which include DB2®, Oracle®, SAP®, WebSphere®, Microsoft® SQL Server, Microsoft Exchange, and Lotus® Domino®
- TotalStorage® Productivity Center for enterprise-wide server and policy administration, LAN-free and Server-free backup and restore, and replication management
- DFSMS™ family of offerings for z/OS on IBM System z™ servers

2.2.3 Platform-specific integration layer

The next layer is the automation and integration of platform-specific commands and procedures to support use of the core technologies.

This layer contains specific integrations by operating system. Currently, four categories of operating systems with integrated automation available are:

- System z
- System p™
- Windows
- Heterogeneous operating systems

Each operating system has specific integration available within the Resiliency Family. Each category has specific products for that operating system. A brief overview is below; each of these solutions will be described in more detail later in this IBM Redbook.

- System z

  Server integration and automation for Business Continuity core technologies are provided by Geographically Dispersed Parallel Sysplex™ (GDPS). GDPS is used in System z environments worldwide to provide automated management of Disaster Recovery and for automated management for high availability in Parallel Sysplex environments. GDPS supports Metro Mirror, GDPS HyperSwap™, Global Mirror, and z/OS Global Mirror.

- AIX and System p

  AIX High Availability Clustered Multi Processors - eXtended Distance (HACMP/XD) is the high availability, clustered server support for System p processors and AIX applications. HACMP™ provides AIX cluster failover, application restart, network takeover, workload management, and automated failback. HACMP/XD supports Metro Mirror on DS6000, DS8000, and SAN Volume Controller.

- Heterogeneous open systems servers

  For the heterogeneous server environment, TotalStorage Productivity Center for Replication provides an enterprise-wide disk mirroring control and management integration package for open systems data on DS6000, DS8000, and SAN Volume

Controller. Without requiring involvement from the application servers, TPC for Replication manages the disk mirroring and automates the consistent recovery point for recovery.

More information about these offerings are in subsequent chapters.

2.2.4 Application-specific integration layer

This layer contains automation and integration packages that provide application-specific commands and procedures to support the use of core technologies, and where necessary, interoperates with the server integration layer.

There are many application-specific offerings in this layer. Some key examples of current IBM offerings include:

- Tivoli Storage Manager for Advanced Copy Services

  This is an integrated solution that combines IBM disk systems, Tivoli Storage Manager software, and the SAP databases on Oracle or DB2 UDB into an end-to-end integrated solution to make nondisruptive point-in-time backups and clones of the SAP database. This offering implements a fully automated replication process, to generate backups or clones of very large databases in minutes instead of hours. This solution eliminates the need for intermediate backups, and helps address the SAP client's key requirement for continuous application availability. It can also provide very fast restores.

- Tivoli Storage Manager for Copy Services

  This solution integrates Microsoft Windows VSS capabilities in disk systems to make fast, nondisruptive snapshots of Exchange databases.

More information about these offerings are in Chapter 5, “Backup and Restore” on page 55.

2.2.5 Summary

Services and skills tie all the technology components of the System Storage Resiliency Portfolio together, leading towards an end-to-end, automated Business Continuity solution.

The strategic value of the System Storage Resiliency Portfolio is in the architecturally defined open standards-based *interfaces* between the layers. These interfaces provide points of technology interlock, exploitation, and discovery for future *services-based* solutions.

Through the provision of an integrated set of Business Continuity solutions, the System Storage Resiliency Portfolio is designed to provide integrated solutions that provide both IT Business Continuity as well as the requisite business value.

2.3 Selecting Business Continuity solutions

There are clearly a large variety of valid IT Business Continuity products and solutions. The fundamental challenge is to select the *optimum, cost-effective blend* of all these Business Continuity products and technologies.

The common problem in the past has been a tendency to view the Business Continuity solution as individual product technologies and piece parts (see Figure 2-3 on page 17).

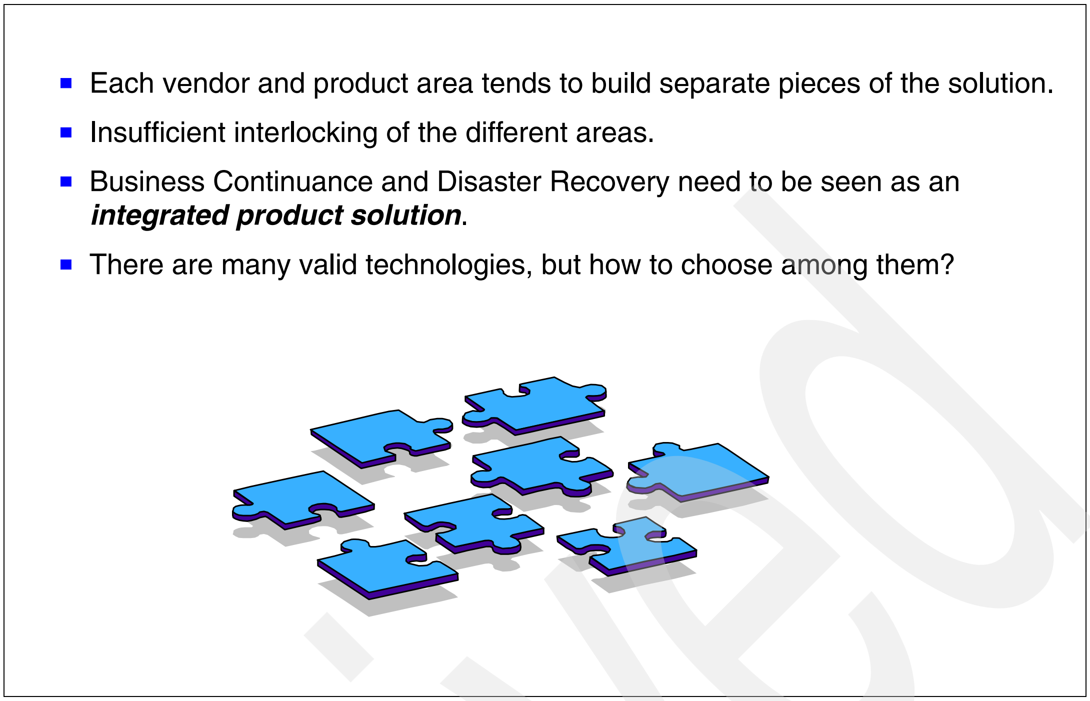

- Each vendor and product area tends to build separate pieces of the solution.
- Insufficient interlocking of the different areas.
- Business Continuance and Disaster Recovery need to be seen as an ***integrated product solution***.
- There are many valid technologies, but how to choose among them?

Figure 2-3 The challenge of selecting the best Business Continuity solution

Instead, Business Continuity solutions need to be viewed as a whole, integrated multi-product solution. In this chapter, we propose a philosophy to select a Business Continuity Solution via solution *segmentation*: using a tier method that sorts, summarizes, and organizes the various business requirements, and maps them to a solution segment based on recovery time.

2.4 The tiers of Business Continuity

We will organize the various IT Business Continuity products and solutions according to the concept of *tiers*. The concept of tiers (which is the commonly used method in today's best practices for Business Continuity solution design) is powerful and central to our selection philosophy.

Definitions:

Recovery Time Objective (RTO): How long can we afford to be without our business systems? What is our desired elapsed time to recover?

Recovery Point Objective (RPO): When our system is recovered, how much data can we afford to recreate?

The tiers concept recognizes that for any given client's Recovery Time Objective (RTO), all Business Continuity products and technologies can be sorted into a *RTO solution subset* that addresses that particular RTO range. The reason for multiple tiers is that as the RTO decreases, the optimum Business Continuity technologies for that RTO must change.

By categorizing Business Continuity technology into tiers according to RTO, we gain the ability to match our RTO time with the optimum price/performance set of technologies. While the technology within the tiers has obviously changed through time, the concept continues to be as valid today as when it was first described by the US SHARE User Group in 1988.

The tiers chart in Figure 2-4 gives a generalized view of some of today's IBM Business Continuity technologies by tier.

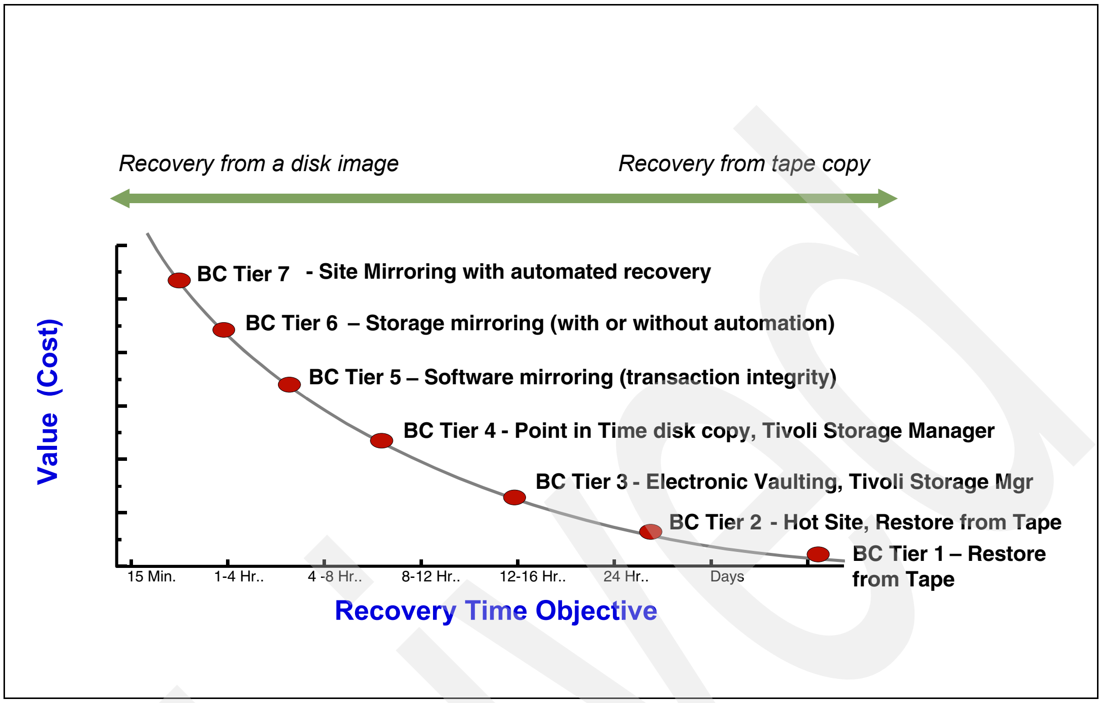

Figure 2-4 Tiers of Business Continuity: technology changes as tiers increase

The concept of the tiers chart continues to apply even as the scale of the application(s) changes. That is, the particular RTO values may increase or decrease, depending on the scale and criticality of the application. Nevertheless, the general relative *relationship* of the various tiers and Business Continuity technologies to each other remains the same. In addition, although some Business Continuity technologies fit into multiple tiers, clearly there is not one Business Continuity technology that can be optimized for all the tiers.

Your technical staff can and should, when appropriate, create a specific version of the tier chart for your particular environment. Once the staff agrees upon what tier(s) and corresponding RTO a solution delivers for your enterprise, then the Business Continuity technical evaluation and comparisons are much easier, and the technology alternatives can be tracked and organized in relation to each other.

2.5 Blending tiers into an optimized solution

To use the tiers to select the best Business Continuity solution, we first recognized that a blend of Business Continuity tiers is the practical method for cost-optimizing the entire solution. The most common result, from an enterprise standpoint, is a strategic architecture of three tiers. Three tiers generally appear as an optimum number, because at the enterprise level, two tiers generally are insufficiently optimized (in other words, overkill at some point and underkill at others), and four tiers are more complex, but usually do not provide enough additional strategic benefit.

To use the tiers to derive a blended, optimized Business Continuity architecture, we suggest:

1. Categorize the business' entire set of applications into three bands: Low Tolerance to Outage, Somewhat Tolerant to Outage, and Very Tolerant to Outage. Of course, while

some applications that are not in of themselves critical, they do feed the critical applications. Those applications would need to be included in the higher tier.

2. Within each band, there are tiers. The individual tiers represent the major Business Continuity *technology choices* for that band. It is not necessary to use all the tiers, and of course, it is not necessary to use all the technologies.
3. Once we have segmented the applications (as best we can) into the three bands, we usually select one best strategic Business Continuity methodology for that band. The contents of the tiers are the *candidate technologies* from which the strategic methodology is chosen.

2.5.1 Three Business Continuity solution segments

IBM Business Continuity solutions in the System Storage Resiliency Portfolio have been segmented for you into these three bands; that segmentation is shown in Figure 2-5.

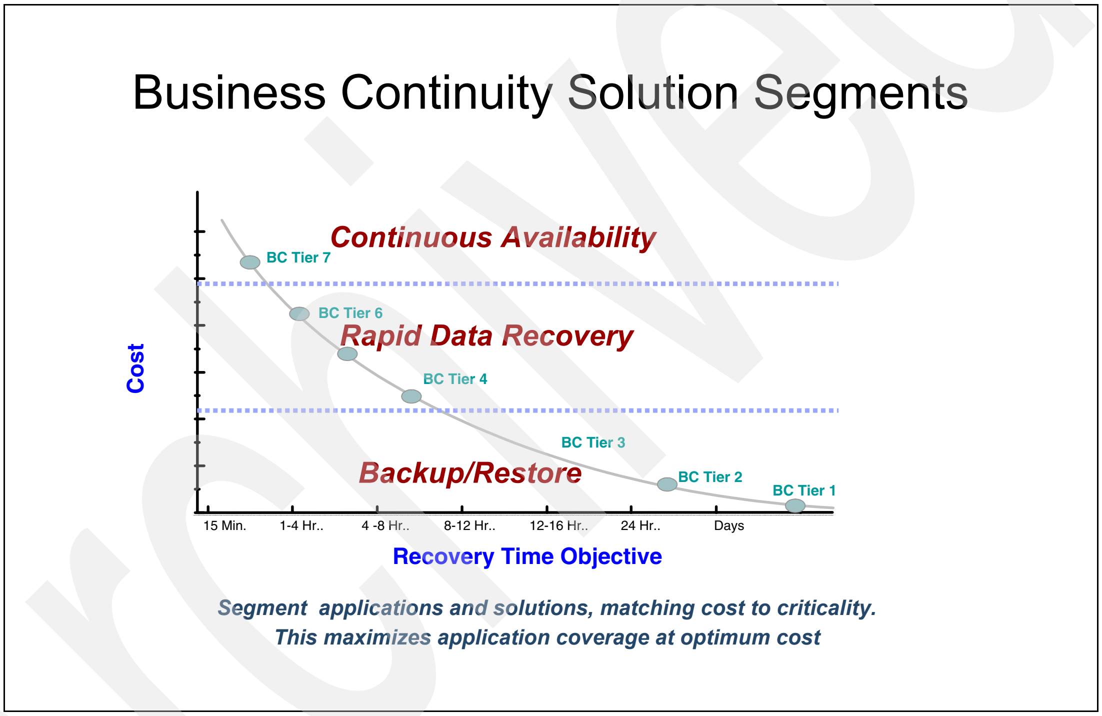
Business Continuity Solution Segments

Segment applications and solutions, matching cost to criticality.
This maximizes application coverage at optimum cost

Figure 2-5 Three Business Continuity solution segments

2.6 Select a solution by segment

Once you have decided upon your desired Recovery Time Objective, find the appropriate solution segment band in the chart above.

IBM System Storage Business Continuity solutions in the *Continuous Availability* solution segment are Tier 7 solutions. They include (but are not limited to):

- System p: (AIX HACMP/XD)
- System z: GDPS

IBM System Storage Business Continuity solutions in the Rapid Data Recovery solution segment are Tier 4 to 6 solutions. They include (but are not limited to):

- TotalStorage Productivity Center for Replication
- Heterogeneous open system disk vendor mirroring: IBM SAN Volume Controller Metro Mirror
- System z: GDPS HyperSwap Manager

IBM TotalStorage Business Continuity solutions in the Backup/Restore solution segment are Tiers 4 to 1. They include (but are not limited to):

- Tivoli Storage Manager and all its associated products
- IBM System Storage DR550
- SMS, DFSMShsm™: DFSMSdss™ for z/OS and DDR for VM volumes (DDR does not invoke FlashCopy)

In this IBM Redbook and in the companion IBM Redbooks *IBM System Storage Business Continuity: Part 1 Planning Guide*, SG24-6547 and *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548, all solutions are identified by tier level. The above set of solutions will be discussed in more detail in subsequent chapters.

2.6.1 Refine with a detailed evaluation team

Having identified a preliminary set of valid candidate Business Continuity solutions and technologies via this segmentation, expect to refine and identify your actual solution by evaluating these candidate solutions with a skilled evaluation team, made up of members qualified to contrast and compare the identified solutions in detail.

The evaluation team will in all likelihood need to further configure the candidate solutions into more detailed configurations to complete the evaluation. In the end, that team will still make the final decision as to which of the identified options (or the blend of them) is the one that should be selected.

2.6.2 Value of Business Continuity solution selection via segmentation

As simple as this sounds, this process of quickly identifying proper candidate Business Continuity solutions for a given set of RTO requirements *is* of significant value.

Much less time and skill is necessary to reach this preliminary solution identification in the evaluation cycle than would otherwise be experienced. This methodology supports the Business Continuity practice of segmenting the Business Continuity architecture into three blended tiers (and therefore three tiers of solutions).

To identify the solutions for the other bands of solutions, you would simply revisit this philosophy and give the lower RTO Level of Recovery for those lower bands and applications, and you would find the corresponding candidate solution technologies in the appropriate (lower) solution segments.

This philosophy is not meant as a substitute for Business Continuity skill and experience, and it is not possible for the philosophy to be a perfect decision tree. While there clearly will be ambiguous circumstances (for which knowledgeable Business Continuity experts will be required), this solution segmentation and selection philosophy provides an efficient process by which the initial preliminary Business Continuity solution selection can be consistently performed.

Part 2

Business Continuity solution offerings

In this part, we describe the Business Continuity solutions and offerings in each of the segments:

- Continuous Availability
- Rapid Data Recovery
- Backup and Restore

3

Continuous Availability

Today’s enterprises can no longer afford planned or unplanned system outages. Even a few minutes of application downtime can cost big financial losses, erode client confidence, damage brand image, and present public relations problems. The on demand data center must be resilient enough to handle the ups and downs of the global market, and it must manage changes and threats with consistent availability and security and privacy, around the world and around the clock.

Continuous Availability solutions are integrations of servers, storage, software and automation, and networking. Most of the solutions we describe are based on some form of operating system server clustering to provide application availability. When an application failure is detected, a Continuous Availability solution will perform a predefined set of tasks required to restart the application on another server.

This chapter describes Continuous Availability in the following environments:

- Geographically Dispersed Parallel Sysplex (GDPS)
- Geographically Dispersed Open Clusters (GDOC)
- HACMP/XD
- Continuous Availability for MaxDB and SAP liveCache
- Copy Services for System i™
- Metro Cluster for N series

For general information about Continuous Availability solutions, see the Web site:

http://www-03.ibm.com/servers/storage/solutions/business_continuity/continuous_availability/technical_details.html.

3.1 Geographically Dispersed Parallel Sysplex (GDPS)

GDPS is a family of IBM Global Services offerings for single site or a multi-site application availability, providing an integrated, end-to-end solution for enterprise IT Business Continuity, integrating software automation, servers, storage, and networking.

GDPS control software manages the remote copy configuration and storage systems, automates System z operational tasks, manages and automates planned reconfigurations, and does failure recovery from a single point of control. GDPS offerings are segmented as Continuous Availability in most cases, though some may be configured, optionally, as rapid recovery, as shown in Figure 3-1 (the GDPS solution has components in the areas denoted by dark shading).

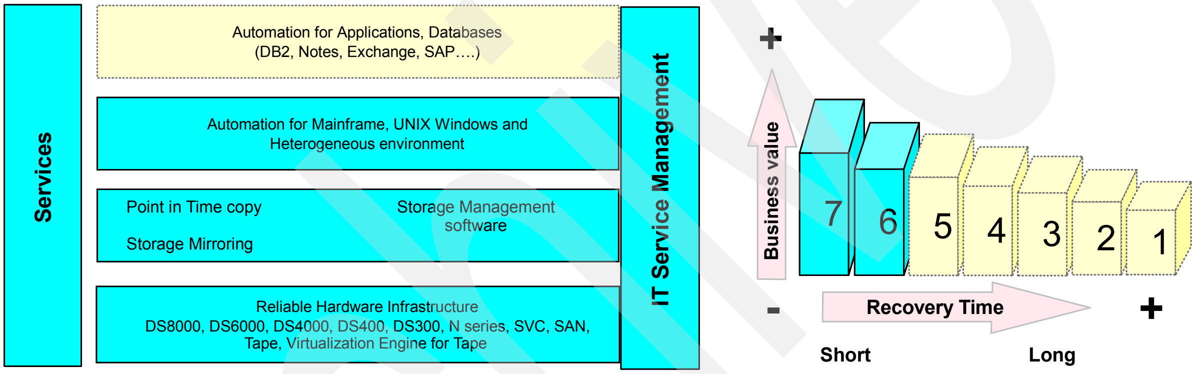
Services
Automation for Applications, Databases
(DB2, Notes, Exchange, SAP….)
Automation for Mainframe, UNIX Windows and
Heterogeneous environment
Point in Time copy
Storage Mirroring
Storage Management
software
Reliable Hardware Infrastructure
DS8000, DS6000, DS4000, DS400, DS300, N series, SVC, SAN,
Tape, Virtualization Engine for Tape
IT Service Management
+
Business value
7
6
5
4
3
2
1
-
Recovery Time
+
Short
Long

Figure 3-1 The positioning of the GDPS family of solutions

GDPS supports both synchronous (Metro Mirror) and asynchronous (z/OS Global Mirror and Global Mirror) remote copy, as well as TS7700 Grid remote copy for tape data in GDPS/PPRC and GDPS/XRC environments.

The GDPS solution is an open technology: it works with any vendor’s disk system that meets the specific functions of the Metro Mirror, z/OS Global Mirror, or Global Mirror architectures required to support GDPS functions.

The GDPS family of System z Business Continuity technologies are:

- GDPS/PPRC solutions, based on IBM System Storage Metro Mirror (Metro Mirror was formerly known as Peer-to-Peer Remote Copy (PPRC)), including:
  - GDPS/PPRC
  - GDPS/PPRC HyperSwap Manager
  - RCMF/PPRC, a remote copy management solution for Metro Mirror

- Asynchronous GDPS technologies, based on System Storage z/OS Global Mirror (z/OS Global Mirror was formerly known as Extended Remote Copy (XRC) and Global Mirror, including:
  - GDPS/XRC
  - GDPS/Global Mirror
  - RCMF/XRC, a remote copy management solution for ZGM

Some of these GDPS offerings can also be combined into three site GDPS solutions, providing Continuous Availability in a multi-site environment.

GDPS/PPRC overview

The physical topology of a GDPS/PPRC, shown in Figure 3-2, consists of a System z base or Parallel Sysplex cluster spread across two sites separated by up to 100 kilometers or 62 miles of fibre, with one or more z/OS systems at each site.

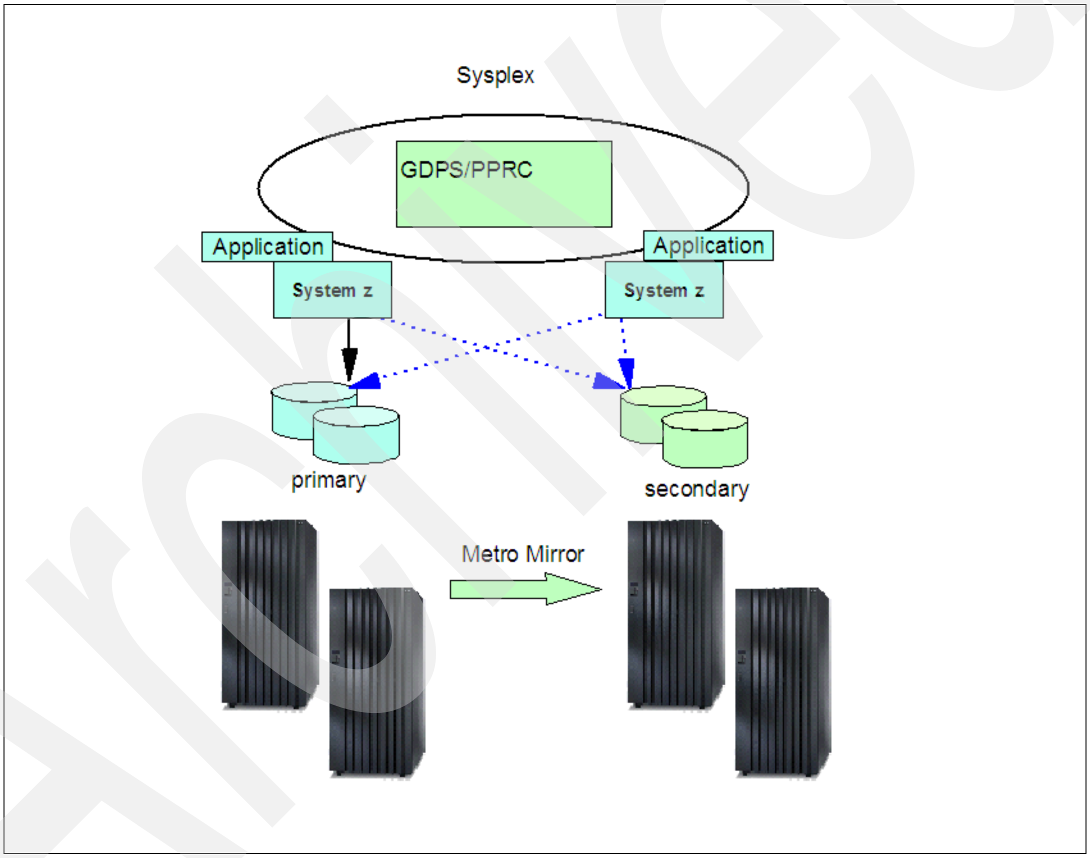

Figure 3-2 GDPS/PPRC topology

GDPS/PPRC provides the ability to perform a controlled site switch for both planned and unplanned site outages, with no or minimal data loss, maintaining full data integrity across multiple volumes and storage subsystems and the ability to perform a normal Data Base Management System (DBMS) restart (not DBMS recovery) in the second site. GDPS/PPRC is application independent, and therefore can cover the client's complete application environment.

Near Continuous Availability of data with HyperSwap

GDPS with Metro Mirror provides HyperSwap functionality, which can help significantly reduce the time needed to switch to the secondary set of disks while keeping the z/OS systems active, together with their applications. In so doing, HyperSwap broadens the near Continuous Availability attributes of GDPS/PPRC by extending the Parallel Sysplex redundancy to disk systems.

GDPS/PPRC management of System z operating systems

In addition to managing images within the base or Parallel Sysplex cluster, GDPS can manage a client's other System z production operating systems; these include z/OS, Linux for System z, z/VM®, and VSE/ESA™. For example, if the volumes associated with the Linux for System z images are mirrored using PPRC, GDPS can restart these images as part of a planned or unplanned site reconfiguration. The Linux for System z images can either run as a logical partition (LPAR) or as a guest under z/VM.

GDPS/PPRC management for open systems LUNs (Logical Unit Number)

GDPS/PPRC can manage a common Metro Mirror Consistency Group for both System z and open systems storage, thus providing data consistency across both z/OS and Open Systems data. This allows GDPS to be a single point of control to manage business resiliency across multiple tiers in the infrastructure, improving cross-platform system management and business processes.

GDPS/PPRC Multi-Platform Resiliency for System z

GDPS/PPRC Multi-Platform Resiliency for System z is especially valuable for clients who share data and storage systems between z/OS and z/VM Linux guests on System z, for example, an application server running on Linux on System z and a database server running on z/OS. GDPS/PPRC provides the reconfiguration capabilities for the Linux on System z servers and data in the same manner as for z/OS systems and data.

GDPS/PPRC summary

GDPS/PPRC can provide a solution for many categories of System z clients, including (but not limited to):

- Clients who can tolerate an acceptable level of synchronous disk mirroring performance impact (typically, an alternate site at metropolitan distance, can be up to 100 KM)
- Clients that need as close to near zero data loss as possible
- Clients that desire a fully automated solution that covers the System z servers, System z Coupling Facilities, System z reconfiguration, and so on, in addition to the disk recovery

GDPS/PPRC can feature:

- A highly automated, repeatable site takeover managed by GDPS/PPRC automation
- High performance synchronous remote copy
- Hardware data loss zero or near zero
- Data consistency and data integrity assured to insure fast, repeatable database restart
- Support for metropolitan distances
- Automation of System z Capacity Back Up (CBU)
- Single point of control for disk mirroring and recovery
- Support for consistency across both open systems and system z data
- Support of the GDPS HyperSwap functionality

GDPS/PPRC HyperSwap Manager overview

The GDPS/PPRC HyperSwap Manager solution is a subset and affordable entry point for the full GDPS/PPRC solution, providing a rapid data recovery solution for enterprise disk-resident data.

GDPS/PPRC HyperSwap Manager (GDPS/PPRC HM) includes the HyperSwap management and Metro Mirror management capabilities.

GDPS/PPRC HyperSwap Manager provides either of these configurations:

1. Near Continuous Availability of data within a single site

   A Parallel Sysplex environment reduces outages by replicating hardware, operating systems, and application components; however, having only one copy of the data is an exposure. GDPS/PPRC HM can provide Continuous Availability of data by masking disk outages caused by disk maintenance or failures. For example, if normal processing is suddenly interrupted when one of the disk systems experiences a hard failure, the applications are masked from this error because GDPS detects the failure and autonomically invokes HyperSwap. The production systems continue using data from the mirrored secondary volumes. Disk maintenance can also be similarly performed without application impact by executing the HyperSwap command.

2. Near Continuous Availability of data and DR solution at metro distances

   In addition to the single site capabilities, in a two site configuration, GDPS/PPRC HM provides an entry-level disaster recovery capability at the recovery site, including the ability to provide a consistent copy of data at the recovery site from which production applications can be restarted. The ability to simply restart applications helps eliminate the need for lengthy database recovery actions.

GDPS/PPRC HM features specially priced, limited function Tivoli System Automation and NetView® software pricing, for a more affordable entry point into the full GDPS automation software. It can be upgraded to full GDPS/PPRC at a later time, preserving the implementation investment and GDPS skills and procedures.

More information about GDPS/PPRC HM is in 4.1, “System Storage Rapid Data Recovery for System z and mixed z+Open platforms (GDPS/PPRC HyperSwap Manager)” on page 46.

3.1.1 GDPS/XRC overview

GDPS/XRC has the attributes of a *disaster recovery* technology. z/OS Global Mirror (ZGM) is a combined hardware and software asynchronous remote copy solution. The application I/O is signaled completed when the data update to the primary storage is completed. Subsequently, a DFSMSdfp™ component called System Data Mover (SDM), typically running in the recovery site (site 2), asynchronously offloads data from the primary storage subsystem's cache and updates the secondary disk volumes.

GDPS/XRC can provide:

- Disaster Recovery solution
- RTO between an hour to two hours
- RPO less than two minutes, typically 3-5 seconds
- Protection against localized as well as regional disasters (distance between sites is unlimited)
- Minimal remote copy performance impact
- Support for z Linux volumes as well as other types of volumes

GDPS/XRC topology

The physical topology of a GDPS/XRC, shown in Figure 3-3, consists of production system(s) which could be a single system, multiple systems sharing disk, or a base or Parallel Sysplex cluster. The recovery site can be located at a virtually unlimited distance from the production site.

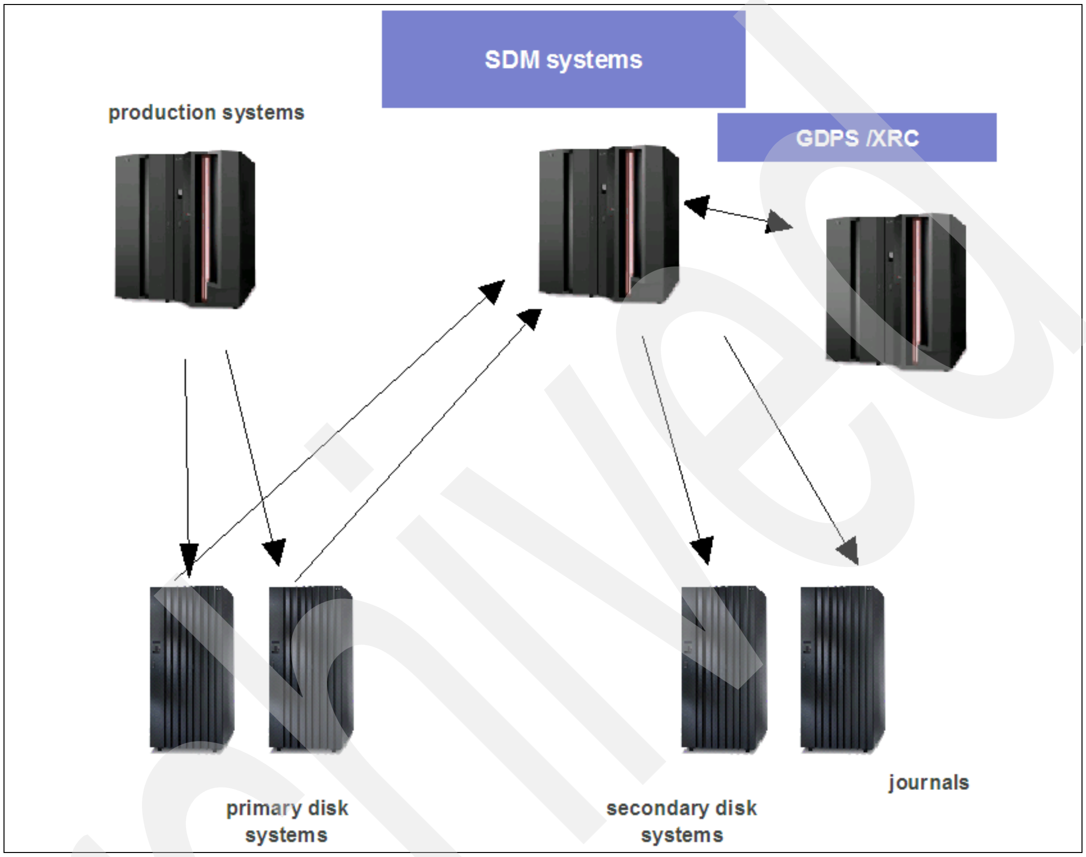

Figure 3-3 GDPS/XRC topology

GDPS/XRC provides a single, automated solution to dynamically manage disk and tape mirroring to allow a business to attain “near transparent” Disaster Recovery with minimal data loss. GDPS/XRC can perform a controlled site switch for an unplanned site outage, maintaining data integrity across multiple volumes and storage subsystems and the ability to perform a normal Data Base Management System (DBMS) restart (not DBMS recovery) in the recovery site. GDPS/XRC is application independent and therefore can cover the client's complete application environment.

3.1.2 GDPS/Global Mirror (GDPS/GM) overview

IBM System Storage Global Mirror is an asynchronous mirroring solution that can replicate both System z and open systems data.

GDPS/GM provides a link into the System z environment in order to enhance the remote copy interface for more efficient use of mirroring with fewer opportunities for mistakes, with the automation and integration necessary to perform a complete Disaster Recovery with minimal human intervention.

GDPS/GM can provide:

- Disaster Recovery technology
- RTO between an hour to two hours
- RPO less than 60 seconds, typically 3-5 seconds
- Protection against localized or regional disasters (distance between sites is unlimited)
- Minimal remote copy performance impact
- Improved and supported interface for issuing remote copy commands
- Maintaining multiple Global Mirror sessions and multiple RPOs

GDPS/GM topology

The GDPS/GM physical topology, shown in Figure 3-4, consists of production system(s), which could be a single system, multiple systems sharing disk, or a base or Parallel Sysplex cluster. The recovery site can be located at a virtually unlimited distance from the production site and, again, is not actually required to be a Parallel Sysplex cluster.

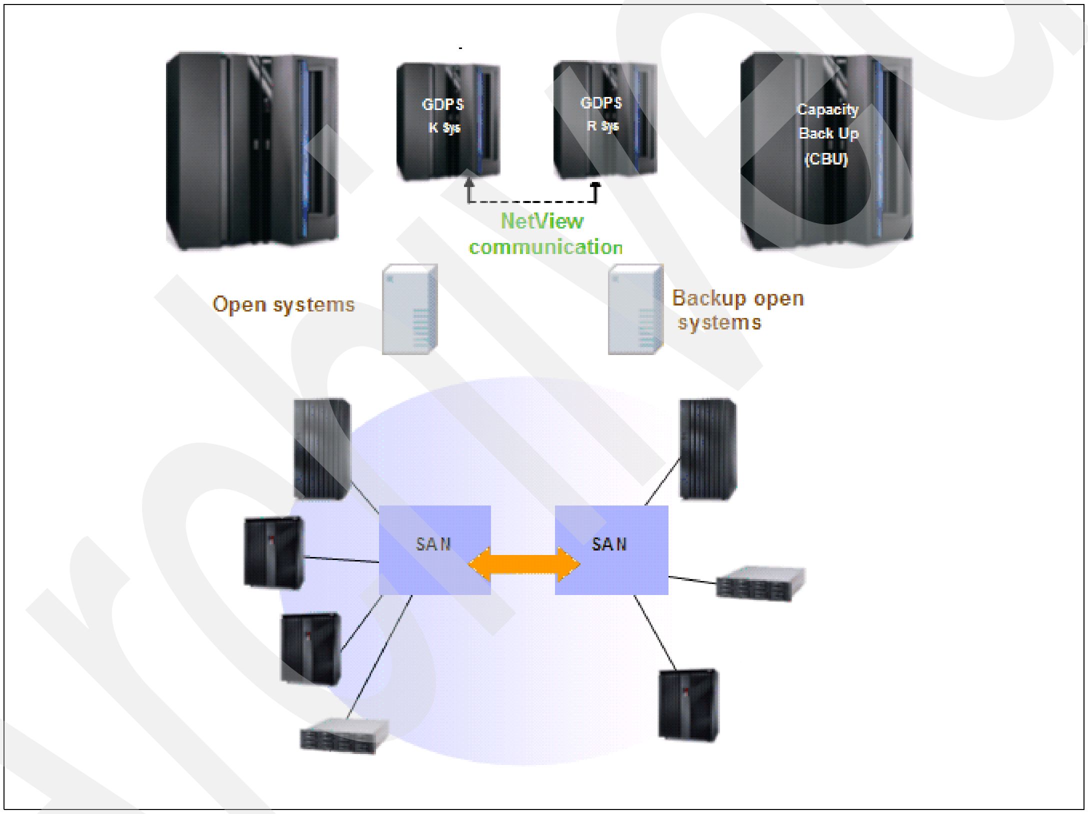

Figure 3-4 Topology of a GDPS/GM environment

The GDPS/GM environment contains a K-System, which is contained within the production site and is the primary control point for interactions with the GDPS environment.

Additionally, the GDPS/GM environment has an R-System, which works with the GDPS/GM K-System by communicating via NetView communications links. By doing this, GDPS/GM can verify the state of operations in each location and, if communications fail, GDPS/GM notifies the user of a potential problem. Importantly, the R-System LPAR does not need to be a part of the recovery site Sysplex; it can act stand alone, but can still use its automation to activate additional engines via CBU and perform reconfiguration of LPARs during a failover.

Open LUN support

As with GDPS/PPRC, GDPS/GM can maintain the mirroring environment for Open Systems data as well as System z data. Data consistency is maintained by the Global Mirror replication technology; GDPS serves as a front end to create an easy to use management environment and single point of control for all mirror related commands.

Also, as with GDPS/PPRC, GDPS/GM only maintains the mirroring environment for Open Systems. Recovery of the servers is left to the administrators.

3.1.3 GDPS three site support

There are two versions of GDPS with three site support:

- GDPS Metro Mirror and z/OS Global Mirror
- GDPS Metro Global Mirror

GDPS Metro Mirror and z/OS Global Mirror

The design is shown in Figure 3-5 with primary, secondary, and tertiary sites. Usually the primary and secondary sites are close to each other, with the tertiary site is located hundreds to thousands of kilometers away.

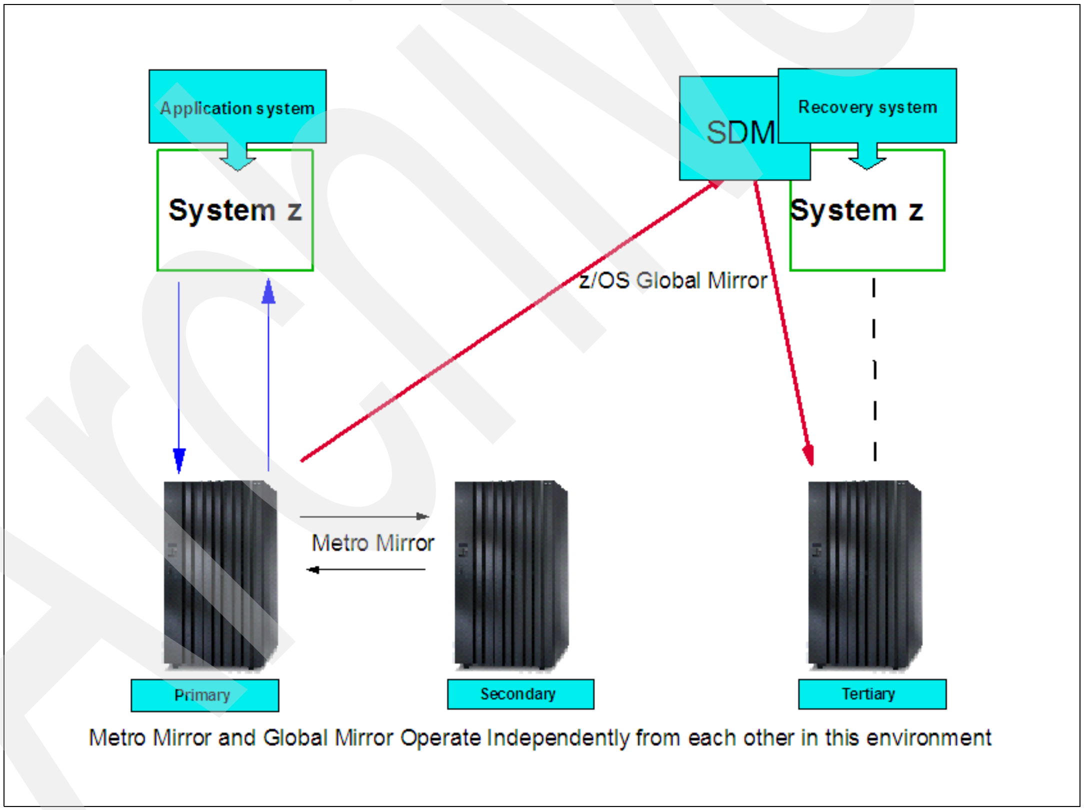

Figure 3-5 GDPS/z/OS Metro Global Mirror

Because they are based on fundamentally different disk mirroring technologies (one that is based on a relationship between two disk systems and another based on a relationship between a disk system and a z/OS server), it is possible to use Metro Mirror and z/OS Global Mirror from the same volume in a z/OS environment. This also means that GDPS Control Software can be used to enhance the solution.

In this case, GDPS/PPRC or GDPS/PPRC HyperSwap Manager would be used to protect the availability of data (through HyperSwap) on disk systems located within 100 km or within the same building. Meanwhile, GDPS/XRC would act as a control point for Disaster Recovery.

GDPS/Metro Global Mirror

The other form of three site mirroring supported by GDPS is based on an enhanced cascading technology, with Metro Global Mirror. The data passes from primary to secondary synchronously and asynchronously from the secondary to the tertiary (see Figure 3-6).

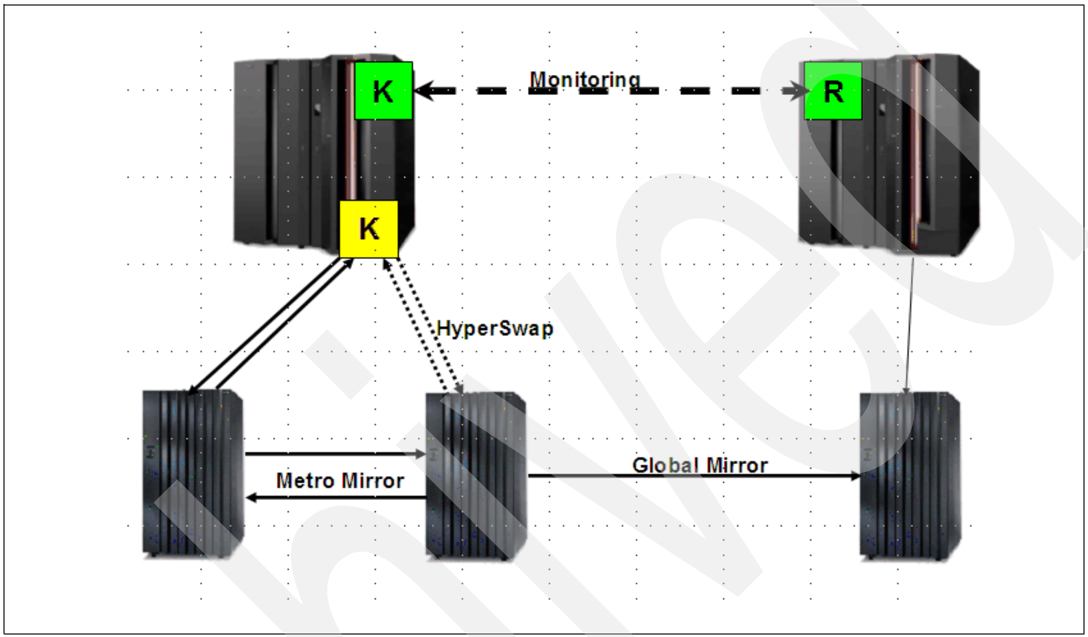

Figure 3-6 GDPS/Metro Global Mirror implementation

In this case, we can also use two forms of GDPS to support the environment. GDPS/PPRC or GDPS PPRC HyperSwap Manager will provide availability to the data via HyperSwap and metro area failover capabilities while GDPS/Global Mirror provides disaster recovery failover capabilities.

3.1.4 Global Technology Services (GTS) offerings for GDPS overview

The following GDPS services and offerings are provided by GTS.

GDPS Technical Consulting Workshop (TCW)

TCW is a two day workshop to explore the business objectives, service requirements, technological directions, business applications, recovery processes, and cross-site and I/O requirements.

GTS specialists present a number of planned and unplanned GDPS reconfiguration scenarios, with recommendations on how GDPS can assist in achieving business objectives.

Remote Copy Management Facility (RCMF)

This service includes installation of RCMF/PPRC or RCMF/XRC automation to manage the remote copy infrastructure, customization of the automation policy, and verification of the automation.

GDPS/PPRC HyperSwap Manager

This service installs and configures GDPS/PPRC HyperSwap Manager and its prerequisites. On-site planning, configuration, implementation, testing, and education are provided.

GDPS/PPRC, GDPS/XRC, and GDPS/Global Mirror

This offering provides planning, configuration, automation code customization, testing, onsite implementation assistance, and training for GDPS/PPRC or GDPS/XRC.

3.1.5 GDPS summary

GDPS is designed to provide not only near Continuous Availability benefits, but it can enhance the capability of an enterprise to recover from disasters and other failures and to manage planned exception conditions. GDPS is application independent and, therefore, can cover the client's comprehensive application environment.

GDPS can allow a business to achieve its own Continuous Availability and disaster recovery goals. Through proper planning and exploitation of IBM GDPS technology, enterprises can help protect their critical business applications from an unplanned or planned outage event.

3.1.6 Additional GDPS information

For additional information about GDPS solutions or GDPS solution components, see:

GDPS home page:

http://www.ibm.com/systems/z/gdps/

System z Business Resiliency Web site:

http://www.ibm.com/systems/z/resiliency

3.2 GDOC

Geographically Dispersed Open Clusters (GDOC) is a multivendor solution for protecting the availability of critical applications that run on UNIX®, Windows, or Linux servers. It is based on an Open Systems Cluster architecture spread across two or more sites with data mirrored between sites to provide high availability and Disaster Recovery.

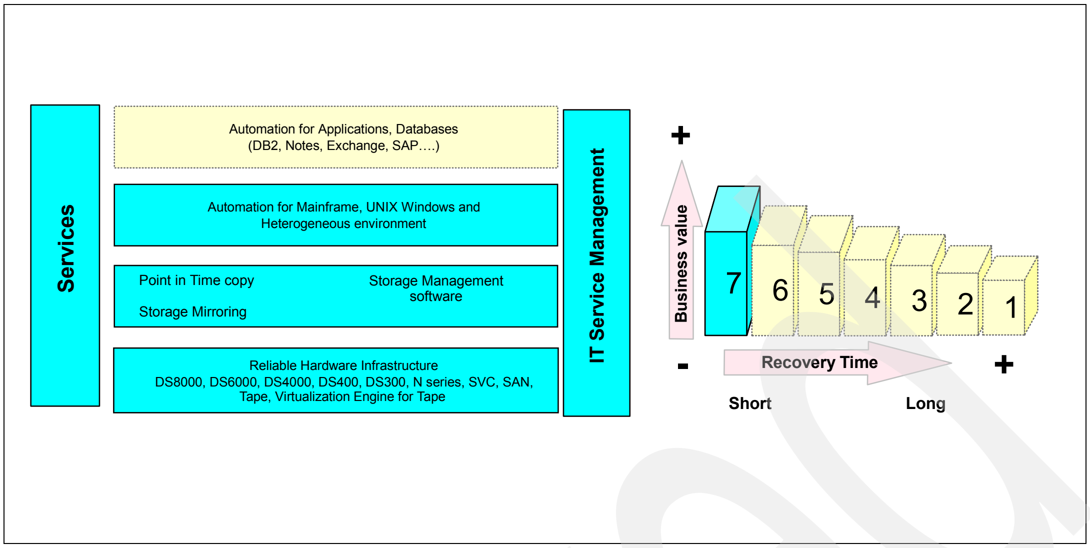
Services
Automation for Applications, Databases
(DB2, Notes, Exchange, SAP….)
Automation for Mainframe, UNIX Windows and
Heterogeneous environment
Point in Time copy
Storage Mirroring
Storage Management
software
Reliable Hardware Infrastructure
DS8000, DS6000, DS4000, DS400, DS300, N series, SVC, SAN,
Tape, Virtualization Engine for Tape
IT Service Management
+
Business value
7
6
5
4
3
2
1
-
Recovery Time
+
Short
Long

Figure 3-7 GDOC positioning within the Resiliency portfolio

Solution description

A GDOC solution consists of two components: GDOC Planning and Deployment. The solution is customized and can be implemented without disrupting existing systems or staff. This is an ideal solution when the potential for downtime or data loss dramatically impacts profits or jeopardizes the ability to conduct business.

The GDOC solution is based on the VERITAS Foundation suite and VERITAS Cluster software products. The currently supported operating system platforms are:

- IBM AIX
- Linux
- Microsoft Windows
- SUN Solaris™
- HP UX

Solution highlights

GDOC:

- Helps improve client service and employee productivity.
- Helps reduce risk of costly compliance or legal or liability issues associated with downtime.
- Protects against lost revenue, productivity, and prestige.
- Provides flexibility to transfer operations between sites with minimal disruption.
- Helps reduce costs via consistent and centralized management across platforms.
- Aims to limit training and errors through ease of use, consistency, and automation.

3.2.1 Solution components

A GDOC solution is based on a combination of IBM Global Services services together with VERITAS software. The services are split between a consulting and planning phase followed by an implementation and deployment phase.

VERITAS software components

There is no predefined set of VERITAS software components because the solution can be adapted to specific client requirements. The main components used are the following:

- VERITAS Cluster Server (VCS)
  - VERITAS Cluster Server Hardware Replication Agent for IBM PPRC
  - VERITAS Cluster Server Global Cluster Option
  - VERITAS CommandCentral Availability (was VERITAS Global Cluster Manager)
- VERITAS Storage Foundation (Optional)
  - VERITAS Volume Manager
  - VERITAS Volume Replicator
  - VERITAS File System (UNIX)

To these basic components, we can add additional software modules that support specific environments, such as SAP or DB2.

3.2.2 GDOC in greater detail

GDOC controls resource and application availability and initiates application failover to alternative servers when needed. Application availability and failover are controlled using VERITAS Cluster Server (VCS) and specialized modules for applications, such as SAP, DB2, and others. Additional cluster modules are available. An example is VERITAS Cluster Server Hardware Replication Agent for IBM PPRC, which is available to control Metro Mirror PPRC-based solutions. Site failover is controlled by VERITAS Cluster Server. VERITAS CommandCentral Availability provides cluster centralization, monitoring, and reporting capabilities. Figure 3-8 shows a diagram of the GDOC functionality.

