IBM

IBM System Storage Business Continuity: Part 1 Planning Guide

- Describes current trends and strategies for Business Continuity
- Explains how to select an appropriate Business Continuity solution
- Presents a step-by-step Business Continuity planning workshop

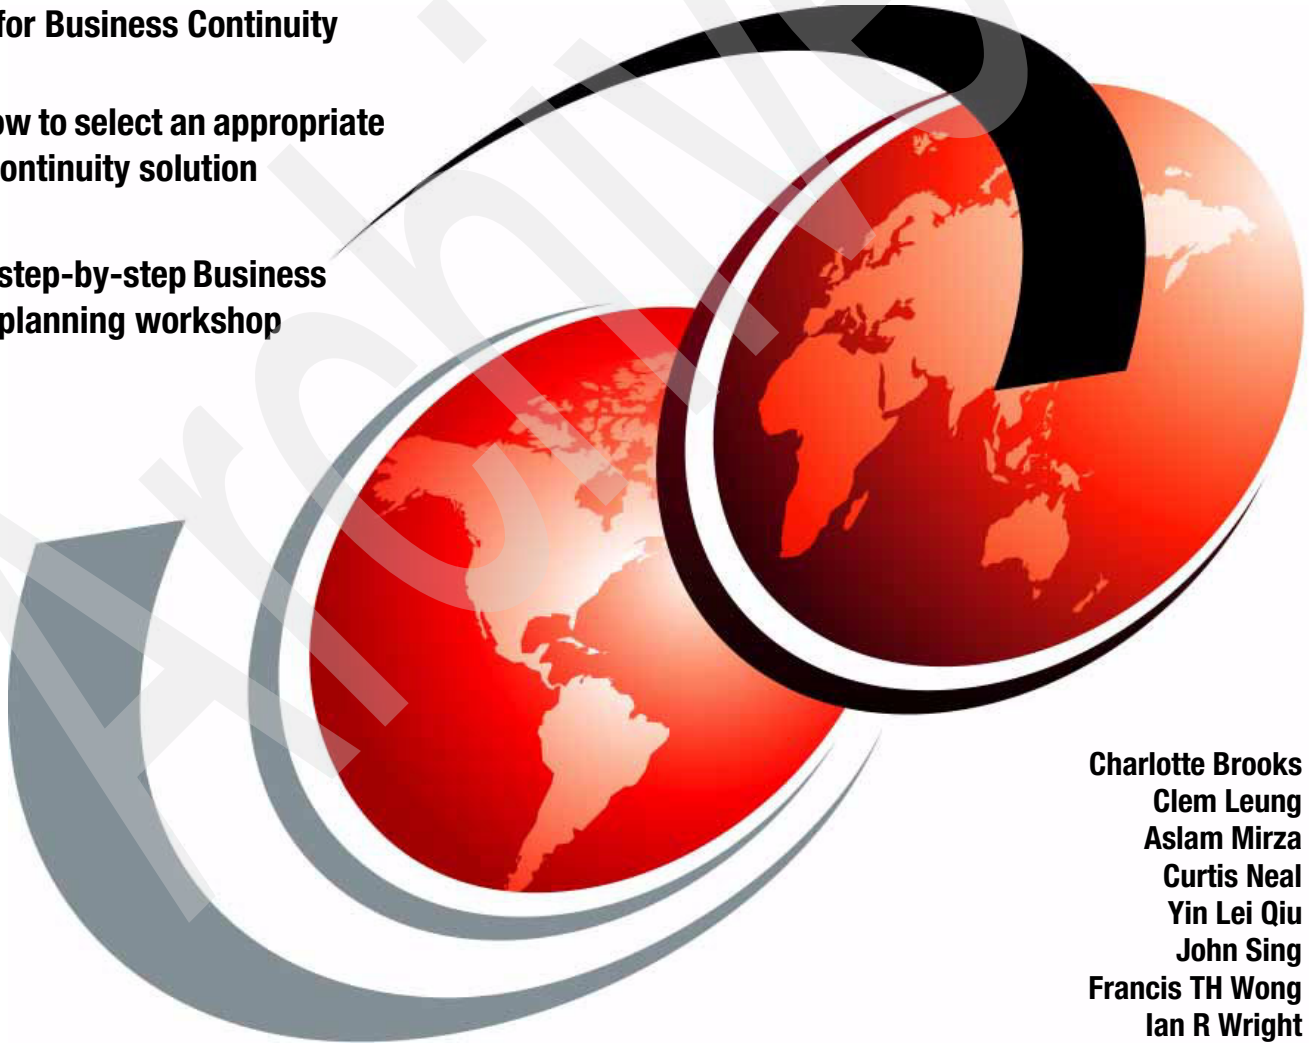

Charlotte Brooks
Clem Leung
Aslam Mirza
Curtis Neal
Yin Lei Qiu
John Sing
Francis TH Wong
Ian R Wright

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM System Storage Business Continuity: Part 1 Planning Guide

March 2007

SG24-6547-03

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

Fourth Edition (March 2007)

This edition refers to the IBM System Storage Resilience Portfolio.

Contents

Notices

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Trademarks

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

Redbooks (logo) ™
e-business Hosting™
z/OS®
AIX®
AS/400®
BladeCenter®
Chipkill™
Domino®
DB2®
DFSMSdfp™
DS4000™
DS6000™
DS8000™
Enterprise Storage Server®
ESCON®
FlashCopy®
FICON®
General Parallel File System™
Geographically Dispersed Parallel Sysplex™
GDPS®
GPFS™
HyperSwap™
HACMP™
Informix®
IBM®
Lotus Notes®
Lotus®
MQSeries®
Notes®
OS/390®
OS/400®
Parallel Sysplex®
Redbooks™
RMF™
RS/6000®
S/390®
System i™
System p™
System x™
System z™
System Storage™
System Storage Proven™
Tivoli Enterprise™
Tivoli®
TotalStorage®
WebSphere®

The following terms are trademarks of other companies:

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

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

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

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

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

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

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

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

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

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

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

Preface

A disruption to your critical business processes could leave the entire business exposed. Today's organizations face ever-escalating customer demands and expectations. There is no room for downtime. You need to provide your customers with continuous service because your customers have a lot of choices. Your competitors are standing ready to take your place. As you work hard to grow your business, you face the challenge of keeping your business running without a glitch. To remain competitive, you need a resilient IT infrastructure.

This IBM® Redbook introduces the importance of *Business Continuity* in today's IT environments. It provides you with a comprehensive guide to planning for IT Business Continuity and can help you to design and to select an IT Business Continuity solution that is right for your business environment.

We discuss the concepts, procedures, and solution selection for Business Continuity in detail, including the essential set of IT Business Continuity requirements that you need to identify a solution. We also present a rigorous Business Continuity Solution Selection Methodology that includes a sample Business Continuity workshop with step-by-step instructions in defining requirements. This IBM Redbook is meant as a central resource book for IT Business Continuity planning and design and is intended for anyone who wants to learn about Business Continuity trends and strategies.

The companion title to this IBM Redbook, *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548, describes detailed product solutions in the System Storage Resiliency Portfolio.

The team that wrote this IBM Redbook

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

Charlotte Brooks is an IBM Certified IT Specialist and Project Leader for Storage Solutions at the ITSO, 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 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, Business Partners, and Clients. He has more than 6 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 Hong Kong and 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 Hong Kong, 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 7 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.

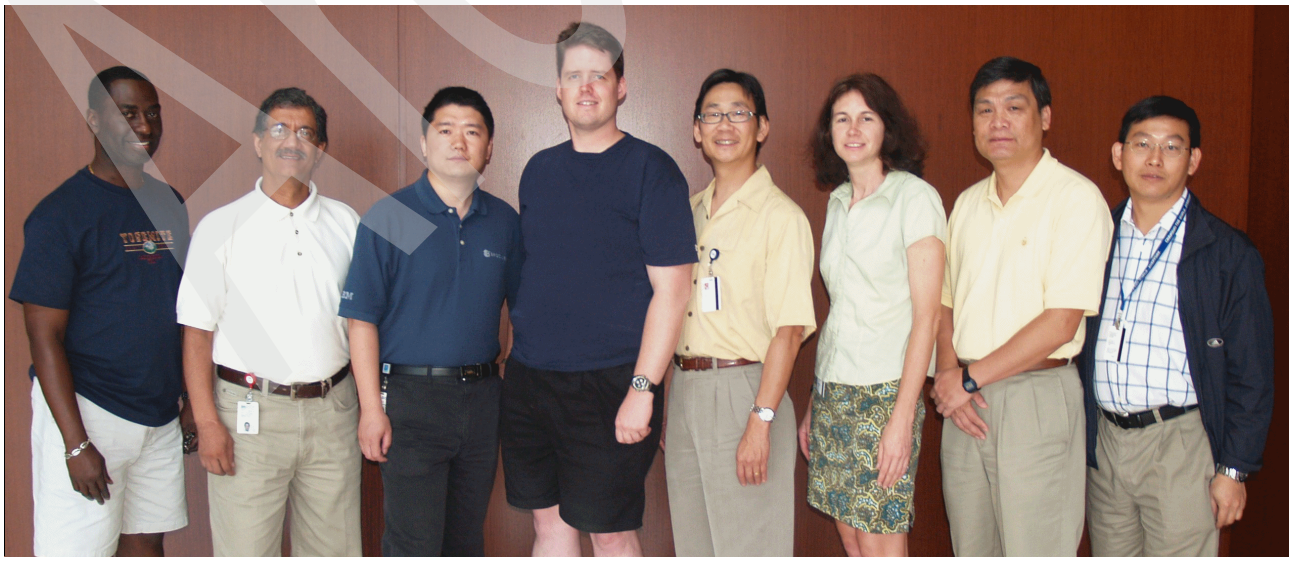

Figure 1 The team: Curtis, Aslam, Yin Lei, Ian, John, Charlotte, Clem, and Francis

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
ITSO, 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, Debbie Willmschen
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

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

Thomas Beretvas
Beretvas Performance Consultants

John Koberlein
Kobe Consulting

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 will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our IBM 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 Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an e-mail to:

  redbooks@us.ibm.com

- Mail your comments to:

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

1

Introduction

This chapter discusses the definition of *Business Continuity* and how the definition is evolving in relation to technological and social trends. Organizations today can never be fully in control of the business environment, and all organizations can face a Business Continuity event at some point.

This IBM Redbook discusses the thought processes, methods, solutions, and product concepts that you can apply to today's Business Continuity requirements. With an increased ability to understand, evaluate, select, and implement solutions that successfully answer today's Business Continuity requirements, enterprises can continue to maintain marketplace readiness, competitive advantage, and sustainable growth.

We recommend that you implement Business Continuity *best practices*. Because all organizations are different, however, specific techniques that work in one organization will not necessarily work in another.

1.1 Business Continuity versus Disaster Recovery

Before we proceed, we need to clarify the terms *Business Continuity* and *Disaster Recovery*. These terms are sometimes used interchangeably, as are *business resumption* and *contingency planning*.

Business Continuity is the ability to adapt and respond to risks, as well as opportunities, in order to maintain continuous business operations. There are three primary aspects of providing Business Continuity for key applications and processes:

- *High availability* is the capability to and processes that provide access to applications regardless of local failures, whether these failures are in the business processes, in the physical facilities, or in the IT hardware or software.
- *Continuous operations* is the capability to keep things running when everything is working properly; where you do not have to take applications down merely to do scheduled backups or planned maintenance.
- *Disaster Recovery* is the capability to recover a data center at a different site if a disaster destroys the primary site or otherwise renders it inoperable. The characteristics of a disaster recovery solution are that processing resumes at a different site and on different hardware.

It is necessary to clarify and differentiate between the terms *Disaster Recovery* and *Business Continuity*. Strictly speaking, Disaster Recovery is the ability to recover data and is only one component of an overall Business Continuity Plan. This book discusses all three aspects of IT Business Continuity.

Clearly, an organization will identify many components in the process of creating a comprehensive Business Continuity plan. Inspired by the concept of database objects, these components have attributes that help define them in terms of their ability to address the basic requirements of Business Continuity. Not surprisingly, they echo the components of most business organizations, as shown in Figure 1-1.

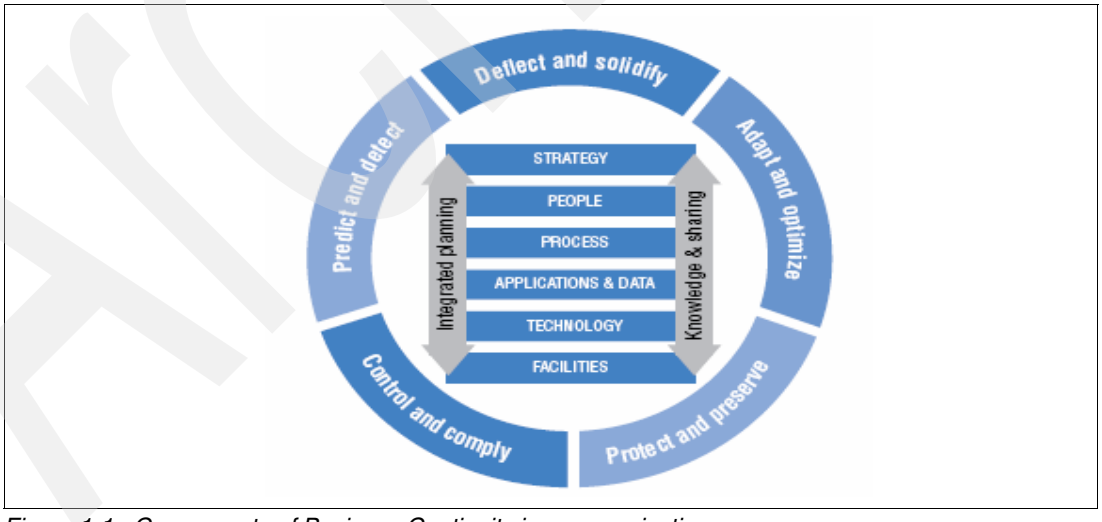

Figure 1-1 Components of Business Continuity in an organization

The components of Business Continuity are:

- **Strategy:** Objects that are related to the strategies used by the business to complete day-to day activities while ensuring continuous operations.
- **Organization**: Objects that are related to the structure, skills, communications, and responsibilities of its employees.
- **Applications and data:** Objects that are related to the software necessary to enable business operations, as well as the method to provide *high availability* that is used to implement that software.
- **Processes:** Objects that are related to the critical business process necessary to run the business, as well as the IT processes used to ensure smooth operations.
- **Technology**: Objects that are related to the systems, network, and industry-specific technology necessary to enable *continuous operations* and backups for applications and data.
- **Facilities:** Objects that are related to providing a *disaster recovery* site if the primary site is destroyed.

There is no doubt that some sort of Business Continuity plan is essential. The Business Continuity plan becomes a source reference at the time of a Business Continuity event or crisis and the blueprint upon which the strategy and tactics of dealing with the event or crisis are designed.

1.2 Evolving definitions of disaster and recovery

We defined *Disaster Recovery* as the capability to recover a data center at a different site if a disaster destroys the primary site or otherwise renders it inoperable.

The definition of what constitutes a disaster is driven by innovations in technology. As the definition of what constitutes a disaster has changed, so too have the methods of assuring that you can recover from the disaster. Disasters occur in many forms, and what you define as a disaster is specific to your business. Disasters can be natural, man-made, or technical events that cause a disruption to your business. As we discuss in this book, for some organizations, the ability to recover data can be of secondary importance to ensuring that their data remains available for use. The needs of each organization differs. A disaster that you define for your business might not be the same disaster for another business segment or might not be within their definition of a disaster.

Business Continuity management is the outcome of the process that started in the early 1970s as disaster recovery planning rather than management. During that time, the disaster recovery activity was controlled by the data processing manager. In those days, if a major event or crisis occurred, the outage could be measured in days rather than hours. Financial organizations, such as banks and insurance companies, invested in alternate sites. Backup tapes were stored at protected sites away from computers. Recovery efforts were almost always triggered by a fire, flood, storm, or other physical devastation.

The 1980s saw the growth of commercial recovery sites offering computer services on a shared basis, but the emphasis was still only on IT recovery. The 1990s witnessed significant change in the IT environment and the move from disaster recovery planning to Business Continuity planning.

Hazards today could be a sudden, unplanned calamitous event that causes damage or loss, making it impossible for an organization's to provide the critical business function for some predetermined period of time. With the globalization of most business and the pervasiveness

of data access, data availability is most critical. Data disruptions for some industries, such as financial institutions, can result in staggering financial losses. Inability to cope or plan for such disruptions can cause businesses huge amounts of financial loss. Reports show that no plans for Business Continuity or improper planning causes businesses to lose their competitive advantage and to lose customers.

The methods available to assist recovery efforts are made more complex by application architectures such as distributed applications, distributed processing, distributed data, and hybrid computing environments.

Then there is the issue of data volume. Applications such as decision support, data warehousing, data mining, and customer resource management can require petabyte-size investments in online storage. Typical yearly growth of new data in an enterprise is in the range of 40% to 70%. More data to manage means more data to recover. It is imperative that you can recover your essential business processes in less time than what was traditionally considered possible. Businesses with a low tolerance level for outages will seek a shorter time to recover than a business with a higher tolerance level. Gradient of tolerance levels can be measured in terms of a business's ability to lose data during the outage period. Revenue loss in some businesses are measured in millions of dollars per hour.

Data recovery no longer lends itself to a one-dimensional approach. The complex IT infrastructure of most installations has just exceeded the ability of most shops to respond in the way they did just a few years ago. The schematic depicted in Figure 1-2 is typical of the diversity and complexity of many, if not most, IT processing environments.

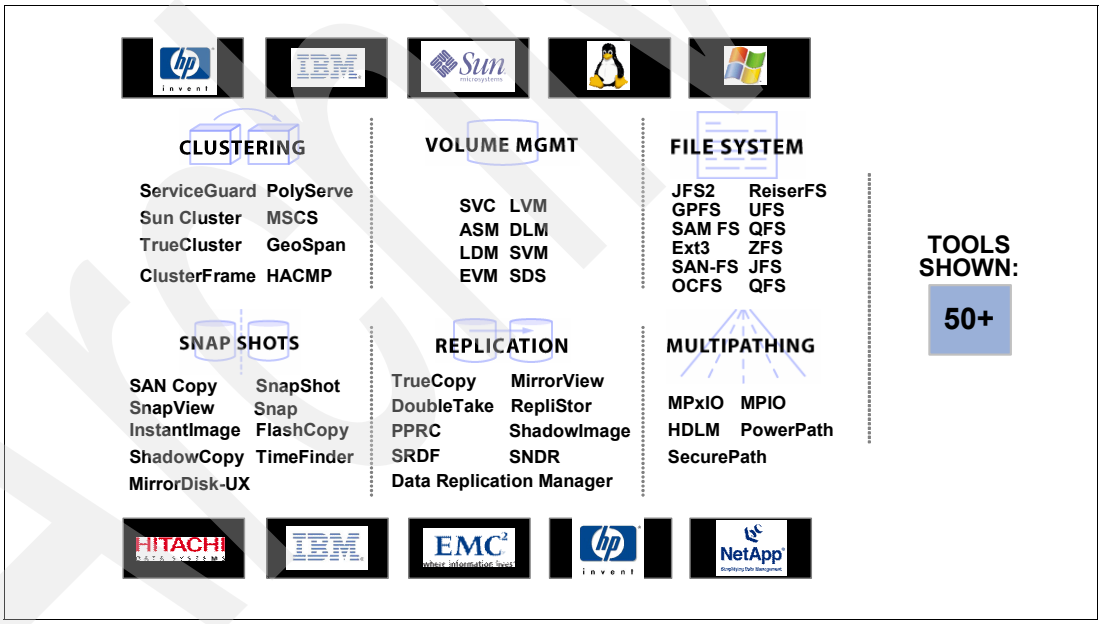

Figure 1-2 Heterogeneous environments necessarily increase complexity

Do all organizations need a real-time recovery strategy in place? Not necessarily. These processes can withstand unplanned outages for a day, or even several days, without severe adverse consequences (in other words, a *business tolerance level* to outages). However, while the loss of phone service for a catalog retailer might be sufferable for a short period of time, as the outage enters day two and three, the revenue losses will grow exponentially. Research studies have shown that without proper planning, businesses that somehow recovered from an immediate disaster event frequently did not survive in the medium term.

1.3 Disasters: Old, new, and planned

Figure 1-3 catalogs various types of unplanned outages. In addition, you can probably add one or two from your own experience. Perhaps we could have made our list more inclusive by including subtopics or by cross-referencing some of our examples. For example, vandalism or sabotage might include hacker and virus threats. We also could have added worms, trojan horses, distributed denial-of-service attacks, and industrial espionage under a hacker subtopic. The list of things that can go wrong is seemingly endless.

For our list of possible outages, we did include relocation delay, which is really a sort of planned outage that includes data center relocations or the migration back to the primary computing facility from a hot or cold site. Typically, when people think of a disaster, they usually think of a surprise event or something unexpected. Here, we want to draw attention to those well-planned outages where even with the best of planning, something goes wrong. In the end, does it really matter to your customer why the data is unavailable to them?

Figure 1-3 A sample list of the potential types of disasters

1.4 Disastrous results

Disasters can cover a wide gamut of definitions, but the results can also vary depending on the nature of the disruption. Although you cannot anticipate that a disaster will happen, you can plan or mitigate a disaster and prevent minimal disruptions to fit your business requirement:

- **Direct losses**

  Direct losses represent the value of any assets that might have been lost and can include your key personnel, data, infrastructure, network, and so forth. The losses are directly related to the value of the assets or tangible assets.

- **Indirect losses**

  Indirect losses or intangible losses are caused by the inability to operate or having to operate in a severely degraded way due to the unavailability of resources such as data access, power, telephone service, and others. Usually these losses continue or become worse until the disruption is alleviated. Losses include lost productivity as well as revenue.

- **Consequential losses**

  Consequential losses or repercussions are caused by the outage but can remain long after recovery or restart. They include organizational image, stock value, and loss of customers and market share which relate to the business going out of business.

1.5 Shapers of recovery and availability strategies

Because each organization is unique and has, therefore, unique requirements for its data recovery and availability strategy, let us examine some of the forces that are currently impacting the recovery strategies that organizations must implement, as illustrated in Figure 1-4.

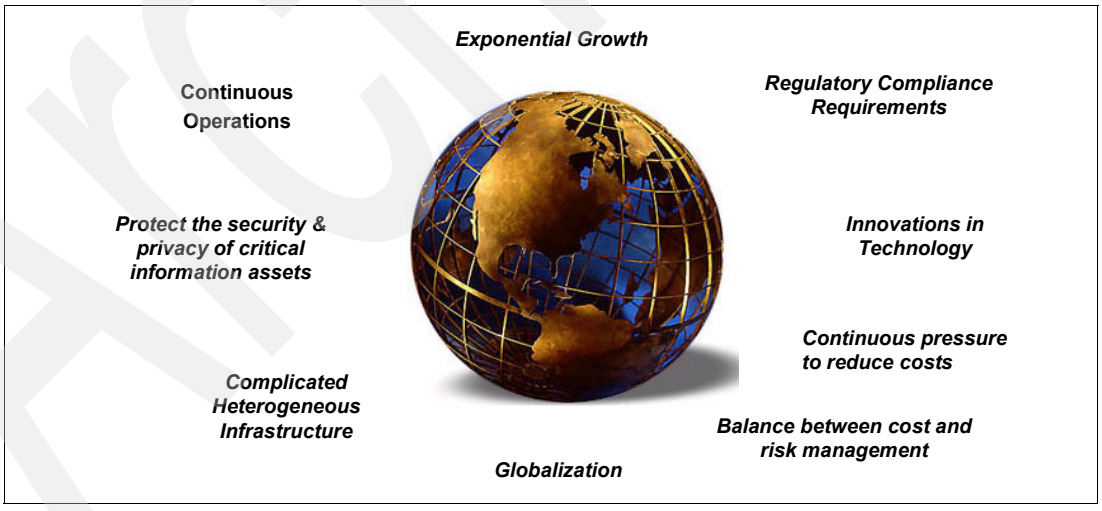

Figure 1-4 Factors that shape recovery and availability strategies

The following factors can impact recovery strategies:

- The exponential growth in the amount of data that must be stored has created the need for a comprehensive program to manage that growth. Data, like any other organizational asset (such as financial, physical, or personnel) must be managed with the same due diligence. The storage management discipline is concerned with managing application

requirements, forecasting future growth needs (capacity planning), performance and tuning, and (our concern in this book) the recovery and availability of data.

- The rate of technological change enormously complicates the recovery planning process. This accelerated rate of technological innovation requires companies to invest heavily in information-based solutions to business processes or run the risk of becoming noncompetitive and losing market share. It seems that almost daily, advances in the networking arena require organizations to address: connectivity issues; Virtual Private Network deployments; faster, more efficient switching technologies; and demand for larger and larger bandwidths. Rapid change and adjustments further complicate staying online or being able to come back online. Typical change rates for processor speed, network bandwidth, and storage density are to double about every two years, three years, and one year, respectively.
- Closely related to the previous point, the mixture of disparate platforms, operating systems, and communication protocols found within most organizations intensifies the already complex task of recovering/preserving data. Today, management must have reliable methods of recovering not only the mainframe, but perhaps multiple flavors of UNIX®, Windows®, Linux®, and IBM System i™, not to mention the network infrastructure. It is not surprising that server and storage consolidation has come to the forefront in many organizations' strategic IT planning, not only to more effectively manage day-to-day operations but to also ease the difficulties of recovery planning.
- A number of constituencies (customers, suppliers, management, and others) expect (in other words, demand) that data be available at anytime from anywhere. Web-based applications, for many, have to be available without exception. Revenue and the preservation of corporate image require it. The explosive growth of the Internet has drastically altered the traditional mechanisms of the marketplace, and have created the need for Web site redirection and load-balancing as availability assurance methods.
- The concept of what is a threat has also changed and now must include terrorism, vandalism from outside as well as from within the organization, industrial espionage, compromised privacy, advanced viruses which can corrupt and modify mission critical data, and fraud. This list is, unfortunately, not all-inclusive, and each can threaten your enterprise's survivability.
- New regulatory requirements have forced many to rethink their approach to data survivability. The Health Insurance Portability and Accountability Act (HIPAA) is an example of a requirement that determines how an entire industry, the US health care industry, must handle and account for patient-related data.
- Globalization of a firm creates new opportunities for some firms, but it can also create additional recovery and availability concerns that did not exist in the centralized computing model. In other words, data can be in any part of the world and different countries have different regulatory guidelines and processes to adhere to, to prevent and to plan for Disaster Recovery. This is particularly true for financial institutions.

1.6 Chapter overview

This chapter provides an overview of the concept of Business Continuity and provides some background on the topics that we discuss in this IBM Redbook. The remainder of this book describes various planning concepts for Business Continuity as they relate to the issues of data availability and recovery. The other chapters in this book include:

- Chapter 2, “Industry Business Continuity trends and directions” on page 9

  Describes some characteristics of the current industry environment.
- Chapter 3, “Business Continuity planning, processes, and execution” on page 43

  Discusses issues that each organization must weigh in formulating their recovery strategy.
- Chapter 4, “Tier levels of Business Continuity solutions” on page 137

  Describes the tier levels of disaster recovery solutions and the most effective way to select the optimum disaster recovery solution for your business.
- Chapter 5, “Business Continuity Solution Selection Methodology” on page 151

  Introduces a method for making sense of all the options available to you and helps you focus on the products that meet your needs and requirements.
- Chapter 6, “The Next Step Business Continuity workshop” on page 173
- Chapter 7, “Next Step Business Continuity workshop: Case Study” on page 233

  Show a step-by-step, applied example of this book's Business Continuity planning, design, and solution methodologies, plus a sample implementation of this methodology.
- Chapter 8, “Planning for Business Continuity in a heterogeneous IT environment” on page 251

  Covers a topic that very few IT sites can ignore today. It examines the concepts of effectively planning Disaster Recovery for a heterogeneous platform environment, and also details current IBM System Storage heterogeneous disaster recovery solutions and their concepts.
- Chapter 9, “Business Continuity for small and medium sized business” on page 273

  Discusses the Small and Medium Business considerations for Business Continuity.
- Chapter 10, “Networking and inter-site connectivity options” on page 281

  Discusses the often overlooked aspects of storage networking for data mirroring planning.
- Chapter 11, “High Availability clusters and database applications” on page 305

  Addresses high availability options on the technical and logical level for databases and applications. These techniques are discussed in the context of the Methodology that we introduce in Chapter 5, “Business Continuity Solution Selection Methodology” on page 151.
- The appendixes include:
  - Appendix A, “Business Continuity Solution Selection Methodology matrixes” on page 329
  - Appendix B, “Terms and definitions” on page 357
  - Appendix C, “Services and planning” on page 365
  - Appendix D, “Networking terminology tutorial” on page 379

2

Industry Business Continuity trends and directions

We live in a dynamic, constantly changing world, and the art and science of Business Continuity continues to change at the same pace. In this chapter, we explore the trends and emerging dynamics that will impact the Business Continuity solution designs of the future. These trends include:

- Justifying the cost of Business Continuity to the business
- Emerging requirements for Business Continuity in an On Demand business world
- Trends in multi-site data center architectures—out of region distances with greater than two sites
- A new data center strategy paradigm—planned workload rotation

2.1 New challenges

Any serious discussion of the future raises a wide range of issues, including the confluence of business and industry developments, consolidation, regulation, industry and business specialization, risk management, changing customer and workforce needs, and emerging technologies. All of these trends place pressure on the IT operating models and raise questions about how we must adapt existing strategies to meet objectives for growth, scalability, service levels, expense management, and added value creation.

Influencing factors from the business world include:

- **Customers continually are redefining the rules of the game:**

  Pronounced shifts in customer and marketplace demographics, attitudes and behaviors, combined with much more ubiquitous information, are giving customers and users the power to demand much greater responsiveness and functionality

- **Large players and focused niche players are squeezing the middle:**

  Large consolidated players are seeking to generate higher growth and revenue through economies of scale. Simultaneously, niche players are aggressively pursuing the most desirable customers by addressing their needs in distinct ways. Businesses and organizations in the middle are being squeezed.

- **Changing workforce and staff composition:**

  An increasingly mobile and diverse workforce is spread across multiple age groups, disciplines, experiences; and even across geographies and cultures. This can raise management complexity and requires flexible, effective approaches to skills availability, skills continuity, skills transfer, and costs.

- **Regulatory burdens intensify:**

  Heightened requirements around privacy, security, partnership risk, and operational risk will require a more proactive enterprise-wide approach to managing compliance and risk issues.

- **Technology improves inexorably to enable breakaway value:**

  Advancing technologies continue to enable customers and users to demand unprecedented offerings and functionality. Emerging IT technologies, such as grid computing, virtualization of data and storage, and predictive intelligence, will cause many entrenched IT philosophies to be forced to change, in favor of a partnership model where specialized skills, specialized organizations, and new relationships are required to survive and thrive.

2.1.1 Shift of focus from technology to business processes

As a result of these new challenges, Business Continuity has taken on new meanings. Business Continuity has traditionally been assigned to specific Business Continuity planners, often at the operational or IT level, with funding levels being set accordingly. This assignment was acceptable as long as information volume and scale was manageable, the span of required recovery was smaller, marketplace pressures and speeds were lower, competition was less rigorous, and the accepted expectation of service was not as high.

Today, however, the presence and interdependence of multiple business processes and processes, spread across multiple lines of business, technology, and staffing has created a paradigm shift. Thus, the IT Business Continuity focus has shifted from an IT technology to a full business process focus, as shown in Figure 2-1.

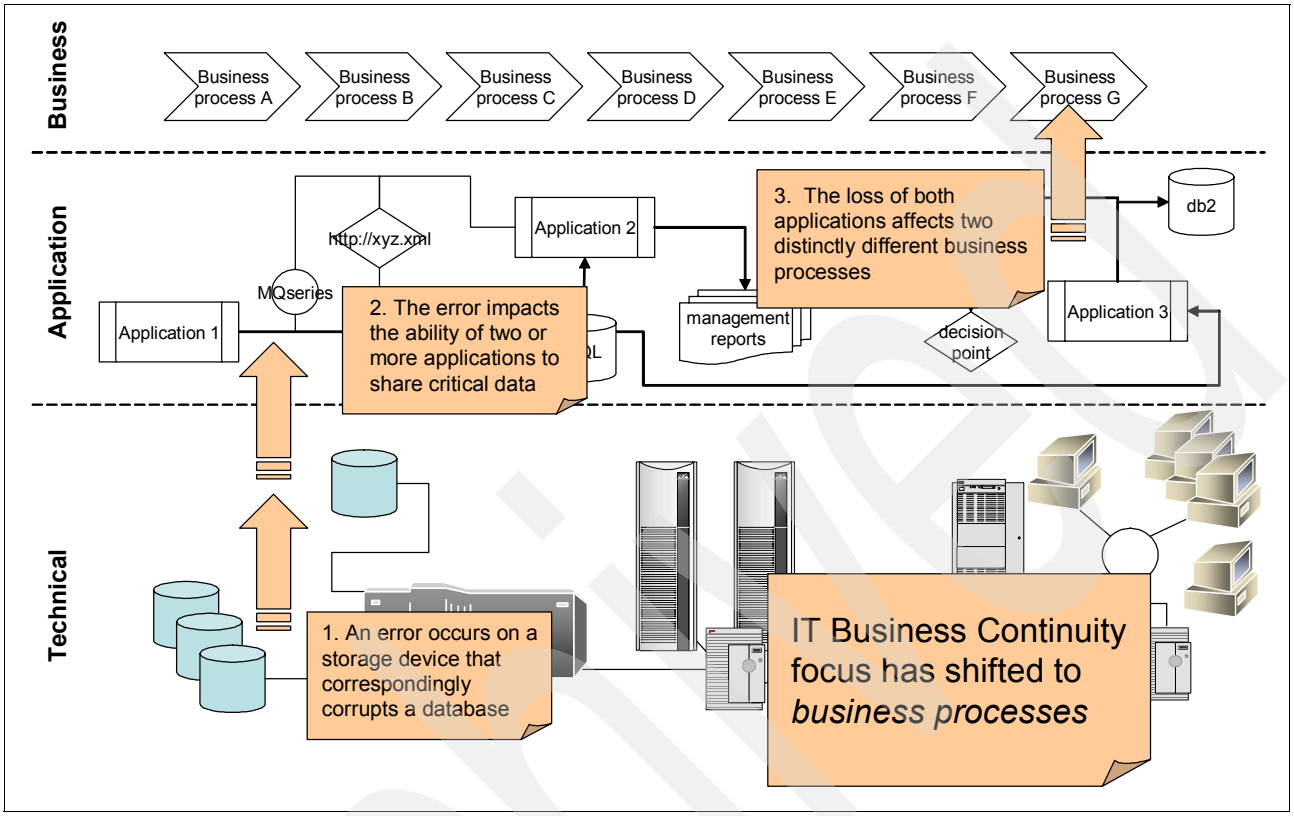

Figure 2-1 The interdependent nature of business processes: An IT outage is also a supply chain impact

With these kinds of interlocked business processes, achieving true IT Business Continuity today must address a significant internal organizational challenge successfully.

2.1.2 An organizational challenge

Typically, there is no one central point of organization and control for all of these business processes. Yet, all interested parties agree that a contingency plan is necessary. Therefore, the challenge can be stated as follows:

- Who is or who should be the actual owner (or owners) of the end-to-end Business Continuity plan?
- What is the appropriate scope and definition of the Business Continuity plan?
- How much IT Business Continuity preventative *functional insurance* is the right amount?
- Who will pay for this *insurance* protection?

We often see today that the IT department, with the Chief Information Officer (CIO) as the responsible party, is seen as the business organization that is best positioned to fulfill the compliance, security, and Business Continuity requirements of the business. In the end, although it is the Lines of Business that have the ultimate risk, IT is often held responsible, although IT is in the support role.

Thus, our discussion begins with how IT might best interlock with the Line of Business (LOB) and gain agreement that the LOB and IT must be in *joint ownership* of an end-to-end Business Continuity plan. We also discuss how IT might best develop a strategy to gain this support from the LOB.

We start by addressing the best ways to justify Business Continuity to the business.

2.2 Justifying Business Continuity to the business

In today's environment, innovation takes many forms, including advances in products and services, markets, operational processes, customer intimacy, new testing, qualification, and delivery strategies. However, innovation is not possible and cannot have the desired impact unless IT organizations can create the requisite conditions for IT innovation development.

In this section, we explore the trends that affect and enable IT innovation development and focus on best practices for *justifying and developing a strategy for IT Business Continuity technology*. Given the expanded scope of IT Business Continuity at the business process level, we provide an overview of principles that describe how you can best position IT Business Continuity technology within the strategic IT funding and decision making process. We also discuss how you can best introduce and integrate IT Business Continuity technology into the overall IT initiatives.

2.2.1 Executive summary

When justifying Business Continuity, we recommend that you not consider it as an end unto itself. Rather, IT Business Continuity is best accomplished when it is included as an *intended by-product* of the overall IT infrastructure streamlining, consolidation, and efficiency strategy.

You can accomplish this goal by making IT infrastructure simplification a necessary prerequisite to cost-effective IT Business Continuity. When this level of strategy integration is done, enacting the IT infrastructure streamlining efforts also produces the necessary IT Business Continuity, security, risk compliance, and audit functions that the business requires.

In the sections that follow, we discuss how you can successfully develop a strategy and implement these recommendations.

2.2.2 Relationship of time to market to Business Continuity

We begin by describing a model of the inter-relationship of business objectives to the IT infrastructure and the IT Data Center strategy.[1] There are two business metrics for this model:

- Time to Market

  A measure of how fast the business is able to deliver new products and services to market. Making Time to Market as fast as possible is essential.

- Time to Customer

  A measure of, from the customer's perspective, how good is the product and service availability, performance, and all other relevant factors.

These metrics compete with Total Cost of Ownership, as shown in Figure 2-2.

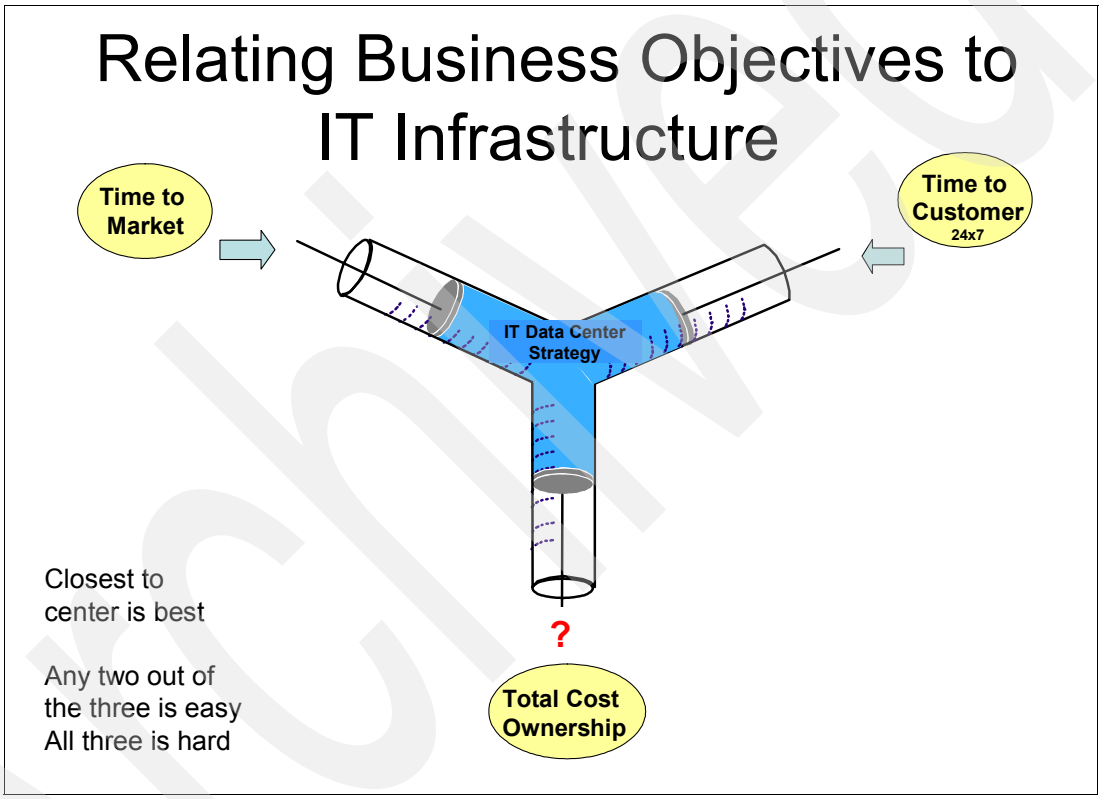

Figure 2-2 Relationship of Time to Market, Time to Customer, and Total Cost of Ownership

Think of this relationship as an example from fluid dynamics, and imagine a liquid inside the three-pronged plunger (or piston). The legs of the plunger represent:

- Time to Market (TTM)
- Time to Customer (TTC), that is application uptime and availability
- Total Cost of Ownership (TCO)

These three legs of the business (TTM, TTC, and TCO) are, of course, inter-related and cannot be truly be considered independently from one another.

[1] The concepts and principles of this section and the subsequent sections, 2.2.3, "The struggle between TTM and TCO" and 2.2.5, "New importance of IT standards and integration", are used with the express permission of the original author, John Koberlein, retired IBM Senior IT Architect and now CEO of Kobe Consulting.

The optimum point of this model would be at the zero point in the middle of the three plungers, that is infinitely fast (zero) time to market, infinitely available applications and services to customers (zero downtime), and zero cost. Obviously, IT technology does not yet have the capability to provide zero cost while providing infinitely fast application delivery (TTM) or infinitely available data, information, and applications (TTC). Yet, it is the business goal to drive each of these three plungers as close to zero as resources allow.

If we assume for the moment that the liquid (that is, the IT Data Center Strategy) in the plunger is not compressible, then to only focus on reducing or pressing inward on TCO, without considering the resulting impact on either TTM or TTC, might mean the organization is unintentionally causing TTM to expand rather than contract or causing TTC to get worse instead of better. This model conveys the correct impression that if nothing is done differently in terms of the IT Data Center infrastructure and strategy, the effect of increased pressure on TTM and TTC will be for TCO to go *out* (that is, increase), rather than moving inward. Clearly, this outcome is not the desired strategic outcome. Thus, focusing on TCO alone is not the best business perspective.

What this model also suggests is that the defining factor in this equation (that is, the liquid) is the IT Data Center strategy. Most importantly, the *manner* in which TTM, TTC, and TCO interact is defined *through the IT Data Center strategy*. Therefore, to achieve the goals of simultaneously compressing these three metrics, it is necessary to change the characteristics (that is the ability to compress the liquid) of the IT Data Center Strategy in a way that allows this three-way compression to occur.

Let us next examine how we might best modify the IT Data Center strategy such that it enables faster TTM and better TTC and, simultaneously, enable lower TCO and justifiable IT Business Continuity funded projects.

2.2.3 The struggle between TTM and TCO

Given that we realize we need to modify the IT Data Center strategy such that we can compress TTM, TTC, and TCO while delivering IT Business Continuity, what IT Data Center architecture and strategy initiatives should we pursue, and how should we interlock those initiative with all the other existing IT projects that also must be done? Because senior management is the real source of strategic funding for Business Continuity, you need to present Business Continuity funding requests in a way that allows senior management to achieve business targets of TTM, TTC, and TCO simultaneously.

Figure 2-3 illustrates the competing nature of TTM and TTC with TCO and IT Business Continuity.

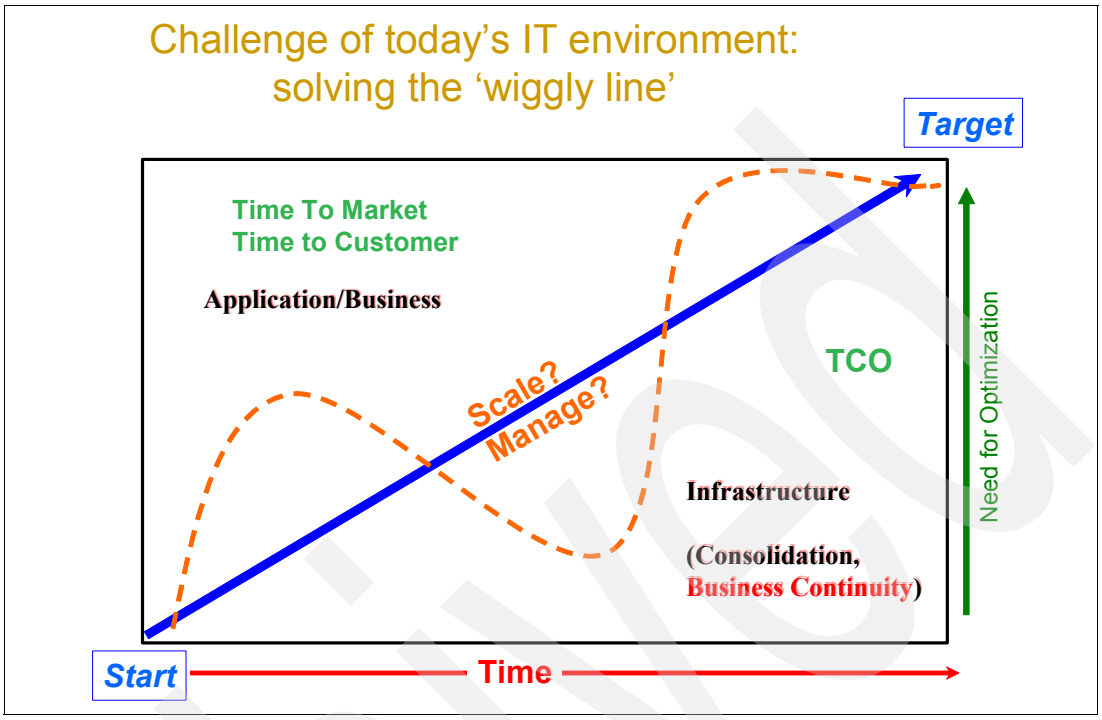

Figure 2-3 The competing nature of TTM, TTC and TCO (or the "wiggly line")

In this figure, note that TTM and TTC are on the upper, left side of the figure, on the application and business side of the Start to Target line.TTM and TTC both attempt to increase the gradient of the Start to Target line (that is, to do more in even less time).

On the opposite side of the equation is TCO, of which Business Continuity is just one cost component. If TCO is viewed as a cost metric, that is a necessary cost to be expended in pursuit of TTM and TTC, then TCO tends to be squeezed and thus tends to drag the Start to Target line downward and to the right (that is, doing less and taking longer to do it).

Because TTM and TTC are on the opposite sides of the Start to Target equation line from TCO, there is a competition between the two sides. Thus, it is not surprising that there is a back forth (the "wiggly line) in an attempt to balance these two competing sides. In such an environment with IT Business Continuity being just one component of IT TCO, IT Business Continuity can get lost in the shuffle. Thus, adequate IT Business Continuity funding is typically a challenge.

2.2.4 The struggle of IT diversity

We must also recognize that today's IT environments have a major compounding problem, especially when it comes to IT Business Continuity. That problem is the diversity and complexity of today's IT environments, and that IT diversity is much more than just an IT Business Continuity problem.

Figure 2-4 illustrates an example of an IT environment where it would be difficult to provide fast TTM, TTC, TCO, and IT Business Continuity due to the large number of components and associated difficulties of managing them all.

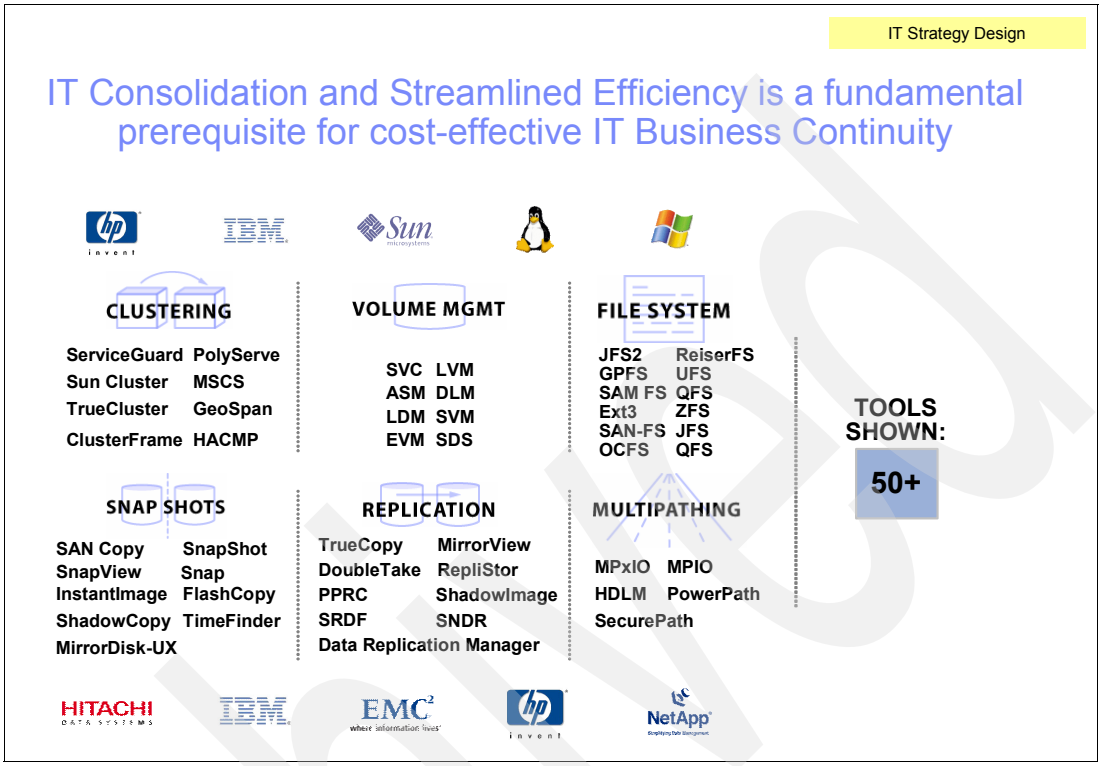

Figure 2-4 A difficult environment to manage and recover due to high diversity

In this environment, it is not that any individual product is at fault. It is just that in a highly diverse and complex environment, there are too many things to recover and too many skills and personnel that must be spread across too many tools. Also in this environment, there is limited ability to share resources and thus provide economies of scale. Coordination is not easily synchronized.

So, how can we develop a strategy to address as a whole TTM, TTC, TCO, and IT Business Continuity issues?

2.2.5 New importance of IT standards and integration

Developing a strategy to address these issues is an age-old IT answer. However, the answer takes on new importance and criticality in today's high scalability, fast paced, rigorous environments that require Business Continuity. To develop a successful strategy, businesses today must accelerate the drive towards *IT standardization and IT integration* throughout the enterprise (as illustrated in Figure 2-5).

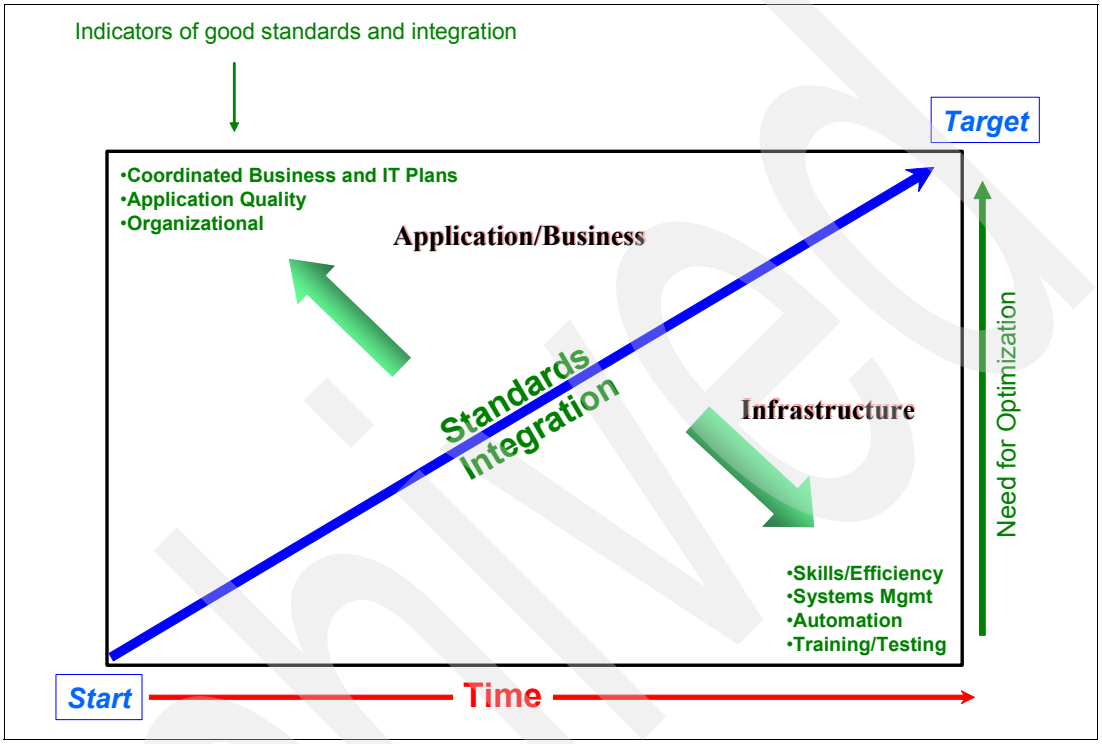

Figure 2-5 IT standardization and integration results

By *standards*, what we mean is *IT standards for your IT department and IT enterprise*. In order to achieve TTM, TTC, TCO, and IT Business Continuity simultaneously, standardization throughout the enterprise has become of paramount importance. Standardization in this manner makes good business sense. As shown in Figure 2-5, the *indicators* of standardization are:

- Coordinated business and IT plans
- Application quality
- Organizational characteristics and responsiveness
- Skills and operational efficiency
- Systems management
- Exploitation of automation
- Resources for training and testing

When these indicators are present, they are proof that sufficient standardization is in place. Conversely, implementing sufficient standardization can assure that each of these indicators are the result.

With a requisite level of IT standardization and simplification in place, your business can make major progress on achieving TTM, TTC, TCO, and IT Business Continuity all at the same time.

Well planned, well governed standards such as these are also the pre-requisites that IT Business Continuity requires. IT Business Continuity demands:

- Coordination of business and IT plans between multiple organizations.
- A smooth team-to-team flow of application quality, application defect fix, and application and business process interlock.
- A high level of interaction between development and infrastructure when rolling out changes.
- A good level of data management, change control, and systems management.

When developing standardization strategies, you must architect two major areas of IT standardization:

- Standardized methodologies in the IT shop
- Appropriate standardization of IT components

We look a little closer at each of these two ideas here. You can find more information about planning and implementing IT standards in Chapter 3, “Business Continuity planning, processes, and execution” on page 43.

Standardized methodologies in the IT department

Standardized methodologies refers to architecting and standardizing how the various IT processes and procedures are done across the enterprise. Standardization opportunities include how backup and restore are done as well as the established standards for backup and restore of software and hardware. Other standardization examples in the process and procedure area might be:

- How you push critical patches out to desktop users
- What support methodologies are employed
- How inter-departmental requests such as a software development team needing new IT infrastructure

Today’s highly-diverse IT environments demands such standardization, as with high TTM and competitive pressures, there is no room for error, no time to waste, and a need to minimize cost. Although a very old IT idea, with today’s pressures of TTM and TCO, standardized methodologies throughout the enterprise take on new meanings to ensure efficiency by minimizing duplication of effort, ensuring higher quality results through consistency, and lowering budgets and maintenance costs in both human cost and in supporting infrastructure terms.

Appropriate standardization of the IT components

It follows that standardized IT components will drive similar cost savings and quality gains. An example might be an enterprise-wide architectural move towards standardization on a single database vendor, thus leading to better code sharing among application developers and a reduction in test workload, with a faster application development cycle. Contrast this kind of streamlined environment with a more difficult heterogeneous environment supporting multiple types of underlying components in hardware, software, and middleware.

Thus, we conclude that cost-justifiable, cost-efficient IT Business Continuity must be built upon the prerequisite of a properly standardized IT environment.

2.2.6 Infrastructure simplification as a prerequisite to IT Business Continuity

We suggest, therefore, that an essential connection must exist between the IT infrastructure simplification and consolidation, as the appropriate and proper prerequisite to IT Business Continuity, as illustrated in Figure 2-6.

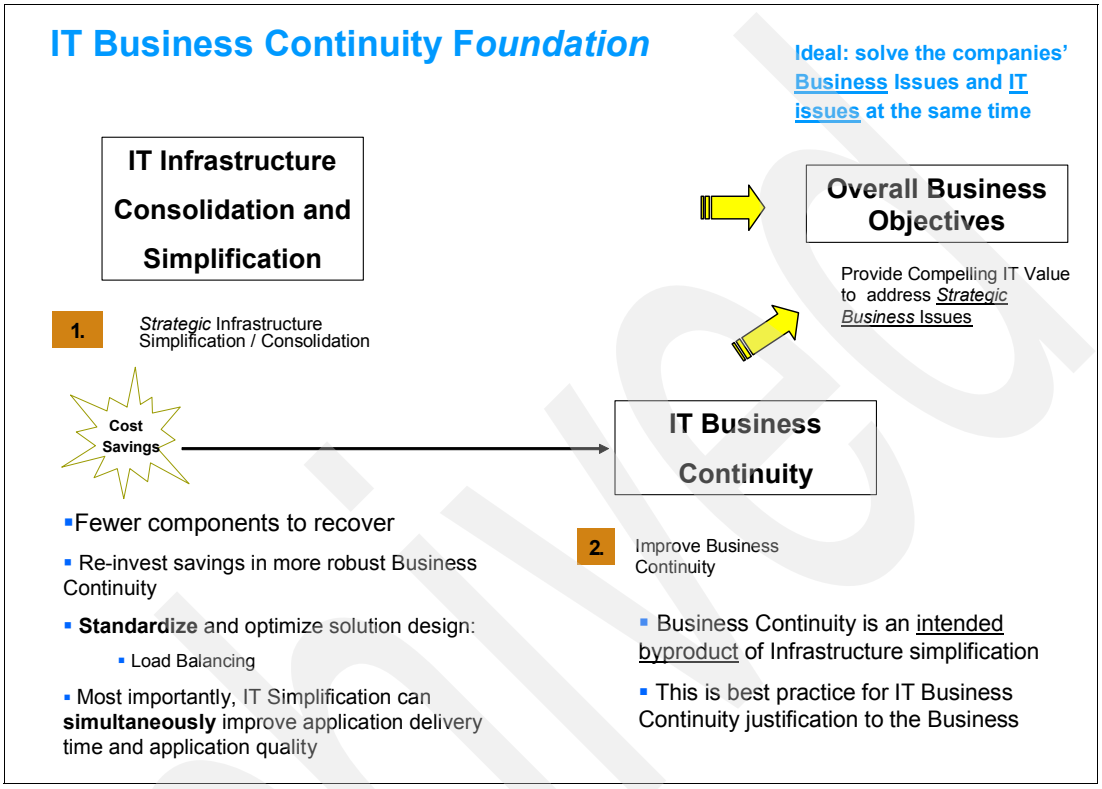

Figure 2-6 IT infrastructure simplification and consolidation as prerequisite to IT Business Continuity

We suggest that when discussing and planning IT Business Continuity, that the conversation always include the key thought that IT Business Continuity is based on an enhanced version of the existing IT consolidation and simplification efforts that are most likely already active in your IT organization. Acceleration of IT consolidation, standardization, and integration of IT methods and components are all ideal objectives in terms of reducing operating cost, TTM and service quality and will also have a positive impact on lowering the cost and speeding the effectiveness of IT Business Continuity functionality.

We suggest that you a solid foundation for solving the IT Business Continuity issue by accelerating and architecting IT consolidation, standardization, and simplification efforts in a way that produces IT Business Continuity as an intended by-product. Some major talking points when presenting this strategy to higher management include:

- It is easier to manage and recover fewer IT entities (for example, servers, operating systems, or applications) than more.
- Cost savings from IT consolidation should be re-invested into the IT Business Continuity expense. Linking cost savings to re-investment can help defray necessary capital and running costs in an IT Business Continuity implementation.
- Given that Business Continuity will be part of your overall IT infrastructure simplification, fold in Business Continuity implementation-related tasks, such as workload balancing, configuration changes, and other adaptation work, into an enhanced version of the IT infrastructure simplification efforts that are already under way.

These discussion points tie IT Business Continuity improvements as an intended by-product of other IT infrastructure improvements that are directly aimed at benefitting the LOB. The LOB executives thus receive the assurance that even as they fund IT to implement new applications and streamline existing business processes, that IT is simultaneously implementing Business Continuity as an intended by-product of that IT simplification.

A vision not a retrofit

We do acknowledge that it might be difficult to retrofit such IT standardization to an existing highly diverse IT environment, and we are not necessarily recommending that you retrofit the entire existing infrastructure. Rather, we recommend that you focus on architecting the *future consolidated, streamlined* IT data center strategy. It is normal that this type of effort will take some time. Thus, you should not attempt all IT Business Continuity projects at the same time. There is a normal time progression in an IT Business Continuity, which we discussed in 3.2, "Typical evolution of a Business Continuity program" on page 44.

These ideas represent an ideal to which you can point your strategic IT vision. By setting this kind of vision and aligning your decisions in that direction, you can improve and maximize the impact of your IT data center strategies. Most importantly, by tying in IT Business Continuity into the large IT infrastructure simplification efforts that enable increases in revenue, service quality, and long-term cost savings, you can drive the business "fluid" dynamics (as illustrated in Figure 2-2 on page 13) closer to that ideal zero point over time.

2.2.7 Value of a combined IT simplification and IT Business Continuity strategy

With this kind of a strategy, IT Business Continuity is raised to the higher level of a *business tool* because it is responsive, flexible, and competitive. In effect, by addressing TTM, TTC, and TCO with the IT simplification, we provide the following business values with the IT Business Continuity solution simultaneously:[2]

- **Provide 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).
- **Provide 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.
- **Provide security, privacy, and data protection:** Help the business protect against threats, both internal and external, and develop a critical information management strategy.
- **Maintaining market readiness:** Helps the business stay competitive by anticipating and quickly responding and 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 business partner quickly and effectively within their industry by becoming a trusted and reliable business partner in their supply chain or value-net.

[2] These business points are used with the permission of the Fact Point Group, which created these points for use by IBM under contract with IBM Global Services.

Strategically, the factors above should be highlighted to senior management as ways that IT infrastructure simplification and IT Business Continuity are important strategic factors in sustainable growth, better profitability, and increased shareholder value.

2.3 Emerging criteria for IT Business Continuity solutions

For any IT Business Continuity solution, regardless of the type of technology or vendor, today's Business Continuity solution criteria have evolved to a new level, as shown in Figure 2-7. These are the business criteria that are fundamental to satisfying the IT Business Continuity solution needs of today's accelerated business environment.

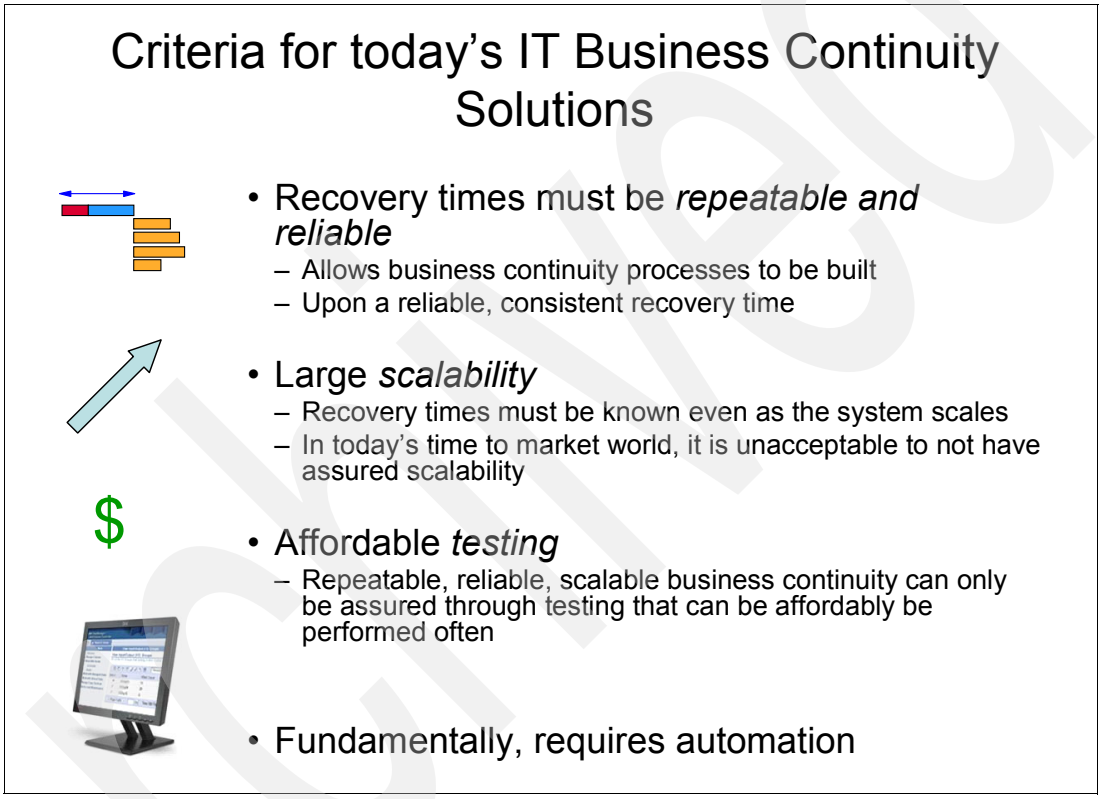

Figure 2-7 Today's IT Business Continuity solution requirements

In the sections that follow, we examine each of these criteria.

2.3.1 Reliability and repeatability

While it seems to be stating the obvious, it is imperative today that recovery times must be repeatable and reliable to a much more rigorous, consistent level.

What we mean by repeatable and reliable recovery times is that regardless of the situation or circumstance, IT needs to be able to deliver an IT Business Continuity solution that provides a very reliable, repeatable, and consistent *Recovery Time Objective* (RTO). Today's best practice is to be able to deliver IT Business Continuity solutions that can reliably, within ideally less than 10% variation, repeatedly deliver a consistent RTO. When this consistency is in place, the end-to-end business processes can then architect themselves for a very specific time window, thus maximizing benefits for a given amount of budget.

If a Business Continuity solution varies too much in recovery time, the net result is that the business eventually has to plan on (and spend the money for) the higher end of the tested recovery times. In that case, the business cannot truly take advantage of the cost advantages when the solution happens to deliver a shorter recovery time. In effect, the business spends extra money to accommodate the longer recovery times and does not necessarily receive the direct cost benefit of the shorter recovery times.

2.3.2 Scalability

An annual growth rate of 35% to 50% in open systems storage is not unusual in many of today's well-managed IT organizations. Therefore, as the business grows and scales at today's large storage growth rates, the business must be able to predict and prove this level of IT Business Continuity reliability and repeatability. It is not viable to have a Business Continuity solution that at some point must be re-architected because of its inability to scale. Especially in today's world with acquisitions and new marketplaces, dramatic workload increases are becoming an occurrence that must be accommodated in order for the business to achieve its economy of scale and cost objectives.

2.3.3 Affordable testability

Reliability, repeatability, and scalability cannot be proven or refined and tune unless we have the ability to perform very frequent testing at a very affordable cost. Therefore, a focus item in the IT Data Center strategy must be methods for performing the testing cycle on a larger scale and on a more frequent basis (ideally nearly continuous testing).

Tip: We discuss some key additional thoughts about providing this affordable testability in 2.5.5, "Planned workload rotation data center model" on page 31.

2.3.4 A fundamental need for automation

In order to meet these objectives within staffing, budget, and time constraints that exist today and continuing into the future, affordable continuous testing can no longer require intensive labor or time away from other implementation tasks. Affordability of frequent testing therefore demands a high degree of automation designed into the IT Business Continuity solution from the start.

Good automation for IT Business Continuity solutions provides the ideal foundation for making testing affordable, by providing the ability for repetitive testing run largely by automation. Automation also has a significant positive effect on the reliability and repeatability of the Business Continuity solution, and it makes it easier and more viable to test for scalability. For all of these reasons, automation is becoming a strategic Business Continuity solution foundation for affordably providing IT Business Continuity that can meet the above business objectives in the years ahead.

2.3.5 Emerging criteria for the human element

The importance of the human element—that is the role, responsibilities, and reaction of people in any Business Continuity plan—has always been high.

Tip: We discuss planning for the human element in Business Continuity in detail in 3.5.1, "Business Continuity program design" on page 67.

World events of the past few years have identified the following additional emerging considerations for the human element. Note that both of these events have a time duration of perhaps two to six months (or more), which is longer than the typical time duration normally planned for with traditional outage scenarios:

- **Pandemic considerations**: Planning should consider the increased possibility of an pandemic event,[3] in which the primary impact is to human resources and there is necessity for humanitarian care and protection. Note, however, the IT planning consideration for a pandemic event: even though the IT infrastructure is intact, the IT infrastructure will deteriorate over time if there is inadequate human resource availability to maintain or perform required updates or problem resolution on the IT infrastructure.
- **Large scale displacement considerations**: Tragic natural disasters such as the December 2004 Asian Tsunami and August 2005 Hurricane Katrina highlighted the need to consider and plan for the effects of a large scale, long term displacement of large population of people. Many basic necessities such as electrical power, fuel, housing, and transportation might be problematical. Even if available, the pricing for these necessities might be abnormally high.

Fortunately, emerging best practice solution approaches to these valid concerns include a strong proactive exploitation of remote access capability and virtual workplace continuity.

2.3.6 Virtual workplace continuity

In the 21st century, the virtualization of the workplace has increased using teleconferences, high speed telecom links (including wireless), voice over IP, and other modern IT tools. Driving factors have included greater needs for business and human resource collaboration, flexibility, cost pressures, competition, and TTM and TTC goals.

From a Business Continuity standpoint, the emerging trend is to actively promote and *exploit* this continuing virtualization of the IT infrastructure. As the workplace is evolving into a virtual workplace, there is an opportunity is to exploit this trend by designing Business Continuity plans, procedures, and IT functionality that proactively equip key human resources with the ability to connect, communicate, and collaborate work from anywhere. Emerging best practices include purposeful design and implementation of Business Continuity remote access, security, and firewalls as well as exploitation of mobility technology, such as having key personnel always carry pre-loaded USB storage devices that can be plugged into any PC to provide secure remote access to the IT infrastructure from anywhere.

Assuring virtual workplace continuity is an emerging and important area of Business Continuity practice for the future. The methods that you choose for your virtual workplace are thus another standardization opportunity for your IT standardization and consolidation strategy, as discussed in 2.2.5, “New importance of IT standards and integration” on page 17 and 2.2.6, “Infrastructure simplification as a prerequisite to IT Business Continuity” on page 19.

[3] The definition of a *pandemic* event is an epidemic (an outbreak of an infectious disease) that spreads across a large region.

2.4 Out of region recovery

In this section, we discuss the ramifications of the growing trend for out of region remote recovery. This discussion then leads to a broader discussion of multi-site data center strategies.

2.4.1 Considerations for distance of to recovery site

Metropolitan distance recovery and out of region recovery are both viable strategies. It is a matter of selecting the best fit for the business and regulatory requirements. Out of region technology is available and implementations are common today. By setting expectations properly as to what is different about out of region recovery versus metropolitan distance recovery, we can assure that the business can simultaneously meet its expectations, its budget, and its IT Business Continuity objectives.

Clearly, as both the business world and the world at large have changed, out of region recovery has become a preferred distance trend in remote data center recovery. The specific distance that is defined to be *out of region* varies by geography. What is important to note is there are additional planning considerations for out of region recovery, that are not present in metropolitan distance recovery.

Major factors to consider for out of region recovery

Here are the factors that you should be aware of regarding out of region recovery:

- There is likely to be an increased staffing and infrastructure cost factor to support the increased distance, because it is likely that there will be some amount of staff duplication required to operate the two geographically distant sites.
- The likelihood of requiring asynchronous data replication is high, because out-of-region distances typically are out of synchronous replication range. Asynchronous replication introduces a larger scope of work effort, capacity planning, and implementation planning for the Business Continuity project compared to synchronous mirroring. So, you need to set time frame expectations and work effort scope accordingly.
- Out of region recovery implies a fast shutdown and a fast restart at the primary site, instead of a *hot failover*. Near continuous availability is unfortunately very difficult or impossible to achieve when switching data centers at out of region distances, due to the fact that there is a lag in data currency at the asynchronous remote site. If the organization has already implemented data center hot failover today at metropolitan distances and wants to retain this capability, then moving to out of region recovery requires either a relaxation of the cutover RTO or potentially a move to a three site configuration.
- Compared to metropolitan distance recovery, continuous replication of data for out of region longer recovery distances might actually increase the ongoing monthly cost of bandwidth for a given amount of MBps. This is because telecom line costs typically have a distance multiplier.
- We project that we will continue to see that bandwidth costs be the limiting factor as to the amount of data that can affordably be replicated, and therefore bandwidth is the delimiter of the *scale* of recovery that can be accomplished at out of region distance.
- In order to accommodate the fact that bandwidth delimits the amount of data that can be transmitted, a segmentation of the business processes is necessary in order to allocate only those business processes' data that requires a high level of remote recovery to the out of region recovery solution.

In summary, compared to metropolitan distance recovery, out of region recovery failover time is likely to be longer, and the overall cost factors of operating the alternate site at out of region

distances will likely be higher. In return, the client is able to have geographic distance separation.

Major factors to consider for metropolitan distance recovery

For comparison, the considerations for metropolitan distance recovery include:

- Often chosen when the smallest possible data or transaction loss at the remote site is desired, because the remote data center is within synchronous data replication range.
- Metropolitan distances might be close enough to support clustering business processes, or IT infrastructure such as server and database memory, thus providing more robust possibilities for improving the business process or IT recovery time capability.
- A single staff might be able to support both metropolitan sites in more scenarios, thus potentially minimizing logistical considerations and reducing cost.
- Intersite communications and telecom line costs for equivalent bandwidth, can typically be less expensive at metropolitan distances compared to out of region distances.
- For true near-continuous availability at the IT data center level, metropolitan distances are the typically required, because IT data center near-continuous availability failover often requires server clustering and database clustering, in addition to synchronous data replication. Server memory clustering and database memory clustering cannot be done at out of region distances.
- Recovery at metropolitan distances might be less expensive overall than comparable data center recovery at out of region distances. All other factors being equal, the recovery time at metropolitan distances is also typically shorter than a comparable out of region distance recovery.

In summary, compared to out of region recovery, metropolitan distance recovery failover time is likely to be faster, and the overall cost factors of operating the alternate site at metropolitan distances will likely be lower. At the same time, the client does not have geographic distance separation.

Tip: We provide a more detailed discussion about considerations for distance to remote site in the “Distance to remote site” on page 87. In addition, we discuss considerations for synchronous data replication versus asynchronous data replication is in “Considerations for selecting synchronous versus asynchronous data replication” on page 113.

Next, we discuss more detail on some other important factors that affect out of region recovery planning.

2.4.2 Out of region bandwidth costs and trends

There is an understandable perception that because broadband Internet lines have become so affordable and commonplace, that the same must be true of high MBps data center grade telecom lines. Unfortunately, this is not the case today. It is important to note that data center grade lines can have a noticeably higher unit cost per MBps than the *commodity* Internet lines.[4]

The reasons for the higher costs are as follows:

- You might think of the analogy of high speed data center grade lines as being similar to a high performance automobile, in that to reacher higher MBps bandwidth speeds with stability, control, and reliability, requires a significant amount of engineering sophistication and cost.

[4] That is, high MBps telecom lines with a high degree of reliability and uptime.

- This engineering and manufacturing cost is reflected in the higher MBps monthly costs for high speed data center grade telecom lines.

Because there is clearly a trend for more out of region recovery, note that there is also a distance factor in the telecom cost equation. For budgeting purposes, as the distance increases, there is a distance multiplier for the higher speed data center lines. In other words, for a given MBps bandwidth, the charges are more for long distance lines than for shorter distance lines.

From an overall trends and directions summary for telecom, we note that the percentage growth rate of high performance mission critical disk storage is higher than the percent decrease in bandwidth costs. Therefore, for the budgeting purposes, we recommend that you set the expectation that raw aggregate expenditure for bandwidth is expected to continue to rise, even though the cost of MBps is declining.

For all of these reasons, even though the unit cost per MBps of bandwidth is clearly decreasing, we project that telecom line costs will continue to be the major limiting cost factor for remote site recovery Business Continuity solution design.

Tip: We provide a more detailed discussion of telecom bandwidth considerations in “Intersite networking and telecom components” on page 108.

2.4.3 Essential need for business process segmentation

In light of these trends and costs in telecom bandwidth, we recommend that you plan on strategic business process segmentation. This plan is necessary in order to best enable a cost-effective Business Continuity architecture that can meet the needs for both out of region recovery and bandwidth requirements.

Tip: We provide a detailed discussion of the need for business process segmentation in “Need for Business Process segmentation” on page 96.

To briefly summarize, the purpose of this business process segmentation is to be able to *affordably* implement the following modern IT trends and directions:

- Tiered storage
- Tiered servers
- Information Life Cycle Management
- Out of region recovery, while keeping the bandwidth costs affordable
- Multi-site data center strategies

By strategically setting out to segment our business processes, we create a construct wherein we are more easily able to map the segments of our IT environment to the appropriate level of service, recovery level, and cost that is desired for that business process segment.

Figure 2-8 illustrates a suggested business process segmentation, which consists of three bands: Continuous Availability, Rapid Data Recovery, and Backup and Restore.

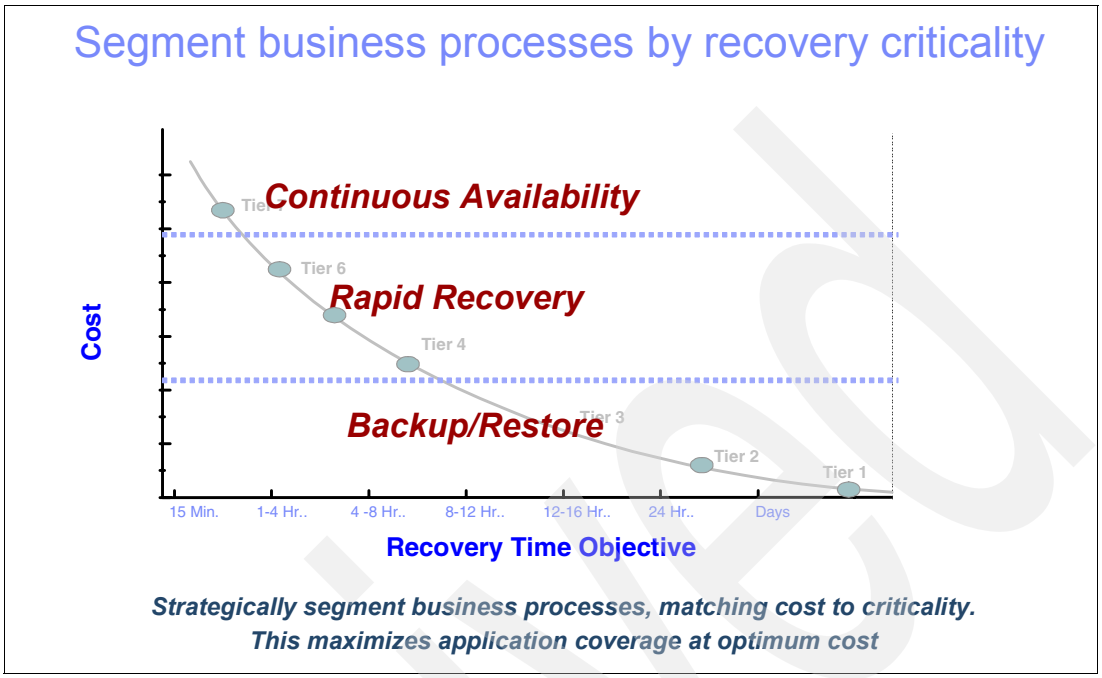

Figure 2-8 Business process segmentation

There are several reasons to implement this segmentation, most important of which is:

- No one set of IT Business Continuity technologies is right for all business processes.
- We need to apply fast recovery technologies to some business processes (accepting the higher cost along with the valuable faster recovery).
- We need to apply less expensive recovery technologies (which are lower in cost) to other business processes.

We recommend that you consider initiating plans to further refine, architect, and enhance your business process segmentation within your strategic Business Continuity planning.

2.5 Emerging data center strategies for multiple sites

Remote recovery data centers have existed for a long time. Today, however, the way in which those geographically remote data centers are used and exploited is rapidly changing. With all of the trends that we have discussed thus far, we now address directions in best practices for IT Data Center strategy.

We discuss methodologies to address new requirements for compliance, testing at a high scale, application quality, application testing, fast TTM and time to implementation, and so forth. We also answer the question: What are the strategic long-range data center strategies that will best support compliance, high availability, and Business Continuity at affordable, justifiable levels to the business data center strategy?

2.5.1 Traditional two data center model

Traditionally, a two-site data center topology was composed of a primary site and a secondary site (which was typically used for test and development), as shown in Figure 2-9.

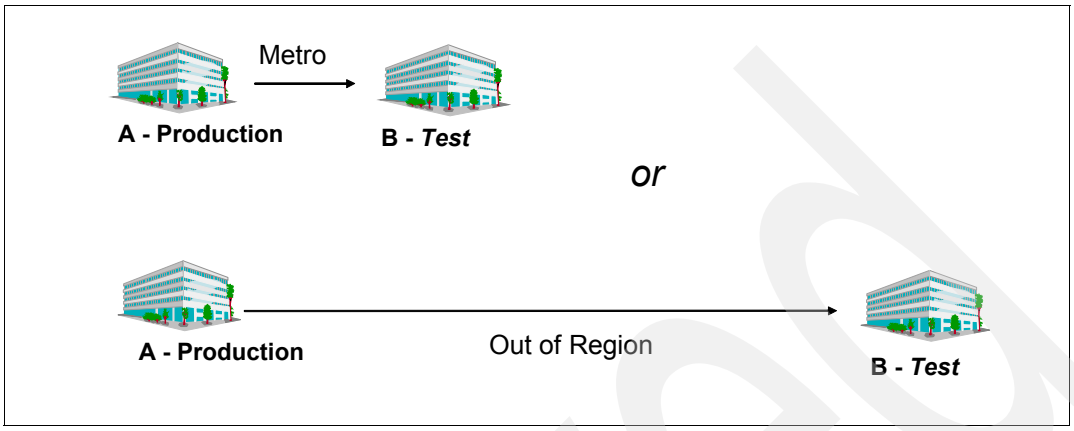

Figure 2-9 Traditional two-site data center model

These designations traditionally were fixed, and production operations were rarely shifted between the two centers (except for major disaster recovery tests). The majority of the time, operations ran in the primary data center. The secondary data center (which might be across campus or geographically separated from the primary) typically ran test and development.

Often, the secondary data center had previous generation components compared to the primary center. The thought was that the likelihood of using the secondary is small, and costs can be saved by taking on some additional risk in equipping the secondary data center with lesser hardware, software and so forth.

There is nothing inherently wrong with such a strategy, as long as it can provide adequate IT scalability and recovery of critical applications at the level that the business requires.

New requirements stress the traditional data center model

However, newer requirements related to TTM, TTC, and TCO, place pressure upon today's data centers due to significantly raised expectations. The traditional data center model (that is, the IT Data Center strategy) needs to be enhanced to address these issues.

A sample of the pressing issues increasingly asked by IT management include:

- What is my frequency of testing? Effort? Failback?

  When was my last full system-wide test or drill?

  How do I know that applications that have been put into production since the last test, how do I prove their full recovery capability, especially from the perspective of a full end-to-end test with all other applications.

  (These questions are more than IT questions. Today, they can become a concern from a business partner, the supply chain, or the customer standpoint.)

  How often is Failback (that is, restoration) from the recovery center back to the primary tested? Is it tested?

- What is my ability to move workloads at will?
- Am I compliant? What is my cost of compliance?

  Proof of regulatory compliance might be related to frequency of testing, especially system-wide testing.
- How is my Business Continuity strategy helping me improve IT Time to implementation?

  How is it helping me improve application quality, problem determination speed?

  How is it helping me handle issues like:

  Test and development resources may not be fully utilized on an ongoing basis. Furthermore, est and development typically does not have the resources to do full scale stress tests, full scale stress problem recreation, or large pools of extra resources for major projects.

  How is it helping me improve my skills and resource productivity?
- Is my Business Continuity strategy helping me do any of these new requirements?

In other words, businesses and IT management today are asking the questions:

- Am I getting full use out of all IT assets?
- Is there a better way to use my IT assets?

In the following sections, we discuss how the IT data center and Business Continuity models must evolve to meet these challenges.

2.5.2 Two-site high availability: Metro distances

Building upon the two-site traditional data center model, we see a rise in the implementation of high availability functionality between two metropolitan distance data centers, as shown in Figure 2-10.

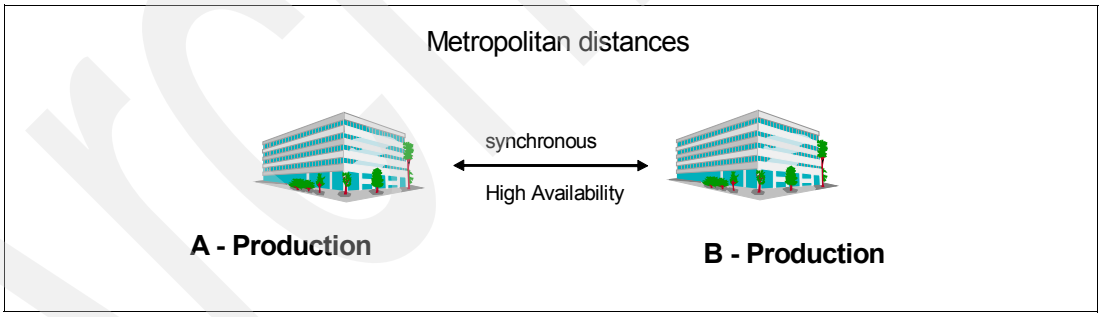

Figure 2-10 Two-site high availability data center model

High availability data centers are implementing functions such as server clustering to provide a higher degree of resiliency, redundancy, and failover/failback capability.

High availability data centers are, by definition, more symmetrical in their hardware and software configurations. High availability data centers are usually close enough together that server memory clustering and database memory clustering, along with synchronous data replication is possible. As these cross-device functionalities get more and more robust, there is a tendency a drive to be closer and closer together, to avoid cross-device communication latency.

Often, the high availability functions are typically implemented on an intra-site basis. Later, portions of the high availability cluster can be re-sited to metropolitan distances, in order to

provide some distance separation, while maintaining various cross-memory, cluster, and heartbeat functions of a tightly clustered high availability server environment.

When high availability server clustering functions are implemented within metropolitan distances, this sets a foundation that begins to answer the questions of IT management. Specifically, they can now start to implement various capabilities to allow workloads to be moved in a transparent manner from one server to another.

However, there are other important issues to be considered, including the need for out of region recovery.

2.5.3 Two-site traditional data center: Out of region recovery

As we discussed in 2.4, “Out of region recovery” on page 24, recovery distances at longer distances are a trend, as shown in Figure 2-11.

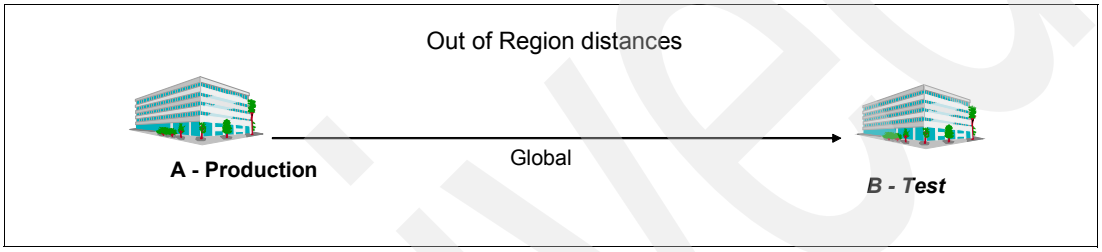

Figure 2-11 Two-site traditional data center - out of region recovery

Because this is a traditional data center model, note that the out of region data center is used for test and development. Traditionally, production workload has not been moved to the out of region data center on an ongoing basis, mainly due to cost and logistical reasons.

For the purposes of our discussion in this section, we briefly review the concepts that we covered in more detail in the “Major factors to consider for out of region recovery” on page 24. The highlights of the key out of region distance considerations (as compared to metropolitan distance High Availability) are:

- By definition, out of region recovery cannot provide true high availability. The out of region distances imply a very fast shutdown and fast restart, but this is not a hot, live cutover.
- Bandwidth costs typically limit what level of out of region recovery is feasible.
- An out of region data center model very much needs to have good business process segmentation in order to assure that bandwidth requirements are affordable. This supports the assertion that we made in 2.4.3, “Essential need for business process segmentation” on page 26 that business process segmentation is a very important strategic element in today’s data center strategy.

Important: Note that if there is a lack of business process segmentation, the exposure is that IT management could find itself in a situation in which their IT environment is too big, within available bandwidth, to be replicated to an out-of-region site. This situation is of particular concern when there are regulatory requirements for out of region recovery.

You should be aware of this kind of situation and work to avoid it.

Out of region recovery or metropolitan distance recovery are both viable strategies. It is a matter of selecting the best fit for the business and regulatory requirements.

2.5.4 Blending the best of High Availability and out of region recovery

With all of these requirements being placed on IT and demanding an upleveling of the IT Data Center strategy, we now examine the emerging best practices evolution of the data center strategy. The objective of this emerging strategy is to be able to blend the best characteristics of:

- Traditional two data center model
- Two-site high availability: Metro distances
- Two-site traditional data center: Out of region recovery

This emerging strategy also needs to successfully address the IT management questions that we described earlier in “New requirements stress the traditional data center model” on page 28.

The *Planned Workload Rotation data center model* that is emerging today is being implemented in major IT organizations to provide benefits that address all of these requirements.

2.5.5 Planned workload rotation data center model

The data center strategy that we see on the strategic horizon to address all of these factors is a planned rotation of workloads between sites on an ongoing basis, as illustrated in Figure 2-12.

Tip: We review this data center model here, because understanding the Planned Workload Rotation data center model is essential in order to fully appreciate, understand, and present the three-site data center business value and three-site data center best practices for business justification.

In Figure 2-12, the clouded area around each data center indicates that workload rotation is in effect.

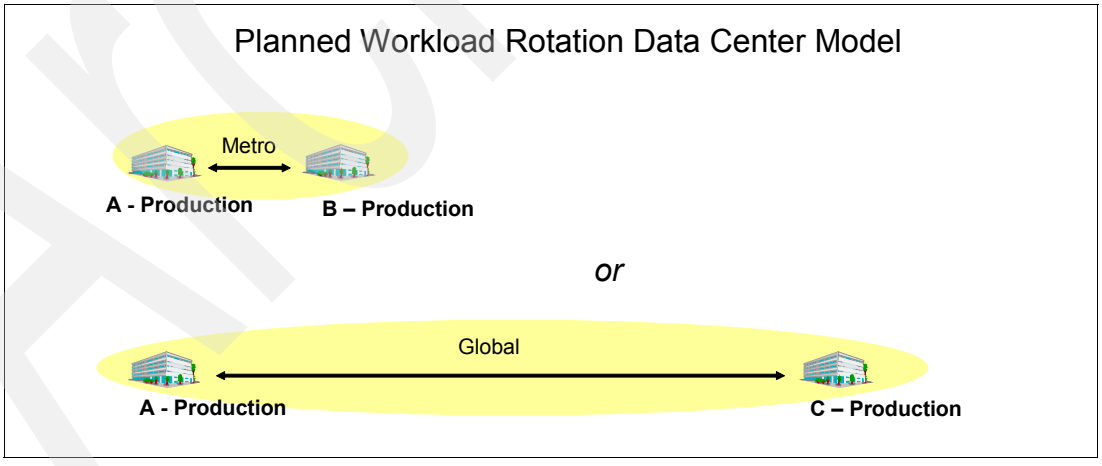

Figure 2-12 Planned workload rotation data center model

The planned workload rotation data center model includes the following actions:

- You implement the strategic path of the data center to eventually perform *planned workload rotation* between sites on an ongoing basis.
- You architect your long range IT infrastructure, business process segmentation, application segmentation, IT department standards, network, testing and development procedures, automation standards, database standards and so forth, to strategically enable, and participate in (over time), in this planned workload rotation data center model.
- You use this architecture in a planned rotation of the data center workloads between the sites on an ongoing basis.
- When you do this workload rotation, you rotate test and development in the opposite direction.
- This bi-directional workload rotation exploits functions such as high availability clustering and automation to provide a fast, reliable, repeatable, scalable, nearly continuous service availability during the workload rotation. You architect non-data center distributed systems to provide service on their own during the planned workload rotation.
- You automate the planned workload rotation using end-to-end automation of all aspects of the failover.

Philosophy of the planned workload rotation data center model

A planned workload rotation data center strategy implies the following factors:

- Symmetrical data centers that provide a 100% mirror of the data center IT environment
- Test and development at this level of scale is intended to enable:
  - TTM and TTC acceleration of application delivery
  - Risk Mitigation and Effort Mitigation for today's large scale projects
- By providing the necessary IT test and development foundation to achieve:
  - Business Process Changes
  - Complete change management
  - Source code management
  - Package application and system changes into release updates
  - Use terminal simulators for stress testing at the system level
  - Have well defined exit criteria from each testing stage
  - Post installation reviews
  - Change quality, problems, and performance
  - Schedule and communicate all changes
  - Automation

The planned workload rotation model leads to a fundamental conceptual reorientation in the way of thinking about how a Data Center does its application development and testing. It becomes the method that is used to respond to the demanding TTM, TTC, and TCO requirements that IT is experiencing today.

The symmetrical data centers are designed to provide a complete controlled testing environment, in which all factors related to timely, quality application delivery are enabled and tested end-to-end. For the large enterprises, a major benefit of the symmetrical data centers is the ability to test and integrate at their required level of very high scalability. Often, vendors do not have test floors that approach the level of scalability that large enterprise clients demand. With this approach, the client's data center will have the scale and IT resources to perform system level application and IT scalability, test, integration, and high system stress-related problem identification, problem reproduction, problem fix, and testing.

In a workload rotation model, our secondary is fully able to run day-to-day production workload, and mirrors the components in our primary. Other than the act of shifting the workload itself, there is then really little or no difference between running on the primary and running on the secondary.

The designation of primary and secondary can therefore easily switch back and forth on a regular basis, and this is what we mean by workload rotations. We intentionally shift workload from primary to secondary under normal operations, on a regular basis. Thus, we also shift the designation of secondary back and forth between the two sites.

You can implement planned workload rotation at either metropolitan distances or at out of region distances. The philosophy of planned workload rotation remains the same, and there are IT organizations that implement planned workload rotation at both metropolitan and out of region distances.

Note: The considerations of distance to the other site do apply. As we discussed in 2.4.1, "Considerations for distance of to recovery site" on page 24, there is greater amount of effort involved in doing planned workload rotation at out of region distances, as compared to metropolitan distances.

IT benefits of the planned workload rotation data center model

Here are the anticipated benefits to the IT department of the planned workload rotation model:

- The secondary site failover is fully tested, and the failover technologies and processes are tried, tested, and vetted out, system-wide, on a regular basis, under controlled conditions. Management and staff are much more likely to execute effectively in the event of an unplanned outage and failover.
- You have a much easier way to demonstrate the quality of the Business Continuity and Disaster Recovery plan from an audit standpoint.
- Beyond Business Continuity, you can also:
  - Provide resources to test and development to model production identically or nearly identically.
  - Perform true scaling tests at a true production level.
  - Reduce the support matrix, because you no longer need to consider differences between components at primary and secondary.
- You have the infrastructure to ease the pain and impact of changes such as massive data migrations, large software and operating system migrations, large infrastructure changes, and so forth.
- You can reduce the cost of compliance because planned workload rotations has the effect of providing ongoing proof points of compliance.

However, not that beyond the implications for Business Continuity, the *major justification* for the planned workload rotation strategy is that it is designed to enable and support the greatly accelerated application delivery TTM requirements that business is demanding for the future.

One of the primary drivers of this strategy are the following justifications, characteristics, and associated benefits designed to address the TTM and TTC issues:

- Faster TTM of applications: Testing and development processes and resources are upleveled as part of this strategy. Because of the need to have the second data center be as symmetrical as possible to the first data center, testing and development acquire resources that allow them to:
  - Perform full-scale stress tests and full-scale problem recreation
  - Significantly expand resources for testing, quality
  - Provide faster time to implementation with higher quality
- Automated data center workload rotations are tested on an ongoing basis
- Automated data center failback workload rotations are tested on an ongoing basis
- Reduce or remove cost of compliance from the future contingency plans

In effect, planned workload rotation data center model is a simultaneous Business Continuity *and* application TTM and Business Continuity configuration that provides actual day-to-day use.

Business benefits of the planned workload rotation data center model

The anticipated business benefits that result from the planned workload rotation data center model are:

- Continuity of Business Operations to help the business become more anticipatory, adaptive, and robust, from IT to all business processes (takes orders, ship, manufacture, and so forth).
- Regulatory Compliance to help the business comply with new government rules and regulations more quickly and cost effectively.
- Reduce the cost of Risk Management to help the business stay competitive by managing risk more efficiently and cost effectively.
- Security, Privacy, and Data Protection to help protect the business against threats, both internal and external, and to develop a critical information management strategy.
- Expertise and Skills (Outsourcing or Training) to help the business obtain and program manage the expertise and skills that are necessary to maintain continuous business operations.
- Maintaining Market Readiness to help the business stay competitive by anticipating and quickly responding and adapting to changes in market requirements and to accelerate research and development to ensure that the business has the right products, at the right place, at the right time.
- Becoming a more attractive Partner to help the business partner more quickly and effectively within the industry by becoming a trusted and reliable business partner in the supply chain or value-net.

A planned workload rotation data center strategy is really an all-inclusive strategy, covering IT infrastructure simplification, IT application development and testing cycle, IT Business Continuity implementation and testing strategy, IT compliance strategy, and more.

In the end, a planned workload rotation strategic goal allows the IT organization to:

- Accelerate application time to implementation
- Significantly improve application quality
- Remove cost of compliance
- Obtain full exploitation of all IT Data Center assets

The model is implemented actively in the IT industry today

The planned workload rotation data center model is being deployed actively today by leading edge IT organizations. Clients embarking on this model are pursuing it for exactly the reasons that we have discussed in this section. It is estimated that it might be possible to achieve as much as a 20% reduction in TTM for developing applications using this data center model.

A vision not necessarily a commitment

It is important to recognize that you should not feel compelled to make a committed delivery date for reaching planned workload rotation in order to reap the benefits. Rather, we recommend that you simply consider that planned workload rotation is the key emerging trend that can provide the best target environment for IT standards and integration, as we discuss in 2.2.5, “New importance of IT standards and integration” on page 17.

You can and should consider planned workload rotation primarily as a *vision* for the data center of the future. By adopting a vision such as planned workload rotation into your data center strategy, you can place your data center on a path to someday reach this vision. Having this model as a vision then allows all other IT activities that are related to architecture, mission, technology selection, and so forth, to align around this vision.

Even if you never achieve the actual state of planned workload rotation, aligning your data center along this path can provide very good benefits, as the intermediate building blocks of this vision can, each in their own way, add value. Each step that you take along this path, benefits because each step is aligned to a greater vision, and all investment made in pursuit of that vision can be protected.

Summary of the planned workload rotation data center model

The objective, benefits, and characteristics of the planned workload rotation strategic data center model are:

- Fully automated high availability data center swaps
- Planned workload movement, ongoing basis
- Foundation for compliance cost avoidance

The important implication of this data center model is that it implies symmetrical data centers. The symmetrical data centers are used to provide increased testing resource and infrastructure, which is used to address strategic business issues:

- IT time to implementation
- IT application quality
- Skills, testing, and procedures
- Significantly remove cost factor from compliance

Planned workload rotation is a key guiding vision for tomorrow’s strategic IT Data Center strategies:

- No need to make an immediate commitment
- However, should start to actively plan and prepare

2.5.6 Three-site data center strategies

Understanding the strategic directions for the planned workload rotation data center model, we now have the foundations to discuss business reasons and business justification for three-site data center topologies. We start by discussing why three-site recovery has become a topic of discussion in many data centers today.

Impetus for three-site recovery

Three-site recovery requirements are occurring in multiple industries today and is no longer only the domain of the finance industry or of only a select few high-end enterprises. The manufacturing, retail, transportation, telecom, and almost all other industries are starting to seriously consider and implement, three-site recovery configurations

Interestingly, IT organizations with local two-site High Availability are interested in adding out-of-region recovery capability, and 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.

The driving factors for three-site recovery include (but are not limited to):

- **Regional, national, and international consolidation**

  Consolidation and integration is occurring, not only of IT systems or of corporations, but also of nations and continents. Entities such as the European Union, China, and many large global corporations continue to form and grow. These kinds of consolidations, with the associated consolidation of key financial, banking, and national interests, drives a very high level of application availability.

- **High criticality applications that require continued resilience, even in event of loss of one site**

  The desire to provide an level of extra protection for business processes and IT systems that provide very mission-critical applications for these large, consolidated organizations and national interests. In the event of outage of the primary site, and failover to one of the two other sites, that a nearly immediate capability to have two site recovery is maintained.

- **Very critical applications that desire zero data loss even at out of region remote site**

  Some of these applications desire that in the event of a loss of the primary site, that zero data loss is still propagated to the asynchronous out of region remote site. To accomplish this, an intermediate, metropolitan distance site within synchronous data replication range is introduced. The intent is that in event of an outage to the primary site, yet the synchronous intermediate site survives, that the intermediate site can then forward the remaining data to the out of region site, thus bringing the out of region site to a zero data loss situation.

- **Existing High Availability recovery configurations desire to add out of region recovery**

  To add geographic distance to their existing recovery capability.

- **Existing out of region recovery configurations desire to add high availability locally**

  Out of region clients desire to implement high availability locally, many with the objective of eventually getting to the benefits of the planned workload rotation model at the local metropolitan distance, as outlined in “Business benefits of the planned workload rotation data center model” on page 34.

Traditional data center models cannot justify three-site recovery

The criteria in the previous section are admirable goals. However, the costs of implementing a second data center was not an inexpensive proposition. The question is, can we make *three*-site recovery affordable and justifiable to implement?

From the discussion in “New requirements stress the traditional data center model” on page 28, we can see that limitations in the traditional data center model imply a certain challenge in justifying just a traditional two-site data center model. Clearly, if the organization is *traditional* today, then it seems quite difficult or impractical to consider implementing a third *traditional* data center if we have not solved the challenges to the traditional two-site model.

In other words, in today's tight budget environment, if a two-site data center model cannot address and resolve the questions listed in "New requirements stress the traditional data center model" on page 28, then it is unlikely that cost justification can be found for a three-site data center that is dedicated to recovery only, as shown in Figure 2-13.

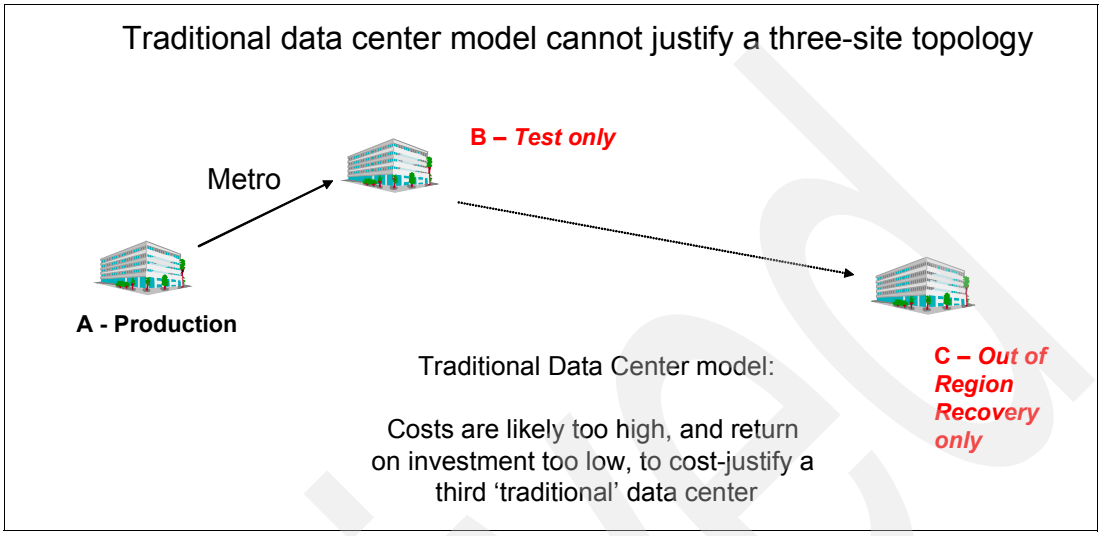

Figure 2-13 Traditional data center model cannot justify the cost of three-site recovery

If the exploitation of resources today are such that the test site B is not fully exploited, it is unlikely that a third site C for only test and recovery reasons can be justified either.

Planned workload rotation provides foundation for three-site recovery

Therefore, we can start see that in order for three-site recovery configurations to be justified to the business, we need to strongly consider including the planned workload rotation data center model as part of the equation. With this model, we can then provide three-site Business Continuity *and* simultaneously address the IT data center business challenges that we posed in "New requirements stress the traditional data center model" on page 28.

Figure 2-14 illustrates the best practices for a business-justifiable, cost-justifiable three-site data center configuration.

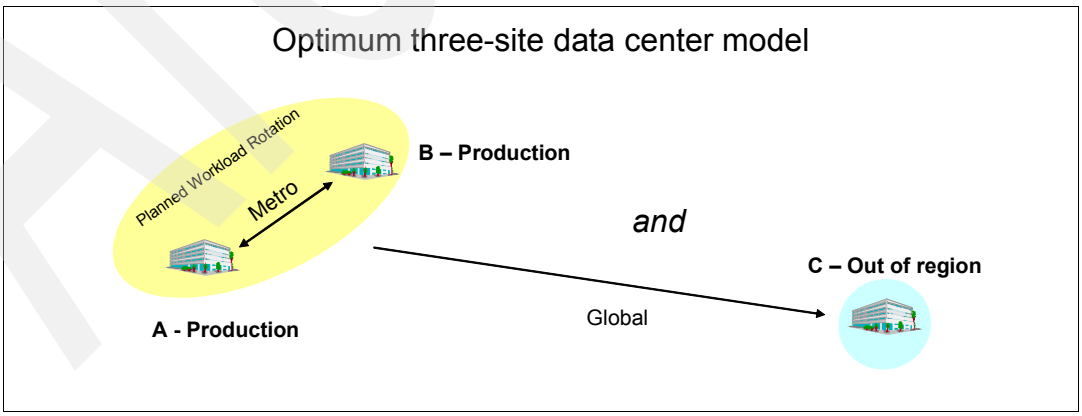

Figure 2-14 Optimum three-site data center model using planned workload rotation

The optimum cost-justifiable three-site data center model would be a combination of:

- Planned workload rotation between two metropolitan distance, synchronous replication sites
- Out of region recovery provided for the *cluster* of the two planned workload rotation sites

This three data center model provides an optimum blend of all the factors that we have discussed in this chapter. Specifically:

- It provides all the benefits of planned workload rotation, using the two metropolitan distance sites.
- It exploits the best blended characteristics of both metropolitan distance recovery and out of region recovery. This confirmation provides the best blend of the factors that we discussed in “Major factors to consider for out of region recovery” and “Major factors to consider for metropolitan distance recovery” on page 25.
- It exploits that fact that it is easier and less expensive to do planned workload rotation between the two halves of the local metropolitan high availability complex, as compared to the higher costs and longer failover time of doing planned workload rotation between out of region sites.

For all the reasons stated in this chapter, it is reasonable therefore to start the three-site implementation when High Availability has been integrated into the planning for the two metropolitan sites. Note that it is not necessary to wait for planned workload rotation to be complete, before initiating the three site data center implementation. Planning for metropolitan distance High Availability, with a future objective of planned workload rotation, provides the best foundation for being able to deliver cost-justification and business-justification for three site recovery to senior management.

There are other variations of the three-site model that might be hybrids of the basic optimized model shown in Figure 2-14.

For certain types of workloads, especially Web-serving workloads, we can implement not only workload rotation, but we might also keep all three sites active and load balance incoming production between them. In this particular permutation, we maintain High Availability and Disaster Recovery protection, while actually reducing aggregate capacity. A three-site all active model can allow to achieve continuous operations with only 150% capacity, including rolling maintenance with no impact to service delivery.

2.5.7 Requirements of three-site data centers

With this understanding of the optimum three site data center architecture, we can then detail the three fundamental IT replication requirements that arise. We recommend that you use these as primary criteria to evaluate three-site recovery solution options.

Assumption: In order to cover all possible outage scenarios while maintaining the criteria listed in this section, a triangle of telecom links with bandwidth sufficient to handle workload and resynchronizations is assumed to exist between all three recovery sites.

Fast failover and failback to any site

Fast failover and failback to any site means that a very fast restart of the production workload is possible at any of the other two sites. Ideally, if we fail over to the other high availability metropolitan distance site, we would want to re-establish recoverability with the out of region distance site as soon as possible.

Fast re-establishment of three-site recovery without production outages

We use the capability to re-establish the three-site recovery capability without quiescing the site where production is running for the following reasons:

- In a planned workload rotation data center model, we use the synchronous Metro Mirror sites to rotate workloads often between the High Availability sites.
- As part of the planned workload rotation, we want to re-establish three-site recovery capability without disruption to the production workload.

Fast re-establishment is the key requirement for a three-site data center recovery with planned workload rotation. If you are unable to do a fast *re-establishment* of three-site recovery without production outages, being unable to do a fast production outages of the planned workload rotation data center model in a three-site environment.

Quickly resynchronize any site with incremental changes only

When returning back to original site or sites or re-establishing three-site recovery, only incremental changes need to be copied back, thus shortening the elapsed time for the re-establishment to take place.

2.5.8 Implementation of the three-site data center model

The implementation cases for a three-site data centers depend on the type of traditional data center model that you are starting from. There are two cases:

- Traditional two-site today at metropolitan distances

  If you have a traditional two-site production and test configuration today at metropolitan distances, the three-site data center implementation steps would be:

  - Implement High Availability between the metropolitan sites, with a target of eventual implementation of planned workload rotation
  - Add out of region recovery to the High Availability metropolitan distance cluster

- Traditional two-site today at out of region distance

  If you have a traditional two-site production and test configuration today at out of region distances, the three-site data center implementation step would be (can be simultaneous) to implement High Availability locally at an additional metropolitan distance site.

In both cases, for the best practices in business justification, we recommend that High Availability between the two metropolitan sites, at a minimum, be part of the strategic plan (and from there, ideally moving towards planned workload rotation).

2.5.9 Strategies involving more than three sites

While still on the planning horizon and not yet in production status, we expect that selected three-site IT data centers with very high levels of critical recovery requirements, will begin to evolve a requirement for more than three sites. The reasoning is as follows:

- If I implement the three-site data center recovery model and if I have an event that causes a loss of both of the metropolitan distance sites, I will failover to the out of region site.
- Given the nature of the outage that caused a loss of both of the metropolitan distance data centers, I most likely will be running at the out of region center for some time.
- *Therefore, until the metropolitan distance data centers are repaired, I have a single point of failure, that is I have no recoverability if I lose the out of region data center.*

- If my applications and business warrants protection against this single point of failure data center consideration, I want to quickly re-establish a recovery capability for the out-of-region data center.

Providing a cost-effective configuration topology for addressing this recovery concern is doable.

The way to address this is to *expand* the concept of what the out-of-region recovery site does. Basically, we uplevel the out-of-region site from a *traditional data center model* to a *planned workload rotation* data center model as well.

Figure 2-15 illustrates this topology. The key point to notice is that the diagram is *not* a four-site replication solution. Rather, it is a *pair* of High Availability and planned workload rotation data centers, with a *portion* of each High Availability and planned workload rotation sites dedicated to providing out of region recovery for the other High Availability and planned workload rotation site.