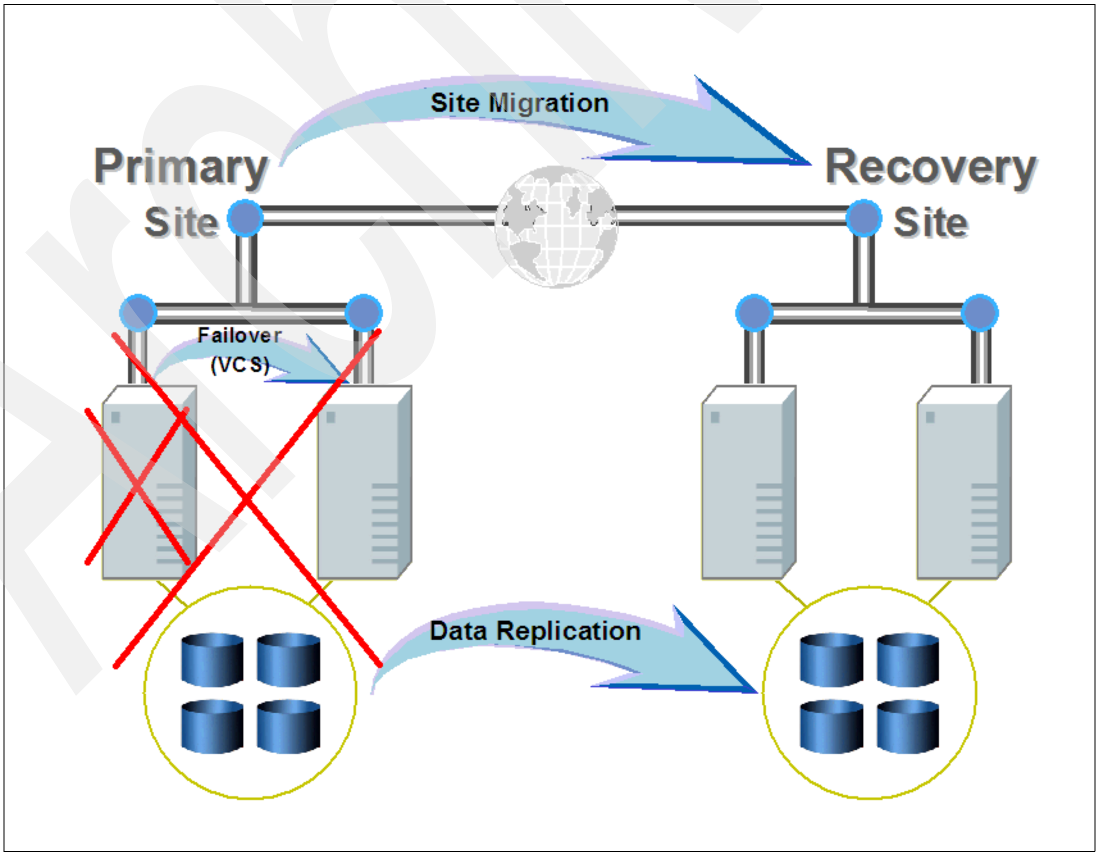

Figure 3-8 GDOC functional overview

Normally applications run on the primary site, with the application data in an external storage system. This allows the application to fail over to a secondary server at the primary site under VCS control. The data is still accessed from the primary storage system.

Data is replicated continuously between the storage systems located at the primary and secondary sites using data replication functions such as DS6000/DS8000 Metro Mirror or VERITAS Volume Replicator. VERITAS Volume Replicator maintains dependent write consistency, and write order fidelity using VERITAS terminology, at the individual server level. If you need to maintain dependent write consistency across multiple servers, then you must consolidate the storage in one or more external storage systems and use Metro Mirror or Global Mirror solutions.

The VERITAS CommandCentral Availability controls failover of applications to the secondary site. This is shown in Figure 3-8 on page 34 as a site migration. A site migration will typically require operator intervention to confirm before performing the site failover.

This solution extends the local High Availability (HA) model to many sites. Dispersed clusters and sites are linked by public carrier over a wide area network or SAN. Each site is aware of the configuration and state of all of the sites (global cluster management). Complete site failover occurs in the event of a catastrophe, and the basis for this failover is the replicated data.

3.2.3 Additional information

This solution is delivered by IBM Global Services. For more information, contact your IBM sales representative.

3.3 HACMP/XD

High Availability Cluster Multiprocessing (HACMP) XD (eXtended Distance) provides failover for applications and data between two geographically dispersed sites. It offers both High Availability and Disaster Recovery across geographically-dispersed HACMP clusters, protecting business-critical applications and data against disasters that affect an entire data center.

HACMP/XD is a Tier 7 solution when used together with the Metro Mirror feature. It can be considered an automated Tier 4 solution when HACMP/XD is used with Global Copy (previously known as PPRC-XD), as shown in Figure 3-9.

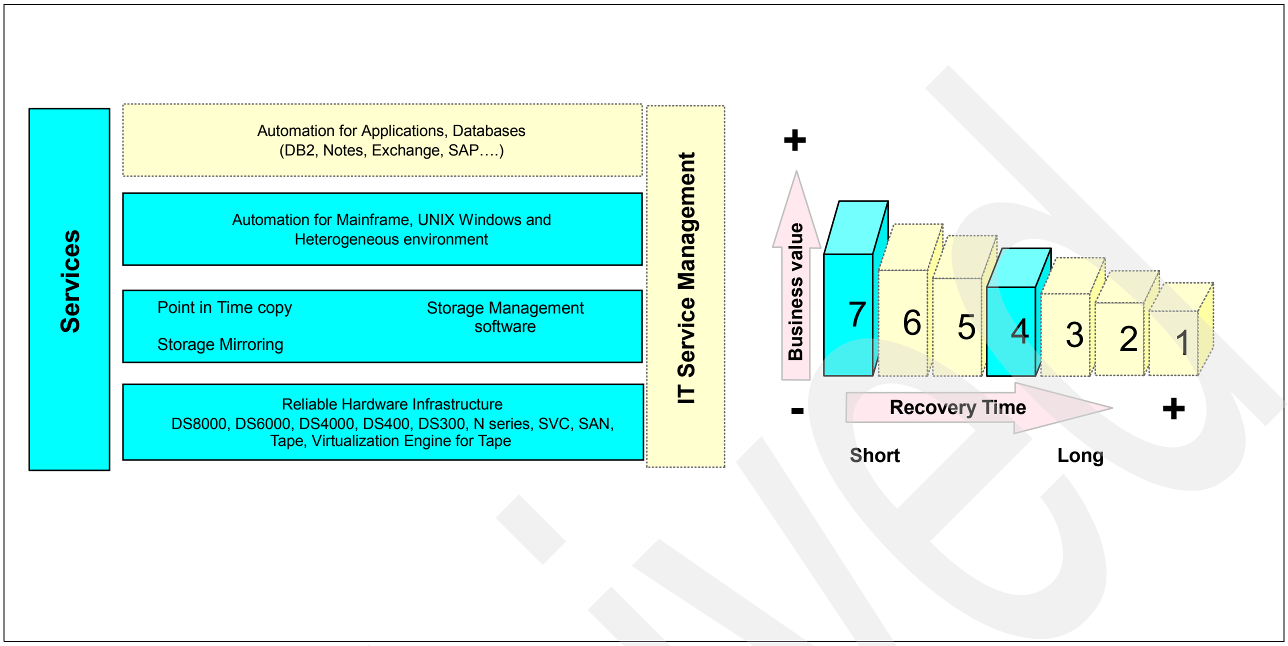
Services
Automation for Applications, Databases (DB2, Notes, Exchange, SAP….)
Automation for Mainframe, UNIX Windows and Heterogeneous environment
Point in Time copy
Storage Mirroring
Storage Management software
Reliable Hardware Infrastructure
DS8000, DS6000, DS4000, DS400, DS300, N series, SVC, SAN, Tape, Virtualization Engine for Tape
IT Service Management
+
Business value
7
6
5
4
3
2
1
-
Recovery Time
+
Short
Long

Figure 3-9 HACMP/XD protection tier

Solution description

- HACMP is high availability clustering software for AIX environments. In a typical HACMP environment, the nodes are all attached to a common disk system and can be active/active or active/passive. In either case, a failure on an active node will trigger a failover of processors and applications to the surviving node. HACMP is typically used to protect availability within a site while the HACMP/XD component deals with failing to an alternate recovery site.
- HACMP/XD is an extension to HACMP that enables the process of failing over to an alternate site. This can be done through storage hardware based mirroring or through server based IP or GLVM Mirroring
- IBM Disk Systems provide a choice of highly reliable and scalable storage, including the DS6000, DS8000, and SAN Volume Controller (SVC).
- Metro Mirror is the IBM name for Synchronous Mirroring technologies. The DS6000 and DS8000 support a maximum distance of 300 km for mirror links (greater distances are supported on special request), and the SVC supports a maximum distance of 100 km.

The components work together as an integrated system to provide:

- Automatic backup and recovery after failures: The recovery of business-critical applications and data after a wide range of system failures is not dependent on the availability of any one component.
- Automated control of data mirrors: The complicated tasks of establishing, suspending, reversing, and resynchronizing data mirrors are automatic, thus reducing the chance of data loss or corruption due to user error.
- Easier execution of planned system outages: Tools for user-controlled operations help to gracefully bring individual system components offline for scheduled maintenance while minimizing the downtime experienced by the users.

Solution highlights

HACMP/XD:

- Improves continuity of business operations and business resiliency
- Provides uninterrupted client service
- Meets Service Level Agreement commitments
- Improves protection of critical business data
- Reduces the risk of downtime
- Provides flexibility for the transfer of operations between sites with minimal disruption
- Improves business resiliency

Figure 3-10 shows a sample two-site configuration using Metro Mirror between the two disk systems with HACMP/XD.

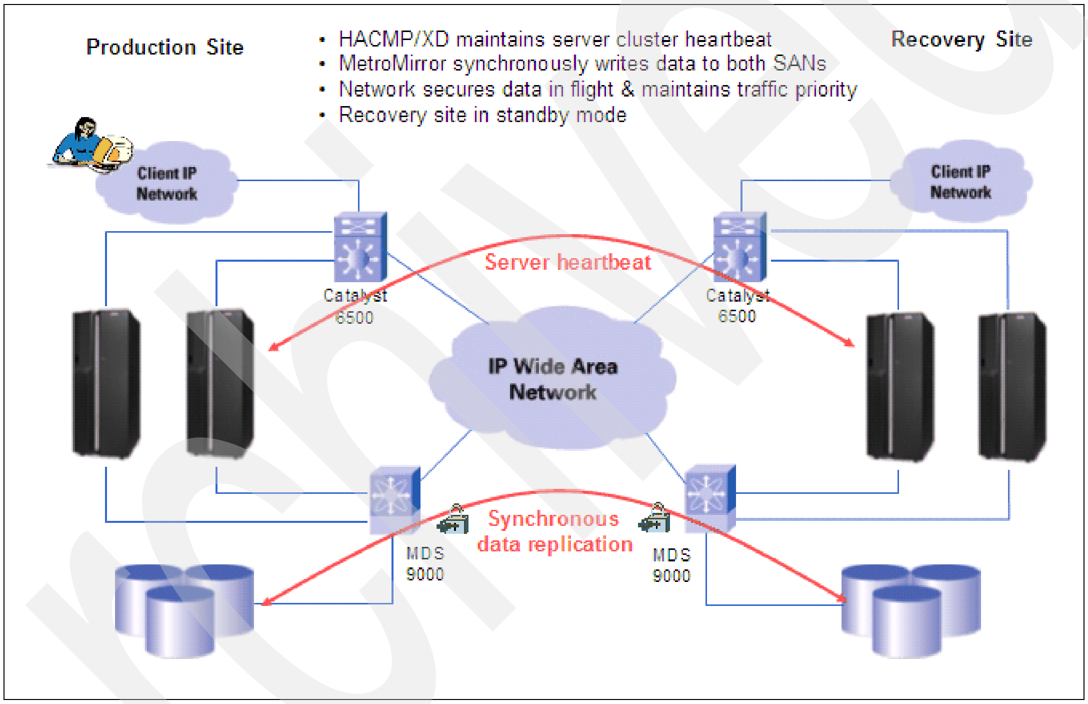

Production Site
• HACMP/XD maintains server cluster heartbeat
• MetroMirror synchronously writes data to both SANs
• Network secures data in flight & maintains traffic priority
• Recovery site in standby mode
Recovery Site
Client IP Network
Client IP Network
Server heartbeat
Catalyst 6500
Catalyst 6500
IP Wide Area Network
Synchronous data replication
MDS 9000
MDS 9000

Figure 3-10 HACMP/XD sample configuration

Solution components

Here we discuss the solution components.

Hardware requirements

HACMP/XD requires:

- IBM System p server
- AIX supported level
- HACMP/XD software
- DS8000, DS6000, or SVC

Additional information

You can find further information at:

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

3.4 Continuous Availability for MaxDB and SAP liveCache

This solution addresses the failover to a standby server and fast restart of SAP liveCache landscapes.

Tier level and positioning within the Resiliency Family

This is a Tier 7 solution, as shown in Figure 3-11. It uses AIX HACMP and FlashCopy to protect SAP SCM 4.1 and above environments. It gives very rapid application failover to a standby server.

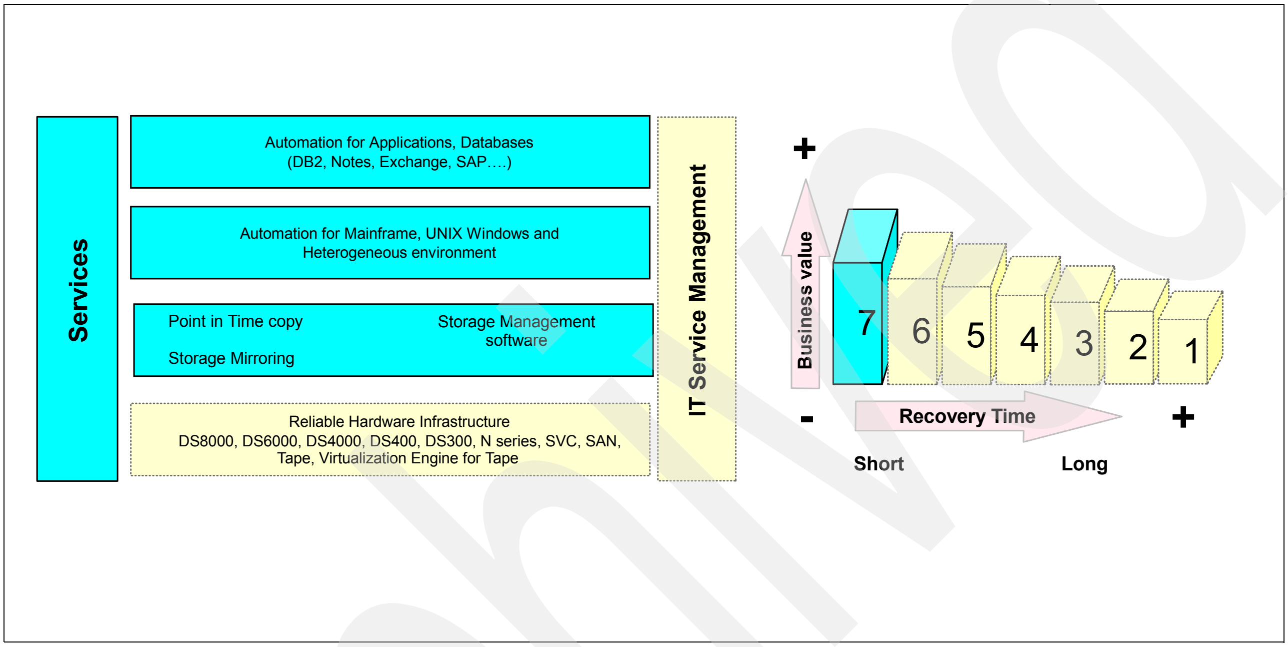
Services
Automation for Applications, Databases
(DB2, Notes, Exchange, SAP….)
Automation for Mainframe, UNIX Windows and
Heterogeneous environment
Point in Time copy
Storage Mirroring
Storage Management
software
Reliable Hardware Infrastructure
DS8000, DS6000, DS4000, DS400, DS300, N series, SVC, SAN,
Tape, Virtualization Engine for Tape
IT Service Management
+
Business value
7
6
5
4
3
2
1
-
Recovery Time
+
Short
Long

Figure 3-11 MaxDB and SAP liveCache hot standby solution positioning

Solution description

SAP liveCache is an SAP database that is preloaded into memory for very fast access. SAP liveCache can build a very large memory cache and perform specially tailored functions against the in-memory data structures.

SAP has introduced hot-standby functionality with liveCache 7.5, available with SCM 4.1, to provide the fastest possible means of recovery. This solution design for the liveCache hot standby uses split mirror and concurrent volume access in the disk system and is closely integrated with the disks' control software via an API. IBM offers this solution on AIX for the SVC, DS8000, and ESS, with HACMP providing the cluster functionality.

In the solution, two AIX nodes share a SAN-connected disk system using HACMP for application monitoring and failover. The second AIX node is automatically initialized with a hot-standby server image with a synchronized copy of the database and memory cache. In the event of failure on the first AIX node, the standby image is automatically failed over to.

Solution highlights

The benefits of this implementation are the following:

- Speed of recovery and return to production
- Coverage of server outage
- Coverage of database outage

- Coverage of data disk failures
- Automated failover and failback
- Designed for minimal or no performance impact on production system
- Ease of management for DB administrators

3.4.1 Solution components

The solution components are the following:

- IBM System p servers
- AIX 5L™ V5.1 or higher
- Databases:
  - MaxDB 7.5.
  - SAP liveCache 7.5 (available with SCM 4.1).
- Storage systems:
  - IBM System Storage DS8000 with Advanced Copy Services.
  - IBM System Storage SAN Volume Controller with Advanced Copy Services.
  - IBM Enterprise Storage Server® with Advanced Copy Services.
- High Available Cluster Multiprocessing (HACMP)
- MaxDB and SAP liveCache hot standby storage dependent library.
  - For DS8000, the DSCLI on all servers and FlashCopy.
  - For SVC, the IBM2145CLI and SecureShell (SSH) are required on all servers.
  - For ESS, the IBM2105CLI on all servers and FlashCopy.

Additional information

For more information, see the following Web site or contact your IBM sales representative:

http://www.ibm.com/servers/storage/solutions/sap/index.html

3.5 Copy Services for System i

Copy Services for System i is required in order to use Metro Mirror or Global Mirror in a System i environment.

It is a Tier 4 solution when FlashCopy is used, and a Tier 7 solution with Metro Mirror or Global Mirror solution together with DS8000 or DS6000, as shown in Figure 3-14 on page 42.

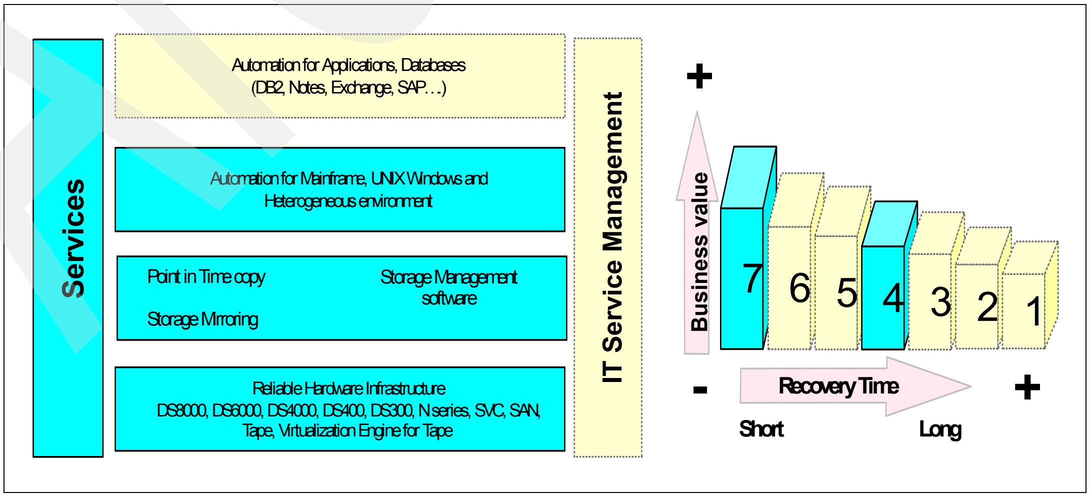
Services
Automation for Applications, Databases
(DB2, Notes, Exchange, SAP…)
Automation for Mainframe, UNIX Windows and
Heterogeneous environment
Point in Time copy
Storage Mirroring
Storage Management
software
Reliable Hardware Infrastructure
DS8000, DS6000, DS4000, DS400, DS300, N series, SVC, SAN,
Tape, Virtualization Engine for Tape
IT Service Management
+
Business value
7
6
5
4
3
2
1
-
Recovery Time
+
Short
Long

Figure 3-12 Copy Services for System i solution positioning

Solution description

Copy Services for System i is an integrated service that allows System i to use the DS8000 / DS6000 FlashCopy, Metro Mirror, or Global Mirror Advanced Copy functions. Copy Services for System i is not a single product or pre-packaged software but, rather, a service offering tailored to the individual environment where it is implemented, and is available only through STG Lab Services.

Its benefits include:

- Ability to use independent auxiliary storage pools (IASPs), which is a collection of disk units that can be brought online or taken offline independently of the rest of the storage on the system
- Automatic failover
- A utility to ensure that all parts of the replicated environment function

Requirements

IASPs require i5/OS® V5R2 or higher and IASP spool files support V5R3 and higher. They are only possible as part of the System i Copy Services offering from STG Lab Services.

3.5.1 Copy Services with IASPs

Metro Mirror, Global Mirror, and FlashCopy are available on System i with IBM System Storage DS6000 and DS8000, as well as with the ESS.

Metro Mirror

While it is technically possible to do Metro Mirror as it would be done in other environments, the fact that System i typically sees only a single level of storage creates a challenge. This is because mirroring in a single level environment means that all data is mirrored. In a synchronous environment like Metro Mirror, this can create performance issues because, in addition to application specific data, the copy service will also mirror the *SYSBAS information. The *SYSBAS information is not needed in the recovery site and doing so creates more transactions that must be synchronously mirrored outside of the relevant applications. As each transaction must be mirrored in turn, this could negatively impact the performance of all applications in that specific System i environment.

Global Mirror

As an asynchronous mirroring technology, Global Mirror will not impact the function of any production applications. However, bandwidth is one of the largest costs involved in any ongoing business continuity effort. As such, it is important to consider whether it is truly necessary to mirror all data.

Through modeling, STG Lab Services has demonstrated a significant cost savings by using IASPs and mirroring only the application data. Again, the *SYSBAS data does not need to be mirrored and doing so can actually complicate recovery efforts. IASPs give the option of separating out that data and only mirroring the application specific data that is needed.

FlashCopy

When issuing a FlashCopy in a non-IASP environment, all applications using the source storage will need to be quiesced in order to produce a consistent, restartable, copy of data in the target storage.

By combining FlashCopy with IASPs, it is possible to eliminate the need to quiesce the applications. Specific applications can have their Point in Time copy made at will, rather than all copies being made at once. This copy can then be used to create tape backups without a long application outage, or for data mining, as needed.

3.5.2 Copy Services for System i

Here we discus Copy Services for System i

Use of IASPs

IASPs provide greater flexibility for data management. Without IASPs, all data is viewed in a single level of storage. This means that any time there is a storage outage, all data must take the same outage. Any time data is mirrored or copied, all data must be mirrored or copied.

Using ASPs allows the data to be broken down into more manageable blocks that can take individual outages or be mirrored independently. It also allows for participation within a SAN environment.

Automatic failover

Copy Services for System i is not just a tool to make more effective use of DS6000/DS8000 copy services within a System i environment. It also provides a true high-availability offering.

Through the Copy Services offering, it is possible to set up mirror pairs along with a server cluster. Doing so allows the environment to fail over some, or all, applications from one location to another. As always, these locations could be within a single physical site or between geographically dispersed sites, depending on the form of disk mirroring in use. Figure 3-13 gives an overview of this environment.

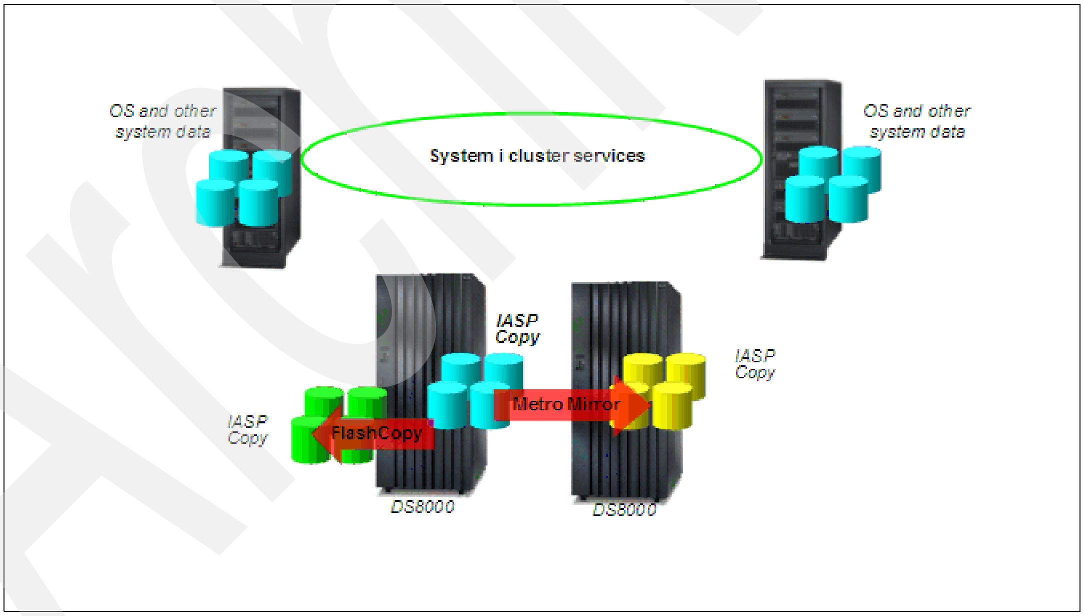

Figure 3-13 Server and storage clustering with Copy Services for System i and Metro Mirror or Global Mirror

Copy Services for System i *can* failover the environment through a single command. The SWPPRC command will ask for confirmation that a site switch is really what is required. Once that confirmation is given, the command will vary off the IASP(s) in their current location, and perform all commands necessary to bring them up in the alternate location. This process will typically occur within a matter of minutes.

3.5.3 Metro Cluster for N series overview

Metro Cluster for N series is a Tier 7 solution, as shown in Figure 3-14.

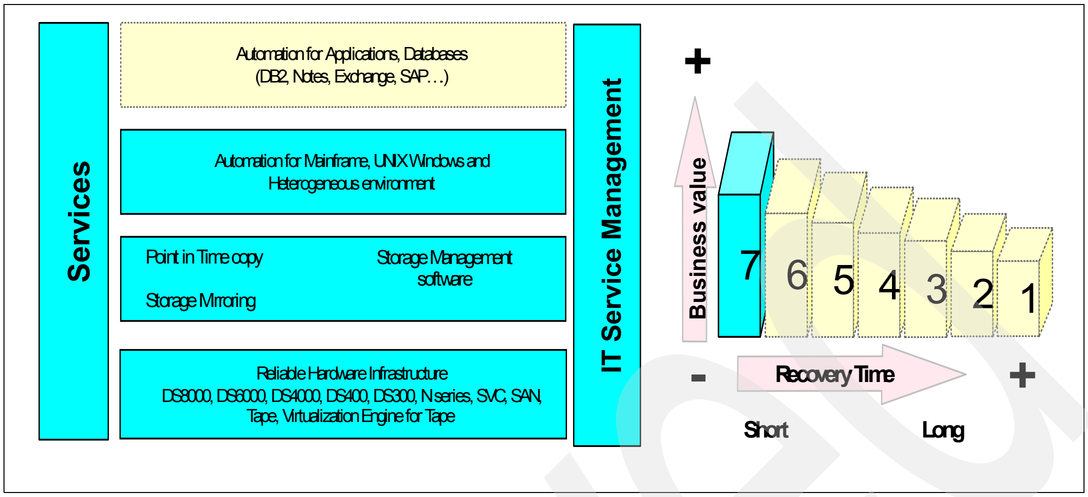
Services
Automation for Applications, Databases (DB2, Notes, Exchange, SAP…)
Automation for Mainframe, UNIX Windows and Heterogeneous environment
Point in Time copy
Storage Mirroring
Storage Management software
Reliable Hardware Infrastructure DS8000, DS6000, DS4000, DS400, DS300, N series, SVC, SAN, Tape, Virtualization Engine for Tape
IT Service Management
+
Business value
7
6
5
4
3
2
1
-
Recovery Time
+
Short
Long

Figure 3-14 Metro Cluster for N series solution positioning

Solution description

For N series, Continuous Availability is implemented using clustering functionality known as Metro Cluster, which builds on synchronous mirror technology in order to move data from one N series to another.

SyncMirror for Metro Cluster

SyncMirror is somewhat similar to logical volume mirroring, and is typically used within an N series, writing to a pair of disks.

In Metro Cluster, as shown in Figure 3-15 on page 43, rather than writing to the disk and then allowing the disk system to synchronously mirror the data to the recovery disk system, the disk controller issues two writes at once: One write goes to the local disk while the other goes to the remote disk. Both N series can be active at the same time and would each write to the other. In order to recover in the case of a failure event, each N series carries a dormant Operating System image of the other N series.

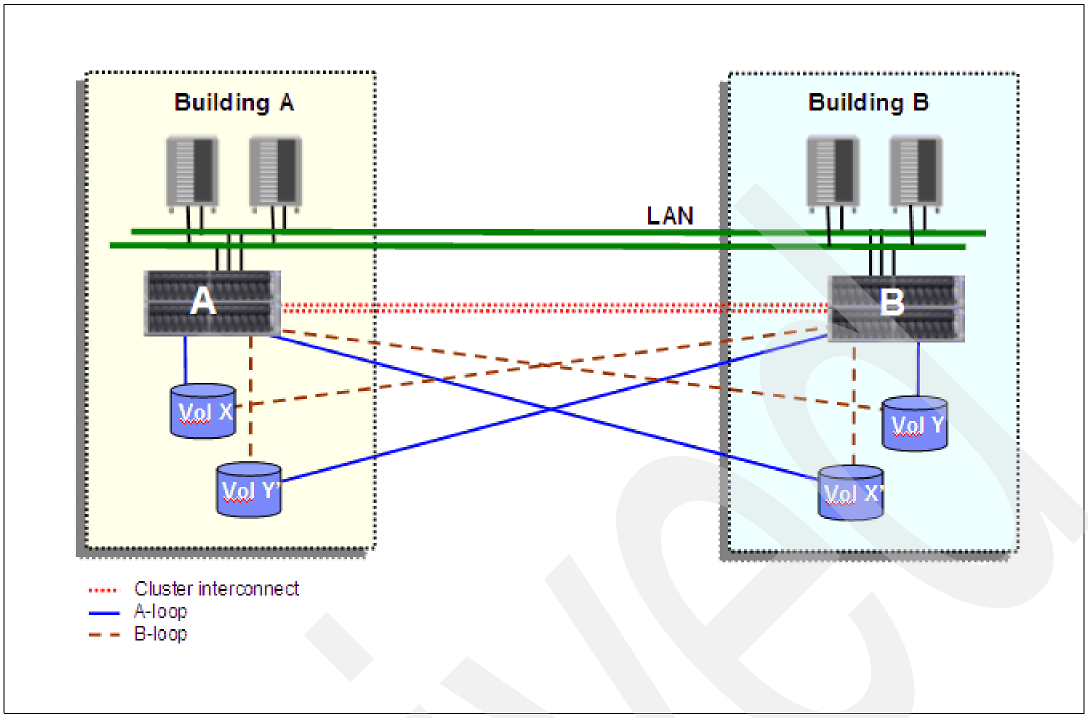

Figure 3-15 Sync Mirror environment for Metro Cluster.

Metro Cluster failover

If a failure occurs, the receiving N series activates the dormant copy of the production system's operating system and accesses the recovery volume locally. Within minutes, the second N series is available and appears identical to the now disabled N series, as shown in Figure 3-16.

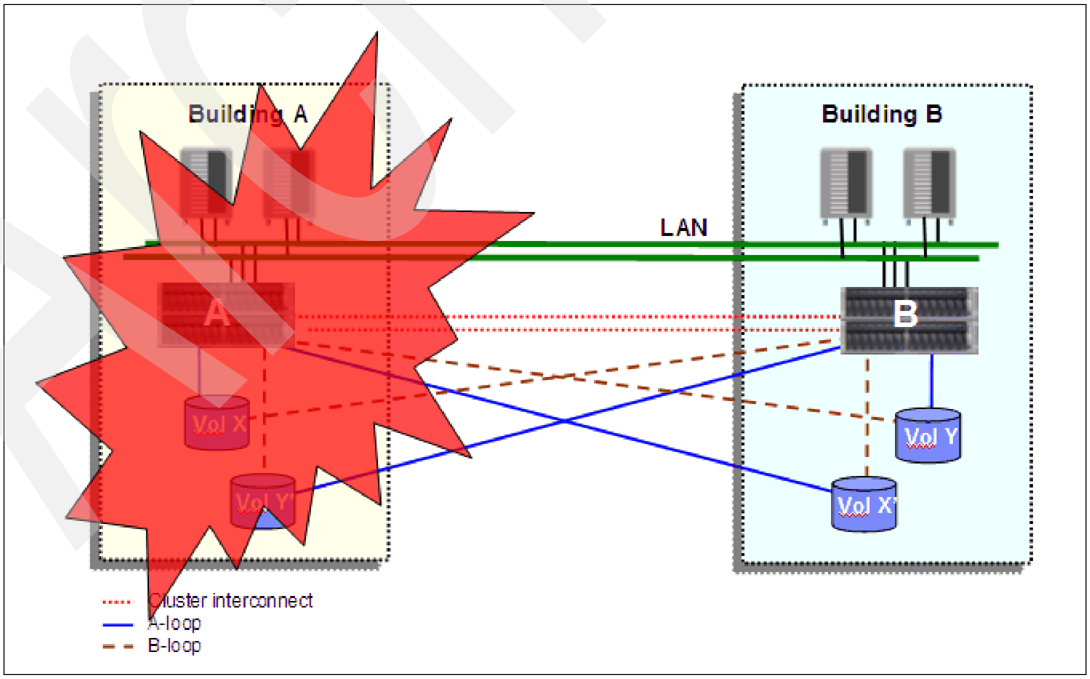

Figure 3-16 Metro Cluster failover to site B

With proper TCP/IP or SAN infrastructure, this may allow applications to continue accessing storage without the need for any particular failover procedure.

Metro Cluster is available on N5000 and N7000 series.

Metro Cluster types

There are two types of Metro Clusters available, depending on the environment and distance to be covered.

- Stretch Cluster is used only for short distances within a campus environment. Stretch Clusters can connect via FCP, a pair of N series systems at a supported distance of up to 300 m.
- Fabric Cluster is used for longer distances. Fabric Clusters allow the N series to be separated by up to 100 km, across a switched FCP infrastructure, as opposed to direct connections permitted in a stretch cluster.

In either case, the Metro Cluster handles the switch only at the disk system level. Any failover that is required at the server level requires separate clustering technology.

4

Rapid Data Recovery

Rapid Data Recovery is based on maintaining a second disk-resident copy of data that is consistent at a point-in-time as close to the time of a failure as possible. This consistent set of data allows for the restart of systems and applications without having to restore data and re-applying updates that have occurred since the time of the data backup. It is possible that there may be a loss of a minimal number of in-flight transactions.

Rapid Data Recovery spans Tier 4 through Tier 6. Rapid Data Recovery is distinguished from Continuous Availability by the fact that the automation required to be a Tier 7 solution is not present.

This chapter provides a high level overview of Rapid Data Recovery in the following environments:

- System z and mixed z + Distributed
- System z, mixed z and distributed, distributed
- TPC for Replication
- SAN Volume Controller
- System i
- FlashCopy Manager and PPRC Migration Manager

4.1 System Storage Rapid Data Recovery for System z and mixed z+Open platforms (GDPS/PPRC HyperSwap Manager)

Rapid Data Recovery for System z is provided by the IBM Global Services service offering, GDPS/PPRC HyperSwap Manager (GDPS/PPRC HM), in the GDPS suite of offerings. It uses Metro Mirror to mirror the data between disk systems. Metro Mirror is a hardware-based mirroring and remote copying solution for IBM System Storage disk solutions.

Used in a two site implementation, GDPS/PPRC HM provides a Tier 6 Business Continuity solution for System z and x+Open data, as shown in Figure 4-1. It falls short of being a Tier 7 solution because it lacks the System z processor, System z workload, and Coupling Facility recovery automation provided by a full GDPS/POPRC implementation.

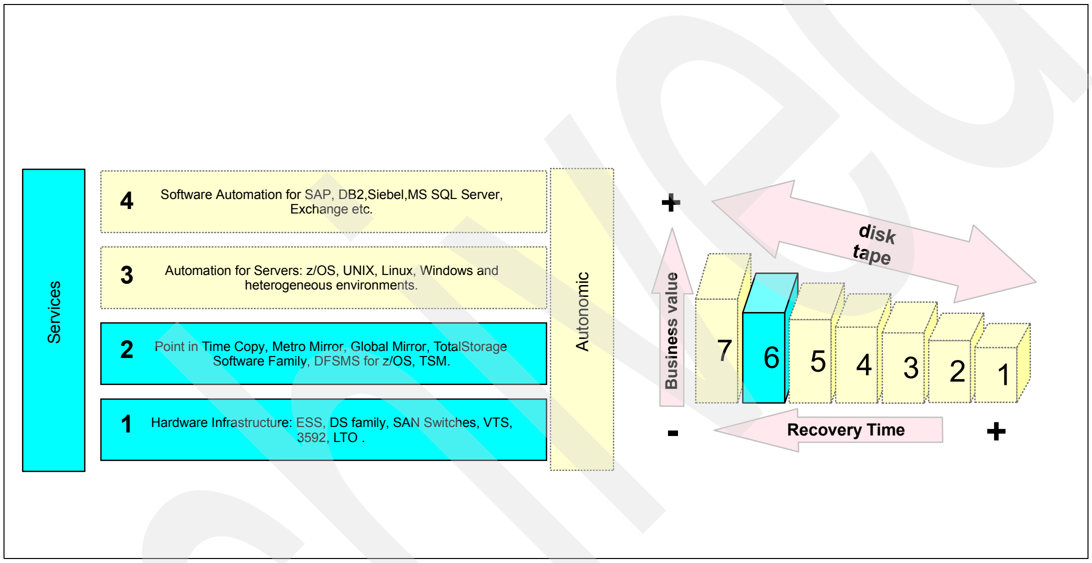
Services
4 Software Automation for SAP, DB2,Siebel,MS SQL Server, Exchange etc.
3 Automation for Servers: z/OS, UNIX, Linux, Windows and heterogeneous environments.
2 Point in Time Copy, Metro Mirror, Global Mirror, TotalStorage Software Family, DFSMS for z/OS, TSM.
1 Hardware Infrastructure: ESS, DS family, SAN Switches, VTS, 3592, LTO .
Autonomic
+
Business value
disk
tape
7
6
5
4
3
2
1
-
Recovery Time
+

Figure 4-1 Tier 6 Business Continuity solution for System z and x+Open data

Continuous Availability for System z data

GDPS/PPRC HyperSwap Manager is primarily designed for single site or multiple site System z environments, to provide Continuous Availability of disk-resident System z data by masking disk outages due to failures. Planned outage support is also included, for example, for planned disk maintenance.

When a disk failure occurs, GDPS/PPRC HM invokes HyperSwap to automatically switch the disk access of System z data to the secondary disk system, as shown in Figure 4-2 on page 47. When a primary disk outage for maintenance is required, user interface panels can be used to invoke a HyperSwap switch of System z data access to the secondary disks. After the disk repair or maintenance has been completed, HyperSwap can be invoked to return to the original configuration.

Disk Reconfiguration with HyperSwap

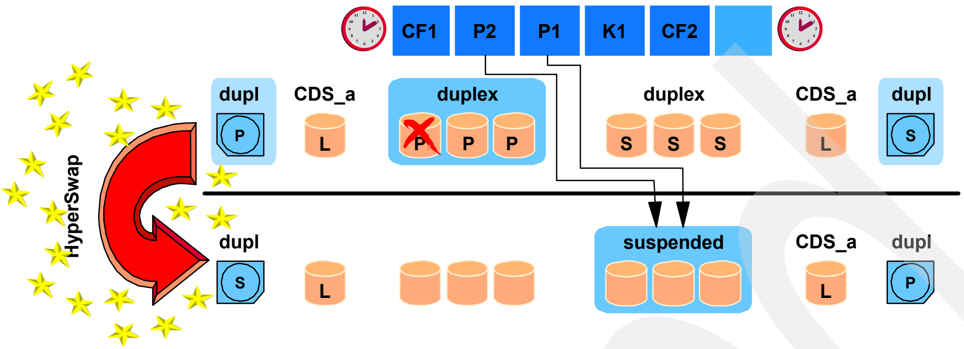

Unplanned

- Parallel Sysplex - P1, P2
- K1 (GDPS controlling system)
- Disk Failure detected
- GDPS automation invokes
  - HyperSwap disk configuration
    - **Failover invoked (secondary disks in suspended state)**
  - After primary disk failure fixed
    - Failback invoked (updates to data copied)
    - Execute HyperSwap again to return to original configuration

Planned

- Parallel Sysplex - P1, P2
- K1 (GDPS controlling system)
- HyperSwap disk configuration
  - Swap primary/secondary disks
  - **Failover invoked (secondary disks in suspended state)**
- After maintenance is completed
  - Failback invoked (updates to data copied)
  - Execute HyperSwap again to return to original configuration

P1, P2, remain active throughout the procedure

Figure 4-2 HyperSwap disk reconfiguration

If a disaster occurs at the primary site, GDPS/PPRC HM can swap to the secondary disk systems, at the same time assuring data consistency at the remote site via GDPS/PPRC HM's control of Metro Mirror Consistency Groups.

4.1.1 GDPS/PPRC Open LUN management

When Metro Mirror is implemented for DS8000 and DS6000 disk systems with System z data, the GDPS Open LUN management function also allows GDPS to manage Metro Mirroring of open systems data within the same primary disk system.

After being HyperSwapped, open system servers will need to be restarted, but the open systems data will be data and time consistent with all other LUNs in the Consistency Group.

When Hyper Swap is invoked due to a failure or by a command, a FREEZE of the open data occurs to maintain data consistency.

4.1.2 Product components

Rapid Data Recovery for System z requires:

- IBM System z servers
- IBM System Storage DS8000, DS6000, (or ESS) with Metro Mirror and FlashCopy
- IBM Tivoli NetView and System Automation
- IBM Global Services for GDPS/PPRC HM and implementation

4.1.3 Additional information

For more information, see the GDPS Web site:

http://www.ibm.com/systems/z/gdps

4.2 System Storage Rapid Data Recovery for System z and mixed z+Distributed platforms using TPC for Replication

Rapid Data Recovery for System z and Distributed systems is also provided by TPC for Replication. TPC for Replication is an effective, user friendly GUI offering for defining and managing the setup of data replication and recovery. TPC for Replication has no dependencies on the operating system, as no client installation is required. TPC for Replication provides a single common GUI to control the entire environment, common across different operating systems.

TPC for Replication provides a Rapid Data Recovery Tier 6 Business Continuity solution for System z and z+Open data. See Figure 4-3.

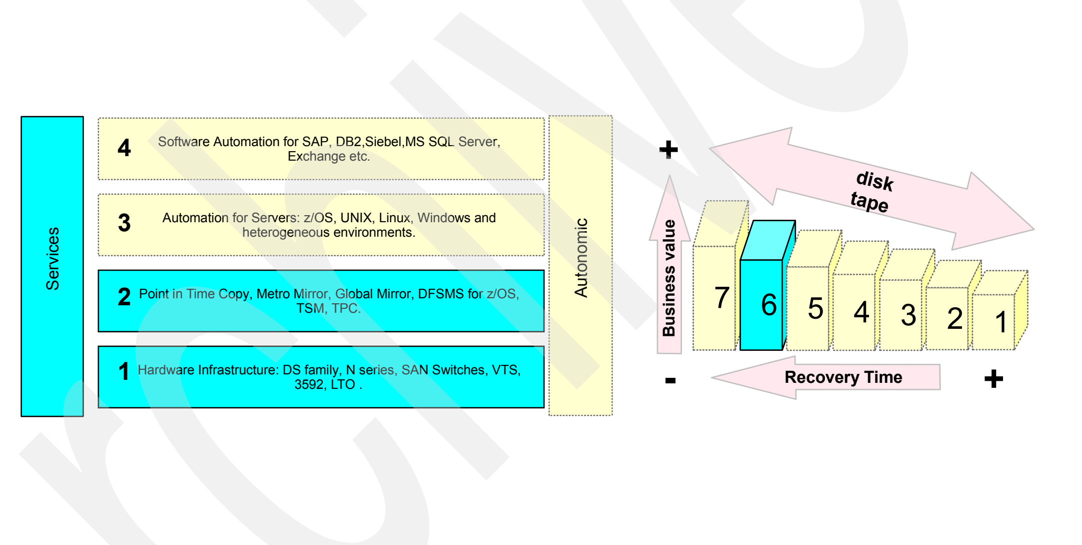
Services
4 Software Automation for SAP, DB2,Siebel,MS SQL Server, Exchange etc.
3 Automation for Servers: z/OS, UNIX, Linux, Windows and heterogeneous environments.
2 Point in Time Copy, Metro Mirror, Global Mirror, DFSMS for z/OS, TSM, TPC.
1 Hardware Infrastructure: DS family, N series, SAN Switches, VTS, 3592, LTO .
Autonomic
+
Business value
disk
tape
7
6
5
4
3
2
1
-
Recovery Time
+

Figure 4-3 TPC for Replication tiers

TPC for Replication manages the FlashCopy, Metro Mirror, and Global Mirror advanced copy services on IBM System Storage DS8000, DS6000, SAN Volume Controller (not Global Mirror at the time of writing), and ESS. It is available in two complementary packages.

- TPC for Replication provides disaster recovery management of advanced copy services of the supported disk systems. It can automate the administration and configuration of these services, provide operational control (starting, suspending, and resuming) of copy services tasks, and monitor and manage the copy sessions.

  TPC for Replication provides management for one way (single direction) data replication to protect you against the primary site failure.

- TPC for Replication Two Site BC requires TPC for Replication as a prerequisite. TPC for Replication Two Site BC provides Disaster Recovery management through planned and unplanned failover and failback automation for the supported disk systems. It also provides a high availability capability for the TPC for Replication server, so that replication can be managed, even if the primary TPC for Replication server fails.

4.3 TPC for Replication functionality

TPC for Replication includes the following functionality:

- Managing and configuring the copy services environment
  - Add/delete/modify storage devices.
  - Add/delete/modify copy sets (a copy set is a set of volumes containing copies of the same data.
  - Add/delete/modify sessions (a container is a container of multiple copy sets managed by a replication manager).
  - Add/delete/modify logical paths (between storage devices).
- Monitoring the copy services environment
  - View session details and progress.
  - Monitor sessions (with status indicators and SNMP alerts).
  - Diagnostics (error messages).

4.3.1 Functionality of TPC for Replication Two Site BC

- All functions of TPC for Replication
- Failover and failback from primary to a Disaster Recovery site
- Support for IBM TotalStorage Productivity Center for Replication high-availability server configuration

4.3.2 Environment and supported hardware

TPC for Replication and TPC for Replication Two Site BC require a separate server for the application (or two if using a standby server). The server can run in Windows, Linux, or AIX. The TPC for Replication Web page is at:

http://www.ibm.com/servers/storage/software/center/replication/index.html

4.4 IBM System Storage SAN Volume Controller (SVC)

Many administrators have to manage disparate storage systems. IBM System Storage SAN Volume Controller (SVC) brings diverse storage devices together in a virtual pool to make all the storage appear as one logical device to centrally manage and to allocate capacity as needed. It also provides a single solution to help achieve the most effective on demand use of key storage resources.

The SVC addresses the increasing costs and complexity in data storage management by shifting storage management intelligence from individual SAN controllers into the network and by using virtualization.

The SVC is a scalable hardware and software solution that facilitates aggregation of storage from different disk subsystems. It provides storage virtualization and thus a consistent view of storage across a SAN.

The IBM SVC provides a resiliency level of Tier 6, when coupled with either Metro Mirror or Global Mirror, as shown in Figure 4-4.

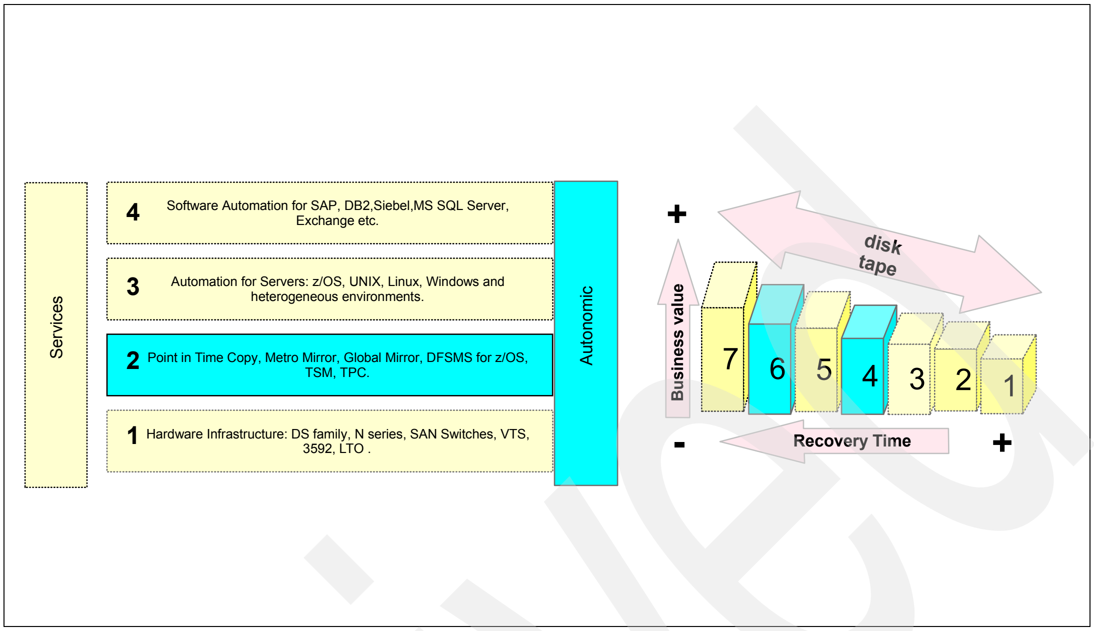
Services
4 Software Automation for SAP, DB2,Siebel,MS SQL Server, Exchange etc.
3 Automation for Servers: z/OS, UNIX, Linux, Windows and heterogeneous environments.
2 Point in Time Copy, Metro Mirror, Global Mirror, DFSMS for z/OS, TSM, TPC.
1 Hardware Infrastructure: DS family, N series, SAN Switches, VTS, 3592, LTO .
Autonomic
+
Business value
disk
tape
7
6
5
4
3
2
1
-
Recovery Time
+

Figure 4-4 SVC tiers

SVC's storage virtualization:

- Consolidates disparate storage controllers into a single view
- Improves application availability by enabling data migration between disparate disk storage devices nondisruptively
- Improves Disaster Recovery and business continuity
- Reduces both the complexity and costs of managing SAN-based storage
- Increases business application responsiveness
- Maximize storage utilization
- Does dynamic resource allocation
- Simplifies management and improves administrator productivity
- Reduces storage outages
- Supports a wide range of servers and storage systems

Rapid Data Recovery is provided by the SVC through the usage of inherent features of the copy services it utilizes, that is, Metro Mirror, Global Mirror, and FlashCopy operations. A principle benefit of the SVC is that it provides a single interface and point of control for configuring, running, and managing these copy services, regardless of what is the underlying disk.

4.4.1 SVC FlashCopy

SVC FlashCopy falls under BC Tier 4. FlashCopy makes a copy of source virtual disks to a set of target virtual disks. After the copy operation completes, the target virtual disks have the contents of the source virtual disks as they existed at a single point-in-time. FlashCopy operations require the use of a Consistency Group; a Consistency Group contains a number of FlashCopy mappings (source to target virtual disk). All the FlashCopy mappings are started simultaneously, so that the point-in-time copy will be consistent across all the mappings in the Consistency Group. This allows data consistency to be preserved across multiple virtual disks.

4.5 SVC remote mirroring

The SVC supports two forms of remote mirroring: synchronous remote copy (implemented as Metro Mirror) and asynchronous remote copy (implemented as Global Mirror).

Metro Mirror: synchronous remote copy

SVC Metro Mirror is a fully synchronous remote copy technique that ensures that updates are committed at both the primary and secondary virtual disks before returning a completion status to the application, as shown in Figure 4-5. Since the write is synchronous, the latency and bandwidth of the remote site connection may have impacts on the applications performance, especially under peak loads. Therefore, there are distance limitations for Metro Mirror of 300 m shortwave and 10 km longwave between the primary and secondary SVC nodes.

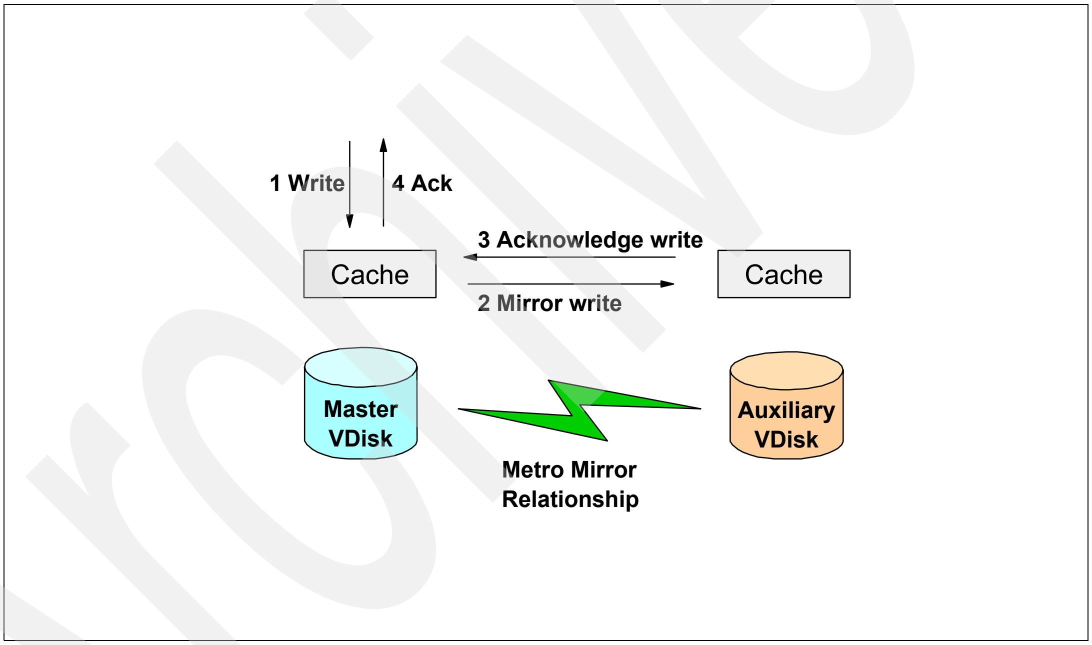

Figure 4-5 Synchronous remote copy

Metro Mirror provides a Tier 6 BC tier solution.

Global Mirror: asynchronous remote copy

With SVC Global Mirror, the application receives a completion status when an update is sent to the secondary site, but before the update is necessarily committed, as shown in Figure 4-6. This means the remote copy can be performed over distances further than those allowed for Metro Mirror.

In a failover situation, where the secondary site needs to become the primary data source, some updates may be missing at the secondary site, because the updates were not fully committed when the failure occurred. The application must have some external mechanism for recovering the missing updates and reapplying them, for example, transaction log replay.

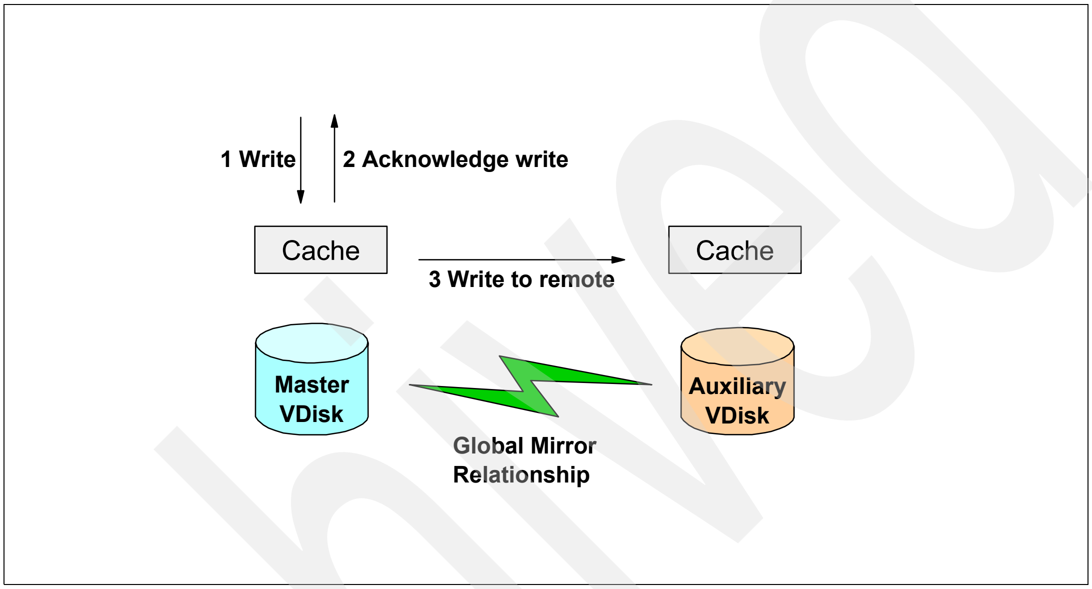

Figure 4-6 ASynchronous remote copy

Global Mirror provides a Tier 6 BC tier solution.

4.6 Rapid Data Recovery with System i

System i utilizes the functional richness of Metro Mirror, Global Mirror, and FlashCopy to provide Rapid Data Recovery

Recovery tools and concepts for System i require careful planning and design phases in order to achieve the desired objectives. This can be done using the System i Copy Services Toolkit and backup, recovery, and media services using Tivoli Storage Management.

More information about the System Copy Services Toolkit is in 3.5, “Copy Services for System i” on page 39.

4.7 FlashCopy Manager and PPRC Migration Manager

FlashCopy Manager and PPRC Migration Manager are IBM Storage Services solutions for z/OS users of FlashCopy and Metro Mirror. For this environment, these packaged solutions are designed to:

- Simplify and automate the z/OS jobs that set up and execute a z/OS FlashCopy, Metro Mirror, or Global Copy environment
- Improve the speed of elapsed execution time of these functions
- Improve the administrator productivity to operate these functions

These two related tools use a common style of interface, operate in a very similar fashion, and are designed to complement each other. A user familiar with one of the offerings will find the other offering easy to learn and use. They are intended for large z/OS remote copy environments in the order of hundreds or thousands of source/target pairs.

FlashCopy Manager is a Tier 4 Business Continuity solution. It is a series of efficient, low overhead assembler programs and ISPF panels that allow the z/OS ISPF user to define, build, and run FlashCopy jobs for any sized FlashCopy z/OS environment.

PPRC Migration Manager provides a series of efficient, low overhead assembler programs and ISPF panels that allow the z/OS ISPF user to define, build, and run DS8000, DS6000, and ESS Metro Mirror and Global Copy jobs. PPRC Migration Manager supports both planned and unplanned outages.

PPRC Migration Manager and FlashCopy Manager also support Global Mirror for z/OS environments.

PPRC Migration Manager is considered a Tier 6 Business Continuity tool when it controls synchronous Metro Mirror storage mirroring, as the remote site will be in synchronous mode with the primary site.

PPRC Migration Manager is considered a Tier 4 Business Continuity tool when it controls non-synchronous Global Copy, as the remote site data will not be in data integrity until the Global Copy *go - to - sync* process is done to synchronize the local site and the remote site.

For more information about these solutions, see:

http://www.ibm.com/servers/storage/services/featured/pprc_mm.html

5

Backup and Restore

Backup and Restore is the most simple and basic solution to protect and recover your data from failure by creating another copy of data from the production system. The second copy of data allows you to restore data to the time of the data backup.

This chapter describes Backup and Restore in the following subjects:

- Overview of Backup and Restore and archive and retrieve
- IBM Tivoli Storage Manager overview
- IBM Data Retention 550
- DFSMShsm, DFSMSdss, and z/VM utilities

5.1 An overview of Backup and Restore and archive and retrieve

Backup is a daily IT operation task where production, application, systems, and user data are copied to a different data storage media, in case they are needed for restore. Restoring from a backup copy is the most basic Business Continuity implementation.

As part of the Business Continuity process, archive data is also a critical data element that should be available. Data archive is different than backup in that it is the only available copy of data on a long term storage media, normally tape or optical disk. The archive data copy will be deleted at a specific period of time, also known as retention-managed data.

5.1.1 What is Backup and Restore

To protect against loss of data, the backup process copies data to another storage media that is managed by a backup server. The server retains versions of a file according to policy, and replaces older versions of the file with newer versions. Policy includes the number of versions and the retention time for versions.

A client can restore the most recent version of a file, or can restore previous retained versions to an earlier point in time. The restored data can replace (over-write) the original, or be restored to an alternative location, for comparison purposes.

5.1.2 What is archive and retrieve

The archive process copies data to another storage media of that is managed by an archive server for long-term storage. The process can optionally delete the archived files from the original storage immediately or at a predefined period of time. The archive server retains the archive copy according to the policy for archive retention time.

A client can retrieve an archived copy of a file when necessary.

5.2 IBM Tivoli Storage Manager overview

IBM Tivoli Storage Manager protects an organization's data from failures and other errors. By managing backup, archive, space management and bare-metal restore data, as well as compliance and Disaster Recovery data in a hierarchy of offline storage, the Tivoli Storage Manager family provides centralized, automated data protection. Thus, Tivoli Storage Manager helps reduce the risks associated with data loss while also helping to reduce complexity, manage costs, and address compliance with regulatory data retention requirements.

Since it is designed to protect a company's important business information and data in case of disaster, the Tivoli Storage Manager server should be one of the main production systems that is available and ready to run for recovery of business data and applications.

Tivoli Storage Manager provides industry-leading encryption support through integrated key management and full support for the built-in encryption capability of the IBM System Storage TS1120 Tape Drive.

This section provides the Tivoli Storage Manager solutions in terms of Business Continuity (BC) and Disaster Recovery. There are six solutions to achieve each BC tier:

- BC Tier 1: IBM Tivoli Storage Manager manual off-site vaulting
- BC Tier 2: IBM Tivoli Storage Manager manual off-site vaulting with a hotsite
- BC Tier 3: IBM Tivoli Storage Manager electronic vaulting
- BC Tier 4: IBM Tivoli Storage Manager with SAN attached duplicates
- BC Tier 5: IBM Tivoli Storage Manager clustering
- BC Tier 6: IBM Tivoli Storage Manager running in a duplicate site

5.2.1 IBM Tivoli Storage Manager solutions overview

These solutions provide protection for enterprise business systems.

Tier level and positioning within the System Storage Resiliency Portfolio

IBM Tivoli Storage Manager solutions support Business Continuity from Tier 1 to Tier 6, as shown in Figure 5-1. These solutions achieve each tier by using hardware, software, and autonomic solutions.

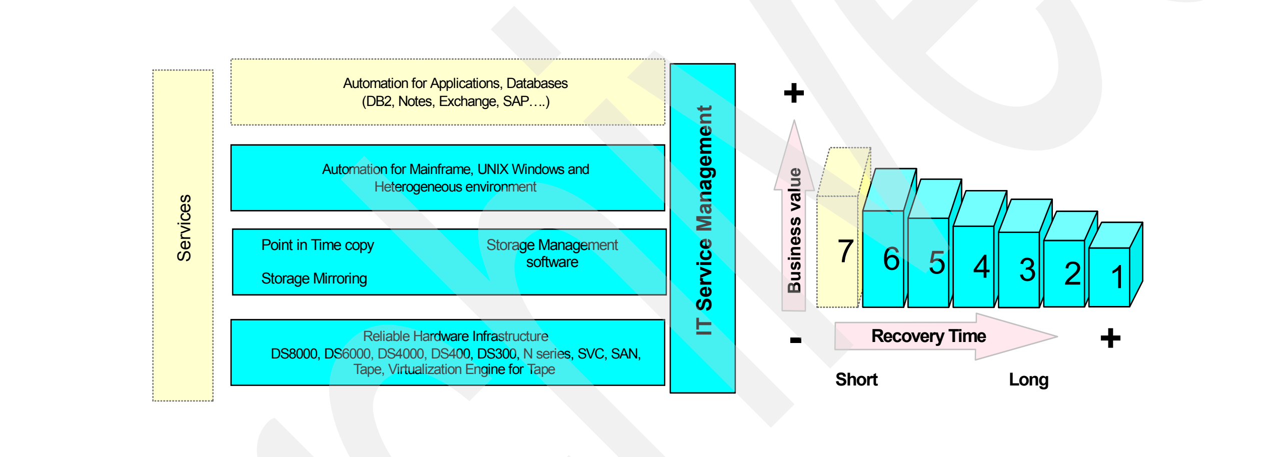
Services
Automation for Applications, Databases
(DB2, Notes, Exchange, SAP….)
Automation for Mainframe, UNIX Windows and
Heterogeneous environment
Point in Time copy
Storage Mirroring
Storage Management
software
Reliable Hardware Infrastructure
DS8000, DS6000, DS4000, DS400, DS300, N series, SVC, SAN,
Tape, Virtualization Engine for Tape
IT Service Management
+
Business value
7
6
5
4
3
2
1
-
Recovery Time
+
Short
Long

Figure 5-1 Tier level and positioning within the Resiliency Family

Solution description

These solutions will enable IBM Tivoli Storage Manager system to achieve Business Continuity for Tier 1 to Tier 6. The solutions provide the ability to minimize the Recovery Time Objective (RTO) and the Recovery Point Objective (RPO) for the client’s Tivoli Storage Manager system.

From BC Tier 1 to Tier 3, the Tivoli Storage Manager BC solutions use features such as Disaster Recovery Manager and server-to-server communication protocol to support tape vaulting and electronic vaulting to an off-site location.

From BC Tier 4 to Tier 6, data storage replication and clustering service are implemented on Tivoli Storage Manager systems. Integration with clustering technology will come into play. Tivoli Storage Manager systems will have the ability to provide high availability and rapid data Backup and Restore service for your enterprise business systems.

Solution highlights

The highlights of these solutions are:

- Continuity of backup and recovery service and meet Service Level Agreement commitments.
- Reduced risk of downtime of Tivoli Storage Manager system.
- Increased business resiliency and maintaining competitiveness.
- Minimize Recovery Time Objective (RTO) of Tivoli Storage Manager system.
- Minimize Recovery Point Objective (RPO) of Tivoli Storage Manager system.

Solution components

The components of these solutions include:

- IBM Tivoli Storage Manager server
- IBM Tivoli Storage Manager product features and functions
  - IBM Tivoli Storage Manager Extended Edition (Disaster Recovery Manager)
  - IBM Tivoli Storage Manager server-to-server communication
  - IBM Tivoli Storage Manager for Copy Services - Data Protection for Exchange
  - IBM Tivoli Storage Manager for Advanced Copy Services - Data Protection for mySAP™
- IBM SystemStorage DS6000, DS8000, and SAN Volume Controller
- IBM High Availability Cluster Multi-Processing (HACMP)
- Microsoft Cluster Server (MSCS)

Additional information

For additional information about these solutions, refer to the following:

- Contact your IBM representative
- See the IBM Redbooks, *IBM Tivoli Storage Management Concepts*, SG24-4877, and *Disaster Recovery Strategies with Tivoli Storage Management*, SG24-6844
- Visit this Web site:

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

5.2.2 IBM Tivoli Storage Manager solutions in detail

This section provides information about solutions and strategies to enable the Tivoli Storage Manager system to achieve the specific BC tiers and meet BC business requirements.

Tier 0 - No off-site data

BC Tier 0 is defined as a single site data center environment having no requirements to back up data to implement a Business Continuity Plan at a different site. See Figure 5-2 on page 59 for an illustration.

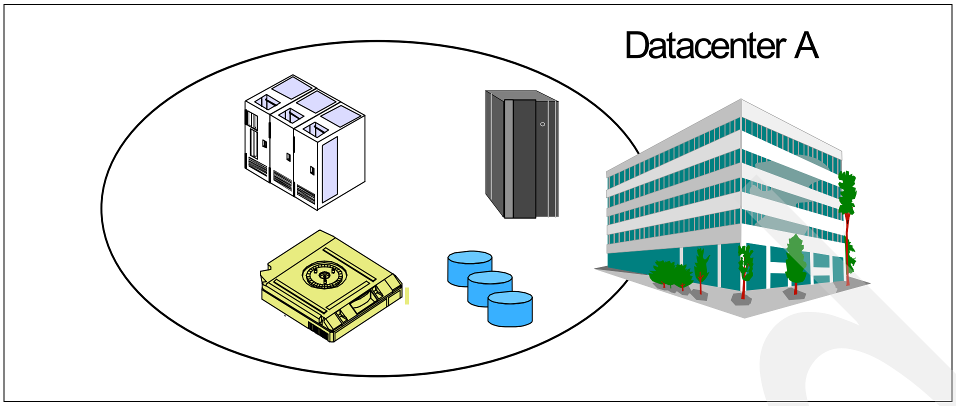

Figure 5-2 BC Tier 0 - Tivoli Storage Manager, no off-site data

For this BC tier, there is no saved information, no documentation, no backup hardware, and no contingency plan. There is therefore no Business Continuity capability at all. Some businesses still reside in this tier. While they may actively make backups of their data, these backups are left onsite in the same computer room, or are only infrequently taken offsite due to lack of a rigorous vaulting procedure. A data center residing on this tier is exposed to a disaster from which they may never recover their business data.

Typical length of time for recovery

The typical length of time for recovery is unpredictable. In many cases, complete recovery of applications, systems, and data is never achieved.

Strategies and operations

The strategies and operations for this solution are the following:

- Normal Tivoli Storage Manager based onsite backups
- No off-site strategy
- No ability to recover from a site disaster except to rebuild the environment

Strategies and operations description

Because BC Tier 0 contains no off-site strategy, the Tivoli Storage Manager system provides Backup and Restore service for the production site only.

Tier 1 - Tivoli Storage Manager backups with manual off-site vaulting

BC Tier 1 is defined as having a Business Continuity Plan (BCP), where data is backed up and stored to a centralized location. Copies of these backups are then manually taken off-site, as shown in Figure 5-3. Some recovery requirements have been determined, including application and business processes. This environment may also have established a recovery platform, although it does not have a site at which to restore its data, nor the necessary hardware on which to restore the data, for example, compatible tape devices.

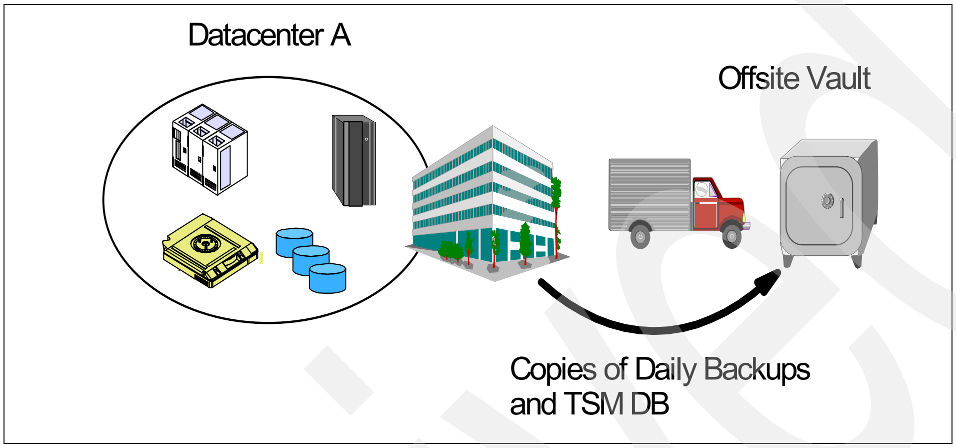

Figure 5-3 BC Tier 1 - Tivoli Storage Manager manual off-site vaulting

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

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

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

Typical length of time for recovery

With this solution, the typical length of time for recovery is normally more than a week.

Strategies and operations

The strategies and operations for this solution are following:

- Use Disaster Recovery Manager (DRM) to automate the Tivoli Storage Manager server recovery process and to manage off-site volumes.
- Recommend creating a Business Continuity Plan and carefully manage off-site volumes.
- Manual storage pool vaulting of copies of data with Tivoli Storage Manager server environment.
- Strategy must include manual vaulting of copies of Tivoli Storage Manager database, volume history information, device configuration information, Business Continuity Plan file, and copy pools for storage at an off-site location.

Strategies and operations description

BC Tier 1 requires off-site vaulting. The following Tivoli Storage Manager components will be sent to off-site:

- Tivoli Storage Manager storage pool copies
- Tivoli Storage Manager database backup
- Tivoli Storage Manager configuration files
  - Volume history file
  - Device configuration file

In this tier, we recommend clients use Disaster Recovery Manager (DRM) to automate off-site volume management and the Tivoli Storage Manager recovery process. The solution diagram is shown in Figure 5-4.

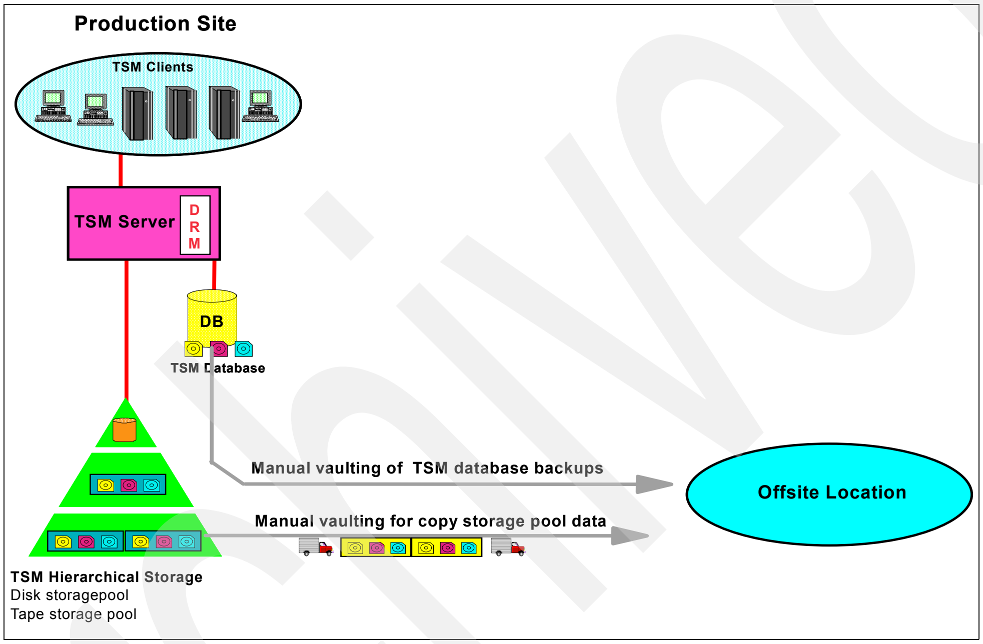

Figure 5-4 BC Tier 1 - Tivoli Storage Manager manual off-site vaulting - solution diagram

When recovery is required, Tivoli Storage Manager system will be recovered by using the off-site tape volumes, off-site database backup volumes, off-site configuration files, and the Business Continuity Plan file with the Tivoli Storage Manager DRM process. Then the Tivoli Storage Manager clients can restore their backup data from the new Tivoli Storage Manager system.

Tier 2 - Tivoli Storage Manager off-site manual vaulting with a hotsite

BC Tier 2 encompasses all the requirements of BC Tier 1 (off-site vaulting and recovery planning) plus it includes a hotsite. The hotsite has sufficient hardware and a network infrastructure able to support the installation's critical processing requirements. Processing is considered critical if it must be supported on hardware existing at the time of the disaster. As shown in Figure 5-5, backups are being taken, copies of these are created, and the copies are being stored at an off-site storage facility. There is also a hotsite available and the copy of the backups can be manually transported there from the off-site storage facility in the event of a disaster.

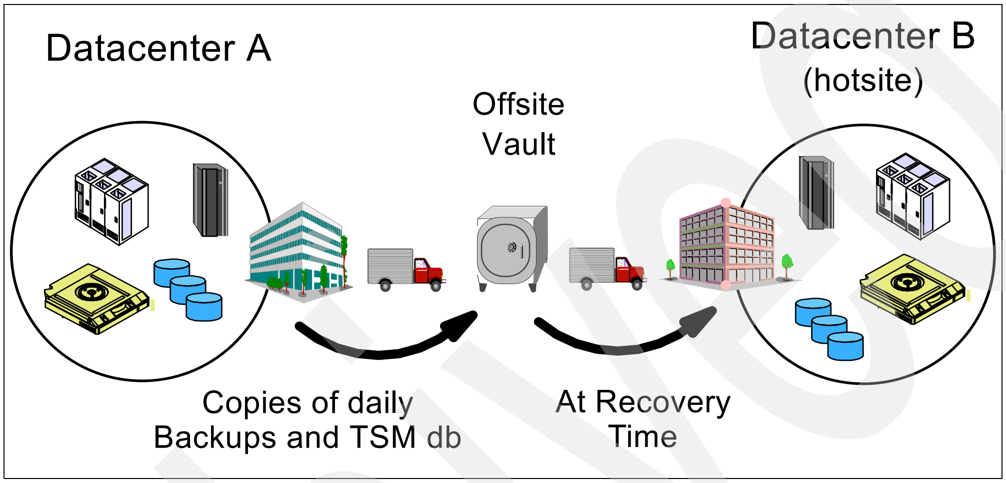

Figure 5-5 BC Tier 2 - Tivoli Storage Manager off-site manual vaulting with a hotsite, copies of daily backup

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

Typical length of time for recovery

With this solution, the typical length of time for recovery is normally more than a day.

Strategies and operations

The strategies and operations for this solution are the following:

- Use DRM to automate the Tivoli Storage Manager server recovery process and to manage off-site volumes.
- Manual vaulting of copies of Tivoli Storage Manager server's database backup and storage pool copies.
- Tivoli Storage Manager server installed at both locations.
- Vault Tivoli Storage Manager database backup, volume history information, device configuration information, Business Continuity Plan file, and storage pool copies at an off-site location.
- Consider the use of Tivoli Storage Manager server-to-server communications to enable enterprise configuration, enterprise event logging, event monitoring, and command routing.

Strategies and operations description

BC Tier 2 requires off-site vaulting with a hotsite; the implementation of this tier is based on Tier 1 implementation (see Figure 5-3 on page 60). In this tier, an additional Tivoli Storage Manager server should be installed at the hot site. We can use Tivoli Storage Manager server-to-server communication to enable enterprise configuration, enterprise event logging, event monitoring, and command routing. The solution is shown in Figure 5-6.

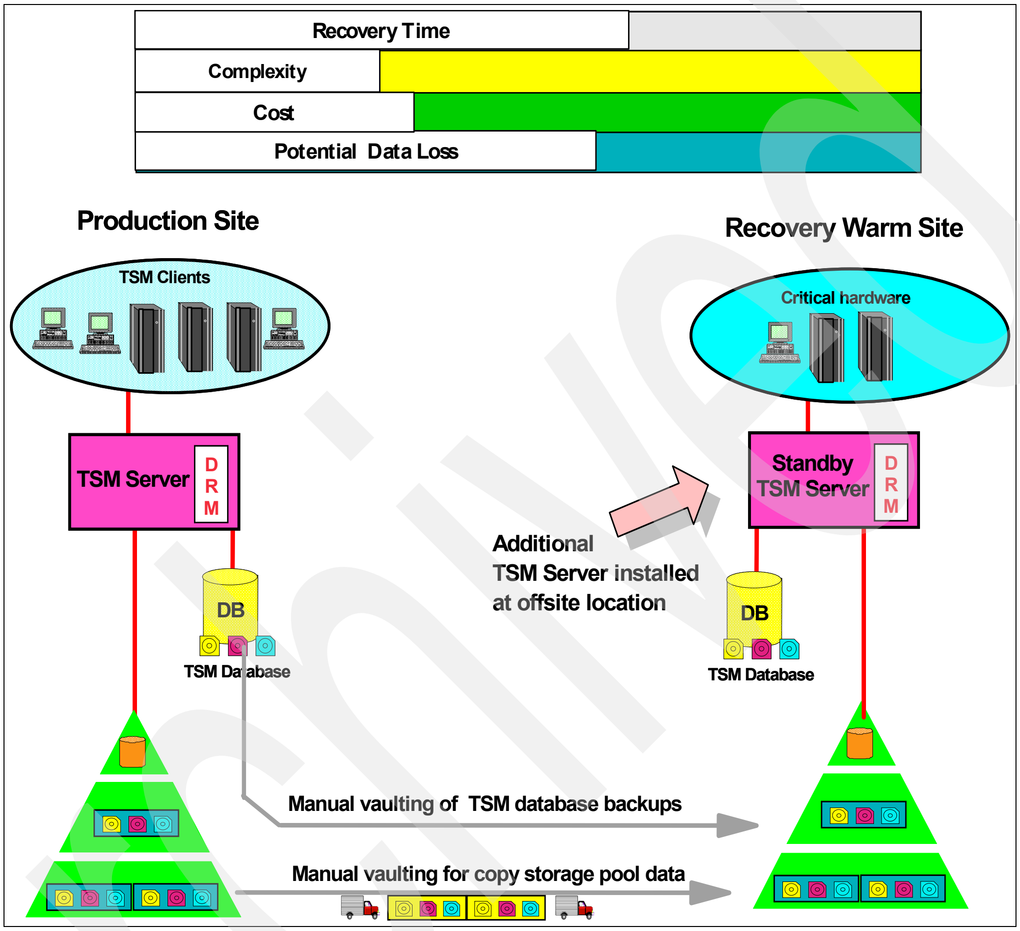

Figure 5-6 BC Tier 2 - Tivoli Storage Manager off-site manual vaulting with a hotsite

When the recovery process is required, the Tivoli Storage Manager recovery process will be run by using the Disaster Recovery Manager on the prepared Tivoli Storage Manager server. The recovery process uses the off-site tape volumes, off-site database backup volumes, off-site configuration files, and Business Continuity Plan file with the Tivoli Storage Manager DRM process, and then the Tivoli Storage Manager clients can restore their data from the new Tivoli Storage Manager system.

Tier 3 - Tivoli Storage Manager electronic vaulting

BC Tier 3 encompasses all of the components of BC Tier 2 (off-site backups, Business Continuity Plan, and hotsite). In addition, it supports electronic vaulting of the backup of the Tivoli Storage Manager database and storage pools. Electronic vaulting consists of electronically transmitting the backups to a secure facility, thus moving business-critical data off-site faster. This is accomplished via Tivoli Storage Manager's virtual vault capability. The receiving hardware must be physically separated from the primary site. As shown in Figure 5-7, backups of the Tivoli Storage Manager database and the entire or a subset of the storage pools can also be optionally made to physical media and manually moved to an off-site storage facility.

.

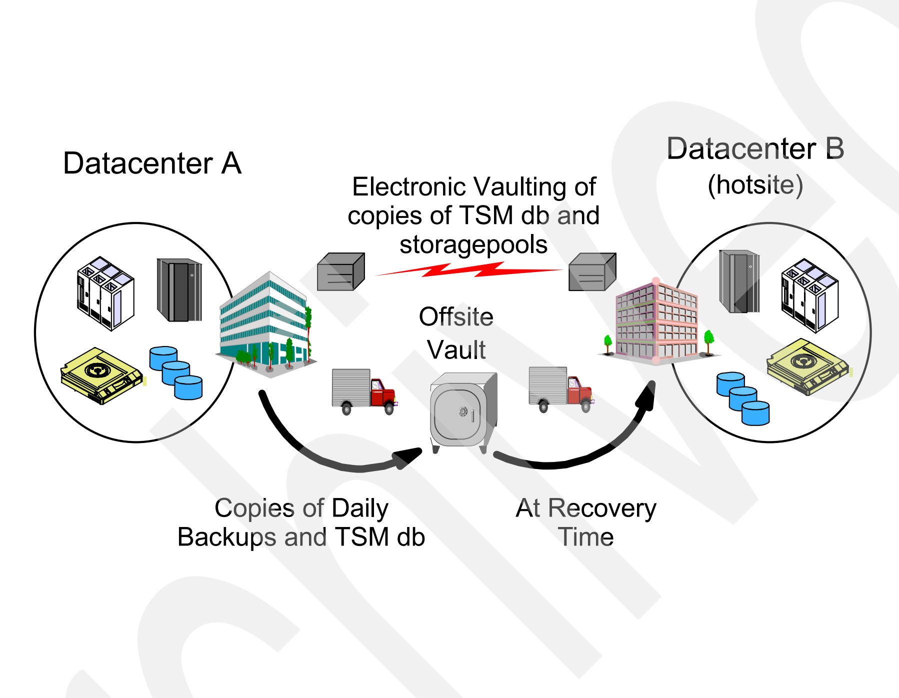
Datacenter A
Electronic Vaulting of copies of TSM db and storagepools
Datacenter B
(hotsite)
Offsite Vault
Copies of Daily Backups and TSM db
At Recovery Time

Figure 5-7 BC Tier 3 - Tivoli Storage Manager with electronic vaulting

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

Typical length of time for recovery

With this solution, the typical length of time for recovery is normally about one day.

Strategies and operations

The strategies and operations for this solution are:

- Tivoli Storage Manager virtual volumes over TCP/IP for storing Tivoli Storage Manager entities (Tivoli Storage Manager database backups, primary and copy storage pool backups, and DRM plan files) on a remote target server.
- A Tivoli Storage Manager server installed at both locations.
- Use DRM to automate the Tivoli Storage Manager server recovery process and to manage off-site volumes.
- Use Tivoli Storage Manager server-to-server communication to enable enterprise management features.

- Optionally, manually vault copies of the Tivoli Storage Manager database backup, volume history information, device configuration information, Business Continuity Plan file, and storage pools copies at an off-site location.

Strategies and operations description

Tivoli Storage Manager lets a server (a *source* server) store the results of database backups, export operations, storage pool operations, and a DRM plan on another server (a *target* server). The data is stored as *virtual* volumes, which appear to be sequential media volumes on the source server but that are actually stored as archive files on a target server. Virtual volumes can be any of the following:

- Database backups
- Storage pool backups
- Data that is backed up, archived, or space managed from client nodes
- Client data migrated from storage pools on the source server
- Any data that can be moved by EXPORT and IMPORT commands
- DRM plan files

Customers may decide to send all of this information or just a subset of it through FTP to the target server. This decision in part is based on the amount of bandwidth available. The data they choose not to send over the network can be taken off-site manually. See Figure 5-8.

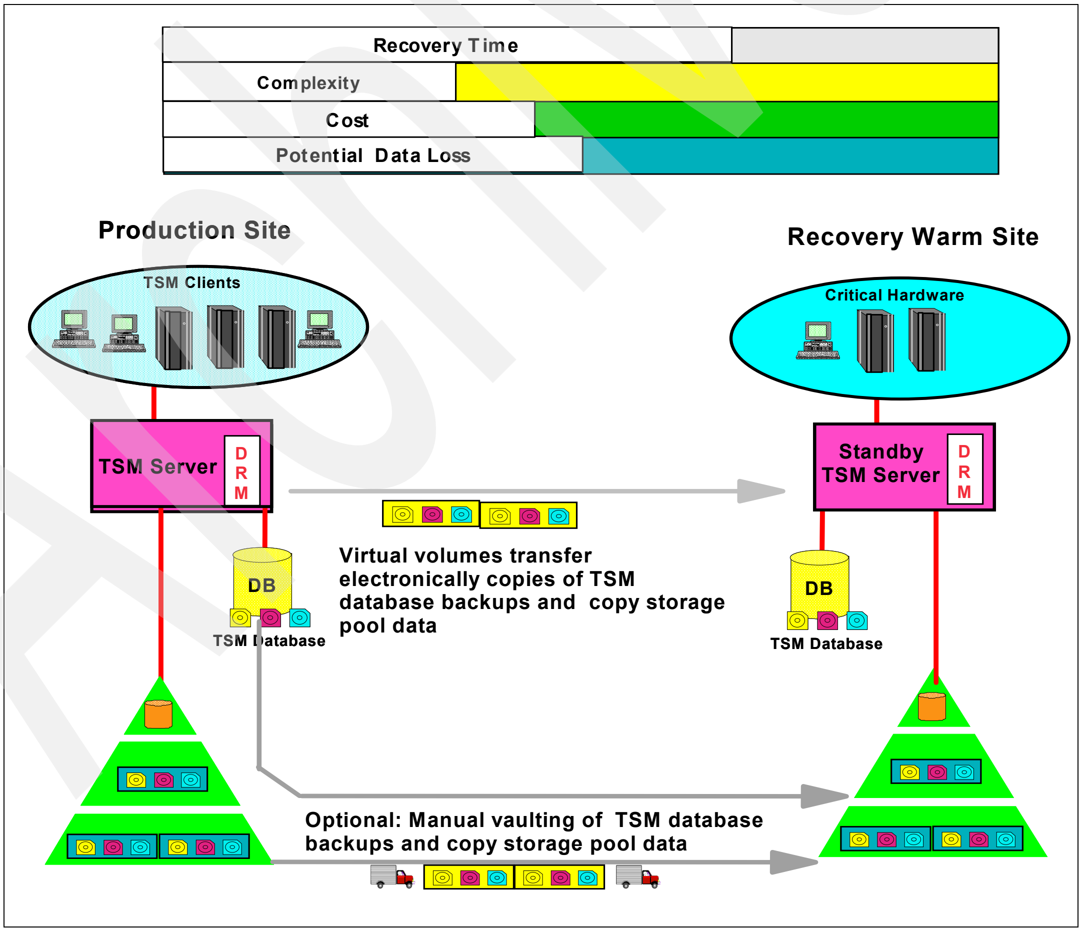
Recovery Time
Complexity
Cost
Potential Data Loss
Production Site
Recovery Warm Site
TSM Clients
Critical Hardware
TSM Server
DRM
Standby TSM Server
DRM
DB
TSM Database
DB
TSM Database
Virtual volumes transfer electronically copies of TSM database backups and copy storage pool data
Optional: Manual vaulting of TSM database backups and copy storage pool data

Figure 5-8 BC Tier 3 - Tivoli Storage Manager with electronic vaulting - solution diagram

BC Tier 4 - Tivoli Storage Manager with SAN attached duplicates

BC Tier 4 is defined by two data centers with SAN attached storage to keep mirrors of the Tivoli Storage Manager database and log. Separate Tivoli Storage Manager storage pools are located on storage residing in the two locations. SAN technology is critical to facilitate the Tivoli Storage Manager database and log mirroring and Tivoli Storage Manager simultaneous data writes occurring to local and remote devices. See Figure 5-9. The hotsite will also have backups of the Tivoli Storage Manager database storage pools stored either as virtual volumes or physical backup tapes to protect against database corruption and ensure the ability to restore the Tivoli Storage Manager server to a different point-in-time.

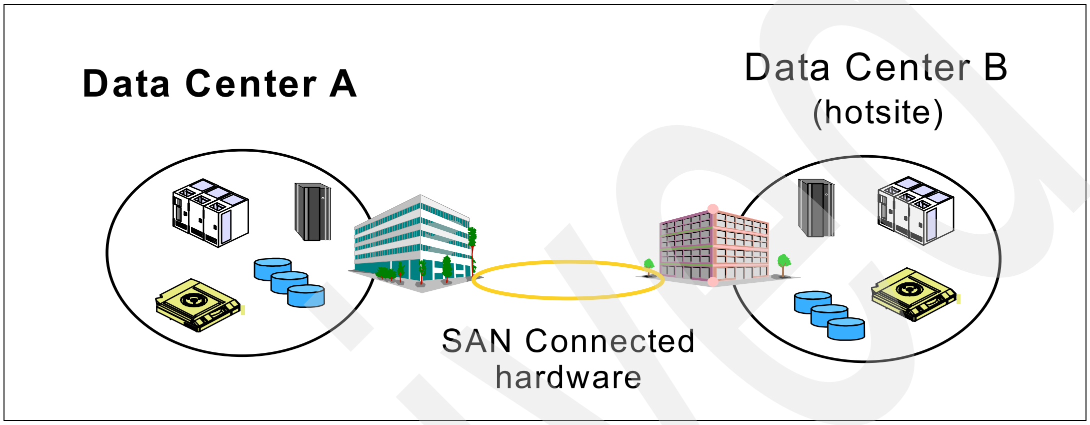

Figure 5-9 BC Tier 4 - Tivoli Storage Manager with duplicate SAN attached hardware

Typical length of time for recovery

With this solution, the typical length of time for recovery is usually up two to four hours.

Strategies and operations

The strategies and operations for this solution are:

- Tivoli Storage Manager database and log volumes mirrored sequentially on SAN hardware at a different location from the primary volumes.
- Use Tivoli Storage Manager's simultaneous copy storage pool write to copy data to a Tivoli Storage Manager storage pool located on SAN hardware that is located at a different locations from the primary storage pool volume.
- Tivoli Storage Manager server code is installed at the second location but is not running.
- Use DRM to automate the Tivoli Storage Manager server recover process and to manage off-site volumes.
- Use Tivoli Storage Manager server-to-server communications to enable enterprise management features.
- Vault copies of Tivoli Storage Manager database backups and storage pool copies either manually (BC Tier 2) or with virtual volumes (BC Tier 3) daily.

Strategies and operations description

BC Tier 4 requires a SAN infrastructure to allow mirroring of the Tivoli Storage Manager database and log (using Tivoli Storage Manager's mirroring capability) to off-site storage. The SAN attached storage pool will be written to concurrently with the local storage pool, using Tivoli Storage Manager's simultaneous copy storage pool write. The simultaneous write will create a copy of the data to both the primary and secondary storage pool at the initial time of write. The Tivoli Storage Manager database and storage pools will still be backed up daily, and either electronically or manually vaulted off-site to protect against corruption to the Tivoli

Storage Manager database, and to allow for the server to be restored to a previous time. See Figure 5-10.

Another Tivoli Storage Manager should be installed at the off-site location, with access to the database and log mirror and to the copy storage pools. However, this Tivoli Storage Manager server is not started, unless the primary Tivoli Storage Manager server fails. The two Tivoli Storage Manager servers should never be run at the same time, as this might corrupt the database and log.

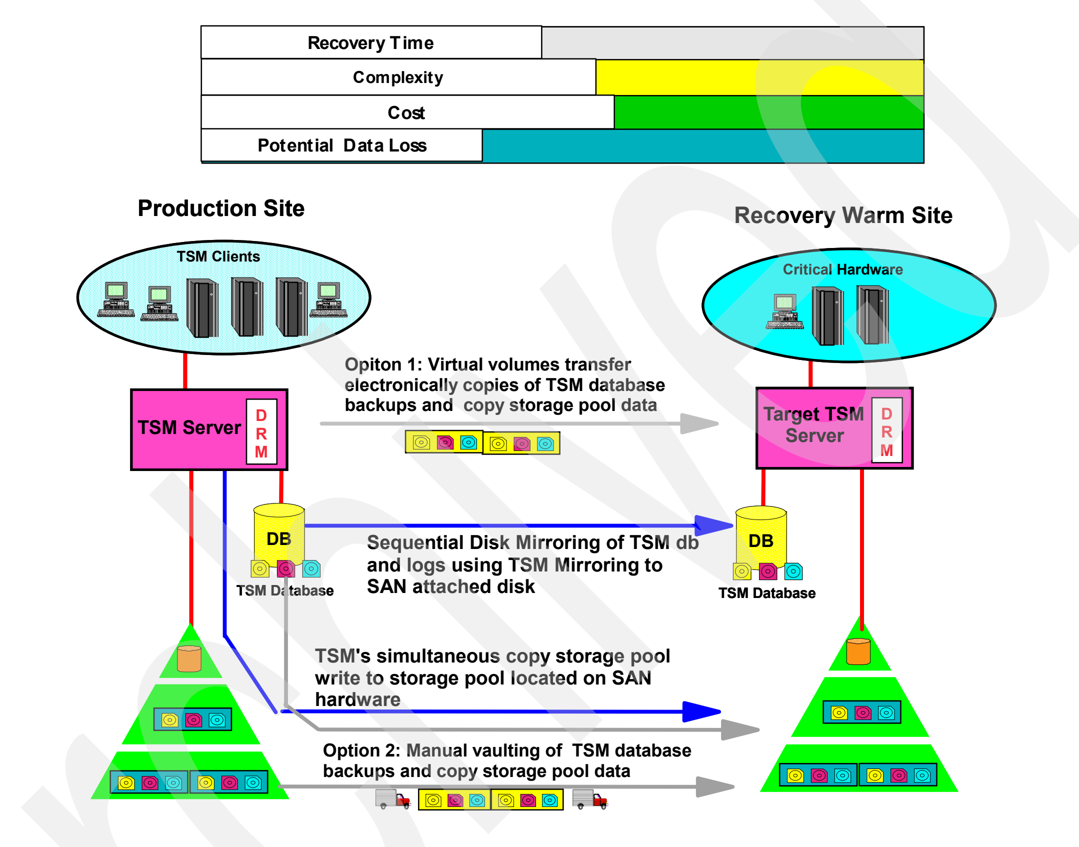
Recovery Time
Complexity
Cost
Potential Data Loss
Production Site
Recovery Warm Site
TSM Clients
Critical Hardware
TSM Server
D
R
M
Opiton 1: Virtual volumes transfer electronically copies of TSM database backups and copy storage pool data
Target TSM Server
D
R
M
DB
TSM Database
Sequential Disk Mirroring of TSM db and logs using TSM Mirroring to SAN attached disk
DB
TSM Database
TSM's simultaneous copy storage pool write to storage pool located on SAN hardware
Option 2: Manual vaulting of TSM database backups and copy storage pool data

Figure 5-10 BC Tier 4 - Tivoli Storage Manager with SAN attached database, log mirrors and copy storage pools

BC Tier 5 - IBM Tivoli Storage Manager clustering

BC Tier 5 is defined by two data centers utilizing clustering technology to provide automated failover and failback capabilities. Clustering applications like HACMP and MSCS provide switchover for the Tivoli Storage Manager server. The hotsite will also have Tivoli Storage Manager database backups and copies of the Tivoli Storage Manager storage pools stored either as virtual volumes or physical backup tapes to protect against database corruption and ensure the ability to restore the Tivoli Storage Manager server back to a different point-in-time. See Figure 5-11.

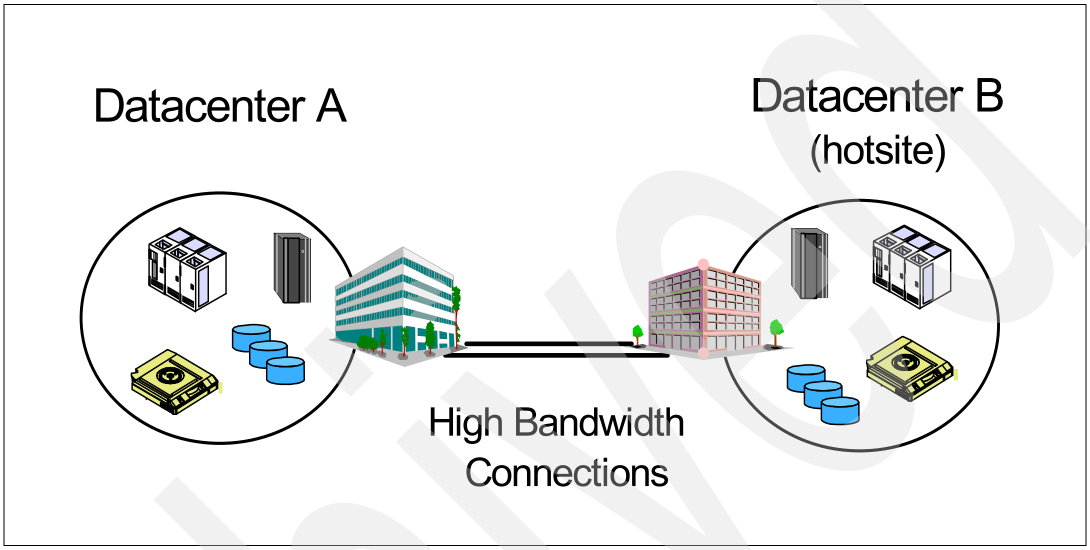

Figure 5-11 BC Tier 5 - Tivoli Storage Manager clustering

Typical length of time for recovery

With this solution, the typical length of time for recovery is normally a few minutes.

Strategies and operations

The strategies and operations for this solution are:

- Utilize HACMP or MSCS to cluster the Tivoli Storage Manager server and clients.
- Install the Tivoli Storage Manager server at the second location in either an active-active, or active-passive cluster.
- Utilize Tivoli Storage Manager's SCSI device failover for SCSI tape libraries, or consider using SAN attached storage pools.
- Use DRM to automate the Tivoli Storage Manager server recover process and to manage off-site volumes.
- Use Tivoli Storage Manager server-to-server communications to enable enterprise management features.
- Vault daily copies of Tivoli Storage Manager database backups and storage pool copies either manually (BC Tier 2) or with virtual volumes (BC Tier 3).

Strategies and operations description

Careful planning and testing is required when setting up a clustering environment. Once the cluster environment is created, Tivoli Storage Manager can utilize the technology to perform automatic failovers for the Tivoli Storage Manager server and clients (including HSM and Storage Agents). See Figure 5-12. Clustering is covered in great detail in the Tivoli Storage Manager manuals and in the IBM Redbook *IBM Tivoli Storage Manager in a Clustered Environment*, SG24-6679.

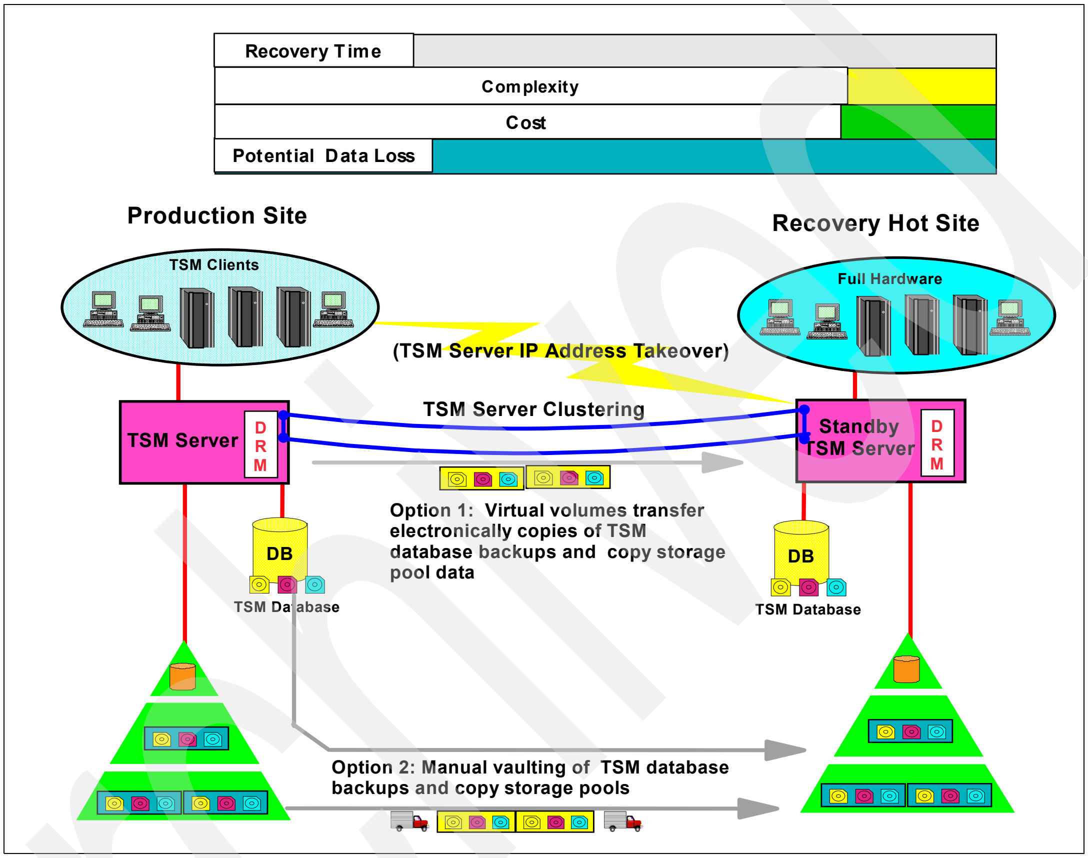

Recovery Time
Complexity
Cost
Potential Data Loss
Production Site
Recovery Hot Site
TSM Clients
Full Hardware
(TSM Server IP Address Takeover)
TSM Server Clustering
TSM Server
D R M
Standby TSM Server
D R M
Option 1: Virtual volumes transfer electronically copies of TSM database backups and copy storage pool data
DB
TSM Database
DB
TSM Database
Option 2: Manual vaulting of TSM database backups and copy storage pools

Figure 5-12 BC Tier 5 - Clustered Tivoli Storage Manager servers and clients

BC Tier 6 - IBM Tivoli Storage Manager running in a duplicate site

BC Tier 6 encompasses zero data loss and immediate, automatic transfer to the secondary platform. The two sites are fully synchronized via a high-bandwidth connection between the primary site and the hotsite. The two systems are advanced coupled, allowing automated switchover from one site to the other when required. See Figure 5-13.

In this tier, there are two independent Tivoli Storage Manager setups. One is dedicated to backing up the data on the primary site, the other is dedicated to backing up the data on the other site. Copies of each locations' Tivoli Storage Manager database and storage pools need to be taken off-site daily either through virtual volumes or manually.

This is the most expensive BC solution, as it requires duplicating, at both sites, all of the hardware and applications, and then keeping the data synced between the sites. However, it also offers the speediest recovery by far.