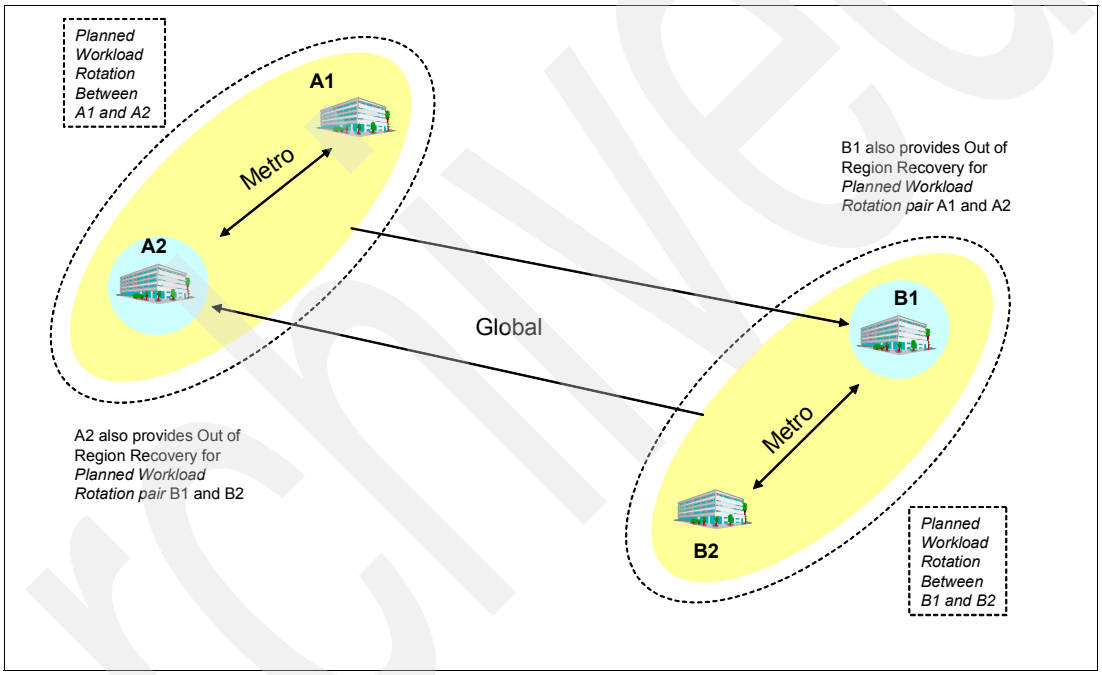

Figure 2-15 Strategies involving more than three sites

In the event of an outage of both A1 and A2, we would fail over to B1. We also would establish B2 as the local metropolitan distance recovery for B1.

The same logic would apply in reverse in the event of an outage of both B1 and B2. We would fail over to A2. We also would establish A1 as the local metropolitan distance recovery for A2.

The risk assessments and associated cost-justifications for this kind of architecture would consider a failover of A1/A2's workload to B1 to be a very rare event. Due to the rarity of such a failover, for cost-justification reasons we can consider the amount of recovery equipment at B1 or A2 to be adequate to provide an acceptably degraded level of performance.

As you might imagine, this level of resiliency is most likely to be required, only after the IT data center strategy reaches the point when the planned workload rotation data center model is already committed to the business.

From a practical standpoint, we highlight these future considerations to give you an awareness of real issues and considerations that are being considered today. There are a selected set of IT clients in the world that are already approaching realistic consideration of these greater than three site requirements.

As an industry, we all watch with interest to see how long-term multi-site data strategies such as these will evolve over time.

2.6 Summary

In this chapter, we outlined a series of the newest Business Continuity and IT Data Center trends and directions. These challenges are broad in scope, but they can be successfully met with proper awareness, insight, experience, and planning.

Briefly, in this chapter we saw that there is a Business Continuity planning focus shift to the business process, and not just the IT level:

- We discussed factors and considerations for justifying Business Continuity to the business. We introduced the terms *Time to Market* and *Time to Customer* and gave a model for relating these crucial business factors to the IT Data Center strategy and to IT Business Continuity.
- We reviewed the importance, more than ever, of IT standardization and integration, which can be accomplished through IT simplification and consolidation, and we also showed how this IT simplification and consolidation is the best prerequisite for Business Continuity. We discussed that this approach allows IT Business Continuity to be an intended by-product of the IT Data Center strategy.
- We reviewed emerging criteria for the Business Continuity solution (reliability, repeatability, scalability, affordable testability, and automation). We discussed out of region recovery and bandwidth.
- We provided a detailed discussion of emerging data center strategies for two, three, or more sites. We highlighted that two site topologies are full exploitable when High Availability and planned workload rotation are the long-range strategic objectives.
- We saw how three-site recovery topologies are expansions of the two-site strategies, and that more than three sites is really an extension of a pair of two-site High Availability and planned workload rotation complexes.

In the end, this chapter is meant to give you a current background on the IT trends and directions that affect how IT Business Continuity, and how IT Business Continuity can provide necessary business values:

- Accelerate application time to implementation
- Significantly improve application quality
- Remove cost of compliance
- Obtain full exploitation of all IT Data Center assets
- Meet recovery time objectives
- Satisfy budget and manpower constraints

With this chapter as a background, you are now in an excellent position to better appreciate the following chapters, where we discuss in detail how to undertake and perform the ideal Business Continuity planning process.

3

Business Continuity planning, processes, and execution

Business Continuity is much more than IT hardware and software infrastructure. People, processes, notification and decision trees, physical facilities, telecom, and much more all come into play. In this chapter, we overview the main aspects of Business Continuity planning, processes, and execution.

The topics in this chapter include:

- Introduce key Business Continuity planning concepts and metrics
- Provide an overall ideal Business Continuity planning cycle
- Provide an overview of various skills and procedures that are necessary for planning a complete Business Continuity solution

In addition, we overview best practices for tiering the recovery, application segmentation, and planning and testing the various phases of a good Business Continuity project.

We present a methodology to apply the concepts of this chapter in Chapter 6, “The Next Step Business Continuity workshop” on page 173.

A key point to remember is that Business Continuity is much more than just Disaster Recovery. Well planned Business Continuity solutions are intended to provide ongoing value to the business on a daily basis, even if a true disaster is never encountered. In this chapter, we see how you can plane for and accomplish that ongoing value as well.

3.1 Introduction to this chapter

As business increasingly depends on IT systems for customer operations, Business Continuity is one of the top priorities in IT organizations. Any outages or data inaccessibility will have severe negative consequences for the business. The ability to keep company information to satisfy governmental regulations is also paramount. As business demands 24x7x365 support, simple data and systems backup and recovery becomes a challenge as planned outage time approaches zero.

At the same time, IT departments are facing unprecedented data center complexity and pressures—more applications, data, servers, storage devices, both at the data centers and remote locations.

With IT budgets under pressure, and sufficient skills difficult to come by, it is not surprising that it is difficult to initiate and drive a major Business Continuity enhancement project to completion.

Because there are a large number of excellent reference materials and expertise available on the topic of Business Continuity planning, this chapter is intended to provide a best practices overview. By reading this chapter, you will be able to understand the essentials of doing Business Continuity planning and how to do that in a methodical, comprehensive manner. We also highlight the best practices for combining the Business Continuity enhancement project with other IT Infrastructure streamlining projects.

3.1.1 Intended audience for this chapter

The intended audience for this chapter is the IT Chief Information Officer, IT Directors, and IT management. The ideal Business Continuity planning process that we present here is inclusive of the entire business, including all non-IT aspects. Business Continuity planners, and non-IT management will also find this chapter of interest.

IT technical management, IT operations management, and IT application development staff can also benefit as well. For this audience, 3.5.2, “IT strategy design” on page 89, is of particular interest.

3.2 Typical evolution of a Business Continuity program

Business Continuity is unique in that it touches almost every part of an organization’s business processes, infrastructure, and people. Accordingly, effecting change in the organization to significantly improve Business Continuity capability, will have a time curve for it to take effect.

We can see a diagram of how this often looks in Figure 3-1.

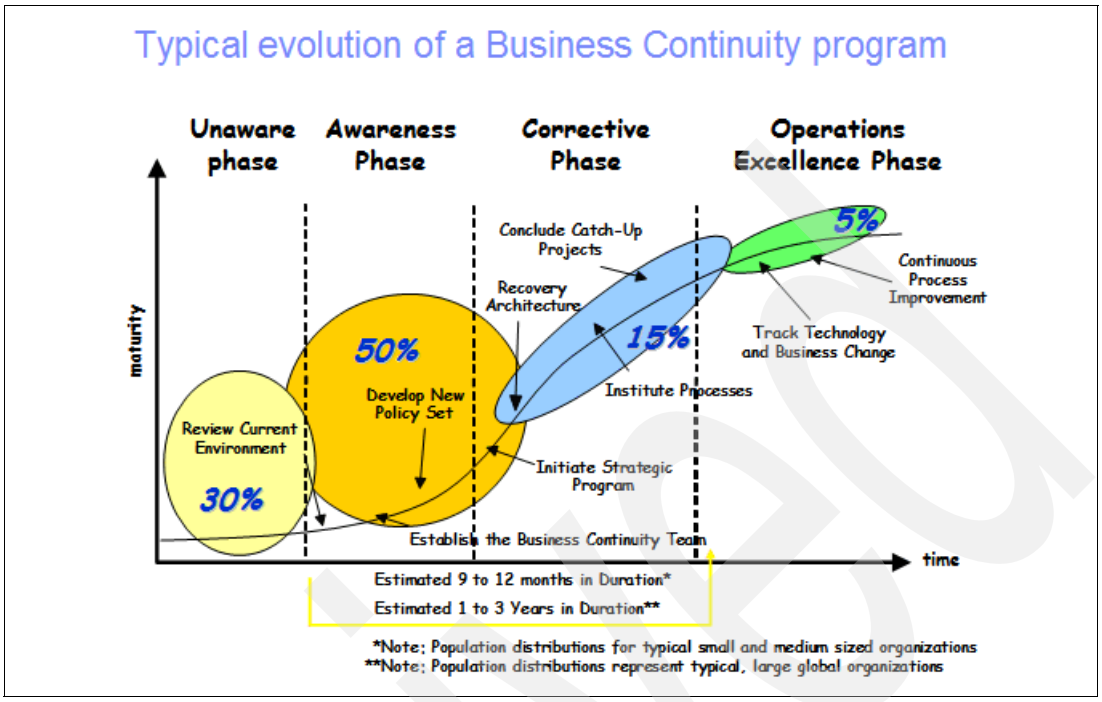

Figure 3-1 Typical evolution of a Business Continuity program

The presence of this time curve is one of the reasons that you need to implement a cost-effective Business Continuity program in phases, that are integrated with the organization's other initiatives related to improving time to market, improving application quality and function, and IT infrastructure streamlining. We discuss the cost-effective justification for Business Continuity in more detail in 2.2, "Justifying Business Continuity to the business" on page 12.

As you start to plan for implementing a Business Continuity program within your IT organization, we discuss how best to handle and plan for the fact that these elapsed times, as shown in Figure 3-1, occur. To manage the time aspect, it is important to:

- Plan on building a step-by-step process that shows incremental value to the business every step of the way. Set expectations appropriately.
- Do not try to solve everything at the same time. Set expectations appropriately with senior management that IT will deliver on a step-by-step, multi-phase approach.
- Plan for, document, and deliver incremental value in every step of the Business Continuity program. In this way, you give senior management the assurance that continued return on investment is visible in every phase of the project.

3.3 Ideal Business Continuity planning process

Figure 3-2 illustrates an ideal Business Continuity planning process, as used by IBM Global Technology Services. We use this planning process as our guideline for this chapter and explore each step of this process in detail.

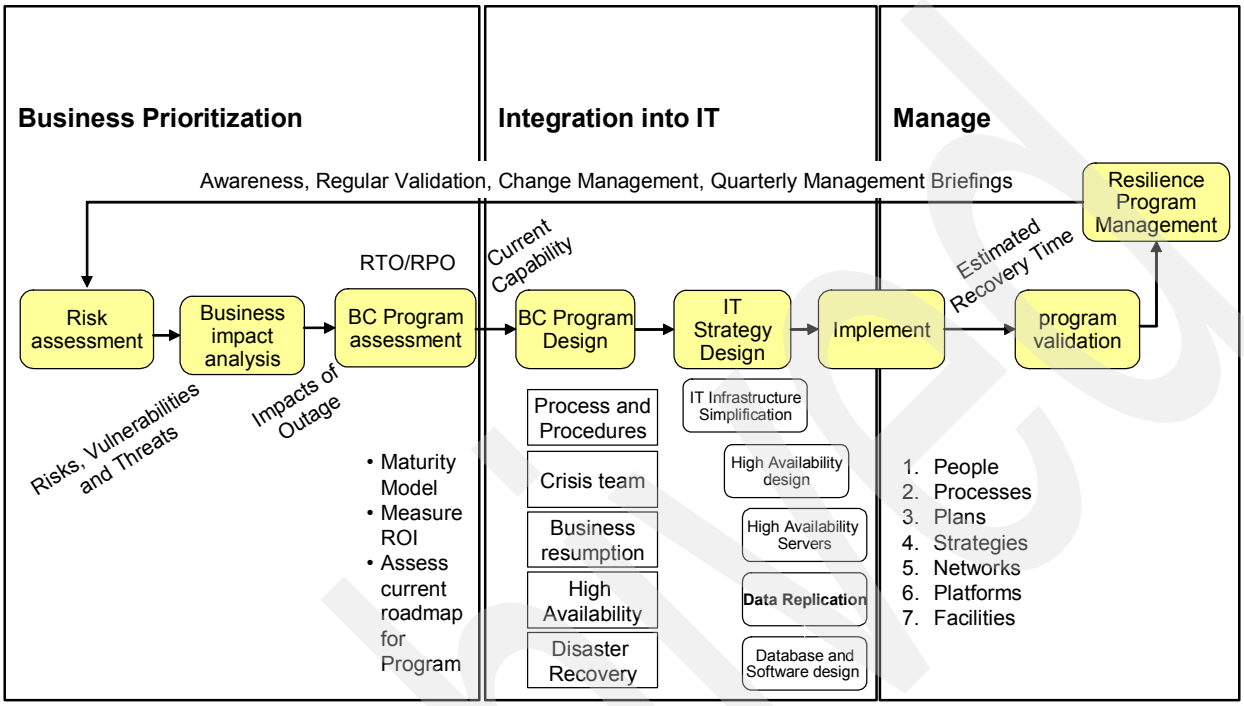

Figure 3-2 Ideal Business Continuity planning process

As shown in Figure 3-2, the ideal Business Continuity planning process is a closed loop that supports continuing iteration and improvement as the objective.

There are three major sections to the planning process:

- Business prioritization
- Integration into IT
- Manage

Note: IBM Global Technology Services as well and IBM Business Partners can provide consulting, expertise, and skills to develop and tailor the kinds of phases, activities, and tasks described in this chapter for a wide variety of Business Continuity projects. You can find more information about the IBM Global Technology Services capabilities that are available in this area at:

`http://www.ibm.com/services/continuity`

We begin our discussion with business prioritization.

3.4 Business prioritization

Business prioritization is the process of identifying the scope for our Business Continuity plan. In business prioritization, we identify various risks, threats, and vulnerabilities and establish priorities. We then compare the priorities with an assessment of our current Business Continuity program to develop a differential baseline from which to do Business Continuity program design.

In this section, we discuss:

- Risk assessment
- Business impact analysis
- Business Continuity program self-assessment

We begin by discussing risk assessment, as shown in Figure 3-3.

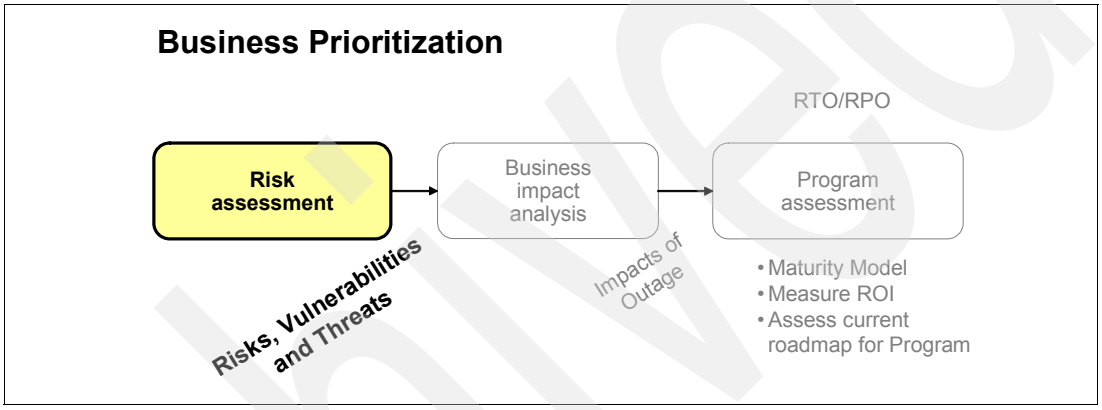

Figure 3-3 Business prioritization - risk assessment phase

Note: The output of the risk assessment phase is input to business impact analysis phase.

3.4.1 Risk assessment

The Business Continuity planning team should prepare a risk analysis and risk management study that includes a range of possible disasters, including natural, technical, and human threats. The output of these studies will be input to the business impact analysis phase.

Risk assessment is the starting point in all initiatives related to Business Continuity, because it determines the scope and scenarios that the company is attempting to mitigate. Major steps in performing risk analysis are:

- Identify the risks, and make an estimation of their likelihood of occurrence
- Evaluate and prioritize the risks
- Create a report of the identified risks and vulnerabilities

Identifying risks, vulnerabilities, and threats

Figure 3-4 shows that IT systems are vulnerable to a variety of disruptions, ranging from mild (such as, short-term power outage or disk drive failure) to severe (such as, equipment destruction or fire). An organization can eliminate many vulnerabilities through technical, management, or operational solutions as part of risk management or security controls; however, typically it is impossible to completely eliminate all risks.

Business Continuity planning in this phase is designed to identify risks and vulnerabilities that you need to add to existing risk management and security activities. In this way, you provide *input* for the scope of the upcoming Business Continuity program design.

Note: Risk assessment identifies potential risks and vulnerabilities, but we do not define the scope of risk assessment in this step. We define the scope in the next step, the business impact analysis phase.

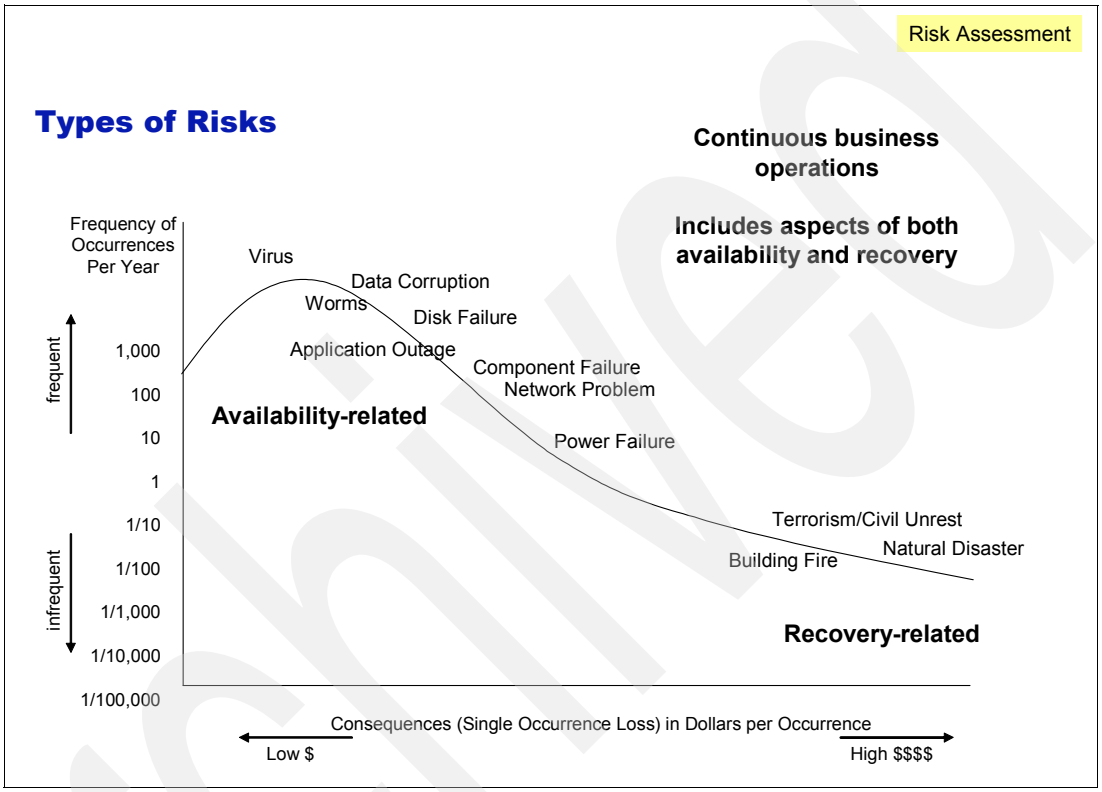

Figure 3-4 Types of risks, frequency, and impact

You need to analyze each functional area of the organization to determine and identify potential risks, their likelihoods, and impacts. In addition, you need to perform the evaluation and identification at the business process level, not just the IT level.

Risks result from a variety of factors, although typically they are classified into three types:

- **Natural**: Hurricane, tornado, flood, and fire
- **People (human)**: Operator error, sabotage, and malicious code
- **Equipment (technology)**: Equipment failure, software error, telecommunications network outage, and electric power failure

Not all risks are present with respect to a given IT department (or site). For example, depending on its location, a system might have no risk of damage by hurricane, but a reasonably high risk of effects from a tornado. It is important to include documentation of the potential impacts and consequences resulting from risks and vulnerabilities which can cause loss of data, information and services.

Performing a risk assessment identifies the potential list of specific risks to a specific set of business processes and systems. The risk assessment is critical because it enables the

management responsible for Business Continuity to focus risk management efforts and resources only on identified risks, in a prioritized manner.

A thorough risk assessment should identify all potential risks and attempt to determine the probability of the risk actually occurring. In the past, fire or natural causes were the greatest threat to an organization. Unfortunately, today, an organization must also consider intentional human destruction. The risk assessment needs to identify both commonplace and worst case situations, such as virus attacks, accidental data corruption, destruction of the main building, and more. It is appropriate to do an initial filtering of the list of potential risks down to a list of identified risks, including the ones with a measurable probability. You then do final refinement is in the next step, business impact analysis.

Items in identifying the probability of a specific vulnerability and risk might include, but are not limited to:

- Geographic location
- Topography of the area
- Proximity to power sources, water bodies, and airports
- Degree of accessibility to the organization
- History of local utility companies in providing uninterrupted services
- History of the area's susceptibility to natural threats
- Proximity to major highways that transport hazardous waste and combustible products
- Proximity to nuclear power plants
- Other factors, such as political instability

Because risks can vary over time and new risks might replace old ones as a system evolves, the risk management process must by ongoing and dynamic. Figure 3-5 is a detailed template for risk assessment.

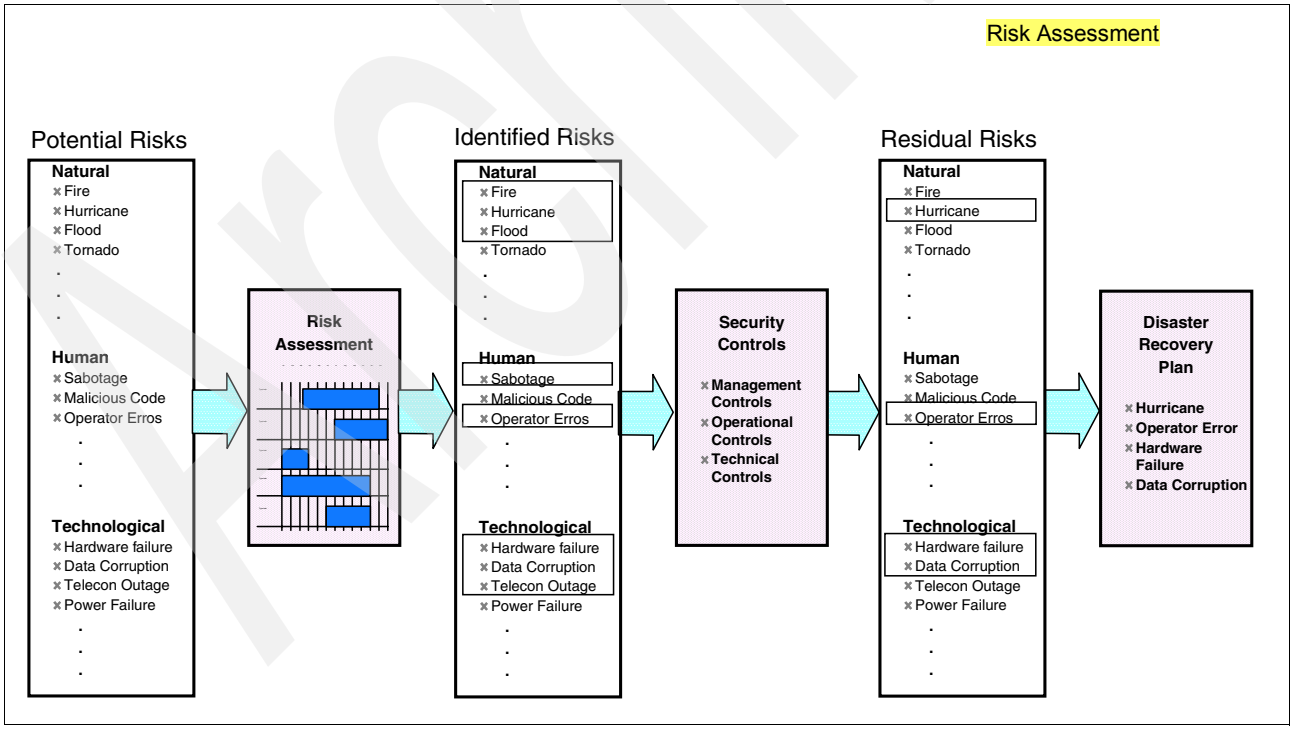

Figure 3-5 Template for risk assessment in the Business Continuity plan

Ideally, all identified risks would be eliminated completely. Of course, this is rarely possible or cost-effective. Rather, you attempt to identify risks to an acceptable level and remain aware of and document residual risks.

Because these identified risks represent the most important set of situations that could affect system performance, availability, integrity, and security, you can reduce the scope of the Business Continuity plan to address only the identified risks. As a result, you can more narrowly focus the Business Continuity plan, conserving company resources while ensuring that efforts are focused on the most important risks to effective system recovery capability.

Figure 3-6 provides a *simplified* output risk assessment template that is suitable for input to a basic business impact analysis.

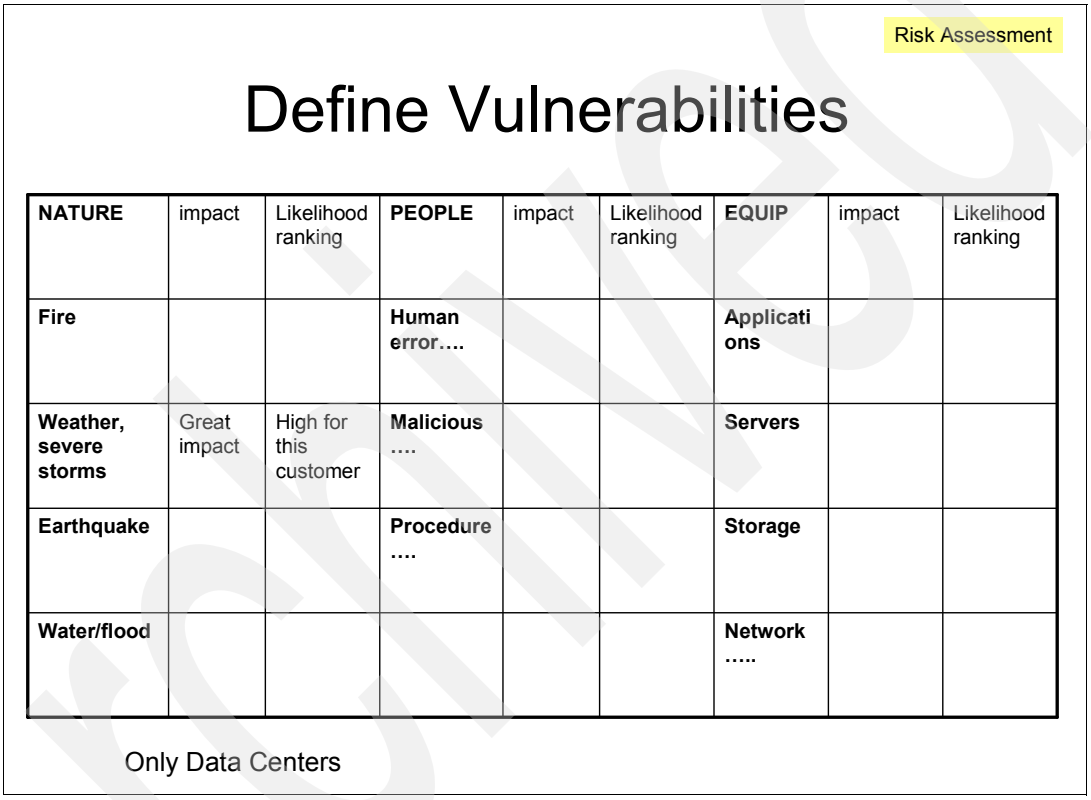
Risk Assessment

Define Vulnerabilities

NATURE	impact	Likelihood ranking	PEOPLE	impact	Likelihood ranking	EQUIP	impact	Likelihood ranking
Fire			**Human error….**			**Applicati ons**		
Weather, severe storms	Great impact	High for this customer	**Malicious ….**			**Servers**		
Earthquake			**Procedure ….**			**Storage**		
Water/flood						**Network …..**		

Only Data Centers

Figure 3-6 Simplified risk assessment output template

You use this simplified template to document a series of identified risks, along with an impact description and a likelihood ranking. The likelihood ranking for the vulnerability considers at least the following factors:

- Historical: Has this event happened in the past?
- Frequency: If this event has happened before, how often has it occurred?
- Duration: What was the duration of the event?

Tip: You can find an applied example of using this simplified risk assessment template in Chapter 6, “The Next Step Business Continuity workshop” on page 173 and in 7.1, “Introduction to the Case Study” on page 234.

Risk assessment summary

Risk assessment is the first step in Business Prioritization in the ideal Business Continuity planning process. Figure 3-7 is a summary flow chart that shows the types of inputs, actions, and outputs in the process of performing a risk assessment.

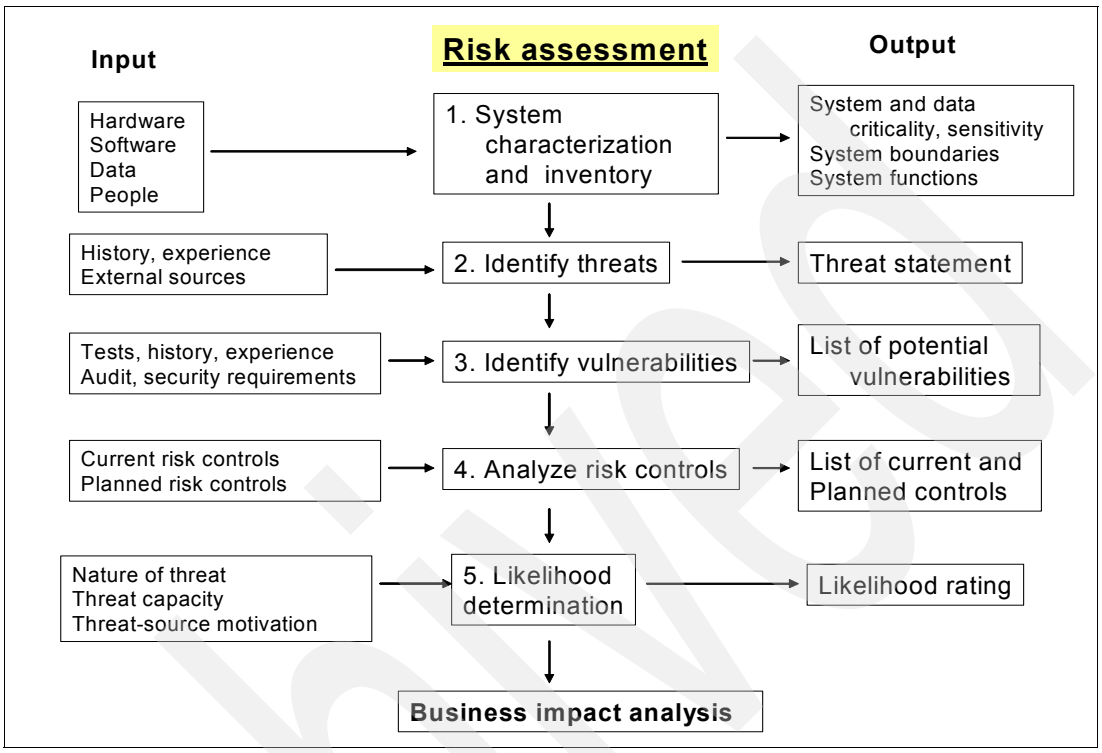

Figure 3-7 Risk assessment summary flow chart

When the risk assessment is complete, you take the output from the risk assessment and feed that output into the next step in Business Prioritization, the business impact analysis.

3.4.2 Business impact analysis

Figure 3-8 shows where the business impact analysis (BIA) phase fits into the Business Prioritization portion of the ideal Business Continuity planning process.

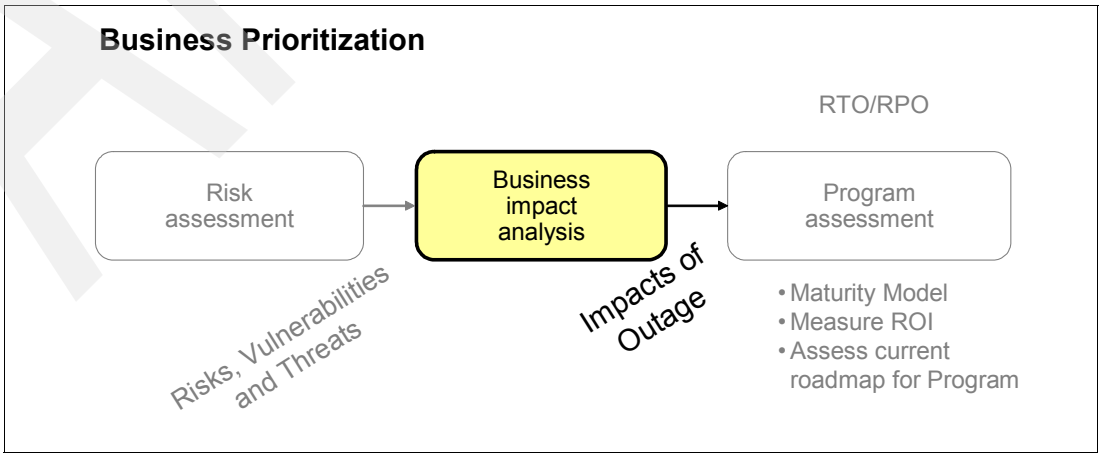

Figure 3-8 Business prioritization - BIA

Taking the list of identified risks from your risk assessment as input, you can begin a detailed assessment of the relative impacts and priority of those risks.

Impacts of outage

The BIA enables the Business Continuity planner to fully characterize the business process requirements, interdependencies and to use this information to determine and refine the *ranking and scope* of the Business Continuity plan. Organizations that have the most revenue, and are the most heavily dependent on online systems, have the highest potential loss of revenue from widespread IT application and network outages. Depending on the industry, the revenue per hours lost by a disaster is different. Thus, the energy, telecommunications, and manufacturing industries as well as financial institutions continue to have the highest revenues per hour.

The percentage of revenue actually lost depends on the criticality of business processes and systems that experiences the outage (such as, degree of customer interaction, existing workarounds, peak periods) and the number of users that are affected by the outage or slowdown. Also, significant immediate losses can result in bad publicity and loss of customer trust that affects future revenues.

Figure 3-9 illustrates a detailed BIA process.

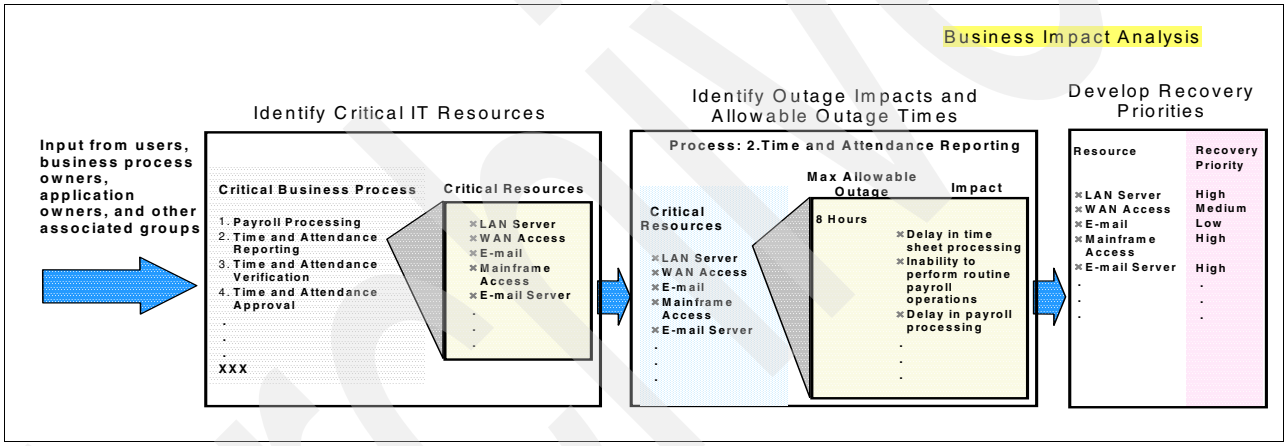

Figure 3-9 Example of a BIA process

The steps in this process are:

1. Identify critical IT resources.

   Evaluate the business processes and the related IT systems to determine the critical functions that are performed by the system and to identify the specific business processes and IT resources (hardware, network, software, and data) that are required to perform them.

2. Identify disruption impacts and allowable outage times.

   Analyze the critical resources that you identified in the previous step and determine the impact on business process operations if a given resource were disrupted or damaged. The analysis should evaluate the impact of the outage in two ways:

   – Over time
   – Across related resources and dependent systems.

3. Develop recovery priorities.

   Characterize the outage impact and allowable outage times. This process is an ongoing process that involves communication with multiple parts of the organization in order to reach a consensus. This communication is essential, in that it defines the recovery strategies that are to be prioritized, developed, and implemented during the activation of the Business Continuity plan. For example, if the outage impacts step determines that the business process must be recovered within four hours, the Business Continuity plan will adopt measures to meet that need.

BIA defines the scope for the Business Continuity plan

The output of the BIA defines the scope of the Business Continuity plan. In other words, you identify and confirm the defined set of conditions, outages, and events that the Business Continuity plan will address.

It is reasonable that this scope should start at a baseline level and then expand over time. This allows us to get started and address important vulnerabilities immediately, while expanding the coverage to be more comprehensive over time. Figure 3-10 is a suggested summary template for documenting the defined risks and scope that the Business Continuity plan will cover over time.

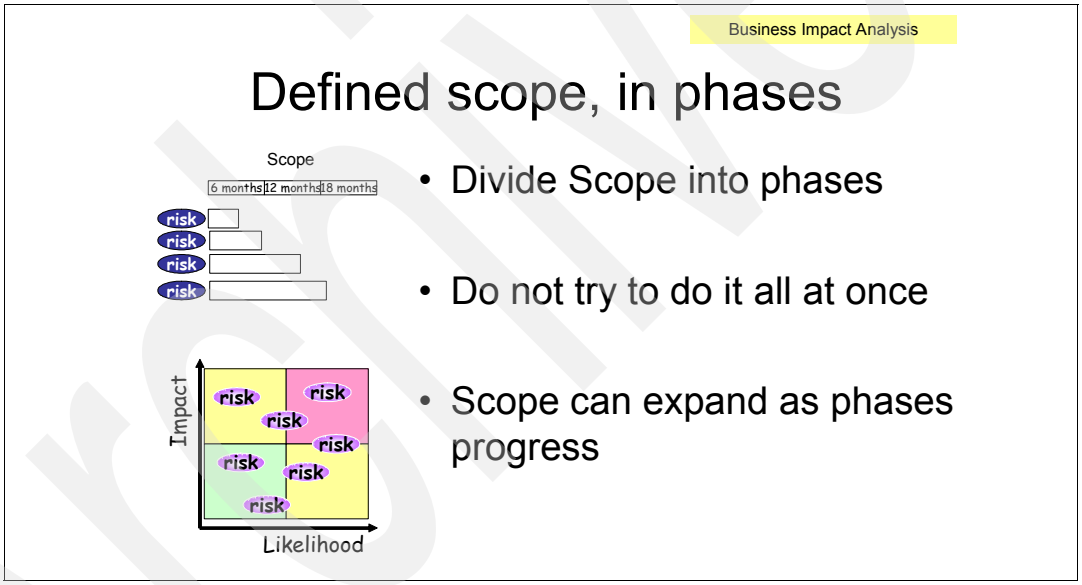

Figure 3-10 BIA defines the scope in phases for the Business Continuity plan

Summary

Figure 3-11 lists a process flow chart for the BIA portion of Business Prioritization, including a list of the inputs and outputs from this step.

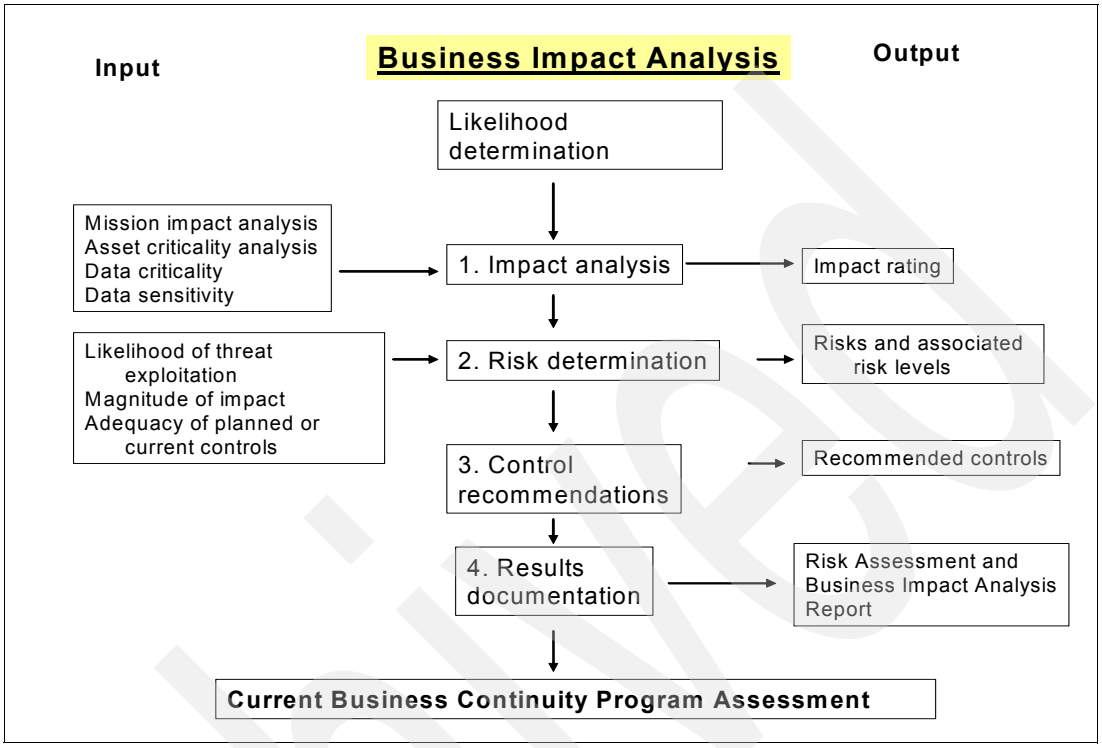

Figure 3-11 BIA flow chart

Next, you use the BIA as input to your assessment of you current Business Continuity program.

3.4.3 Program assessment

Figure 3-12 shows where the assessment phase fits into the Business Prioritization portion of the ideal Business Continuity planning process.

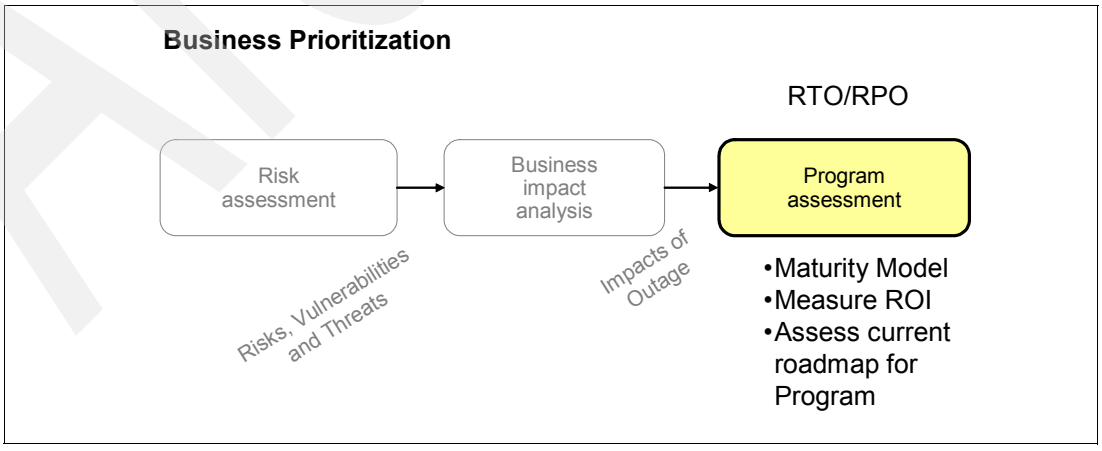

Figure 3-12 Business Prioritization - program assessment

In this step, you do a significant amount of data collection to determine and document the baseline status of current Business Continuity program.

Data collection to assess current Business Continuity environment

To determine the current status of the organization's Business Continuity program, you document the various functions performed within each department. The depth of your analysis can vary, depending on the level and scope of the Business Continuity project. Two weeks to one month might be required to fully document all the principle functions that are performed inside and outside a department and to assist in identifying the necessary current Business Continuity status for the departments to conduct their daily operations satisfactorily.

For circumstances where this level of investigation is not required, Figure 3-13 provides a simplified data collection template.

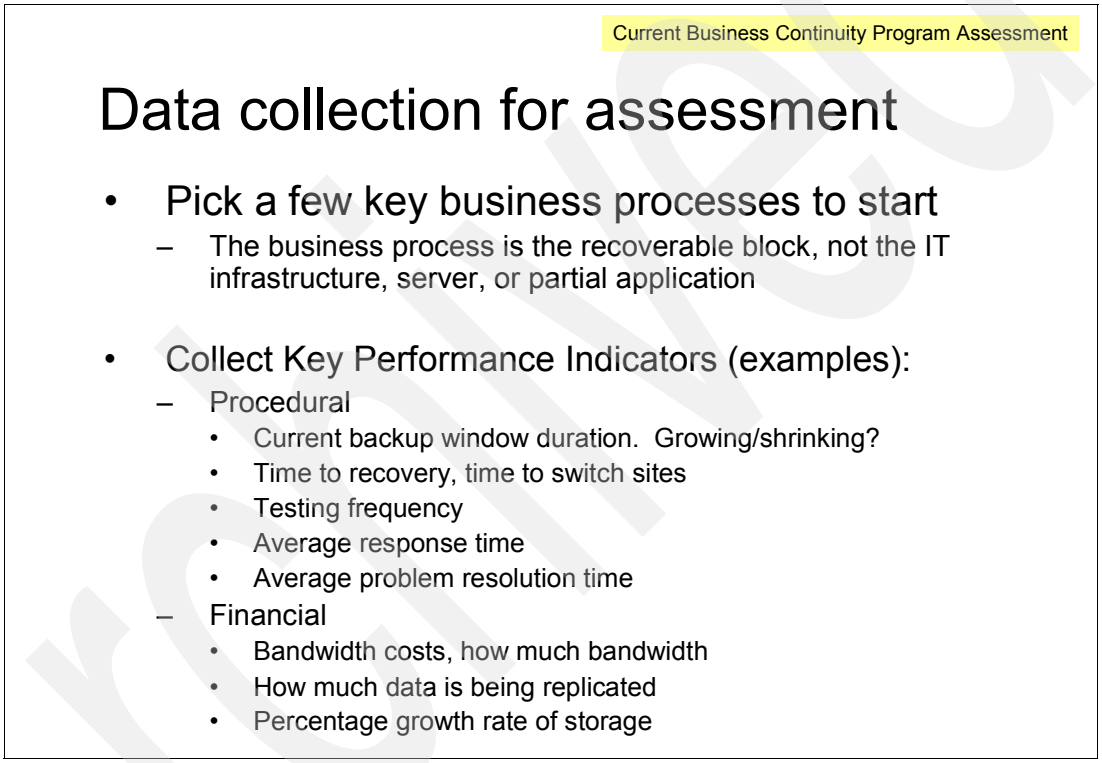

Figure 3-13 Simplified data collection template for current Business Continuity program assessment

Note that in this simplified template, we recommend that you select just a few critical business processes to study. Because Business Continuity planning should be an ongoing, iterative process, it is not necessary to document all the business processes initially. You will have time to document the processes in subsequent iterations of the planning process. We suggest that it is more important that you get started building skills and understanding the methodology. Therefore, a good way to start is to gather important information about just one, two, or at most three business processes.

Key Performance Indicators

Key Performance Indicators (KPIs) is a business management concept that is particularly applicable to a Business Continuity project. In this section, we suggest some starter IT Business Continuity KPIs. You can choose and modify these as you see fit.

KPIs reflect your client's IT Business Continuity requirements success factors. The Business Continuity solution KPIs must accurately portray the organization's Business Continuity

objectives, they must accurately measure the keys to success, and they must be measurable and quantifiable.

When defining or suggesting Business Continuity solution KPIs, consider the following:

- Should be considerations for the long-term
- The definition of what they are and how they are measured should not change often
- Must be accurately defined and measured
- Must have quantifiable targets and time frames

The following list illustrates a less-than-good example of a Business Continuity KPI:

- **Title of Key Performance Indicator**: Improve Backup and Recovery Window
- **Defined**: Reduce recovery times from month to month
- **Measured**: Speed of recovery
- **Target**: Improve each month

So, what is missing from this Business Continuity KPI? Consider the following questions:

1. Does this measure ask us to improve backup and recovery time by days, minutes, or hours?
2. If by minutes, what are we measuring? Does it measure planned outages or unplanned outages?
3. Are we talking about time from the application quiesce to resumption of the application? Are we speaking of RTO, RPO, both, or neither?
4. How do we test and validate this speed of recovery?
5. How much, by percentage or time, do we want to improve the backup and recovery time each month?

Now, here is a good example of a Business Continuity KPI.

- **Title of Key Performance Indicator:**
  - Reduce elapsed time for Backup Window
- **Defined:**
  - Daily planned outage, the total amount of time in minutes that Business Processes 1, 3, and 7 need to be quiesced in order for the daily system backups to be made.
  - The window starts when the Business Processes pause accepting live transactions from a system-wide user standpoint and ends when the business processes resume.
- **Measured:**
  - Each nightly backup window is measured and documented in the Systems Management and Performance application.
  - Associated data such as amount of data backed up, amount of other workload on system that could influence the backup window, are also measured and documented.
  - Trend analysis is automatically performed and the results posted on a weekly basis to the employee billboard.
  - Improvements are tested at the development site, targets set, and then validated as soon as the improvement reaches production.
- **Target:**
  - Reduce backup window to:
    - 30 minutes by 4Q06
    - 10 minutes by 1Q07
    - 1 minute by 2Q07

Starter set of KPIs for IT Business Continuity

Figure 3-14 shows a template for the metrics of Business Continuity KPIs, including a suggested starter set.

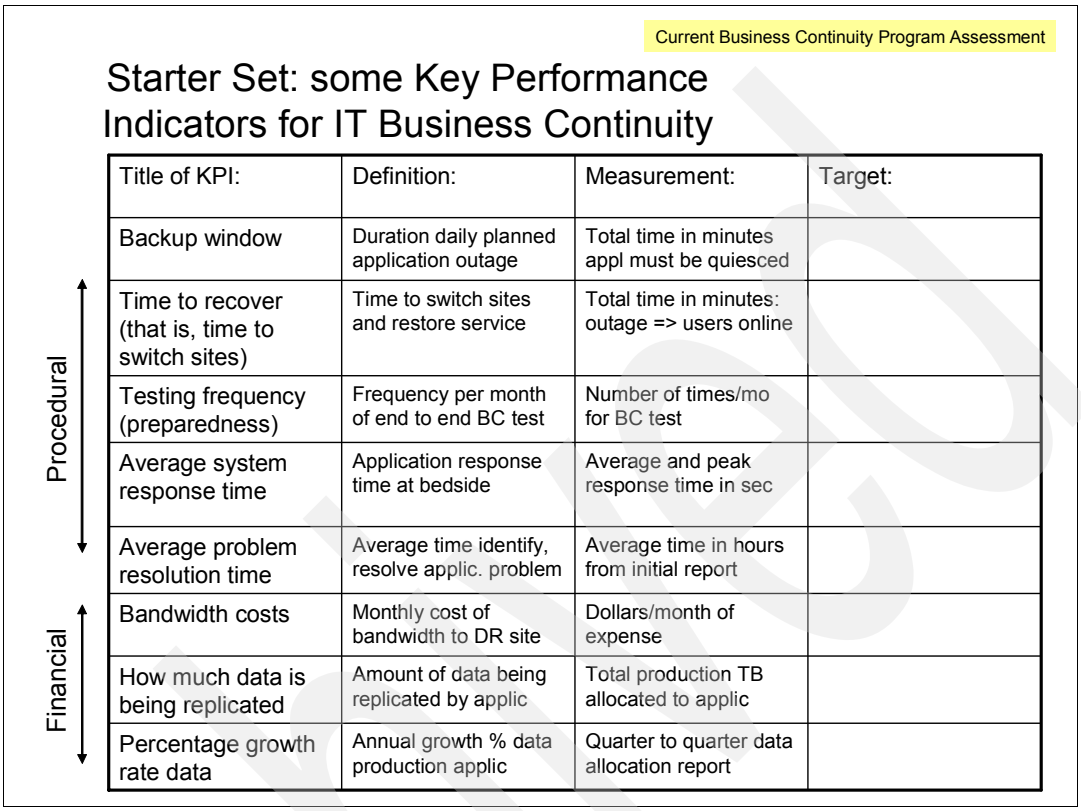

Title of KPI:	Definition:	Measurement:	Target:
Backup window	Duration daily planned application outage	Total time in minutes appl must be quiesced	
Time to recover (that is, time to switch sites)	Time to switch sites and restore service	Total time in minutes: outage => users online	
Testing frequency (preparedness)	Frequency per month of end to end BC test	Number of times/mo for BC test	
Average system response time	Application response time at bedside	Average and peak response time in sec	
Average problem resolution time	Average time identify, resolve applic. problem	Average time in hours from initial report	
Bandwidth costs	Monthly cost of bandwidth to DR site	Dollars/month of expense	
How much data is being replicated	Amount of data being replicated by applic	Total production TB allocated to applic	
Percentage growth rate data	Annual growth % data production applic	Quarter to quarter data allocation report	

Figure 3-14 Step 1 - document current Key Performance Indicators for Business Continuity

The use of KPIs provides a key management tool. They give everyone in your organization a clear vision of what is important and of what they will need to make happen.

Tip: See Figure 3-63 on page 134 for an applied example of KPIs. Also, note how the KPIs are used in quarterly management briefings and review.

We suggest that you publicize and post the KPIs in places throughout your organization where all employees can have access to them: in the lunch room, on the walls of conference rooms, on your company's intranet, and even on the company's Web site for some of them. Show the target for each KPI, and show the progress toward that target. With good project management, your staff will be motivated and feel pride in reaching the KPI targets.

After you have defined your KPIs, you need collect information about the IT environment and components that underpin the selected business process.

Component inventory information to be collected

For the logistics and preparation for these tasks, we provide templates for organizing the component information that you need to gather. Figure 3-15 lists the IT components about which you need to gather information, so that you can properly assess the Business Continuity program's ability to recover the business processes that you have selected for your scope.