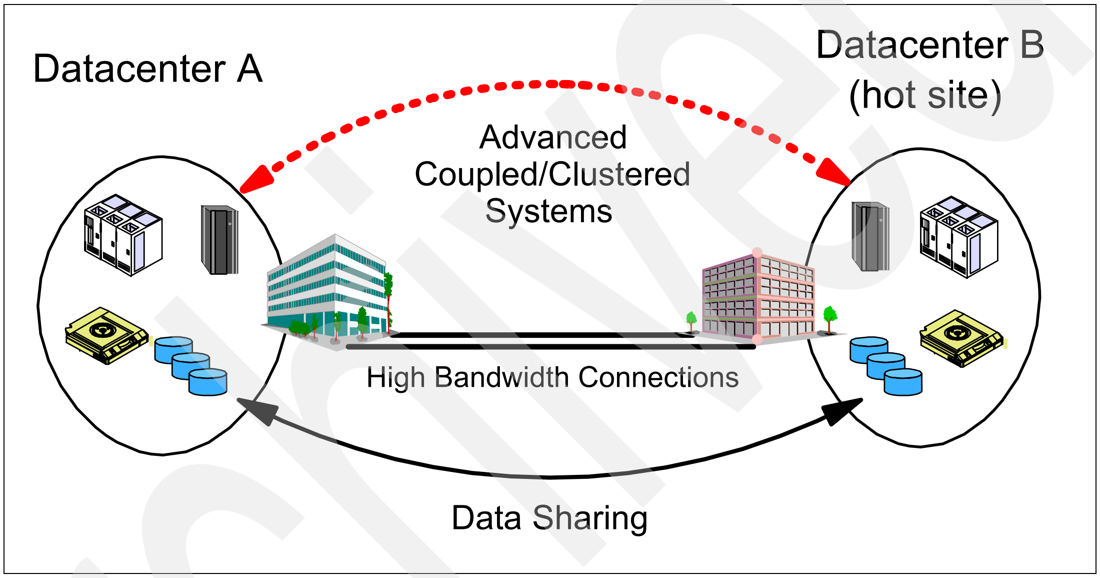

Figure 5-13 BC Tier 6 - Dual production Tivoli Storage Manager servers running in a zero data loss environment

Typical length of time for recovery

With this solution, the typical length of time for recovery is normally almost instantaneous.

Strategies and operations

The strategies and operations for this solution are following:

- Independent Tivoli Storage Manager servers installed at both locations and backing up that locations' data.
- Tivoli Storage Manager virtual volumes over TCP/IP connection to allow storage of Tivoli Storage Manager entities (Tivoli Storage Manager database backups, primary and copy storage pool backups, and DRM plan files) on a remote target server.
- Use DRM to automate the Tivoli Storage Manager server recover process and to manage off-site volumes.
- Use Tivoli Storage Manager server-to-server communications to enable enterprise management features.
- Optionally, manually vault copies of Tivoli Storage Manager database backup, volume history information, device configuration information, Business Continuity Plan file, and storage pools copies at an off-site location.

Strategies and operations description

The data at primary and off-site location are fully synchronized utilizing a high-bandwidth connection. The two systems are advanced coupled, allow an automated switchover from one site to the other when required. A Tivoli Storage Manager setup is installed at each location and run independently, creating dual production systems. Each Tivoli Storage Manager server is protected using BC Tier 2 or BC Tier 3 technology. See Figure 5-14.

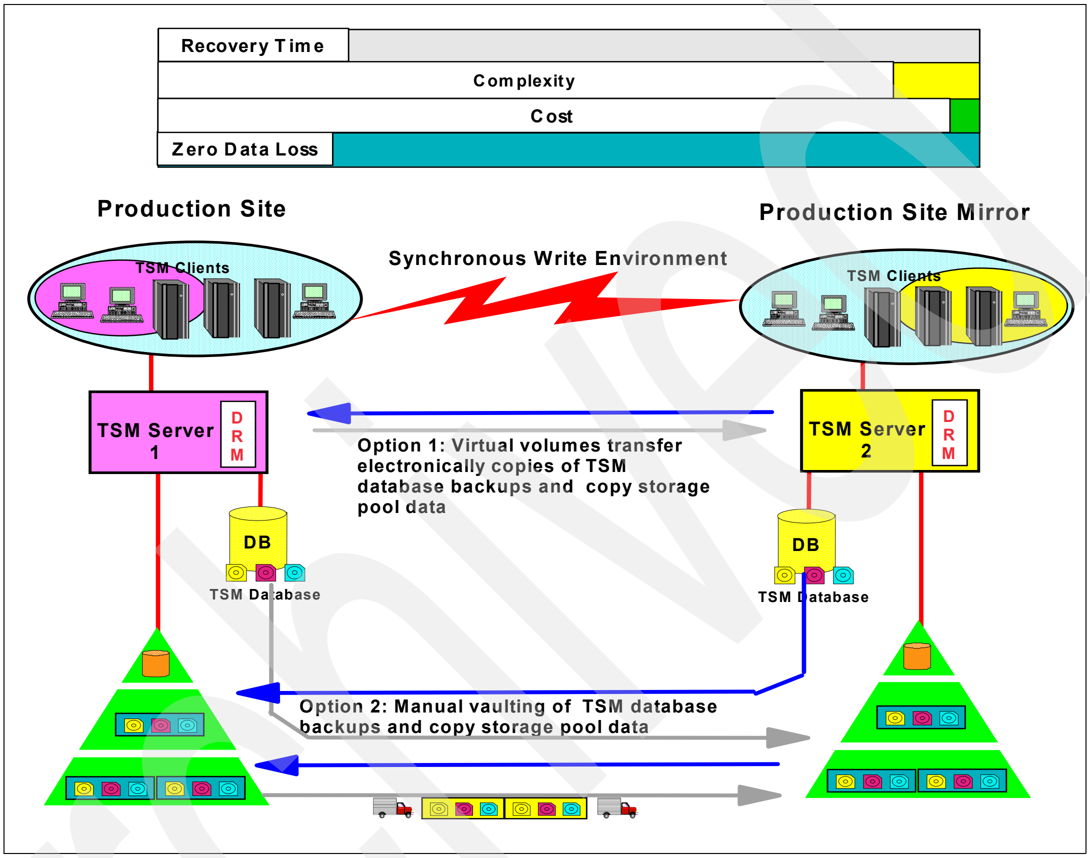
Recovery Time
Complexity
Cost
Zero Data Loss
Production Site
Production Site Mirror
Synchronous Write Environment
TSM Clients
TSM Clients
TSM Server 1
DRM
TSM Server 2
DRM
Option 1: Virtual volumes transfer electronically copies of TSM database backups and copy storage pool data
DB
TSM Database
DB
TSM Database
Option 2: Manual vaulting of TSM database backups and copy storage pool data

Figure 5-14 BC Tier 6 - Dual production Tivoli Storage Manager servers running in a zero data loss environment

5.2.3 Integrated solution capabilities of IBM Tivoli Storage Manager

Additional optional IBM Tivoli Storage Manager products are available that provide integrated solution capabilities. These include:

- IBM Tivoli Storage Manager for Copy Services - Data Protection for Exchange

  This provides integration of Tivoli Storage Manager with VSS providers on supported storage systems to make fast, snapshot-based backups of Microsoft Exchange databases. The snapshot can then be copied to a Tivoli Storage Manager server for longer term, policy-managed storage. Optionally, the Tivoli Storage Manager backup of the snapshot can be performed by a separate server, freeing up CPU resources on the production Exchange server, known as offloaded backup. On a SAN Volume Controller, *instant restore* is possible, where a snapshot can be rapidly restored back to the original source disks using the hardware FlashCopy restore.

- IBM Tivoli Storage Manager for Advanced Copy Services, with the following modules:
  - Data Protection for IBM Disk Storage and SAN Volume Controller for mySAP with DB2
  - Data Protection for IBM Disk Storage and SAN Volume Controller for mySAP with Oracle
  - Data Protection for IBM Disk Storage and SAN Volume Controller for Oracle
  - DB2 UDB Integration Module and Hardware Devices Snapshot™ Integration Module
  - Data Protection for ESS for Oracle
  - Data Protection for ESS for mySAP DB2 UDB
  - Data Protection for ESS for mySAP

  These products provide SAP certified integrated backup for SAP (on DB2 or Oracle) databases, using FlashCopy functions of the DS6000, DS8000, SAN Volume Controller, or ESS. They can also backup DB2 or Oracle databases with FlashCopy.

5.2.4 Further information

For more information about products, solutions and implementation, we recommend you refer to these IBM Redbooks and Web sites:

- IBM Redbooks
  - *Disaster Recovery Strategies with Tivoli Storage Management*, SG24-6844
  - *IBM Tivoli Storage Management Concepts*, SG24-4877
  - *IBM Tivoli Storage Manager Version 5.3 Technical Guide*, SG24-6638
  - *Using IBM Tivoli Storage Manager to Back Up Microsoft Exchange with VSS*, SG24-7373
- Web sites:

  http://www.ibm.com/software/tivoli/products/storage-mgr/
  http://www.ibm.com/storage/th/disk/ds6000/
  http://www.ibm.com/storage/th/disk/ds8000/
  http://www.ibm.com/servers/aix/products/ibmsw/high_avail_network/hacmp.html
  http://www.ibm.com/servers/aix/products/ibmsw/high_avail_network/hageo_georm.html

More information about Tivoli Storage Manager is in Appendix C, “Tivoli Storage Manager family overview” on page 151.

5.2.5 IBM Tivoli Storage Manager summary

At an enterprise software level, Tivoli Storage Manager policy must meet overall business requirements for data availability, data security, and data retention. Enterprise policy standards can be established and applied to all systems during the policy planning process. At the systems level, RTO and RPO requirements vary across the enterprise. Systems classifications and data classifications typically delineate the groups of systems and data along with their respective RTO/RPO requirements.

In view of Business Continuity, enhancement of the Tivoli Storage Manager system to support the highest BC tier is a key for continuity for the Backup and Restore services. The tier-based IBM Tivoli Storage Manager solutions given above are guidelines to improve your current Tivoli Storage Manager solution to cover your business continuity requirements based on the BC tier. These solutions allow you to protect enterprise business systems by having a zero data lost on business backup data and have continuity of backup and recovery service.

5.3 IBM Data Retention 550

This section discusses the need for storage archiving with the IBM System Storage DR550 and DR550 Express, that can address data retention and other regulatory compliance requirements.

More detail on the overall business, legal, and regulatory climate, which is the underlying driving force behind the growth in retention-managed data, can be found in *Understanding the IBM System Storage DR550*, SG24-7091.

5.3.1 Retention-managed data

Beyond laws and regulations, data often needs to be archived and managed simply because it represents a critical company asset.

Examples of such data include contracts, CAD/CAM designs, aircraft build and maintenance records, and e-mail, including attachments, instant messaging, insurance claim processing, presentations, transaction logs, Web content, user manuals, training material, digitized information, such as check images, medical images, historical documents, photographs, and many more.

The characteristics of such data can be very different in their representation, size, and industry segment. It becomes apparent that the most important attribute of this kind of data is that it needs to be retained and managed, thus it is called *retention-managed data*.

Retention-managed data is data that is written once and is read rarely (sometimes never). Other terms abound to describe this type of data, such as reference data, archive data, content data, or other terms that imply that the data cannot be altered.

Retention-managed data is data that needs to be kept (retained) for a specific (or unspecified) period of time, usually years.

Retention-managed data applies to many types of data and formats across all industries. The file sizes can be small or large, but the volume of data tends to be large (multi-terabyte to petabytes). It is information that might be considered of high value to an organization; therefore, it is retained near-line for fast access. It is typically read infrequently and thus can be stored on economical disk media such as SATA disks; depending on its nature, it can be migrated to tape after some period.

It is also important to recognize what does not qualify as retention-managed data. It is not the data that changes regularly, known as *transaction data* (account balance, inventory status, and orders today, for example). It is not the data that is used and updated every business cycle (usually daily), and it is not the backup copy of this data. The data mentioned here changes regularly, and the copies used for backup and business continuity are there for exactly those purposes, meaning backup and business continuity. They are there so that you can restore data that was deleted or destroyed, whether by accident, a natural or human-made disaster, or intentionally.

5.3.2 Storage and data characteristics

When considering the safekeeping of retention-managed data, companies also need to consider storage and data characteristics that differentiate it from transactional data.

Storage characteristics of retention-managed data include:

- Variable data retention periods: Usually a minimum of a few months, up to forever.
- Variable data volume: Many customers are starting with 5 to 10 TB of storage for this kind of application (archive) in an enterprise. It also usually consists of a large number of small files.
- Data access frequency: Write once, read rarely, or read never. See data life cycle in the following list.
- Data read/write performance: Write handles volume; Read varies by industry and application.
- Data protection: Pervasive requirements for non-erasability, non-rewritability, and destructive erase (data shredding) when the retention policy expires.

Data characteristics of retention-managed data include:

- Data life cycle: Usage after capture, 30 to 90 days, and then near zero. Some industries have peaks that require access, such as check images in the tax season.
- Data rendering after long-term storage: Ability to view or use data stored in a very old data format (say after 20 years).
- Data mining: With all this data being saved, we think there is intrinsic value in the content of the archive that could be exploited.

5.3.3 IBM strategy and key products

Regulations and other business imperatives, as we just briefly presented, stress the need for data retention processes and tools to be in place.

IBM provides a comprehensive and open set of solutions to help. IBM has products that provide content management, data retention management, and sophisticated storage management, along with the storage systems to house the data.

To specifically help companies with their risk and compliance efforts, the IBM Risk and Compliance framework is another tool designed to illustrate the infrastructure capabilities needed to help address the myriad of compliance requirements. Using the framework, organizations can standardize the use of common technologies to design and deploy a compliance architecture that may help them deal more effectively with compliance initiatives.

> **Important:** The IBM offerings are intended to help customers address the numerous and complex issues relating to data retention in regulated and non-regulated business environments. Nevertheless, each customer’s situation is unique, and laws, regulations, and business considerations impacting data retention policies and practices are constantly evolving. Customers remain responsible for ensuring that their information technology systems and data retention practices comply with applicable laws and regulations, and IBM encourages customers to seek appropriate legal counsel to ensure their compliance with those requirements. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer is in compliance with any law.

For more detailed information about the IBM Risk and Compliance framework, visit:

http://www.ibm.com/software/info/openenvironment/rcf/

Key products of the IBM data retention and compliance solutions are IBM System Storage Archive Manager and IBM DB2 Content Manager, along with any needed disk-based and tape-based storage:

- IBM DB2 Content Manager for Data Retention Compliance is a comprehensive software platform combining IBM DB2 Content Manager, DB2 Records Manager, and DB2 CommonStore, and services that are designed to help companies address the data retention requirements of SEC, NASD, and other regulations. This offering provides archival and retention capabilities to help companies address the retention of regulated and non-regulated data, providing increased efficiency with fast, flexible data capture, storage, and retrieval.
- IBM System Storage Archive Manager offers expanded policy-based data retention capabilities that are designed to provide non-rewritable, non-erasable storage controls to prevent deletion or alteration of data stored using IBM System Storage Archive Manager. These retention features are available to any application that integrates with the open IBM System Storage Manager API.

Key technologies for IBM data retention and archiving solutions include:

- IBM SystemStorage DS4000 with Serial Advanced Technology Attachment (SATA) disk drives to provide near-line storage at an affordable price. It is also capable of Enhanced Remote Mirroring to a secondary site.
- Write Once Read Many (WORM) media technology (in supported IBM tape drives and libraries). Once written, data on the cartridges cannot be overwritten (to delete the data, the tape must be physically destroyed). This capability is of particular interest to clients that need to store large quantities of electronic records to meet regulatory and internal audit requirements.

Figure 5-15 shows where the DR550 series, built on some of these key products and technologies, fits.

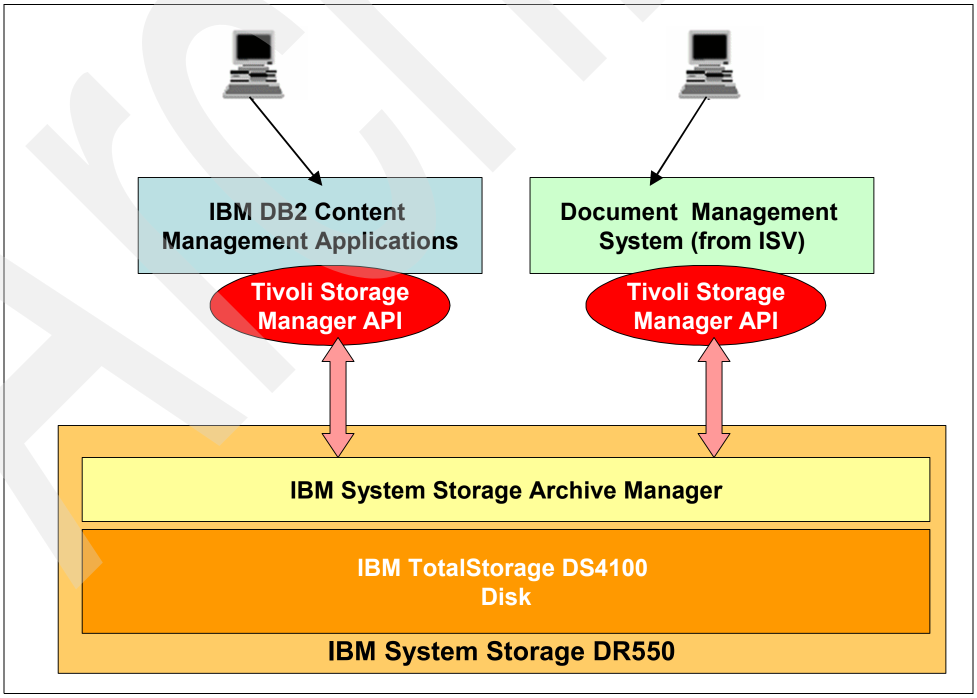
IBM DB2 Content Management Applications
Document Management System (from ISV)
Tivoli Storage Manager API
Tivoli Storage Manager API
IBM System Storage Archive Manager
IBM TotalStorage DS4100 Disk
IBM System Storage DR550

Figure 5-15 DR550 context diagram

The main focus of the IBM System Storage DR550 is to provide for a secure storage system, where deletion or modification of data is completely disallowed except through a well-defined retention and expiration policy.

The IBM System Storage DR550 is the repository for regulated business information. It does not create the information. A complete solution includes applications that gather information, such as e-mail and instant messaging, content management applications, such as IBM DB2 Content Manager, or other document management systems from independent software vendors that can index, archive, and later present this information to compliance reviewers, and finally storage systems that retain the data in accordance with regulations.

5.3.4 IBM System Storage DR550

IBM System Storage DR550 (Figure 5-16), one of the IBM Data Retention offerings, is an integrated offering for clients that need to retain and preserve electronic business records. It is designed to help store, retrieve, manage, and retain regulated and non-regulated data. In other words, it is not just an offering for compliance data, but can also be an archiving solution for other types of data.

Integrating IBM System p5™ servers (using POWER5™ processors) with IBM System Storage hardware products and IBM System Storage Archive Manager software, this system is specifically designed to provide a central point of control to help manage growing compliance and data retention needs.

The DR550 brings together off-the-shelf IBM hardware and software products. The hardware comes already mounted in a secure rack. The software is preinstalled and to a large extent preconfigured. The system's compact design can help with fast and easy deployment and incorporates an open and flexible architecture.

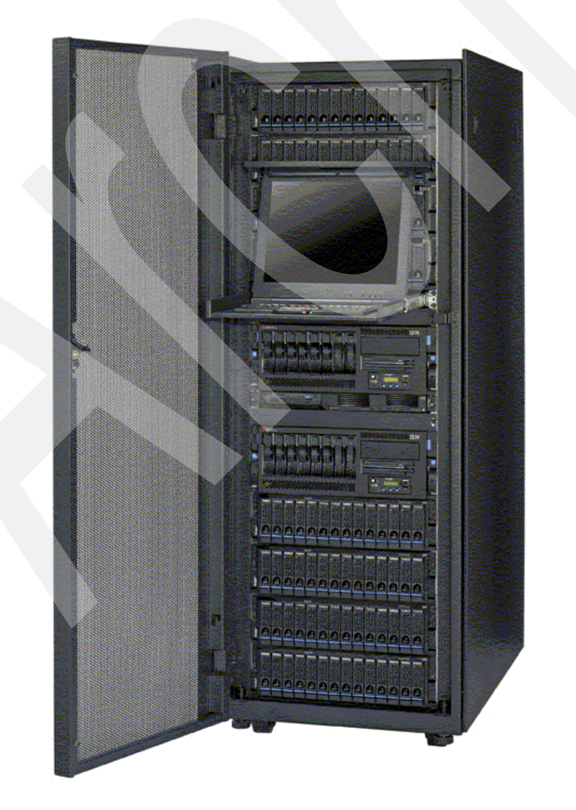

Figure 5-16 DR550 dual node with console kit

5.4 Summary

Typically, the backup and restore segment has been positioned in Business Continuity Tier 1, Tier 2, and Tier 3.

In 5.2, “IBM Tivoli Storage Manager overview” on page 56, we showed that Backup and Restore can be applied in Business Continuity Tier 4, Tier 5, and Tier 6 as well.

Furthermore, archive data also plays an important role for business continuity, which we have discussed in 5.3, “IBM Data Retention 550” on page 73.

The Backup and Restore can be flexibly implemented with Enterprise Data Management software, such as Tivoli Storage Manager. Enterprise Data can be policy managed to achieve various business needs and objectives on business continuity.

VSS snapshot and FlashCopy can be easily exploited with Tivoli Storage Manager for Copy Services and Tivoli Storage Manager for Advanced Copy Services.

Advanced WORM Tape or WORM Disk Technology can be easily exploited with Tivoli Storage Manager on storage pooling with WORM media.

Finally, with the availability of Encrypted Key Management capability of Tivoli Storage Manager, sensitive data that requires higher security protection can now be easily deployed on TS1120 tape. Backup and Restore is making another step forward on business deployment with Information Technology.

5.5 System z backup and restore software

There are very well established products and methods to back up System z environments, within the DFSMS family of products. We will simply summarize these here and refer the reader to IBM Redbooks, such as *Z/OS V1R3 and V1R5 DFSMS Technical Guide*, SG24-6979.

5.5.1 DFSMSdss

The primary function of DFSMSdss is to move and copy data. It can operate at both the logical and physical level and can move or copy data between volumes of like and unlike device types. DFSMSdss can make use of the following two features of the DS6000 and DS8000:

- FlashCopy: A point-in-time copy function that can quickly copy data from a source location to a target location.
- Concurrent Copy: A copy function that generates a copy of data while applications are updating that data.

DFSMSdss does not communicate directly with the disk system to use these features; this is performed by a component of DFSMSdfp, the System Data Mover (SDM).

5.5.2 DFSMShsm

Hierarchical Storage Manager (DFSMShsm) is a disk storage management and productivity tool for managing low-activity and inactive data. It provides backup, recovery, migration, and space management functions as well as full function Disaster Recovery. DFSMShsm improves disk use by automatically managing both space and data availability in a storage hierarchy. DFSMShsm can also be useful in a backup/restore situation. At a time specified by the installation, DFSMShsm checks to see whether data sets have been updated. If a data set has been updated, then it can have a backup taken. If a data sets are damaged or accidentally deleted, then it can be recovered from a backup copy. There can be more than one backup version, which assists in the recovery of a data set which has been damaged for some time, but this has only recently been detected.

DFSMShsm also has a feature called Fast Replication that invokes FlashCopy for volume-level replication.

5.5.3 z/VM Utilities

VM utilities for backup and restore include:

- DDR (DASD Dump and Restore), a utility to dump, copy, or print data that resides on z/VM user minidisks or dedicated DASDs. The utility may also be used to restore or copy DASD data, which resides on z/VM user tapes
- DFSMS/VM™
- z/VM Backup and Restore Manager

Refer to your z/VM operating system documentation for more information.

Part 3

End-to-end Business Continuity solution components

In addition to the three segments of Business Continuity solutions, there are other key components that are needed. These include:

- Networking
- Application and Database Considerations
- Small and Medium Business Considerations

6

Networking and inter-site connectivity options

A properly configured network infrastructure is critical for the efficient transmission of data from one location to another. Typically, Continuous Availability and rapid recovery IT Business Continuity solutions use a two site implementation, providing a fully redundant configuration mirroring the mission critical environment. A redundancy site allows immediate recovery of system and application functionality.

Today's trends are moving to Three-Site Recovery, where two of the sites are separated by up to 100 km of fiber to provide Continuous Availability and Business Continuity protection against metropolitan events. The third site can be an out-of-region DR site separated by unlimited distances to provide DR protection against regional events.

This chapter provides basic information regarding networking options that should be considered prior to designing Two-site or Three-site network infrastructures. It is an often underestimated, but critically important aspect of any storage networking or data mirroring infrastructure. It includes the basics of network transport as they apply to Disaster Recovery and covers Network Topology, Fiber Transport Options, Wavelength Division Multiplexing (WDM), and Channel Extension.

The importance of network design and sizing cannot be emphasized enough. Having sufficient bandwidth for any data replication solution, even an asynchronous solution, is absolutely critical to success. It is beyond the scope of this IBM Redbook to provide detailed guidance, but a network expert should be an early and constant part of the BC solution design team.

6.1 Network topologies

Remote storage and remote backup are key components in either a Continuous Availability, Rapid Data Recovery, or Backup and Restore solution. Establishing a remote storage solution presents challenges and decisions as to which extension method is used to connect the Business Continuity sites. Depending on distance, transmission latency could have a significant impact on if Metro Mirror, Global Mirror, or a cascaded implementation is selected.

Figure 6-1 shows two-site and three-site recovery (cascaded) topologies that are occurring in multiple industries today. It is no longer only the domain of the finance industry or of only a select few high end enterprises implementing this high-end IT Business Continuity solution. We are seeing three site requirements from manufacturing, from retail, from transportation, from telecom, and from other industries. Interestingly, IT organizations with local two-site high availability are interested in adding out-of-region recovery capability, while IT organizations with out-of-region recovery capability are interested in adding local high availability. As consolidation and IT infrastructure streamlining continues, the interest in three-site recovery continues to grow.

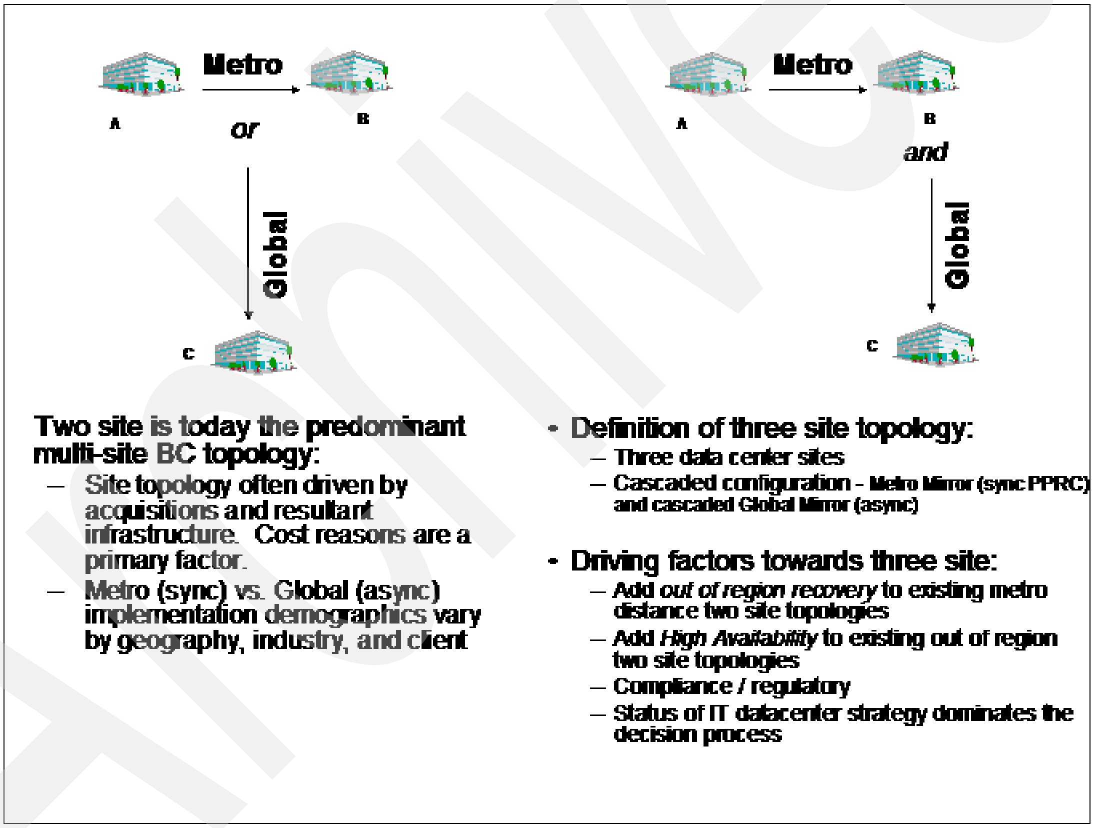

Two site is today the predominant multi-site BC topology:
- Site topology often driven by acquisitions and resultant infrastructure. Cost reasons are a primary factor.
- Metro (sync) vs. Global (async) implementation demographics vary by geography, industry, and client

- **Definition of three site topology:**
  - Three data center sites
  - Cascaded configuration - Metro Mirror (sync PPRC) and cascaded Global Mirror (async)
- **Driving factors towards three site:**
  - Add *out of region recovery* to existing metro distance two site topologies
  - Add *High Availability* to existing out of region two site topologies
  - Compliance / regulatory
  - Status of IT datacenter strategy dominates the decision process

Figure 6-1 Two- and three-site topologies

If you are involved in designing, building, or testing network technologies or channel extension for a IT Business Continuity solution, you must be familiar with network protocols used. The data transport selected will depend on your answers to the following question:

- Distance: How far away is the remote site?
- Bandwidth: How much data is transported in what time frame?
- Recovery Point Objective: How much data can you afford to lose?
- Recovery Time Objective: How long can you afford to be without your systems?
- Network Recovery Objective: What are the applications’ response time requirements?

- Storage-over-distance options: What does your hardware support?

The remainder of this chapter provides basic information about networking options to be considered for designing site-to-site channel extension communication.

6.2 Fiber transport

This section defines the two major options in fiber transport technology: dedicated fiber and SONET.

6.2.1 Dedicated fiber

Dedicated fiber is fiber that is not shared, and thus not lit by other users. This is typically a privately owned fiber route that, unlike shared SONET rings, only has light passing through it when its owner connects devices. Because it is privately owned and there are no competing users, the business gets full use of the bandwidth provisioned at any time. The downside is that it tends to be an expensive solution and usually requires some form of multiplexing to control fiber transport costs as the business grows. Due to technical limitations, dedicated fiber can only be used for relatively short distances, especially when compared to SONET. Dedicated fiber is often used in Data Center Point-to-Point configurations.

We can contrast this with dark fiber, which is dormant, unused (that is, unlit) fiber.

6.2.2 SONET

Synchronous Optical NETwork (SONET) (SDH in Europe) is a standard for transporting data across public telecommunication rings. Typically, a business will contract with a telecommunications or Network Connectivity provider for a specific amount of bandwidth. The customer’s building is then provided leads (the specific type of lead depends on how much bandwidth has been contracted for) into the telecommunication’s network that is set up as a series of ring configurations. This allows the channel extension devices to connect the equipment in the data centers into the high speed network.

Because it is based on ring or mesh topologies, SONET infrastructure is usually *self-healing*. If any point in the SONET ring is disabled, traffic will be re-routed in the opposite direction. As a result, this technology provides very high availability without necessitating that a secondary route be provisioned.

Figure 6-2 shows an example of a SONET ring. Here you see the leads entering the public network through leads with names such as OC1, OC3, and Gigabit Ethernet. These names refer to an amount of bandwidth available in that pipe which is then mirrored in the amount of bandwidth that has been provisioned for that pipe in the network. So, as an example, a business leasing a single OC3 through IBM Global Services would be provided with a link into its facility that is capable of handling 155 Megabits per second (Mbps). This pipe would then lead into the larger public network where that business would receive 155 Mbps of the total shared bandwidth available.

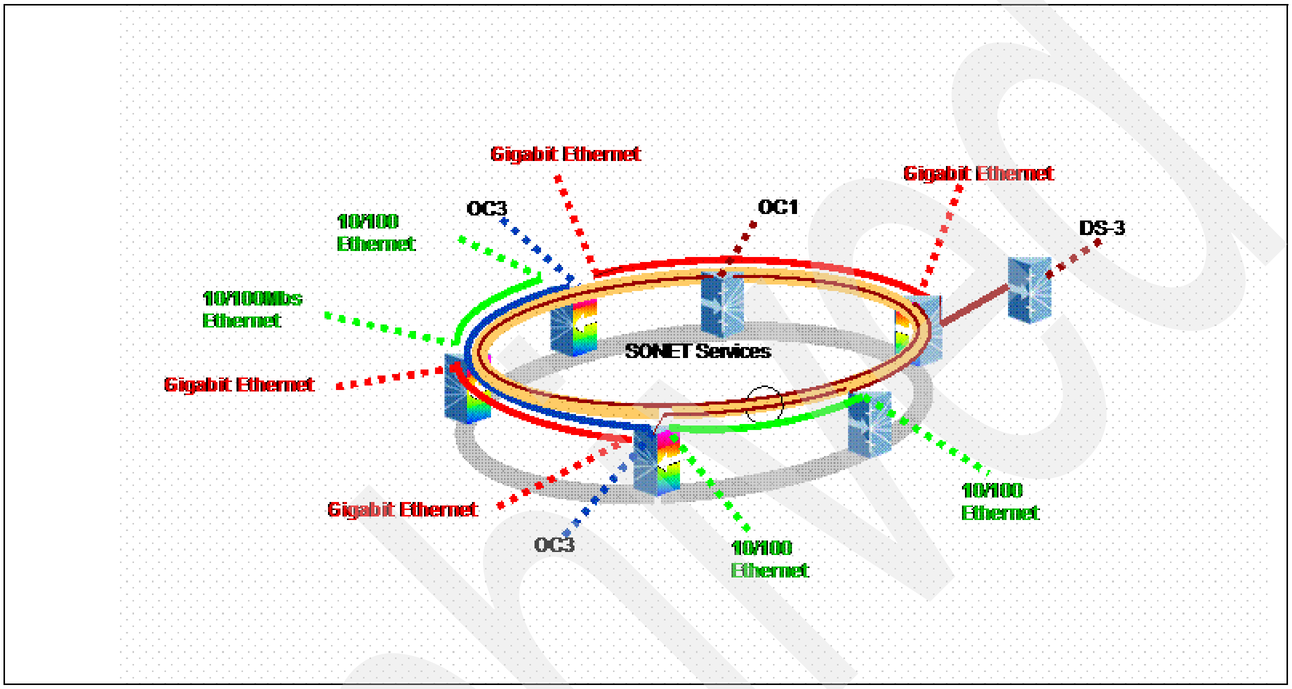

Figure 6-2 Synchronous Optical NETwork

6.2.3 Data transport speed, bandwidth, and latency

The speed of a communication link determines how much data can be transported and how long the transmission will take. The faster the link, the more data can be transferred within a given amount of time. Bandwidth is the throughput of a network, its capacity to move data as measured in millions of bits per second (Mbps) or a billions of bits per second (Gbps). Latency is the time that it takes for data to move across a network from one location to another and is measured in milliseconds. See Figure 6-3 on page 85 for a comparison of the various transport links available.

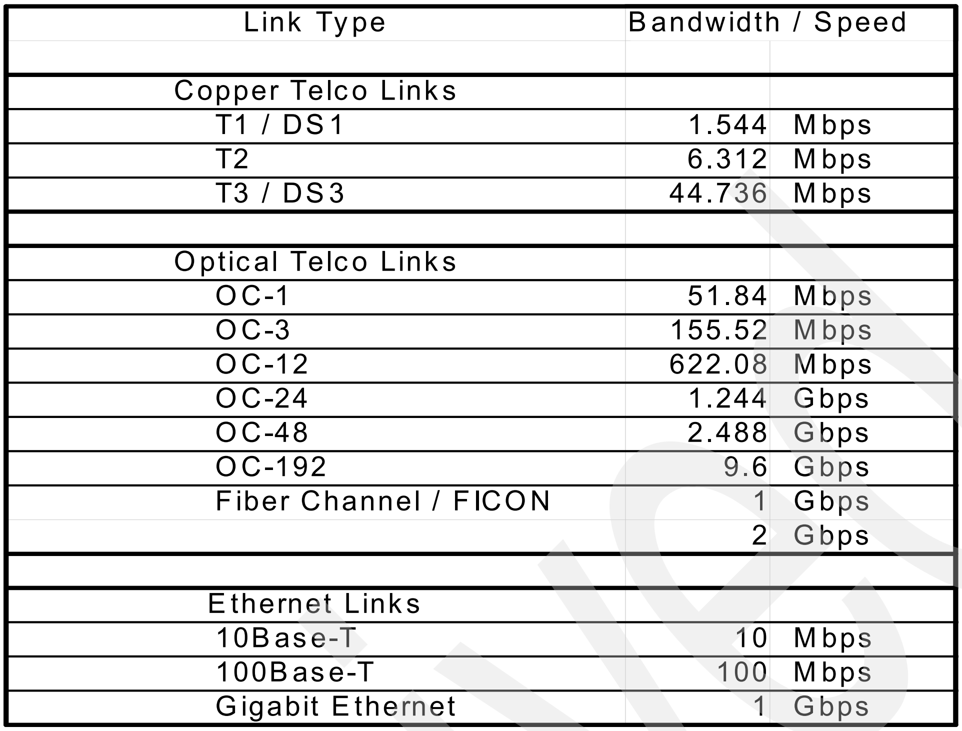

Link Type	Bandwidth / Speed	
Copper Telco Links		
T1 / DS1	1.544	Mbps
T2	6.312	Mbps
T3 / DS3	44.736	Mbps
Optical Telco Links		
OC-1	51.84	Mbps
OC-3	155.52	Mbps
OC-12	622.08	Mbps
OC-24	1.244	Gbps
OC-48	2.488	Gbps
OC-192	9.6	Gbps
Fiber Channel / FICON	1	Gbps
	2	Gbps
Ethernet Links		
10Base-T	10	Mbps
100Base-T	100	Mbps
Gigabit Ethernet	1	Gbps

Figure 6-3 Data transport link bandwidth and speed

It is also useful to compare some of the relative line speeds as shown in Table 6-1.

Table 6-1 Line speed comparison

	Mbps	Approximate MBps	Equivalent T1 lines
T1	1.544	.1544	1
T3	44.746	4.4746	28
OC3	155	15.5	100
OC12	622	62.2	400
OC48	2488	248.8	1600

The bits of data travel at about two-thirds the speed of light in an optical fiber. However, some latency is added when packets are processed by switches and routers and then forwarded to their destination. While the speed of light may seem infinitely fast, over continental and global distances, latency becomes a noticeable factor. There is a direct relationship between distance and latency, propagated by the speed of light. For some synchronous remote copy solutions, even a few milliseconds of additional delay may be unacceptable. Latency is a particularly difficult challenge because, unlike bandwidth, spending more money for higher speeds will not reduce latency.

6.2.4 Technology selection

Figure 6-4 presents a very generalized summary of three network technologies, the scenario for their use, and the network provider. Network technology selection is based on distance, type of traffic, traffic volume, speed, access, cost, and other factors. While the selection might be straight forward in some cases, generally the selection process requires significant networking expertise. This is a strategic infrastructure decision for an enterprise. Selecting an inappropriate technology might prove unworkable upon implementation or could negatively impact an enterprise's ability to expand in the future.

IBM has designed, built, and managed huge pervasive networks not only for internal use, but also for others such as large financial institutions. IBM has relationships and sourcing contracts with every major networking and telecom provider in the industry today. The intellectual capital base and wealth of expertise accumulated from these endeavors and relationships is available from IBM Global Services networking services.

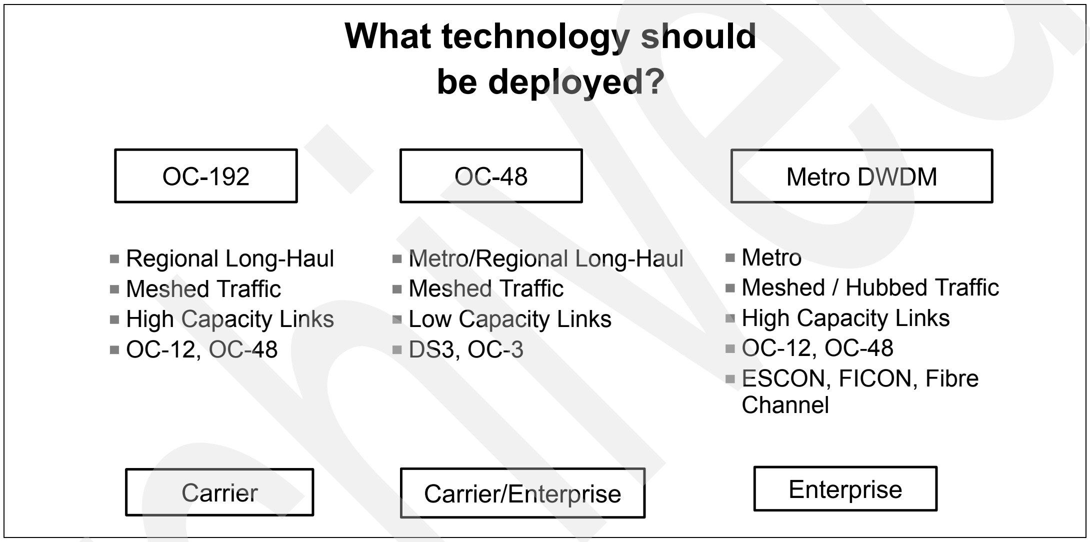

What technology should be deployed?

OC-192

- Regional Long-Haul
- Meshed Traffic
- High Capacity Links
- OC-12, OC-48

Carrier

OC-48

- Metro/Regional Long-Haul
- Meshed Traffic
- Low Capacity Links
- DS3, OC-3

Carrier/Enterprise

Metro DWDM

- Metro
- Meshed / Hubbed Traffic
- High Capacity Links
- OC-12, OC-48
- ESCON, FICON, Fibre Channel

Enterprise

Figure 6-4 Network transport technology selection

> **High availability and distance:** When planning to install any sort of network equipment it is important to bear in mind that most devices will include a high availability function to protect from failures internally or in the fiber path. This will require that a second path be available to be certain that events occurring outside of the data center (such as construction that could accidentally cut through a fiber route) do not become an obstacle to your ability to mirror data.
>
> Bear in mind that just as fiber routes do not necessarily follow a short, straight line, the secondary fiber route will almost definitely be longer than the primary and should be used as the measurement when determining what, if any, application impact results from distance. This is because the distance of the secondary route will represent your *worst case scenario.*
>
> Your fiber or managed network provider will be able to put in writing what distance the data is actually traveling over the primary and secondary routes (or rings).

6.3 Wavelength Division Multiplexing (WDM)

Wavelength Division Multiplexing (WDM) is a method of improving the efficiency of dedicated fiber by condensing, or *multiplexing*, the transmitted channels. To visualize this process, imagine an inverted prism.

White light passing through a prism is split into a wider spectrum of colors, each with its own wavelength. In the WDM world, this is reversed and each device connected to the WDM is given a wavelength of light (known as a lambda), similar to the different colors in the spectrum. The WDM then takes these wavelengths and allows them to pass together across the fiber in the same area of the light spectrum, around 1550 nanometers. Figure 6-5 demonstrates the flow as the channels connect to the WDM and enter the fiber strand.

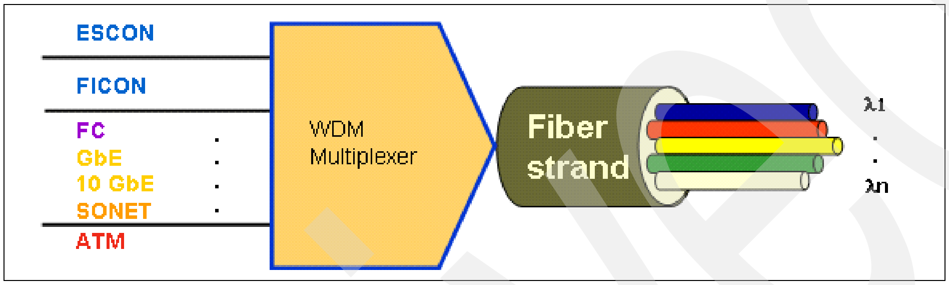

Figure 6-5 WDM conceptual flow

At the remote location the light passes through a second WDM. This device takes in the *white light* transmission of data and divides it back into individual wavelengths connecting to devices as though it were attached by its own fiber strand. Because this allows many devices to connect over a single pair of single-mode fiber strands, this represents a great improvement in efficiency over direct connections through individual fiber strands and can represent a major cost savings given the expense involved in building up a fiber infrastructure.

Additionally, WDM technology represents a long term investment protection on the fiber infrastructure. Improvements in WDM technology often are able to continue to improve the efficiency of fiber infrastructure over time, reducing or removing the need to add additional fiber strands when more channels can be added to the existing WDM infrastructure.

Remember: Not every device is supported in any given configuration. Before committing yourself to a specific WDM or Channel Extension device for a solution (such as GDPS), make sure to discover whether it is supported in your configuration.

6.3.1 Optical amplification and regenerative repeaters

Under normal circumstances, current WDM technology is able to transport data to a range of at least 50 km. Depending on the specifics of the device and certain optional equipment (such as more powerful lasers), some devices are capable of connecting at a distance of 80 to 90 kilometers. Because there is no encapsulation or decapsulation involved in the process, this is at full, native speeds. As a result, data is kept current in the remote disk with minimal latency in synchronous technologies and asynchronous technologies minimize the risk of data recreation.

In some cases, however, the native range of a WDM falls under the range desired or required for a given disaster recovery solution, but dedicated fiber is still seen as necessary or desirable. In these cases, optical amplifiers should be considered.

Optical amplifiers are devices set between WDM units that take in the signal from the sending WDM and strengthen it. As a result, the effects of distance are minimized and the supported range between a pair of WDM devices increases substantially to a maximum supported distance of 300 km. Longer distances are possible via special request.

Regenerative repeaters can be used to further extend the distance between sites. A regenerative repeater converts incoming optical signals to an electrical signal, which is cleaned up to eliminate noise. The electrical signal is then converted back to an optical signal and retransmitted to the target site.

It is important to note that not all protocols are created equal where distance is concerned. Some older protocols do not fare as well over long distances. FICON® and other Fibre Channel based protocols, however, perform very solidly even at a range of 300 km. As such, for longer ranged solutions over dedicated fiber, it will generally be safer and more efficient to make use of the newer technologies.

Vendors of WDM devices, optical amplifiers, and regenerative repeaters should be contacted regarding attachment and distance capabilities and line quality requirements. The vendors should also provide hardware and software prerequisites for using their products.

6.3.2 CWDM versus DWDM

There are two forms of WDM technology available. While the basics are the same for both forms, it is important to understand each has potential options that affect the cost and scalability of solutions.

Coarse Wavelength Division Multiplexing (CWDM)

CWDM spreads out the lightwaves instead of trying to keep them closer together. The result is a smaller increase in the number of lambdas available through the fiber strands. Although this is, in and of itself, not an advantage of CWDM technology when compared to DWDM, it tends to be significantly less expensive. As a result, it is an appropriate technology for businesses who use dedicated fiber but only have a limited number of cross site links. CWDM is normally capable of sending eight channels of data across one pair of fiber strands. This, however, can be increased through subrate multiplexing.

Wavelength Division Multiplexing (DWDM)

The most common form of multiplexing, DWDM tries to keep the lambdas as close together as possible. As a result, the number of channels that can pass through a pair of fiber is greatly increased. Current standards allow for up to 32 lambdas to pass through one pair of fiber, and this can be multiplied through technologies such as subrate multiplexing.

6.3.3 Subrate multiplexing

Current standards in WDM technology allow devices to split the light within one pair of fiber strands into eight or 32 lambdas. Subrate Multiplexing (also known as Time Division Multiplexing or TDM) is the concept of making more efficient use of each of those lambdas by sending multiple channels of information over a single optical interface. Multiple devices are able to connect to a single optical interface card and their data is transmitted with slight time variations. As a result, these cards act as a multiplier on your WDM infrastructure allowing for the possibility of transmitting two or more times as many channels within a single frame.

Because of the time variance, however, not every device will be supported equally with TDM cards. Some particularly sensitive devices, such as sysplex timers, may not be supported due to the minor variances in timing. Other channel types are far more tolerant.

6.4 Channel extension

Channel extension is a distance transport technology in which the device connects to SONET optical channels. These devices then take the data from all connecting devices and send it across the SONET route using however much bandwidth has been provisioned by the telecommunications provider. Unlike DWDM, this method does not give each device its own wavelength, but it is suitable for much longer distances and, depending on the particular business and configuration, may make better financial sense than using dedicated fiber. Figure 6-6 shows how channel extension can be used to transfer large amounts of data between data centers.

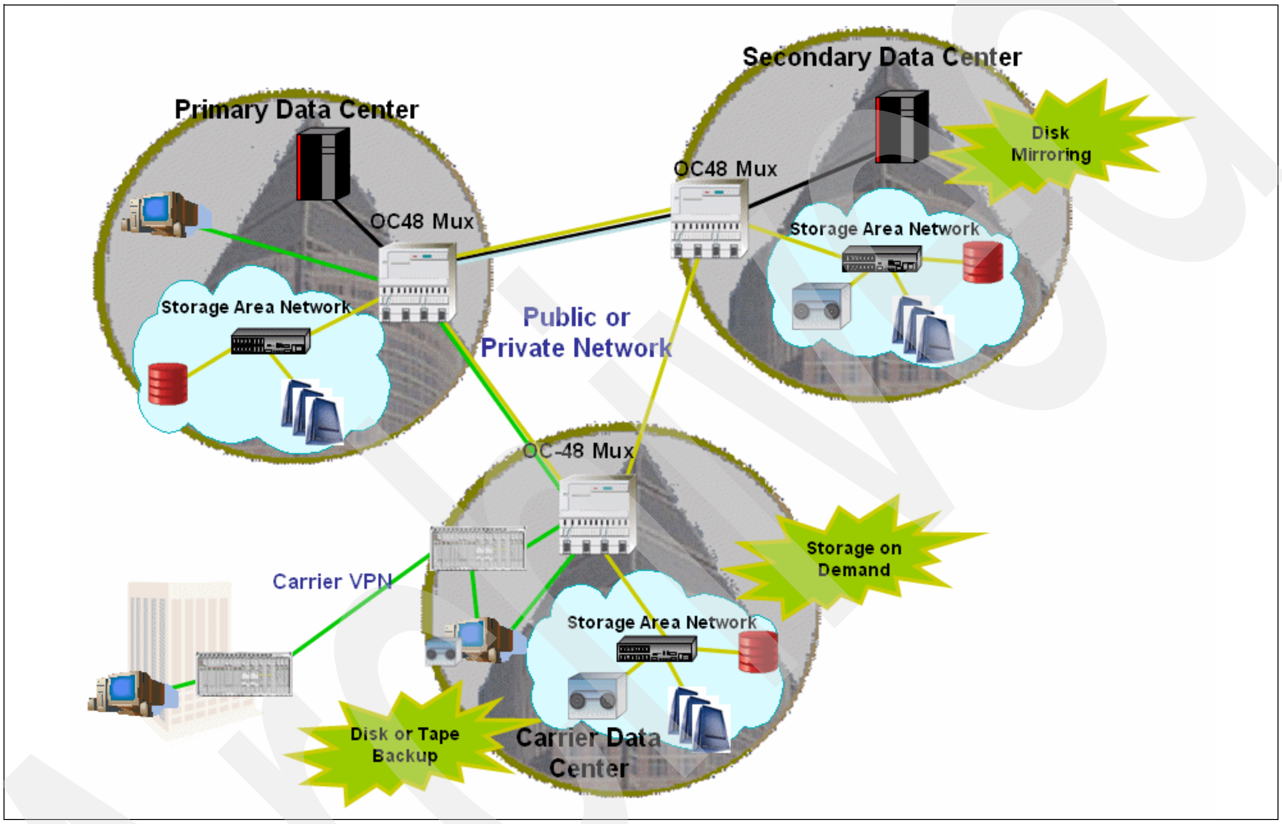

Figure 6-6 Example of channel extension implementation utilizing SONET

6.4.1 Methods of channel extension

Because SONET is used over such long distances, special methods must be used to ensure that devices do not time out while data is still in-flight. This is typically handled in one of two ways.

Traditionally, data is encapsulated in IP packets and placed into Asynchronous Transfer Mode (ATM) cells before being transmitted into the network. This has been and continues to be a viable method of transmitting data over longer distances. Under this form of transmission, the channel extension equipment spoofs or tricks the receiving devices into thinking that data has been received. Meanwhile, the data being transmitted is broken and encapsulated into IP packets that are then placed into ATM cells. The cells are received in the remote location and are decapsulated into their original format, so that they can be properly received by the remote devices.

A second, newer form of channel extension is to augment the buffer credits (the number of unacknowledged frames that can accumulate before data transmission stops) on both sides of the telecommunications line. By doing so, the line is constantly filled and remains open to the receiving devices while data information is transmitted without encapsulation.

FICON and Fibre Channel over distance

While using ESCON® over SONET has always been acceptable, even at extremely long distances, due to the way that channel extension breaks the protocol with encapsulation, the protocol exchanges still do slow down the connections somewhat. Additionally, ESCON faces some limitations when transmitting over longer distances. Because FICON and Fibre Channel Protocol (FCP) perform so well over longer distances, their connectivity has become important to channel extension device vendors.

Technologies have been developed recently that make FICON and FCP over SONET a possibility and vendors have started to offer it on channel extension devices. As such, when planning a Disaster Recovery solution, due consideration should be given to the performance and cost efficiencies of the newer FICON or FCP technologies rather than staying with the solid, but less efficient, ESCON.

Fibre Channel over IP

Another option for FCP over long distances is known as Fibre Channel over IP (FCIP). This technology allows a user to encapsulate data in an IP packet but rather than sending it through a SONET ring, the data travels through the business's existing IP network. Many businesses have already spent a great deal of time and money building up their IP networks and in doing so have developed a strong skill base in IP networking. In these cases, FCIP may prove to be more cost effective than other channel extension technologies.

There is one word of caution when investigating Fibre Channel over IP, though: large disk mirroring implementations may take up a sizable amount of the available bandwidth and their use may result in the need to further expand an IP network. If the intent is to share the same LAN for both data replication and standard IP network uses, stress testing must be performed to ensure that any application impact in either use would be acceptable. However, the potential cost and benefit associated with such a decision is something that should be reviewed on a case by case basis.

6.5 Testing

Testing a proposed Disaster Recovery solution is important before, as well as after, implementation. This is because introducing distance into your applications could have unforeseen effects and these must be understood to properly prepare. While it is unlikely that you would be able to appropriate temporary fiber in the ground, there are methods that will allow you to test within a single room if desired.

If you plan to use dedicated fiber, you should investigate whether your fiber provider can arrange for you to have access to *fiber suitcases*. These are spools of fiber optic cable with the cladding removed. They are placed in a suitcase-like container and may be connected together to simulate the actual distance between locations as well as test failover for high availability functions. Additionally, attenuators can be connected to simulate the level of *line noise* that is anticipated on the actual fiber run. The attenuators reduce the signal which in turn changes the signal-to-noise ratio.

Fiber spools may not always be a workable solution, especially in solutions that transport over continental distances. In these cases, an alternative may be introducing devices that introduce additional latency into the environment. By doing so, the data takes longer to be received, just as it would be in its real implementation due to the speed of light over distance.

7

High-availability clusters and database applications

This chapter addresses the concepts used for creating high availability clusters and their use with regard to databases and other applications. After reading this chapter, you should have a basic understanding of high availability functions available through clustering and database applications.

It is not our intention to cover all possible database and application backup and recovery technologies. We want to point out that besides technical and logical data availability, the most important consideration is business continuance. IBM Global Services offers dedicated services for Business Continuity that are not the subject of this IBM Redbook, but we refer to them when it is appropriate.

This chapter first explains the methods on how to make systems highly available. Then we will introduce you to the technologies and functions suitable for Disaster Recovery and fault-tolerant solutions for databases.

7.1 High availability

When we refer to high availability in an IT infrastructure, we need to look in detail at how high availability can be achieved for each of the represented solution building blocks. When considering Business Continuity, we must have a basic understanding about the principal differences between high availability, disaster tolerance, and Business Continuity.

High availability: In general, high availability for IT systems is provided by the following characteristics:

- No single point-of-failure (SPOF) within the system or the systems' components by using redundant components.
- Automatic failover/failback.
- Masking and elimination of downtime by using redundant components.
- System availability might be degraded or unavailable during planned and unplanned maintenance.
- Usually local to one site.
- Fault resilience.
- Hardware based.

Disaster tolerant: In general, disaster tolerance for IT systems is defined by the following characteristics:

- High availability of all system components.
- Possibility to restore and enable the IT infrastructure and the current and valid data within a given time objective (RTO).
- Planned and unplanned maintenance windows shall not impact the system availability.
- IT infrastructure is usually dispersed over two or more sites.

Business Continuity: Business Continuity is defined by the following characteristics:

- A highly automated disaster tolerant IT infrastructure.
- Organizational and business related infrastructure is available.
- The company's business continues with little or no degradation.

Solutions that utilize technologies that provide Continuous Availability will provide the necessary IT infrastructure as a building block for business continuance depending on the availability requirements, as shown in Figure 7-1 on page 96.

7.2 Clustering technologies

In a highly available and redundant IT infrastructure, availability can be achieved with redundancy of systems by introducing clustering technologies. A clustered system is defined as a group of individual server systems that share some resources and act as a single system.

Server clusters are designed to keep applications available, rather than keeping data or the business process available. Thus, any kind of hardware based clustering is just a way to protect the physical infrastructure, but does not provide protection for logical or other data related disasters (data corruption, data loss, viruses, human errors, intruders, or hackers). To protect the data and the business process against those kind of threats, organizations need solid data protection for their companies' data!

Therefore, besides technical data protection and physical access barriers to the core systems, proven recovery plans and technologies must be in place. Cluster technology cannot protect the data against failures caused by viruses, software corruption, or human error. With clustered systems, we differentiate between four types of clusters:

- Shared nothing cluster
- Common shared cluster
- Application clusters
- Geographical Dispersed Clusters

In the following sections of this chapter, we describe the basic cluster technologies and how they can be deployed as highly available building blocks in disaster tolerant solutions in conjunction with databases.

7.2.1 Shared nothing clusters

A shared nothing cluster is defined as a system of independent servers (two or more) that have access to a common storage system but operate on their own resources. A resource in a shared nothing cluster can be, for example, a disk, an application, or a file share. The clients, however, see the cluster as one entity and not as *n* individual servers. Any connection from a user or application will only be executed on one individual node within the cluster.

Figure 7-1 shows a shared nothing cluster.

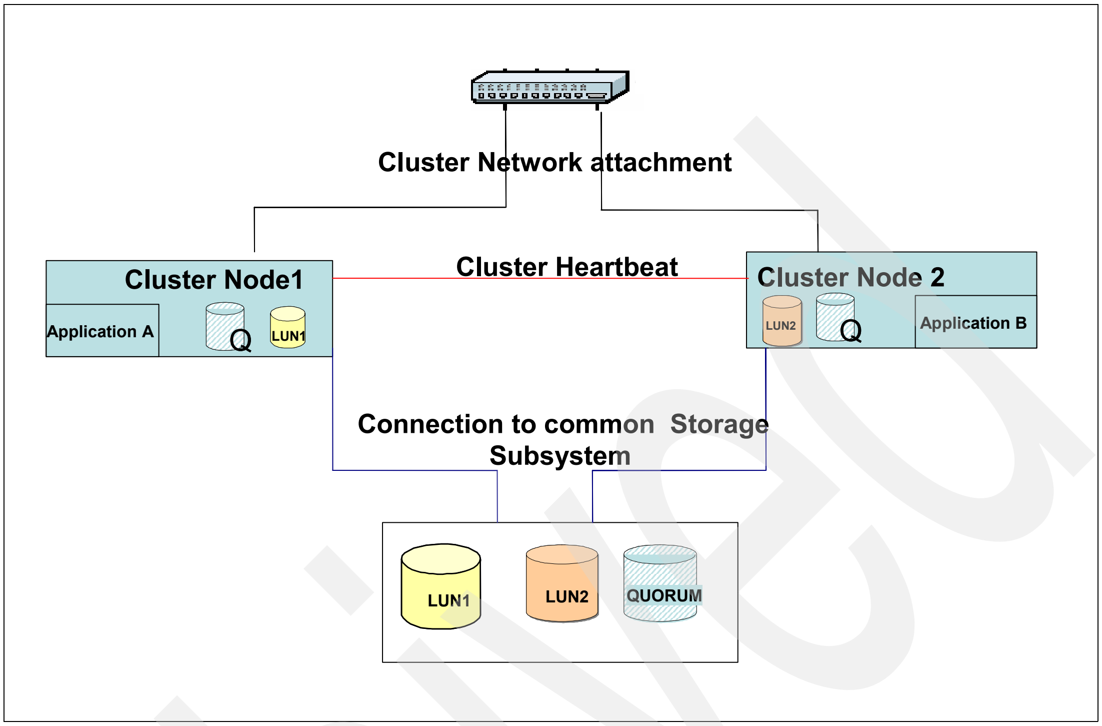

Figure 7-1 High availability shared nothing cluster

The nodes can be configured in two ways, either active/active or active/passive. With an active/active cluster, both nodes operate on their own data and applications and in the event of a problem on either cluster node, the resources will fail over to the surviving node (see Figure 7-2 on page 97). The cluster heartbeat is a dedicated network that checks the availability of the cluster nodes and resources. High availability, shared nothing clusters are usually operating system based clusters. With Windows.Net 2003 Enterprise Edition and Windows.Net 2003 Datacenter Edition, a cluster can contain up to eight nodes. IBM AIX HACMP is another example of a shared nothing cluster.

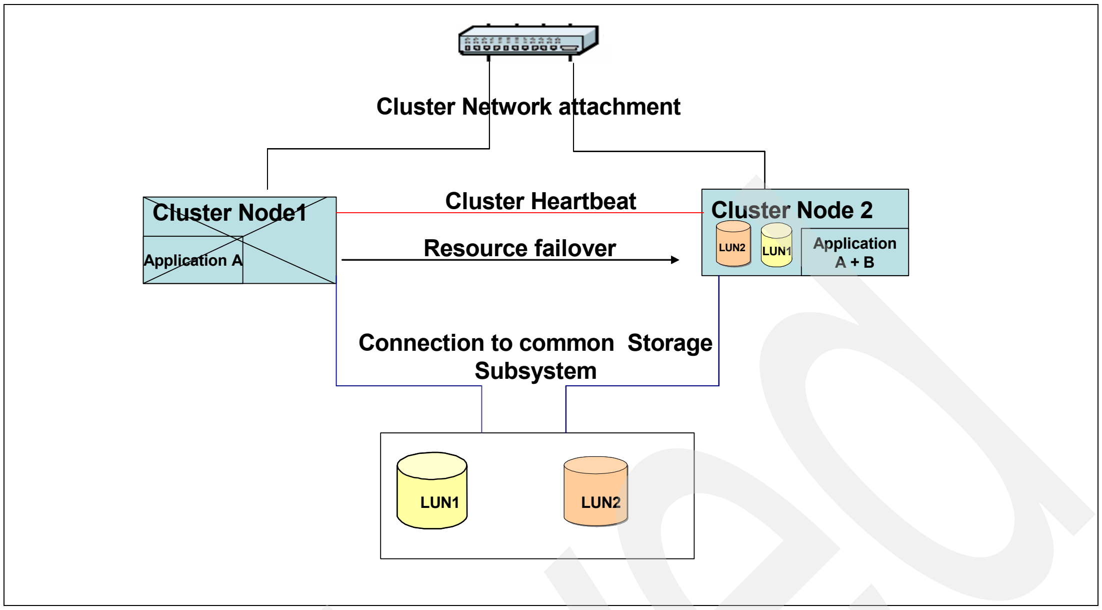

Figure 7-2 Failover within a shared nothing cluster

7.2.2 Common shared cluster

Physically, a common shared cluster may consist of the same or similar building blocks as the shared nothing cluster discussed before. The major difference between the two cluster types is that within a common shared cluster, all nodes have access to the same data and the same resources at a time. However, as a process modifies data on the application, the cluster must provide a specific functionality that enables data locking. DB2 for z/OS Datasharing would be a prime example of a common shared cluster.

Figure 7-3 shows a common shared cluster.

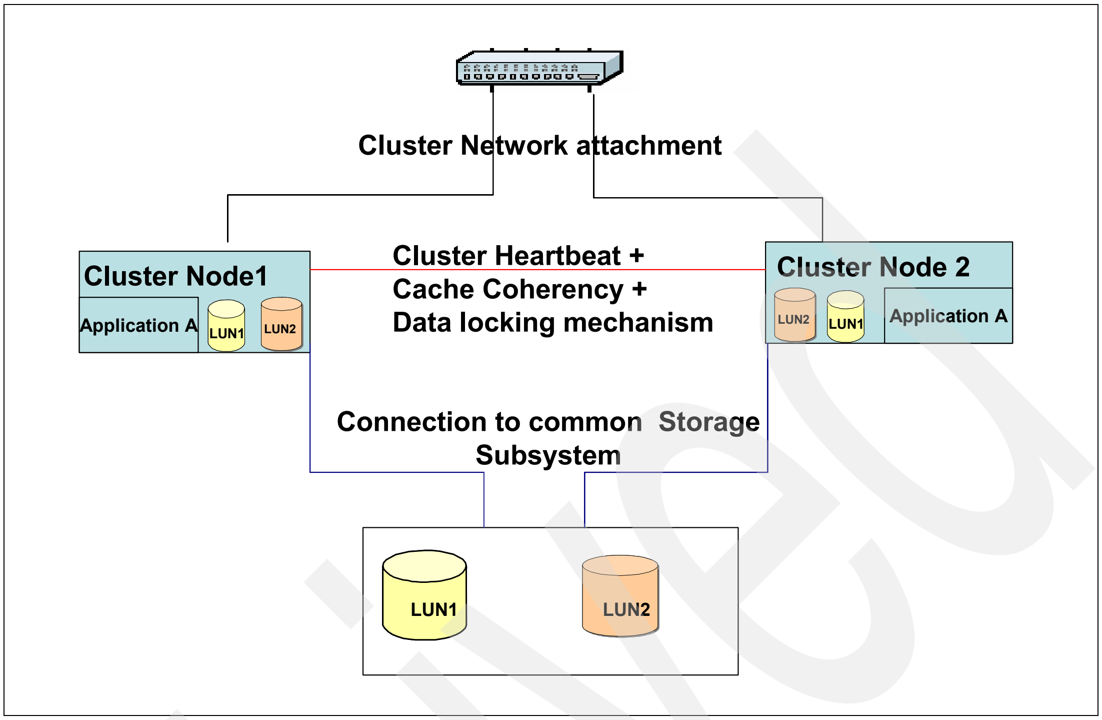

Figure 7-3 Highly available and scalable common shared cluster

7.2.3 Shared nothing application cluster

In a shared nothing application cluster, the users are linked to a farm of independent cluster nodes that act as one single system. Each node has its own data and the load between the nodes is being distributed and coordinated by the applications load balancing facility. In case of a node crash, the users can reconnect to the server farm and continue on the remaining nodes within the cluster. However, transactions and data until the crash of the individual node might be lost depending on the application's behavior. An example for an application server farm would be CITRIX or IBM WebSphere.

Figure 7-4 on page 99 shows a shared nothing application cluster.

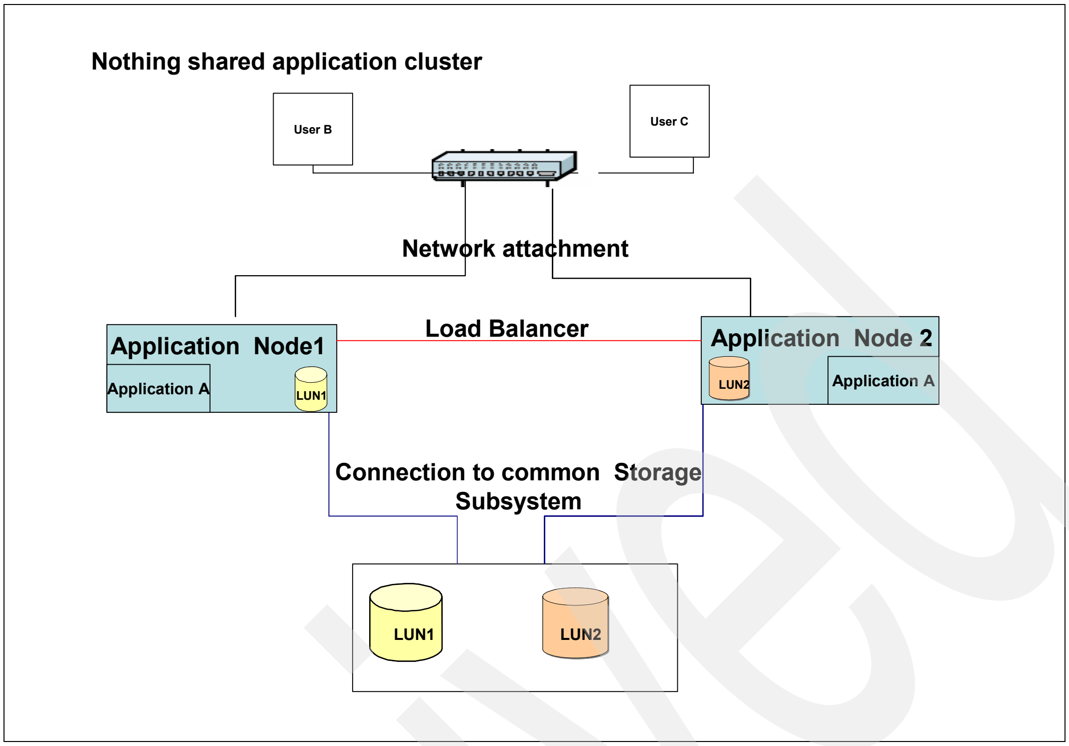

Figure 7-4 Nothing shared application cluster

7.2.4 Geographically dispersed clusters

With the introduction of disaster protection, we have to keep in mind that one type of disaster might be a complete site loss due to fire, water, earthquake, or other severe type of destruction. In this scenario, a local high availability solution of any kind will not provide the availability and continuance of the data and the business process. Therefore, alternatives have to be discussed on how to protect the physical infrastructure, the data, and the applications available after the event of a site failure. Geographically dispersed clusters provide the physical possibilities to extend the cluster nodes (the servers) and the storage subsystems across distances.They do not protect the data from logical or human errors.

As indicated in Figure 7-5, there are two general possibilities within Geographical Dispersed Clusters for synchronization and data mirroring of the storage systems between the sites:

- Host based mirroring: Such as with the Logical Volume Manager
- Storage based mirroring:
  - Metro Mirror for DS6000, DS8000, ESS, and SVC
  - Enhanced Remote Mirroring (ERM) with DS4000
  - Sync-Mirror for N series

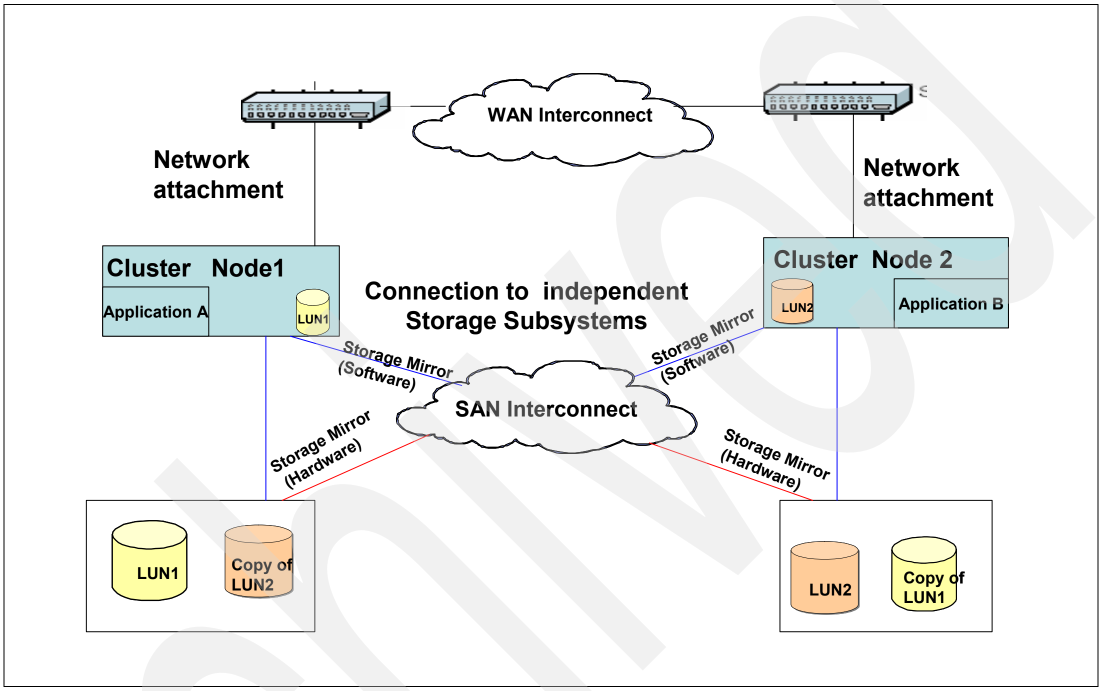

Figure 7-5 Geographically Dispersed high availability cluster

Either method has its advantages and disadvantages.

Examples of Geographically Dispersed High Availability Clusters include GDPS/PPRC and “Copy Services for System i” System i Clusters.

7.2.5 Backup and recovery considerations for databases

Any type of hardware based cluster solution in combination with databases on its own provides just *one building block* for the overall Business Continuity solution. We must also have a basic understanding of how databases work. This knowledge is important to select the appropriate method and technology for storage or server based copy functions in the context of the database protection requirements.

Thus, backup and recovery plans together with the appropriate technologies are a key component within the disaster protection and data recovery methodology. Two general methods of backing up databases are to be considered: *online* and *offline backup*.

- Offline backup usually means to shut down the database and take a full backup of the data at the file or raw volume level.
- Online backup allows the database to continue running during the backup. The database manager software provides an orderly way of backing up the database files, control files, and archived transaction logs.

Before we look in detail at online backup, we need to describe how the database components interact with each other. A transaction oriented online database of any kind consists of three major disk storage components:

- Data: Data is organized in flat files, tables, or tablespaces depending on the database.
- Log files: Keep the records of all the transactions that have been written to the database. As log files fill, they will switch the processed transactions to archive logs, as shown in Figure 7-7 on page 102.
- Archived log files: Contain already-processed transactions written from the log files.

Figure 7-6 shows the process flow of a database write operation with transactional logging enabled. This approach is used to guarantee that the data in the database will be write consistent in the event of a server crash.

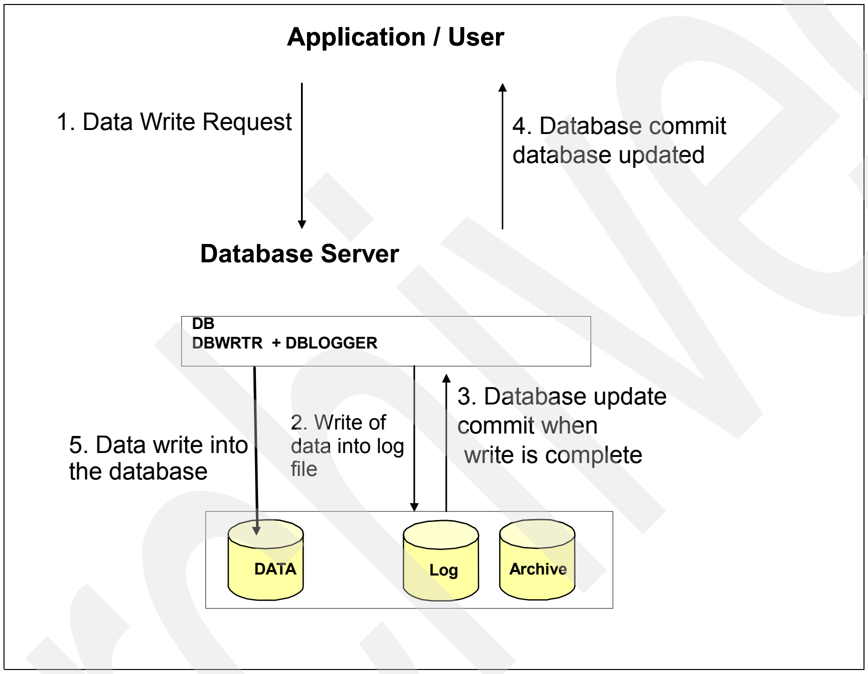

Figure 7-6 Simplified transaction based database write

Once a log file has filled up with transactions, a switch to a new log file will occur. In the mean time, the transactions from the old log file are copied to the Archive Logs.

Figure 7-7 shows a simplistic database log file switch.

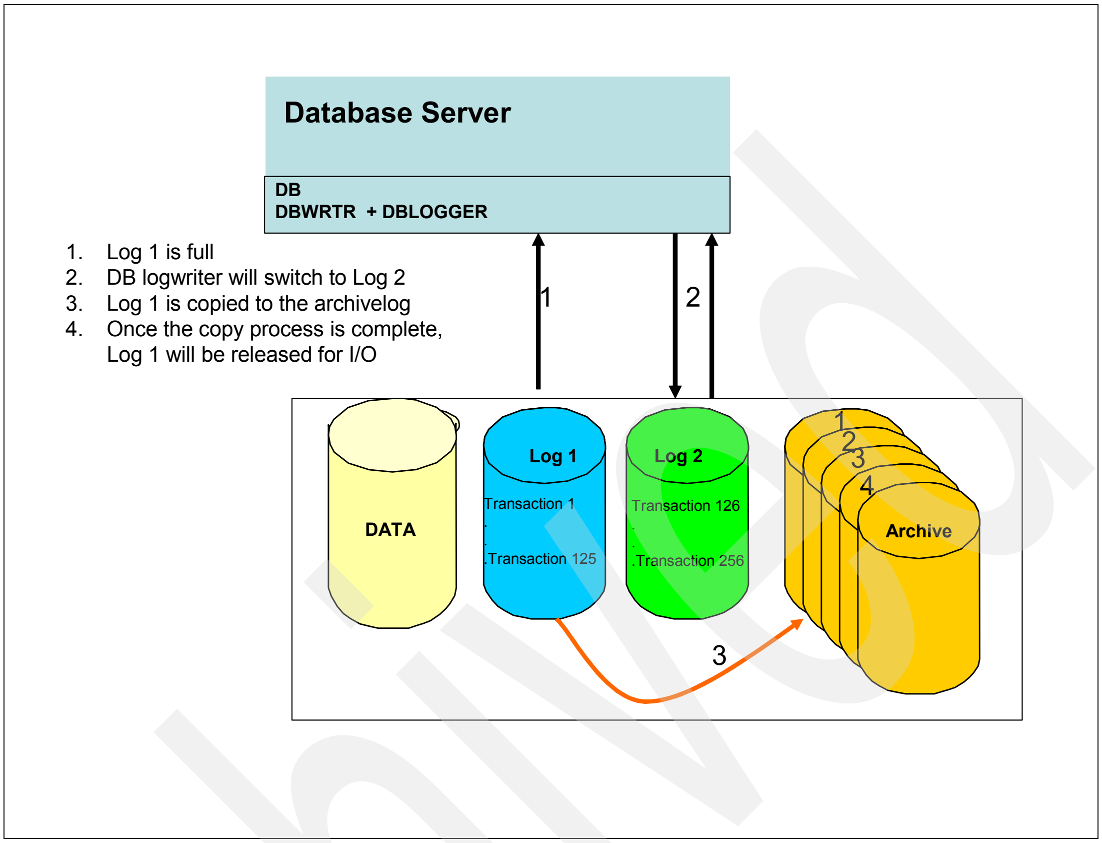

1. Log 1 is full
2. DB logwriter will switch to Log 2
3. Log 1 is copied to the archivelog
4. Once the copy process is complete, Log 1 will be released for I/O

Figure 7-7 Simplistic database log file switch

For offline backups, people usually focus on how long the backup takes. However, the more important question is “How long can I be without access to my data, during a physical data restore from tape?”

Besides the physical speed of the tape, or recovery methods from disk, the time for database recovery depends on the time to apply the Archive Transaction Logs and other housekeeping functions to the restored database. You can calculate the database recovery Recovery Time window as follows:.

- Recovery Time (Rt): Time window from physical restore to data availability
- SetupTime (St): Administrative time to prepare the restore
- Physical Recovery Time (PRt): Time to restore data from tape or other backup storage
- Number of Transactions/Second (nTS): Number of performed transactions/sec on DB server
- Number of Transactions (nT): Means the number of transactions recovered contained in the Archive Files from backup media which have to be applied against the Database

Thus the formula for the recovery time (Rt) is:

Rt = St + PRt + nT / nTs

Important: To recover a database to a certain point-in-time, the log files + the database + the archive files must be in sync! Without the archive and log files being available, and in sync with the database data, a recovery might be impossible.

Figure 7-8 shows how a simplified database point-in-time recovery works. We consider a disaster on Friday, where we have to recover from tape. If the backup policy was a full database backup including the database, the log files, and the archives every Sunday, then we need to keep backups of the log files and archives for the week to recover to the point-in-time. This implies that after the physical restore of the database and the corresponding log files and archives, all transactions that are stored in the archives have to be redone on the database! In our scenario, this means we need to redo the Archive Logs from Monday through Friday to recover the database to the Friday point-in-time level.

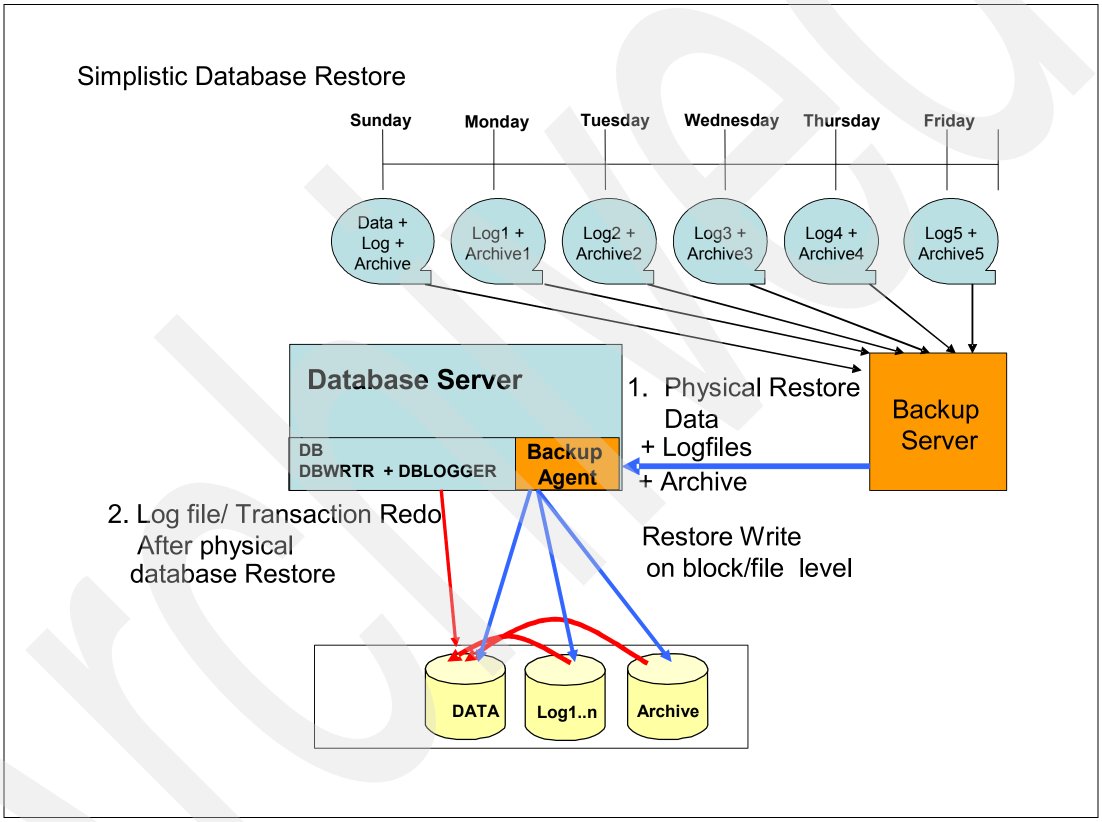

Figure 7-8 Simplistic Database Restore - Point in Time

Considerations

When the database is online, the log files and the database tables are also open for I/O. Due to the fact that these components are open for data I/O, one must consider the behavior of a backup system and how it interacts with the file and data access.

Online backup

There are two ways of online database backup: logical online backup or physical online backup.

In *Logical online backup*, the database and the backup software communicate via backup agents and APIs. The backup software is closely integrated with the database.

As indicated in Figure 7-9, the Backup Agent will trigger the backup. While the online backup is running, all write activity against the database will be rerouted to the dedicated backup log. This will guarantee that the database and the log files are always in sync for a point-in-time recovery. The disadvantage of such a solution might be the impact on the write access during the backup and after the backup has completed.

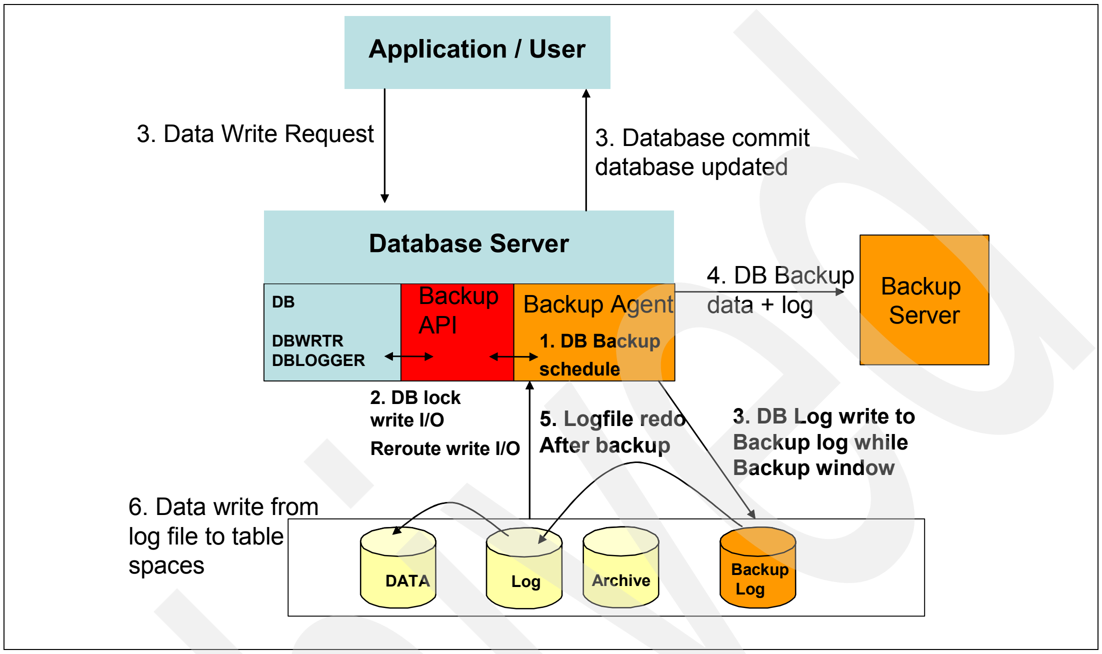

Figure 7-9 Logical Online Backup Method (LOB)

If heavy database write activity is expected, then the performance of the database will decrease during the backup window. The overall time for the full database backup depends on the size of the database and the backup technology in place.

Physical online backups are considered as a physical image copy of the database and its related log files, but executed on the storage (*physical*) level. This implies proper I/O quiescence for the database write access from within the database and the operating system layer during the copy process.

> **Important:** The database must flush its internal buffers, such as DB2 bufferpools or Oracle SGA, to the underlying file system or raw device. If a file system is involved and has lazy write enabled, where writes do not go synchronously to disk but some time later, it will be necessary to flush the file system buffers with specific commands for file system to disk synchronization.

The copy process is independently executed on the storage site to avoid inconsistencies on the database level.

As shown in Figure 7-10 on page 105, the copy process itself is executed on the storage layer versus the logical online backup where the backup is executed on the *server layer*. Once the *FlashCopy process* has been initiated, which will take just a couple of seconds, the actual copy process of the data to the FlashCopy volume will be executed in the background. This allows the database to resume write access. Figure 7-11 on page 105 shows FlashCopy copy on write.

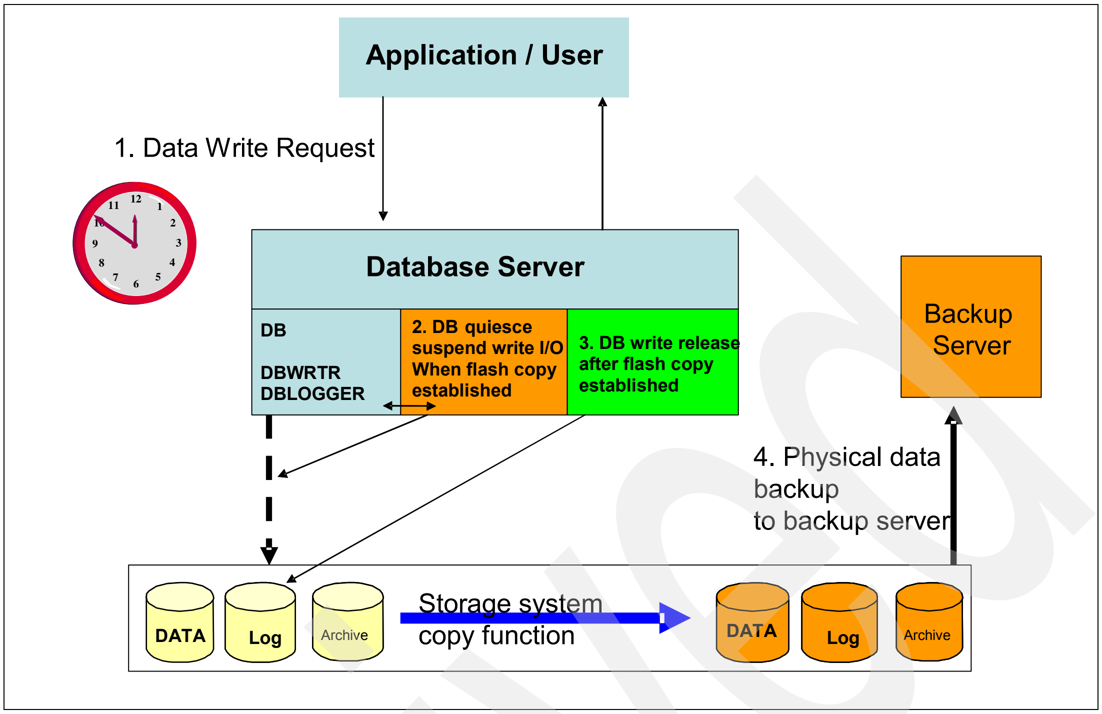
Application / User
1. Data Write Request
Database Server
DB
DBWRTR
DBLOGGER
2. DB quiesce suspend write I/O When flash copy established
3. DB write release after flash copy established
Backup Server
4. Physical data backup to backup server
DATA
Log
Archive
Storage system copy function
DATA
Log
Archive

Figure 7-10 Physical database copy using storage copy function

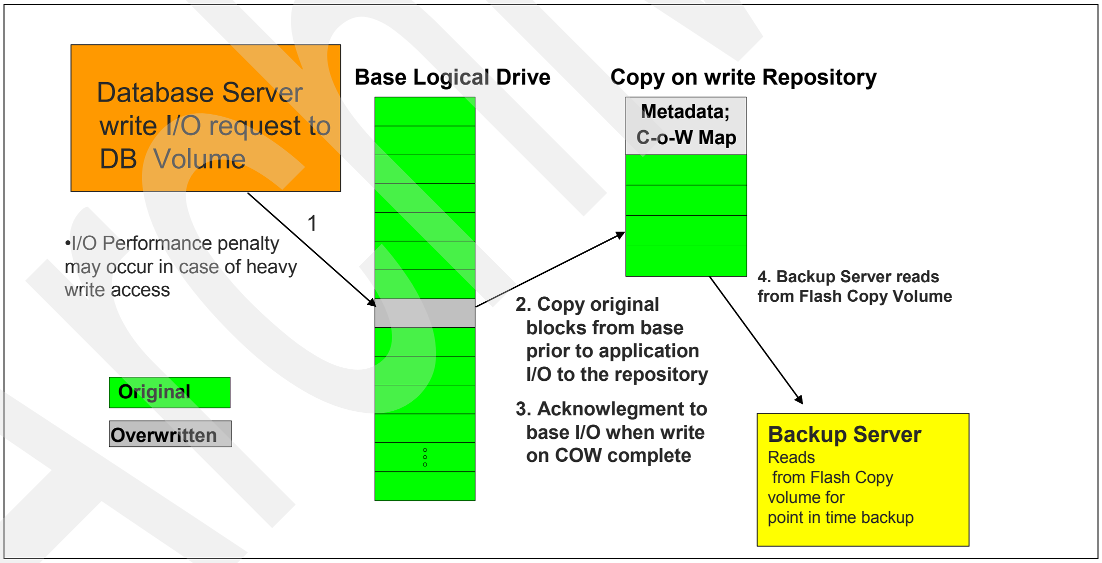
Database Server write I/O request to DB Volume
Base Logical Drive
Copy on write Repository
Metadata; C-o-W Map
1
•I/O Performance penalty may occur in case of heavy write access
2. Copy original blocks from base prior to application I/O to the repository
3. Acknowlegment to base I/O when write on COW complete
4. Backup Server reads from Flash Copy Volume
Original
Overwritten
Backup Server
Reads from Flash Copy volume for point in time backup

Figure 7-11 Simplistic FlashCopy operation

There are some unique considerations for Windows and AIX with FlashCopy, in particular, target LUNs for FlashCopy should *not* be mounted by any server. After the Flash, they can be mounted and dumped. See *IBM System Storage DS8000 Series: Copy Services in Open Environments*, SG24-6788 for more information.

Conclusion

Any kind of online database backup will avoid interrupting production without shutting the database down. It will also avoid long restart activities and the restoration of the database buffers. Thus it is important to select the appropriate technology depending on the business needs and in conjunction with the IBM Business Continuity Solution Selection Methodology.

The following sections will describe examples of how database security can be achieved by using certain technologies and methods. We will also describe the advantages and disadvantages of each solution.

However, these scenarios are just examples; this IBM Redbook cannot cover all possible combinations of databases, operating systems, and hardware platforms. Therefore, we always recommend using the basic IBM Business Continuity Solution Selection Methodology and compare any projected solution with the basic solutions for high availability and database backup and recovery, as described earlier in this chapter.

8

Business Continuity for small and medium sized business

This chapter provides a Business Continuity solution overview specific to the small and medium sized business (SMB) environment.

SMB enterprises, which play a vital role in our worldwide economy, are small or medium *only* in relation to the size and scale of large multi-national corporations. SMBs are often quite large within their regional or local geography, and certainly SMB enterprises are not small at all in terms of dynamic ideas, innovation, agility, and growth.

In many ways, SMB companies have IT needs and concerns similar to large enterprises. Yet in other ways, SMB companies have key differences.

This chapter provides a discussion of these differences. If your business exhibits characteristics of this vital type of enterprise, this discussion of key Business Continuity differences for SMB should be useful.

8.1 Small and medium sized business overview

Small and medium businesses (SMB) play a very important role in the worldwide economy.

As an example, according to US Government Small Business Administration data, SMB companies range from home based entrepreneurs (small business) to those with a thousand employees (medium business). Their annual revenues can run from thousands to billions of dollars. The same data says that in the US, SMB represents nearly 99.7% of all employers, responsible for nearly three quarters of the net new jobs created to the US economy in the last three years. SMB accounts for over half of the United States private work force and drives over 40% of private sales.

While SMB statistics will vary according to geography and economic conditions, clearly SMB companies have specific requirements for Business Continuity. SMB companies, depending on the nature of their business, have different business continuity requirements. As computing technologies are becoming affordable, SMB businesses can take advantage of emerging business continuity technologies to help drive growth and profits for their business.

8.1.1 SMB company profiles and business continuity needs

Recent natural disasters and terrorist threats have put Business Continuity (BC) as a top priority for enterprises worldwide. This is a lot of urgency within SMB companies to gear up their BC capabilities, as many if not most are behind in this area. IBM customer surveys indicate that Business Continuity is the number one IT issue since 2003 for SMB companies. These are the four key drivers:

- As SMB companies increasingly rely on their IT systems to support their business, any system downtime and data loss has severe negative impacts on revenues, profits, and client satisfaction. Extended outages or the inability to recover critical data can cause permanent damage to companies.
- Recent government compliance regulations, such as the United States' Sarbanes-Oxley Act and HIPAA, also push data backup, restore, retention, security, auditability, and Disaster Recovery requirements to top priorities for public SMB companies.
- Customers and IBM Business Partners increasingly require reliable and highly available systems as prerequisites for doing business.
- The ability to minimize risks is important for some maturing SMB companies. Planning for Business Continuity is similar to buying insurance; recent events make this type of insurance a must rather than a luxury as in the past.

Most SMB IT management finds Business Continuity complex and resource intensive, so BC planning usually is an afterthought. With increasing pressure from the lines of business and BC technologies and solutions becoming more affordable and simple, SMB IT management is moving BC projects forward. In some cases, BC can be leveraged to drive additional revenues and profits.

With limited financial and technical resources, IT staff face the following challenges:

- Ever diminishing data backup window time: With more servers and storage devices coming on line, dramatic growth of data and the push for 24x7x365 system uptime, planned outage windows are smaller by the day, affecting the ability to back up systems, applications, and data adequately. Some data may be backed up at all, exposing the business to liabilities and losses.
- Inefficient tools: Since most off the shelf applications bought by SMB use their own backup and restore tools to support their data only, it is common for SMB IT staff to run numerous

backup jobs daily, straining the system and staff resources, eating up precious backup window time; the trend is to have more applications so the situation will only get worse.

- Limited staffing and time: Backup jobs usually are run after work hours and staff has to be around late to support these jobs, in addition to their day duties, resulting in low staff morale, jobs run poorly, or not consistently
- Lack of experiences and skills: BC is still fairly new to SMB and experiences and skills in this area are usually not top priorities with the IT staff; a good example is systems management discipline, including change and problem management, which affect system availability.
- Limited budgets and resources: SMB constantly reprioritize their projects, evaluate trade-offs, and balance resources to achieve their business goals. BC usually is not a top priority funded item until a systems outage actually impacts the business. The actions are usually reactive and can be costly in the long run, such as a total revamp of systems and hiring of outside consultants. Proper planning and funding is essential to successful BC implementation.

In this chapter, we hope to answer the following questions related to the “successful” planning, evaluation, and implementation of BC solutions for SMB companies:

- What are the key BC components for SMB, and how do they affect my business?
- What are the steps in planning for a successful SMB BC project?
- How much SMB BC can I afford for my company?
- Which SMB BC solutions are suitable for me?

8.1.2 SMB company IT needs as compared to large enterprises

SMB companies have IT needs and concerns similar to large enterprises. They need Enterprise Resource Planning (ERP), Supply Chain Management (SCM), and back-office systems (such as e-mail, accounting, and so on). The key differences are scale and costs.

SMB growth rate tends to be steep. The capacity to start from very small and then scale big is a key requirement, without massive changes to existing systems.

SMB companies tend to be more cost sensitive. Maximizing value is a common mantra among SMB. It extends from the purchase and upkeep of necessary computing assets to engaging IT staff who perform a wide range of tasks including operations, support, and maintenance. At the same time, most SMB companies have more flexibility in terms of leveraging standardized IT offerings with less customization and stringent service level requirements to keep their costs low, compared to large enterprises.

To deal with short term financial pressures, many SMB companies follow an IT purchasing strategy of choosing price over performance, and relying on platforms that staff members are familiar with, rather than alternatives that may offer features better suited to the company’s actual business and technical needs. Recent surveys show that increasingly, SMB companies are starting to look at overall costs of ownership at a system level as compared with hardware or software components only in the past. Just as with large enterprises, SMB companies appreciate IT vendors who can demonstrate complete solutions (combination of hardware, software, services, and the ability to integrate into their existing environments) and provide the best IT value in the long term.

8.1.3 SMB IT data center and staff issues

Most SMB IT staff have to support numerous IT platforms and a great variety of data center tasks, ranging from hardware and software installation, daily operations, help desk to troubleshooting problems. Because of the heavy load of fire fighting, IT management and staff usually spend little time on planning and procedures, impacting their overall productivity and the service levels to customers. These challenges and complexity tend to expand exponentially as the company grows, making it increasingly necessary and expensive to engage specialized contract services. Increasingly, SMB IT management pays more attention to planning and procedures to address these issues, especially in data center operations, such as backup, restore, and Disaster Recovery. This area of expertise is usually not of high priority to the SMB IT staff.

8.2 Business Continuity for SMB companies

The basic definition of Business Continuity is the ability to conduct business under any circumstances. From an IT standpoint, it is the ability to provide systems and data for business transactions to a set of service levels based on end-to-end availability, performance (such as response times), data security and integrity, and other factors. Service level agreements (SLAs) usually drive the BC design and budgets. For a variety of reasons, most SMB IT management does not have SLAs with the lines of business. As more SMB companies are leveraging their IT capabilities to drive revenues and profits, SLAs are increasingly required.

8.2.1 Major SMB business continuity design components

Particularly in SMB environments, these are the major BC design components:

- Prevention Services
- Recovery Services

Since budget and value are the decision criteria for SMB companies, recovery services are usually the starting points for BC. As prevention services are becoming more affordable, usually BC solutions consist of a combination of the two, depending on the company’s needs.

The three aspects of Business Continuity are:

- High Availability
- Continuous Operations
- Disaster Recovery

Let us examine how an SMB enterprise usually views these three aspects.

Prevention Services

Prevention Services are the ability to avoid outages or minimize down time by anticipation, systems planning, and high availability technology and solution deployment

High availability: Build reliability and redundancy into systems infrastructure, including hardware, software, and networks, to eliminate single points of failure; it also includes some automatic switchover or restart capabilities to minimize down time

Continuous operations: Minimize data center operation impacts on up time that include hardware and software maintenance and changes, backup, systems management, speedy problem resolution, virus and spam attacks prevention, security measures, and so on. The solutions usually involve management and process automation, systems and data consolidation (less to support), and improved efficiency of operations. More information about Continuous Availability solutions is in *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548.

Recovery services

Recovery services are the ability to recover the system and data speedily in whole, partial, or degraded modes when outages occur. Here we would examine:

Disaster Recovery: Invoked when the primary operation site is no longer operable and the alternate site is the only option

System component or operation recovery: Invoked when an individual component or group of components fail, or when human errors occur during operation

Service level targets will dictate the degrees of prevention and recovery services required and the budgets supporting them. Usually it is a decision on risks: a balance between the avoidance costs and BC solutions investments.

8.2.2 Business Continuity impacts on SMB business

Business Continuity impacts are usually measured in potential revenue and profits loss, staff productivity loss, customer and IBM Business Partner satisfaction and loyalty loss, and so on. Revenue and profit loss can be calculated by dollars lost by the inability to conduct business due to a system outage for a time frame. Other impacts can be estimated by industry averages. A risk assessment of the potential costs and the odds of the outages will be the primary factors for the BC measure necessity, design, and budgets.

8.3 Successful SMB Business Continuity planning and implementation

Here are the recommended planning and implementation steps, especially for the SMB enterprise.

1. Conduct a risk assessment to develop a set of BC service targets and IT metrics for key business processes with lines of business: The assessment results should determine BC priorities, scope, goals, budgets, and success criteria; the service targets can include end-to-end systems availability and response time, Disaster Recovery objectives, and so on.
2. Assess present attainment of these service targets and metrics: Establish a base line for comparison and understanding of the challenges to meeting the targets
3. Develop and evaluate technology and solution options: The success criteria should drive the evaluation and priority; the technology and solutions are fairly standard these days.
4. Develop an architecture and roadmap to support the solution implementation: BC solutions usually take some time to implement based on budgets and resource availability; a base architecture on which the solutions can build is critical
5. Develop an overall BC strategy and plan: It is important that the IT BC plan coordinates with the overall business plan.

8.3.1 SMB business continuity implementation steps

Here are the recommended steps for implementing Business Continuity, especially for the SMB enterprise.

1. Simplify, consolidate, standardize and centralize infrastructure: Reduce the number of servers, storage, and network equipment footprints, reduce the number of application instances and operating systems to be supported, reduce the complexity of backup and

management, deploy technologies such as server and storage virtualization, clustering and centralization, including SAN and NAS.

2. Build well documented and tested data center systems management procedures: The ability to minimize human errors and preventable outages is the key to minimizing down time.
3. Acquire systems management tools to monitor, prevent outages, automate diagnostics and recovery, and report to stakeholders: Tools are important to prevent and predict outages and avoid them.
4. Make BC a strategic part of application and IT infrastructure planning: Business Continuity, based on SLA targets (both IT internal and lines of business external) must be key system acquisition and design criteria.

8.3.2 SMB Business Continuity affordability

There are two major factors in assessing affordability:

- How much one can afford to lose
- How much one can afford to pay

Basically, this is a risk and investment assessment. It is somewhat similar to a home owner's insurance. Although BC is more than loss recovery, it can be used to drive the positive aspects of the business. It can be leveraged to increase business, improve staff productivity, and build confidence and trust with customers and partners.

Calculating affordability

Here are the steps we recommend for calculating affordability, especially for the SMB enterprise.

Recovery objective

Determine:

- How much downtime can your business tolerate before it starts to hurt your bottom line (potential revenues and profits loss, customer satisfaction or defection, staff morale, business partnership breakage, and government regulatory liabilities)? Is the affordable downtime in seconds, minutes, hours, or days?
- How much data loss and what data loss will start to hurt bottom line? For what period of time?

Budget objective

Determine:

- How much money loss can be attributed to the outages the business can afford?
- What are the odds of outage occurring?
- What is the percentage of the potential loss the business is willing and can afford to pay? The ratios vary by industries and business types (reliance on IT). They can range from 10 to >1% of the total IT annual budget (ongoing and capital).

8.4 SMB Business Continuity solution components

The following tables list the components that typically make up a cost-effective SMB BC solution, at differing levels of recovery.

While your results will vary according to your specific requirements, this chart will give a good beginning guideline. You may use it to build your own specific chart for your enterprise.

In Figure 8-1, we diagram the typical BC solution components, according to their tier level of recovery.