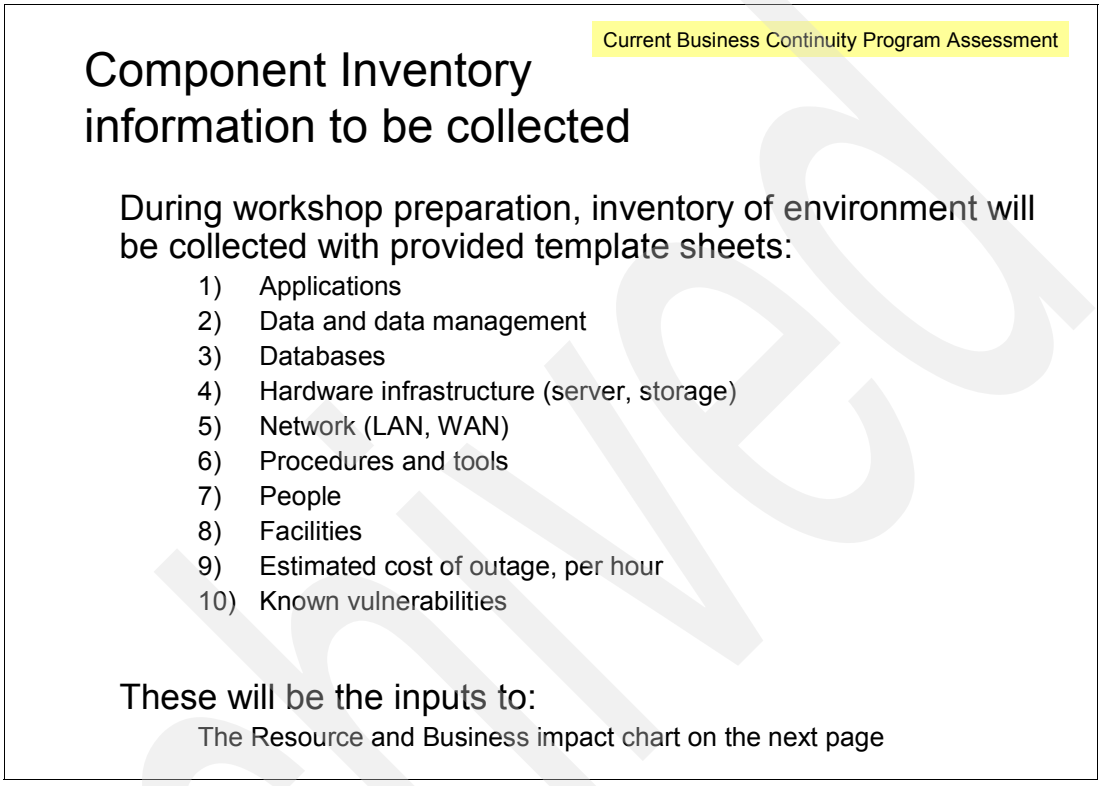

Figure 3-15 Step 1 - component inventory information to be

Here are some important points about this component information. A high-level description is sufficient. You can drill down for more detail as appropriate during the planning process:

- **Applications**: List the various applications that make up the business process.
- **Data and data management**: For these applications, describe at a high level, the current data allocation, data management, data backup policies, and the infrastructure for this data.
- **Databases**: List the databases used to manage the data, and any relevant configuration facts.
- **Hardware infrastructure** (server, storage): List the IT hardware infrastructure on which the business processes, applications, data, and databases reside.
- **Network** (LAN, WAN): Describe the networking infrastructure that connects these business processes.
- **Procedures and tools**: Describe operational procedures, and tools used.
- **People**: What and who are the staff that provides the skills and operates the business procedure. Where are they located, and what are the issues related to them?
- **Facilities**: What are the physical facilities and locations that make up this business process?

- **Estimated cost of outage, per hour**: Describe the estimated cost to the business of this business process is not available
- **Known vulnerabilities**: Describe known vulnerabilities that are desired to be addressed.

Figure 3-16 provides a template for collecting this information.

Current Business Continuity Program Assessment

Scope - Resource and Business Impact

Component	Business process 1	Business process 2	Business process 3
Applications			
Data and data management policies			
Databases			
Servers			
Storage			
Network (LAN, WAN)			
Procedures and tools			
People			
Facilities and physical location			
Known vulnerabilities			
Estimated cost of outage, per hour			
Business Impact			
Ranking			

Figure 3-16 Simplified template for resource and business impact

You can use this simplified template to summarize the affected components for the identified critical business processes, for a basic BIA.

Tip: You can see an applied example of this simplified Resource and Business Impact template in Figure 7-2 on page 239.

Here are some more detailed question examples that you can ask when using this template:

- If an outage occurred, how long could the department function without the existing equipment and departmental organization?
- What are the high-priority tasks, including critical manual functions and business processes in the department? How often are these tasks performed, such as, daily, weekly, monthly, and so on?
- What staffing, equipment, forms, and supplies would be necessary to perform the high priority tasks?
- How would the critical equipment, forms and supplies be replaced in a outage situation?
- Does any of this information require long lead times for replacement?
- What reference manuals and operating procedure manuals, are used in the department? How would these be replaced in the event of a disaster?

- Should any forms, supplies, equipment, procedure manuals or reference manuals from the department be stored in an off-site location?
- Identify the storage and security of original documents. How would this information be replaced in the event of a disaster? Should any of this information be in a more protected location?
- What are the current computer backup procedures? Have the backups been restored? Should any critical backup copies be stored off-site?
- What would the temporary operating procedures be in the event of a disaster?
- How would other departments be affected by an interruption in the department?
- What effect would a disaster at the main computer have on the department?
- What outside services or vendors are relied on for normal operation?
- Would a disaster in the department jeopardize any legal requirements for reporting?
- Are job descriptions available and current for the department?
- Are department personnel cross-trained?
- Who would be responsible for maintaining the department's Business Continuity plan?
- Are there other concerns related to planning for Disaster Recovery?

In addition to these questions, additional recommended data gathering materials and documentation to be included: backup position listing, critical telephone numbers, communications inventory, distribution register, documentation inventory, equipment inventory, forms inventory, insurance policy inventory, main computer hardware inventory, master call list, master vendor list, microcomputer hardware and software inventory, notification checklist, office supply inventory, off-site storage location inventory, software and data files backup/retention schedules, telephone inventory, temporary location specifications, other materials, and documentation.

Identifying current controls

Because a goal of Business Continuity planning is to ensure the safety of personnel and assets during and following a disaster, another important aspect of the current Business Continuity assessment is to identify and assess the preparedness and preventive measures and controls in place at any point-in-time. When the potential areas of high exposure to the organization are identified, additional preventative measures and controls can be considered for implementation.

Assessment of prevention techniques typically includes two categories: Procedural prevention and physical prevention.

- Procedural prevention

  Procedural prevention relates to activities performed on a day-to-day, month-to-month, or annual basis that relate to security and recovery. Procedural prevention begins with assigning responsibility for overall security of the organization to an individual with adequate competence and authority to meet the challenges. The objective of procedural prevention is to define activities necessary to prevent various types of disasters and ensure that these activities are performed regularly.

- Physical prevention

  Business Continuity preparedness begins when a site is constructed. It includes special requirements for building construction, as well as fire protection for various equipment components. Special considerations include: the computer area, fire detection and extinguishing systems, records protection, air conditioning, heating and ventilation,

electrical supply and uninterruptible power supply systems, dual power substation feeds, emergency procedures, vault storage areas, and archival systems.

Assess current recovery time capability

In today's pervasive IT environment, customers, providers, and employees are often clearly aware if an outage has happened to the central computer facility. Usually such an outage is measured in two different non-overlapping components:

- **Service restoration (Recovery Time Objective)**

  Service restoration represents the elapsed time that is experienced from the moment of outage up to the moment when the system has been recovered. This time to recover is typically specified as the *Recovery Time Objective* (RTO).

- **Data loss (Recovery Point Objective)**

  Data loss represents the actual loss of data that is experienced or how much data has to be recreated after the system is recovered to reach the same level of data as before the outage.

Applications and business processes are not all equal, and we find that recovery time objectives and recovery point objectives can and should be categorized and segmented by recovery time.

Another component for consideration is the *Recovery Distance Objective* (RDO) which represents how far away copies of data need to be located. This will be influenced, among other factors, by the risk assessments of the local geography (for example, earthquake fault lines, the likelihood of hurricanes, and so forth) and regulatory requirements.

Although these segments for business process RTO and RPO can vary from company to company, some common guidelines that we use in this IBM Redbook are as follows, to help in qualifying the business process or application recovery time required:

- **Continuous Availability:** RTO < 1 hour after the outage occurred and RPO within 2 minutes of the time of outage
- **Rapid Data Recovery:** RTO < 8-10 hours after the outage occurred and RPO start of the day of the outage
- **Backup/Restore:** RTO > 10 hours after the outage occurred and RPO within 24 hours of the time of the outage

As always, these are general guidelines. You are free to modify and change these suggestions to meet your specific needs.

Summary: Each business unit should declare the maximum time of recovery of their application. There are two pieces of information that should be delivered: The Recovery Time Objective (RTO), how long should it take to recover to the point for a user to be online again; and the Recovery Point Objective (RPO), how much data the user can afford to recreate after the disaster. Usually department personnel are biased in setting higher than needed figures for recovery. The other key piece of information is the Network Recovery Objective (NRO). This is the time to recover the network and move the users to the disaster site. There is no point in recovering the applications in 10 minutes if it takes several hours to connect the users to the new site. It is extremely helpful to develop pre-formatted forms to facilitate the data gathering process.

When you have documented these critical needs, expressed in terms RTO and RPO, IT management has the necessary input to begin assessing the overall recoverability of the organization.

Definition: RTO and RPO

We refer to the diagram in Figure 3-17 multiple times throughout this book. It shows the timeline of an IT recovery and illustrates in detail how the RTO and RPO concepts are defined.

Here are some things to note and to assess your organization's capabilities in the execution of the timeline of an IT recovery:

1. After the outage occurs, the first step is to assess the ability for *management* to *assess* the outage (this incurs elapsed time). Because there is a significant capital cost to declaring a disaster and executing a recovery, management needs to be sure that the situation warrants committing their organization to that expense. After management has decided to declare a disaster, then they activate the Business Continuity plan.
2. The first stage of the Business Continuity plan is to recover the hardware, operating systems, and the data itself. Operations, networking, telecommunications, physical facilities, and associated staff are involved in this process.

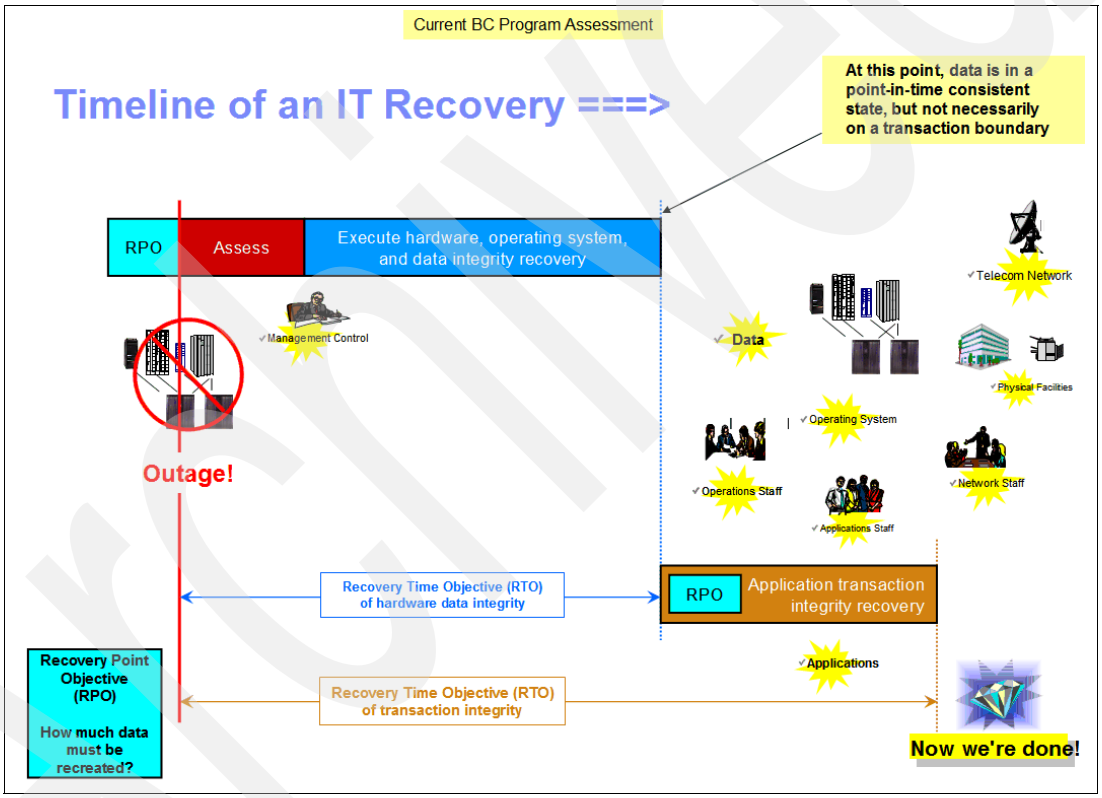

Figure 3-17 Timeline of an IT recovery

3. At the end of this stage, the operating systems and the data are recovered. The data ideally is accurate and data consistent to a point-in-time prior to the outage. The time duration to this point is the RTO of hardware data integrity.
4. However, we are not recovered from a user standpoint. The servers and storage are only capable of accurate *byte movement* (storage) and *proper data write sequencing* (servers and server clustering). Hardware data integrity is not the same as database integrity. The storage and servers cannot know what the logical database relationship is between multiple data blocks in the database. Therefore, when the first stage is complete, the transaction integrity recovery must next be performed by the applications staff, on the application and database.

5. The applications staff performs transaction integrity recovery. Hopefully, this is a database restart and *not* a database recovery. This process will back out incomplete logical units of work, and restore the database to logical integrity as of the most recent time possible.
6. When the transaction integrity recovery (rollback or roll forward) is complete, we now have the application and databases ready for user access. This duration is the RTO of transaction integrity.
7. Note the difference in elapsed time between RTO of hardware data integrity and RTO of transaction integrity. When discussing the RTO, it is important to distinguish which of these two is being referred to. Operations and application staff can have differing perceptions of RTO depending on whether the RTO is assumed to be at the hardware recovery level, or at the application recovery level. The fact that there are these two different RTOs is essential to understanding and planning for heterogeneous Business Continuity.
8. Finally, observe how RPO is depicted in Figure 3-18. RPO (which is the amount of data recreation required prior to the point of outage) is shown as the time offset before the outage occurred. Note that the RPO data recreation happens in the transaction integrity recovery stage. RPO data recreation cannot be done in the hardware and operating system recovery, as the server and storage components do not have knowledge of the logical relationships between multiple application and database blocks of data.

Figure 3-18 provides a template format for documenting the current Business Continuity program assessment. You can collect the relevant information about each of the various steps required in today's recovery plan and document them using this template.

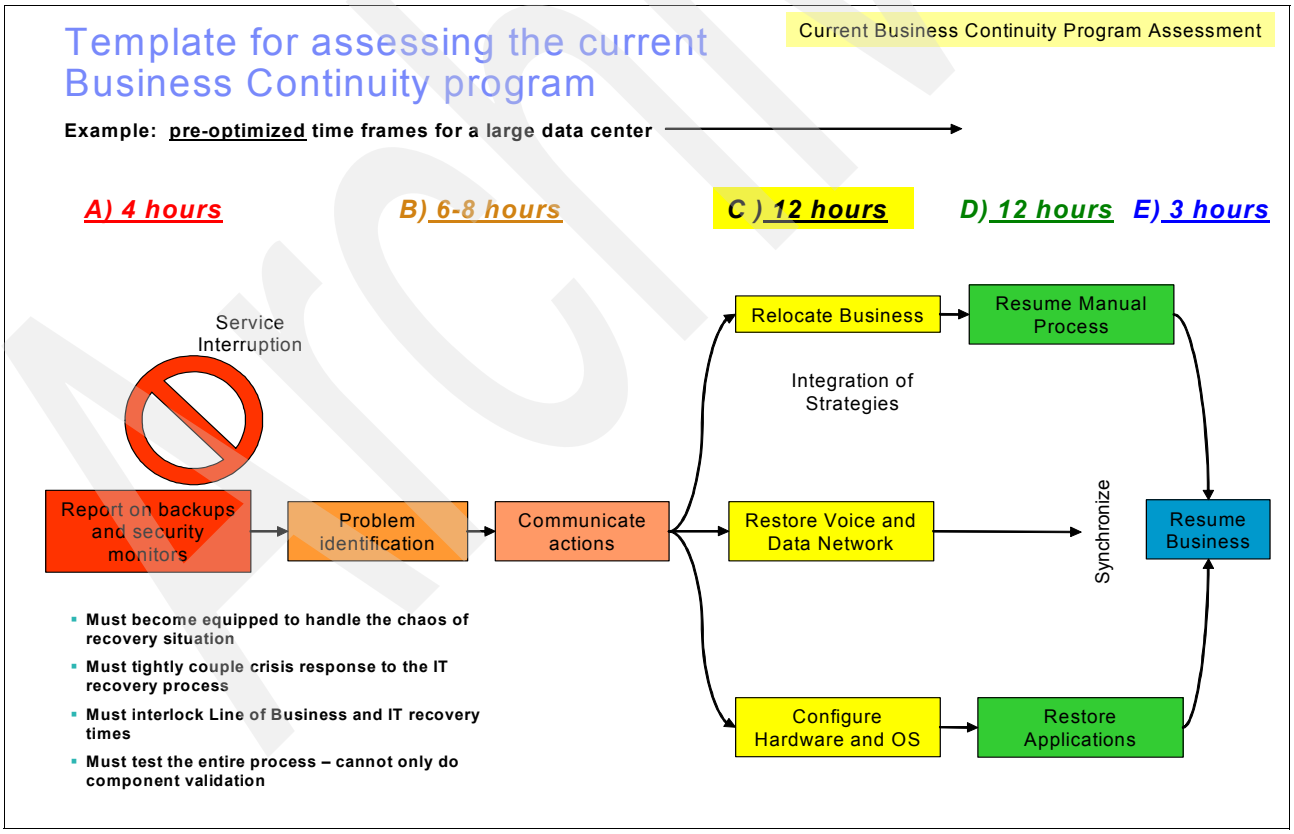

Figure 3-18 Template for assessing the overall Business Continuity program

The template shows that for this business process recovery, the assessment documents the following:

- **Step A - 4 hours**: Triage, problem identification, finding decision makers and decision on declaration.
- **Step B - 6 to 8 hours**: Identify IT and business staff to start recovery, communicate actions and communicate to customers and employees.
- **Step C - 12 hours**: Move tapes to location, contact network provider, relocate staff and begin configuration.
- **Step D - 12 hours**: Application and database redo logs. Begin data integrity validation. Integrate end-user connectivity.
- **Step E - 3 hours**: Validate functionality of critical applications.

You can use these templates as input to the Business Continuity program design phase.

Tip: It is almost certain that you will need new processes and procedures as you implement Business Continuity technology. You should allow for and plan for the time and effort required for these new processes and procedures.

Assess current Business Continuity maturity level

Simultaneously with all of these inputs, it is important to realize that you also need to assess the current *maturity level* of the Business Continuity program. It is normal for an organization to mature over time. Being at a lower or higher maturity level is not of a matter of good or bad, it is simply a point on a continuum. Your objective is to assess where you are, so that you can best know what steps to take to move ahead over time.

Figure 3-19 expresses typical phases in Business Continuity program maturity.

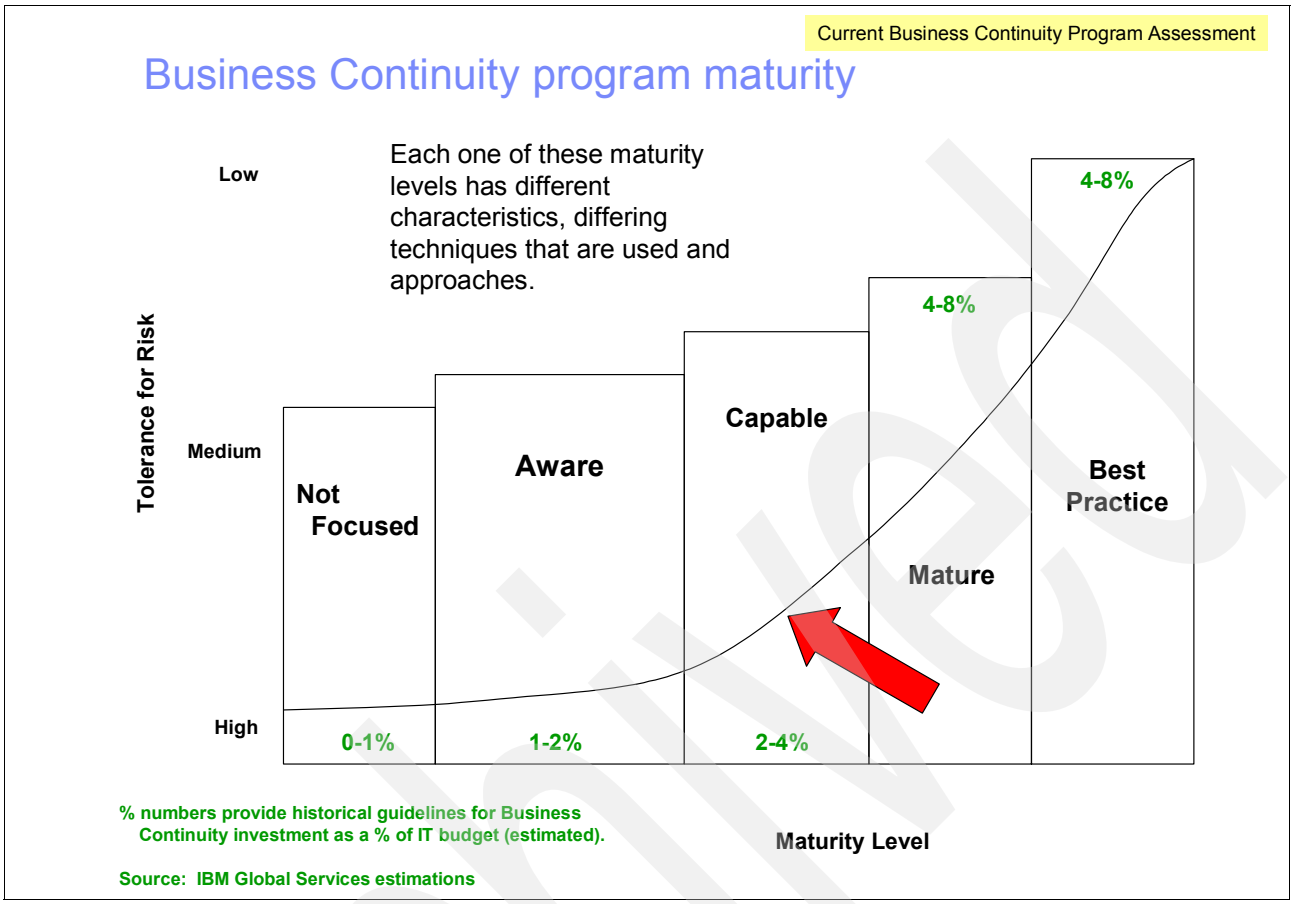

Figure 3-19 Business Continuity program maturity levels

Definitions and good guidelines for the level of Business Continuity maturity are:

- **Not focused**: Program is not well defined; processes are not developed or followed; change management between production and recovery needs improvement.
- **Aware**: Knowledgeable about disaster recovery and High Availability; Business Continuity plans documented; strategies developed for some systems.
- **Capable**: Recovery requirements are clearly understood; Business Continuity program in place; comprehensive strategies and plans.
- **Mature**: Regular testing; executive understanding; good crisis management process; in compliance; governance model in place.
- **World class**: Documented, validated and tested in an integrated manner, ITIL® methods implemented, environment is monitored, High Availability solutions have been implemented for top tier; compliant and auditable.

KPI: Note that as shown in Figure 3-19, as the maturity level increases, the percentage investment in Business Continuity, as a percentage of the overall IT budget, increases. You can compare your Business Continuity budget percentages with these guidelines to see how your budgetary expenditures compare.

Assess funding and return on investment (ROI)

In this step, we assess the funding levels and justification models for the current Business Continuity program funding. The objective is to determine how the funding is done, the criteria for existing or new funding, and whether Business Continuity is viewed as a cost, a competitive advantage, or something in between.

Note: Business Continuity projects should not be considered in the same ROI model as other traditional IT projects. Other IT projects are designed to bring a specific return to the business on an ongoing basis. However, by definition, Business Continuity programs are designed as preventative measures that provide value only if there is a successful recovery or an avoidance from an outage.

The purpose of this step is to understand the current funding levels and justification models and to compare them with the best practices. The outcome of this step, therefore, is to understand the gap that exists and to begin to evolve your Business Continuity program towards those best practices.

Tip: You can review a detailed discussion on best practices for Business Continuity funding and ROI in detail in 2.2, “Justifying Business Continuity to the business” on page 12.

Roadmap status for the Business Continuity program

In this assessment step, you document the current roadmap of improvements for the Business Continuity plan, because this is one of the inputs into the next step, Business Continuity program design. Obviously, if the existing Business Continuity roadmap is robust or in need of improvement, you use that as input and take appropriate steps in the Business Continuity program design.

3.4.4 Summary of Business Prioritization

You have now completed the first section of the ideal Business Continuity planning process. You completed the following important steps:

1. **Risk assessment**: You identified the set of risks and vulnerabilities that you will address in your plan and created an initial scope for your Business Continuity plan.
2. **Business impact analysis**: You took the identified risks and vulnerabilities and ranked them according to the needs of your business, thus further defining and refining the scope of your Business Continuity plan.
3. **Program assessment**: You assessed your current Business Continuity program to understand your current baseline status from which you will build the enhanced Business Continuity program, using the enhanced scope defined by the risk assessment and business impact analysis.

You now have a rich set of documented input for the next section in our planning process, and you can move from Business Prioritization to Integration into IT.

3.5 Integration into IT

Figure 3-20 shows where Integration into IT fits in the ideal Business Continuity planning process.

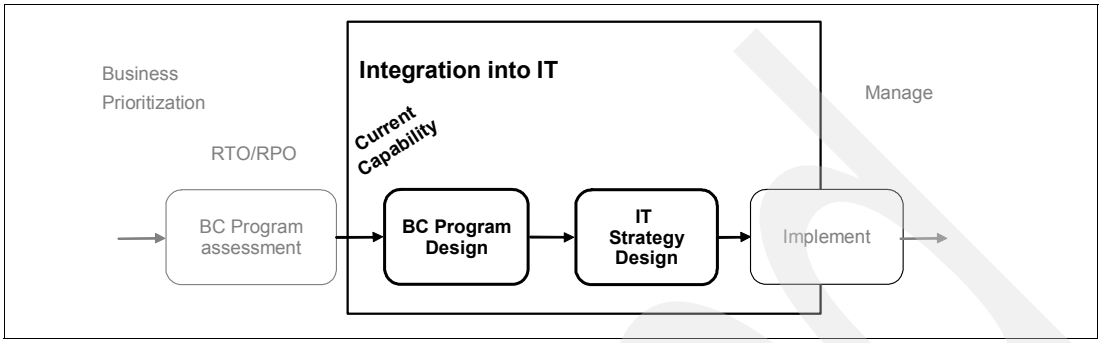

Figure 3-20 Integration into IT - overview of steps

In this section, we take the input from Business Prioritization (specifically, the Business Continuity program assessment) and perform overall end-to-end Business Continuity program design. This program design then provides the requirements for which you complete the IT Business Continuity strategy design.

3.5.1 Business Continuity program design

You begin the process with Business Continuity program design, as shown in Figure 3-21.

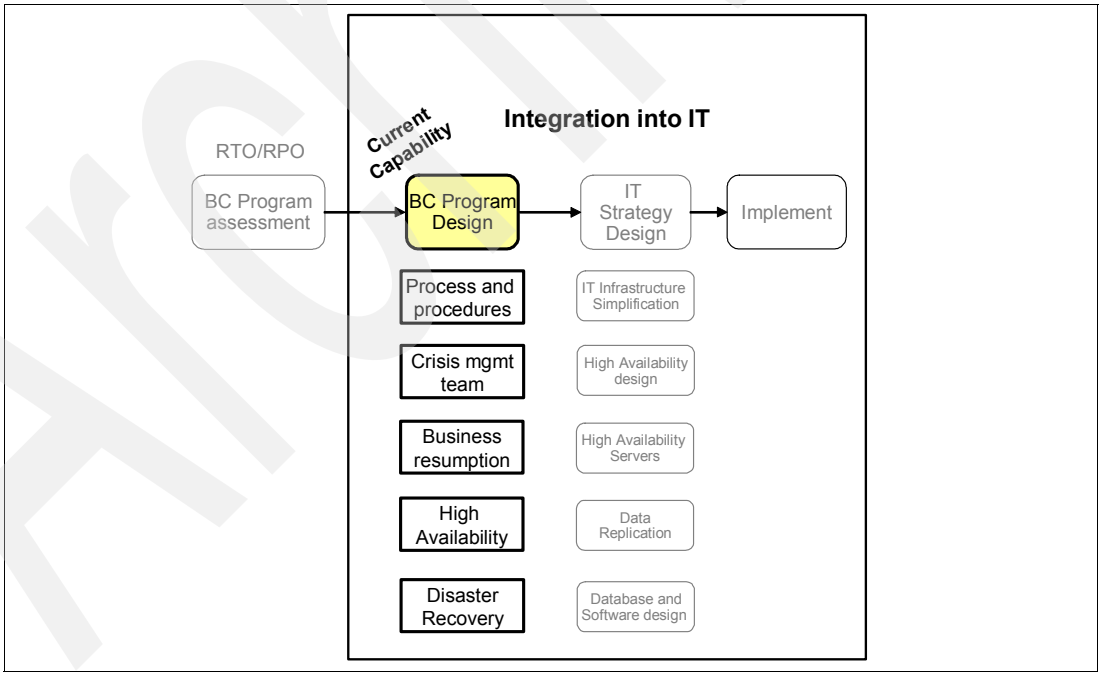

Figure 3-21 Integration into IT - Business Continuity program design

In this step, you design or enhance your existing end-to-end Business Continuity program. Business Continuity program design is differentiated from IT strategy design in that you include all aspects of the business, particularly those *outside of IT*, including:

- Business processes and procedures (non-IT)
- Crisis team management (IT and non-IT)
- Definitions of how non-IT portions of the business will resume operation
- The external, non-IT *business process* aspects of High Availability and Disaster Recovery

In this section, we examine each of the steps shown in Figure 3-21 in detail.

Process and procedures

In addition to IT technology, the Business Continuity plan also covers people, processes, and procedures. Figure 3-22 is an example of the categories of questions that the plan addresses.

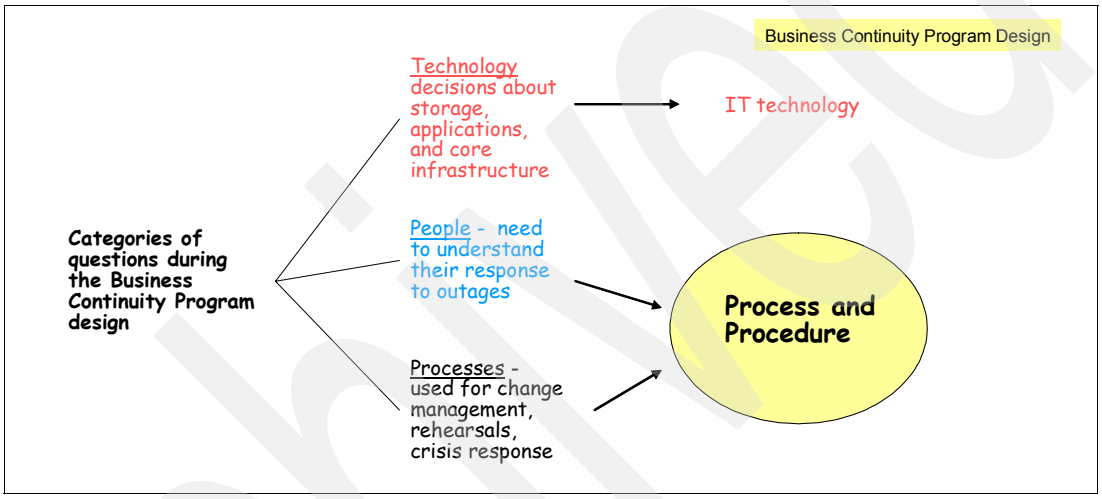

Figure 3-22 Business continuity program - categories of questions that are addressed

The primary design objective of the end-to-end Business Continuity plan is to answer the people and process questions in your Business Continuity program design.

Note: You primarily address the technology questions in the IT strategy design portion of your planning, as discussed in 3.5.2, “IT strategy design” on page 89.

Before you can complete the IT technology design, however, you must consider the important people and process-related questions and the requirements of the Business Continuity program, as illustrated in Figure 3-23.

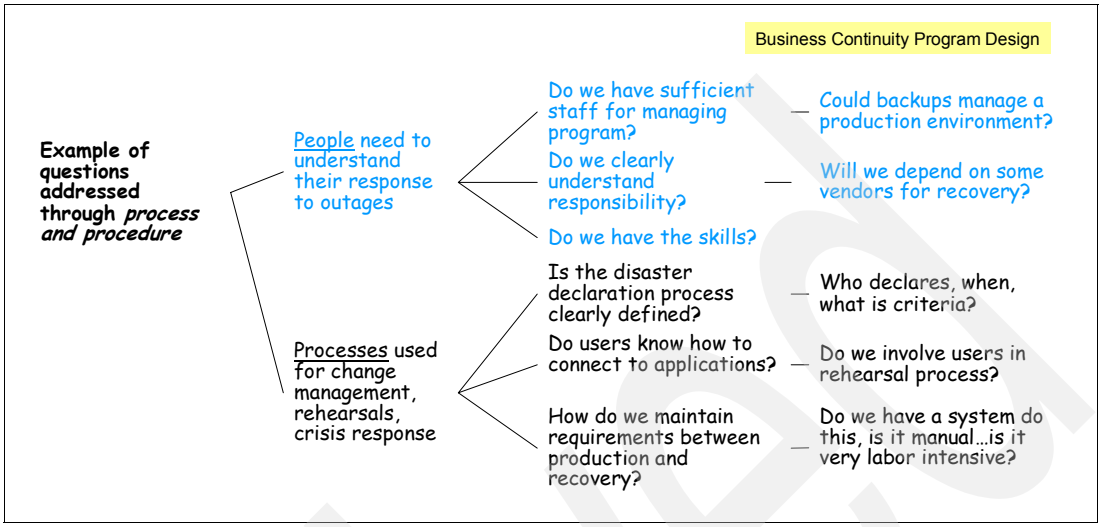

Figure 3-23 Example of questions addressed through process and procedure enhancements

Let us go into more detail about these questions.

Introduction to Business Continuity plan process and procedure

The objective of this step of the ideal Business Continuity planning process is to design the people-related processes and procedures necessary to reliably and repeatably protect the organization in the event that all or part of its operations or computer services are rendered unusable. A Business Continuity plan establishes the procedures and actions to be done when exercising functions for local High Availability and Continuous Operations procedures, as well as Disaster Recovery functions done only in the event of an unplanned outage. Clearly, it is fundamental that you can *manage* an incident successfully, assisted by a high-quality, testable, repeatable, and reliable set of processes and procedures.

You design the process and procedure first, so that you can then select technology that can best serve the desired process and procedure in the Business Continuity plan.

The Business Continuity plan involves more than off-site storage or backup processing actions. It also includes all non-IT related functions necessary to restore operations. As such, this phase of the planning will design and specify the process and procedures that will address all the critical operations and functions of the business. The plan should include documented and tested procedures, which, if followed, ensure the ongoing availability of critical resources and continuity of operations.

Important: Process and procedure is critically important, because as the planner, you want to know precisely how your staff and the IT systems will react to a disaster which does strike. Predictability of the *reaction to a disaster* is the goal, because it is impossible to predict the disaster itself. This can only be accomplished by having a combination of automation functions and well documented and regularly tested procedures. As a result, affordable frequent testing is a KPI. We cannot wait until a disaster occurs to determine whether the plan will work.

Preparedness is the key. Properly done, the Business Continuity planning process will minimize the disruption of operations and ensure an acceptable level of organizational stability and an orderly recovery after a unplanned outage.

Business Continuity program plan scope

The Business Continuity plan applies to all events (both planned and unplanned), that could deny access to the normal IT infrastructure for an extended period. Planned outages should be included as well (an example would be the need for the a workload rotation to an alternate site, to enable an orderly shutdown of the primary site for a major upgrade).

A Business Continuity plan design aims to provide a comprehensive statement of consistent actions to be taken *before*, *during*, and *after* a disaster. As we saw earlier in 3.4.2, “Business impact analysis” on page 51, the BIA defines the *scope* for our Business Continuity program design.

You then use this scope to define what level of planning, process, and procedure you need to design and test to ensure the continuity of operations and availability of critical resources. Usually these business processes have required years to create and establish, but in the event of an outage, management must be able to reestablish these functions within hours or days.

This is a difficult problem, and reestablishing the complex business environment in a timely manner requires a well thought-out plan in place ready to be executed. This plan also predicts actions to be executed after the reestablishment of the damaged site, in order to return the original situation.

Role of risk management

Risk management encompasses a broad range of activities to identify, control, and mitigate risks to a business or to an IT system. Risk management activities can be considered to have two primary functions:

- **Preventative:** Actions to prevent or reduce the likelihood of damaging incidents by reducing or eliminating risks. These preventative measures typically form the security controls and High Availability features that protect and mask a system and its users from feeling the effects of natural, human, and technological events or threats.
- **Recovery:** Actions to reduce or limit the consequences of threats in the event that they successfully disrupt a system. These measures are developed in anticipation of a possible event, are executed after that event has occurred, and form the basis for the Business Continuity plan.

Figure 3-24 illustrates the inter-relationship between risk management, security control implementation, and Business Continuity planning.

Figure 3-24 Risk management as an element of the Business Continuity program design

Data protection

Data protection refers to all the organization's measures to safeguard assets, ensure the accuracy and reliability of records, and encourage operational efficiency and adherence to prescribed procedures. The system of internal controls also includes the measures adopted to safeguard computer systems and to determine which data needs which level of protection.

The nature of internal controls is such that certain control procedures are necessary for a proper execution of other control procedures. This interdependence of control procedures might be significant because certain control objectives that appear to have been achieved might, in fact, not have been achieved because of weaknesses in other control procedures upon which they depend.

Security is an increasing concern because computer systems are increasingly complex. Particular security concerns result from the proliferation of PCs, wireless networks, and online systems that allow more access to the various servers. Modern technology provides computer thieves with powerful new infiltration tools.

Important areas of concern related to general computer internal controls include: organization controls, systems development and maintenance controls, documentation controls, access controls, data and procedural controls, physical security, password security systems, and communications security.

Insurance coverage

A Business Continuity Plan might have insurance as one of the tools in place to mitigate the effects of an unplanned outage. Insurance can provide a certain level of confidence in knowing that if a major catastrophe occurs, it will not result in financial or organizational disaster. However, just having an insurance policy by itself is not enough, because it cannot compensate sufficiently for the incalculable loss of business during the interruption or for the business that goes bankrupt because of the disaster.

Adequate insurance coverage is a key tool to consider when developing a Business Continuity plan. Although having a Business Continuity plan and testing it regularly might not, in itself, lower insurance rates in all circumstances, a good plan can reduce risks and address many concerns of the underwriter.

Most insurance companies specializing in business interruption coverage can provide the organization with an estimate of anticipated business interruption costs. Many organizations

that have experienced a disaster indicate that their costs were significantly higher than expected in sustaining temporary operations during recovery.

Most business interruption coverages include lost revenues following a disaster. Extra expense coverage includes all additional expenses until normal operations can be resumed. However, coverages differ in the definition of resumption of services. As a part of the risk analysis, these coverages should be discussed in detail with the insurer to determine their adequacy and applicability.

To provide adequate proof of loss to an insurance company, the organization might need to contract with a public adjuster who can charge between 3% and 10% of recovered assets for the adjustment fee. Asset records become extremely important as the adjustment process takes place.

Types of insurance coverages to be considered include: computer hardware replacement, extra expense coverage, business interruption coverage, valuable paper and records coverage, errors and omissions coverage, fidelity coverage, and media transportation coverage.

With estimates of the costs of these coverages, management can make reasonable decisions on the type and amount of insurance to carry. These estimates also allow management to determine to what extent the organization should self-insure against certain losses.

Data and records classification

In the process of this planning, we will consequently identify the types of *data and records* that are the major critical resources needed by the business processes. Clearly, just like business processes, not all data nor records have the same level of recovery requirement. Some vital data and records are irreplaceable. Some critical data and records can be obtained or reproduced at considerable expense, and only after considerable delay. Some non-critical data and records would cause inconvenience if lost but can be replaced without considerable expense.

As part of the planning process, we identify how vital and critical records and data should be duplicated, and how and where they should stored, in an area protected from the defined risks and threats identified in the risk analysis and BIA steps of the planning process.

All non-electronic data and records that are pertinent to the recovery must be included and classified as well. For these records, this could range from gathering non-electronic papers hastily and exiting quickly, to an orderly securing of documents in a fire-proof vault. Identifying data, records and information properly, and then keeping that identification current, is part of Business Continuity planning for ensuring the continuity of business operations.

Today's compliance environment dictates that the appropriate data and records will also need to be affordably retained in a way that insures compliance with legal and statutory requirements. In addition, businesses must also satisfy retention requirements as an organization and employer. These kinds of data and records are used for independent examination and verification of sound business practices.

Federal and state requirements for data record retention must be analyzed. Each organization should have its legal counsel approve its own retention schedule.

When replicating data in real time using today's data replication techniques, data integrity of the remote site data is also a major consideration, because of the common type of disaster that has the most impact: the **rolling disaster**. The rolling disaster is defined as a disaster that unfolds along a significant amount of time (as in fire or water flood). Modern IT

technology recognizes the existence and pervasiveness of the rolling disaster issue in Business Continuity, and techniques should be applied to prevent the effects. You can find more information about the rolling disaster and techniques that you can use to mitigate it in “Rolling disasters” on page 259.

Types of Business Continuity contingency plans

In this section, we discuss different types of Business Continuity contingency planning and what the objectives of these plans should be.

The Business Continuity program design represents a broad scope of activities designed to sustain and recover critical business processes and IT services following an emergency. Business Continuity program design covers a broad range of emergency preparedness environments, including Organizational and Business Process Continuity and Recovery Planning. Ultimately, an organization would use a suite of plans to properly prepare response, recovery, and continuity activities for disruptions affecting the organization's IT systems, business processes, and the facility. Because there is an inherent relationship between an IT system and the business process it supports, there should be coordination between each plan during development and updates to ensure that recovery strategies and supporting resources neither negate each other nor duplicate efforts.

In general, universally accepted definitions for contingency planning and these related planning areas have not been available. Occasionally, this has led to confusion regarding the actual scope and purpose of various types of plans. To provide a common basis of understanding regarding Business Continuity program planning, this section identifies several other types of plans and describes their purpose and scope relative to Business Continuity program planning. Because of the lack of standard definitions for these types of plans, in some cases the scope of actual plans developed by organizations can vary from the descriptions that we include here.

Business Continuity planning confronts the likelihood of a disaster, how the disaster interrupts the business process, and how the business can continue in operation. An interruption could be something related to a winter storm, the loss of electricity to the general area, or the complete inaccessibility of a facility for an extended period of time. The cause of the interruption does not matter; what matters is gaining management control and processing capacity just after the interruption.

Business Continuity planning relates to considerations that effect the whole of the business process. If a building is unusable where will workers be relocated? How will their personal business tools be replaced? How will these workers access the system by the network?

The Business Continuity plan (BCP) focuses on sustaining an organization's business functions during and after a disruption. An example of a business function might be an organization's payroll process or consumer information process. A BCP can be written for a specific business process or can address all key business processes. Information systems are considered in the BCP only in terms of their support to the larger business processes. In some cases, the BCP might not address long-term recovery of processes and return to normal operations, but only cover interim Business Continuity requirements.

We discuss a variety of necessary contingency plans which make up the component parts of the BCP. The inter-relationship of these plans can be diagrammed as shown in Figure 3-25.

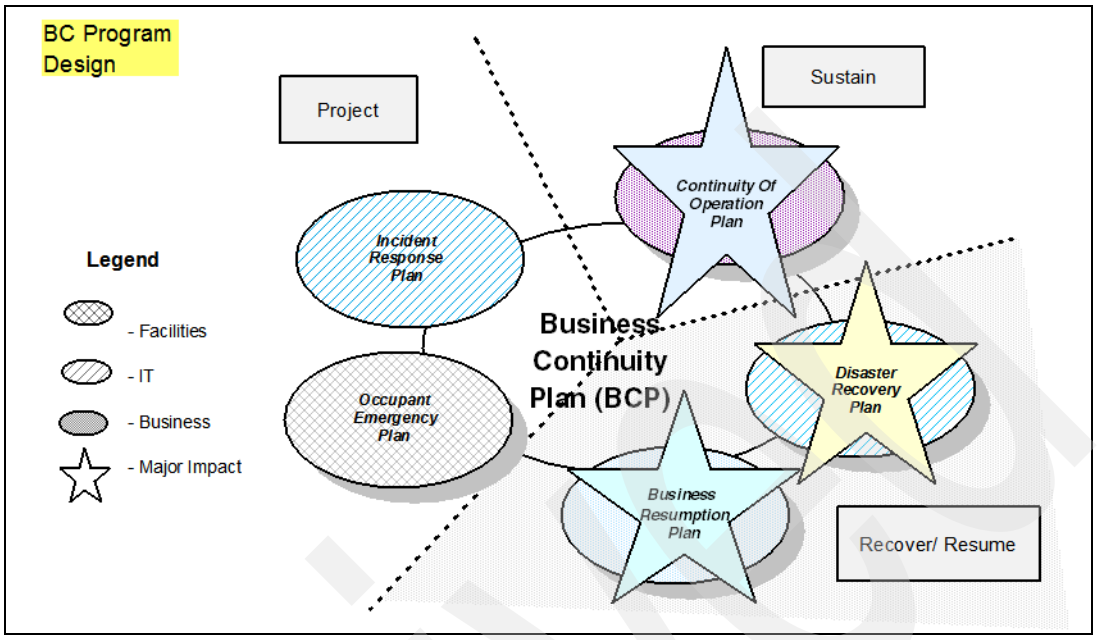

Figure 3-25 Interrelationship of contingency plans

Let us now discuss definitions of these common types of contingency plans. It is likely that the BCP should have most if not all of these in place.

- Business Resumption Plan (BRP)

  The BRP addresses the restoration of business processes after an emergency. The BRP is similar to the BCP, but unlike that plan, the BRP typically lacks procedures to ensure continuity of critical processes *during* an emergency or disruption.

- Continuity of Operations Plan (COOP)

  The COOP focuses on restoring an organization's (usually the headquarters element) essential functions at an alternate site and performing those functions for up to 30 days before returning to normal operations.

- Incident Response Plan (IRP)

  The IRP establishes procedures to address cyber-attacks against an organization's IT server and workstation systems. These procedures are designed to enable security personnel to identify, mitigate, and recover from malicious computer incidents, such as unauthorized access to a system or data, denial of service, or unauthorized changes to system hardware or software (e.g. malicious logic such as a virus, worm, or Trojan horse).

- Occupant Emergency Plan (OEP)

  The OEP provides the response procedures for occupants of a facility in the event of a situation posing a potential threat to the health and safety of personnel, the environment, or property. Such events would include a fire, hurricane, criminal attack, or a medical emergency. OEPs are developed at the facility level, specific to the geographic location and structural design of the building.

- Disaster Recovery Plan (DRP)

  As suggested by its name, the DRP applies to major events that deny access to the normal facility for an extended period. Frequently, the DRP refers to an IT-focused plan designed to restore operability of the target system, application, data, or computer facility

at an alternate site after an emergency. The DRP scope can overlap that of a BCP; however, the DRP is narrower in scope and does not address minor disruptions that do not require relocation.

Nevertheless, logical data errors such as viruses, have no need for relocation, but can also be seen as a big disaster for a client. These cases are covered in the IRP and are very much linked with the design of disaster recovery solutions.

Crisis management team

A major part of the processes and procedures is to ensure the *management* of the organization and its resources during the recovery. Major aspects of designing a good Business Continuity crisis and management team include:

- Top management commitment

  Top management must support and be involved in developing, coordinating, and ensuring the disaster recovery effectiveness within the organization. Adequate time and resources must be committed to the development of an effective plan. Resources could include both financial considerations and the effort of all personnel involved.

- An established planning committee

  You need to appoint a planning committee to oversee the development and implementation of the plan, and the planning committee must include representatives from all functional areas of the organization. Key committee members should include the operations manager and the data processing manager. The committee also should define the scope of the plan. The committee should inform top management regularly, because of the very sensitive nature of the subject and because the expenses (in money and people) can be quite high.

- A defined scope

  Although most Business Continuity plans address only data processing related activities, a comprehensive plan will also include areas of operation outside data processing. The plan should have a broad scope if it is to effectively address the many disaster scenarios that could affect the organization.

Assumptions about the crisis management team

When developing the plan, you need to consider a *worst case scenario* where the main or primary facility is destroyed. Because the plan is written based on this premise, less critical situations can be handled by using only the needed portions of the plan, with minor (if any) alterations required.

Every BCP has a foundation of assumptions on which the plan is based. The assumptions limit the circumstances that the plan addresses. The limits define the magnitude of the disaster the organization is preparing to address. The assumptions can often be identified by asking the following questions:

- What equipment or facilities have been destroyed?
- What is the timing of the disruption?
- What records, files, and materials were protected from destruction?
- What resources are available following the disaster:
  - Staffing?
  - Equipment?
  - Communications?
  - Transportation?
  - Hot site or alternate site?

When writing the BCP, consider the following typical planning assumptions:

- The main facility of the organization has been destroyed.
- Staff is available to perform critical functions defined within the plan.
- Staff can be notified and can report to the backup site or sites to perform critical processing, recovery, and reconstruction activities.
- The plan is current.
- Subsets of the overall plan can be used to recover from minor interruptions.
- An alternate facility is available.
- An adequate supply of critical forms and supplies are stored off-site, either at an alternate facility or off-site storage.
- A backup site is available for processing the organization's work.
- The necessary long-distance and local communications lines are available to the organization.
- Surface transportation in the local area is possible.
- Vendors will perform according to their general commitments to support the organization in a disaster.

This list of assumptions is not all-inclusive, but it is intended as a thought-provoking process in the beginning stage of planning. The assumptions themselves will often dictate the makeup of the plan, therefore, management should carefully review them for appropriateness.

Teams within the crisis management team and their responsibilities

You need to construct your BCP using a *team approach*. You need teams responsible for administrative functions, facilities, logistics, user support, computer backup, restoration, and other important areas in the organization. Then, you can assign specific responsibilities to the appropriate team for each functional area of the company.

The structure of the contingency organization might not be the same as the existing organization chart. The contingency organization is usually structured in teams that are responsible for major functional areas. For example, the teams can include:

- Management team
- Business recovery team
- Departmental recovery team
- Computer recovery team
- Damage assessment team
- Security team
- Facilities support team
- Administrative support team
- Logistics support team
- User support team
- Computer backup team
- Off-site storage team
- Software team
- Communications team
- Applications team
- Human relations team
- Marketing/customer relations team

It is not necessary for the organization to have separate teams created exactly as those that we list here, but we strongly recommend that you create functions that are related to each of these organizational teams.

You need to chose personnel to staff these teams based on their skills and leadership. Ideally, teams are staffed with the personnel that are responsible for the same or similar operation under normal conditions (for example, the computer recovery team members include the server administrators). Team members must understand not only the contingency plan purpose but also the procedures necessary for executing the recovery strategy. Teams should be sufficient in size to remain viable even if some members are unavailable to respond (such as, due to vacations), or spare team members can be designated. Similarly, team members should be familiar with the goals and procedures of other teams to facilitate inter-team coordination.

Each team is led by a leader who directs overall team operations and acts as the team's representative to management and liaisons with other team leaders. The team leader disseminates information to team members and approves any decisions that must be made within the team. Team leaders should have a designated substitute to act as the leader if the primary leader is unavailable.

The most important team in disaster recovery planning is the management team, which provides overall guidance following a major system disruption or emergency. The management team is typically led by a senior management official, such as the Chief Information Office (CIO), who has the authority to make decisions regarding spending levels, acceptable risk, and intercompany coordination.

Major functions of such a team are:

- Responsible for activating a contingency plan and supervising the execution of contingency operations.
- Facilitates communications among other teams and supervises plan tests and exercises.
- Takes the lead in specialized contingency teams.
- Coordinates the recovery process.
- Assesses the disaster, activates the recovery plan, and contacts team managers.
- Oversees, documents, and monitors the recovery process.
- Makes final decision in setting priorities, policies, and procedures.

Another important team is the disaster recovery management team, who is called into action under the authority of the top administrative committee.

Notification and activation procedures

As a part of crisis management, we need to highlight the need for good notification and activation procedures, which are the initial actions taken after a system disruption or emergency has been detected or appears to be imminent. This phase includes activities to notify recovery personnel, assess system damage, and implement the plan. At the completion of this phase, recovery staff will be prepared to perform contingency measures to restore system functions on a temporary basis.

An event can occur with or without prior notice. For example, advanced notice is often given that a hurricane will affect an area or that a computer virus is expected on a certain date. However, there might be no notice of equipment failure or a criminal act. Notification procedures should be documented in the plan for either type of situation. The procedures should describe the methods used to notify recovery personnel during business and non-business hours.

Prompt notification is important for reducing the effects on the IT system, and in some cases, it can provide enough time to allow system personnel to shut down the system gracefully to avoid a hard crash. Following the disaster event, notification should be sent to the team responsible for the damage assessment, so that it can determine the status of the situation and the appropriate next steps. When damage assessment is complete, the appropriate recovery and support teams should be notified.

Notifications can be accomplished through a variety of methods, including telephone, pager, work or personal e-mail, or cell phone. Notification tools that are effective during widespread disasters include radio and television announcements and Web sites. The notification strategy should define procedures to be followed in the event that certain personnel cannot be contacted. You need to document notification procedures clearly in the Business Continuity plan.

A common notification method is a *call tree*. A call tree involves assigning notification duties to specific individuals, who in turn are responsible for notifying other recovery personnel. The call tree needs to account for primary and alternate contact methods and needs to discuss procedures to be followed if an individual cannot be contacted. You need to identify personnel to be notified clearly in the contact list that is appended to the plan. This list should identify personnel by their team position, name, and contact information, including home, work, and pager numbers, e-mail addresses, and home addresses. Additionally, the primary personnel might not be available due to the disaster, so you need to contract for alternate personnel at the recovery site to execute the Business Continuity plan.

Note: Be aware, that with some incidents of a physical site disaster, telephone lines or cellular telephone networks can be damaged or overloaded for hours or even days. Thus, you need to consider alternative plans for notification.

Notification should also be sent to points of contact (POCs) for external organizations or interconnected system partners that might be adversely affected if they are unaware of the situation. Dependent on the type of disruption, the POC might have recovery responsibilities. Therefore, for each system interconnection with an external organization, a POC should be identified to the extent that the organizations will assist each other and the terms under which the assistance will be provided. These POCs should also be listed in an appendix to the plan.

Figure 3-26 provides one example of a call tree that defines the chain of information flow in the event of a crisis.

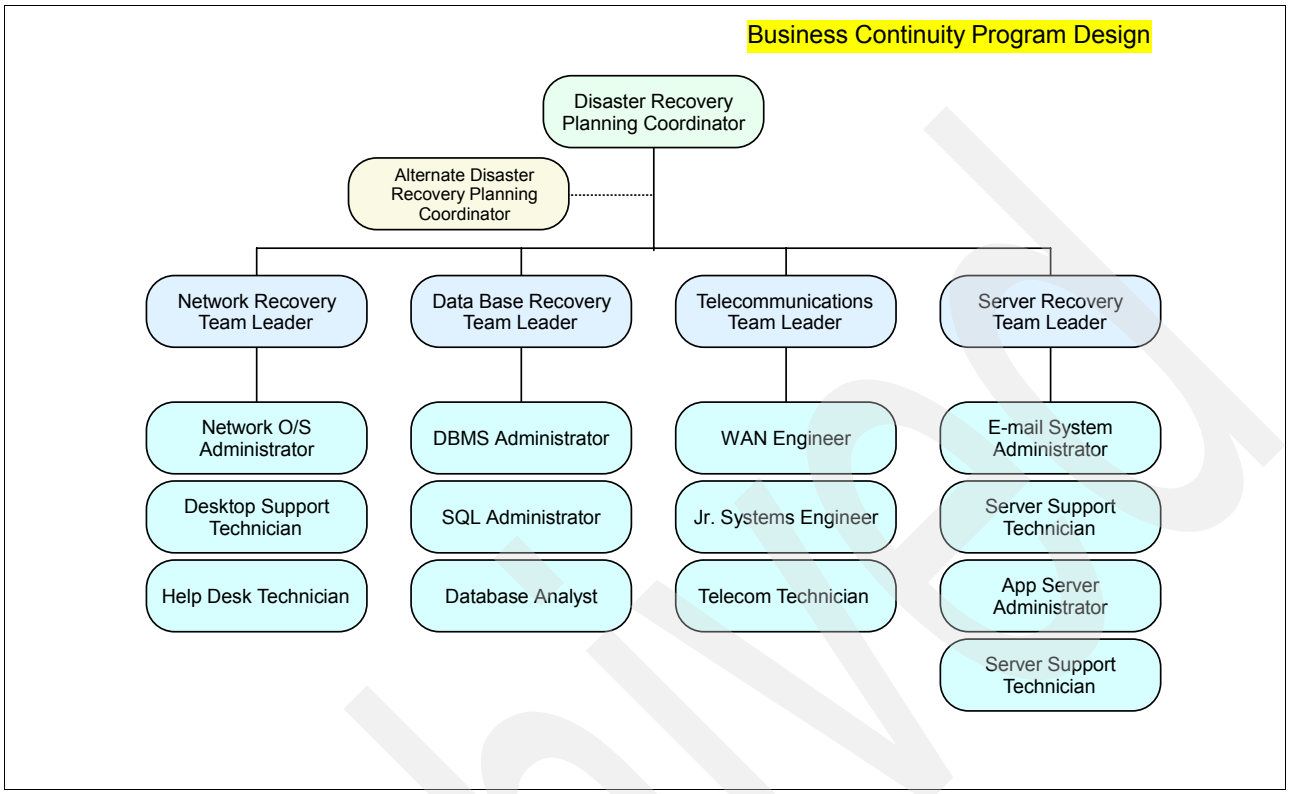

Figure 3-26 Notification: A sample call tree

You need to document in the plan the type of information that is relayed to those notified. The amount and detail of information relayed can depend on the specific team that is notified. As necessary, the notification information can include:

- Nature of the incident that has occurred or is impending
- Loss of life or injuries
- Any known damage estimates
- Response and recovery details
- Where and when to convene for briefing or further response instructions
- Instructions to prepare for relocation for estimated time period
- Instructions to complete notifications using the call tree (if applicable)

Activation plan

The Business Continuity plan should be activated only when the damage assessment indicates that some of the activation criteria for that system are met. If an activation criterion is met, the disaster recovery planning coordinator should activate the plan. Activation criteria for events are unique for each organization and should be stated in the disaster recovery planning policy statement and can be based on:

- Safety of personnel or extent of damage to the facility
- Extent of damage to system (physical, operational, or cost)
- Criticality of the system to the organization's mission (such as, critical infrastructure protection asset)
- Anticipated duration of disruption

After the system damage has been characterized, the disaster recovery planning coordinator can select the appropriate recovery strategy, and the associated recovery teams can be notified.