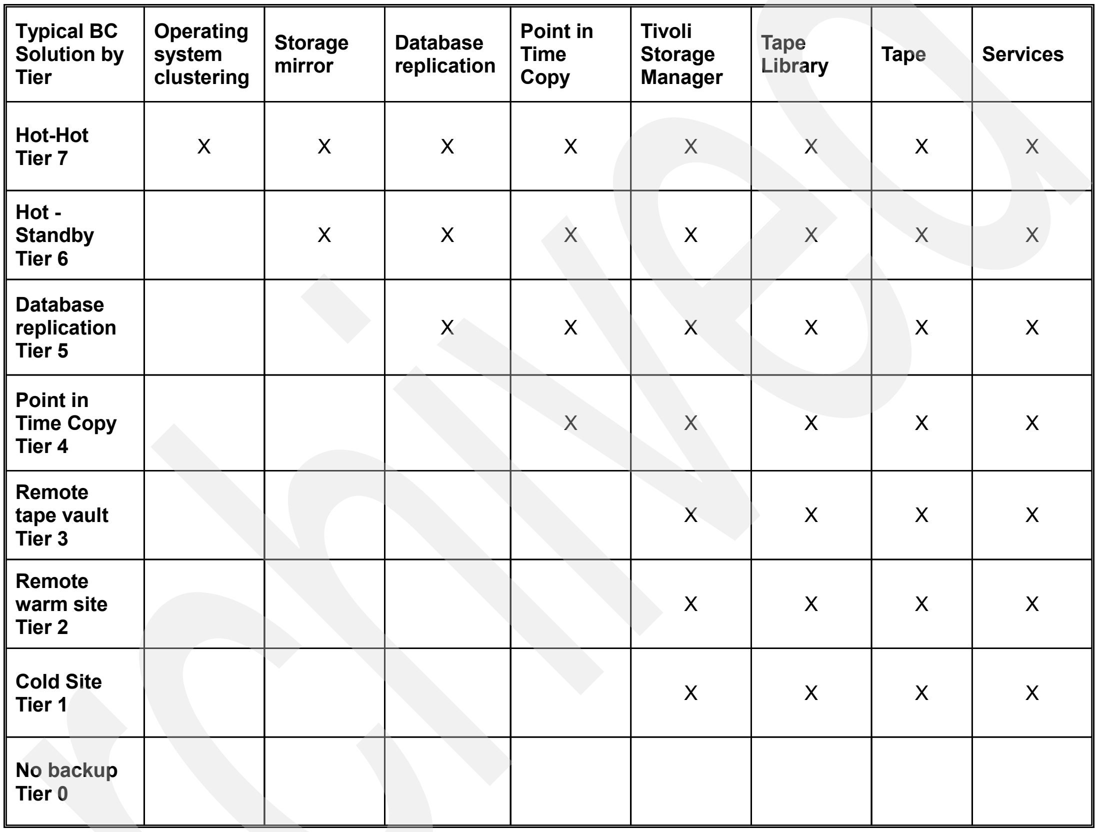

Typical BC Solution by Tier	Operating system clustering	Storage mirror	Database replication	Point in Time Copy	Tivoli Storage Manager	Tape Library	Tape	Services
Hot-Hot Tier 7	X	X	X	X	X	X	X	X
Hot - Standby Tier 6		X	X	X	X	X	X	X
Database replication Tier 5			X	X	X	X	X	X
Point in Time Copy Tier 4				X	X	X	X	X
Remote tape vault Tier 3					X	X	X	X
Remote warm site Tier 2					X	X	X	X
Cold Site Tier 1					X	X	X	X
No backup Tier 0								

Figure 8-1 Small Medium Business typical business continuity solution components

8.4.1 Typical SMB BC solutions: performance and downtime

The following are the typical performance and downtime characteristics of typical BC solutions in the SMB environment.

While your results will vary according to your specific requirements, this chart should give a good beginning guideline to what you may expect at differing tier levels of recovery.

With the sample chart shown in Figure 8-2, you may build your own specific characteristics chart for your enterprise's BC solution.

Components for typical SMB business continuity solutions are described in their respective chapters.

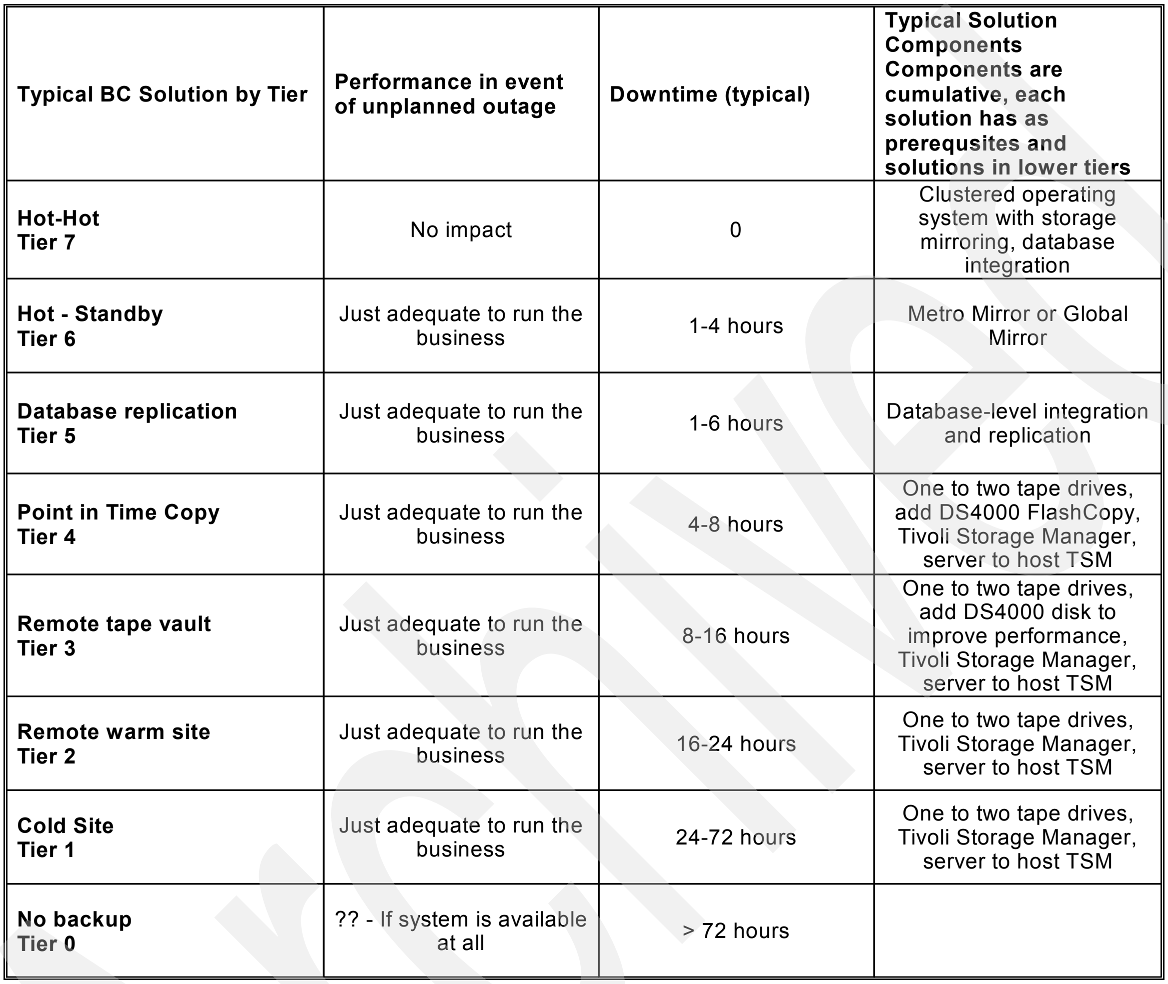

Typical BC Solution by Tier	Performance in event of unplanned outage	Downtime (typical)	Typical Solution Components Components are cumulative, each solution has as prerequsites and solutions in lower tiers
Hot-Hot **Tier 7**	No impact	0	Clustered operating system with storage mirroring, database integration
Hot - Standby **Tier 6**	Just adequate to run the business	1-4 hours	Metro Mirror or Global Mirror
Database replication **Tier 5**	Just adequate to run the business	1-6 hours	Database-level integration and replication
Point in Time Copy **Tier 4**	Just adequate to run the business	4-8 hours	One to two tape drives, add DS4000 FlashCopy, Tivoli Storage Manager, server to host TSM
Remote tape vault **Tier 3**	Just adequate to run the business	8-16 hours	One to two tape drives, add DS4000 disk to improve performance, Tivoli Storage Manager, server to host TSM
Remote warm site **Tier 2**	Just adequate to run the business	16-24 hours	One to two tape drives, Tivoli Storage Manager, server to host TSM
Cold Site **Tier 1**	Just adequate to run the business	24-72 hours	One to two tape drives, Tivoli Storage Manager, server to host TSM
No backup **Tier 0**	?? - If system is available at all	> 72 hours	

Figure 8-2 Typical SMB business continuity solution: performance and downtime characteristics

The definitions of these terms, the tiers, and the various types of solution components, are in the other chapters of this IBM Redbook.

The components shown above are not the only components or products that may be part of the solution; these are meant as general guidelines. The products shown are typical, and may be substituted as specific client requirements dictate.

Other variations of these typical solutions can include network attached storage (NAS) devices, centralized tape libraries, and other products for specific circumstances. SMB companies, just like larger enterprises, can scale up the tiers by deploying additional solutions and technologies as the business grows.

8.5 Summary

Business Continuity is, and will be, a key requirement for SMB to conduct business. Solutions and technologies continue to improve and be affordable to SMB companies. It is important for SMB IT management to incorporate Business Continuity planning in their strategy and building systems and applications from the beginning. It will cost less and help drive their business objectives from the start.

A

Tiers of Business Continuity

In this appendix, we provide a further definition of the tiers of Business Continuity.

Further definition of the tiers of Business Continuity

In this IBM Redbook, all of the Business Continuity technologies we describe are listed by tier level.

Assignment of tier levels is subjective and an approximation, not an exact science. Large systems will tend to shift the tier levels to the right; small scale systems will tend to shift the tiers to the left. As part of the planning process, you should create a tiers of Business Continuity chart specific to your organization, and use it as a specific organization tool within your own organization. We offer here a generalized view of the tiers.

BC Tier 0: No off-site data

Businesses with a Tier 0 BC solution have no Business Continuity Plan.

- There is no saved information, no documentation, no backup hardware, and no contingency plan.
- Typical recovery time: The length of recovery time in this instance is unpredictable. In fact, it may not be possible to recover at all.

8.5.1 BC Tier 1: Data backup with no hot site

Businesses that use Tier 1 BC solutions back up their data typically with tape, at an off-site facility. Depending on how often backups are made, they are prepared to accept several days to weeks of data loss, but their backups are secure off-site. However, this tier lacks the systems on which to restore data.

Tier 1 Business Continuity solutions:

- Pickup Truck Access Method (PTAM): Sending physical tapes
- Disk Subsystem or Tape based mirroring to locations without processors
- IBM Tivoli Storage Manager

BC Tier 2: Data backup with a hot site

Businesses using Tier 2 BC solutions make regular backups on tape. This is combined with an off-site facility and infrastructure (known as a hot site) in which to restore systems from those tapes in the event of a disaster. This tier of solution will still result in the need to recreate several hours to days worth of data, but it is faster and more consistent, compared to BC tier 1 in recovery time.

Tier 2 Business Continuity solutions:

- PTAM with hot-site available
- IBM Tivoli Storage Manager

BC Tier 3: Electronic vaulting

Tier 3 BC solutions utilize components of Tier 2. Additionally, some mission critical data is electronically vaulted. This electronically vaulted data is typically more current than that which is shipped via PTAM. We also add higher levels of automation; as a result there is less data recreation or loss.

Tier 3 Business Continuity solutions:

- Electronic Vaulting of Data (that is, remote tape via channel extension)
- IBM Tivoli Storage Manager (Disaster Recovery Manager)

BC Tier 4: Point-in-time copies

Tier 4 BC solutions are used by businesses requiring both greater data currency and faster recovery than users of lower tiers. Rather than relying largely on shipping tape, as is common on the lower tiers, Tier 4 solutions begin to incorporate more disk based solutions, typically using point in time disk copy. Several hours of data loss is still possible, but it is easier to make such point-in-time (PIT) copies with greater frequency than data can be replicated through tape based solutions.

Tier 4 Business Continuity solutions:

- Batch/Online Database Shadowing and Journaling
- Global Copy
- FlashCopy and FlashCopy Manager
- Peer-to-Peer Virtual Tape Server
- Metro/Global Copy
- IBM Tivoli Storage Manager (Disaster Recovery Manager)
- IBM Tivoli Storage Manager for Copy Services
- IBM Tivoli Storage Manager for Advanced Copy Services
- TotalStorage Productivity Center for Replication
- N series Snapshot with N series software

BC Tier 5: Transaction integrity

Tier 5 BC solutions is a BC Tier reserved for application software and database replication at the transaction level. This BC solution are used by businesses with a requirement for consistency of data between production and recovery data centers. There is little to no data loss in such solutions; however, the presence of this functionality is entirely dependent on the applications in use.

Tier 5 Business Continuity solutions:

- Software, two-phase commit, such as DB2 remote replication and MQ Series
- Oracle Data Guard

BC Tier 6: Zero or little data loss

Tier 6 BC solutions maintain the highest levels of data currency. They are used by businesses with little or no tolerance for data loss and who need to restore data to applications rapidly. These solutions have no dependence on the applications to provide data consistency; typically, these solutions use real-time storage mirroring or server mirroring.

Tier 6 Business Continuity solutions:

- Metro Mirror
- Global Mirror
- z/OS Global Mirror
- GDPS HyperSwap Manager
- Peer-to-Peer VTS with synchronous write
- PPRC Migration Manager
- TotalStorage Productivity Center for Replication
- HACMP/XD with Logical Volume Mirroring

BC Tier 7: Highly automated, business integrated solution

Tier 7 BC solutions include all the major components being used for a Tier 6 BC solution with the additional integration of end-to-end automation for servers, storage, software, and applications and network. This allows a Tier 7 BC solution to ensure consistency of data above that which is granted by Tier 6 BC solutions. Additionally, recovery of the applications is automated, allowing for restoration of systems and applications much faster and more reliably than would be possible through manual Business Continuity procedures.

Tier 7 Business Continuity solutions:

- GDPS/PPRC with or without HyperSwap
- GDPS/XRC
- GDPS/GM
- AIX HACMP/XD with Metro Mirror
- System i High Availability Business Partner Software
- System i Copy Services Toolkit

B

IBM System Storage Components and Copy Services

This appendix provides an overview of the IBM System Storage components (DS8000, DS6000, ESS, DS4000, SVC, N series, Open Virtual Tape (TS7510), and TS7740) Grid that are used in some of the solutions described in this IBM Redbook, including:

- Disk storage systems (DS8000, DS6000, and DS4000)
- N series
- SAN Volume Controller (SVC)
- TS7740 Grid
- TS7510 Virtualization Engine

This also includes the Advanced Copy Services core technologies that are used:

- Point-in-Time copy
- Disk mirroring
  - Synchronous
  - Asynchronous
  - Three-site mirroring

IBM System Storage DS Family

The IBM System Storage DS™ Family includes:

- IBM System Storage DS8000 Series
- IBM System Storage DS6000 Series
- IBM System Storage DS4000 Series

IBM System Storage DS8000 series

The IBM DS8000 series is the top of the line high performance disk system for business critical enterprise workloads. DS8000 supports all the advanced features of the DS6000 and ESS, plus significant enhancements. This new, high-end member of the IBM System Storage family offers a significant step forward in performance, virtualization capability, and capacity.

The DS8000 uses IBM POWER5 processors and IBM Virtualization Engine technology to provide logical partitioning. The DS8000 scales from 1.1 TB to 320 TB, and is architected to scale to over a petabyte. Up to 256 GB cache memory and up to 128 FICON/Fibre Channel ports can be installed to accommodate a wider range of workloads.

IBM System Storage DS8000 supports heterogeneous environments, such as IBM z/OS, IBM OS/400®, IBM AIX, Linux, UNIX, Microsoft Windows, HP-UX, and SUN Solaris.

IBM System Storage DS6000

The IBM DS6000 is a mid-range disk storage system with the resiliency, performance, and many key features of the larger IBM DS8000, in a rack-mountable 3U-high unit. It can be a lower cost alternative for a secondary remote mirror site, or used in a Test/Development environment.

IBM System Storage DS6000 supports heterogeneous environments, such as IBM z/OS, IBM OS/400, IBM AIX, Linux, UNIX, Microsoft Windows, HP-UX, and SUN Solaris.

The IBM TotalStorage ESS

The IBM TotalStorage Enterprise Storage Server (ESS 750/800) was the precursor for the DS6000 and DS8000. Much of its technology forms the basis for the new products. Although it is no longer actively sold by IBM, the ESS remains widely deployed in IT environments.

Advanced Copy Services for DS6000 and DS8000

This section describes the core data replication technologies available on the DS6000 and DS8000 storage systems. These are:

- FlashCopy, which is a Point-in-Time Copy function
- Remote Mirror and Copy functions, which include:
  - Metro Mirror
  - Global Copy
  - Global Mirror
- z/OS Global Mirror
- Metro/Global Copy
- Metro/Global Mirror
- z/OS Metro/Global Mirror

FlashCopy

FlashCopy is a fast point-in-time (T0) disk to disk copy technique. It can be used to copy full volumes or z/OS data sets. It supports multiple targets for a source on different volumes/LSSs (see Figure B-1). FlashCopy can be used for cloning for test environments, database query, or data mining; and making rapid data backups, minimizing user downtime.

When FlashCopy is invoked, the internal microcode initializes pointers to the T0 copy. This happens very quickly (only a matter of seconds in many cases). When the initialization is complete, the source volumes and the copied volumes are available for full read and write, thus making an efficient point-in-time copy on the disk system.

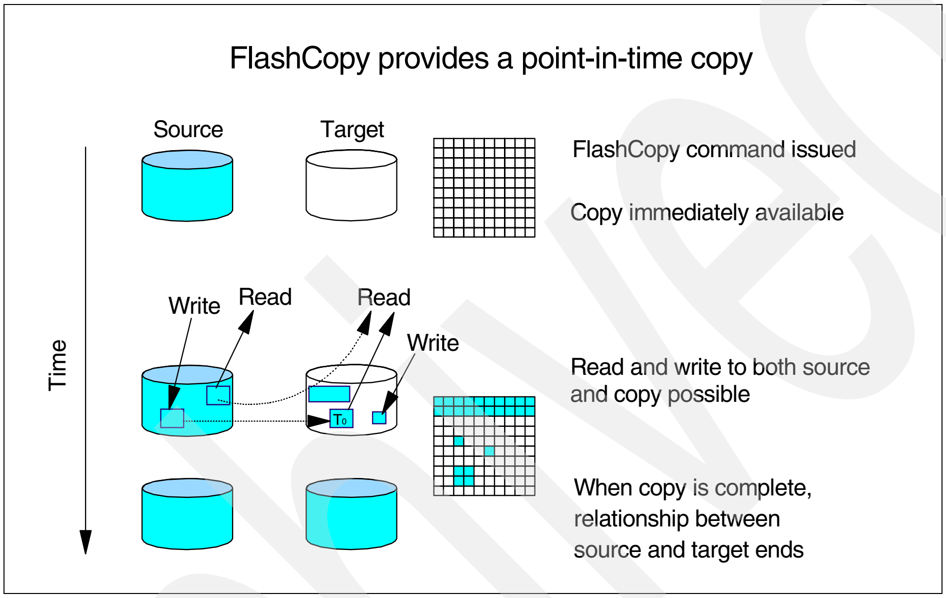
FlashCopy provides a point-in-time copy
Source
Target
FlashCopy command issued
Copy immediately available
Time
Write
Read
Read
Write
T0
Read and write to both source and copy possible
When copy is complete, relationship between source and target ends

Figure B-1 FlashCopy

Metro Mirror

Metro Mirror (see Figure B-2) is a synchronous real-time copy from one volume (LUN) to another volume. The volumes can be in the same storage system or on different ones, up to a distance of 300 km using DWDM. Greater distances are supported on special request.

The Metro Mirror protocol enforces data currency between the primary and the secondary volumes by acknowledging that the I/O to the secondary has been written successfully, though the distance between the primary and secondary impacts the performance.

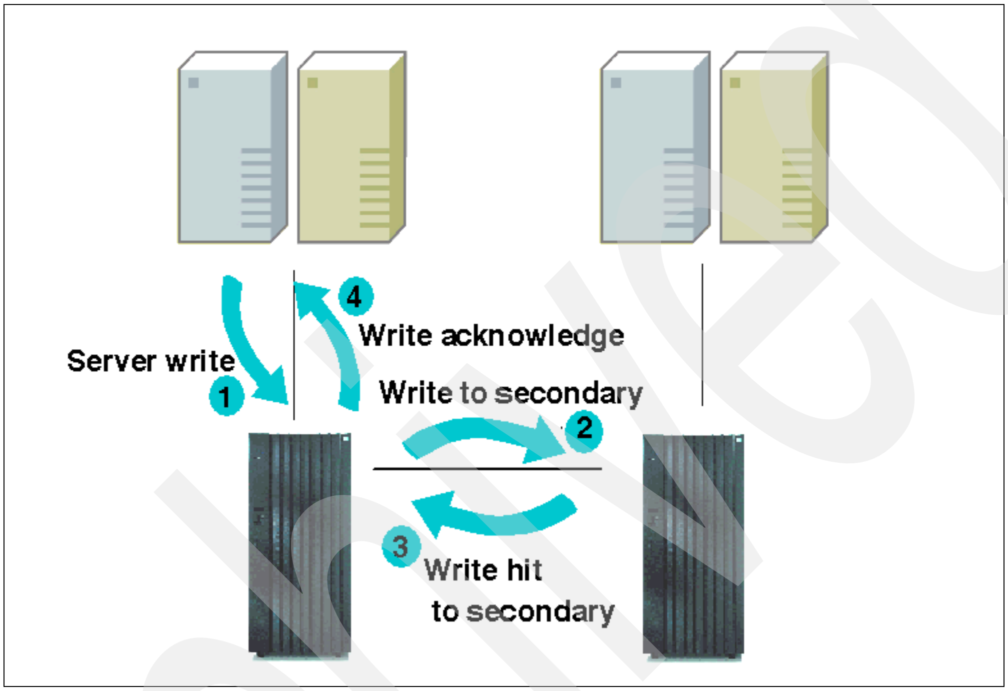

Figure B-2 Metro Mirror

Global Copy

Global Copy (see Figure B-3) is an asynchronous remote copy function for z/OS and open systems for long and continental distances. With Global Copy, write operations are acknowledged as complete on the primary storage system before they are received by the secondary storage system.

Global Copy is designed to prevent the primary system's performance from being affected by wait time from writes on the secondary system. Therefore, the primary and secondary copies can be separated by any distance depending on your storage networking.

This function is appropriate for remote data migration, off-site backups, and transmission of data at virtually unlimited distances.

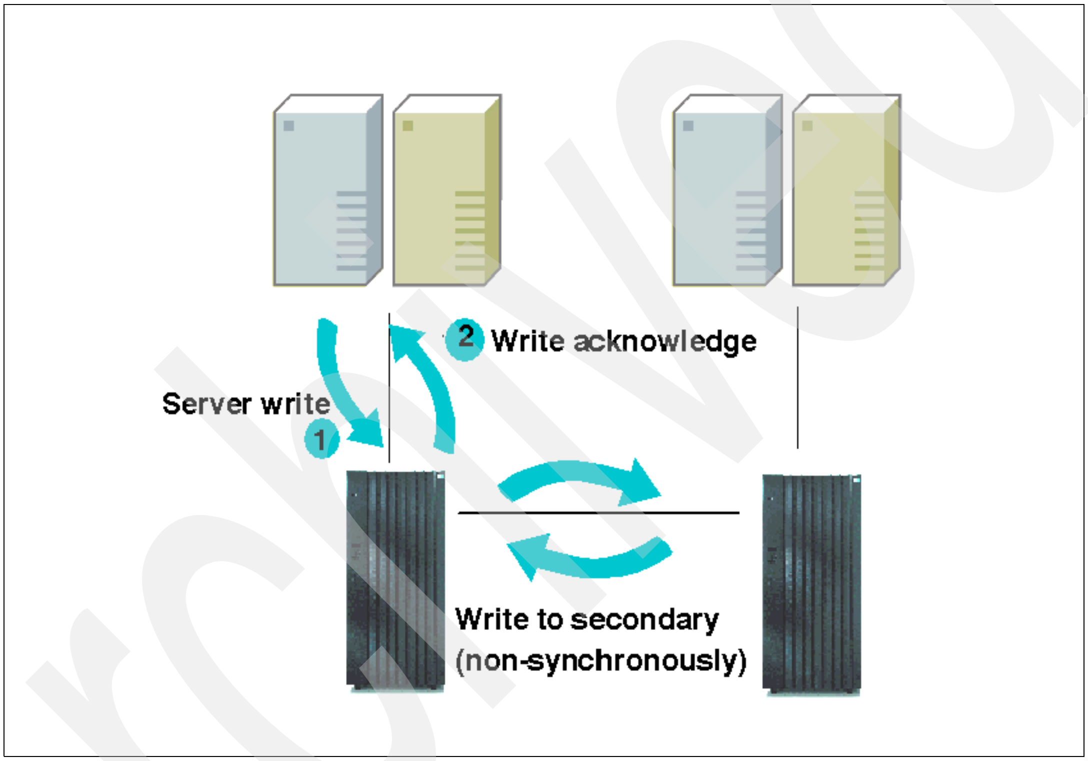

Figure B-3 Global Copy

Global Mirror

Global Mirror provides a two-site, extended distance remote mirroring function for z/OS and open systems servers.

With Global Mirror (see Figure B-4), the data that the host writes to the storage unit at the local site is asynchronously shadowed to a storage unit at the remote site. The volumes can reside on several different storage units. A consistent copy of the data is then automatically maintained on the storage unit at the remote site. This two site data mirroring function is designed to provide a high-performance, cost-effective global distance data replication and Disaster Recovery solution.

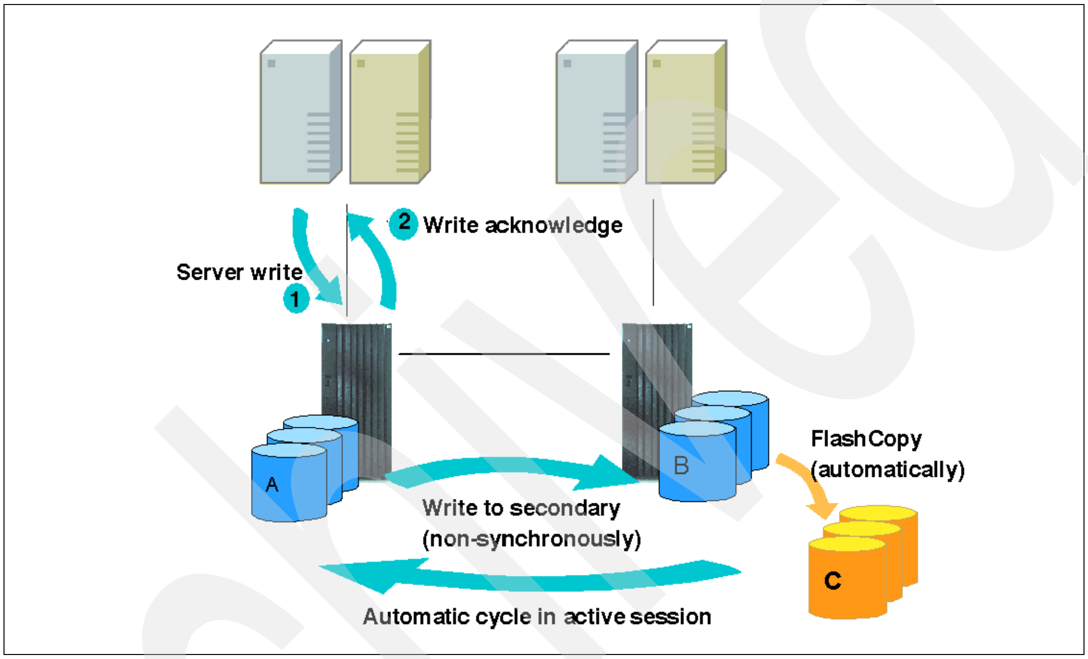

Figure B-4 Global Mirror

z/OS Global Mirror

z/OS Global Mirror (see Figure B-5 on page 127), only supports System z hosts, and mirrors data on the storage unit to a remote location for disaster recovery. It protects data consistency across all volumes defined for mirroring. The volumes can reside on several different storage units. The z/OS Global Mirror function can mirror the volumes over several thousand kilometers from the source site to the target recovery site. With z/OS Global Mirror, you can suspend or resume service during an outage. You do not have to terminate your current data-copy session. You can suspend the session, then restart it. Only data that changed during the outage needs to be re-synchronized between the copies.

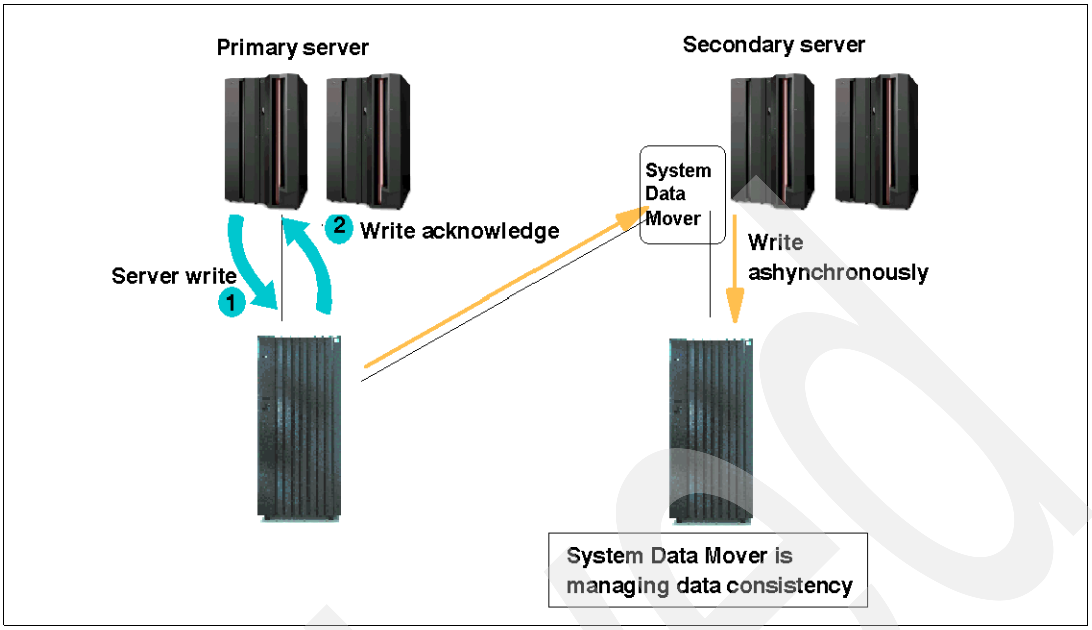

Figure B-5 z/OS Global Mirror

Metro/Global Copy (3-site Metro Mirror and Global Copy)

Metro/Global Copy (see Figure B-6) is a cascaded three-site disk mirroring solution. Metro Mirror is used between production site A and intermediate site B. Global Copy is used between the intermediate site B and the remote site C. Metro/Global Copy is often used for migration purposes in a two site mirroring environment. Global Copy keeps the cascaded copy (which can be located at either the remote or local site) nearly current with the running two-site disk mirroring configuration.

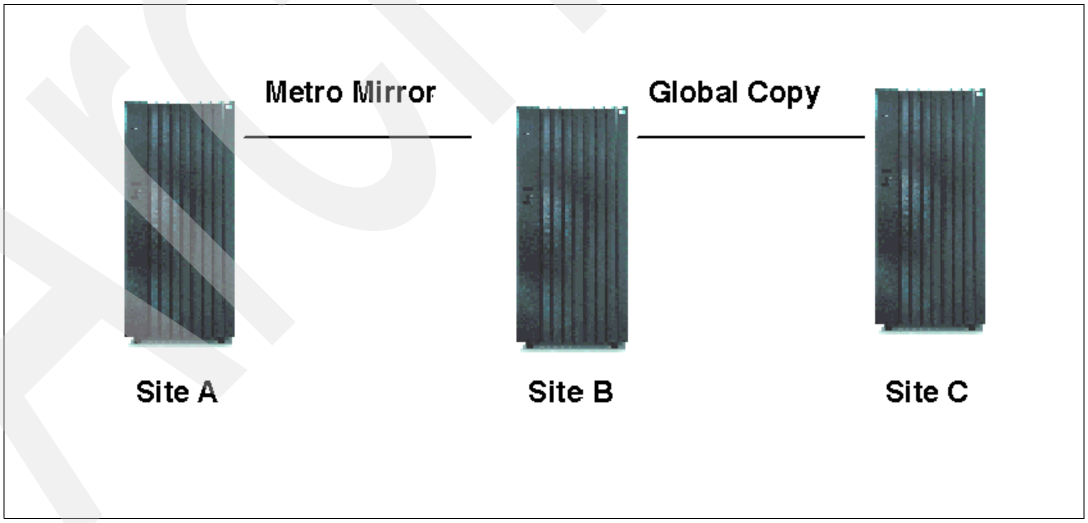

Figure B-6 Metro/Global Copy

Metro/Global Mirror

Metro/Global Mirror (see Figure B-7) is a cascaded three-site disk mirroring solution. Metro Mirror is used between production site A and intermediate site B. Global Mirror is used between the intermediate site B and the remote site C. In the event of a loss of access to intermediate site B, the license for DS8000 Metro/Global Mirror provides new functionality to incrementally resynchronize and establish Global Mirror from production site A to remote site C, without application impact to site A, thus maintaining out of region Disaster Recovery. When intermediate site B access returns, that site can be re-inserted into the three site cascading topology without impact to production site A applications. In all cases, only incremental changes need to be sent for resynchronization.

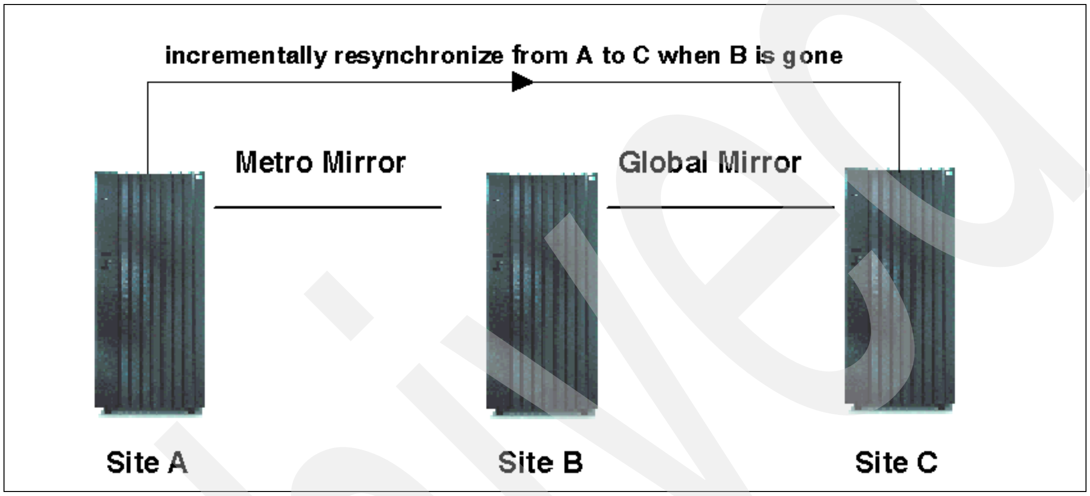

Figure B-7 Metro/Global Mirror

z/OS Metro/Global Mirror

z/OS Metro/Global Mirror (see Figure B-8) uses z/OS Global Mirror to mirror primary site data to a location that is a long distance away and also uses Metro Mirror to mirror primary site data to a location within the metropolitan area. This enables a z/OS 3-site high availability and Disaster Recovery solution for even greater protection from unplanned outages. The z/OS Metro/Global Mirror function is an optional function.

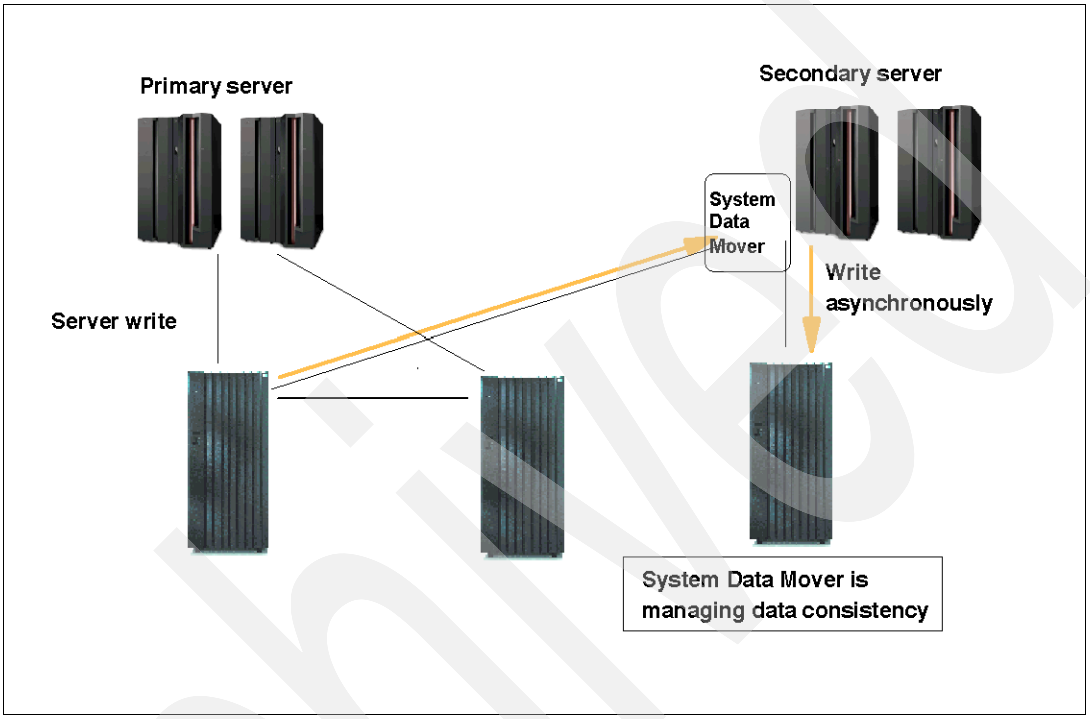

Figure B-8 z/OS Metro/Global

IBM System Storage DS4000 Family

The IBM System Storage DS4000 Series, is a scalable, flexible series of disk storage products, which supports a wide range of environments, including IBM AIX, Linux, UNIX, Microsoft Windows, HP-UX, and SUN Solaris.

The various DS4000 models available, and their hardware specifications, are shown in Figure B-9, including the connectivity, scalability, and performance capabilities.

The DS4000 series offers exceptional performance, robust functionality, and unparalleled ease of use. The DS4000 series offers 4 Gbps FC technology while retaining compatibility with existing 1 and 2 Gbps devices.

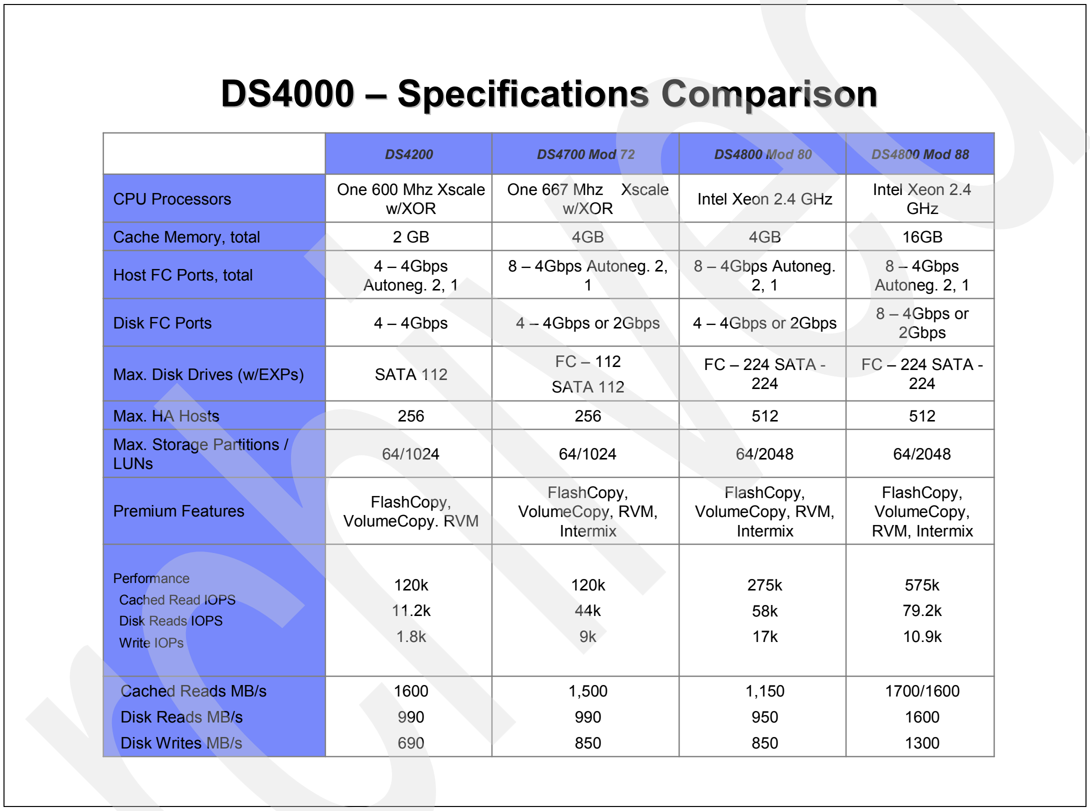

DS4000 – Specifications Comparison

	DS4200	DS4700 Mod 72	DS4800 Mod 80	DS4800 Mod 88
CPU Processors	One 600 Mhz Xscale w/XOR	One 667 Mhz Xscale w/XOR	Intel Xeon 2.4 GHz	Intel Xeon 2.4 GHz
Cache Memory, total	2 GB	4GB	4GB	16GB
Host FC Ports, total	4 – 4Gbps Autoneg. 2, 1	8 – 4Gbps Autoneg. 2, 1	8 – 4Gbps Autoneg. 2, 1	8 – 4Gbps Autoneg. 2, 1
Disk FC Ports	4 – 4Gbps	4 – 4Gbps or 2Gbps	4 – 4Gbps or 2Gbps	8 – 4Gbps or 2Gbps
Max. Disk Drives (w/EXPs)	SATA 112	FC – 112 SATA 112	FC – 224 SATA - 224	FC – 224 SATA - 224
Max. HA Hosts	256	256	512	512
Max. Storage Partitions / LUNs	64/1024	64/1024	64/2048	64/2048
Premium Features	FlashCopy, VolumeCopy. RVM	FlashCopy, VolumeCopy, RVM, Intermix	FlashCopy, VolumeCopy, RVM, Intermix	FlashCopy, VolumeCopy, RVM, Intermix
Performance Cached Read IOPS Disk Reads IOPS Write IOPs	 120k 11.2k 1.8k	 120k 44k 9k	 275k 58k 17k	 575k 79.2k 10.9k
Cached Reads MB/s Disk Reads MB/s Disk Writes MB/s	1600 990 690	1,500 990 850	1,150 950 850	1700/1600 1600 1300

Figure B-9 DS4000 comparative specifications

Advanced Copy Services for DS4000

This section describes the core data replication technologies available on the DS4000 family storage subsystems. These include:

- FlashCopy/VolumeCopy
- Enhanced Remote Mirroring
  - Metro Mirror
  - Global Copy
  - Global Mirror

FlashCopy

A FlashCopy logical drive, shown in Figure B-10, is a point-in-time image of a logical drive. It is the logical equivalent of a complete physical copy, but it is created much more quickly than a physical copy. It also requires less disk space. On the other hand, it is not a real physical copy, because it does not copy all the data. Consequently, if the source logical drive is damaged, the FlashCopy logical drive cannot be used for recovery.

In the DS4000 Storage Manager, the logical drive from which you are basing the FlashCopy, called the *base logical drive,* can be a standard logical drive or the *secondary logical drive* in a Remote Mirror relationship. Typically, you create a FlashCopy so that an application (for example, an application to take backups) can access the FlashCopy and read the data while the base logical drive remains online and user-accessible. In this case, the FlashCopy logical drive is no longer needed (it is usually disabled rather than deleted) once the backup completes.

You can also create multiple FlashCopies of a base logical drive and use the copies in write mode to perform testing and analysis. Before you upgrade your database management system, for example, you can use FlashCopy logical drives to test different configurations. Then you can use the performance data provided during the testing to help decide how to configure the live database system.

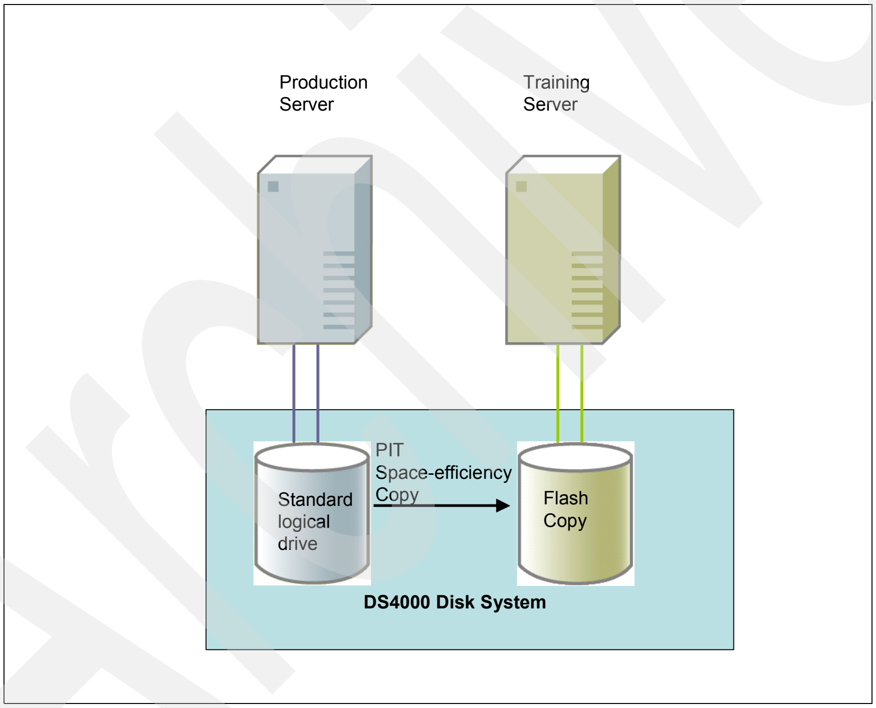

Figure B-10 FlashCopy

VolumeCopy

The VolumeCopy premium feature copies data from one logical drive (source) to another logical drive (target) in a single disk system (Figure B-11). The target logical drive is an exact copy or *clone* of the source logical drive. This feature can be used to copy data from arrays that use smaller capacity drives to arrays that use larger capacity drives, to back up data, or to restore a FlashCopy to the base logical drive.

The VolumeCopy premium feature must be enabled by purchasing a feature key. FlashCopy must be installed as well. VolumeCopy is only available as a bundle that includes a FlashCopy license.

VolumeCopy is a full point-in-time replication. It allows for analysis, mining, and testing without any degradation of the production logical drive performance. It also brings improvements to backup and restore operations, making them faster and eliminating I/O contention on the primary (source) logical drive.

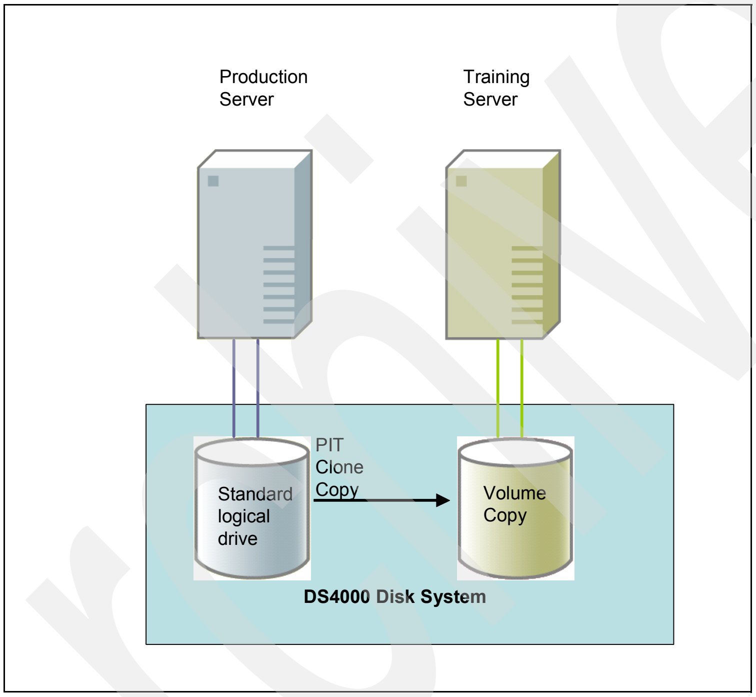

Figure B-11 VolumeCopy

Enhanced Remote Mirroring

The Enhanced Remote Mirroring (ERM) option is a premium feature of the IBM DS4000 Storage Manager software. It is enabled by purchasing a premium feature key.

Enhanced Remote Mirroring is used for online, real-time replication of data between DS4000 systems over a remote distance. In the event of disaster or unrecoverable error at one DS4000, you can promote the second DS4000 to take over responsibility for normal I/O.

Data can be replicated between using different mirroring modes. In the remainder of this section, we explain the write operation sequences and processes involved for the different ERM mirroring modes. These are:

- Metro Mirroring
- Global Copy
- Global Mirroring

In all cases, data on the secondary logical drive of a mirror relationship can only be changed through the mirroring process. It cannot be changed by the host, or manually.

The read operations are identically treated in all three modes because, for a read request, there is no data exchange between the primary and secondary logical drives.

Metro Mirroring (synchronous mirroring)

Metro Mirroring is a synchronous mirroring mode; the controller does not send an I/O complete status to the host until the data has been copied to both the primary and secondary logical drives.

When a primary controller receives a write request from a host, the controller first logs information about the write request on the *mirror repository logical drive*. In parallel, it writes the data to the primary logical drive. The controller then initiates a remote write operation to copy the affected data blocks to the secondary logical drive at the remote site. When the remote write operation is complete, the primary controller removes the log record from the mirror repository logical drive. Finally, the controller sends an I/O completion indication back to the host system

When write caching is enabled on either the primary or secondary logical drive, the I/O completion is sent when data is in the cache on the site (primary or secondary) where write caching is enabled. When write caching is disabled on either the primary or secondary logical drive, then the I/O completion is not sent until the data has been stored to physical media on that site.

Figure B-12 shows the data flow for Metro Mirroring.

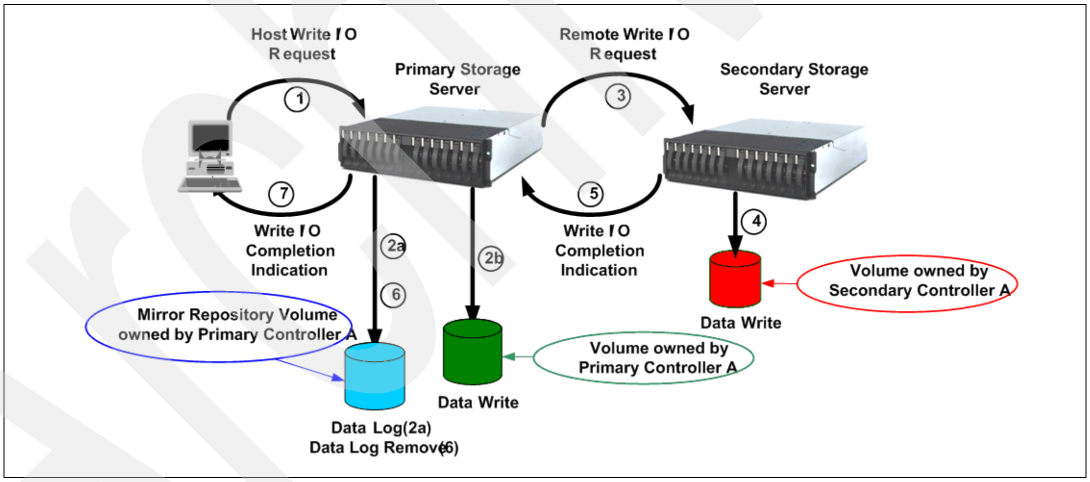

Figure B-12 Metro Mirroring Mode (Synchronous Mirroring) data flow

When a controller receives a read request from a host system, the read request is handled on the primary disk system and no communication takes place between the primary and secondary disk systems.

Global Copy (asynchronous mirroring without write consistency group)

Global Copy is an asynchronous write mode. All write requests from host are written to the primary (local) logical drive and immediately reported as complete to the host. Regardless of when data was copied to the remote disk system, the application does not wait for the I/O write request result from the remote site. Global Copy does not ensure that write requests at the primary site are processed in the same order at the remote site. As such, it is also referred to as *asynchronous mirroring without write consistency group*.

When a primary controller receives a write request from a host, the controller first logs information about the write request on the *mirror repository logical drive*. In parallel, it writes the data to the primary logical drive (or cache). After the data has been written (or cached), the host receives an I/O completion from the primary controller. The controller then initiates a background remote write operation to copy the corresponding data blocks to the secondary logical drive at the remote site. After the data has been copied to the secondary logical drive at the remote site (or cached), the primary controller removes the log record on the mirror repository logical drive.

When multiple mirror relationships are defined on the disk system, the background synchronization of affected data blocks between the primary and secondary controller for the different relationships are conducted in parallel (a multi-threaded process). Thus, the write order for multiple volumes (for example, write requests to a database volume and a database log volume on a database server) is not guaranteed with the Global Copy mode.

See Figure B-13 for a logical view of the data flow.