Business Resumption

The recovery and restoration plan phase deals with recovery strategies that provide a mean to restore IT operations quickly and effectively following a service disruption. Specifically, these plans address:

- What happens when the organization has reached a relatively stable mode of recovery operation
- What happens to prepare for and then execute a successful return to the original site

The strategies should address residual risks identified in the BIA. You need to consider several alternatives when developing the strategy, including cost, allowable outage time, security, and integration with larger, organization-level disaster recovery plans.

We recommend that you design the return to the original site to use the same procedures, technologies, and methodologies that were used to failover to the recovery site. If this is not possible, then the testing of those different set of procedures would be necessary, introducing additional complexity and risk in the return process. The recovery strategy that you select should address the potential impacts that are identified in the BIA and should be integrated into the system architecture during the design and implementation phases of the system life cycle.

For this phase automation is a key feature. It is too demanding for personnel staff to be able to execute all the needed functions at the verge of a disaster. However, all the automation policies and people interaction must be rehearsed to exhaustion to achieve success.

Consider the following topics when developing the recovery plans:

- Systems recovery

  A systems recovery plan should address mission-critical application hosts, both centralized and distributed.

- Network recovery

  Network recovery plans are developed to provide for the restoration of:

  - Internal LAN and peripheral network support for mission-critical business processes
  - External WANs and telecommunications services
  - Communications between recovered systems and users

- Activation of Business Continuity disaster recovery teams

  The Business Continuity management team must coordinate all activities of all emergency recovery teams. Quick and valid decisions are key in an exceptional situation. The management team must have the authority for all decisions, including financial decisions. A notification plan must be available to activate all emergency teams. The experience of the 9/11 disaster demonstrated, that when a major disaster occurs, many people can be affected and not available in the emergency teams. A successful plan will take this into account to ensure that the business can be recovered and maintained even if key people are not available.

- End user recovery

  End user recovery plans, often ignored in traditional disaster recovery planning, consist of strategies to provide users with a mechanism for interacting with restored systems and networks to perform useful work.

Recovery procedures

Recovery operations begin after the Business Continuity plan has been activated, damage assessment has been completed (if possible), personnel have been notified, and appropriate teams have been mobilized. Recovery Phase activities focus on contingency measures to restore temporary IT processing capabilities, while activities executed during the Reconstitution Phase in “Restoration phase” on page 82 are directed to repair damage to the original system and restore operational capabilities at the original or new facility. At the completion of the recovery, the system will be operational and performing the functions designated in the plan. Depending on the recovery strategies defined in the plan, these functions could include temporary manual processing, recovery and operation on an alternate system, or relocation and recovery at an alternate site. Teams with recovery responsibilities should understand and be able to perform these recovery strategies well enough that if the paper plan is unavailable during an event, they can still perform the necessary activities.

Sequence of recovery activities

When recovering a complex system, such as a WAN involving multiple independent components, recovery procedures should reflect system priorities identified in the BIA. The sequence of activities should reflect the system's allowable outage time to avoid significant impacts to related systems and their application. Procedures should be written in a sequential format so that system components can be restored in a logical manner. For example, if a LAN is being recovered after a disruption, the most critical servers should be recovered before other, less critical devices, such as printers. Similarly, to recover an application server, procedures should first address operating system restoration and verification before the application and its data are recovered. The procedures should also include instructions to coordinate with other teams when certain situations occur, such as:

- An action is not completed within the expected time frame
- A key step has been completed
- Items must be procured
- Other system-specific concerns

If conditions require the system to be recovered at an alternate site, you will need to transfer or procure certain materials. These items can include shipment of data backup tapes from off-site storage, hardware, copies of the recovery plan, and software programs.

Procedures should designate the appropriate team or team members to coordinate shipment of equipment, data, and vital records. References to applicable appendixes, such as equipment lists or vendor contact information, should be made in the plan where necessary. Procedures should clearly describe requirements to package, transport, and purchase materials required to recover the system.

Recovery procedure highlights

To facilitate the recovery phase operations, the Business Continuity plan should provide detailed procedures to restore the system or system components. Given the extensive variety of system types, configurations, and applications, this planning guide does not provide specific recovery procedures.

You need to assign procedures to the appropriate recovery team and address the following actions:

- Notify internal and external business partners associated with the system
- Obtain necessary office supplies and work space
- Obtain and install necessary hardware components
- Obtain and load backup tapes or media
- Restore critical operating system and application software
- Restore system data

- Test system functionality including security controls
- Connect system to network or other external systems
- Operate alternate equipment successfully

Write recovery procedures in a straightforward, step-by-step style. To prevent difficulty or confusion in an emergency, do not assume or omit procedural steps. Use a checklist format for documenting the sequential recovery procedures. A checklist format is also useful for troubleshooting issues if the system cannot be recovered properly.

Restoration phase

In the restoration phase, recovery activities are terminated and normal operations are transferred back to the organization's original facility. During the recovery phase, as the contingency activities are performed, restoration of the original site should be under way. If the original facility is unrecoverable, the activities in this phase can also be applied to preparing a new facility to support system processing requirements.

When the original or new site is restored to the level that it can support the IT system and its normal processes, you can restore the system to the original or to the new site. Until you restore and test the primary system, you need to continue to operate the contingency system.

The restoration phase should specify teams responsible for restoring or replacing both the site and the IT system. The following major activities occur in this phase:

1. Ensuring adequate infrastructure support, such as electric power, water, telecommunications, security, environmental controls, office equipment, and supplies.
2. Installing system hardware, software, and firmware. This activity should include detailed Restoration procedures similar to those followed in the Recovery phase.
3. Establishing connectivity and interfaces with network components and external systems.
4. Testing system operations to ensure full functionality.
5. Backing up operational data on the contingency system and uploading to the restored system.
6. Shutting down the contingency system.
7. Terminating contingency operations.
8. Removing or relocating all sensitive materials at the contingency site.
9. Arranging for recovery personnel to return to the original facility.

Note: The teams within the crisis management teams need to understand and be able to perform their required functions without a paper plan in the event such documentation is unavailable.

High availability

High availability is defined as providing redundancy in an existing business process or infrastructure, such that loss or outage of a individual component is masked from causing an outage of that business process or infrastructure. We discuss this topic in “IT High Availability design” on page 92.

You can plan for and apply the concept of High Availability to all non-IT aspects of the Business Continuity plan as well. Examples include (but are not limited to):

- Physical facilities
- Telecom
- Accommodations for staff if relocation
- Voice communications

Disaster Recovery

A *disaster* is an unplanned outage that requires recovery at a different physical site on different physical hardware. You need to research and evaluate the most practical alternatives for processing in case of a disaster. It is important to consider all aspects of the organization, such as facilities, hardware, software, communications, data files, customer services, user operations, user systems, and other processing operations (as shown in Figure 3-27). The alternatives (dependent upon the evaluation of the computer function) can include hot, warm, or cold sites reciprocal agreements, two data centers, multiple computers, service centers, consortium arrangement, vendor supplied equipment, or combinations of all of these.

Although major disruptions with long-term effects might be rare, you need to account for these occurrences in the Business Continuity plan. Thus, the plan must include a strategy to recover and perform system operations at an alternate facility for an extended period. In general, three types of alternate sites are available:

- Dedicated site that is owned or operated by the company
- Reciprocal agreement with an internal or external entity
- Commercially leased facility

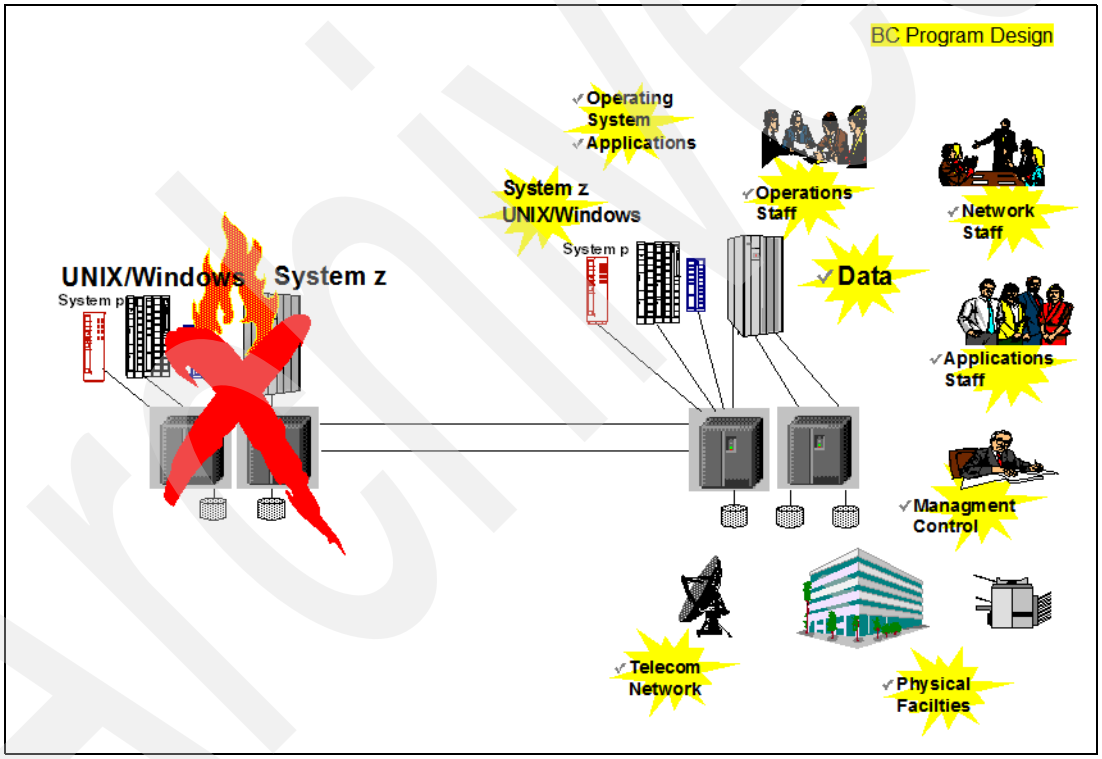

Figure 3-27 Alternate sites - require IT equipment, people, physical facilities, and telecom

Regardless of the type of alternate site that you chose, the facility must be able to support system operations as defined in the Business Continuity plan. The three alternate site types also might be categorized in terms of their operational readiness. Based on this factor, sites can be classified as cold sites, warm sites, hot sites, mobile sites, and mirrored sites, progressing from basic to advanced:

- **Cold sites**

  Typically consist of a facility with adequate space and infrastructure (electric power, telecommunications connections, and environmental controls) to support the IT system. The space can have raised floors and other attributes suited for IT operations. The site

does not include IT equipment and usually does not include office automation equipment, such as telephones, facsimile machines, or copiers. The organization using the cold site is responsible for providing and installing the necessary equipment and telecommunications capabilities.

- **Warm sites**

  Partially equipped office spaces that contain some or all of the system hardware, software, telecommunications, and power sources. The warm site is maintained in an operational status ready to receive the relocated system. The site might need to be prepared before receiving the system and recovery personnel. In many cases, a warm site can serve as a normal operational facility for another system or function, and in the event of a contingency plan activation, the normal activities are displaced temporarily to accommodate the disrupted system.

- **Hot sites**

  Office spaces appropriately sized to support system requirements and configured with the necessary system hardware, supporting infrastructure, and support personnel. Hot sites are typically staffed 24 hours a day, 7 days a week. Hot site personnel begin to prepare for the system arrival as soon as they are notified that a contingency plan has been activated.

  If the RTO is short, you can mirror the data to the site over fibre links either synchronously or asynchronously, depending on how far apart the two sites are. If the RTO is long, then you can store the data on backup tapes and install it prior to the site going live.

- **Mobile sites**

  Self-contained, transportable shells custom-fitted with specific telecommunications and IT equipment necessary to meet system requirements. These are available for lease through commercial vendors. The facility often is contained in a tractor-trailer and can be driven to and set up at the desired alternate location. In most cases, to be a viable recovery solution, mobile sites should be designed in advance with the vendor, and a service-level agreement (SLA) should be signed between the two parties. This is necessary because the time required to configure the mobile site can be extensive, and without prior coordination, the time to deliver the mobile site can exceed the system's allowable outage time.

- **Mirrored sites**

  Fully redundant facilities with full, real-time information mirroring. Mirrored sites are identical to the primary site in all technical respects. These sites provide the highest degree of availability because the data is processed and stored at the primary and alternate site simultaneously. These sites typically are designed, built, operated, and maintained by the organization.

  A requirement for these mirrored sites is the availability of synchronous remote copy among the I/O controllers. Usually companies that implement mirrored sites use both sites to produce work to decrease the expenses, not just a hot stand-by site. Depending on the distance, you can also install a Geographically Dispersed Parallel Sysplex™ (GDPS) between the two sites. In this case, you can have workload distribution and continuous availability on top of the disaster recovery solution.

There are obvious cost and setup-time differences among these different options. The mirrored site is the most expensive choice, but it ensures virtually 100% availability. Cold sites are the least expensive to maintain; however, they require substantial time to acquire and install the necessary equipment. Partially equipped sites, such as warm sites, fall in the middle of the spectrum. In many cases, mobile sites can be delivered to the desired location within 24 hours. However, installation time can increase this setup time. Table 3-1 summarizes the criteria that can be employed to determine which type of alternate site meets the organization's requirements. You need to analyze sites further based on the specific requirements that are defined in the BIA.

As sites are evaluated, the disaster recovery planning coordinator should ensure that the system's security, management, operational, and technical controls (such as firewalls and physical access controls) are compatible with the prospective site.

Table 3-1 Alternate site criteria selection

Site	Cost	HW equipment	Telecommunications	Setup time	Location
Cold site	Low	None	None	Long	Fixed
Warm site	Medium	Partial	Partial/full	Medium	Fixed
Hot site	Medium/high	Full	Full	Short	Fixed
Mobile site	High	Dependent	Dependent	Dependent	Not fixed
Mirrored site	High	Full	Full	None	Fixed

The organization can own or operate these alternate sites (internal recovery), or it can contract them commercially. If contracting for the site with a commercial vendor, adequate testing time, work space, security requirements, hardware requirements, telecommunications requirements, support services, and recovery days (how long the organization can occupy the space during the recovery period) must be negotiated and clearly stated in the contract.

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

Two or more organizations with similar or identical IT configurations and backup technologies might enter a formal agreement to serve as alternate sites for each other or to enter a joint contract for an alternate site. This type of site is set up through a reciprocal agreement or memorandum of understanding. A reciprocal agreement should be entered into carefully because each site must be able to support the other, in addition to its own workload, in the event of a disaster. This type of agreement requires the recovery sequence for the applications from both organizations to be prioritized with a joint perspective. Testing should be conducted at the partnering sites to evaluate the extra processing thresholds, compatible system and backup configurations, sufficient telecommunications connections, and compatible security measures, in addition to functionality of the recovery strategy.

You need written agreements for the specific recovery alternatives that you select, including the following special considerations:

- Contract duration
- Termination conditions
- Testing
- Costs
- Special security procedures

- Notification of systems changes
- Hours of operation
- Specific hardware and other equipment required for processing
- Personnel requirements
- Alternate personnel available if necessary to do the recovery at the recovery site
- Physical space at the recovery site
- Circumstances constituting an emergency
- Process to negotiate extension of service
- Guarantee of compatibility
- Availability
- Non-mainframe resource requirements
- Contracted network switchover to recovery site
- Priorities
- Other contractual issues

The strategy should include a combination of methods that complement one another to provide recovery capability of a full spectrum of identified risks. A wide variety of recovery approaches can be considered; the appropriate choice depends on the type of systems and their operational requirements.

Cost considerations for alternative sites

The planning coordinator should ensure that the strategy chosen can be implemented effectively with available personnel and financial resources. The cost of each type of alternate site, equipment replacement, and storage option under consideration should be weighed against budget limitations. The coordinator should determine known planning expenses, such as alternate site contract fees, and those that are less obvious, such as the cost of implementing a company-wide contingency awareness program and contractor support.

The budget must be sufficient to encompass software, hardware, travel and shipping, testing, plan training programs, awareness programs, labor hours, other contracted services, and any other applicable resources (such as, desks, telephones, fax machines, pens, and paper). The company should perform a cost-benefit analysis to identify the optimum recovery strategy.

Table 3-2 provides a template for evaluating cost considerations.

Table 3-2 Recovery strategy budget planning template

	Costs	Vendor	Hardware	Software	Travel/ Shipping	Labor/ Contractor	Testing	Supply
Alternate Site	*Cold Site*							
	Warm Site							
	Hot Site							
	Mobile Site							
	Mirrored Site							
Off-site Storage	*Commercial*							
	Internal							
Equipment Replace-ment	*SLAs*							
	Storage							
	Existing Use							

Distance to remote site

When configuring a remote recovery site, Figure 3-28 shows considerations for metropolitan distances versus out of region distances. (We discussed this topic in “Major factors to consider for out of region recovery” on page 24.)

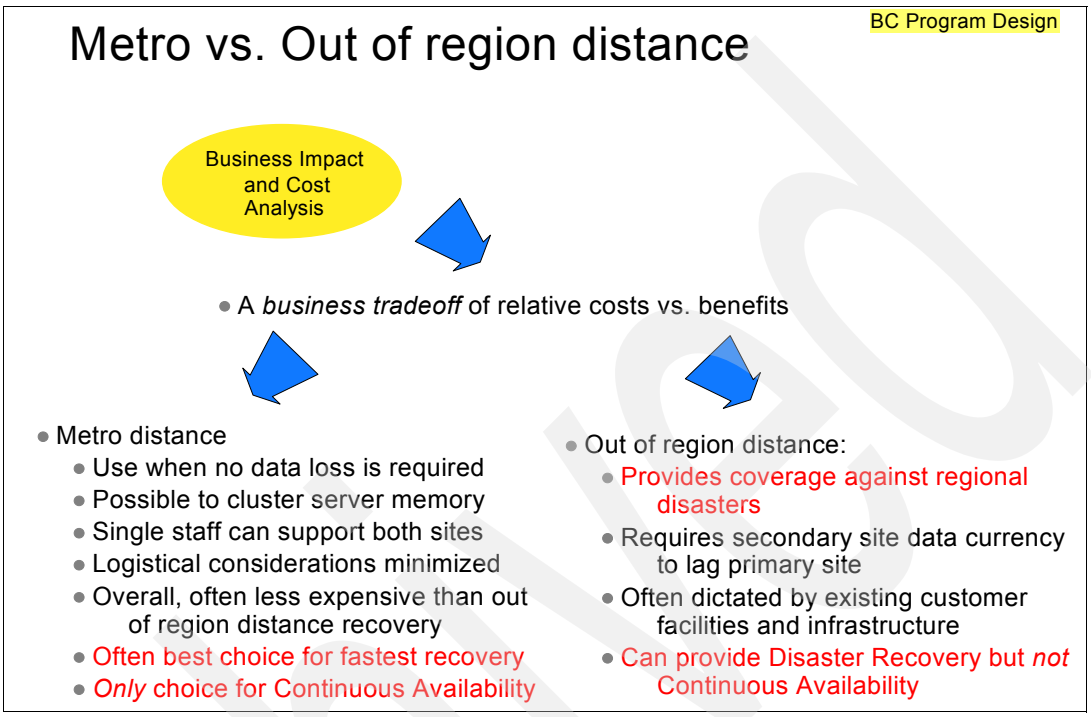

Figure 3-28 Distance considerations - metro distance versus out of region distance

Considerations for metropolitan distances include:

- Often chosen when the smallest possible data or transaction loss at the remote site is desired, because the remote data center is within synchronous data replication range.
- Metropolitan distances might be close enough to support clustering business processes, or IT infrastructure such as server and database memory, thus providing more robust possibilities for improving the business process or IT recovery time capability.
- A single staff might be able to support both metropolitan sites in more scenarios, thus potentially minimizing logistical considerations and reducing cost.
- Intersite communications and telecom line costs for equivalent bandwidth, can typically be less expensive at metropolitan distances compared to out of region distances.
- For true near-continuous availability at the IT data center level, metropolitan distances are typically required. The reason is: IT data center near-continuous availability failover will often require server clustering and database clustering, in addition to synchronous data replication. Server memory clustering and database memory clustering cannot be done at out of region distances.
- Recovery at metropolitan distances might be less expensive overall than comparable data center recovery at out of region distances. All other factors being equal, the recovery time at metropolitan distances is also typically shorter than a comparable out of region distance recovery.

Considerations for out of region distances include:

- Clearly, the industry trend is towards out of region recovery, in order to provide distance separation.
- In order to implement out of region recovery, the secondary site or sites data currency needs to lag the primary data center. This amount of data currency lag (which will be variable) will need to be managed, and data recreation implemented into the recovery design. Note that these are issues can be successfully handled with appropriate Business Continuity program design.
- Because it is quite expensive to build a new data center, an out of region location of the second data center might be indicated by the need to use existing data center infrastructure and locations.
- An out of region recovery at the data center level usually necessitates a shutdown and very fast restart. The distance separation means that server memory clustering or database memory clustering is not possible.

Note: If a combination of near-continuous availability *and* out of region recovery is required, this will need to be architected at the *software application* level, typically using duplexed standby databases and a message queuing middleware architecture. You can find more information about this kind of IT Business Continuity architecture in Chapter 11, “High Availability clusters and database applications” on page 305.

Neither metropolitan distance or out of region distance is universally better than the other. Base your final choice on a combination of business factors, Business Continuity program design, and IT strategy design.

Summary of Business Continuity program design

In this section, we described considerations for end-to-end Business Continuity program design planning, including (but not limited to):

- The importance of process and procedure: Being able to know precisely how your staff and your IT infrastructure will react to planned outages, as well as unplanned outages if they strike. Predictability of the *reaction to a disaster* is the goal,
- Different types of contingency plans in the BCP.
- Crisis management team, notification and activation procedures, call tree.
- High Availability and Disaster Recovery considerations, including alternate site selection considerations.

With a reasonable amount of end-to-end Business Continuity planning under way, you next turn to specific IT strategy and design considerations.

3.5.2 IT strategy design

You are now at the point of designing the specific IT strategies, architectures, and technologies that support the end-to-end Business Continuity program design that you created in the previous step (Figure 3-29).

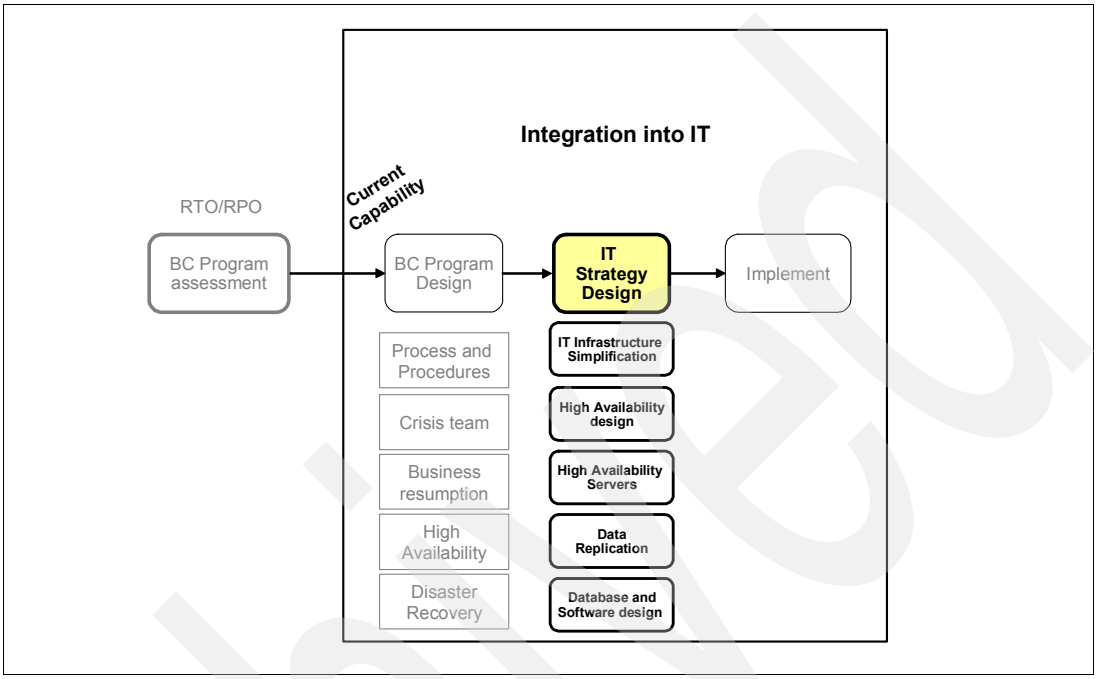

Figure 3-29 IT strategy design overview

In this section, we explore the best practices steps, which are composed of:

- IT infrastructure simplification as a prerequisite for IT Business Continuity
- High Availability design
- High Availability servers
- Data replication
- Database and software design

As a part of this design, we also explore major technology considerations for choosing from among the many candidate technologies.

Note: Because this book discusses IBM System Storage, we also provide key strategic System Storage Business Continuity planning considerations in “Data replication” on page 101.

IT infrastructure simplification as a prerequisite to Business Continuity

We recommend that your first step to IT strategy design should be to assess the existing IT infrastructure from a standpoint of consistency of operations, consistency of tools, and consistency and effectiveness of data management. A highly diverse environment, as shown in Figure 3-30, is not in and of itself bad, but clearly a large number of tools, policies, and skill sets makes recovery of such a diverse environment more difficult.

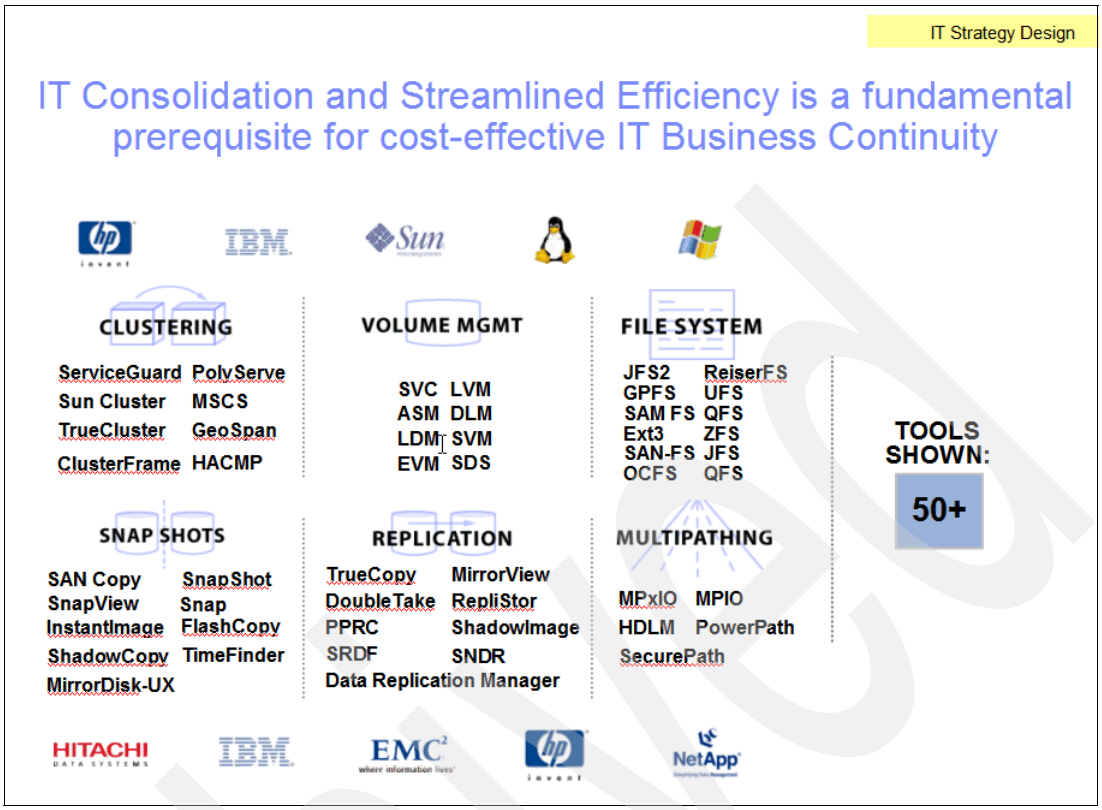

Figure 3-30 An IT infrastructure that might be difficult to do recovery - due to the number of tools

As we discussed in 2.2.4, “The struggle of IT diversity” on page 15, today’s IT Business Continuity requirements are best based upon a consolidated, standardized IT infrastructure. Such standardization and consolidation most effectively meets the needs of the business, not only from a Business Continuity standpoint, but also a TTM, TTC, application quality, flexibility, and scalability standpoint.

Therefore, we recommend that an important initial prerequisite for any Business Continuity solution is to base it upon, and blend it with, the existing IT infrastructure consolidation and standardization efforts already under way within your IT infrastructure, as illustrated in Figure 3-31.

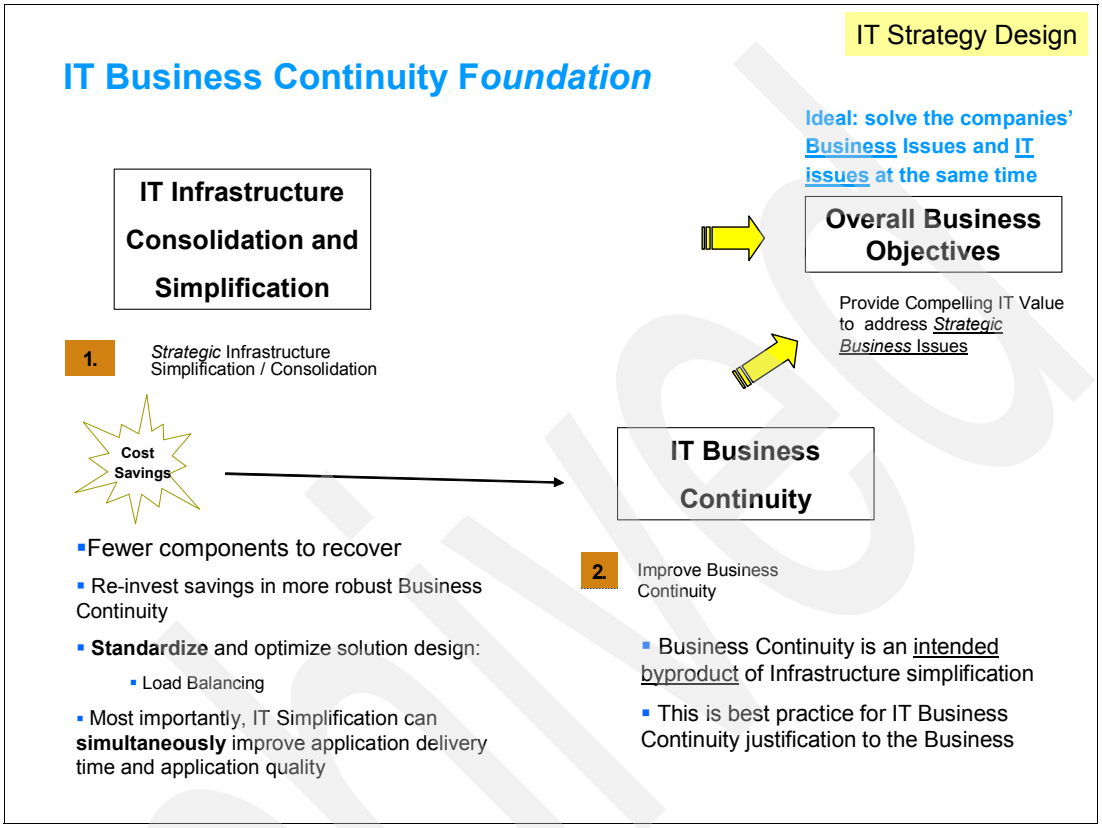

Figure 3-31 IT Simplification as the foundation for IT Business Continuity

The benefits of this approach are:

- It is easier to manage and recover fewer components than many.
- You can and should reinvest cost savings from IT consolidation into the IT Business Continuity project.
- You can sychronize necessary modifications and changes to the IT infrastructure to implement Business Continuity within other IT consolidation efforts. Necessary IT Business Continuity implementation tasks, such as improvements in workload management, data management, systems management, change control, and so forth, are very similar or ideally exactly the same type of efforts that need to be done for IT consolidation and simplification.

With an appropriate amount of IT Consolidation and Simplification under way upon which to base the Business Continuity effort, you next turn your attention to IT High Availability design.

IT High Availability design

IT High Availability design consists of three related, yet distinct aspects, as illustrated in Figure 3-32.

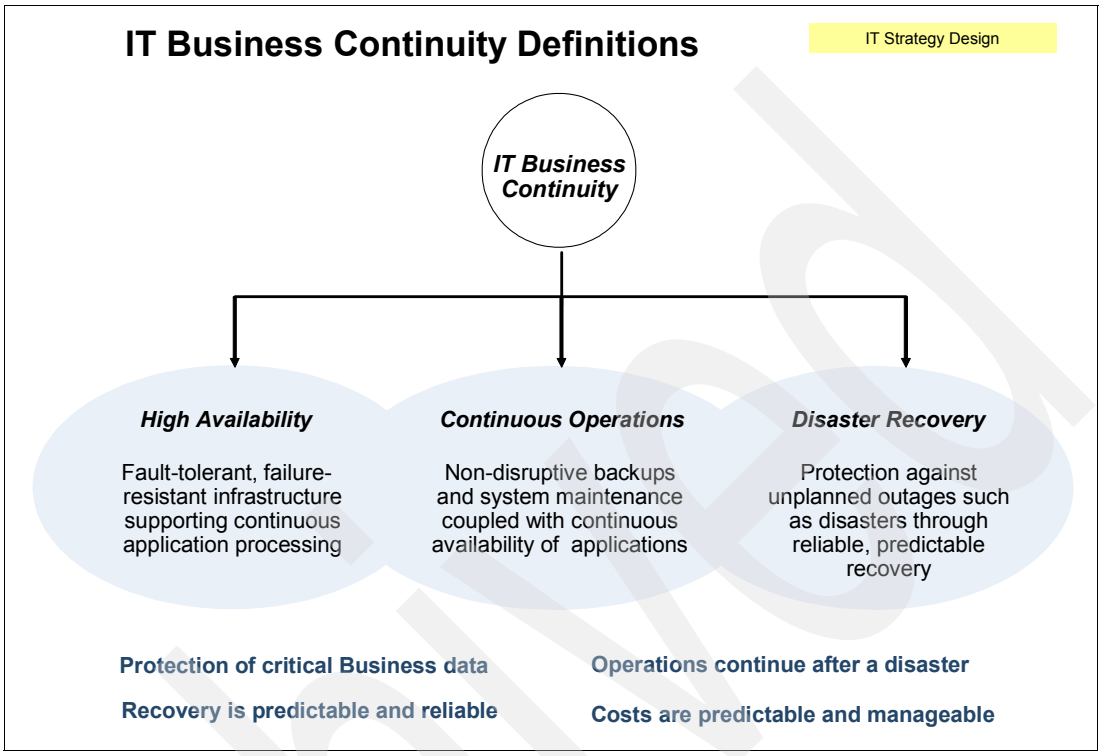

Figure 3-32 Inter-relationship of High Availability, Continuous Operations, and Disaster Recovery

These design concepts are defined as:

- **High Availability**

  High Availability builds redundancy into the local infrastructure with the intent of keeping individual local component failures from impacting the users (Figure 3-33).

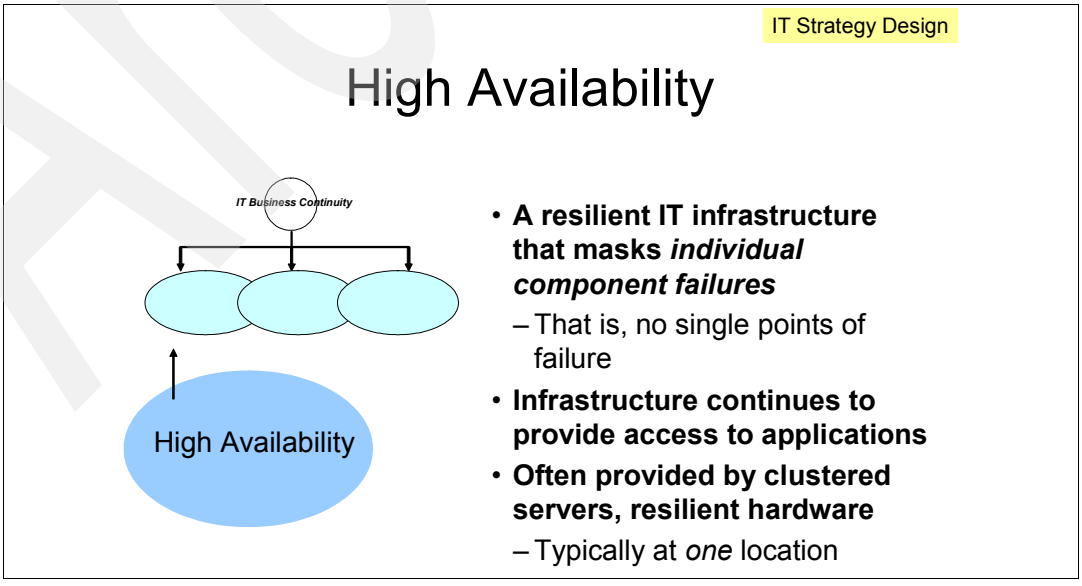

Figure 3-33 Definition of High Availability

IT High availability is often provided by server clustering solutions that work within operating systems, and coupled with hardware infrastructure, to provide an IT infrastructure that has no single points of failure. If a server that is running an application suffers a failure, the application is picked up by another server in the cluster, and users see minimal or no interruption.

Today's servers and storage systems are also built with fault-tolerant architectures to minimize application outages due to hardware failures. In addition, there are many aspects of security imbedded in the hardware from servers, to storage, to network components to help protect unauthorized access.

You can think of IT High Availability as resilient IT infrastructure that masks failures and thus continues to provide access to applications.

- **Continuous Operations**

Sometimes you must take important applications down to upgrade them or take backups. Fortunately, technology for online backups has improved greatly in recent years. However, even with these advances, you must at times take down applications as planned outages for maintenance or upgrading of servers or storage.

Continuous Operations provides the capability, when everything is working properly, to maintain access to the IT applications and infrastructure at all times (Figure 3-34). In other words, you do not have to take applications down merely to do scheduled backups or planned maintenance.

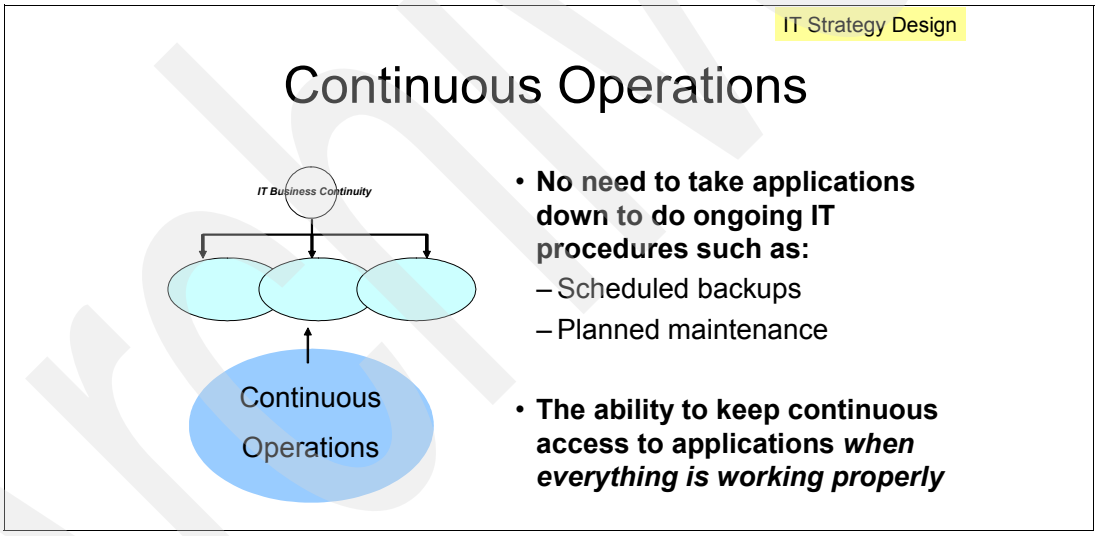

Figure 3-34 Definition of Continuous Operations

- **Disaster Recovery**

  Disaster Recovery is the ability to recover an IT data center at a different site if a disaster destroys the primary site or otherwise renders the site inoperable. The characteristics of a disaster recovery solution are that processing resumes at a different site and on different hardware (Figure 3-35).

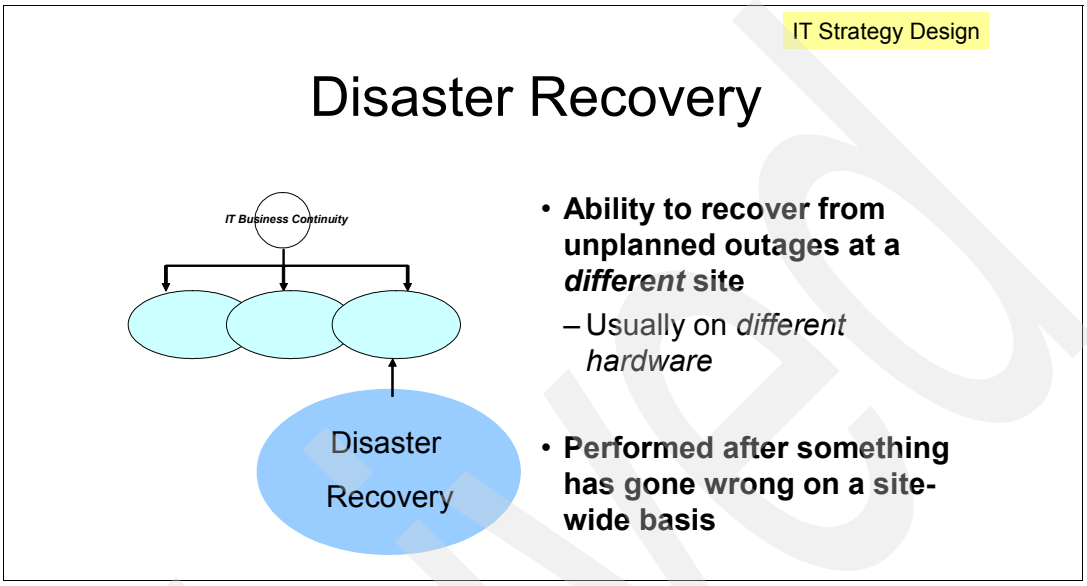

Figure 3-35 Definition of Disaster Recovery

Note: By this definition, a non-disaster issues, such as a corruption of a key customer database, might indeed be a catastrophe for a business, but it is not *by definition* a disaster unless processing must be resumed at a different location and on different hardware.

Disaster recovery is the ability to recover from unplanned outages at a different site—something that you do after something has gone wrong.

IT Business Continuity tiers of technology

There are many possible types of valid technologies that you can select as part of your IT strategy design. It can be a challenge to understand the best practices to make sense of them all. We use the concept of Business Continuity tiers as a primary tool to organize these many multiple valid products and technologies. We provide an overview of the Business Continuity tiers here and describe them in detail in Chapter 4, “Tier levels of Business Continuity solutions” on page 137.

The concept of *Business Continuity tiers* continues to be as valid today as when it was first created and described by the US SHARE User Group in 1988. While the technology within the tiers has obviously changed over the years, the conceptual value of organizing various recovery technologies according to their recovery speed has stood the test of time.

The Business Continuity tiers help you to organize various technologies into *useful subsets* that are much easier to evaluate and manage. The tiers concept recognizes that for a given customer RTO, all Business Continuity products and technologies can be sorted into a solution subset that addresses that particular RTO range.

Thus, by categorizing Business Continuity technology choices by RTO into the various Business Continuity tiers, you have the capability to more easily match your desired RTO time with the optimum set of technologies. Note that as the RTO time decreases, the optimum Business Continuity technologies for RTO must change. For any given RTO, there are always a particular set of optimum price or performance Business Continuity technologies.

Figure 3-36 illustrates the Business Continuity tiers chart.

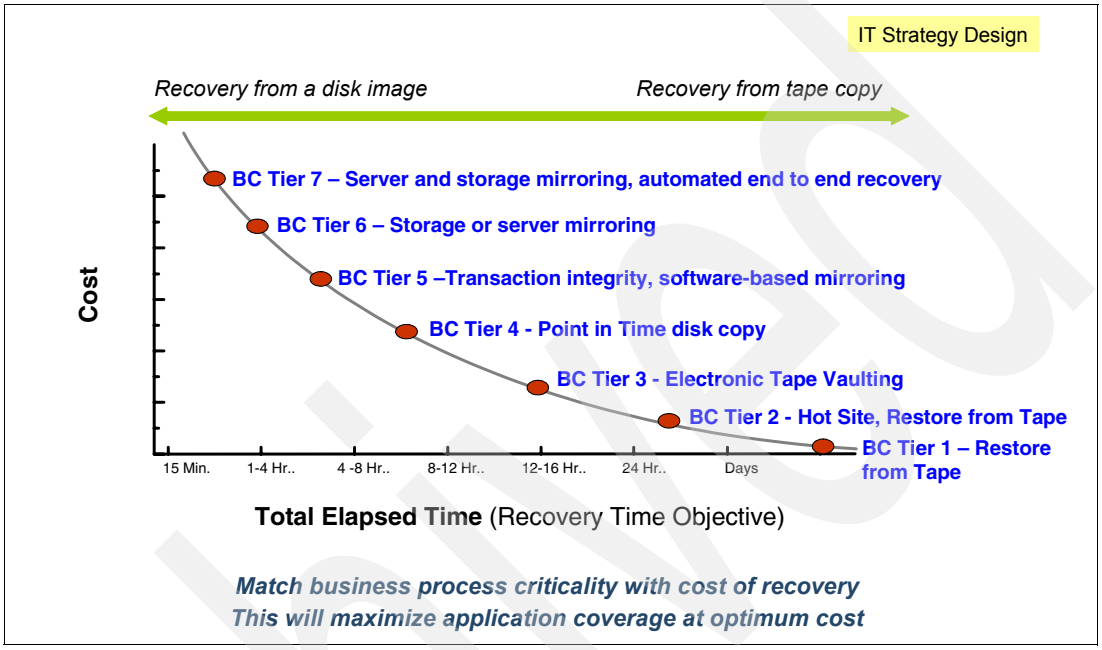

Figure 3-36 Introducing the Business Continuity tiers chart

As the recovery time decreases, more aggressive Business Continuity technologies must be applied to achieve that RTO (carrying with them their associated increase in value and capital cost).

Note: The tiers concept is flexible. The recovery time of a given Business Continuity tier is relatively fixed. As products and functions change and improve over time, you can update the Business Continuity tiers chart by the addition of that new technology into the appropriate tier and RTO.

The concept of the tiers chart continues to apply even as the scale of the application(s) changes. That is, the particular RTO values can 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.

We recommend that as part of the IT strategy design, your technical staff create your own internal version of the Business Continuity tiers chart that is specific to your particular environment. This chart can become a powerful and useful tool for internal Business Continuity planning, communication, and decision making. Creating and gaining agreement to such a tiers chart will tend to move into motion the requisite risk analysis, business impact analysis, and Business Continuity program design that is necessary to build this chart for your enterprise.

With such a chart in place, the business has a powerful tool to assign what tier or tiers and corresponding RTO an application requires and to readily determine what solutions might be viable.

Need for Business Process segmentation

Next, you need to do a segmentation of the organization's business processes and IT applications onto the Business Continuity tiers. There are many reasons to do this segmentation, including:

- No single set of IT Business Continuity technologies is right for every business process. We need to apply fast recovery technologies to some business processes (accepting the higher cost along with the valuable faster recovery), and we need to apply less expensive recovery technologies (which are lower in cost) to other business processes.
- Best practices therefore has evolved to segment the applications and business processes into (ideally) three segments, according to the speed of recovery that is required.
- You can then map these business process segments onto the Business Continuity tiers chart.
- This segmentation provides an excellent, in fact, essential foundation for implementing many of today's IT tiering concepts, such as tiered storage, tiered servers, and Information Life Cycle Management of data.
- Telecom bandwidth costs, as discussed in "Intersite networking and telecom components" on page 108, continue to limit the amount of data that can affordably be mirrored to a remote site. To facilitate identifying this data (and associated business processes and applications), business process segmentation provides an essential foundation.

We recommend that you identify and architect this business process segmentation as part of the IT strategy design. We also recommend that you consider (ideally) three business process or application segments as a guideline for an optimum number. For most IT shops, two tiers might be insufficiently optimized (in other words, overkill at some point and underkill at others) and four tiers are more complex but generally do not provide enough additional strategic benefit.

Applied to the Business Continuity tiers, this business process and application segmentation typically looks similar to Figure 3-37.

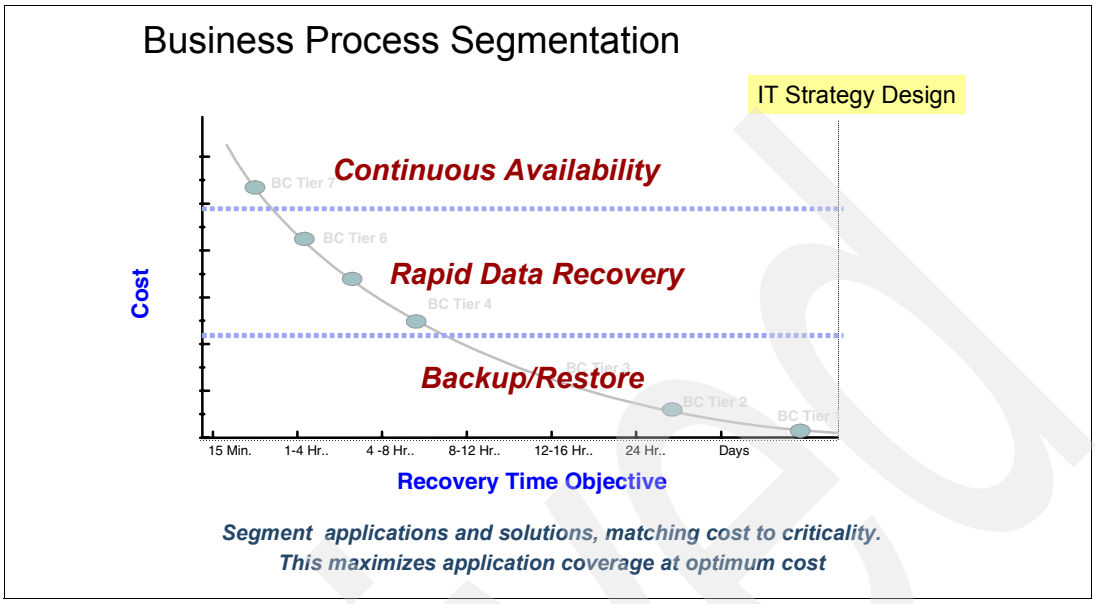

Figure 3-37 Business process segmentation - mapped to the Business Continuity tiers

Important: The reason for this step is to map the business process segments onto the Business Continuity technologies and, thus, to select the appropriate technology to provide the desired recovery level. In this way we link the *recovery of the business process* to the IT Business Continuity technologies.

Some general guidelines and definitions for the recovery capability of the recommended three business process segments include:[1]

- **Continuous Availability**
  - 24x7 application and data availability (server, storage, network availability)
  - Automated failover of total systems / site failover
  - Very fast and transparent recovery of servers, storage, network
  - Ultimate Disaster Recovery: Protection against site disasters, system failures
  - General RTO guideline: minutes to < 2 hours
- **Rapid Data Recovery**
  - High availability of data and storage systems (storage resiliency)
  - Automated or manual failover of storage systems
  - Fast recovery of data/storage from disasters or storage system failures
  - Disaster Recovery from replicated disk storage systems
  - General RTO guideline: 2 to 8 hours
- **Backup/Restore**
  - Backup and restore from tape or disk
  - Disaster Recovery from tape
  - RTO = 8 hours to days

[1] The stated RTOs in the chart's Y-axis and given in the definitions are *guidelines* for comparison only. RTO can and will vary depending on the size and scope of the solution.

With each of these business process segments defined, we can proceed to mapping the appropriate Business Continuity technology to the appropriate business process.

Tip: You can find more information about the types of technologies that are available to map to each of these business process segments in *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548.

Each segment builds upon the foundation of the preceding segment

The Business Continuity functionality of each business process or application segment is built upon the technology foundation of the segment that is below it. In other words, Backup/Restore technologies are the necessary foundations for more advanced technologies that deliver Rapid Data Recovery. And Rapid Data Recovery technologies are the necessary foundations for the more advanced technologies that deliver continuous availability. So, it is really a matter of building upwards upon the foundations of the technologies of the previous segment. Best practices for Business Continuity implementation roadmaps are to create a multiple phase project in which the overall Business Continuity solution is built step-by-step upon the foundation of the previous segment's technology layer.

To match this best practice of segmentation, IBM Business Continuity solutions are mapped into these three segments, as are the various IBM System Storage Business Continuity solutions that we describe in this IBM Redbook and in *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548.

Strategic value of business process segmentation

The strategic value of business process segmentation becomes clear as we look forward over time. Consider the diagram shown in Figure 3-38.

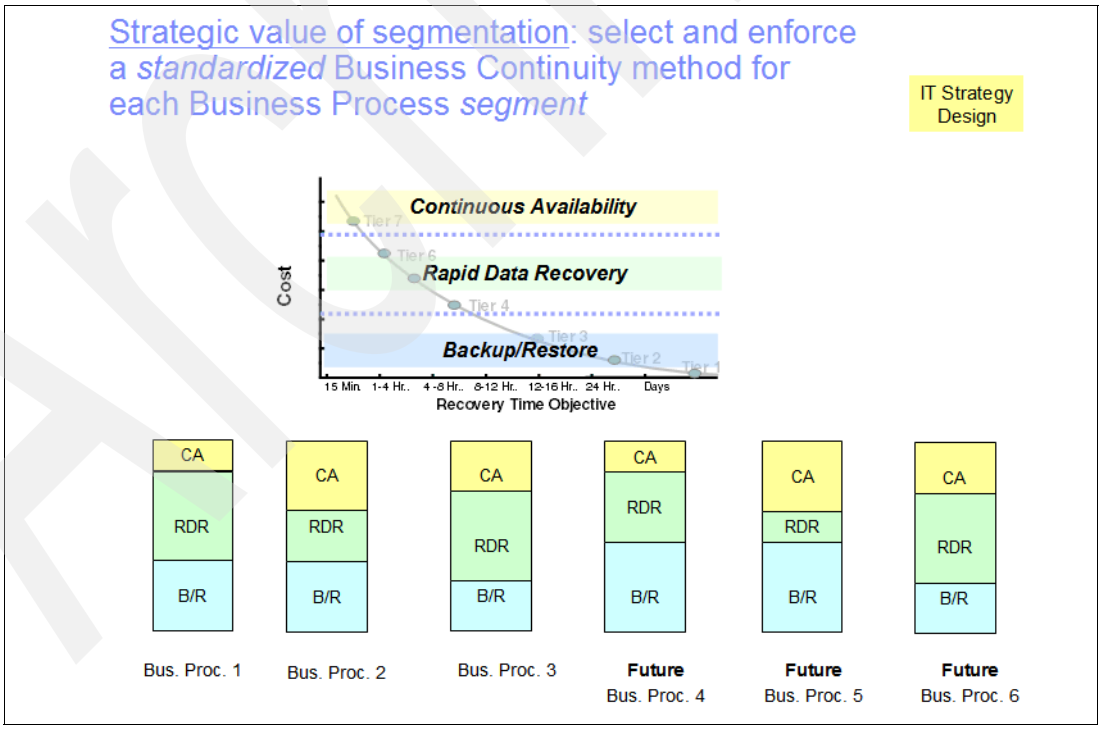

Figure 3-38 Business process segmentation - extend horizontally over time

These boxes represent business processes for which we are strategically designing segmentation. These boxes could map to lines of business or whatever makes sense for your organization.