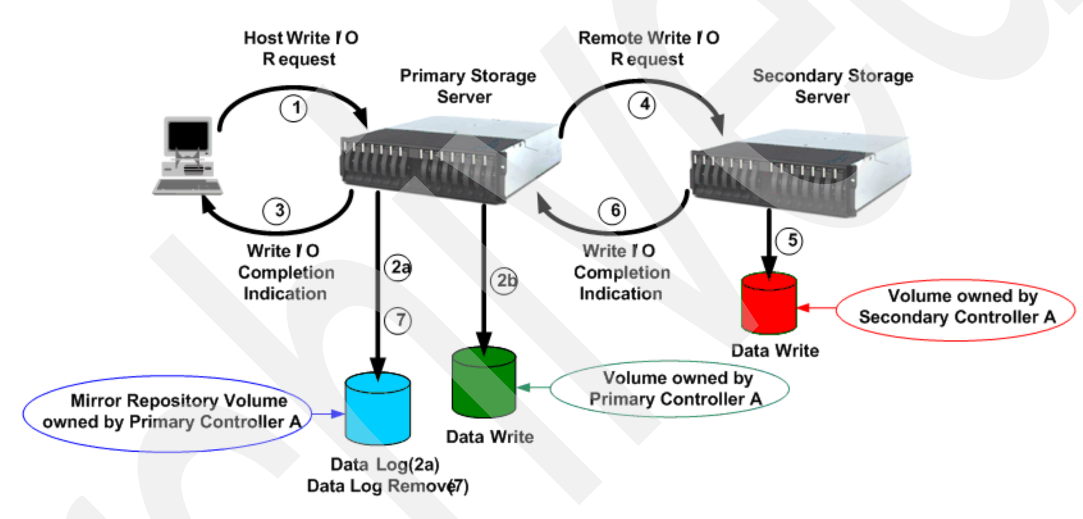
Host Write I O Request
Primary Storage Server
Remote Write I O Request
Secondary Storage Server
1
3
4
6
5
2a
7
2b
Write I O Completion Indication
Write I O Completion Indication
Volume owned by Secondary Controller A
Data Write
Volume owned by Primary Controller A
Mirror Repository Volume owned by Primary Controller A
Data Write
Data Log(2a)
Data Log Remove(7)

Figure B-13 Global Copy Mode (Asynchronous Mirroring) data flow

When write caching is enabled on either the primary or secondary logical drive, the I/O completion is sent when data is in the cache on the site (primary or secondary) where write caching is enabled. When write caching is disabled, then the I/O completion is not sent until the data has been stored to physical media on that site.

When a controller receives a read request from a host system, the read request is handled on the primary disk system and no communication takes place between the primary and secondary disk systems.

Global Mirroring (asynchronous mirroring with write consistency group)

Global Mirroring is an asynchronous write mode where the order of host write requests at the primary site is preserved at the secondary site. This mode is also referred as asynchronous mirroring with write consistency group.

To preserve the write order for multiple mirrored volumes, Global Mirroring uses the *write consistency group* functionality. It tracks the order of the host write requests, queues them, and sends them to the remote controller in the same order.

The volumes for which the write request order must be preserved have to be defined as members of a Write Consistency Group.

When a primary controller receives a write request from a host, the controller first logs information about the write on the *mirror repository logical drive*. It then writes the data to the primary logical drive. The controller then initiates a remote write operation to copy the affected data blocks to the secondary logical drive at the remote site. The remote write request order corresponds to the host write request order.

After the host write to the primary logical drive is completed and the data has been copied to the secondary logical drive at the remote site, the controller removes the log record from the mirror repository logical drive. Figure B-14 shows a logical view of the data flow.

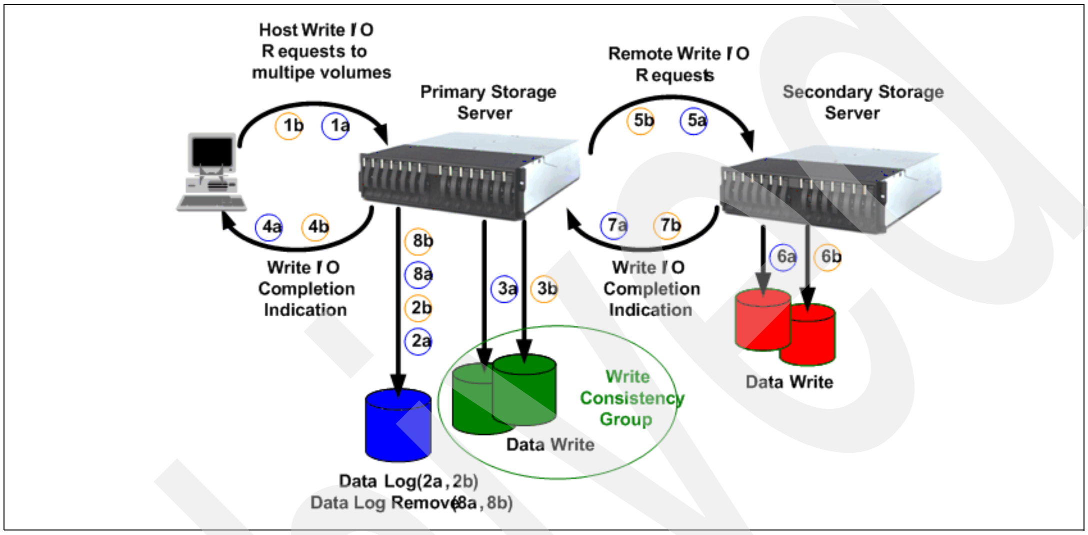

Figure B-14 Global Mirroring logical data flow

IBM System Storage N series

The N series products (Figure B-15) provide network attached storage to a broad range of host and client systems using multiple network access protocols, including CIFS, NFS, and block I/O protocols (including iSCSI and FCP), all from a single hardware platform, simultaneously. N series supports both Fibre Channel and SATA disk drives.

All N series systems utilize a single operating system (Data ONTAP®) across the entire platform and offer a combination of multiple advanced function software features for comprehensive system management, storage management, onboard and outboard copy services, virtualization technologies, and Disaster Recovery and backup solutions.

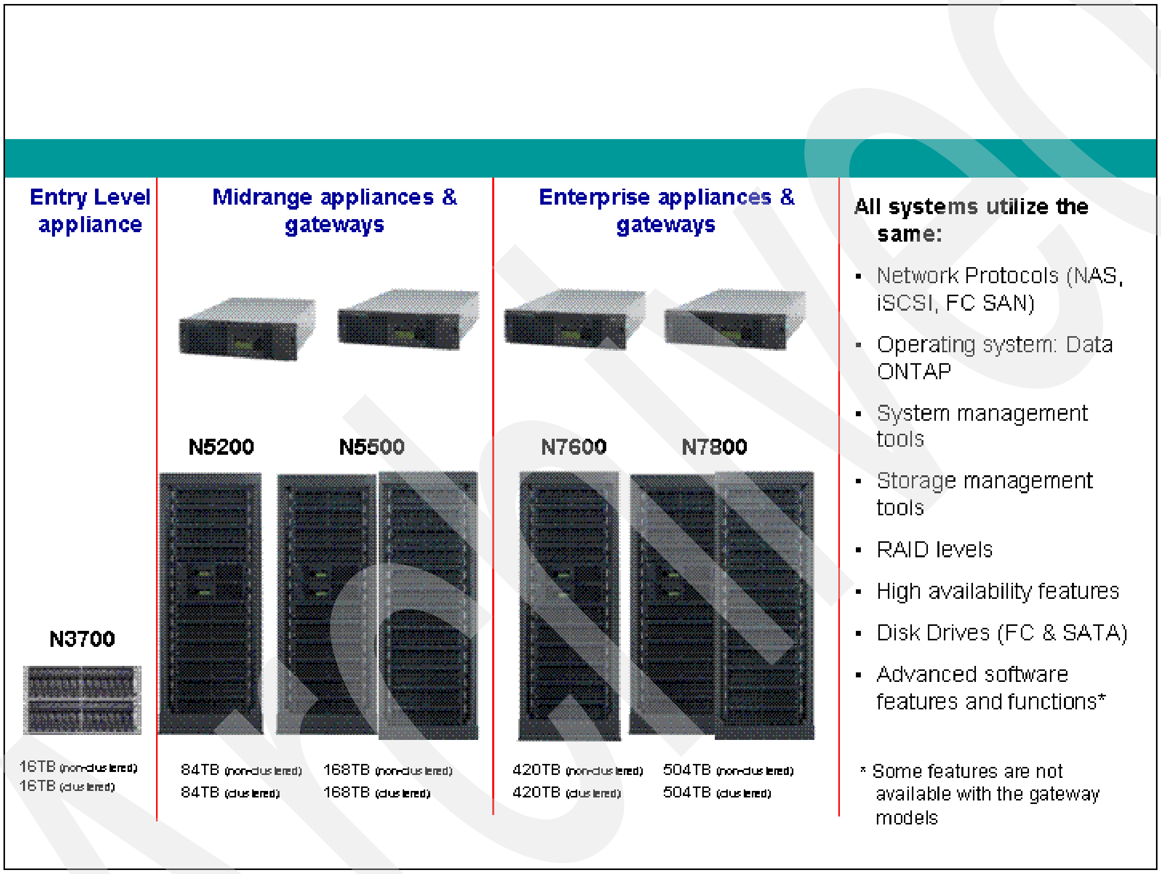

All systems utilize the same:

- Network Protocols (NAS, iSCSI, FC SAN)
- Operating system: Data ONTAP
- System management tools
- Storage management tools
- RAID levels
- High availability features
- Disk Drives (FC & SATA)
- Advanced software features and functions*

* Some features are not available with the gateway models

Figure B-15 Completion of N series Portfolio

Advanced Copy Services for N series

The core data replication technologies available on the N series are:

- SnapShot
- SyncMirror
- SnapMirror
- Metro Cluster

SnapShot

A SnapShot copy is a locally retained point-in-time image of data. SnapShot technology is a feature of the Write Anywhere File Layout (WAFL®) storage virtualization technology that is a part of Data ONTAP. A SnapShot is a “frozen,” read-only view of a WAFL volume that provides easy access to old versions of files, directory hierarchies, or LUNs.

A SnapShot (Figure B-16) can take only a few seconds to create, regardless of the size of the volume or the level of activity on the N series. After a SnapShot copy has been created, changes to data objects are reflected in updates to the current version of the objects, as though SnapShot copies did not exist. Meanwhile, the SnapShot version of the data remains completely unchanged. A SnapShot copy incurs little performance overhead; depending on available space, users can store up to 255 SnapShot copies per WAFL volume, all of which are accessible as read-only, online versions of the data.

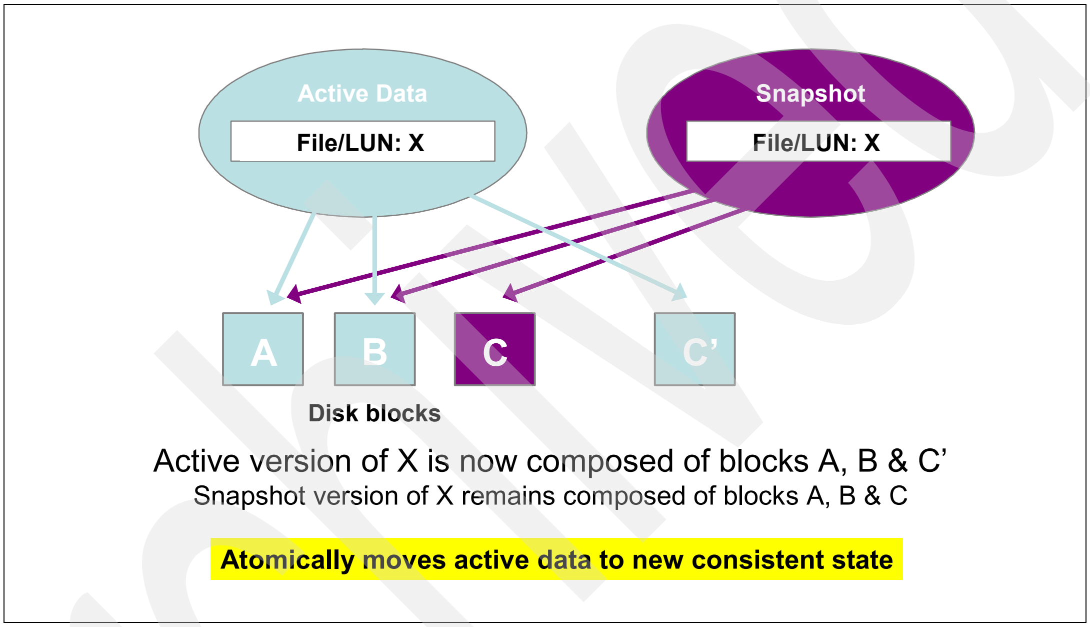

Active version of X is now composed of blocks A, B & C’

Snapshot version of X remains composed of blocks A, B & C

Atomically moves active data to new consistent state

Figure B-16 N series SnapShot - How it works

A SnapShot copy can be used to provide frequent, low-impact, user-recoverable backups of files, directory hierarchies, LUNs, or application data. A SnapShot copy can significantly improve the frequency and reliability of backups, since it is designed to avoid performance overhead and can be created on a running system.

A SnapShot supports near-instantaneous, user-managed restores. Users can directly access SnapShot copies to recover from accidental deletions, corruptions, or modifications of their data.

SyncMirror

The SyncMirror functionality, shown in Figure B-17, provides synchronous local mirroring from one volume to another volume attached to the same filer. It maintains a strict physical separation between the two copies of the mirrored data. In case of an error in one copy, the data is still accessible without any manual intervention.

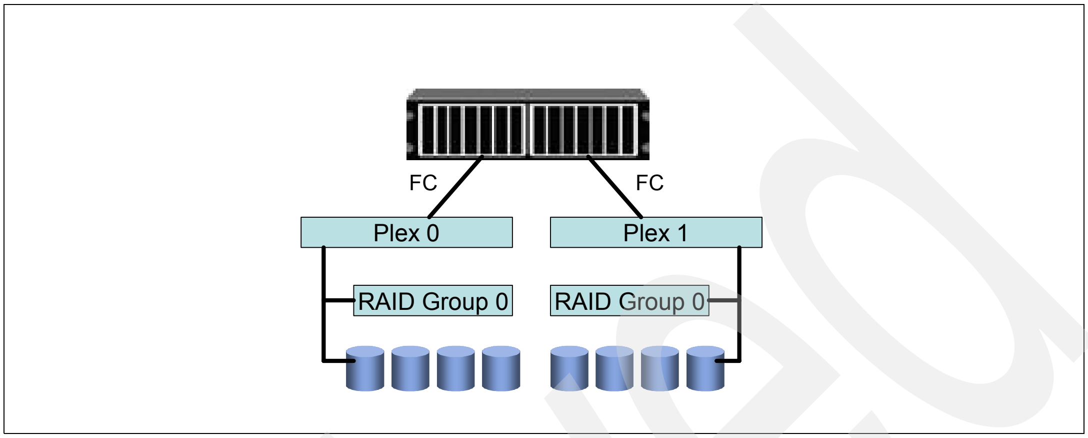

Figure B-17   N series SyncMirror

With SyncMirror, filers can tolerate multiple simultaneous disk failures across the RAID groups within the WAFL file system. This redundancy goes beyond typical mirrored (RAID-1) implementations. Because each SyncMirror RAID group is also RAID-4 or RAID-DP protected, a complete mirror could be lost and an additional single drive loss within each RAID group could occur without data loss.

SnapMirror

SnapMirror, shown in Figure B-18 on page 139, replicates data sets between N series over a network for backup or disaster recovery. After an initial baseline transfer of the entire data set, subsequent updates only transfer new and changed data blocks from the source to the destination, which makes SnapMirror highly efficient in terms of network bandwidth utilization. The destination file system is available for read-only access, or the mirror can be “broken” to enable writes to occur on the destination. After breaking the mirror, it can be reestablished by synchronizing the changes made to the destination back onto the source file system. In the traditional asynchronous mode of operation, updates of new and changed data from the source to the destination occur on a schedule defined by the storage administrator. These updates could be as frequent as once per minute or as infrequent as once per week, depending on user needs. Synchronous mode is also available, which sends updates from the source to the destination as they occur, rather than on a schedule. This can guarantee that data written on the source system is protected on the destination even if the entire source system fails due to natural or human-caused disaster. A semi-synchronous mode is also provided, which minimizes data loss in a disaster while also minimizing the performance impact of replication on the source system.

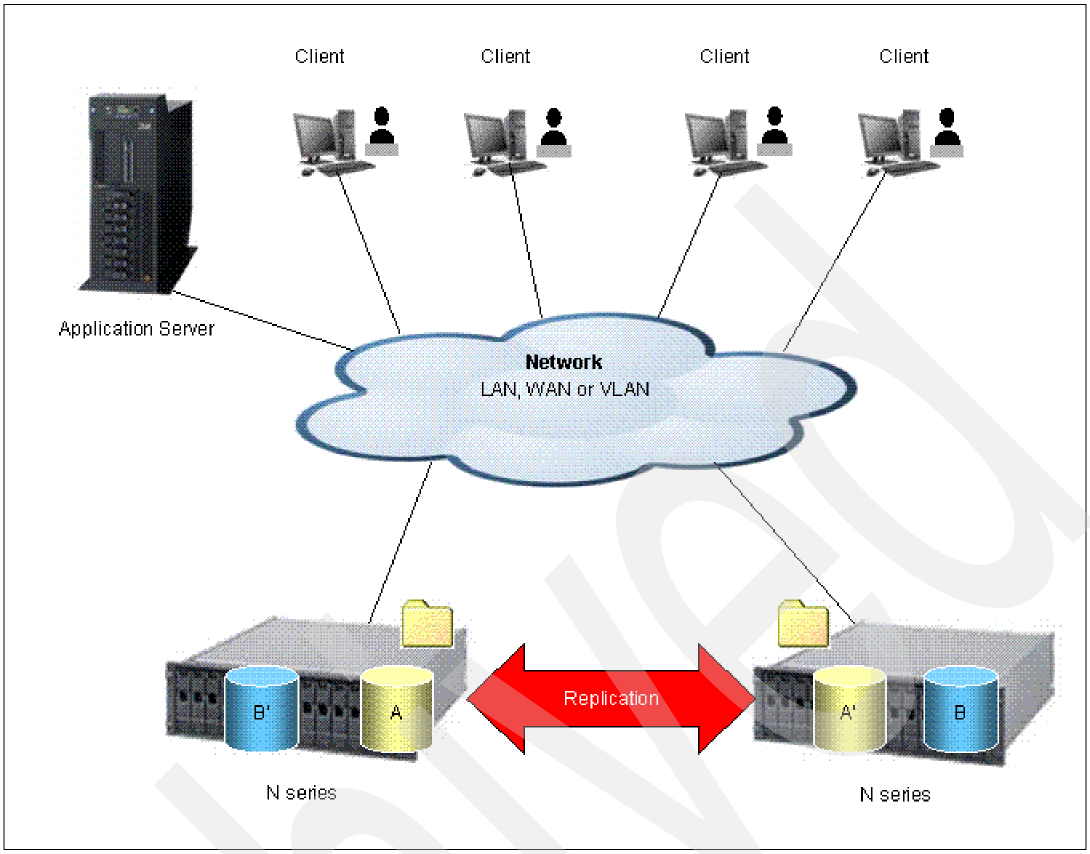

Figure B-18 IBM N series SnapMirror

Metro Cluster

MetroCluster feature is an integrated, high availability and business-continuance and allows clustering of two N5000 or N7000 storage controllers at distances up to 100 kilometers.

With MetroCluster, the active/active configuration can be spread across data centers up to 100 kilometers apart. In the event of an outage at one data center, the second data center can assume all affected storage operations. SyncMirror is required as part of MetroCluster to ensure an identical copy of the data exists in the second data center should the original data center be lost.

1. MetroCluster along with SyncMirror extends active/active Clustering across data centers up to 100 kilometers apart (Figure B-19).
2. MetroCluster and SyncMirror provide the highest level of storage resiliency across a local region.
3. Highest levels of regional storage resiliency ensure continuous data availability in a particular geography.

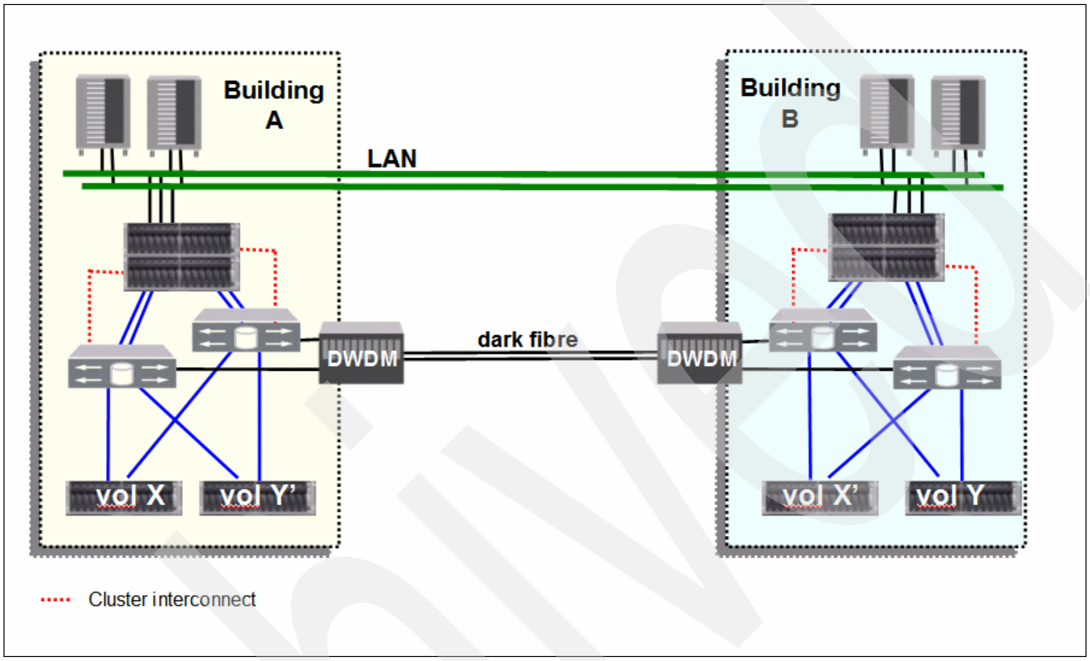

Figure B-19 MetroCluster

IBM System Storage SAN Volume Controller

In this section, we introduce some of the storage components of the IBM System Storage SAN Volume Controller.

The IBM System Storage SAN Volume Controller (SVC) provides block aggregation and volume management for disk storage within the SAN. The SVC manages a number of heterogeneous back-end storage controllers and maps the physical storage within those controllers to logical disk images (known as virtual disks, or *VDisks*), that can be seen by application servers in the SAN. The SAN is zoned in such a way that the application servers cannot see the back-end storage, preventing any possible conflict between SVC and the application servers both trying to manage the back-end storage.

SVC supports heterogeneous server and storage environments, including:

- IBM AIX, Linux, UNIX, Microsoft Windows, HP-UX, and SUN Solaris.
- IBM DS4000/6000/8000/ESS/N series, EMC DMX/CLARiiON, HP XP/EVA/MA/EMA, and Hitachi TagmaStore/Thunder/Lightning

For the complete list of supported servers and storage, see:

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

Advanced Copy Services for SVC

Advanced copy services for SVC include:

- FlashCopy
- Metro Mirror
- Global Mirror

FlashCopy

FlashCopy (see Figure B-20) is a point-in-time copy of a virtual disk on an SVC.

FlashCopy works by defining a FlashCopy mapping consisting of one source VDisk together with one target VDisk. Multiple FlashCopy mappings can be defined and PiT consistency can be observed across multiple FlashCopy mappings using consistency groups.

When FlashCopy is started, it makes a copy of a source VDisk to a target VDisk, and the original contents of the target VDisk are overwritten. When the FlashCopy operation is started, the target VDisk presents the contents of the source VDisk as they existed at the single point in time (PiT) the FlashCopy was started.

When a FlashCopy is started, the source and target VDisks are instantaneously available. This is because when started, bitmaps are created to govern and redirect I/O to the source or target VDisk, respectively, depending on where the requested block is present, while the blocks are copied in the background from the source to the target VDisk.

Both the source and target VDisks are available for read and write operations, although the background copy process has not yet completed copying across the data from the source to target volumes.

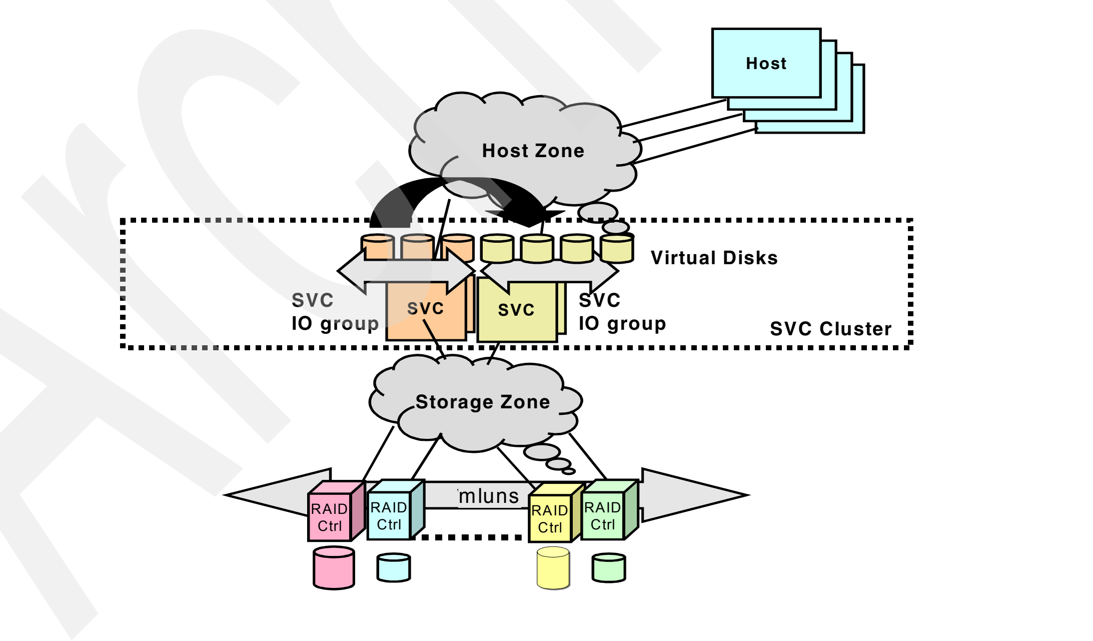

Figure B-20 IBM SAN Volume Controller - FlashCopy

Metro Mirror

Metro Mirror is a fully synchronous remote copy technique that ensures that updates are committed at both primary and secondary VDisks before the application is given completion to an update.

Figure B-21 shows how a write to the master VDisk is mirrored to the auxiliary disk's cache before a write acknowledgement is sent back to the host issuing the write. This ensures that the secondary is real-time synchronized, in case it is needed in a failover situation.

However, this also means that the application is fully exposed to the latency and bandwidth limitations of the communication link to the secondary site. This might lead to unacceptable application performance, particularly when placed under peak load. This is the reason for the distance limitations when applying Metro Mirror.

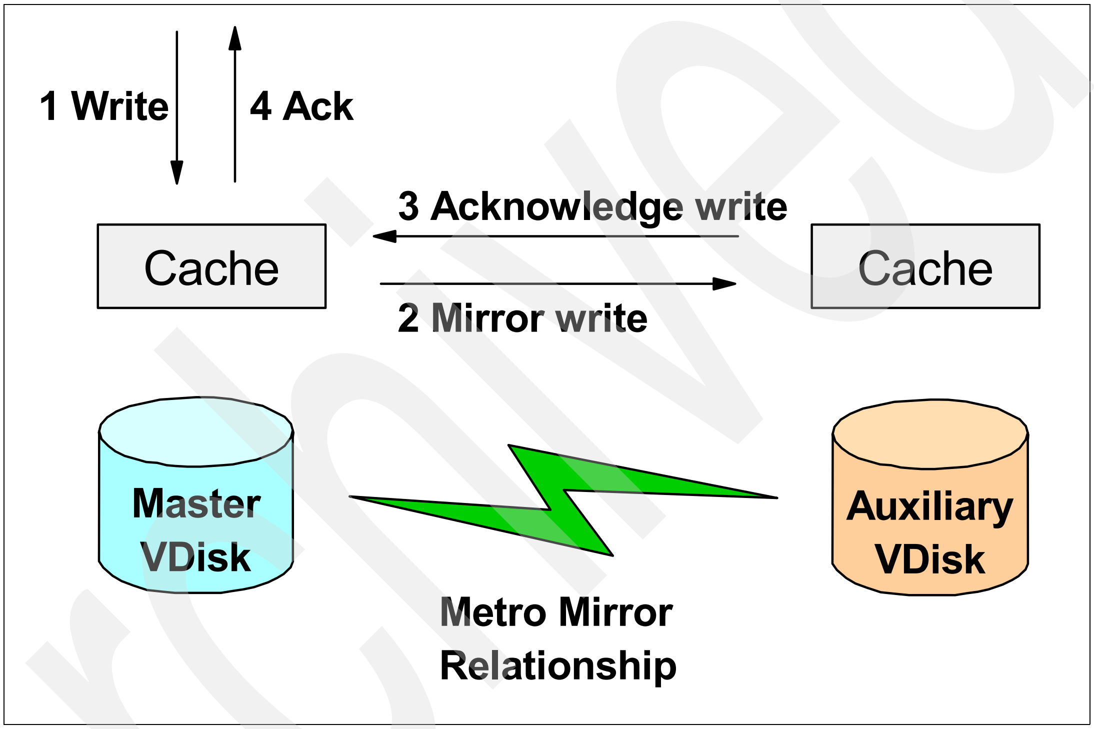

Figure B-21 Write on VDisk in Metro Mirror relationship

Global Mirror

Global Mirror works by defining a relationship between two equally sized VDisks that maintains data consistency in an asynchronous manner. Therefore, when a host writes to a source VDisk, the data is copied from the source VDisk cache to the target VDisk cache. At the initiation of that data copy, confirmation of I/O completion is transmitted back to the host.

SVC supports Global Mirror (consistent asynchronous secondary copy of data) at nearly unlimited distances with minimal performance impact on production servers. Figure B-22 provides an overview of how this works.

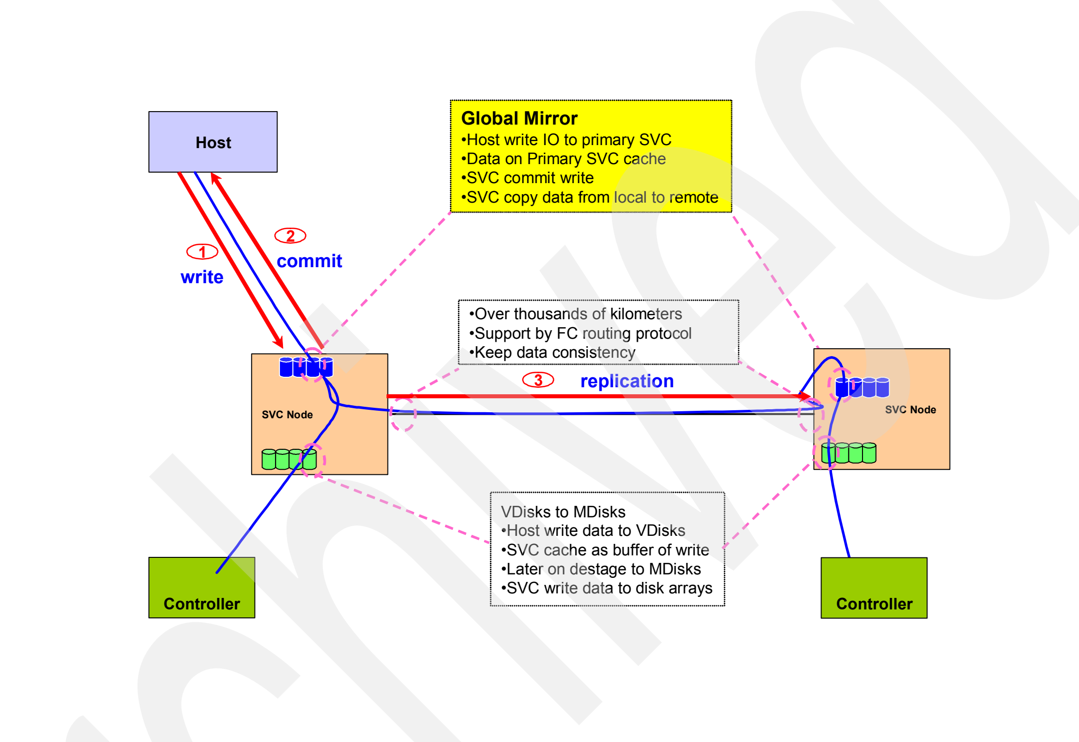

Global Mirror

- Host write IO to primary SVC
- Data on Primary SVC cache
- SVC commit write
- SVC copy data from local to remote

- Over thousands of kilometers
- Support by FC routing protocol
- Keep data consistency

VDisks to MDisks

- Host write data to VDisks
- SVC cache as buffer of write
- Later on destage to MDisks
- SVC write data to disk arrays

Figure B-22 Business continuity with SVC Global Mirror

IBM System Storage Virtualization Engine TS7500 and TS7700

IBM virtualized tape products are:

- IBM Virtualization Engine TS7510 (open systems Virtual Tape Library support)
- IBM Virtualization Engine TS7740 (mainframe Virtual Tape Library support)

IBM Virtualization Engine TS7510

Figure B-23 shows the TS7510. The included application software provides tape library and tape drive emulation including virtual volumes. The TS7510 provides tape automation and cross site support, which enables a BC Tier 5 implementation.

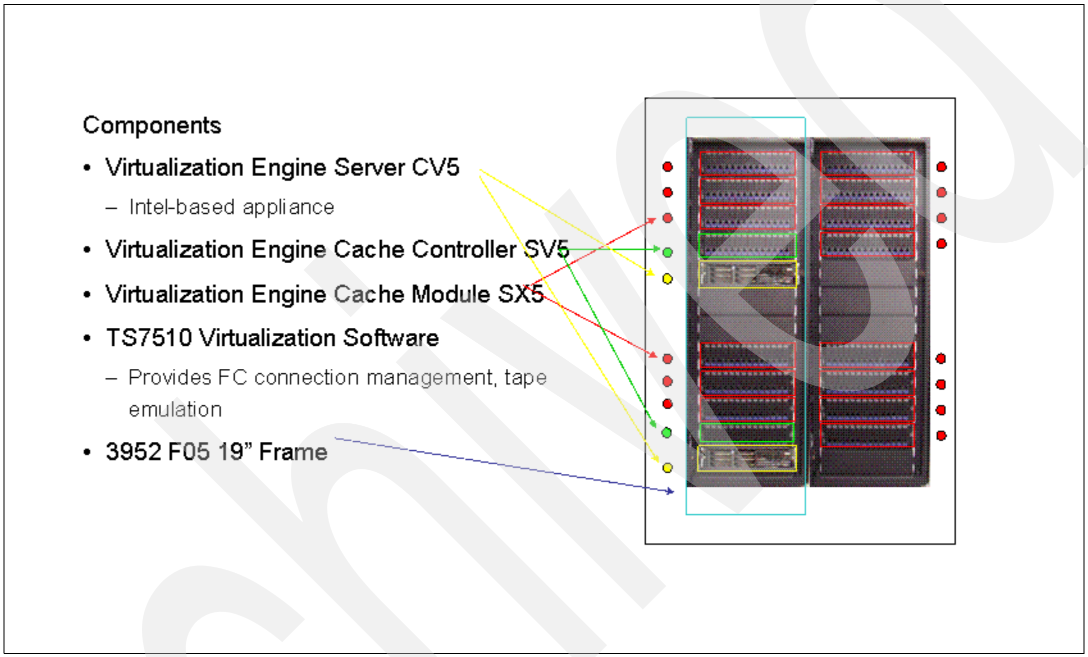

Components

- Virtualization Engine Server CV5
  - Intel-based appliance
- Virtualization Engine Cache Controller SV5
- Virtualization Engine Cache Module SX5
- TS7510 Virtualization Software
  - Provides FC connection management, tape emulation
- 3952 F05 19” Frame

Figure B-23 IBM Virtualization Engine TS7510 - Virtual Tape Library Solution Components

The following describes the TS7510 cross site tape support.

Network Replication

Network Replication provides a method to recover from complete data loss by sending copies of data off site. There are three methods of Network Replication: Remote Copy, Replication, and Auto Replication. To use the Network Replication function, you need two IBM Virtualization Engine TS7510s:

- The primary TS7510 that serves virtual tape drives to your backup servers
- A disaster recovery/remote Virtualization Engine TS7510

Remote Copy

Remote Copy is a manually triggered, one-time replication of a local virtual tape. After the Remote Copy is complete, the tape resides on both the primary TS7510 and in the remote TS7510’s Replication Vault.

You cannot perform Remote Copy on a tape that has already been configured for Auto Replication. When using Remote Copy, the copied tape can reside either in one of the virtual tape libraries, in a virtual tape drive, or in the virtual vault. The Remote Copy option preserves the barcode from the Virtualization Engine that the remote copy initiated.

Figure B-24 illustrates the Remote Copy movement. The primary Virtualization Engine is on the left, and the remote backup is on the right.

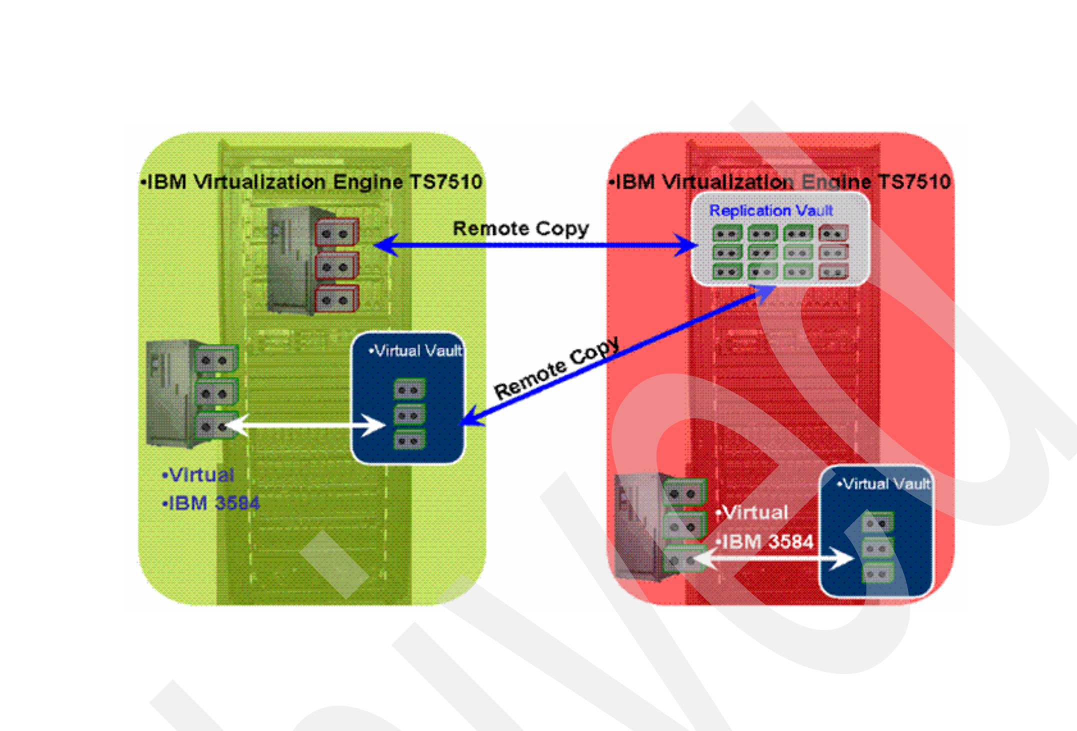

Figure B-24 Remote Copy Data Movement

Replication

The Replication process is either triggered by a scheduled event or when the virtual volume reaches a certain predetermined size. When Replication is configured, a primary virtual volume is created and linked to the virtual replica on the remote Virtualization Engine. A replica tape is always linked to the original virtual tape. It cannot be used by any virtual library or for import/export by the remote Virtualization Engine until this linked relationship is broken. This condition is also known as *promoting a replica*. Its only purpose is to maintain an in-sync copy of the primary tape.

The replica tape simply gets incremental changes from the source tape, ensuring the two tapes are always in-sync at the end of a replication session. This is why it is a *dedicated relationship*.

Data traveling across the replication path can be compressed, encrypted, or both. Additional license codes are required to activate these features, which we explain later. If the replica is promoted, it is placed in the virtual vault on the remote Virtualization Engine, with the same barcode label as the source virtual tape. It can then be used like any other virtual tape.

Figure B-25 illustrates replication movement. The left TS7510 is the primary engine, and the right TS7510 is the backup. Data replicates from the primary to the backup utilizing the replication process. When the primary engine fails in order to use a replica that is on the backup engine, the virtual replica sitting in the replication vault is promoted to a virtual volume and moved to the virtual vault. It is either placed in a virtual library on the backup or copied back to the primary.

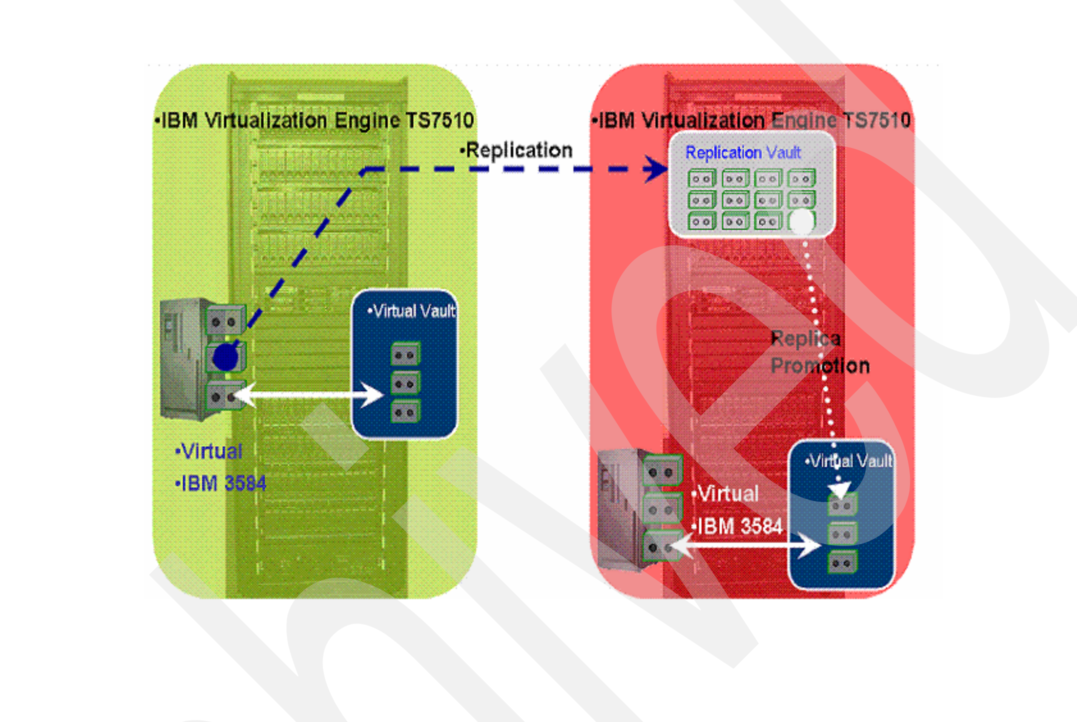

Figure B-25 Replication Data Movement

Auto Replication

As we mentioned previously, Auto Archive involves a one-time copy or move of the virtual tape as soon as the backup software has sent an **eject** command. Auto Replication provides for the same, one-time copy or move after the eject, but the destination is to the remote Virtualization Engine instead of to a local physical 3584 or 3494 tape library.

Figure B-26 on page 147 shows the Auto Replication process. The left side shows the primary engine, and the right side shows the backup engine. The primary initiates the Auto Replication function. Also, a one-time copy or move after the **eject** command is sent to the backup Virtualization Engine. The virtual volume is then placed in the replication vault.

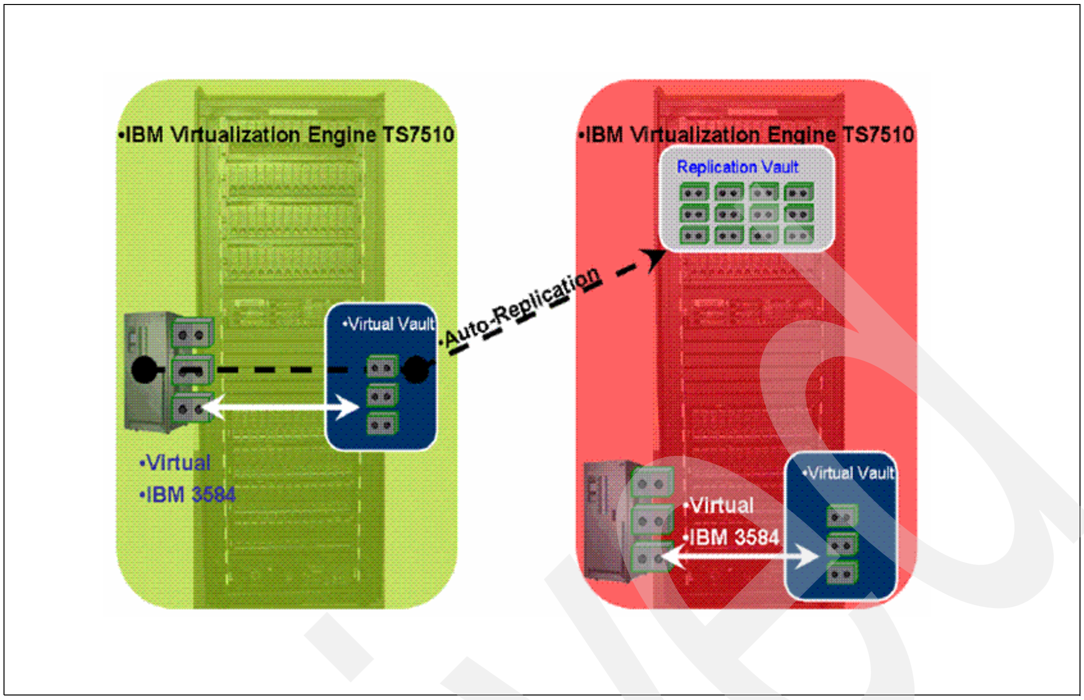

Figure B-26 Auto Replication Data Movement

IBM Virtualization Engine TS7740

The TS7740 can form a TS7740 Grid, which enables a seamless VTL integration. This is the most robust virtual tape implementation that can support BC Tier 7 or Tier 5.

High Availability Configuration with TS7740 Grid

Figure B-27 shows that the two TS7740 Tape Library cluster can be configured as high availability tape library in a single tape library image.

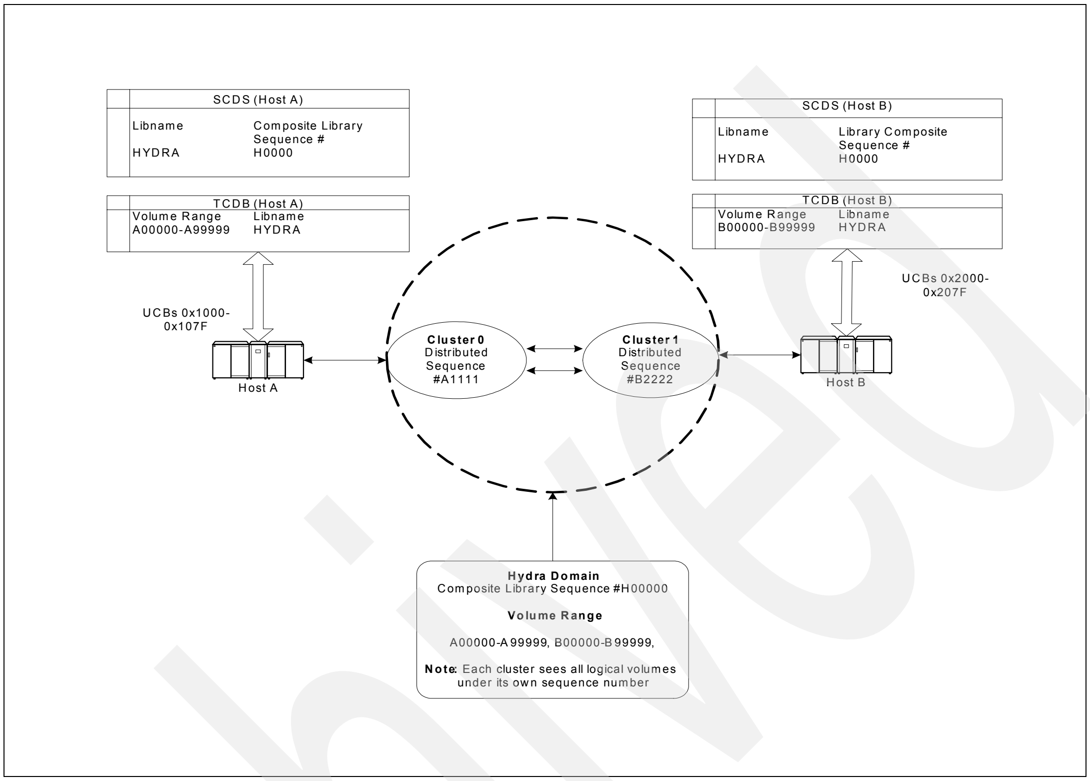

SCDS (Host A)	
Libname	Composite Library Sequence #
HYDRA	H0000

TCDB (Host A)	
Volume Range	Libname
A00000-A99999	HYDRA

SCDS (Host B)	
Libname	Library Composite Sequence #
HYDRA	H0000

TCDB (Host B)	
Volume Range	Libname
B00000-B99999	HYDRA

Hydra Domain
Composite Library Sequence #H00000

Volume Range

A00000-A99999, B00000-B99999,

Note: Each cluster sees all logical volumes under its own sequence number

Figure B-27 Example of TS7740 Dual Cluster Grid in a High Availability Configuration

Figure B-28 on page 149 shows the interconnect of TS7740 Grid, where two of the TS7740s are inter-connected to LAN/WAN.

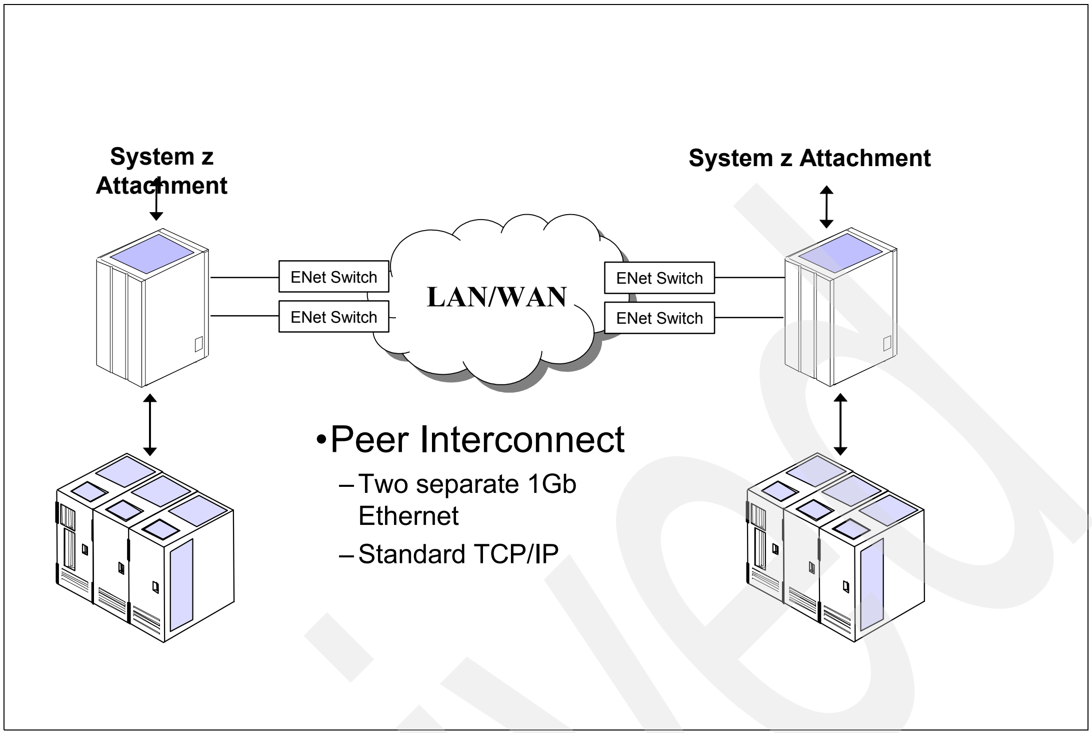

Figure B-28 TS7740 Dual Cluster Grid Configuration

GDPS / PPRC with TS7700 Grid Configuration

Figure B-28 shows how TS7740 Grid interconnects across two remote sites. This can be deployed to provide BC Tier 7 GDPS/PPRC support with GDPS/PPRC (Figure B-29). Automated cross-site recovery can be deployed together with a TS7740 Grid for more robust automation continuous operation.