- Architecting Continuous Availability

  We would start by identifying the continuous availability portion of your data and applications, which are defined to absolutely critical and in need of the continuous availability technology to do an immediate takeover in the event of a failure. Having defined these applications and business processes, we assign resources to them. These resources are designated to be the continuous availability pool of resources. In other words, we would define specific servers, specific storage pools, specific database and application software, as the strategic standards for this continuous availability segment.

- Architecting Rapid Data Recovery

  We would continue to architect the next level, the Rapid Data Restore pool. We would define specific servers, specific storage pools, specific database and application software, as the strategic standards and pool of resources for this Rapid Data Recovery segment. This set of resources might be a less costly level of storage, but satisfactory for the recovery and service level requirements of this segment.

- Architecting Backup/Restore

  Finally, we would continue to architect the final level, the Backup/Restore pool. We might define specific servers, specific storage pools, specific database and application software, as the strategic standards for this Rapid Data Recovery segment.

- Architecting a standardized Business Continuity environment

  What we are architecting by this business process segmentation, is the strategy that we use, in our IT Consolidation and Simplification, to standardize our IT infrastructure, horizontally across multiple business processes or even Lines of Business.

  If we are able to do this, we end up strategically architecting and standardizing, and creating a defined set of resource standards and pools for our enterprise, that are consistent, in a horizontal manner, across multiple business processes.

  We thus create a methodology that has strategic economies of scale. By enforcing this architecture over time with good policies and IT governance, we are able to create an expandable Business Continuity architecture for the future, based on our Business Process segmentation.

- Strategic value

  Over time, as new applications and business processes are defined, with such an architecture, the decision as to what type of Business Continuity recovery capability is to be implemented is now standardized.

The questions become:

1. What level of recovery capability does this application require?
2. What is the defined business process segmentation that fits that level of recoverability?
3. What are the shop standards and resource pools for server, storage, databases, that are assigned to that level of recoverability?

In this way, this kind of business process segmentation Business Continuity architecture specifies, defines, standardizes, and governs a *flexible, scalable*, and *consistent* Business Continuity solution set that will stay consistent over time.

As more and more new applications come online, they adhere to these standards., making the lines very crisp and clean. Thus, over time, as the applications and infrastructure evolves,

it will evolve into an ever more clean, crisp, consistently recoverable Business Continuity capability.

Implementation of the concept of Business Process Segmentation is an ongoing process, and is one key component in the overall IT Business Continuity strategy that we discuss in this chapter.

Summary of High Availability design

We discussed in this section the use of the Business Continuity tiers and business process segmentation concepts to derive a blended, optimized enterprise Business Continuity architecture.

We reviewed the following general steps:

1. Categorize the business' entire set of business processes into three segments: Low Tolerance to Outage, Somewhat Tolerant to Outage, and Very Tolerant to Outage. Of course, while some business processes that are not in and of themselves critical, they do feed the critical business processes. Therefore, those applications would need to be included in the higher tier.
2. Within each segment, there are multiple Business Continuity tiers. The individual tiers represent the major Business Continuity technology *choices* for that band. It is not necessary to use all the Business Continuity tiers, and of course, it is not necessary to use all the technologies.
3. After we have segmented the business processes and 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 for that application segment.

Designing this three-segment blended Business Continuity architecture has the effect of optimizing and mapping the varying application recovery time requirements with an appropriate technology, at an optimized cost. The net resulting blended Business Continuity architecture provides the best possible application coverage for the minimum cost.

High Availability servers

An extensive amount of information is available for configuring servers for High Availability. Because this is an IBM Redbook that focuses on System Storage, we give a brief overview of High Availability server concepts and then refer to the specific server reference material for the servers you use in your IT organization.

You should consult with your server specialists, and jointly plan with them to assure that a you have a good baseline level of consolidation and resiliency within your servers and operating systems. With an appropriate amount of fault tolerance, redundancy, resiliency, and clustering capability, your servers are then ready to participate in the larger IT Business Continuity design for an end-to-end recovery capability.

In general, High Availability servers are designed to mask outages through a degree of internal fault tolerance that is designed to avoid single points of failure. This design allows the servers to operate without interruption. Usually, this design includes components, CPUs, memory, and disks that have a resilient design.

Multiple servers are then set up for High Availability with a combination of hardware and software components configured to work together to ensure automated recovery in case of failure, with a minimal acceptable downtime. In such systems, the operating system detects problems in the environment, and manages application survivability by restarting it on the

same or on another available machine (taking over the identity of the original machine or node).

In this kind of High Availability server environment, it is very important to eliminate all single points of failure (SPOF) related to the external server connections. For example, if the server has only one network interface (connection), a second network interface should be provided in the same node to take over in case the primary interface providing the service fails.

A primary tool for server High Availability is clustered servers, which are loosely-coupled collections of independent servers (often called nodes) or Logical Partitions (LPARs) organized into a network for the purpose of sharing server resources and communicating with each other. Often, data and storage in a clustered environment is resident on shared disk arrays, accessible from any server in the cluster.

The same concepts previously discussed about the value and need for consolidation and standardization apply to servers. Building a cost-effective High Availability server environment is much easier when a requisite amount of server consolidation, server standardization, and server systems and change management, is in place.

Exploitation of server virtualization capability is a major trend, and should be exploited for servers to gain better control, more efficiency in management, much in the same way that storage virtualization provides similar benefits for storage.

In the next section, we focus in more detail on how to protect the actual data, particularly data that is resident on storage. We will consider data replication. Servers offer capabilities for data replication,

Data replication

In IT Business Continuity strategy design for multiple sites, the fundamental question from a remote recovery standpoint is: how best to replicate data to the desired location? In this section, we examine the IT strategy and design considerations for this question. We discuss data replication in general. Because this is a book about IBM System Storage, we also provide and emphasize key detailed information about the roles that IBM System Storage plays in providing tools, functions, and mechanisms to for data replication.

There are two basic types of data protection that we are concerned with when we speak of data replication:

- **Data corruption:** Most often occurs due to human or application errors. For this kind of data protection, *local point-in-time copies* are the typical recovery technology, whether those be disk copies or tape copies.
- **Data loss:** Most often occurs because of loss of access, physical failure, or physical destruction of the local storage or server devices. To provide these kinds of data protection, *remote data replication* to a remote site is the typical recovery technology.

Valid data replication can be done by the:

- Software application or database
- Server and operating system
- Storage device

Each of these data replication points has its advantages and considerations. Figure 3-39 shows these different data replication methods.

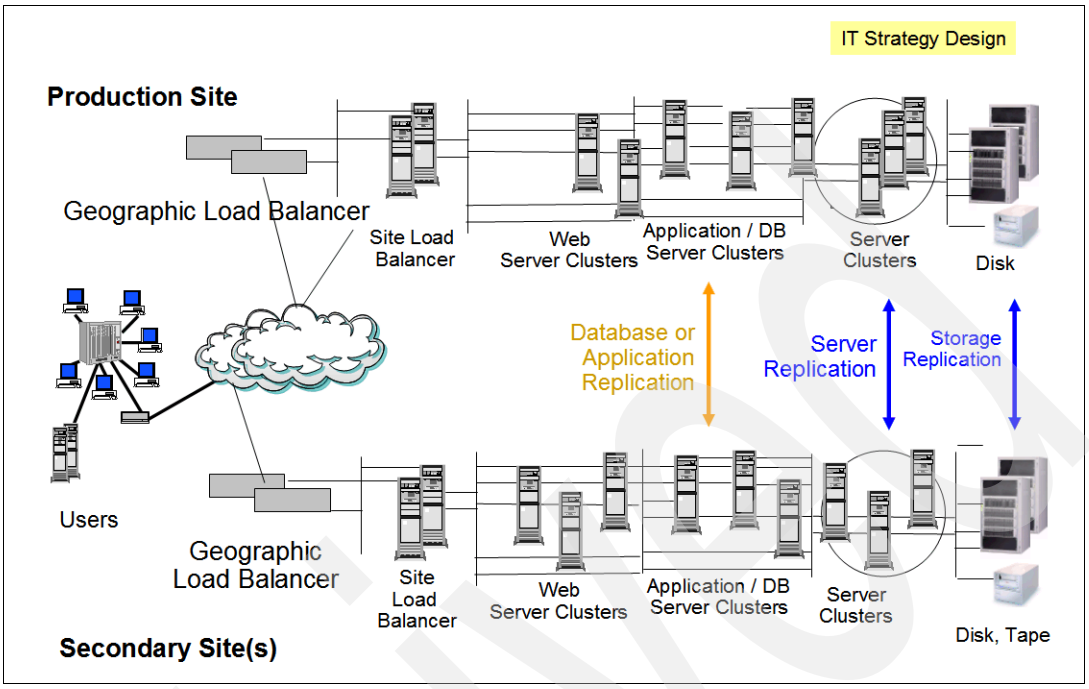

Figure 3-39 Three different, valid methods of data replication

Figure 3-40 summarizes the general considerations for selecting the appropriate data replication technology for your requirements.

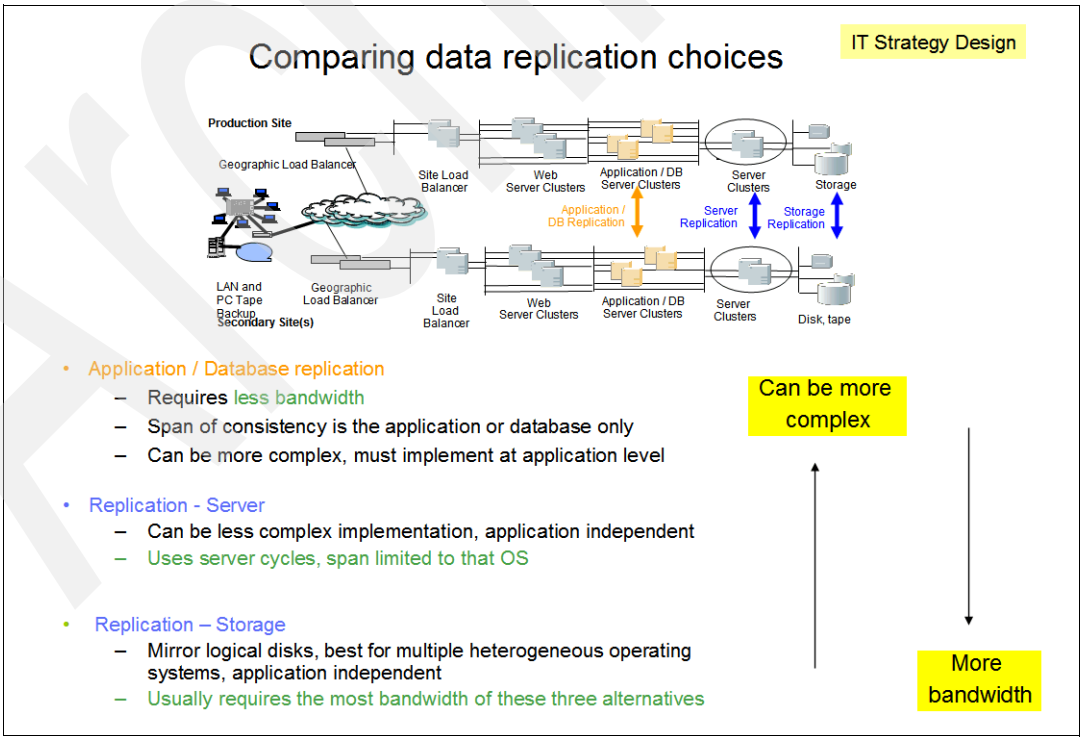

Figure 3-40 Comparing data replication choices

Let us see what the considerations are for each.

Replication: Application or database

Most major application and database software today provides facilities to replicate data to a remote site. As shown in Figure 3-41, this type of replication usually consists of the application or database software forwarding data (usually at the level of a logical unit of work or a transaction) to a companion instance of that software at a remote site. This often takes the form of log files. At the remote site, those log files can be stored and then applied to a shadow copy of the database.

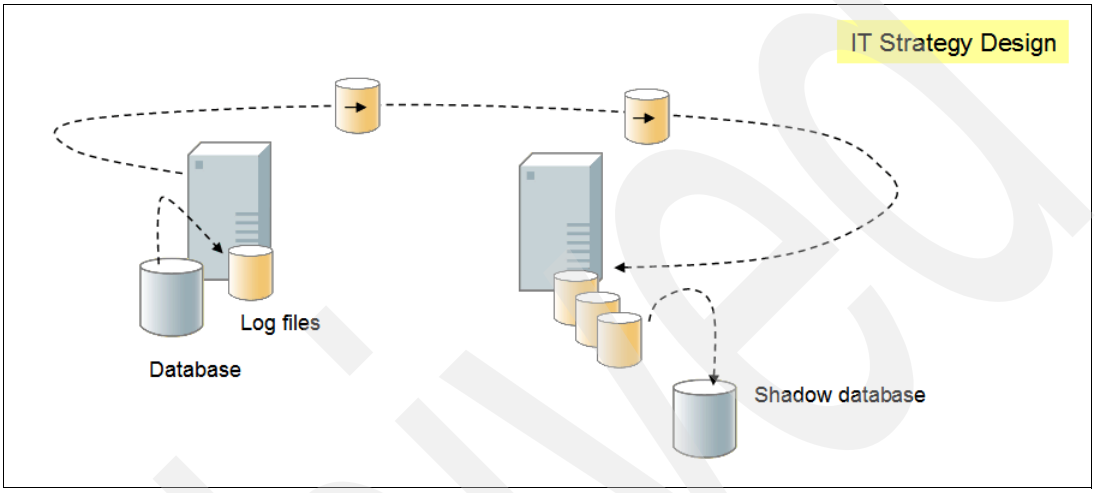

Figure 3-41 Application or database software data replication - transaction LUW level

Data replication using software-based replication is very useful especially where the requirement to minimize bandwidth is very important. Because the application or database has full knowledge of the data and the transactions, the software has the intelligence to strip out unnecessary information before forwarding the data to the remote database. As a result, application or database replication typically has the lowest bandwidth requirements for a given amount of transaction processing.

The span of consistency and recovery in a this environment must be within the application or database itself. If you were to strategically choose to use IBM DB2® remote replication, then all objects that you want to recovery would need to be under DB2 control. This issue might or might not be a problem. It just depends on your requirements.

For organizations with a highly homogenous IT environment from an application or database standpoint, and know that in their strategic future they can continue to standardize on this one or few major applications or databases, then choosing application or database replication makes a great deal of sense.

Tip: Examples of a highly homogenous IT environment from an application or database standpoint, might be: “The strategic scope of our workload is only SAP®” or “We are standardized on DB2 and do not anticipate strategic non-DB2 applications.”

The characteristics of this type of data replication are:

- Application or database software is in control of the replication, so the software reduces the amount of data transmitted to just the essentials that it needs. Often, this can be at the sub-record level.
- This kind of replication tends to use the lowest amount of bandwidth tends to transmit data in Logical Unit of Work.
- At the remote site, roll-forward and roll-back are performed for transaction integrity.
- Uses server cycles to move the data.

The maximum span of recovery is within this application software or database software. In other words, recovery typically will be able to span federated instances of this software application or databases, but you will not be able to synchronize this data replication with another different data replication mechanism that is outside the span of control of this software.

Examples of application or database replication include (but are not limited to):

- IBM DB2 replication
- Oracle® replication
- IBM WebSphere® MQSeries®
- And many more

There is no one right answer for all scenarios. Each is a value judgement based on your requirements and your strategic IT data center strategy.

Application or database replication is often chosen when lower bandwidth requirements are a must. It is also often chosen when the recovery span is deemed to always be within the boundaries of this application or database software.

Replication: Server

Server-level replication is defined as an operating system or a server software (such as a file system) replicates *block-level* write IOs to the remote site.

The difference between application or database software replication and server-level replication is that:

- Application or database software replicates data on a transaction or logical unit of work boundary
- Server replication software replicates data at the block level, that is, at the write I/O level, but it has no knowledge of the transaction or logical unit of work boundaries.

Figure 3-42 illustrates server replication.

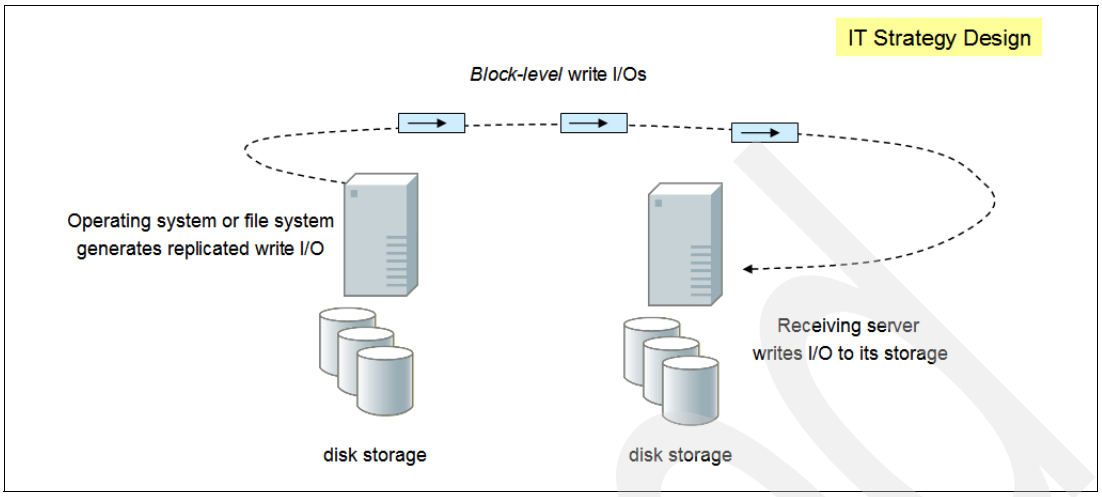

Figure 3-42 Server block-level replication

In this environment, because a common server platform has full knowledge of all applications and databases running under its control, the server-based replication (usually some form of logical volume mirroring or server clustering) can forward the changed data to the remote site. This choice will be independent of the applications and data themselves, and thus can be advantageous when there are too many different applications to manage individually on this server platform, or when there are too many different types and amounts of data to manage individually. Often, server-based is chosen when that is the skill base that is available within the existing staff.

The span of consistency and recovery in a this environment must be within the same server platform. For example, if you were to strategically choose to use IBM System p™ remote replication, then all objects that you want to recover, would need to be under System p control. This might or might not be a issue. It just depends on your requirements.

Server-based mirroring does not have as intimate knowledge of the logical meaning of the data being transferred, so compared to application or database replication, there is likely to be a higher bandwidth requirement.

Examples of server-based replication include (but are not limited to):

- IBM System p HACMP™
- IBM System p Logical Volume Manager and GLVM
- IBM System z™ System Data Mover and z/OS® Global Mirror
- And many more

The characteristics of server data replication are:

- Operating system or server software is in control of the replication, and replicates at the block write I/O level.
- This has the important effect of being agnostic of the application or database software. So there is an ability to have the span of replication cover more than one unrelated software or application set.
- However, the data at the remote site, if needed for a restart, will represent a *power-outage crashed image* of the data. Database or applications will need to do appropriate transaction backout on the data upon restart

- As the operating system or replication software is replicating block write IO's, it does not have intelligent knowledge of the logical meaning of the data. Therefore, the raw amount of data transmitted will be higher, and the amount of bandwidth required will be higher.
- Uses server cycles to move the data.

The span of replication can be any applications or databases running under the same server platform (examples: System p, System i, and Wintel) or under the same file system platform (example: VERITAS file system).

In summary, replication at the server level can be less complex to implement and is application independent. It uses server cycles to move the data and the span of recovery limited to that server platform.

Replication: Storage

Data replication using storage-based replication is very useful, especially when the server environment is highly *heterogeneous*. Storage replication by the storage system has gained significant popularity since its introduction in the mid-1990s on disk systems. Today, some tape systems (typically virtual tape) also do storage replication.

Figure 3-43 illustrates storage replication.

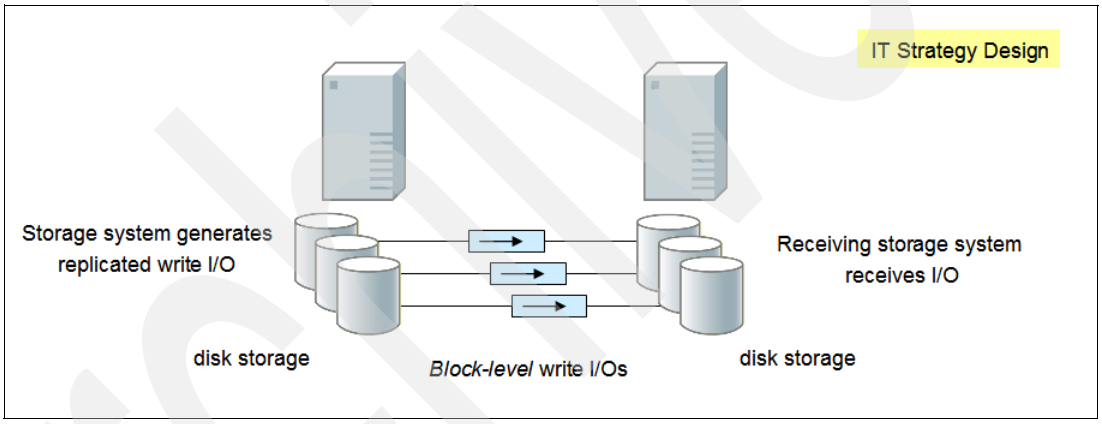

Figure 3-43 Storage system replication

Storage-based data replication is often chosen when the span of recovery must span multiple operating system platforms. It is also often chosen when the scale of the replication is large, because then the need to offload cycles from the servers is the greatest. Finally, storage replication is often chosen when fastest recovery is required, as the data is being replicated directly without intervention by the server.

In this environment, because a common set of storage platforms holds all the data for all of the servers, applications and databases, the storage-based replication (usually some form of disk or tape mirroring) can be used to provide a common mechanism to forward changed data to the remote disk or tape systems. This choice will be independent of the server platforms themselves, and thus can be advantageous when there are too many different servers, server platforms, and applications to manage individually.

The span of consistency and recovery in this environment must be within the same set of storage platforms, and typically, must be from the same vendor and within the same family of storage systems. This is because storage replication typically involves the transfer of a storage controller cache image from one controller to the next - therefore, the storage controllers must have the same internal cache structure (which normally occurs only within the same vendor and same family of storage controllers).

For example, if you were chose to use IBM SAN Volume Controller to do remote replication, then you are also setting the standard that all data that you want to replicate within the same recovery, should be under the control of the SAN Volume Controller. This might or might not be an issue. It just depends on your requirements.

The characteristics of storage data replication are:

- The storage system (disk or tape) is in control of the replication, and replicates at the block write I/O level.
- This has the important effect of being agnostic of the operating system platforms and application or database software. So there is an ability to have the span of replication now cover all server platforms attached to a storage platform
- As with server replication, the data at the remote site, if needed for a restart, will represent a *power-outage crashed image* of the data. Database or applications will need to do appropriate transaction backout on the data upon restart.
- As the storage system is replicating block write I/Os, it does not have intelligent knowledge of the logical meaning of the data. Therefore, the raw amount of data transmitted will be higher, and the amount of bandwidth required will be higher. In many cases, the bandwidth required for storage replication might be even higher than server replication, because the storage has even less knowledge about the logical meaning of the transmitted data than the servers.
- Uses storage controller cycles to move the data, thus offloading the server.

Examples of storage-based replication include (but are not limited to):

- IBM Metro Mirror
- IBM Global Mirror
- IBM N series SnapMirror®
- And many more

In summary, storage-based replication can provide the widest span of recovery, across multiple server and applications. It uses storage cycles to move the data, and the span of recovery is such that storage can replicate to other storage from that same vendor and vendor family. Consider the following key points:

- As shown in the Timeline of an IT Recovery in Figure 3-17 on page 62, data replication is just *one* component of the many activities that will be necessary in a successful Timeline of an IT Recovery.
- Data replication can be best thought of as the data transport portion of the overall solution. Other tools and procedures will need to be added to this data transport to assure that the applications and databases using the replicated data will find the data to be logically consistent. Therefore, one should not think of data replication as a solution in itself, it is just one component of the overall solution.

Thus, there is not one *best* way to replicate data. Rather, you need to examine the requirements, compare those requirements with the overall IT Business Continuity and Data Center strategy, and then select strategic data replication methods that synergize with those strategies.

Determining the best data replication for your organization

Determining the best type of data replication technology for your requirements depends on your IT recovery requirements and your strategic data center strategies. Figure 3-44 lists the essential IT requirements questions that can assist you in determining what is the optimum data replication technology for your particular needs. (We also discuss these questions in detail in Chapter 5, “Business Continuity Solution Selection Methodology” on page 151.)

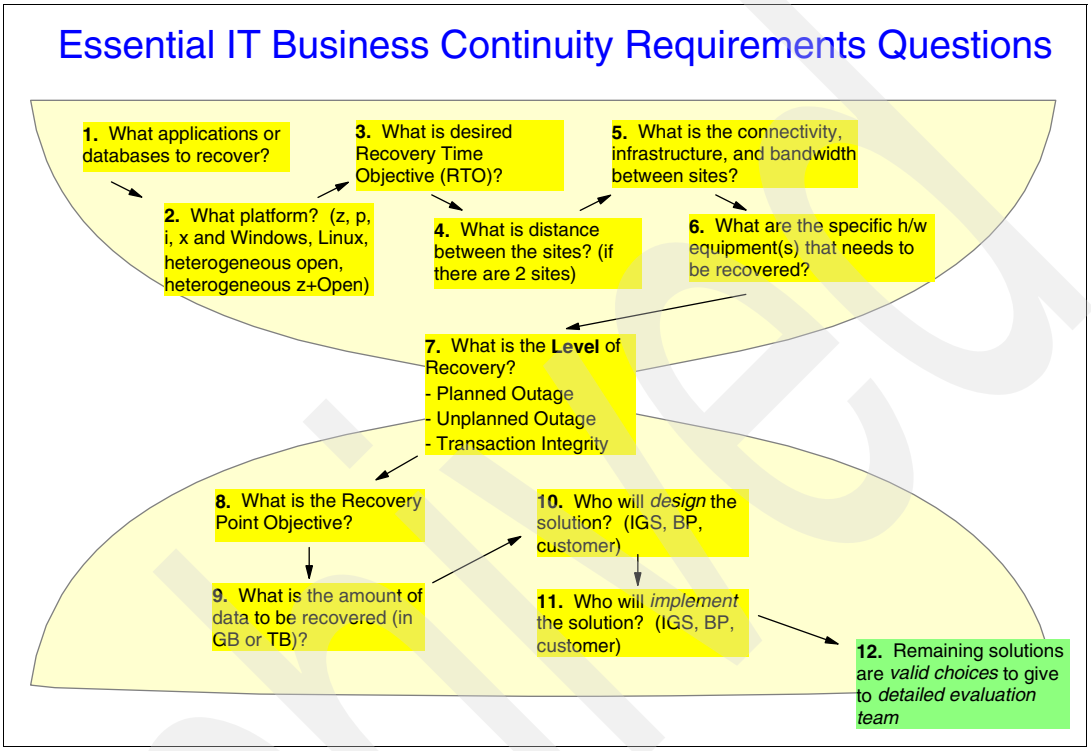

Figure 3-44 Essential IT Business Continuity requirements questions

Intersite networking and telecom components

In any remote data replication Business Continuity solution, you must have sufficient bandwidth and network infrastructure to carry the data to the remote site. The cost of bandwidth has historically been the major delimiter of how much data can be remotely replicated. We believe this will continue to be the case for the foreseeable future because although telecom costs are coming down, the annual percentage growth in data and storage continues to exceed the annual percentage decrease in telecom MBps of bandwidth.

Note: The fact that the major data replication limiting factor is telecom bandwidth is one of the key drivers of the need for Business Process segmentation, as described in “Need for Business Process segmentation” on page 96.

To properly size bandwidth, we diagram the inter-site data replication components as shown in Figure 3-45.

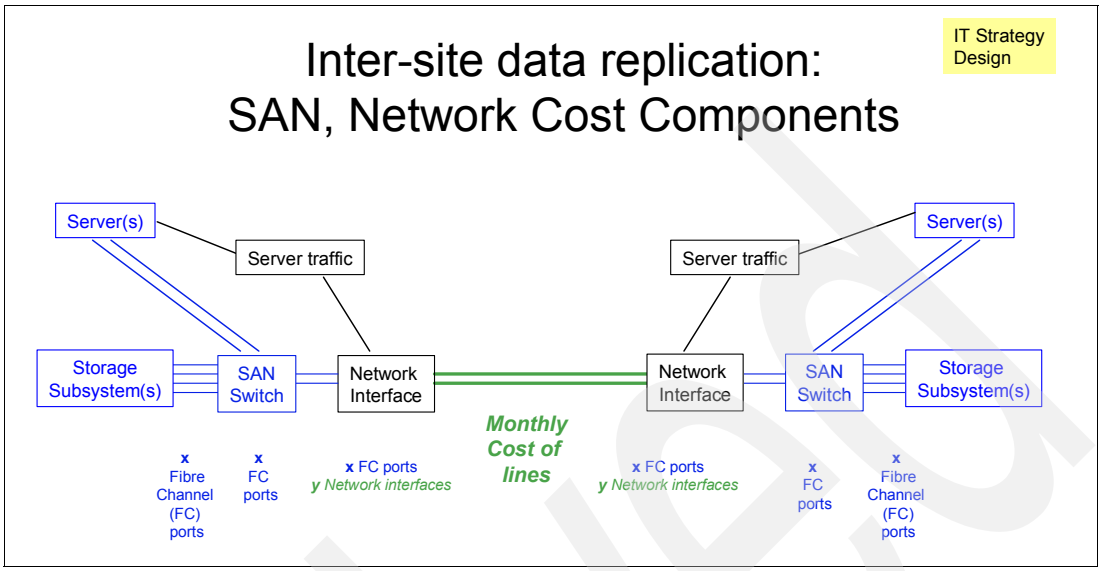

Figure 3-45 Inter-site networking components

Using this chart, you can diagram and track the major components of the end-to-end network configuration required to connect sites to each other. When planning the inter-site data replication network, you need to consider:

- SAN switches and network interface equipment will provide link aggregation

  For example, the SAN switches in Figure 3-45 take the incoming Fibre Channel workload on multiple SAN ports from both the servers and storage system and are able to multiplex that workload onto a fewer number of outbound SAN ports going into the network. This action assumes, of course, that enough link bandwidth is available on the outgoing SAN or network interface ports.

- As we discussed in “Replication: Storage” on page 106, disk mirroring has a fairly intensive telecom bandwidth requirement.

  In a continuously replicated disk mirroring environment, we recommend that you make a planning assumption that ongoing telecom costs will most likely be the largest TCO component of the entire data replication Business Continuity solution (often > 50% on a three year TCO basis)[2]

  Therefore, when considering any remote data mirroring, we recommend working with your network and telecom providers to have the costs of telecom available, so that you can plan an affordable balance between the amount of data to be mirrored, and the budget for bandwidth costs.

  It is common to segment business processes, data, and applications in such a way that disk replication is applied only to those business processes, data, and applications that require it. Furthermore, it is common to decide to reduce the amount of data to be replicated in order to fit into existing budget for bandwidth.

Note: An extensive discussion and definition of network transport options is discussed in Chapter 10, “Networking and inter-site connectivity options” on page 281. In this section, we will provide an overview and summary.

[2] Based on ongoing IBM IT industry studies of bandwidth costs for high speed data center telecom lines.

Telecom bandwidth capacity: Tips and a rule of thumb

In a data replication environment, it is necessary to perform a detailed workload analysis and bandwidth study to determine the amount of telecom bandwidth that will be required. We recommend that you use the smallest time granularity that can be afforded and measured in your workload analysis and bandwidth study.

In an ideal world, we recommend that you analyze your workload with a granularity as small as 30 seconds. Although this level of granularity might not often be feasible because of the amount of overhead, if it can be afforded and tools are available that can measure at this level, we recommend doing so.

The reason for this very small granularity, is that we are looking for *short duration* write bursts when performing data replication workload analysis. These peak write burst workloads cause the greatest stress on the remote data replication technology, and we will need to identify these peaks in order to properly size the telecom capacity.

The result of a workload and bandwidth analysis will be three bandwidth metrics, usually expressed in MBps:

- Steady state MBps requirement
- Peak MBps requirement
- The ratio of steady state to peak bandwidth requirement

Tip: For any given workload, the ratio of steady state to peak bandwidth requirement can vary depending on the data replication technique and specific implementation by vendor. As telecom is purchased in steady state MBps increments, it is very useful to consider the steady state to peak ratio of bandwidth requirement when doing evaluation of different data replication alternatives or vendors.

Remember that obtaining a workload analysis and bandwidth study is often a significant task and that you should plan your time and effort expectation accordingly.

Therefore, Figure 3-46 provides a general rule of thumb for continuous data replication bandwidth requirements. This rule of thumb provides a useful starting place for initial data and capacity sizing, and can used to move the IT Strategy Design process forward while the more detailed workload analysis and bandwidth study is completed in the background.

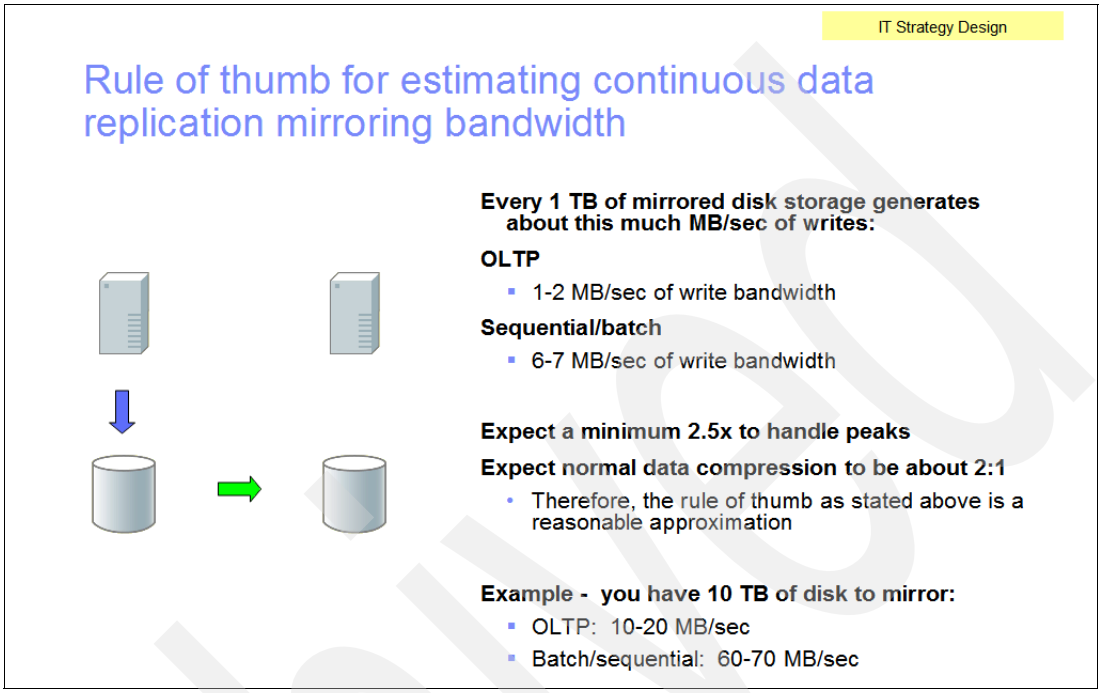

Figure 3-46 Rule of thumb for estimating continuous data replication bandwidth

Example 3-1 shows the derivation of this rule of thumb. You can examine this calculation and its assumptions and modify the rule of thumb using your own assumptions for block size, read/write ratio, and access density.

Example 3-1 Derivation of the rule of thumb for estimating continuous data replication bandwidth

```
The telecom ROT is based on worldwide averages for Access Density (I.e. number of
I/Os per second per GB)

The current reasonable worldwide average Access Density for production IT
workloads is: every 1 GB of disk data produces, on average, a little less than 1
I/O per second per GB (this is reasonable for both mainframe and Open)

We use Access Density to derive the Rule of Thumb (ROT)

1. Assume you have 1 TB of production data ( = 1000 GB).

2. Therefore, worldwide average Access Density says that you have, on average,
about 1000 Ops/sec to the disk storage.

3. Next, note that worldwide database Read/Write ratio is on average 3:1 (i.e.,
out of every 4 I/O's,  3 are reads, 1 is write). With this assumption, we then
derive that on average, 1/4 of all I/O's are write I/O's. Therefore:

4. If 1 TB generates 1000 ops/sec, then we have =  750 reads/sec, 250 writes/sec
```

5. We can then estimate the write bandwidth assuming certain write block sizes (OLTP below is **"O**n**L**ine **T**ransaction **P**rocessing**"**):

250 writes/sec * 4K block size OLTP = **1 MB/sec**

250 writes/sec * 27K block size BATCH = **6.75 MB/sec**

The above assumptions for block sizes are based on standard IBM performance measurement methods and techniques. You may adjust the block size, the Read/Write ratio, or the access density in the above calculations to better suit your own environment.

Examples for 10 TB: what is the estimated write MB/sec?

Example 1: daytime OLTP, assume average write is 4K (reasonable for both z/OS and Open)

Compute write MB/sec for OLTP:

10TB is 10000 GB, so that is 10000 ops/sec
With a 3:1 Read/Write ratio, that is 2500 writes/sec
2500 writes/sec * 4K = 10000K/sec = **10 MB/sec**

Example 2: night time batch, assume average write is 27K (reasonable for both z/OS and Open)

Compute write MB/sec for batch:

10 TB is 10000 GB, so that is 10000 ops/sec.

If we assume the Read/Write ratio is still 3:1, then

that would be 2500 writes/sec * 27K = **67.5 MB/sec**

Notes and assumptions:

- This rule of thumb is not intended to replace detailed capacity planning.
- The rule of thumb's calculation is linear. In other words, the calculation's amount of bandwidth scales directly with the amount of storage being mirrored.
- Scalability related considerations, queuing considerations, and so forth, are not considered in this rule of thumb.
- Batch and sequential read/write ratios can often be lower (that is, more write intensive) than online OLTP. If your batch R/W ratios are lower, this would increase the bandwidth required to mirror the batch, compared to the rule of thumb.
- Many techniques for optimizing the way that data replication is done in batch environments are available and should be considered. Some of those considerations are documented in the following sections in this chapter.

In summary, we recommend:

- Whenever doing data replication design, work together with your networking staff and networking vendors to have available an estimation of the monthly networking costs at differing levels of MBps.
- Expect to do a detailed workload analysis and bandwidth study, with as small a time interval granularity as feasible.
- You can use the rule of thumb shown in Figure 3-46 on page 111 to determine a reasonable initial assumption and rough estimate of bandwidth requirements, in order to keep moving forward with IT strategy design.

You can find more information about telecom and networking in:

- Chapter 10, “Networking and inter-site connectivity options” on page 281
- Chapter 6 “Storage networking for IT Business Continuity” in *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548
- *SAN Multiprotocol Routing: An Introduction and Implementation*, SG24-7321

Considerations for selecting synchronous versus asynchronous data replication

Selecting synchronous versus asynchronous data replication is related to, but not exactly the same thing as, selecting metropolitan distance versus out of region data center location.

Tip: We discussed the topic of selecting distance for remote IT data center locations in “Distance to remote site” on page 87.

Let us examine these selection considerations in more detail, as shown in Figure 3-47. The final selection is always a trade-off of relative costs versus benefits.

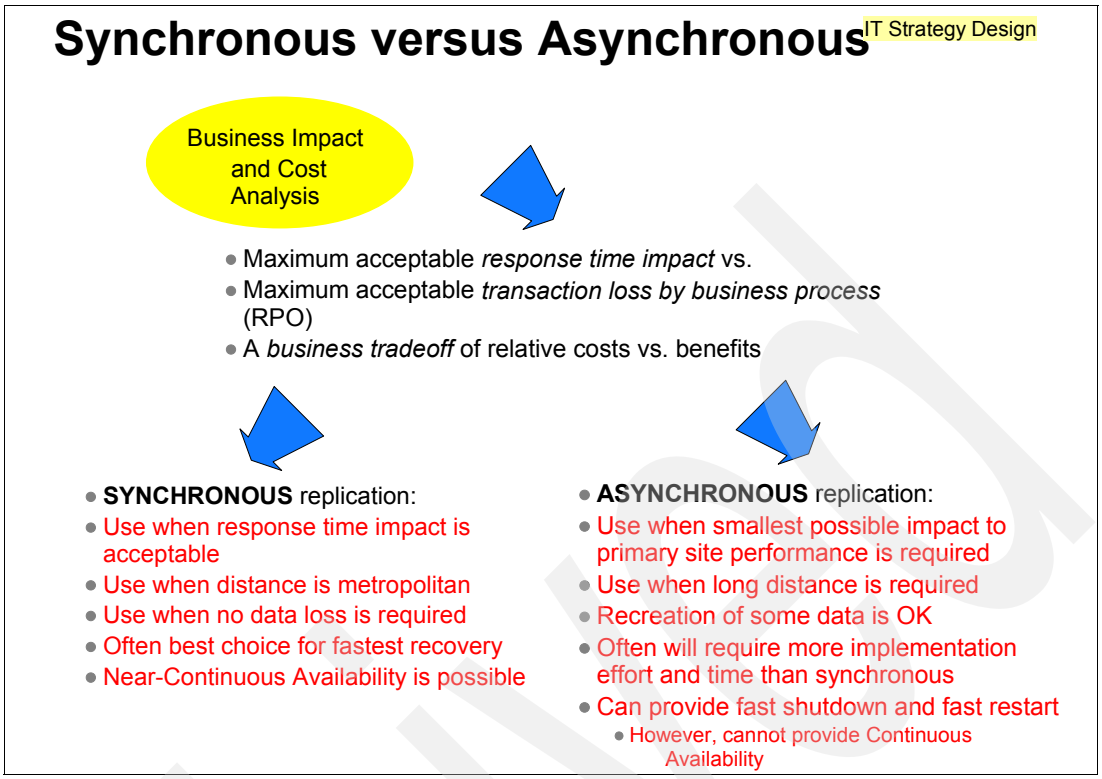

Figure 3-47 Considerations: Synchronous versus asynchronous data replication

Considerations for synchronous data replication include:

- Can use when the response time impact is acceptable.
- Can use when the distance is metropolitan.
- Use when no data loss, at the block storage level, is required.
- Synchronous replication can provide the foundation for near-continuous availability, because the remote site is in a synchronized state that can support a nearly transparent swap to the remote site.
- Therefore, synchronous is often the best choice for the fastest recovery.

Considerations for asynchronous data replication include:

- Often chosen when the smallest possible impact to the production applications is desired, regardless of distance.
- Chosen when out of region distance is required.
- When recreation of data at remote site (that is, a lag in data currency) is OK and is recoverable.
- Asynchronous replication can provide a fast shutdown and a fast restart.

Note: A near-continuous availability failover is *not* possible when using asynchronous data replication because the remote site is lagging behind the primary site in data currency. Thus, the remote asynchronous site is not in a synchronized state (that is, both data and system information (catalogs, VTOCs, indexes, and so forth) are lagging in currency and, thus, a transparent swap is not possible.

Neither method is inherently better than the other. It is a matter of choosing the tool that best fits the Business Continuity program design.

Consideration for implementation effort, synchronous versus asynchronous

Note that asynchronous data replication (from any vendor) usually requires more implementation effort and elapsed project time than synchronous data replication for the following reasons:

- Unlike synchronous, the data currency can *lag* at the remote site

  A certain amount of additional data recreation, over and above that present in fully synchronized data replication, is necessary when recovering at the remote site. Depending on the data lag, this lag can also introduce additional recovery ripple effects in other related business processes, over and above what happens in a fully synchronized environment

- Unlike synchronous, the data currency can *vary* at the remote site

  This variation means that the recovery processes and procedures need to account for a more variable interaction between the different business processes that are recovered. Capacity planning and testing need to account for and quantify the variability and frequency of the data lag. Testing and planning time need to resolve differing test results that can occur at differing levels of data lag.

  Techniques for handling these asynchronous considerations are well proven. However, we recommend that you add an appropriate amount of additional project planning and IT integration time when selecting asynchronous replication.

Disk replication, data consistency, and consistency group

When using disk replication technology in your IT strategy design, remember the following key disk mirroring technology concepts.

- Volume/LUN level versus File Level

  Disk systems typically mirror at the LUN or volume level. Examples of these kinds of disk mirroring systems are the IBM DS8000™, DS6000™, ESS, SAN Volume Controller, and DS4000™.

  File level mirroring can also be done by certain disk systems, such as the IBM N series network-attached storage. The N series can replicate both at the volume level as well as the file level.

- Power-outage, crash-consistent, point-in-time image of the data

  When disk system mirrored volumes are recovered at a remote site, they present to the application and database team, a *point-in-time* set of data-consistent, power-outage consistent, *crash-consistent* data LUNs / volumes.

Figure 3-48 illustrates this concept of the Timeline of an IT recovery chart.

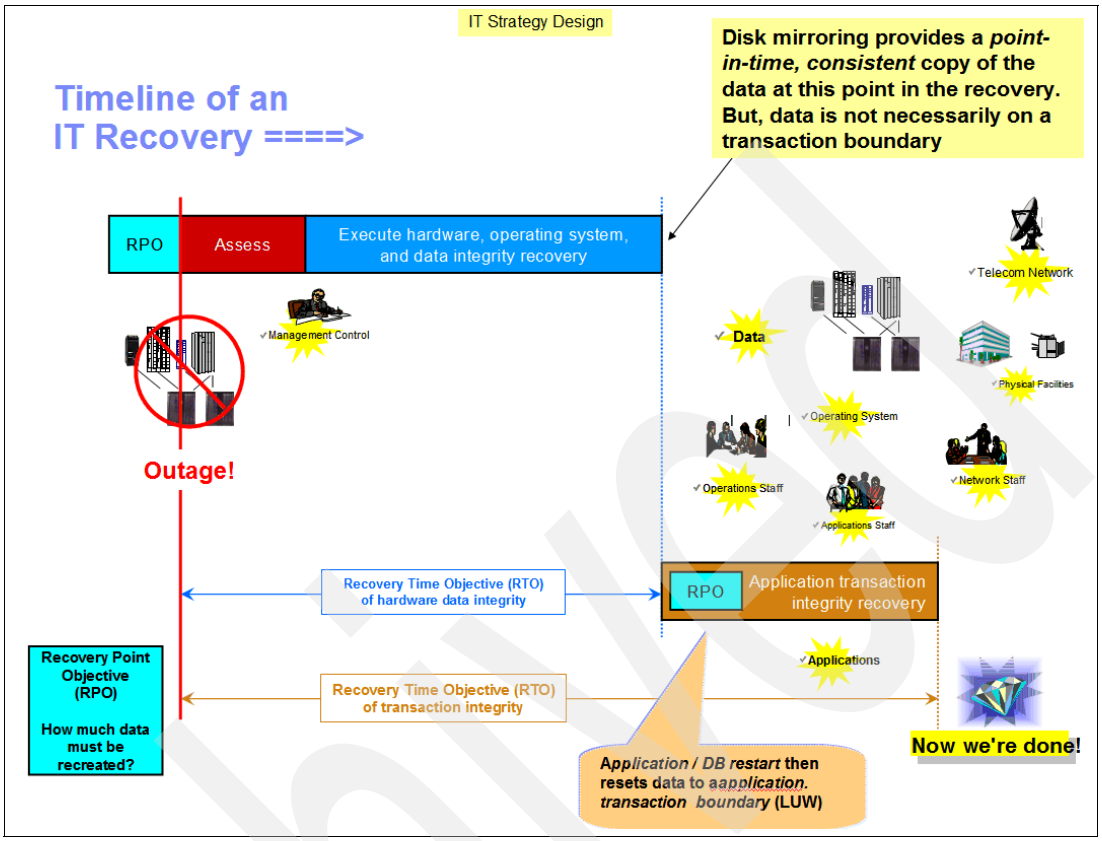

Figure 3-48 Timeline of an IT recovery

In this diagram, disk mirroring provides (only) a point-in-time set of data-consistent, power-outage consistent, and crash-consistent data LUNs and volumes to the applications. The application or database software then takes this set of point-in-time LUNs or volumes, and during the emergency restart, *resets* the data to the most recently available transaction boundary.

- Consistency Group assures data consistency across multiple mirrored LUNs/volumes

  This concept of a power-outage, crash-consistent point-in-time image of the data, spread across multiple LUNs or volumes, is essential to insuring a repeatable, reliable recovery of of IT applications and databases. The disk microcode function that keeps these multiple LUNs or volumes in a consistent point in time image with each other, is called a Consistency Group.

Figure 3-49 shows the concept of a Consistency Group.

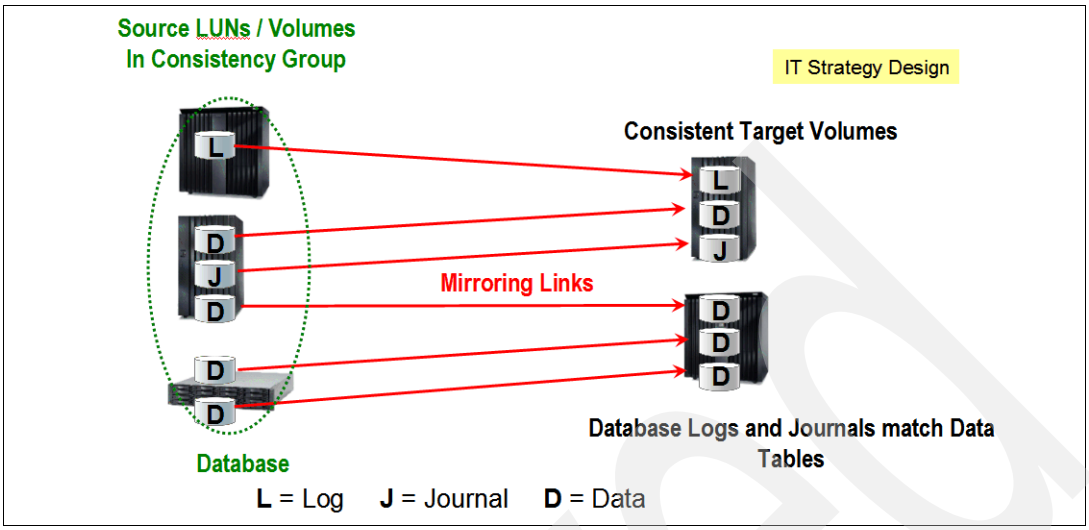

Figure 3-49 Consistency groups assure data consistency across multiple mirrored LUNs/volumes

- Maximum Consistency Group span determines scalability of the disk mirroring

  The maximum number of LUNs/volumes that can be defined in a Consistency Group, therefore determines the maximum size of the application set or databases that can be at one consistent point in time with each other in a disk mirrored environment. Different disk systems have differing levels of scalability. We recommend when selecting disk mirroring, to always insure that your disk mirroring vendor of choice provides Consistency Groups of adequate size for your applications.

  Figure 3-50 lists planning considerations for the maximum size of a Consistency Group for current IBM disk systems.

IBM Disk Mirroring
Consistency Group Scalability

IT Strategy Design

	IBM Metro Mirror	IBM Global Mirror	IBM z/OS Global Mirror
IBM DS8000, DS6000, ESS	•Unlimited •Multiple subsystems or LSSs require automation software	•Standard: up to 8 disk subsystems •RPQ: up to 17 primary side subsystems	•Unlimited •Requires System z server System Data Mover cycles for Consistency Group
IBM SAN Volume Controller	Up to 1,024 LUNs	Up to 1,024 LUNs	n/a
IBM DS4000	•Does not support Metro Mirror Consistency Group •Use DS4000 Global Mirror if Consistency Group is needed	Up to 64 LUNs	n/a

Figure 3-50 IBM disk mirroring consistency group scalability

The Consistency Group size is appropriate for the typical data replication technique that is used in different product ranges:

- In the network attached storage (NAS) or midrange disk systems, it is more common for data replication to be done by the application software or the servers. Therefore, the NAS or midrange storage systems typically are not required to mirror as much data as the larger disk subsystems.
- The SAN Volume Controller provides up to 1,024 LUNs in a Consistency Group. This is appropriate given SVC's role as a storage virtualization and consolidation controller for a heterogeneous environment.
- Finally, the enterprise class DS8000, DS6000, and ESS offer unlimited Consistency Group size capability.

Tip: You can find a more detailed discussion for the need for disk mirroring Consistency Groups and their implementation in , "Rolling disasters" on page 259.

Recovery of batch processing in a disk mirrored environment

When considering recovery in a disk replication environment, there are two major workload types to consider:

- Online Transaction Processing (OLTP)
- Batch/sequential

In this section, we focus on considerations for recovering batch processing. Batch has a different workload and recovery personality than OLTP, and we need to consider this when creating the IT strategy design for the overnight batch processing.

Background

Typically, OLTP software has integrated index files, log files, and roll-forward or roll-back procedures. The software has integrated facilities to recover from a power-outage, data-consistent crashed image of their data. In effect, the OLTP software has its own internal checkpoint/restart functionalities, and the software automatically uses those facilities when doing an emergency restart.

If you are disk mirroring an OLTP transaction system, the storage mirroring will provide a point-in-time, data consistent image of the data. Thus, upon restart in the remote site, the OLTP system will be able to do an emergency restart, and use its internal facilities to use the point-in-time, power-outage, crash-consistent copy of the data to restart at the point of failure.

Tip: See Figure 3-48 on page 116 to review of the concept of a power-outage, crash-consistent copy of the data.at the remote site.

However, in most cases, batch processes and batch jobs typically implement a lower (in some cases, a *much* lower) amount of checkpoint or restart capabilities. Checkpoint or restart in batch processing is not an automatic batch function, it needs to be designed into the batch process by the application developers.

The issue for batch recovery in a disk mirrored environment

Let us see what the planning issue is for batch recovery in a disk mirrored environment. Consider a typical batch update job in progress. Suppose that this batch job does not have checkpoint or restart implemented. Figure 3-51 illustrates what happens if an outage occurs part of the way through the batch run.

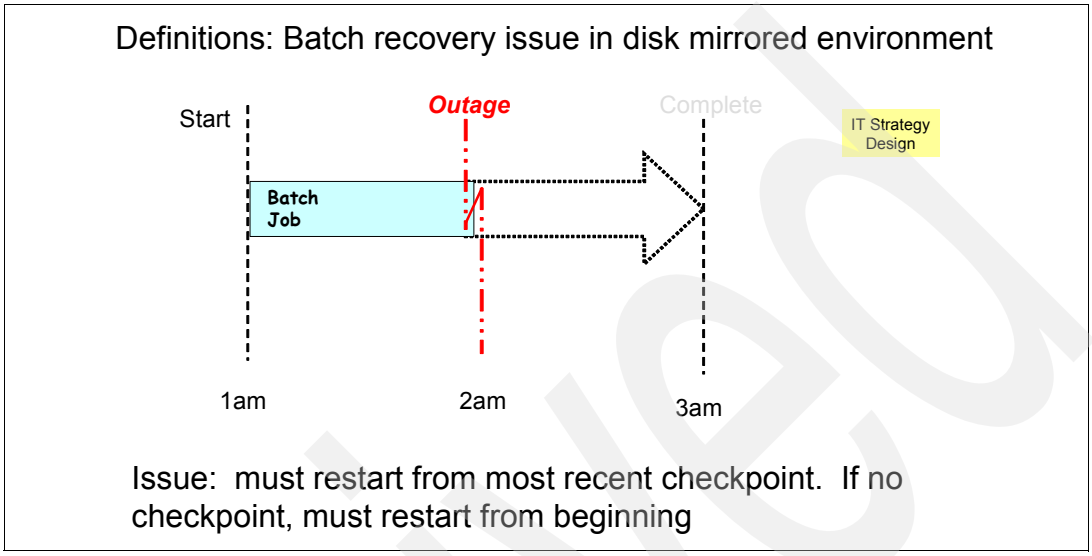

Figure 3-51 Batch job in progress - recovery issue

When the system is restarted at the remote site, the batch sequential files are a point-in-time, power-outage consistent, crash consistent copy. However, because the outage was an unplanned outage, the sequential files that were being updated were never closed. The end-of-file marker is pointing still to the data as it was in the sequential file as of the beginning of the job, (that is, 1:00 a.m.).Therefore, even though the batch job had written one hour's worth of data onto the batch job's LUNs/volumes, you need to restart this batch job from that beginning point (that is, at from 1:00 a.m.).

To avoid this issue, checkpoint or restart capabilities for batch and sequential workloads are available at either the operating system level or the application program level (usually through either system facilities or ISV software). If we implement checkpoint or restart functions into the batch processing as shown in Figure 3-52. Then, in the event of an outage and recovery at the remote site on the mirrored disk, you can restart this batch job at 1:30 a.m.

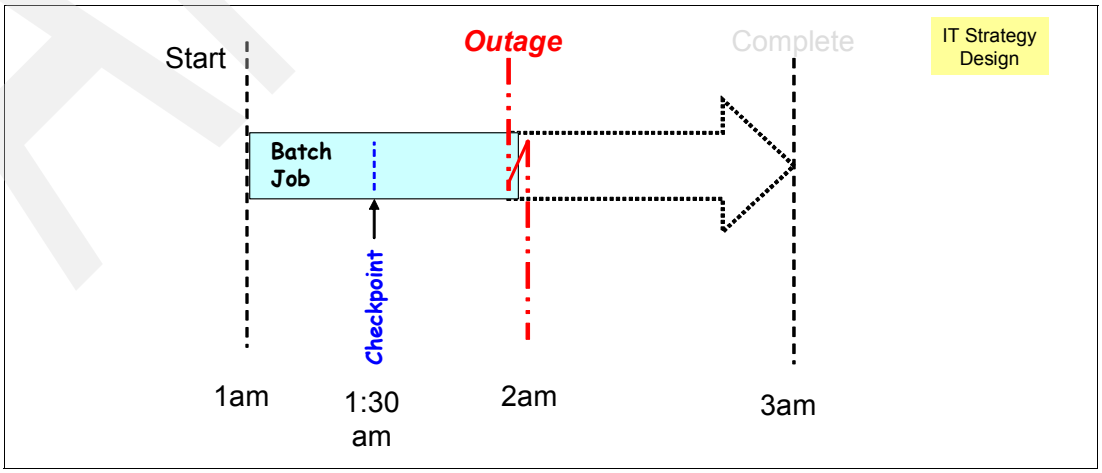

Figure 3-52 Use of checkpoint in the batch job or batch process

There is an additional issue. A batch process typically requires a certain number of inputs, from a variety of sources, in order to run. If your batch job consists of a complex input stream, you will want to consider and design a strategy for reproducing that input stream as well. As an example, consider Figure 3-53 where there are multiple batch processes, running in parallel, that feed each other and, therefore, create dependencies.

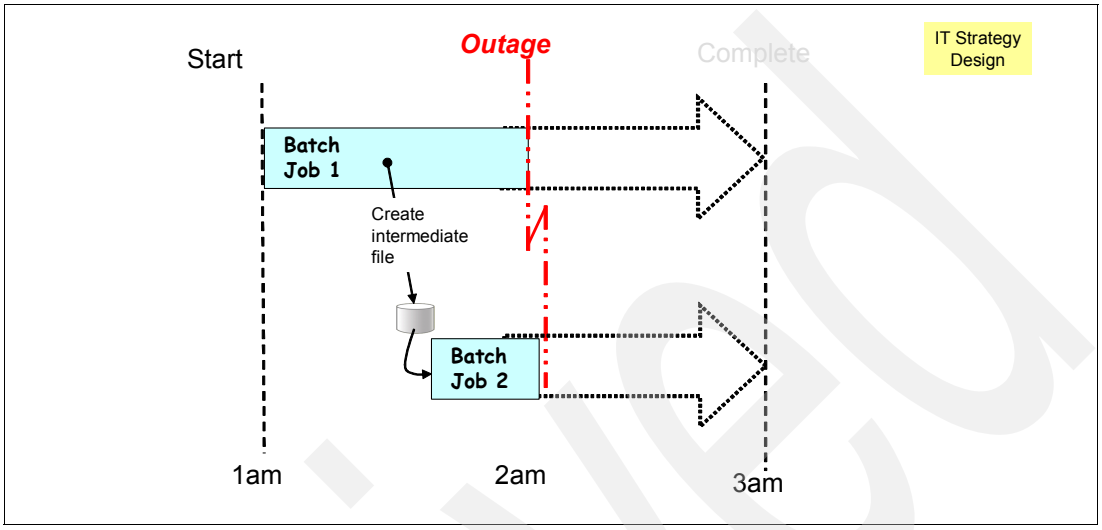

Figure 3-53 Batch process inter-dependencies

In essence, after we have identified these checkpoints and interdependencies, we need primarily need to replicate those checkpoints, that is *transport the contents* of those major checkpoints, to the remote site, as shown in Figure 3-54.

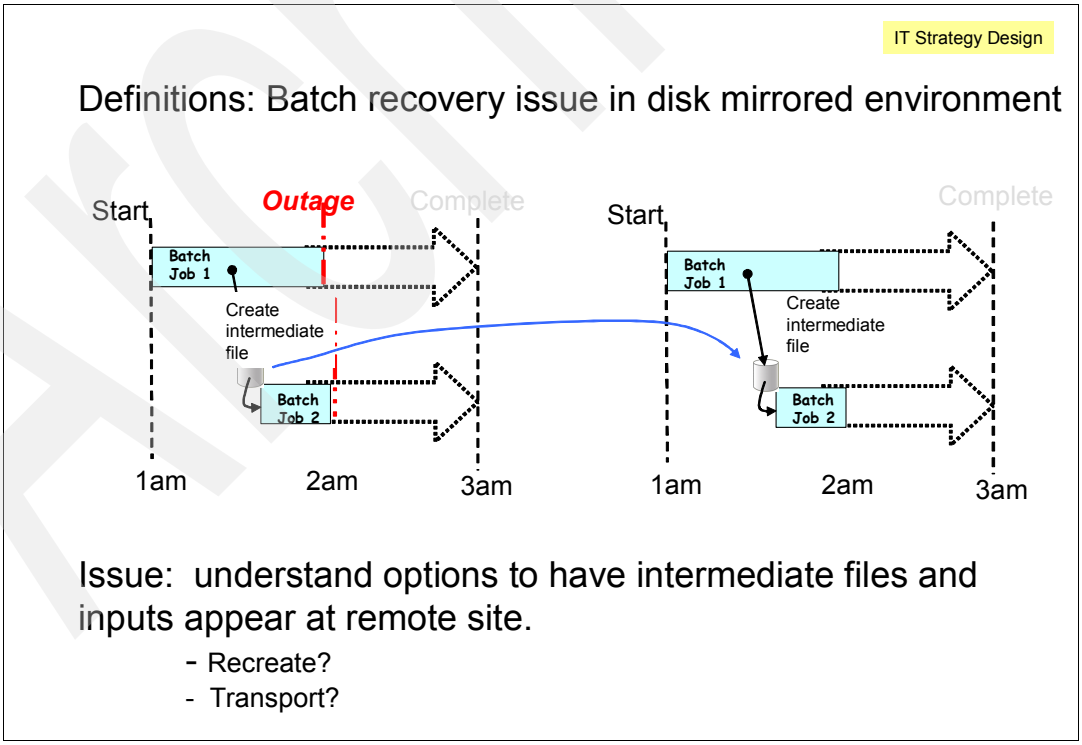

Figure 3-54 Transporting the contents of a checkpoint to the remote site

In consideration of these realities of batch processing in a disk mirrored environment, it is not always necessary to mirror the disk *during* the batch run.

It is likely that your batch recovery in a disk mirrored environment will be strategized on the concept of taking major system-wide checkpoints at the beginning of the batch run or at important milestone points in the middle of the batch run.

Philosophy and guidelines for planning batch recovery in a disk mirrored environment

When implementing disk mirroring, we recommend that you follow these guidelines when designing a batch process recovery strategy in a disk mirrored environment:

1. Understand your existing batch job processes, including the locations and frequencies of checkpoints that you might have today. Document how these jobs run (job scheduler), and how you would do a local recovery in the event of a power-outage.
2. Understand the inputs to the batch processes, and document how these inputs would be saved or recreated if the batch process were needed to be restarted
3. Do not expect to re-engineer the batch process all at the same time.
4. Instead, first take the time to define overall goals and metrics of how you *eventually and ideally* would like your batch process to be recovered in the disk mirrored environment.

   You would envision a future optimized batch processing *checkpoint* environment to give you a vision to which to aim. This batch checkpoint strategy would include identifying and saving the appropriate inputs that you would need to restart at the checkpoint. Note that you will not implement this strategy right away, this is for defining the end vision and goal.

   Your design will include your vision of how you might want to *transport* the necessary input streams and checkpoint data to the remote site.

   Your design should include necessary data management and data allocation policies. Upleveling of the data management and data allocation policies are likely to be necessary to support the checkpoints and transporting of the checkpoints
5. With this master disk mirrored batch recovery strategy in place, you can then scope out a series of phases to reach your final goal.
   - In general, you should *start* with establishing *one* system-wide initial batch checkpoint.
   - The strategy is to bracket (that is, establish and implement) a known master starting point at the beginning of your batch run. Define your ending completion point as well.
   - Identify all necessary inputs that need to be available at the remote recovery site for you to restart your batch successfully at this initial batch checkpoint.
   - Develop a strategy to transport these inputs to the remote site.
     - If the checkpoints and inputs are being made on tape, design a method of transporting the contents of those tapes to the remote site (this might or might not mean staging the data onto disk). You can use tape remote vaulting, other transports, and so forth.
     - If the checkpoints and inputs are being made on disk, design a strategy to remote mirror or otherwise transport the contents of those disks to the remote site.
     - You can use the disk mirroring to establish a known point in time copy of the system, at the remote site.
     - To do this, you establish a *checkpoint replication pool* that contains all necessary input and checkpoint information for the beginning of the batch run.

- You can assign to every necessary application team, their individual responsibility to load their batch cycle checkpoint inputs into this replication pool, by a defined cutoff time that you specify.
- The replication pool is replicated to the remote site.
- After the remote site has successfully saved this set of restart data, you can begin the batch cycle.

– Make an initial level of new data management and data allocation policies for loading these identified batch cycle initial inputs into the checkpoint replication pool. By focusing on this initial master checkpoint, you narrow the scope and enable doing just the first portion of the upleveled data management and data allocation policies that you will find necessary.

– This becomes your first phase of your disk mirrored batch recovery project. At this initial stage, if there is a failure in the batch run, you have successfully replicated everything you need to the remote site to restart the batch run from the beginning.

– After this is well under way, then you can repeat this process to introduce one major intermediate checkpoint, and thereby shorten the batch rerun time that would be required in event of a primary site outage.

– When that is done, you can proceed to implement additional checkpoints. We recommend to follow this strategy of sequentially adding intermediate checkpoints in phases to help you keep progressing towards your final goal.

Figure 3-55 is a useful template for planning batch recovery in a disk mirroring implementation.

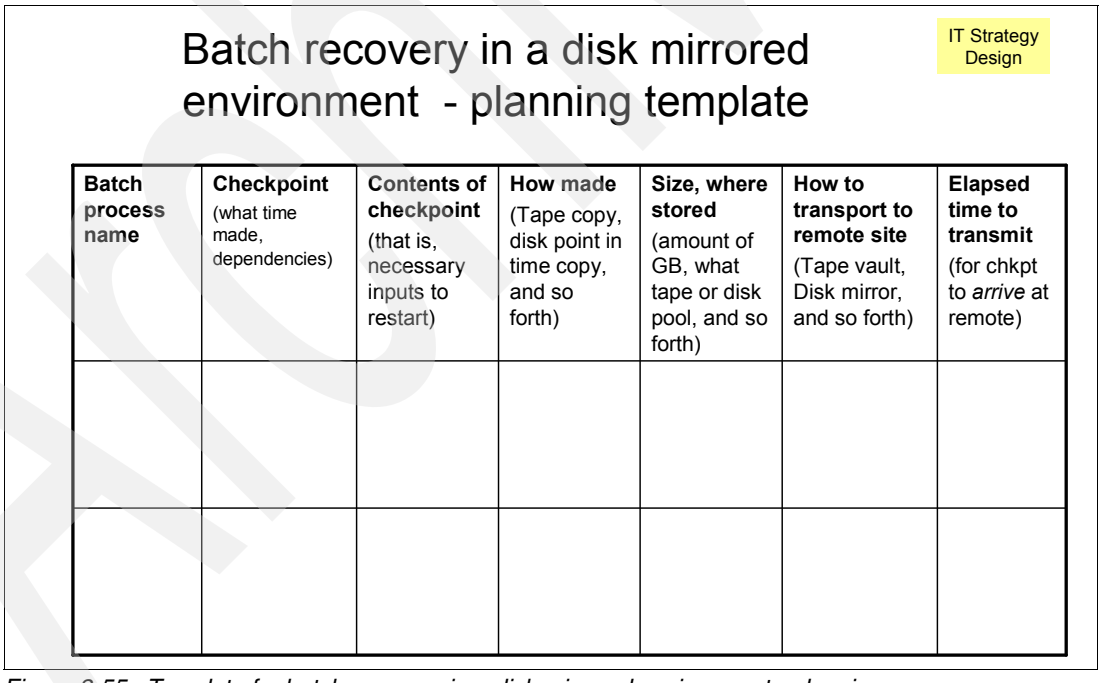

Figure 3-55 Template for batch recovery in a disk mirrored environment - planning

Batch recovery in a disk mirrored environment can be easily handled, when you know in advance of the above issues, and allocate adequate time to consider their impact and design your strategy to handle them in your disk mirrored environment.

Point in Time disk copy considerations in a disk mirrored environment

In a disk mirrored environment, use of point in time disk copy capabilities has some considerations.Disk system internal point-in-time copy facilities make fast logical copies of data, internally in the disk system, by storage controller microcode that replicate pointers to the data (and thus, not needing to copy the data itself).

This has the effect that large amounts of changed data can “suddenly appear” on a LUN or many LUNs, in just a matter of seconds. In a non-disk-mirrored environment, this is not a problem. However, in a disk mirrored environment, this interaction between point in time copy and disk mirroring needs to be examined:

- Some disk systems allow point in time copies to land on in-session primary mirrored volumes.
- If this is allowed, there can be planning or performance implications that must be considered.
- Other disk systems might only allow point in time copies to land on in-session primary mirrored volumes in certain circumstances, or in certain combinations or modes.
- For some disk systems, this might not be allowed at all.

Being aware of this is important, especially where the *existing* non-disk-mirrored environment is making significant use of Point in Time disk copy for:

- Checkpoints for batch recovery
- Fast backups of databases and applications, especially if the volume of data is large

We can consider three possible situations, as shown in Figure 3-56.

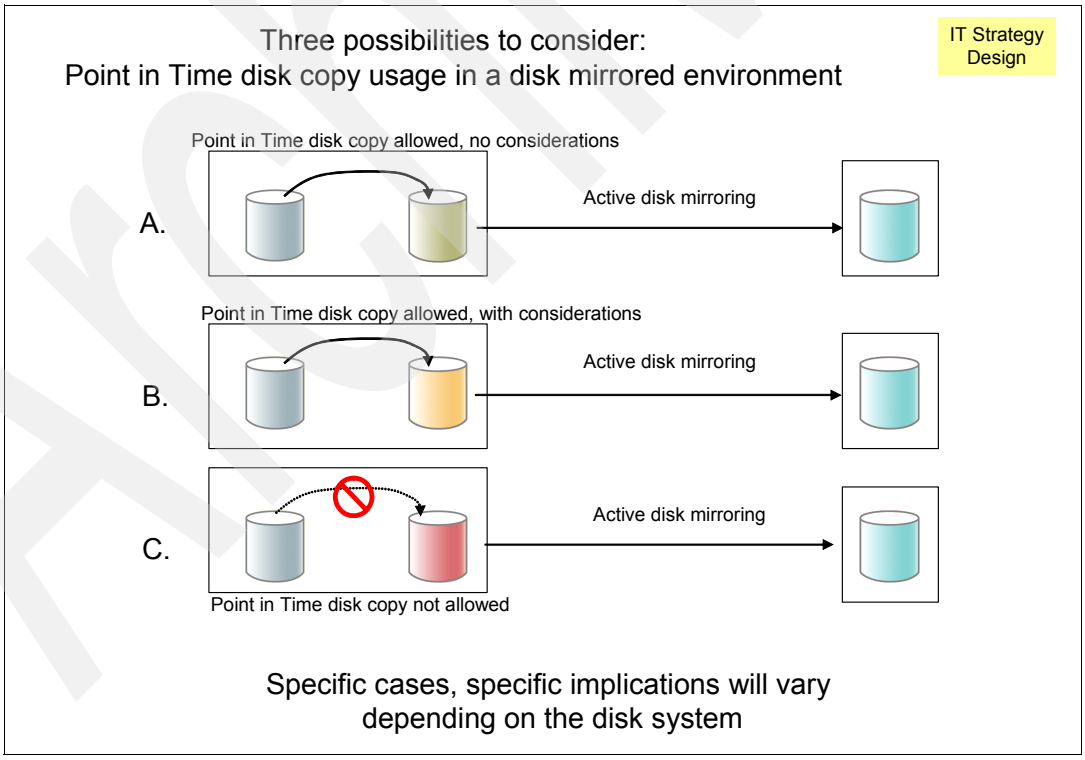

Figure 3-56 Three planning possibilities for point in time copy and disk mirroring interaction

If you are currently using disk point in time disk copy today in a non-mirrored environment, and you want to move to disk mirroring, inquire and be aware of any considerations.

In case B, the typical consideration is that you are allowed to use point in time disk copy onto the primary mirrored volume, but that there might be a time lag window while the newly changed Point in Time data is being copied from to the remote volume.

- In this case, there will be a time lag before the Point in Time copy is complete and available at the remote site.
- In this case, you should design and plan your operational use of Point in Time disk copy such that you are still able to acceptably recover at the remote site, if there is an outage during the time lag window.

Database and software design

We provide a complete separate chapter on database and software considerations when planning for the Business Continuity IT strategy. See Chapter 11, “High Availability clusters and database applications” on page 305.

Summary of IT strategy design

In this section, we discussed:

- IT Infrastructure consolidation and simplification is an important prerequisite to affordable, cost-effective IT Business Continuity.
- High Availability Design consists of:
  - Designing a resilient IT infrastructure that simultaneously can provide IT High Availability, IT Continuous Operations, and IT Disaster Recovery
  - Use of the Business Continuity Tiers of technology for organizing and sorting the possible technology choices by Recovery Time capability
  - Application of Business Process segmentation into three segments
    - Continuous Availability
    - Rapid Data Recovery
    - Backup / Restore
  - Mapping each business process segment onto the appropriate Business Continuity Tier to obtain a mapping of the recovery requirements with the appropriate Business Continuity technology that can provide it
  - Each segment builds upon the foundation of the underlying segment
- High availability servers were overviewed
- Data replication IT strategy planning considerations were reviewed, including:
  - Types of data replication: application / database, server, storage
  - Determining the best type of data replication for me
  - Intersite networking and telecom components
  - A rule of thumb for bandwidth
  - Disk replication, data consistency, and consistency groups
  - Considerations for selecting synchronous versus asynchronous data replication
  - Recovering batch processing
- Database and software design considerations were overviewed

3.6 Manage

Figure 3-57 shows where we are in our ideal Business Continuity planning process. We are the final major portion to *Manage* what we have assessed and designed so far.

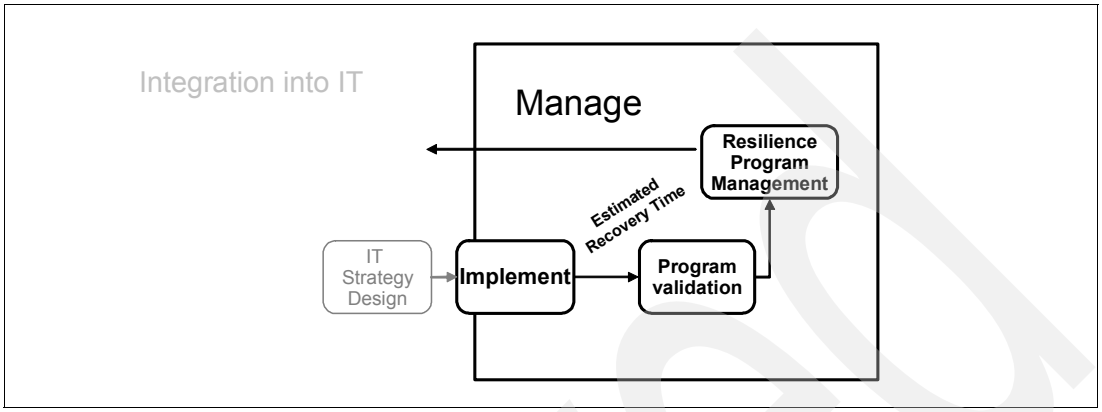

Figure 3-57 Manage: Implement, validate, and refine the program

3.6.1 Implement

After you have created an IT strategy design in the Integration into IT segment of the ideal Business Continuity planning cycle, you can move to the actual execution of the design through a normal project plan that implements each of the following aspects of the Business Continuity plan in a phased manner:

- People
- Processes
- Plans
- Strategies
- Networks
- Platforms

Figure 3-58 shows where we are in the ideal Business Continuity planning cycle.

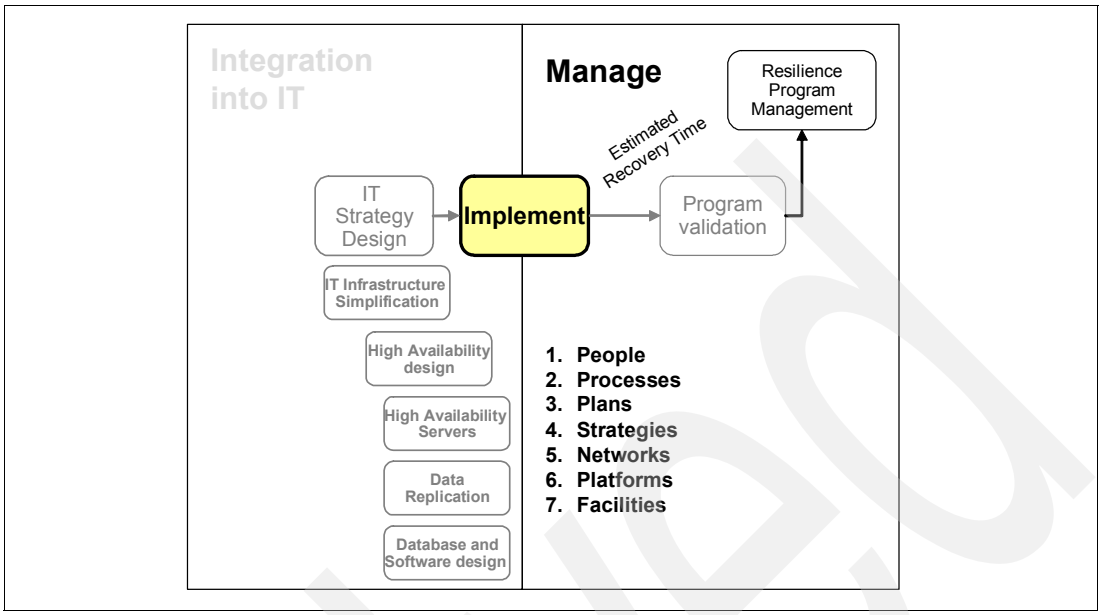

Figure 3-58 Implementation

As we execute this plan, here are some key principles to keep in mind:

- Include IT consolidation and simplification as a foundation for Business Continuity
- Do not try to do everything at the same time

Plan for building up the IT Business Continuity step-by-step:

- Start with where you are today and plan a multi-phase project to build upwards
- Incrementally build towards the final objective
- Each stage builds necessary foundation for the next step

Put into a flow chart, this step-by-step concept might look as shown in Figure 3-59.

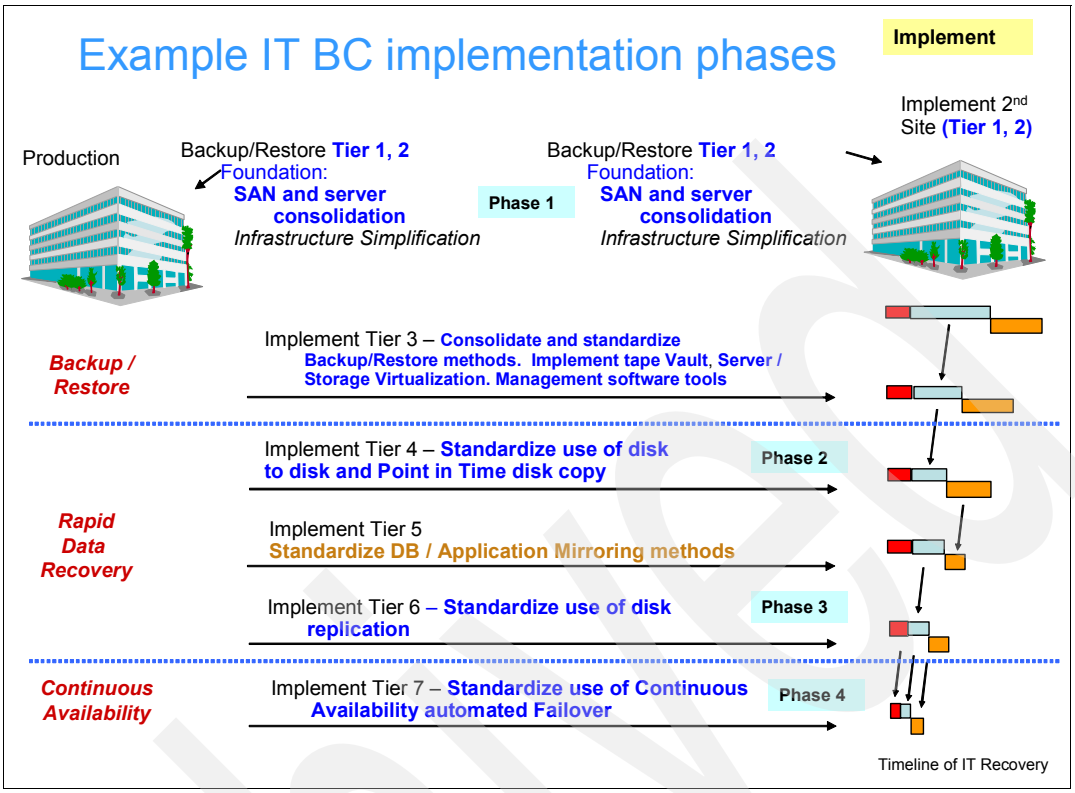

Figure 3-59 Implementation concepts to keep in mind

You can execute a detailed project plan, using normal project planning methodologies, at this point in the ideal Business Continuity planning cycle.

Estimated recovery time

From the Business Continuity program design and IT strategy design, you have an estimated recovery time, likely in the format of KPIs, as we discussed in “Key Performance Indicators” on page 55. These KPIs state a series of estimated recovery times and other recovery metrics.

Note: KPIs can change and improve as each corresponding phase of the project completes and you move to the next phase.

Figure 3-60 is an example of an estimated recovery time KPI. This is only an example. You will need to develop your own KPIs to fit your specific environment.

Example: Key Performance Indicators for Estimated Recovery Time

Implement

	Title of KPI:	Definition:	Measurement:	Target:
Procedural	Backup window	Duration daily planned application outage= 20 min	Total time in minutes appl must be quiesced	Reduce to from 20 min to 1 min by 1Q2007
	Time to recover (i.e. time to switch sites)	Time to switch sites and restore service	Total time in minutes: outage => users online	Reduce to 30 minutes by YE2007
	Testing frequency (preparedness)	Frequency per month of end to end BC test	Number of times/mo for BC test	Improve from 2x/yr to 1x per month
	Average system response time	Application response time at bedside	Average and peak response time in sec	In DR mode, maintain no more than 40% increase in resp. time
	Average problem resolution time	Average time identify, resolve applic. problem	Average time in hours from initial report	Reduce average time to 2 hours by YE2007
Financial	Bandwidth costs	Monthly cost of bandwidth to DR site	Dollars/month of expense	Keep % annual expense growth < 20%
	How much data is being replicated	Amount of data being replicated by PR,Radio	Total production TB allocated to PR, Radio.	Maintain at 20% of total production TB
	Percentage growth rate data	Annual growth % data production PR, Radio.	Quarter to quarter data allocation report	% growth = 10% less than patient growth %

Figure 3-60 Example estimated recovery times documented using Key Performance Indicators

You use these estimated objectives in the next section.

3.6.2 Program validation

Figure 3-61 illustrates where we are in the ideal Business Continuity planning process.

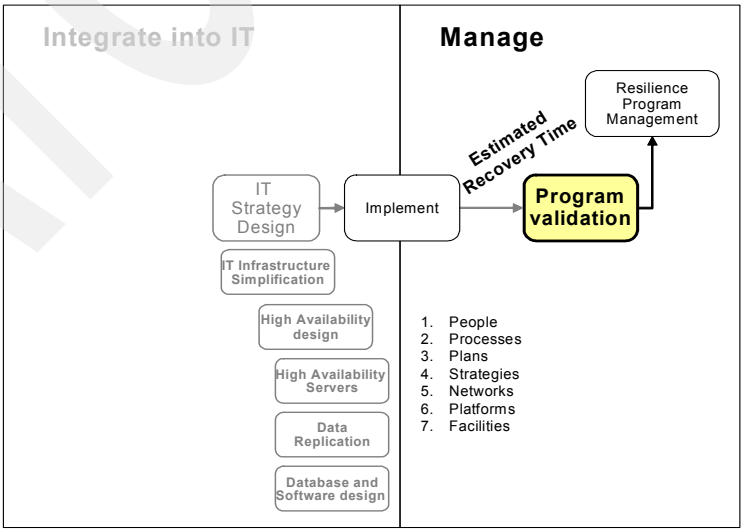

Figure 3-61 Program validation

Let us examine some key points about program validation.

Enabling affordable testing

To review briefly, as we saw in 2.3, “Emerging criteria for IT Business Continuity solutions” on page 21, the following are the emerging requirements for IT Business Continuity solutions:

- Reliability
- Repeatability
- Scalability

Most importantly, these three aspects cannot be proven or refined and tuned unless we have the ability to perform very frequent testing at a very affordable cost. In order to meet these three objectives, within staffing, budget, and time constraints, affordable continuous testing can no longer require intensive labor or time away from other implementation tasks. In the program validation step, we should exploit where ever possible, a high degree of automation.

Good automation for IT Business Continuity solutions provides the ideal foundation for making testing affordable, by providing the ability for repetitive testing on a frequent basis. Eliminating the people requirement for testing can have a significant effect on the reliability and repeatability of the Business Continuity solution; the solution is also easier to test for scalability.

Testing

Testing of both the Business Continuity plan (including the IT Business Continuity plan) is a critical element of a viable Business Continuity strategy capability. Testing enables plan deficiencies to be identified and addressed. Testing also helps evaluate the ability of recovery staff to implement the plan quickly and effectively. Each element of the Business Continuity plan should be tested to confirm the accuracy of the individual recovery procedures and the overall effectiveness of the plan. You need to address the following areas in a contingency test:

- System recovery on an alternate platform from backup tapes
- Coordination among recovery teams
- Internal and external connectivity
- System performance using alternate equipment
- Restoration of normal operations

To derive the most value from the test, explicit test objectives and success criteria should be identified. For example, one test objective might be the recovery of a database, database server, and operating system at an alternate site within eight hours, and database recovery with no errors. The use of test objectives and success criteria enable the effectiveness of each plan element and the overall plan to be assessed. Test results and lessons learned should be documented and reviewed by test participants and other personnel as appropriate. Information collected during the test and post-test reviews that improve plan effectiveness should be incorporated into the plan.

The testing process would have the primary disaster recovery personnel execute the Business Continuity plan on a regular basis. The time interval between testing cycles is shrinking, and should be assumed and planned to continue to shrink.

During each testing cycle, all error conditions would be noted, and driven into a cycle for fixing and resolution. In doing so all of the documentation for any steps which might have been missed or changed over time need also be updated.

Validating RTOs with operations staff only

Validation of the recovery plan also must accommodate the realities of what will occur should actual execution of the plan be required. You need to design test scenarios and conduct test exercises that include a variety of *non-ideal* situations.

An excellent best practices is to test execution of the Business Continuity plan with key Business Continuity design and subject matter experts *declared on "vacation" or otherwise unavailable*. These key personnel can be physically present to observe the test in progress, but they can be marked with a black armband or other means to signify to the executing operational staff, that these persons cannot actually participate in the test.

In this situation, the Business Continuity drill would be executed wholly by the group of alternate personnel at the recovery site following the current documented process. Alternate personnel might have to be utilized at the time of a disaster as the primary team might not be available. It is desirable to have the plan be *customer environment skills neutral*.

Automation of the recovery procedures further assists in this process, as automation can reduce the dependency on people skills, and thus reduce the impact to the Business Continuity plan if some of the disaster recovery team are unavailable due to the disaster.

Training

Training for personnel with Business Continuity plan responsibilities should complement testing. Training should be provided at least annually; new hires who will have plan responsibilities should receive training shortly after they are hired. Ultimately, contingency plan personnel should be trained to the extent that they are able to execute their respective Recovery procedures without aid of the actual document. This is an important goal in the event that paper or electronic versions of the plan are unavailable due to the extent of the disaster situation. Recovery personnel should be trained on the following plan elements:

- Purpose of the plan
- Cross-team coordination and communication
- Reporting Procedures
- Security requirements
- Team-specific processes (Activation/Notification, Recovery, and Reconstitution Phases)
- Individual responsibilities (Activation/Notification, Recovery, and Reconstitution Phases)

3.6.3 Resilience program management

Finally, we discuss the completion of our ideal Business Continuity planning process. In this section we review the following key aspects:

- Raising awareness
- Regular validation
- Change and systems management
- Quarterly management briefings

In addition, we discuss other aspects that are designed to assure the continuous refinement of our Business Continuity plan through ongoing Resilience Program Management, as shown in Figure 3-62.

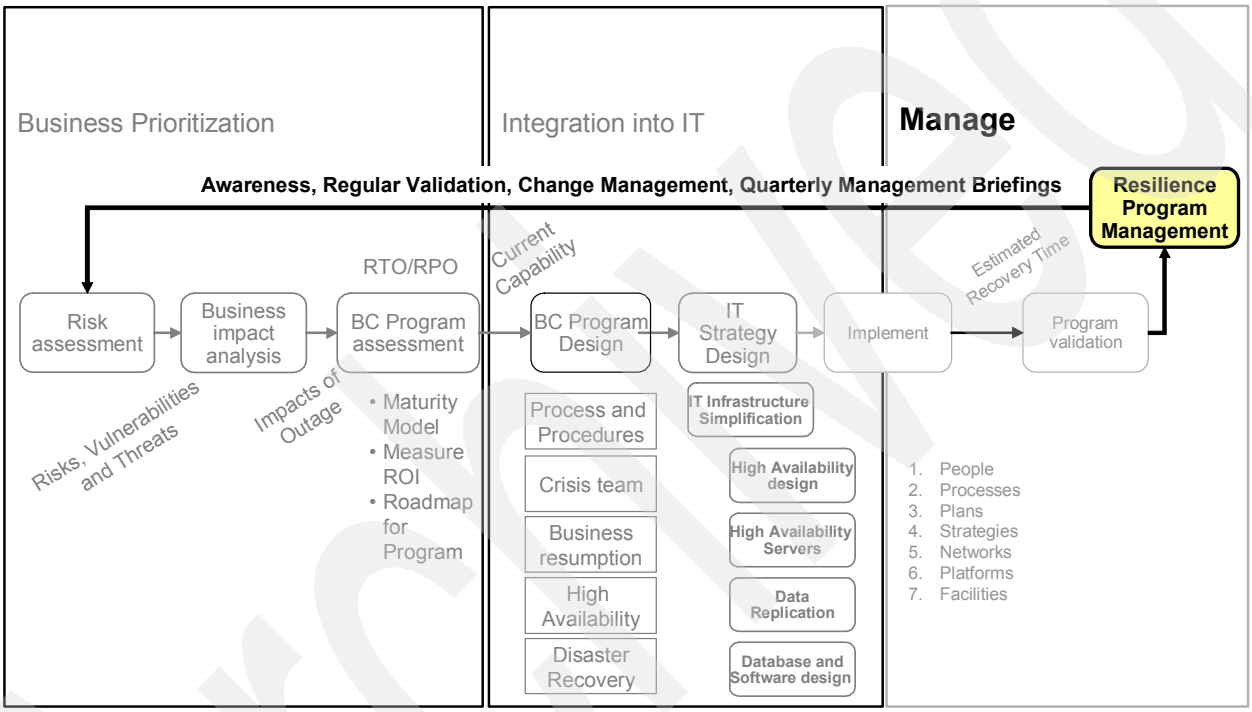

Figure 3-62 Resilience program management

The key point is that each of the steps in this phase is designed to create awareness, new requirements, and input for returning to the beginning of the cycle. As staff and management review through Resilience Program Management, the progress of the implementation and results of program validation, we create new inputs for the cycle to repeat. A continuous reiteration of new Risk Assessments, leading to refined Business Impact Analysis, and so forth is the end result of the ideal Business Continuity planning cycle.

Let us now examine aspects of Resilience Program Management.

Awareness, regular validation

KPIs play an essential role in publicizing the performance and validating the ongoing effectiveness of the Business Continuity plan. As we suggested in “Key Performance Indicators” on page 55, publicizing the KPIs for the Business Continuity plan in key locations within the organization. Show the target for each Business Continuity plan KPI, and show the progress toward that target. With good project management, your staff will be aware of the organization’s validation status in achieving the KPIs. progress in validating the motivated and feel pride in reaching the KPI targets.

Change management, systems management

To be effective, the Business Continuity plan must be maintained in a ready state that accurately reflects system requirements, procedures, and policies. IT systems undergo frequent changes because of shifting business needs, technology upgrades, or new internal or external policies. Therefore, it is essential to review and regularly update the plan to ensure new information is documented and contingency measures are revised if required.

As a general rule, the plan should be reviewed for accuracy and completeness at least on a quarterly basis, whenever significant changes occur to any element of the plan. Certain elements will require frequent reviews, such as application inventory lists, application release and patch levels, server and storage inventory lists. People and process updates will include call tree accuracy, contact lists, and so forth. Automation can help greatly in this area to assist in evaluating lan contents and procedures more frequently.

At a minimum, plan reviews should focus on the following elements:

- Operational requirements
- Security requirements
- Technical procedures
- Hardware, software, and other equipment (types, specifications, and amount)
- Names and contact information about team members
- Names and contact information about vendors, including alternate and off-site POCs (Points of Contact)
- Alternate and off-site facility requirements
- Vital records (electronic and hardcopy)

Because the Business Continuity plan contains potentially sensitive operational and personnel information, its distribution should be marked accordingly and controlled. Typically, copies of the plan are provided to recovery personnel for storage at home and office. A copy should also be stored at the alternate site and with the backup tapes. Storing a copy of the plan at the alternate site ensures its availability and good condition in the event local plan copies cannot be accessed as a result of the disaster.

The Business Continuity planning coordinator should maintain a record of copies of the plan and to whom they were distributed. Other information that should be stored with the plan include contracts with vendors (SLAs and other contracts), software licenses, system user manuals, security manuals, and operating procedures.

Changes made to the plan, strategies, and policies should be coordinated through the disaster recovery planning coordinator, who should communicate changes to the representatives of associated plans or programs, as necessary. The disaster recovery planning coordinator should record plan modifications using a Record of Changes, which lists the page number, change comment, and date of change. A sample of a Record of Changes is shown in Table 3-3.

Table 3-3 Sample record of changes

Record of Changes			
Page number	**Change comment**	**Date of change**	**Signature**

The Business Continuity planning coordinator should coordinate frequently with associated internal and external organizations and system points-of-contact (POC) to ensure that impacts caused by changes within either organization will be reflected in the Business Continuity plan. Strict version control should be maintained by requesting old plans or plan pages to be returned to the planning coordinator in exchange for the new plan or plan pages.

The planning coordinator also should evaluate supporting information to ensure that the information is current and continues to meet system requirements adequately. This information includes the following:

- Alternate site contract, including testing times
- Off-site storage contract
- Software licenses
- Memorandum of understanding or vendor SLAs
- Hardware and software requirements
- Security requirements
- Recovery strategy
- Contingency policies
- Training and awareness materials
- Testing scope

Although some changes can be quite visible, others will require additional analysis. The Business Impact Analysis should be reviewed periodically and updated with new information to identify new contingency requirements or priorities. As new technologies become available, preventive controls can be enhanced and recovery strategies can be modified.

This phase must be totally connected with the IT Change Management team, which is in charge of informing the disaster recovery team of eventual updates in the system, which cause modification in the Recovery Plans.

Quarterly management briefings

KPIs are a valuable tool to brief senior management on the status and progress of the Business Continuity plan. We recommend that at least quarterly, a major update of the Business Continuity plan status should be presented and reviewed to senior management, especially including the Lines of Business.

Tip: A case study which can provide an excellent example and template of material for a quarterly management briefing, is in this book in Chapter 7, “Next Step Business Continuity workshop: Case Study” on page 233.

During the quarterly management briefing, the following key business and process points should be reiterated:

- IT is presenting the Business Continuity status to the Lines of Business, to assure that LOB and IT are in joint ownership and partnership of the of end-to-end Business Continuity plan.
- Use KPIs to *bottom-line* the briefing, and to restate the business value and justification for the Business Continuity plan.
- In the end, it is the Line of Business that has the ultimate risk, and IT ultimately is in the support role.
- IT is taking responsibility, however, IT wants Line of Business to be in partnership for Business Continuity.

Reiterating these key aspects also assists IT in assuring that the value-add of the Business Continuity plan is visible to the Lines of Business. This also assists IT in assure that Line of Business support for adequate funding of the IT Business Continuity plan is in place.

Repeating the previous chart showing KPIs, we suggest these are also key tools to be used in the quarterly presentation to management, as shown in Figure 3-63.

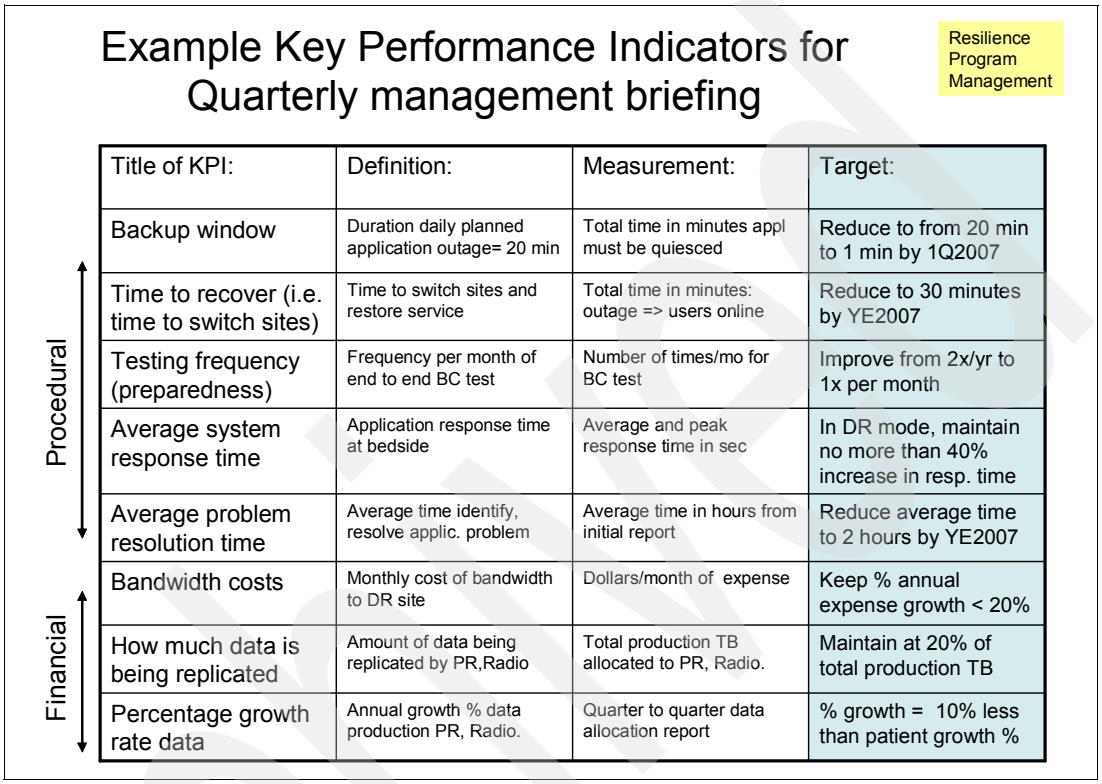

Title of KPI:	Definition:	Measurement:	Target:
Backup window	Duration daily planned application outage= 20 min	Total time in minutes appl must be quiesced	Reduce to from 20 min to 1 min by 1Q2007
Time to recover (i.e. time to switch sites)	Time to switch sites and restore service	Total time in minutes: outage => users online	Reduce to 30 minutes by YE2007
Testing frequency (preparedness)	Frequency per month of end to end BC test	Number of times/mo for BC test	Improve from 2x/yr to 1x per month
Average system response time	Application response time at bedside	Average and peak response time in sec	In DR mode, maintain no more than 40% increase in resp. time
Average problem resolution time	Average time identify, resolve applic. problem	Average time in hours from initial report	Reduce average time to 2 hours by YE2007
Bandwidth costs	Monthly cost of bandwidth to DR site	Dollars/month of expense	Keep % annual expense growth < 20%
How much data is being replicated	Amount of data being replicated by PR,Radio	Total production TB allocated to PR, Radio.	Maintain at 20% of total production TB
Percentage growth rate data	Annual growth % data production PR, Radio.	Quarter to quarter data allocation report	% growth = 10% less than patient growth %

Figure 3-63 Example of key performance indicators for quarterly management briefings

As staff and management review Resilience Program Management, they create new inputs for the Business Continuity planning cycle to repeat.

3.7 Summary

Figure 3-64 illustrates the ideal Business Continuity planning process.

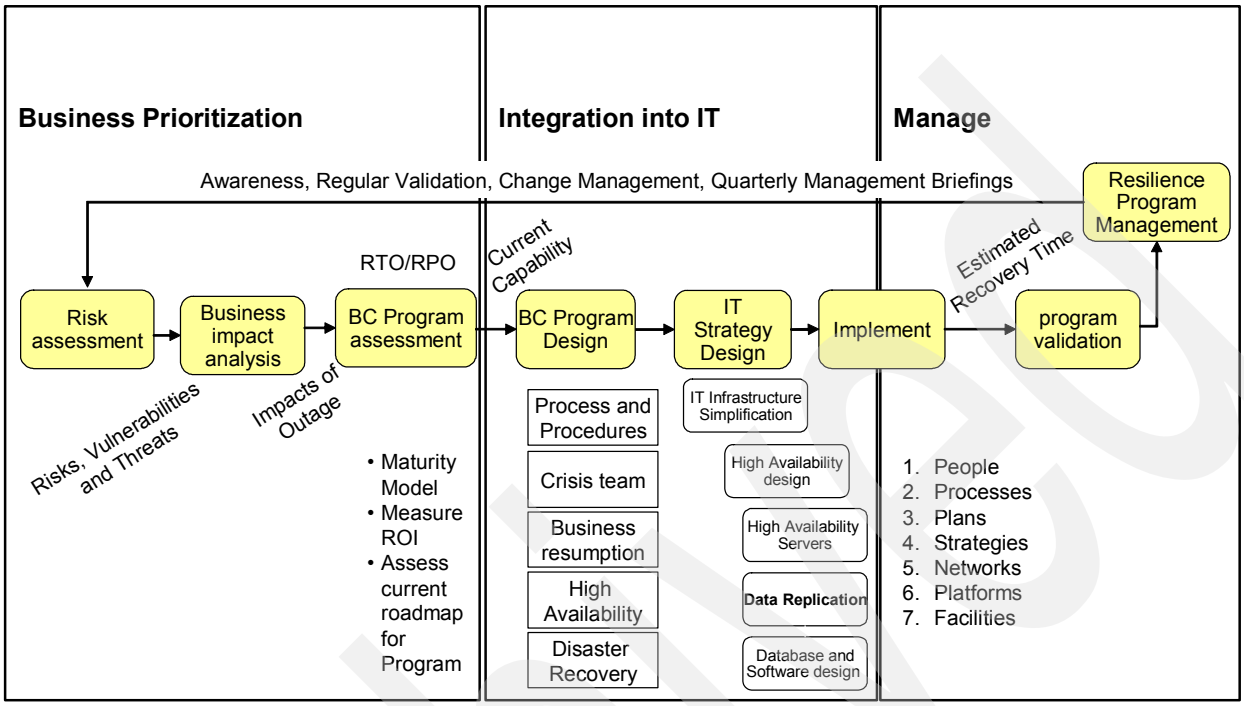

Figure 3-64 Ideal Business Continuity planning process

This process is based on methodologies used by IBM Global Technology Services, Business Continuity and Recovery Services. The ideal Business Continuity planning process is a closed loop, supporting continuing iteration and improvement as the objective.

We reviewed the three major sections to the planning process:

- Business prioritization
- Integration into IT
- Manage

We also reviewed in detail each of the steps in the flow chart shown in Figure 3-64.

Tip: To bring this chapter into practical, usable focus, we provide a distilled, step-by-step applied example of all of the concepts of this chapter. We also present this information in a workshop format and a full case study in the following chapters:

- Chapter 6, “The Next Step Business Continuity workshop” on page 173
- Chapter 7, “Next Step Business Continuity workshop: Case Study” on page 233

We highly recommend that you review the Case Study in particular, because it brings to life all of the concepts that we have discussed in this chapter.

In the remaining chapters of this book, we go into more planning details specific to the IT implementation of IBM System Storage Business Continuity technology.

4

Tier levels of Business Continuity solutions

This chapter describes the tier levels of Business Continuity solutions, including the seven tiers of Business Continuity solution definition. The seven tiers of Business Continuity can be mapped to the three Business Continuity solution segments. With this approach, you can develop your Business Continuity solution on various business applications according to your business needs. (For information about the methods and considerations for planning Business Continuity solutions using the IBM Business Continuity planning method, see Chapter 5, “Business Continuity Solution Selection Methodology” on page 151.)

With the levels of Business Continuity solutions, we also look in detail at the hierarchical dependencies of the system layer architecture. This tiered approach applies also to heterogeneous IT platforms as discussed in Chapter 8, “Planning for Business Continuity in a heterogeneous IT environment” on page 251.

4.1 Seven Business Continuity tiers

This section describes the seven Business Continuity tiers and gives a more detailed breakdown of the possibilities that exist under each tier (as shown in Figure 4-2 on page 139).

The goal of any Business Continuity planning is to protect the most business critical processes and to minimize unplanned downtime. Keep in mind that planning for any type of an impact tolerant solution is always subject to balancing the solution design versus the downtime versus the cost, as illustrated in Figure 4-1.

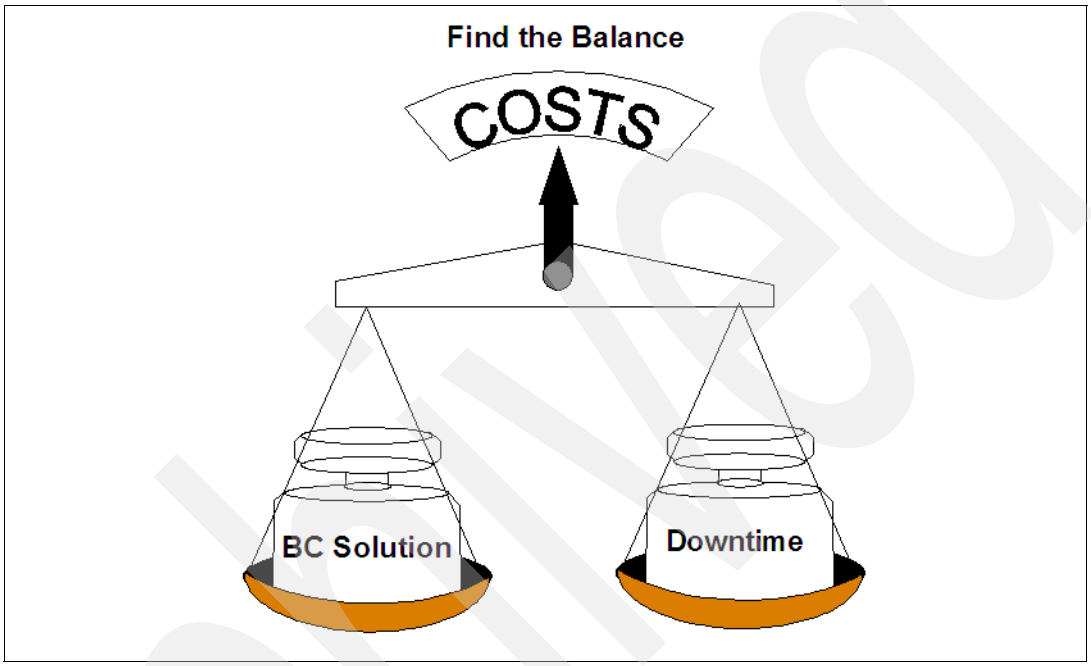

Figure 4-1 Find the balance

The recovery time of any of the seven Business Continuity tiers depends heavily on:

- Recovery time for the data availability
- Recovery of the IT infrastructure
- Restoring the operational processes
- Restoring the business processes

We discuss the usage of the Business Continuity tiers as a tool for designing Business Continuity solutions in Chapter 5, “Business Continuity Solution Selection Methodology” on page 151.

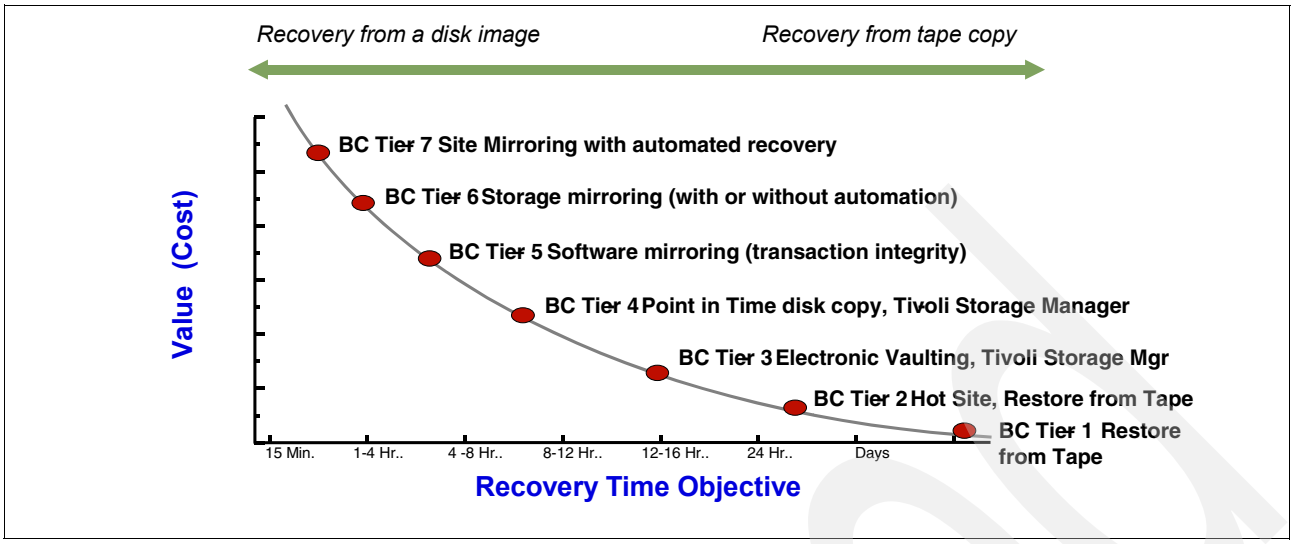

Figure 4-2 Seven tiers of Business Continuity solutions

Note: The RTO timeline of Figure 4-2 is a general guideline for understanding what the Business Continuity technology can do at these Business Continuity tier levels. We discuss RTO in Chapter 3, “Business Continuity planning, processes, and execution” on page 43.