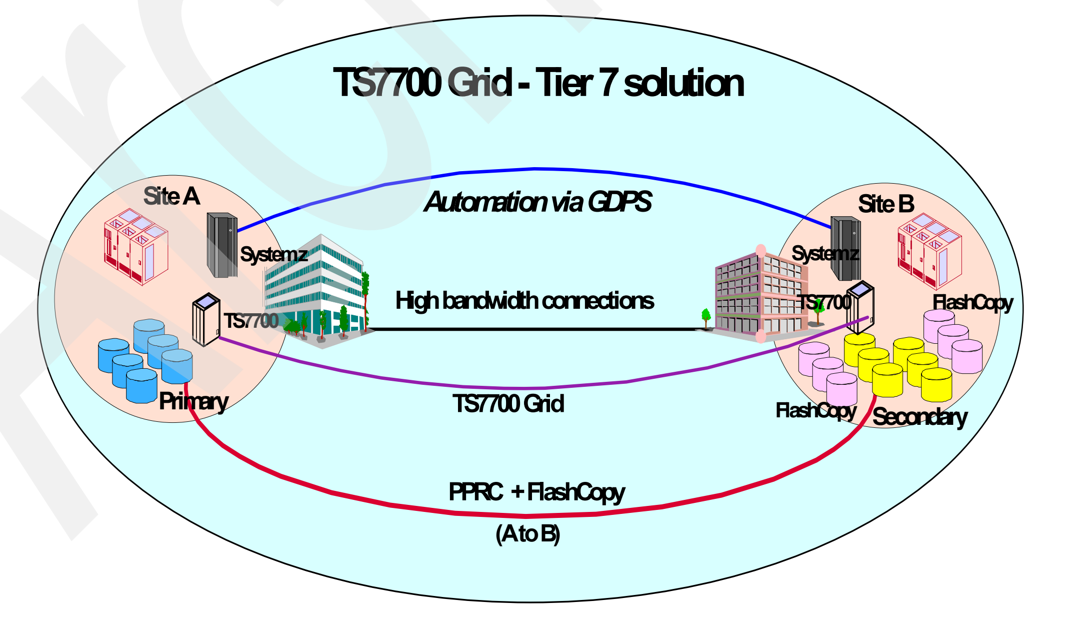

Figure B-29 A TS7740 grid - Tier 7 solution with GDPS / PPRC in a System z Environment

C

Tivoli Storage Manager family overview

This chapter provides a product overview of the Tivoli Storage Manager family of products that may be used with the solutions in this IBM Redbook. The products include:

- IBM Tivoli Storage Manager
- IBM Tivoli Storage Manager for Databases
- IBM Tivoli Storage Manager for Mail
- IBM Tivoli Storage Manager for ERP
- IBM Tivoli Storage Manager for Copy Services
- IBM Tivoli Storage Manager for Advanced Copy Services
- IBM Tivoli Storage Manager Disaster Recovery Manager
- IBM Tivoli Storage Manager for System Backup and Recovery

IBM Tivoli Storage Manager

IBM Tivoli Storage Manager enables you to protect your organization's data from failures and other errors by storing backup, archive, space management, and bare-metal restore data, as well as compliance and Disaster Recovery data in a hierarchy of offline storage. Because it is highly scalable, Tivoli Storage Manager can help protect computers running a variety of different operating systems, on hardware ranging from notebooks to mainframe computers and connected together through the Internet, wide area network (WAN), local area network (LAN), or storage area network (SAN).

It uses Web-based management, intelligent data move-and-store techniques, and comprehensive policy-based automation, which work together to help increase data protection and potentially decrease time and administration costs. Because it is highly scalable, Tivoli Storage Manager can also help protect computers running a variety of different operating systems. A sample window of Tivoli Storage Manager Backup and Archive Graphic User Interface Client is shown in Figure C-1.

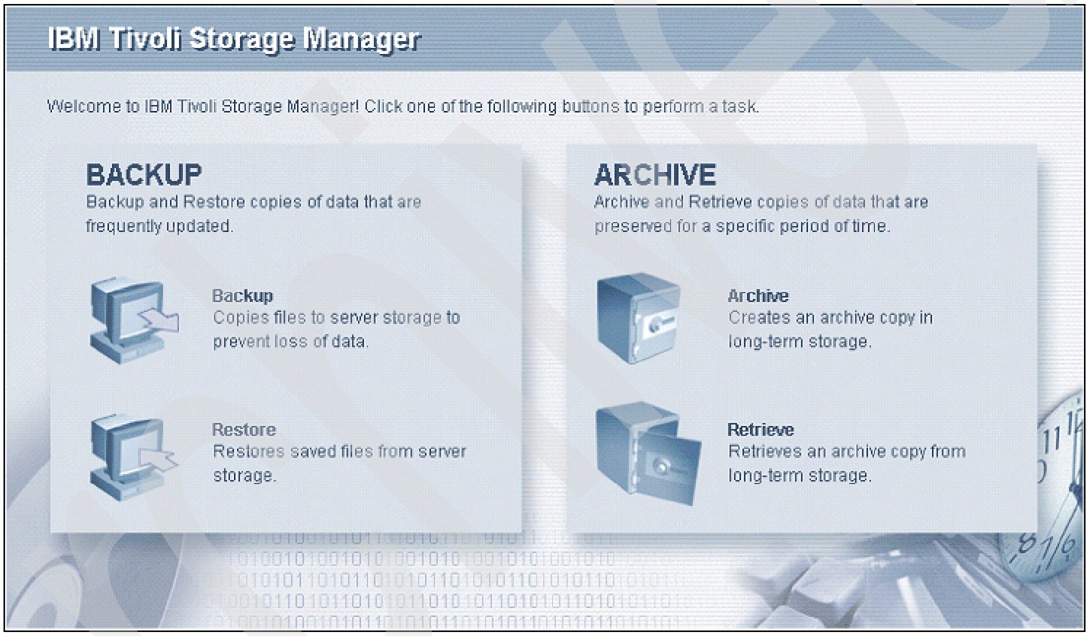

IBM Tivoli Storage Manager

Welcome to IBM Tivoli Storage Manager! Click one of the following buttons to perform a task.

BACKUP

Backup and Restore copies of data that are frequently updated.

Backup

Copies files to server storage to prevent loss of data.

Restore

Restores saved files from server storage.

ARCHIVE

Archive and Retrieve copies of data that are preserved for a specific period of time.

Archive

Creates an archive copy in long-term storage.

Retrieve

Retrieves an archive copy from long-term storage.

Figure C-1 IBM Tivoli Storage Manager Backup Archive Client: GUI client

IBM Tivoli Storage Manager's core functions include:

- Backup and recovery management
- Archive management

IBM Tivoli Storage Manager Extended Edition adds additional support through:

- Disaster preparation planning and recovery (Disaster Recovery Manager)
- NDMP backup for Network Attached Storage
- Small and large tape libraries

Attributes that set Tivoli Storage Manager apart include:

- Easy management of multiple types of inactive data in a hierarchical repository
- Lower storage cost through intelligent hierarchy of storage
- Centralized, comprehensive management
- Reduced network bandwidth through intelligent data movement
- Policy-based automation

IBM Tivoli Storage Manager family of offerings include:

- IBM Tivoli Storage Manager for Application Servers
- IBM Tivoli Storage Manager for Databases
- IBM Tivoli Storage Manager for Enterprise Resource Planning
- IBM Tivoli Storage Manager for Copy Services
- IBM Tivoli Storage Manager for Advanced Copy Services
- IBM Tivoli Storage Manager for Mail
- IBM Tivoli Storage Manager for System Backup and Recovery
- IBM Tivoli Storage Manager for Data Retention

Optional Additions to Tivoli Storage Manager include:

- IBM Tivoli Storage Manager for Storage Area Networks
- IBM Tivoli Storage Manager for Space Management

For more product information, see the following Web site:

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

Disaster Recovery Manager (DRM)

DRM is a component of IBM Tivoli Storage Manager Extended Edition. It provides detailed tracking of the additional copies of your backed-up, archived, and space managed data that IBM Tivoli Storage Manager creates for safekeeping at an off site location. IBM Tivoli Storage Manager DRM also prepares and keeps up-to-date a text file with detailed recovery steps and automated computer-scripts to form the *recovery plan*. Should a disaster strike and destroy your storage and computers, this plan and the off-site data copies will get your business back up and running quickly.

The core functions include:

- Automated generation of a customized server disaster recovery plan.
- Detailed off-site recovery media management.
- An inventory of machine information is required to recover the server and its clients.
- Centralized management of the disaster recovery process.
- Executable scripts that assist in recovery automation.
- Electronic vaulting of storage pool and database backups.

IBM Tivoli Storage Manager for Databases

IBM Tivoli Storage Manager for Databases is a software module that works with IBM Tivoli Storage Manager to protect a wide range of application data via the protection of the database's management system. Tivoli Storage Manager for Databases exploits the backup-certified utilities and interfaces provided for Oracle, and Microsoft SQL Server. In conjunction with Tivoli Storage Manager, this module automates data protection tasks and allows database servers to continue running their primary applications while they back up and restore data to and from offline storage.

> **Note:** This same functionality is included in the IBM DB2 Universal Database™ and Informix® (latest release) package, allowing it to work directly with Tivoli Storage Manager without the need to buy any additional modules.

Regardless of which brand of database is used, Tivoli Storage Manager for Databases allows the centralized and automated data protection capabilities of Tivoli Storage Manager to be applied to up and running database servers.

For more product information, see the following Web site:

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

IBM Tivoli Storage Manager for Mail

IBM Tivoli Storage Manager for Mail is a software module for IBM Tivoli Storage Manager that automates the data protection of e-mail servers running either Lotus Domino or Microsoft Exchange. This module utilizes the application program interfaces (APIs) provided by e-mail application vendors to perform online *hot* backups without shutting down the e-mail server and improves data-restore performance. As a result, it can help protect the growing amount of new and changing data that should be securely backed up to help maintain 24x7x365 application availability.

For more product information, see the following Web site:

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

IBM Tivoli Storage Manager for Enterprise Resource Planning

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

Specifically designed and optimized for the SAP R/3 environment, IBM Tivoli Storage Manager for ERP provides automated data protection, reduces the CPU performance impact of data backups and restores on the R/3 server, and greatly reduces the administrator workload necessary to meet data protection requirements. Tivoli Storage Manager for ERP builds on the SAP database, a set of database administration functions integrated with R/3 for database control and administration. The Tivoli Storage Manager for ERP software module allows multiple R/3 servers to utilize a single Tivoli Storage Manager server to automatically manage the backup of R/3 data. As the intelligent interface to the R/3 database, Tivoli Storage Manager for ERP is SAP certified in heterogeneous environments, supporting large-volume data backups, data recovery, data cloning, and Disaster Recovery of multiple SAP R/3 servers.

For more product information, see the following Web site:

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

IBM Tivoli Storage Manager for Copy Services

IBM Tivoli Storage Manager for Copy Services helps protect business critical Microsoft Exchange databases that require 24x7 availability. It offers options to implement high-efficiency backup of business-critical applications while virtually eliminating backup-related impact on the performance of the MS Exchange production server. This is done by integrating the snapshot technologies of the storage system with Tivoli Storage Manager's database protection capabilities for Microsoft Exchange to support a “near zero-impact” backup process. The product also works with IBM SAN Volume Controller in a virtualized storage environment and provides a unique feature leveraging SVC's FlashCopy function to allow for Instant Restores.

Tivoli Storage Manager for Copy Services takes advantage of a supported storage system's built-in snapshot capabilities, together with the Volume Shadowcopy Service provided in

Microsoft Windows 2003, to provide fast online backups of Exchange databases. The snapshot backups can be retained on disk, or optionally copied to Tivoli Storage Manager for longer term storage. Because the snapshot operation itself is rapid, the impact of backup on the database is minimal. To further reduce the overhead of backup, the copy operation to Tivoli Storage Manager can optionally be performed by a separate server known as an offloaded backup server.

IBM Tivoli Storage Manager for Advanced Copy Services

Tivoli Storage Manager for Advanced Copy Services software provides online backup and restore of data stored in mySAP, DB2, and Oracle applications by leveraging the copy services functionality of the underlying storage hardware. Using hardware-based copy mechanisms rather than traditional file-based backups can significantly reduce the backup/restore window on the production server. Backups are performed through an additional server called the *backup server*, which performs the actual backup. Since 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 application users' performance is not affected during backup.

Specifically, Tivoli Storage Manager for Advanced Copy Services provides FlashCopy integration on the DS6000, DS8000, SVC, and ESS for split mirror backup and optional FlashBack restore of mySAP, DB2 UDB, and Oracle databases.

IBM Tivoli Storage Manager for System Backup and Recovery

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

For more product information, see the following Web site:

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

D

Services and planning

In this appendix, we discuss the following service related topics:

- What services are required and who can provide them
- IBM Global Services services families and solution life cycle
- On demand services
- IBM Global Services Business Resiliency services
- Network Consulting and Integration services

Services and services providers

Most of the solutions illustrated in this IBM Redbook rely on interactions between multiple components that can come from multiple vendors. The components may span:

- Hardware infrastructures, such as disk and tape devices
- Hardware and software enabling technologies, such as DS8000 Copy Services
- Automation components, such as AIX HACMP or Tivoli System Automation
- Software automation, such as DB2 support for FlashCopy backups

Integrating the components so they operate in a coherent manner and perform a desired function or task is one of the reasons for acquiring external services. Implementing complex solutions alone can be a time consuming and labor intensive endeavour. Because of this, services are often the *glue* that tie together the components or pieces of a solution. The IBM Global Services organization has traditionally delivered services to clients, and integration services targeted at both IBM and OEM products. The scope of these services spans from simple product installation and implementation services to IT and business transformation consulting engagements, and to out-tasking of the management of IT infrastructures.

In recent years, IBM has also started to rely on IBM Business Partners to deliver services. Many IBM Business Partners exist around the world; they offer a wide variety of services that can be beneficial in building the desired solution. The major difference between IBM Global Services and services from IBM Business Partners is size and scope. IBM Global Services has a worldwide organization and presence while IBM Business Partners tend to be more local or regional. IBM Global Services in general has a more comprehensive offering in terms of supported platforms and technologies, and delivers services that range from business process consulting to solution component implementation and support activities.

When do you choose IBM Global Services or an IBM Business Partner? This question is not easily answered. It depends on the size and complexity of the solution being built. When a client contacts IBM directly for a solution, IBM will evaluate the complexity and effort and give a recommendation to the client as whether to involve IBM Global Services or an IBM Business Partner. In the following section, we will give an overview of IBM Global Services services.

IBM Global Services services families and life cycle

IBM Global Services addresses a complex engagement or service in a series of individual steps, defined in the IBM Global Services Method. The IBM Global Services Method provides a single method to enable a common language among all practitioners delivering business solutions. IBM Global Services Method is the work product based method for IBM Global Services Practitioners. At the base of the Method is a core set of Work Product Descriptions (WPDs) that can be shared by all practitioners using the Method. Work Products are tangible, reusable artifacts produced as a result of one or more tasks performed on an engagement.

The Method also provides guidance for how engagements should be conducted. This guidance is delivered through Engagement Models that represent many of the typical projects conducted in IBM Global Services. Each Engagement Model provides guidance for the phases, activities, and tasks required (often called the work breakdown structure (WBS)), the Work Products that are produced, the roles required to perform the work, and any applicable techniques that should be used for one or more of the tasks.

In storage and resiliency services, the IBM Global Services Method consulting approach is divided into five steps plus the engagement management step, as shown in Figure D-1 on page 159.

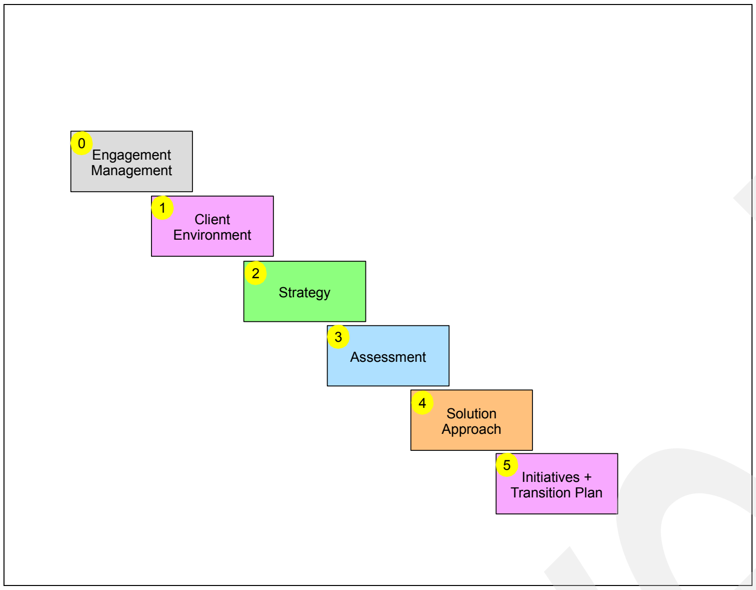

Figure D-1 IBM Global Services Method for resiliency and storage related services

The engagement management step, or step 0, is where the scope of the services is defined. What will be the deliverables and what is the project plan for performing the services? After this step, this methodology defines the following steps. Each step uses outputs (Work Products) of the previous steps and answers a set of questions:

1. What are the business and IT issues facing the client or opportunities the company wants to pursue? What are the IT constraints that need to be considered?
2. What is the target environment and the storage solution(s) that will enable the business to achieve its objectives?
3. What are the areas that the business needs to focus on to move towards the target environment from its current environment? In what order of priority must these areas be addressed?
4. What are the recommended approaches or paths that the business needs to follow to implement the strategy in the most cost-effective manner?
5. What actions do you need to take to establish the target environment that is required to achieve the stated objectives? What is required to implement them in terms of projects, resources, and schedules? What value will they bring in terms of value propositions and a sustainable business case?

To put it in simpler terms, step 1 defines where we are, step 2 states where we want to go, step 3 defines what needs to be changed, and steps 4 and 5 define the path to reach the goal of where we want to go.

The approach illustrated above applies to a more comprehensive consulting type of engagement. Another way of defining and classifying services is illustrated in Figure D-2. This is a more solution life cycle oriented classification.

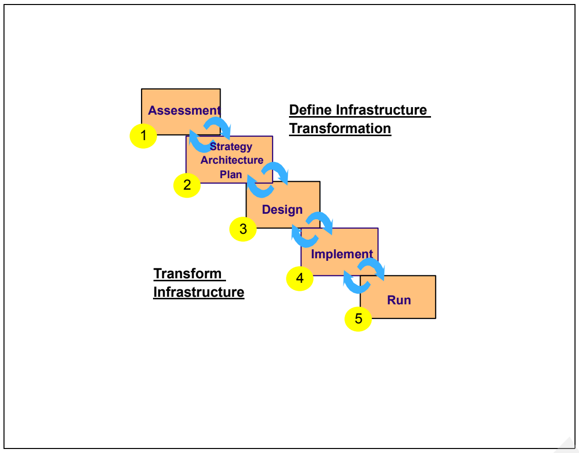

Figure D-2 IBM Global Services classification

Steps 1 and 2 are the steps where you decide what business needs you have to address. Once these needs are understood and documented, we can proceed to defining a target storage environment architecture. This phase may also include a cost and benefit analysis. It certainly will include a transition plan. The transition plan should state and document the solution and solution component providers, the projects and approach required to implement the solution, the services providers, and an evaluation should be made of whether to out-task.

The next steps are about performing the transformation itself:

- Step 3 - Design: Defines the detailed physical and logical design of different components of the architecture. Step 3 may also include testing to evaluate and prototype solutions (for example, SAN) for reliability, performance, and security.
- Step 4 - Implementation: Addresses the proof of concept and deployment of recommended solutions and solution components, such as DS8000 Copy Services or script development for automation and so on.
- Step 5 - Run and support: Addresses storage management activities. IBM Global Services Managed Storage Services offer clients the opportunity of out-tasking storage management related activities to IBM Global Services.

You may decide to rely on services for any or all of the steps we have illustrated. If you have already chosen the solution components and do not require IBM Global Services to manage the solution you are building, you may well choose only steps 2 and 3, design and implement.

Using this five step approach, we can attempt a first classification of the vast array of services offered by IBM Global Services.

On Demand services

IBM and IBM Global Services have recently realigned major initiatives to solve specific client problems or sets of problems. Figure D-3 on page 161 shows the On Demand initiative alignment. The On Demand framework addresses all client major focus areas: Virtualization, Automation, and Integration.

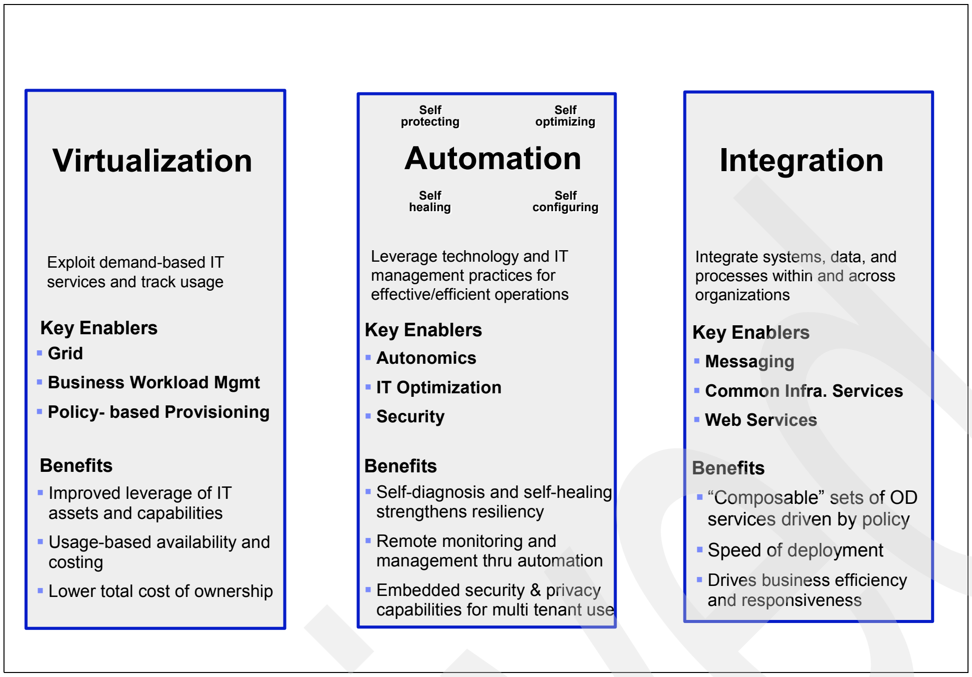

Virtualization

Exploit demand-based IT services and track usage

Key Enablers

- Grid
- Business Workload Mgmt
- Policy- based Provisioning

Benefits

- Improved leverage of IT assets and capabilities
- Usage-based availability and costing
- Lower total cost of ownership

Self protecting | Self optimizing

Automation

Self healing | Self configuring

Leverage technology and IT management practices for effective/efficient operations

Key Enablers

- Autonomics
- IT Optimization
- Security

Benefits

- Self-diagnosis and self-healing strengthens resiliency
- Remote monitoring and management thru automation
- Embedded security & privacy capabilities for multi tenant use

Integration

Integrate systems, data, and processes within and across organizations

Key Enablers

- Messaging
- Common Infra. Services
- Web Services

Benefits

- “Composable” sets of OD services driven by policy
- Speed of deployment
- Drives business efficiency and responsiveness

Figure D-3 On Demand initiatives

Storage and storage related services fall mainly into the Virtualization and Automation initiatives. IBM Global Services has defined the following storage related initiatives, shown in Figure D-4.

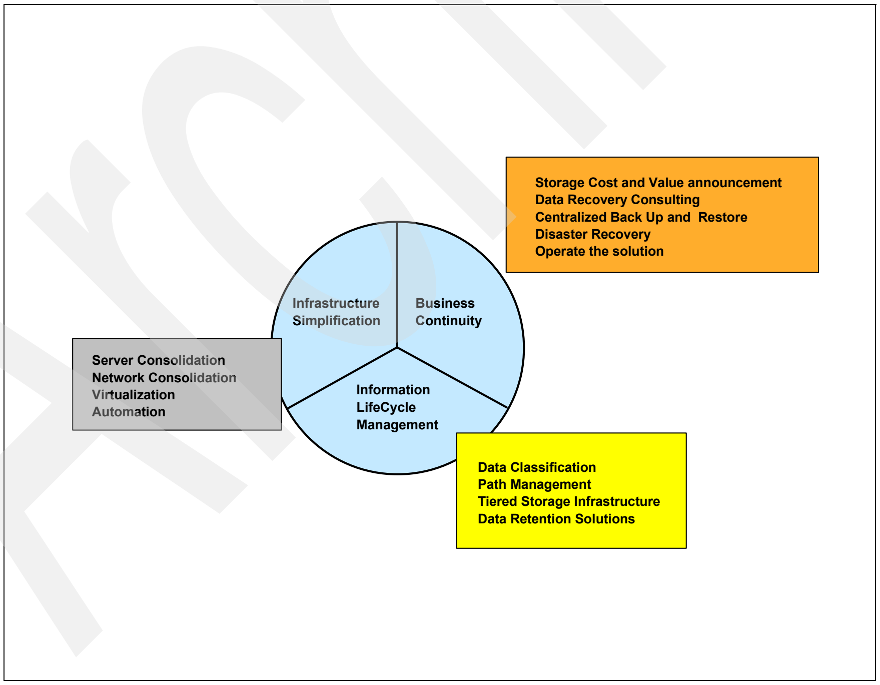
Infrastructure Simplification
Business Continuity
Information LifeCycle Management

Storage Cost and Value announcement
Data Recovery Consulting
Centralized Back Up and Restore
Disaster Recovery
Operate the solution

Server Consolidation
Network Consolidation
Virtualization
Automation

Data Classification
Path Management
Tiered Storage Infrastructure
Data Retention Solutions

Figure D-4 IBM Global Services solution classification for TotalStorage

Business Continuity solutions address application and service resiliency. Information life cycle management solutions address classification (what should go where) and data retention solutions for legal compliance. Infrastructure simplification, often a prerequisite to the preceding two solutions, addresses consolidation and virtualization aspects.

IBM Global Services solutions for resilient infrastructures

We will now illustrate a sample of IBM Global Services solutions that relate to Business Continuity and Resiliency, as shown in Figure D-5. IBM Global Services offers a vast portfolio of services that can be tailored to meet virtually any client requirement.

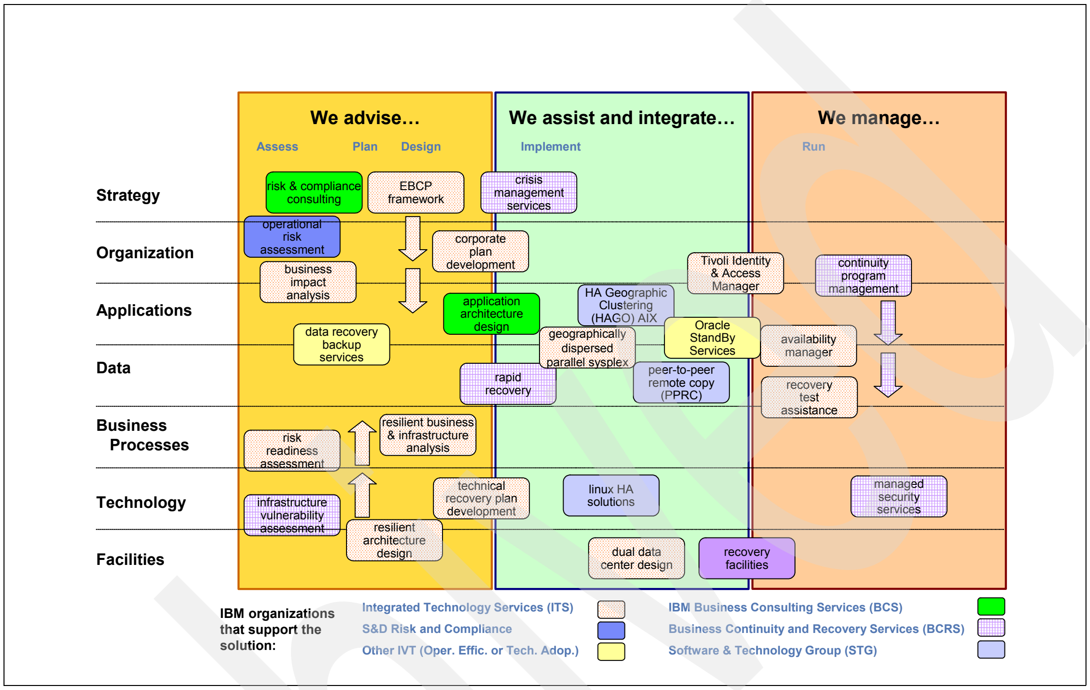
We advise…
We assist and integrate…
We manage…
Assess
Plan
Design
Implement
Run
Strategy
Organization
Applications
Data
Business Processes
Technology
Facilities
risk & compliance consulting
EBCP framework
crisis management services
operational risk assessment
corporate plan development
business impact analysis
Tivoli Identity & Access Manager
continuity program management
HA Geographic Clustering (HAGO) AIX
application architecture design
Oracle StandBy Services
availability manager
data recovery backup services
geographically dispersed parallel sysplex
peer-to-peer remote copy (PPRC)
rapid recovery
recovery test assistance
resilient business & infrastructure analysis
risk readiness assessment
technical recovery plan development
linux HA solutions
managed security services
infrastructure vulnerability assessment
resilient architecture design
dual data center design
recovery facilities
IBM organizations that support the solution:
Integrated Technology Services (ITS)
S&D Risk and Compliance
Other IVT (Oper. Effic. or Tech. Adop.)
IBM Business Consulting Services (BCS)
Business Continuity and Recovery Services (BCRS)
Software & Technology Group (STG)

Figure D-5 IBM Global Services Sample Business Resiliency services

Figure D-6 shows a suggested set of services that could be used to build and manage a resilient infrastructure.

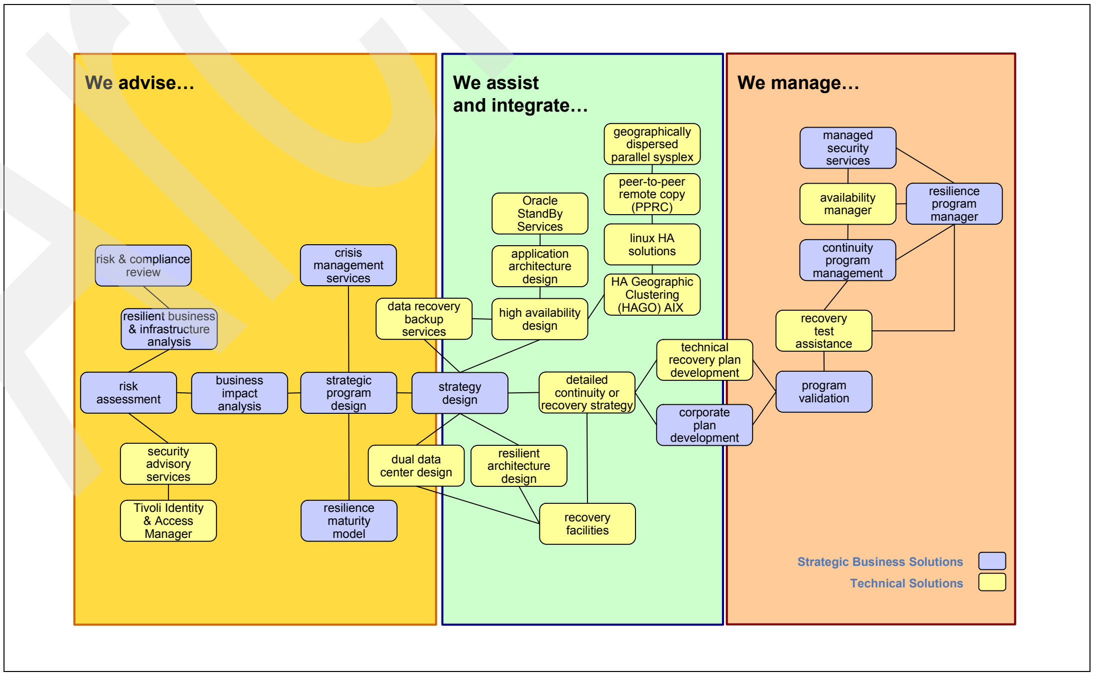
We advise…
We assist and integrate…
We manage…
geographically dispersed parallel sysplex
managed security services
peer-to-peer remote copy (PPRC)
Oracle StandBy Services
availability manager
resilience program manager
linux HA solutions
application architecture design
continuity program management
risk & compliance review
crisis management services
HA Geographic Clustering (HAGO) AIX
data recovery backup services
high availability design
recovery test assistance
resilient business & infrastructure analysis
technical recovery plan development
risk assessment
business impact analysis
strategic program design
strategy design
detailed continuity or recovery strategy
program validation
corporate plan development
security advisory services
dual data center design
resilient architecture design
Tivoli Identity & Access Manager
resilience maturity model
recovery facilities
Strategic Business Solutions
Technical Solutions

Figure D-6 Possible IBM Global Services services for a resilient infrastructure

For more information about IBM Global Services services, see the following Web site or contact your IBM representative or IBM Business Partner:

http://www.ibm.com/services/us/index.wss/gen_it

Network Consulting and Integration services

IBM Global Services offers the broadest range of services offerings in the industry for networking solutions. Our services capabilities, skills depth and breadth, and reach and range are simply unmatched in the networking services industry.

There are two basic dimensions to our networking services offerings:

- The Network Lifecycle dimension
- The Network Technology dimension (for example, routing and switching, VPN/Security, wireless, Optical Networking, and so forth)

The Integrated Technology Services (ITS) networking services offering portfolio covers both dimensions. Comprehensive information about IBM Network Services can be found at:

http://www.ibm.com/services/us/index.wss/az#N

Optical/Storage Networking

Here we discuss Optical/Storage Networking service offerings.

IBM Network Consulting for Optical Networks

This offering provides a set of supporting materials (sales positioning presentations, statement of work (SOW), technique papers, a spec sheet, and Web cast) for optical networking consulting engagements. This offering enhancement extends IBM's capability to deliver Metropolitan Area Networks and other optical networking solutions using technologies such as Gigabit Ethernet, Dense Wave Division Multiplexing (DWDM), Coarse Wave Division Multiplexing (CWDM), Synchronous Optical Networks / Synchronous Digital Hierarchy (SONET/SDH), and Channel Extensions.

IBM Network Integration Deployment Services for Optical Networking

Enterprise optical networking services from IBM deliver comprehensive solutions for high-speed, high-availability networking to improve performance, optimize information technology (IT) investments, and enhance client service. Our team can help you streamline your voice, video and data networks with advanced fiber optic technologies through a variety of assessment, strategy, architecture, and design services. We evaluate your systems and make recommendations that are fully integrated with your business objectives, helping you expand capacity as your business grows, so you can stay competitive in the marketplace.

This offering provides a set of supporting materials (CIO presentation, sample SOWs, guides, pricing and sizing advice, a spec sheet, and a Web cast) for engagements that involve planning for and implementing optical networks for enterprises.

IBM Network Consulting Services for Storage Networking

IBM Consulting Services for Storage Networking can help you determine the storage technology that best fits your needs, whether that technology is server attached storage (SAS), network attached storage (NAS), or storage area network (SAN). We can then define the integration process based on your company's particular access and storage methods. Our networking experts work closely with your staff to help determine the optimal data paths between users and data storage devices for efficient storage networking.

This offering provides a set of presentations, guides, and other supporting materials for marketing and executing a storage network strategy and conceptual design engagement. This is a vendor neutral approach that looks at multiple options, including traditional server attached storage (SAS), network attached storage (NAS), and storage area networks (SANs). The client's storage and networking environments are examined in relationship to its business and IT strategies. This analysis is used as input to a networking storage strategy and a conceptual design. Individual guides address input to the strategy, developing the strategy, the design options, and issues that should be considered when developing a conceptual design for storage networking.

The documentation includes an executive (CIO) presentation, SOWs, and technique papers covering assessment, strategy, and architecture/high level design. A Web cast provides an explanation of the service and an overview of the documentation. A marketing guide and resource matrix are also provided.

IBM Network Integration Deployment Services for Storage Networking

IBM Networking Services for Storage offers adaptable solutions for improving storage throughput and meeting specific networked storage needs. For example, if your company has several branches distributed on a metropolitan area network, we can integrate your data center environments using high-speed, dark-fiber services, which gives each branch direct, full-speed connections to centralized storage via Dense Wavelength Division Multiplexing (DWDM) technology.

This offering provides supporting materials (a spec sheet, sample SOW, and a presentation) for engagements for integrating storage devices and servers within a client's existing network. The materials are intended for marketing and performing on engagements concerned with storage area networks or network attached storage:

- Networking between a client's remotely installed storage area networks (SAN) in a metropolitan area network (MAN) using dense wavelength division multiplexing (DWDM) technology
- Integration of a client's existing Fibre Channel-based SAN with its existing TCP/IP network in a local area network (LAN) or a wide area network (WAN) using the Internet SCSI (iSCSI) technology
- Addition of storage devices and servers into a client's existing network, using network attached storage (NAS) technology
- Detailed design and implementation for OEM partners

IBM Network Consulting for Resilient Networks

IBM is uniquely positioned to provide you with a comprehensive approach to operational resilience that spans every aspect of your business. IBM can apply expertise from any field to match your business requirements. As a leading networking services provider, IBM has extensive knowledge in state-of-the-art networking technologies, and has helped clients worldwide establish tailored, resilient network solutions. Working with leading network equipment providers and network service providers, IBM Network Consulting can bring you the most cost-effective network plan to meet your operational resilience requirements.

This offering provides a set of supporting materials (CIO and marketing presentations, SOW, technique papers, pricing and sizing advice, a spec sheet, and a Web cast) for resilient networking consulting engagements. This offering enhancement focuses on network elements needed to maximize the ability of an enterprise to maintain continuity of operation in case of unplanned events that threaten its operation, regardless of their origin. The assessment, strategy, and design steps recommended in the enhancement focus on recommendations for an enterprise-wide set of technologies and processes aimed at ensuring the flow of information whenever and wherever needed to keep a business operating.

Network Deployment

Here we discuss Network Deployment service offerings.

IBM Rapid Network Deployment for e-business

Rapid Network Deployment for e-business includes:

- Project management
- Strategy and standards definition
- Planning and implementation
- Ordering support and procurement
- Logistics
- Predelivery preparation
- Site preparation
- Cabling services
- Configuration and installation
- Asset control

This engagement portfolio provides several different types of intellectual capital that will assist in both the sales processes and the actual implementation of Rapid Network Deployment opportunities. While the emphasis is on “rapid”, all of the materials are useful for any type of network deployment that may include procuring hardware, configuration, installation, cabling, site services, site surveys, and so on, for network upgrades, enhancements, or new installations at multiple sites both within a single country and multi-nationally.

Network Management

Here we discuss Network Management service offerings.

IBM Network Consulting and Integration Services: Network Management

This offering provides a guide and sample SOW to provide an approach to creating a network management architecture and design that combines IBM and OEM products. The method includes:

- Assessing the client's physical infrastructure, network management tools, tracking and reporting indicators, supervision environment, and support structure
- Developing network management requirements and norms
- Recommending a network management architecture and design
- Identifying gaps in the network management and support environment
- Recommending a network management process and people role improvements
- Modeling and validating the proposed network management platform and supporting environment
- Planning for implementation

While this material is tailored to IBM and OEM products, much of the content applies to network management architecture and design engagements in general, especially if two or more vendor products may be used to implement the design.

The guide contains suggested areas for follow-on engagements, including capacity planning, QOS audits, and out-tasking of network management functions.

Monitoring and Performance Analysis of the Network Infrastructure

IBM Network Management Services offers a comprehensive solution for your network monitoring and management needs. We combine powerful network management tools with our proven methodologies and processes to track your network's performance and optimize its availability at a fixed cost. Our suite of remote management services is provided through a Network Operations Center (NOC) and includes monitoring, performance management, problem management, change management, configuration management, and security management.

Network Cabling

Here we discuss Network Cabling service offerings.

Optical Fiber Cabling Solutions for the Enterprise Data Center

This offering provides a statement of work, a physical and configuration planning guide, and supporting materials for data center cabling services using Fibre Channel Protocol (FCP) and a 50 micron multimode fiber. Storage area network (SAN) products are a special focus, as is how to integrate and deploy 50 micron multimode fiber within existing systems.

Internet Data Center Server Farm Cabling

IBM Networking and Connectivity Services Cabling Services for Internet Data Center Server Farms provides high-density yet flexible cabling solutions that incorporate state-of-the-art communication components with full integration of both copper and fiber optic cabling. These solutions support your business objective of having a network solution that is reliable, highly available, adaptable, easy to use and manage, and cost-effective.

This offering provides a set of guides, presentations, and supporting materials for marketing and performing a cabling engagement at an Internet data center. These materials cover the planning, design, and installation of cabling infrastructure for Internet data center server farms. The documentation includes a statement of work (SOW), a spec sheet, a design guide, a technology and component guide, an installation guide, and education material.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this IBM Redbook.

IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 168. Note that some of the documents referenced here may be available in softcopy only.

General

- *IBM System Storage Business Continuity: Part 1 Planning Guide*, SG24-6547
- *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548
- *IBM System Storage Solutions Handbook*, SG24-5250

Software

- *Disaster Recovery Using HAGEO and GeoRM*, SG24-2018
- *IBM System Storage SAN Volume Controller*, SG24-6423

Disk

- *DS4000 Best Practices and Performance Tuning Guide, SG24-6363*
- *IBM System Storage DS4000 Series, Storage Manager and Copy Services*, SG24-7010
- *IBM System Storage DS6000 Series: Architecture and Implementation*, SG24-6781
- *The IBM TotalStorage DS6000 Series: Concepts and Architecture*, SG24-6471
- *IBM System Storage DS6000 Series: Copy Services with IBM System z*, SG24-6782
- *IBM System Storage DS6000 Series: Copy Services in Open Environments*, SG24-6783
- *IBM System Storage DS8000 Series: Copy Services with IBM System z*, SG24-6787
- *IBM System Storage DS8000 Series: Copy Services in Open Environments*, SG24-6788
- *The IBM System Storage N Series*, SG24-7129
- *The IBM TotalStorage DS8000 Series: Concepts and Architecture*, SG24-6452
- *IBM TotalStorage Enterprise Storage Server Model 800*, SG24-6424
- *Implementing Linux with IBM Disk Storage*, SG24-6261

Tape

- *IBM Tape Solutions for Storage Area Networks and FICON,* SG24-5474
- *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632
- *IBM TotalStorage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide,* SG24-6115
- *The IBM TotalStorage Tape Libraries Guide for Open Systems,* SG24-5946
- *IBM TotalStorage Virtual Tape Server Planning, Implementing and Monitoring,* SG24-2229

- *Implementing IBM Tape in UNIX Systems*, SG24-6502

SAN

- *Designing an IBM Storage Area Network*, SG24-5758
- *IBM TotalStorage: Implementing an Open IBM SAN*, SG24-6116
- *IBM SAN Survival Guide*, SG24-6143
- *IBM SAN Survival Guide Featuring the Cisco Portfolio*, SG24-9000
- *IBM SAN Survival Guide Featuring the IBM 3534 and 2109*, SG24-6127
- *IBM SAN Survival Guide Featuring the McDATA Portfolio*, SG24-6149
- *Introduction to Storage Area Networks*, SG24-5470

Tivoli

- *Deploying the Tivoli Storage Manager Client in a Windows 2000 Environment*, SG24-6141
- *Disaster Recovery Strategies with Tivoli Storage Management*, SG24-6844
- *Get More Out of Your SAN with IBM Tivoli Storage Manager*, SG24-6687
- *IBM Tivoli Storage Management Concepts*, SG24-4877
- *IBM Tivoli Storage Manager for Advanced Copy Services*, SG24-7474
- *IBM Tivoli Storage Manager: Bare Machine Recovery for AIX with SYSBACK*, REDP-3705
- *IBM Tivoli Storage Manager Implementation Guide*, SG24-5416
- *IBM Tivoli Storage Manager Version 5.3 Technical Guide*, SG24-6638
- *IBM Tivoli Workload Scheduler Version 8.2: New Features and Best Practices*, SG24-6628
- *Using IBM Tivoli Storage Manager to Back Up Microsoft Exchange with VSS*, SG24-7373

How to get IBM Redbooks

You can search for, view, or download IBM Redbooks, IBM Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy IBM Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

A

application integration 13
asynchronous 143
Asynchronous Transfer Mode (ATM) 89

B

backup
  software 103
    close integration 103
backup server 155
bandwidth 84
base logical drive 131
  multiple FlashCopies 131
bitmaps 141
BRMS
  Tivoli Storage Manager 52
building block 94
Business Continuity 3, 107
  assessment 8
  basic definition 110
  causes of outages 7
  impact to business 7
  IT roadmap 12
  key components 109
  metrics of success 8
  selecting a solution 11, 16
  services 93
  specific requirements 108
  successful implementation 109
  tiers 17
  value of Solution Selection methodology 20
business process 95

C

CDWM 88
channel extension 89
  equipment spoof 89
  FCIP
  FICON and Fibre 90
  implementation 89
  methods 89
clustering technologies 95
  common shared cluster 97
  geographical dispersed clusters 99
  shared nothing clusters 95
Coarse Wavelength Division Multiplexing see CWDM
content management
  application 76
Continuous Availability 23
  GDPS 24
  HACMP/XD 35
Continuous Availability for MaxDB and SAP liveCache 23, 38
  solution components 39
  solution description 38
  solution highlights 38
  tier level 38
Continuous Availability for open systems see GDOC
Continuous Availability for zSeries data
  GDPS/PPRC HyperSwap Manager 46
Continuous Availability for zSeries see GDPS
continuous operations 5
  definition 6
Copy Services for System i 39
  solution description 40
core technologies 12
critical data 108
CWDM 88

D

dark fiber 83
data
  retention-managed 73–74
data availability 93
Data Backup with no Hot-Site 118
Data Center
  Point-to-Point configuration 83
data consistency 120, 143
data life cycle 74
data loss 95, 108, 112
  Several hours 119
data mining 74
data mirroring 100
Data ONTAP 136
data protection 95
data rendering 74
Data Retention 73–75
  Compliance 75
  Policy 74
  System Storage Archive Manager 75
data shredding 74
data transport speed 84
database
  Disaster Recovery considerations 100
DDR 78
dedicated fiber 83
Dense Wavelength Division Multiplexing see DWDM
DFSMS/VM 78
Disaster Recovery 5, 108, 110
  automated management 15
  definition 7
  differentiating Business Continuity 7
disaster recovery 73
disk drive 75
disk subsystem 46
distance limitations 142
DS6000 122

DS8000 122
DWDM 88

E
electronic vaulting 64, 118
  Tivoli Storage Manager 64
Enhanced Remote Mirroring 100
eRCMF 45
  solution description 48
ERP 154–155
ESS 122
EXP100 unit 75

F
failover 94
failover situation 142
FCIP 90
feature key 132
fiber transport 83
  dark fiber 83
  SONET 83
Fibre Channel (FC) 88, 90
  protocol 90
Fibre Channel over IP see FCIP
FICON 90
FlashCopy 122–123, 130–132, 136, 141
  volume 104
FlashCopy logical drive 131
FlashCopy Manager 45, 53

G
GDOC 23, 34
  consulting and planning 34
  solution components 33
  solution description 33
  solution highlights 33
  VERITAS 34
GDPS 15, 23–24
  IBM Global Services 31
  summary 32
  Technical Consulting Workshop 31
  three site solutions overview 30
GDPS/PPRC 24
  IBM Global Services 32
  management for Open Systems LUNs 26
  Multi-Platform Resiliency for zSeries 26
  Near Continuous Availability of data with HyperSwap 26
  overview 25
  Tier 7 Near Continuous Availability 26
GDPS/PPRC HyperSwap Manage
  overview 27
GDPS/PPRC HyperSwap Manager 46
  IBM Global Services 32
  Resiliency Family positioning 49
GDPS/XRC 25
  IBM Global Services 32
  overview 27–28
  topology 28–29
Geographically Dispersed Parallel Sysplex see GDPS
Global Copy 122, 125, 130
Global Mirror 122, 126, 130, 143

H
HACMP 36
HACMP/XD 23, 35–36
  hardware requirements 37
  solution components 37
  solution description 36
  solution highlights 37
  tier level 35
hardware infrastructure 12
high availability 5
  definition 6
highly automated, business integrated 120
hot site 118
human error 95, 111

I
IBM AIX
  HACMP 96
IBM DB2 Content Manager 75
IBM Global Services 93
  GDPS 31
  GDPS Technical Consulting Workshop 31
  GDPS/PPRC 32
  GDPS/PPRC HyperSwap Manager 32
  GDPS/XRC 32
  RCMF 31
IBM Global Services solutions
  resilient infrastructures 162
IBM System Storage Resiliency Family 9
IBM TotalStorage DR550 73
  main focus 76
IBM TotalStorage Resiliency Family 13
  strategic value 16
IP packet 89
iSeries
  introduction 52

L
latency 84
local mirroring
  N series
    local mirroring 138
Log file 101, 103
logical drive
  primary 132
logical online backup (LOB) 103
  storage layer 104
logical volume 100

M
mapping 141
MaxDB 38
Metro Cluster for N Series 42

solution description 42
Metro Mirror 46, 100, 122, 124, 130
Metro/Global Copy 126
mining 74

N
N series
Data ONTAP 136
synchronous local mirroring 138
SyncMirror 138
NASD 75
Network Consulting and Integration 163
network topologies 82
network transport
technology selection 86
testing 90
Number of Transactions (NT) 102
Number of Transactions/Second (NTS) 102

O
operating system (OS) 96
optical amplification 87
optical amplifiers 88
overwritten 141

P
Physical Recovery Time (PRT) 102
Point-in-Time 132
point-in-time (PIT) 119
point-in-time copy (PTC) 119
point-in-time image 131
PPRC Migration Manager 45, 53
Tier 6 53
primary 132
PTAM 60

R
Rapid Data Recovery 45
Rapid Data Recovery for UNIX and Windows
eRCMF 45
SAN Volume Controller 45
RCMF
IBM Global Services 31
RCMF/PPRC 24
RCMF/XRC 25
real-time synchronized 142
Records Manager 75
Recovery Point Objective see RPO
recovery time 102
Recovery Time Objective see RTO
Redbooks Web site 168
Contact us xii
regenerative repeaters 88
relationship 143
Remote Mirror and Copy 122
Remote Volume Mirror 132
rendering 74
retention period 74
retention policy 74
retention-managed data 73–74
RPO 11, 17
RTO 8, 11, 17, 94

S
SAN File System 56
SAN Volume Controller 45, 49
Global Mirror 143
SAP 38
SAP liveCache 23
SATA 73, 75
SEC 75
Serial Advanced Technology Attachment (SATA) 75
server integration 13
service level 110
services
Network Consulting and Integration 163
Network Deployment 165
Network Management 165
services and planning 157
shredding 74
single point 110
single point-of-failure (SPOF) 94
SMB
business continuity affordability 112
Business Continuity implementation 111
business continuity implementation steps 111
Business Continuity planning 111
continuous operations 110
Disaster Recovery 110
High Availability (HA) 110
IT data center and staff 109
IT needs 109
prevention services 110
recovery services 111
tier levels 113
SMB company 107–108
Business Continuity 110
business continuity solutions 109
top priorities 108
SMB enterprise 107
solution component 113
SONET 83
SONET ring 83
source logical drive 132
source virtual disks 141
STANDARD LOGICAL Drive 131
storage area network (SAN) 49
storage device 108
Storage Management
software 131
storage subsystem 100
standard logical drive 131
storage system 95
storage virtualization 49
subrate multiplexing 88
Supply Chain Management (SCM) 109
Suspend Mirror 136
synchronous local mirroring 138

SyncMirror 138
SysBack 155

T
tape-based storage 75
target logical drive 132
target virtual disks 141
TDM 88
  card 88
Tier 0 118
Tier 1 57, 118
Tier 2 57, 118
Tier 3 57, 118
Tier 4 57, 119
Tier 5 57, 119
Tier 6 57, 120
  PPRC Migration Manager 53
Tier 7 120
tiers
  blending 18
  Business Continuity 17
  levels 113
Time Division Multiplexing see TDM
Tivoli Storage Manager 56, 152
  backup 60
  BRMS 52
  client 61
  clustering 57, 68
  DRM 61
  DRM process 63
  electronic vaulting 57
  family 56
  family overview 151
  manual off-site vaulting 57
  off-site manual vaulting 62
  server 56
  solutions 58
  solutions overview 57
Tivoli Storage Manager (TSM) 57
Tivoli Storage Manager for Database 153
Tivoli Storage Manager for Enterprise Resource Planning 154–155
Tivoli Storage Manager for Mail 154
Tivoli Storage Manager for System Backup and Recovery 155
TotalStorage Resiliency Family
  application-specific integration 16
  hardware infrastructure 14
  product layers 14
transaction integrity 119

V
VERITAS 34

W
WAFL 137
WDM 87
  optical amplification 87
  subrate multiplexing 88
write caching 133–134
Write Once Read Many (WORM) 75

Z
z/OS Global Mirror 122, 126
z/OS Metro/Global Mirror 129
z/VM 78
zero data loss 120
zSeries 45

IBM Redbooks

IBM System Storage Business Continuity Solutions Overview

(0.2”spine)
0.17”<->0.473”
90<->249 pages

IBM

IBM System Storage Business Continuity Solutions Overview

Redbooks

Anticipating and responding to risks quickly and cost-effectively

Reviewing of Business Continuity expertise and skills

Using IBM System Storage products

This IBM Redbook presents an overview and descriptions of IBM System Storage Business Continuity Solutions for Backup / Restore, Rapid Data Recovery, and Continuous Availability.

IT Business Continuity concepts are discussed; advice, tips, and a roadmap for selecting the optimum IT Business Continuity solution for your organization is provided.

IBM System Storage products are described that can fulfill your IT Business Continuity solution design. This IBM Redbook is a summary of the detailed information contained in *IBM System Storage Business Continuity: Part 1 Planning Guide*, SG24-6547 and *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548.

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

SG24-6684-01 ISBN 0738489719