4.2 A breakdown of the seven tiers

In 1992, the SHARE user group in the United States together with IBM defined a set of Business Continuity tier levels to address the need to properly describe and quantify various different methodologies for successful mission-critical computer system Business Continuity implementations. Accordingly, within the IT Business Continuance industry, the tier concept continues to be used, and it is very useful for describing today's Business Continuity capabilities. The tiers, therefore, need only to be updated for today's specific Business Continuity technologies and associated RTO and RPO.

The seven tiers of Business Continuity solutions offer a a simple methodology for defining your current service level, the current risk, and the target service level and target environment. Each tier builds upon the foundation of the previous tier. In this way, a properly designed Business Continuity solution protects investment and enhances the Business Continuity solution step-by-step over time. See 4.3, “The relationship of Business Continuity tiers and segments” for more information.

Note: The Business Continuity tiers should not be confused with tiered storage. The concepts are completely separate and independent. Business Continuity tiers are a logical construct for classifying levels of system recoverability. Tiered storage is a term used in Information Life Cycle Management (ILM) to describe different classes of storage hardware used for different types of data.

4.2.1 Business Continuity Tier 0: No off-site data

Businesses with a *Tier 0 Business Continuity* solution have no Business Continuity plan. For these businesses, the following facts are true:

- There is no saved information, no documentation, no backup hardware, and no contingency plan.
- Typical recovery time is unpredictable. In fact, it might not be possible to recover at all.

4.2.2 Business Continuity Tier 1: Data backup with no hot site

Businesses that use *Tier 1 Business Continuity* 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. Examples of Tier 1 Business Continuity solutions include:

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

4.2.3 Business Continuity Tier 2: Data backup with a hot site

Businesses using *Tier 2 Business Continuity* solutions make regular backups on tape. This solution 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 in response time when compared to Tier 1 Business Continuity solutions. Examples of Tier 2 Business Continuity solutions include:

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

4.2.4 Business Continuity Tier 3: Electronic vaulting

Tier 3 Business Continuity solutions use components of Tier 2 solutions. Additionally, some mission critical data is electronically vaulted. This electronically vaulted data is typically more current than that which is shipped with PTAM. These solutions also add higher levels of automation, as a result there is less data recreation or loss. Examples of Tier 3 Business Continuity solutions include:

- Electronic Vaulting of Data, (i.e. remote tape via channel extension)
- IBM Tivoli Storage Manager (Disaster Recovery Manager)

4.2.5 Business Continuity Tier 4: Point-in-time copies

Tier 4 Business Continuity 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. Examples of Tier 4 Business Continuity solutions include:

- 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

4.2.6 Business Continuity Tier 5: Transaction integrity

Tier 5 Business Continuity solutions is a Business Continuity Tier reserved for application software and database replication at the transaction level. This Business Continuity 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. Examples of Tier 5 Business Continuity solutions include:

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

4.2.7 Business Continuity Tier 6: Zero or little data loss

Tier 6 Business Continuity 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. Examples of Tier 6 Business Continuity solutions include:

- 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

4.2.8 Business Continuity Tier 7: Highly automated, business integrated solution

Tier 7 Business Continuity solutions include all the major components being used for a Tier 6 Business Continuity solution with the additional integration of complete automation for servers, storage, software and applications and network. This allows a Tier 7 Business Continuity solution to ensure consistency of data above that which is granted by Tier 6 Business Continuity 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. Examples of Tier 7 Business Continuity solutions include:

- 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

4.3 The relationship of Business Continuity tiers and segments

This section provides a simplified approach of Business Continuity solutions. Even though each of the Business Continuity tiers Business Continuity has a different RTO, a collective of tiers to Business Continuity mapping will simplify the decision making process. Thus, you can choose the segment of the Business Continuity solution, then you can decide on the most appropriate Business Continuity solution for the respective Business Continuity tier.

4.3.1 The mapping of seven tiers to the three Business Continuity segments

We defined the three Business Continuity segments in Chapter 3, “Business Continuity planning, processes, and execution” on page 43.

Figure 4-3 describes how you can map these three segments to the seven tiers of Business Continuity solutions. Each tier builds upon the previous tier.

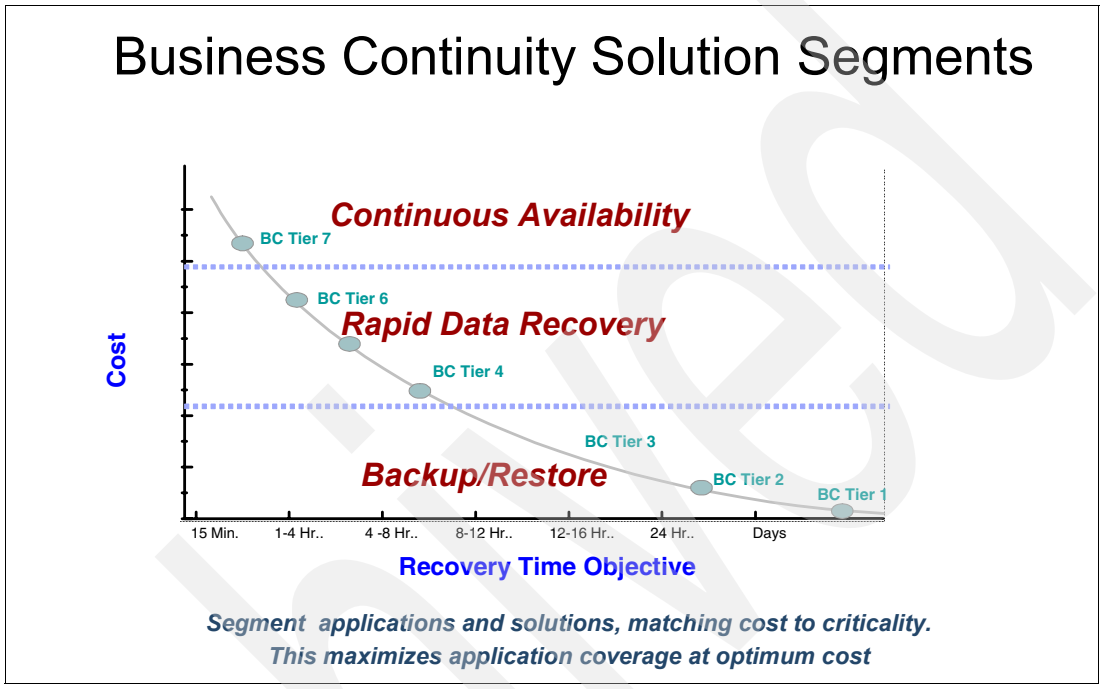

Figure 4-3 Business Continuity solution segments

Figure 4-4 illustrates how the Backup and Restore segment maps to Business Continuity tiers 1, 2, and 3,

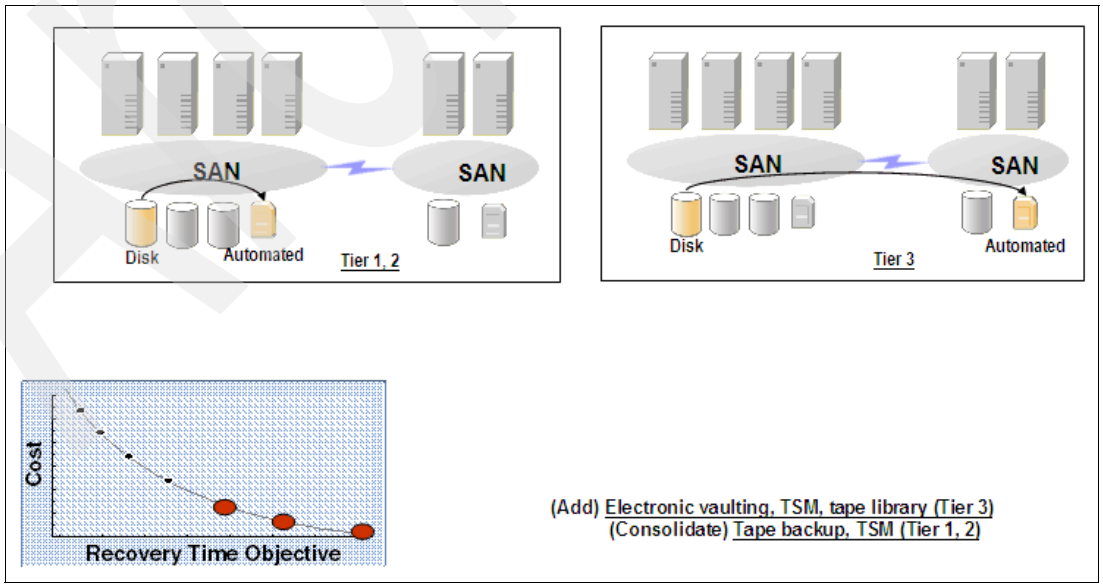

Figure 4-4 Backup and Restore segment and Business Continuity solution Tier 1, 2, and 3

Figure 4-5 illustrates how the Rapid Data Recovery segment maps to Business Continuity tiers 4, 5, and 6.

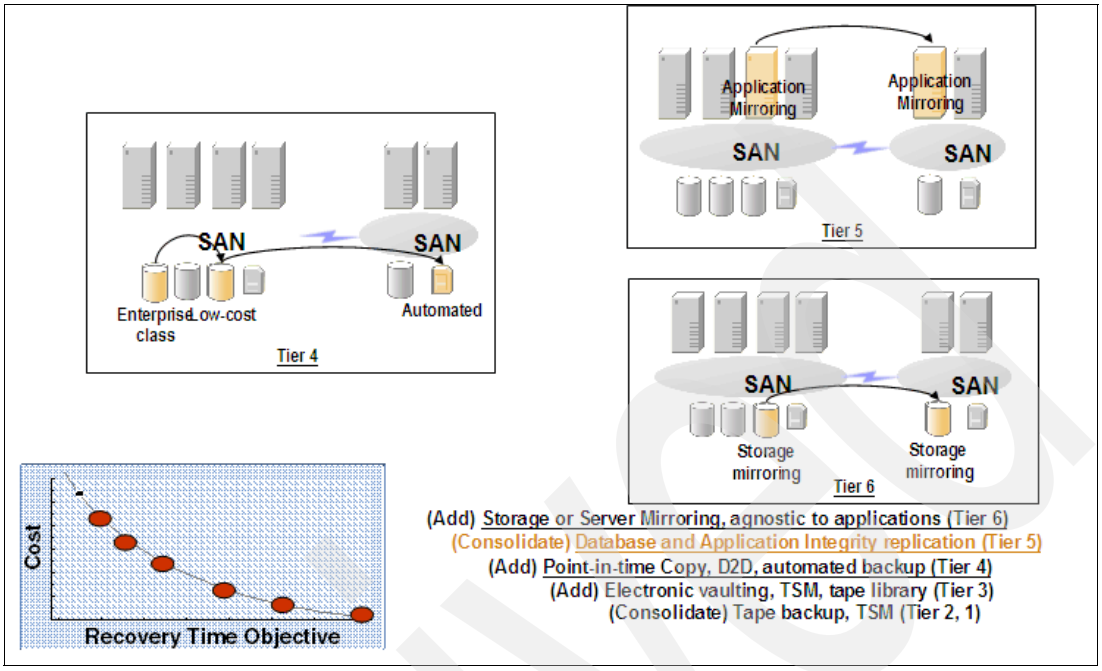

Figure 4-5 Rapid Recovery segment and Business Continuity solution Tier 4, 5, and 6

Figure 4-6 illustrates how the Continuous Availability segment maps to Business Continuity solution Tier 7.

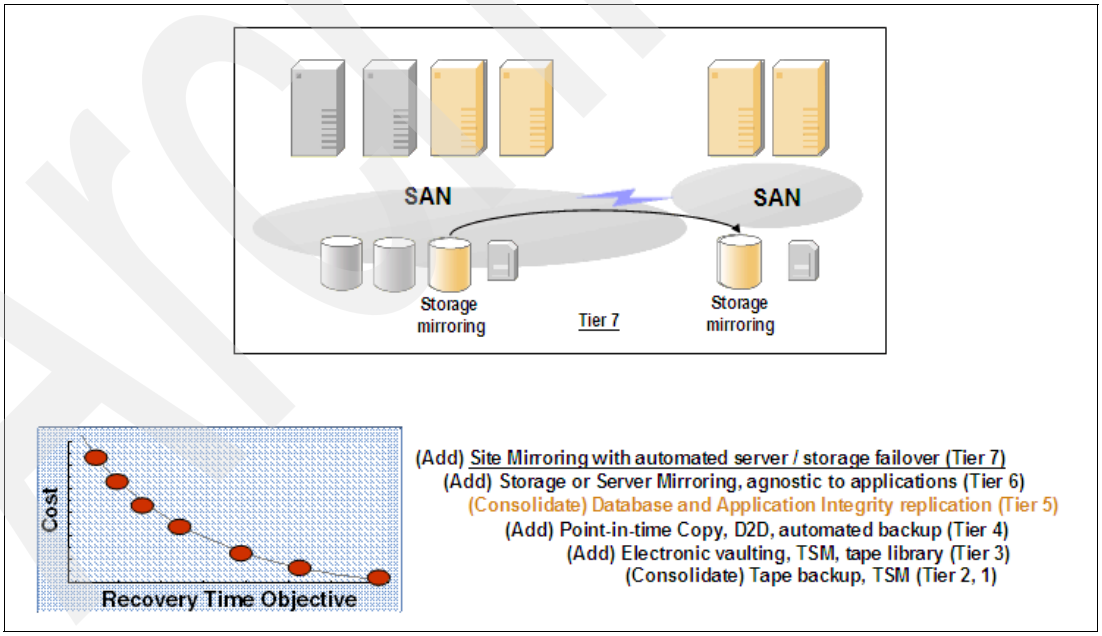

Figure 4-6 Continuous Availability segment and Business Continuity solution Tier 7

Based on the business assessment that we discussed in 3.4.1, “Risk assessment” on page 47 and the segment to tier mapping that we present here, you can determine the appropriate storage technology to achieve the business objectives on the Business Continuity.

4.4 Selecting the optimum Business Continuity solution

This section includes an overview of how to select the optimum Business Continuity solution for your organization. You can find a more detailed discussion in Chapter 5, “Business Continuity Solution Selection Methodology” on page 151. It is important to understand that the cost of a solution must be in reasonable proportion to the business value of IT. You do not want to spend more money on a Business Continuity solution than the financial loss you would suffer from a disaster.

Based on the following objectives, it is relatively simple to decide, as a business, which solution to select according to how much you can afford to spend and the speed at which you need your data recovered. The quicker the recovery the higher the cost:

- RTO: How long can you afford to be without your systems?
- RPO: When it is recovered, how much data can you afford to recreate?
- Degraded operations objective (DOO): What will be the impact on operations with fewer data centers?
- Network Recovery Objective (NRO): How long to switch over the network?
- Recovery Distance Objective (RDO): How far away the copies of data need to be located?

Normally all the components that make up continuous availability are situated in the same computer room. The building, therefore, becomes the single point-of-failure. While you must of course be prepared to react to a disaster, the solution that you select might be more of a recovery solution than a continuous-availability solution. A recovery solution must then be defined by making a trade-off among implementation costs, maintenance costs, and the financial impact of a disaster, resulting from performing a business impact analysis of your business.

Figure 4-7 represents a simple starting point for determining your Business Continuity solution.

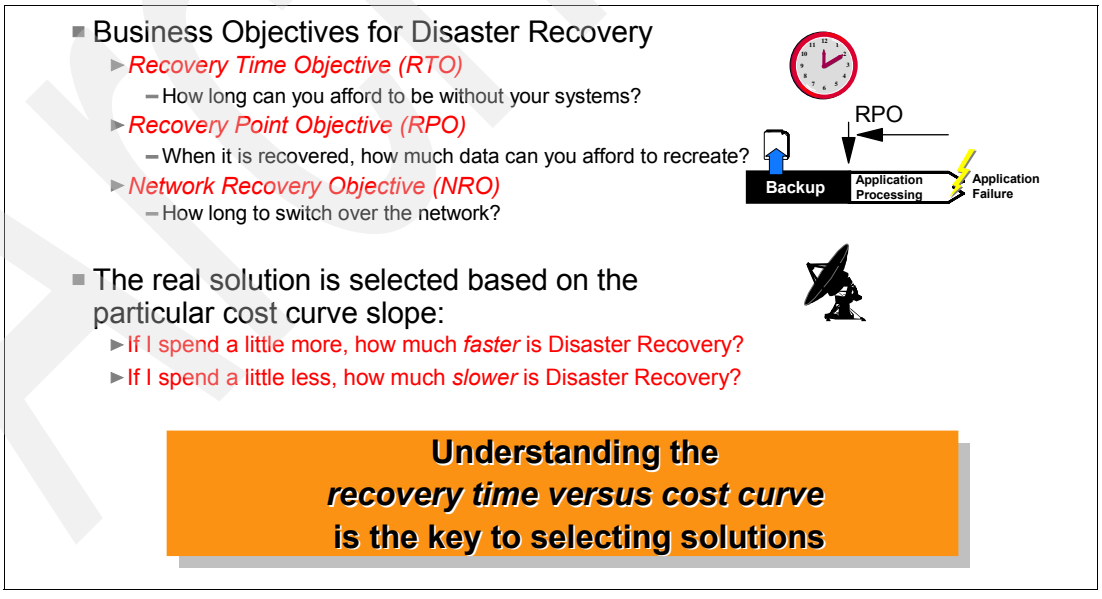

Figure 4-7 Business Continuance objectives

4.4.1 Four key objectives

The typical recovery time associated with each tier is just a rough indication of the time that an installation usually needs to restore its computing services. In an actual disaster, however, there are many other considerations. For example, some installations can tolerate resuming their services after a longer period but with maximum data currency. Other installations must resume their services as soon as possible, regardless of data currency. Others need both a short recovery time and maximum data currency.

By using the four key objectives listed in Figure 4-8 and by assessing your actual Business Continuity requirements pertinent to your site, you have an excellent starting point for your Business Continuity Plan. See Chapter 3, “Business Continuity planning, processes, and execution” on page 43 for more details on developing a Business Continuity plan.

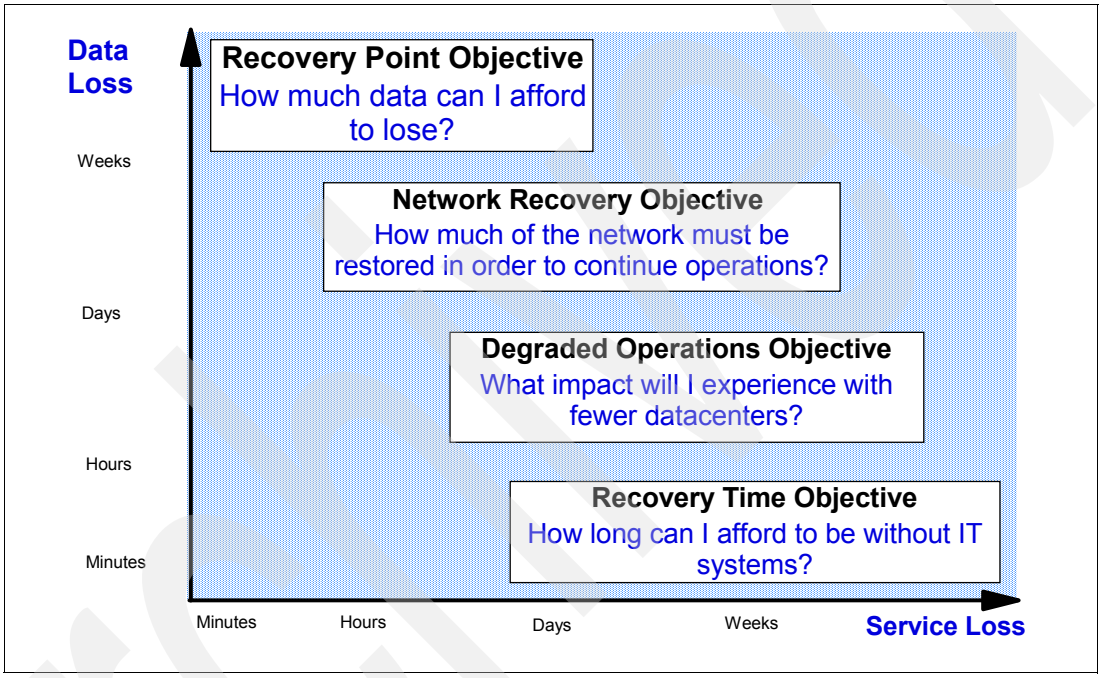

Figure 4-8 Four key objectives

4.4.2 Cost of outage versus cost of solution

The cost of a Business Continuity solution is worked out according to how quickly you need to recover your data versus how much it will cost the company in lost revenue due to being down and unable to continue normal business processing. The shorter the time period required to recover the data to continue business processing the higher the cost, as indicated in Figure 4-9.

The other important point is that the longer a company is unable to process transactions the more expensive the outage is going to be for the company. The optimal solution is, therefore, where the cost of the solution curve and the cost of the outage curve intersect in the cost/time window.

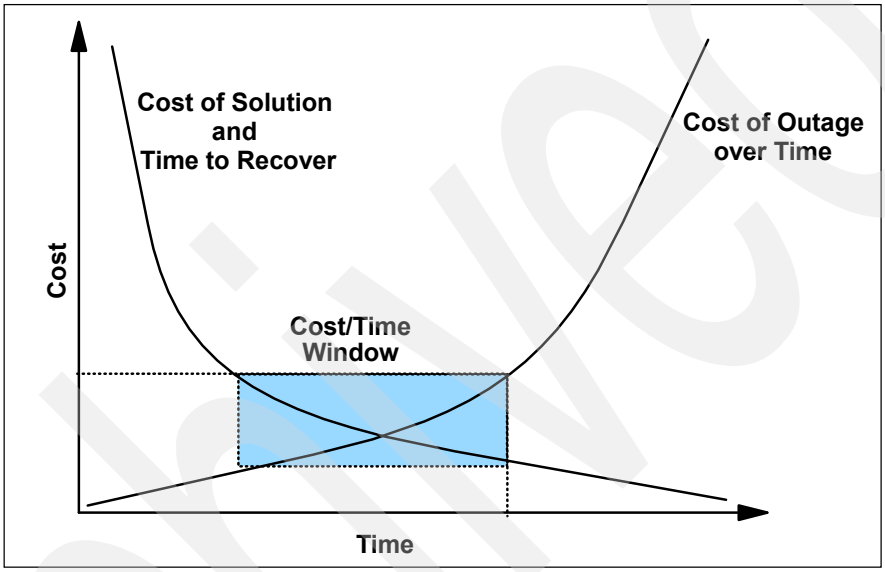

Figure 4-9 Cost of outage versus cost of solution

4.4.3 Hierarchical dependencies of the system layer architecture

Considering the seven tiers of Business Continuity solutions, we have to look in detail at the individual levels within the system components. To have a reliable and proven infrastructure for Business Continuity, redundancy of the individual system level components is mandatory, as shown in Figure 4-10 and Figure 4-11.

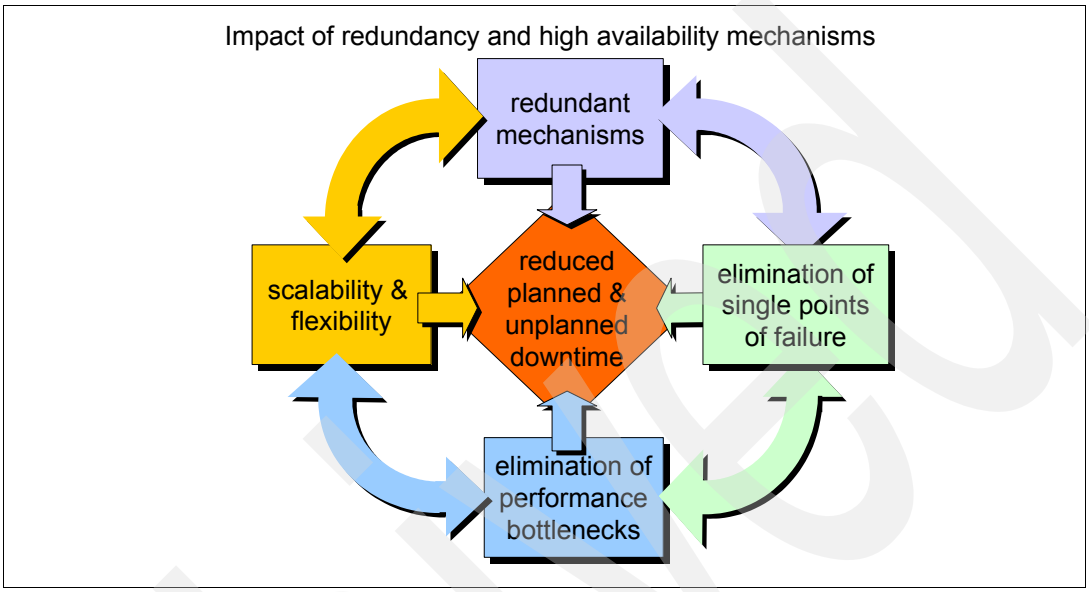

Figure 4-10 Impact of redundancy and high availability

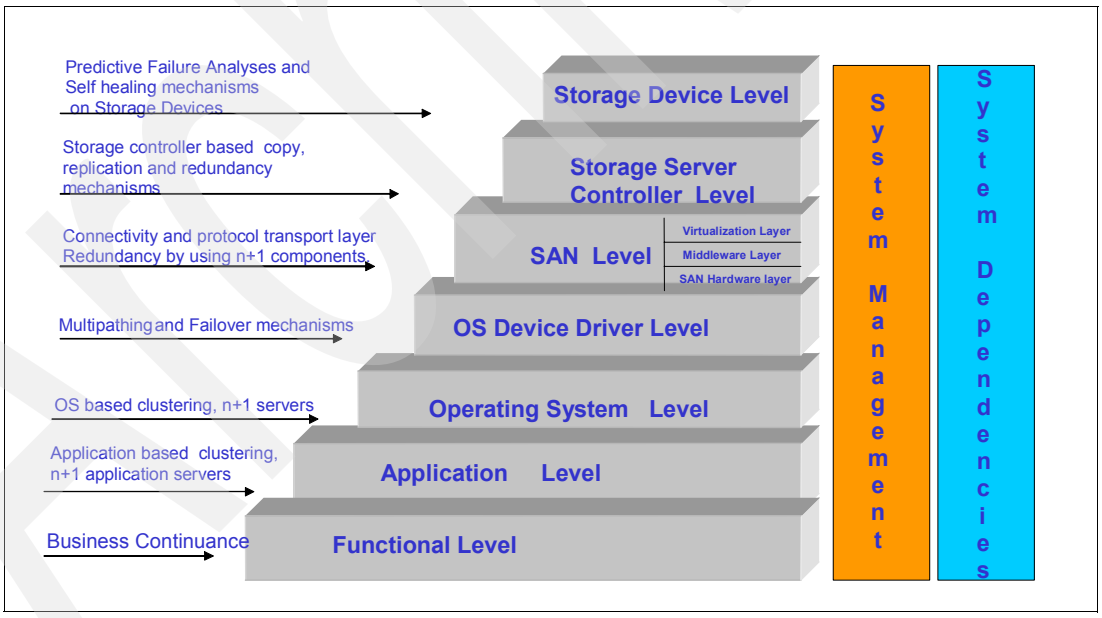

Figure 4-11 Components of a highly available system landscape

Storage device level

At the storage device level (storage objects) we refer to the physical entity, such as, the disk drive or a tape. For availability the device level will be protected on the built in functionality of the device itself. This functionality can be implemented with redundancy by using spare tracks (disks) or by special writing mechanisms (tape) as with IBM LTO cartridges and IBM 3592 technology.

Storage Server Controller (storage system) level

The Storage Controller Level itself interfaces between the SAN or the server or servers and the storage device or devices. The storage controller is responsible for the execution of all storage related built in functionality:

- RAID level
- Built in copy functionality such as Point-in-time Copy, Remote Mirroring
- Built in high availability mechanisms (redundancy, failover)

Storage Area Network level

Storage Area Network (SAN) level redundancy can be achieved by using either redundant SAN switches (n+1) with redundant components such as redundant power supplies or by the usage of director class SAN building blocks. The major difference between a SAN switch and a director comes in the nature of serviceability and availability. The ability to perform concurrent microcode/firmware upgrades with no loss of service is a standard feature in director class products. In case of hardware errors, usually only parts of a director must be changed, or with a switch, it has to be exchanged in total in several error cases. This being said, the boundary between director and switch class products, appears to be constantly shifting.

Operating system device driver level

The device driver interacts between the storage system, the operating system of the server, and the Host Bus Adapter. The device driver is responsible for all hardware related functionality presented to the operating system and the communication to the storage system, such as multipathing in a Fibre Channel environment and path failover capabilities.

Operating system level

On the operating system level high availability can be achieved by using operating system based clustering techniques such as HACMP for AIX, STEELEYE for LINUX, and Microsoft® Windows clustering. We consider clustering techniques as building blocks for disaster protection. A cluster on its own does not represent the appropriate infrastructure for disaster protection.

Application level

Achieving redundancy at the application level is very much dependent on the type of the application that is being considered. For example in a three tier SAP environment, the application tier can be made highly available by using multiple application servers. If any application server fails, the load will be redistributed to the surviving servers and the operation continues.

Functional level

The functional level is actually the most important level within the system layer architecture model and is dependent on the following availability levels:

- IT infrastructure availability (Operating System + Server + Storage + Network)
- Application availability (Application + Data) + IT infrastructure availability
- Business Process availability (Application availability + external dependencies)

The planning for disaster protection on the functional level must include all external dependencies, such as, collaborative business scenarios between different enterprises.

4.5 Summary

With your understanding of tier levels of Business Continuity solutions and the hierarchical layers of the system architecture, you can now move to the development of a Business Continuity solution for your IT environment.

5

Business Continuity Solution Selection Methodology

This book describes a wide variety of System Storage Business Continuity technologies and solutions. Each are very powerful in their own way, and each has their own unique characteristics. How do you select the optimum combination of solutions? How do you organize and manage all these valid Business Continuity technologies? Traditionally, developing the skills to perform this selection function effectively has been time consuming and incomplete. In addition, it can be difficult to transfer these skills to other colleagues.

In this chapter, we offer a suggested *Business Continuity Solution Selection Methodology* that can provide assistance to this problem. This methodology allows you to navigate the seemingly endless permutations of Business Continuity technology quickly and efficiently and to identify initial preliminary, valid, cost-justified solutions.

This methodology is not designed to replace in-depth skills. Instead, it is meant as a guideline and a framework. Proper application of this methodology can significantly reduce the time and effort required to identify proper solutions and, therefore, accelerate the selection cycle.

5.1 The challenge in selecting Business Continuity solutions

From an IT infrastructure standpoint, there are a large variety of valid Business Continuity products. The fundamental challenge is to select the optimum 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 (Figure 5-1). Instead, Business Continuity solutions need to be viewed as a whole, integrated multi-product solution.

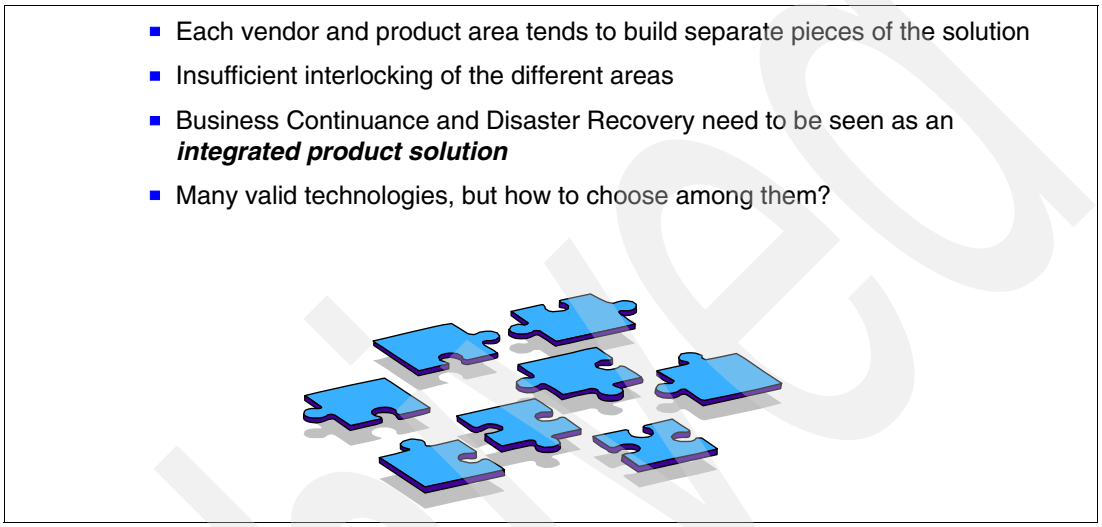

Figure 5-1 Historical challenge in selecting Business Continuity solutions

In this chapter, we propose a Business Continuity Solution Selection Methodology that you can use to sort, summarize, and organize the various business requirements in a methodical way. Then, we methodically use those business requirements to efficiently identify a proper and valid subset of Business Continuity technologies to address the requirements.

We assume that before you enter the Solution Selection Methodology that we describe in this chapter, that you have performed the requisite Risk Assessment, Business Impact Analysis, and Assessment of the current environment and have used the information to gain an initial desired Recovery Time Objective (RTO). You can find information about performing these steps in Chapter 3, “Business Continuity planning, processes, and execution” on page 43.

With the desired RTO, we discuss how to use the concepts of the Tiers of Business Continuity and Solution Segmentation to methodically identify the appropriate *candidate* Business Continuity solutions from among the continuum of today’s Business Continuity technologies.

5.1.1 The nature of Business Continuity solutions

To combine and properly select among multiple products, disciplines, and skills to effect a successful IT Business Continuity solution, you can categorize all valid Business Continuity IT technologies into five component domains:

- Servers
- Storage
- Software and automation
- Networking and physical infrastructure
- Skills and services that are required to implement and operate these components

The IT infrastructure that is necessary to support the Business Continuity solution can be inserted into one of these five components, as shown in Figure 5-2.

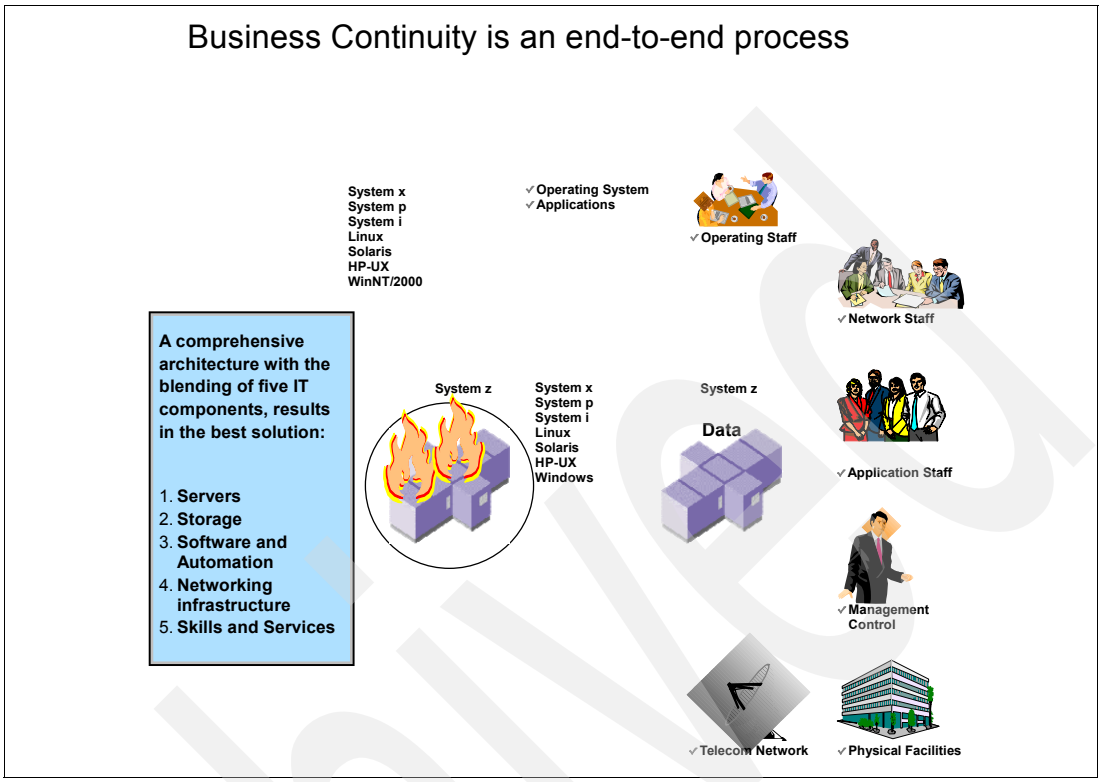

Figure 5-2 The five components of IT Business Continuity

You can use components from each of the five categories in the process of building the Business Continuity solution. Our solution selection methodology is designed to facilitate and simply the process of identifying candidate solution technologies for a given set of IT Business Continuity requirements.

5.2 The tiers of Business Continuity

We use the Business Continuity tiers as the primary tool to organize the many multiple valid products and technologies from which we can select. We described the tiers of Business Continuity in detail in Chapter 4, “Tier levels of Business Continuity solutions” on page 137.

The Business Continuity concept continues to be as valid today as when it was first created and described by the US SHARE User Group in 1988. While the technology within the tiers has obviously changed over the years, the conceptual value of organizing various recovery technologies according to their recovery speed has stood the test of time.

The Business Continuity tiers assist us in organizing the various technologies into *useful subsets* that are much easier to evaluate and manage. The concept of Business Continuity tiers is powerful and central to our selection methodology, because the tiers concept recognizes that for a given customer RTO, all Business Continuity products and technologies can be sorted into a solution subset that addresses that particular RTO range.

Thus, by categorizing Business Continuity technology choices by RTO into the various Business Continuity tiers, you can more easily match your desired RTO time with the optimum set of technologies.

Note: The reason there are multiple Business Continuity tiers is that as the RTO time decreases, the optimum Business Continuity technologies for RTO must change. For any given RTO, there are always a particular set of optimum price or performance Business Continuity technologies.

Figure 5-3 illustrates the Business Continuity tiers chart and gives a generalized view of some of today's IBM Business Continuity technologies by tier. As the recovery time becomes shorter, then more aggressive Business Continuity technologies must be applied to achieve that RTO (carrying with them their associated increase in value and capital cost).

The tiers concept is flexible. The recovery time of a given Business Continuity tier is relatively fixed. As products and functions change and improve over time, you can update the Business Continuity tiers chart by adding the new technology into the appropriate tier and RTO.

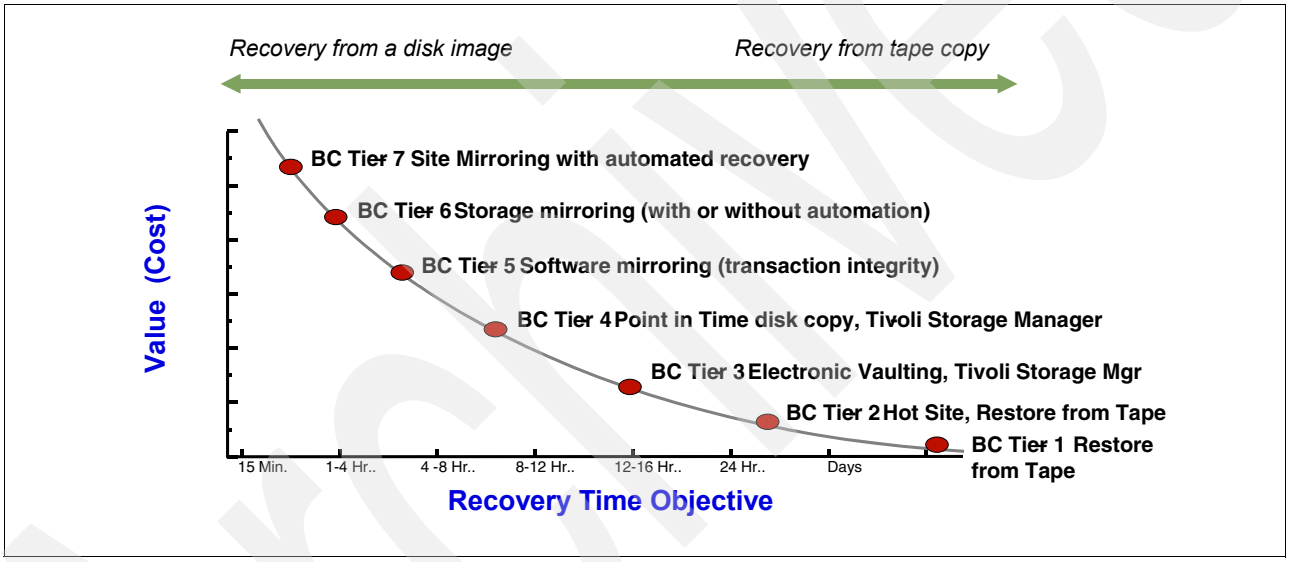

Figure 5-3 Tiers of Business Continuity

The concept and *shape* of the tiers chart continues to apply even as the scale of the application or applications changes. Large scale applications tend to move the curve to the right, while small scale applications tend to move the curve to the left. However, in both cases, the general *relationship* of the various tiers and Business Continuity technologies to each other, remains the same. Finally, although some Business Continuity technologies fit into multiple tiers, there is not one Business Continuity technology that can be optimized for all the tiers.

We recommend that your technical staff create your own internal version of the Business Continuity tiers chart that is specific to your particular environment. This version can become a powerful and useful tool for internal Business Continuity planning, communication, and decision making. Creating and gaining agreement to such a tiers chart tends to move into motion the requisite risk analysis, business impact analysis, and Business Continuity program design that is necessary to build this chart for your enterprise.

With such a chart in place, the business has a powerful tool to assign what tier or tiers and corresponding RTO an application requires and to readily determine what solutions might be viable.

5.3 Application segmentation for Business Continuity

Next, we need to do a mapping the organization's business processes and IT applications onto the Business Continuity tiers. We do this mapping to draw the appropriate linkages between the business processes), the supporting IT applications and infrastructure, and the necessary Business Continuity technologies to provide the appropriate resiliency. The best practices today in doing this business process mapping is to segment the applications and business processes into (ideally) three segments, according to the speed of recovery that is required.

Applied to the Business Continuity Tiers, this business process and application segmentation typically appears as shown in Figure 5-4.

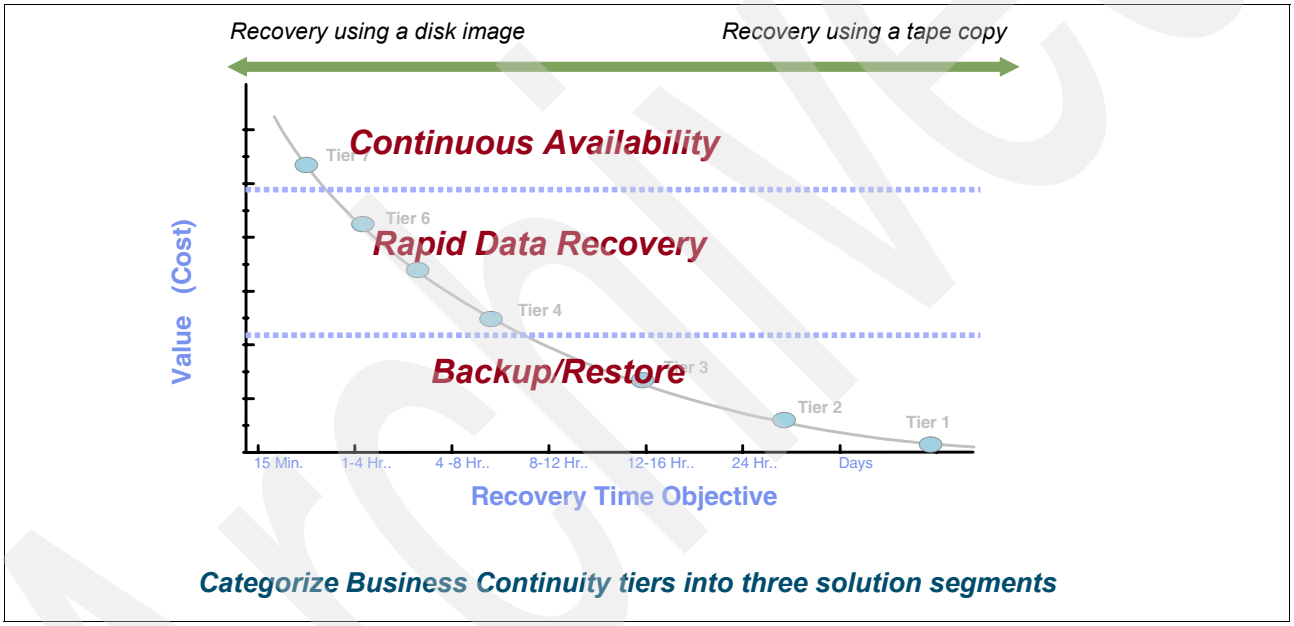

Figure 5-4 Example of a three segment IT Business Continuity architecture

Three business process or application segments 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 generally do not provide enough additional strategic benefit.

Some general guidelines and definitions for the recovery capability of the recommended three business process segments include:[1]

- **Continuous Availability**
  - 24x7 application and data availability (server, storage, and network availability)
  - Automated failover of total systems or site failover
  - Very fast and transparent recovery of servers, storage, network
  - Ultimate Disaster Recovery: Protection against site disasters, system failures
  - General RTO guideline: minutes to < 2 hours
- **Rapid Data Recovery**
  - High availability of data and storage systems (storage resiliency)
  - Automated or manual failover of storage systems
  - Fast recovery of data or storage from disasters or storage system failures
  - Disaster Recovery from replicated disk storage systems
  - General RTO guideline: 2 to 8 hours
- **Backup/Restore**
  - Backup and restore from tape or disk
  - Disaster Recovery from tape
  - RTO = 8 hours to days

5.3.1 Each segment builds upon foundation of the preceding segment

The Business Continuity functionality of each business process or application segment is built upon the technology foundation of the segment that is below it. In other words, Backup/Restore technologies are the necessary foundations for more advanced technologies that deliver Rapid Data Recovery. And Rapid Data Recovery technologies are the necessary foundations for the more advanced technologies that deliver continuous availability. So, it is really a matter of building upwards upon the foundations of the technologies of the previous segment. Best practices for Business Continuity implementation roadmaps are to create a multiple phase project in which the overall Business Continuity solution is built step-by-step upon the foundation of the previous segment's technology layer.

To match this best practice of segmentation, IBM Business Continuity solutions are mapped into these three segments, as are the various IBM System Storage Business Continuity solutions that we describe in this IBM Redbook and in *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548.

5.3.2 Application segmentation summary

To use the Business Continuity tiers and segmentation concepts to derive a blended, optimized enterprise Business Continuity architecture, we suggest the following steps:

1. Categorize the business' entire set of business processes into three segments:
   - Low Tolerance to Outage
   - Somewhat Tolerant to Outage
   - Very Tolerant to Outage

   Of course, while some business processes that are not in and of themselves critical, they do feed the critical business processes. Therefore, those applications need to be included in the higher tier.

[1] The stated RTOs in the chart's Y-axis and given in the definitions are *guidelines* for comparison only. RTO can and will vary depending on the size and scope of the solution.

2. Within each segment, there are multiple Business Continuity tiers of technology. The individual tiers represent the major Business Continuity technology *choices* for that band. It is not necessary to use all the Business Continuity tiers, and of course, it is not necessary to use all the technologies.
3. After we have segmented the business processes and 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 for that application segment.

Designing this three-segment blended Business Continuity architecture has the effect of optimizing and mapping the varying application recovery time requirements with an appropriate technology, at an optimized cost. The net resulting blended Business Continuity architecture provides the best possible application coverage for the minimum cost.

The methodology that we discuss in this chapter is designed to help you identify the appropriate Business Continuity technology solution for each business process or application *segment* that you define in your environment.

5.4 Using tiers and segmentation as a communication tool to management

Building and documenting Business Continuity tiers chart and business process segmentation for your organization is also very useful as a communication tool for discussing and evaluating Business Continuity solution recommendations to others within IT and especially with senior business management. The tiers and segmentation concept is simple enough that non-technical personnel can see the bottom line RTO end result of technical evaluations in a straightforward fashion. Senior management does not need to understand the technology that is inside the tier or segment, but they can clearly see the RTO and the associated cost versus RTO trade-off.

This ability *to communicate the bottom line* allows senior management to understand the recommendation and the trade-offs and, therefore, to make a decision quickly and efficiently. Because of the clarity of the decision alternatives, it can be more likely that management understands the choices and reaches decisions more quickly. This clarity of the choices and the associated financial cost should result in a higher likelihood of adequate funding for the Business Continuity project.

5.4.1 The use of tiers in this book

We have categorized the product information in this book by tiers. You can recognize the tier or tiers for any given technology easily. In this way, you can categorize where you can use a particular Business Continuity tool and solution in your environment. Figure 5-5 provides a partial summary list of the many IBM technologies that we categorize in this book.

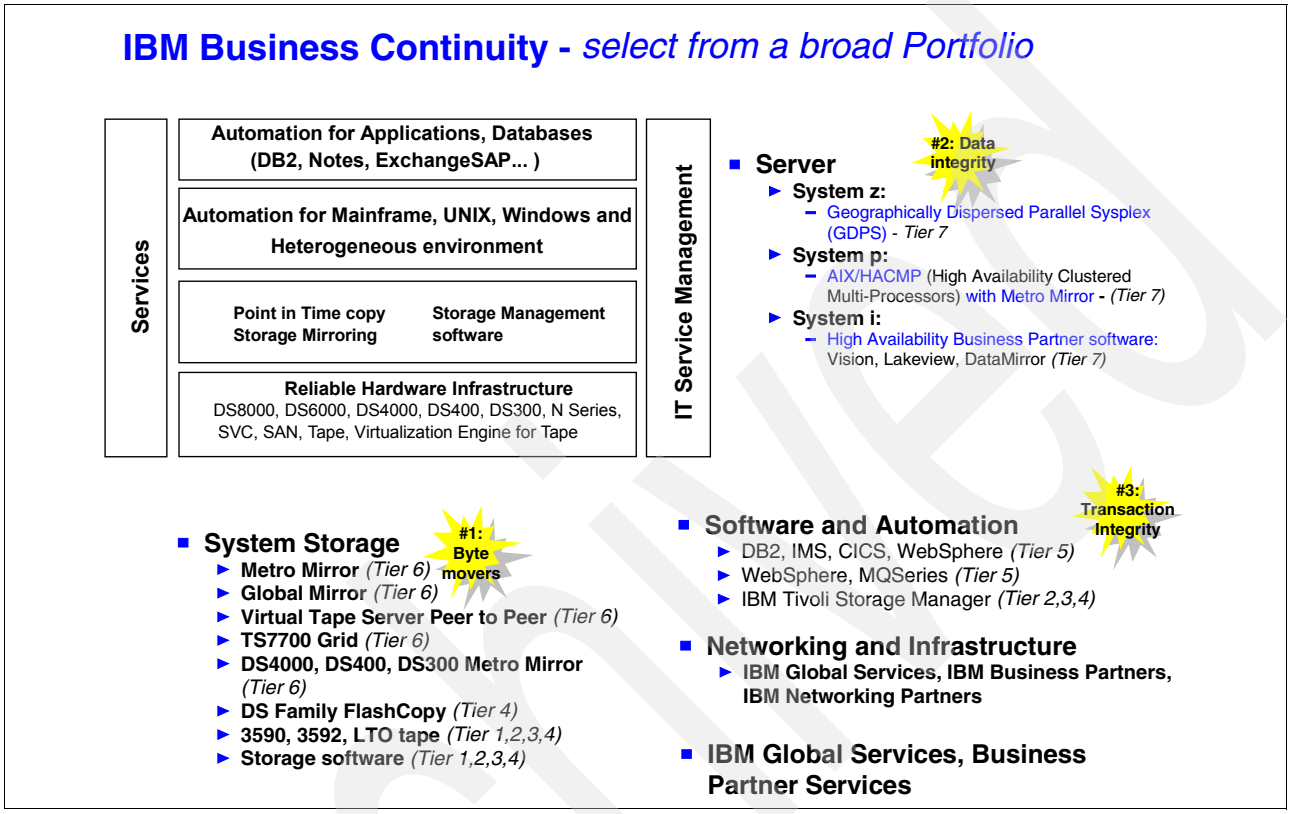

Figure 5-5 Portfolio of IBM IT Business Continuity tools

5.5 Business Continuity Solution Selection Methodology

This section describes how to use this methodology. It begins with an overview and then provides some examples of how to use the methodology.

5.5.1 Flow chart of the methodology

The Business Continuity Solution Selection Methodology is designed to provide a clear, understandable, flexible, and repeatable method to efficiently subset, organize, and select initial preliminary Business Continuity solution recommendations from the wide possible portfolio of technologies.

Figure 5-6 is a flow chart of our suggested methodology.

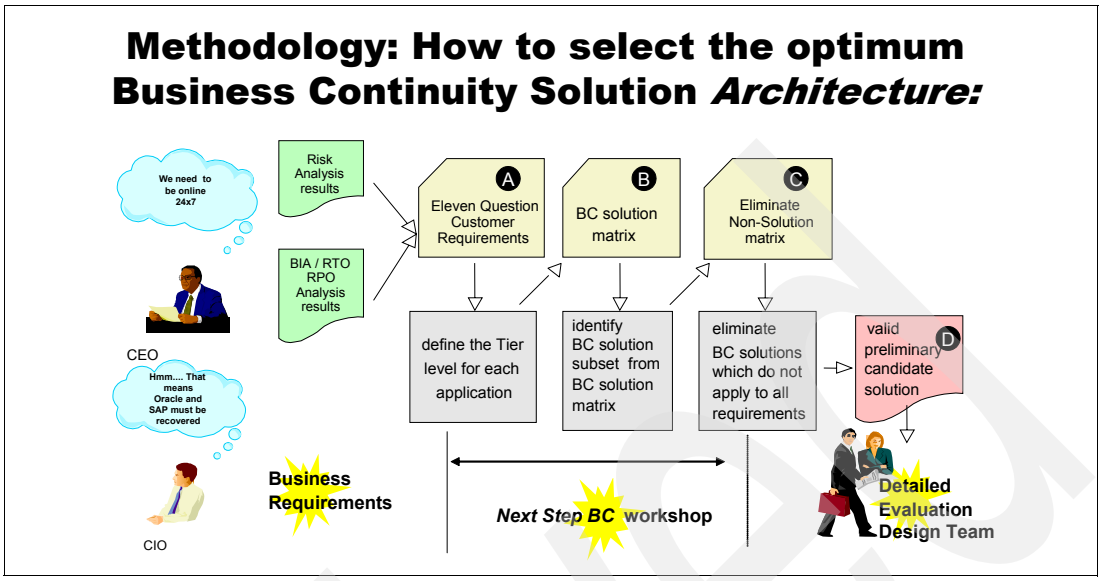

Figure 5-6 Flow chart of the Business Continuity Solution Selection Methodology

Note: The prerequisite to entering the methodology is having already performed and reached organizational agreement on the business requirements: Risk Analysis, Business Impact Analysis, application segmentation, and associated RTOs and Recovery Point Objectives (RPOs).

The Business Continuity Solution Selection Methodology is designed to:

- Provide a methodology to quickly and repeatedly identify valid initial configuration options, which is intended to accelerate determining the final solution.
- Enforce asking of correct business and IT requirements questions for a proper Business Continuity configuration.
- Provide a convergence discussion methodology for the multiple products and IT disciplines that must make up an integrated Business Continuity solution.
- Be easily extendable as products and technologies evolve.
- Capture basic expert Business Continuity intellectual capital in a teachable, repeatable way, and provide a framework to consistently propagate these basic skills to a worldwide audience, remote or local.

5.5.2 Intended usage and limitations of the methodology

The Business Continuity Solution Selection Methodology is intended to be used *after* you complete a requisite amount of Risk Assessment, Business Impact Analysis, and current environment assessment. We also assume that you have established an agreed upon RTO to be used in this methodology.

You can use this methodology at each segment level and repeat the methodology for each higher level segment. In other words, as we described in 5.3, “Application segmentation for Business Continuity” on page 155, you can use an identified RTO in the methodology to identify a strategic Business Continuity solution for Backup/Restore. Then, for the subset of Backup/Restore applications with a faster *rapid data recovery* RTO, you can then use the

methodology to identify the appropriate solutions. Finally, for the subset of Rapid Data Recovery applications that require the fastest *continuous availability* RTOs, you can then use the methodology to identify those appropriate solutions.

In all cases, we suggest that you use the methodology *early* in the selection cycle to establish a generalized vision of the requirements and to establish the kinds of solutions and Business Continuity technologies so that you can begin to investigate any particular set of Business Continuity requirements (Figure 5-7).

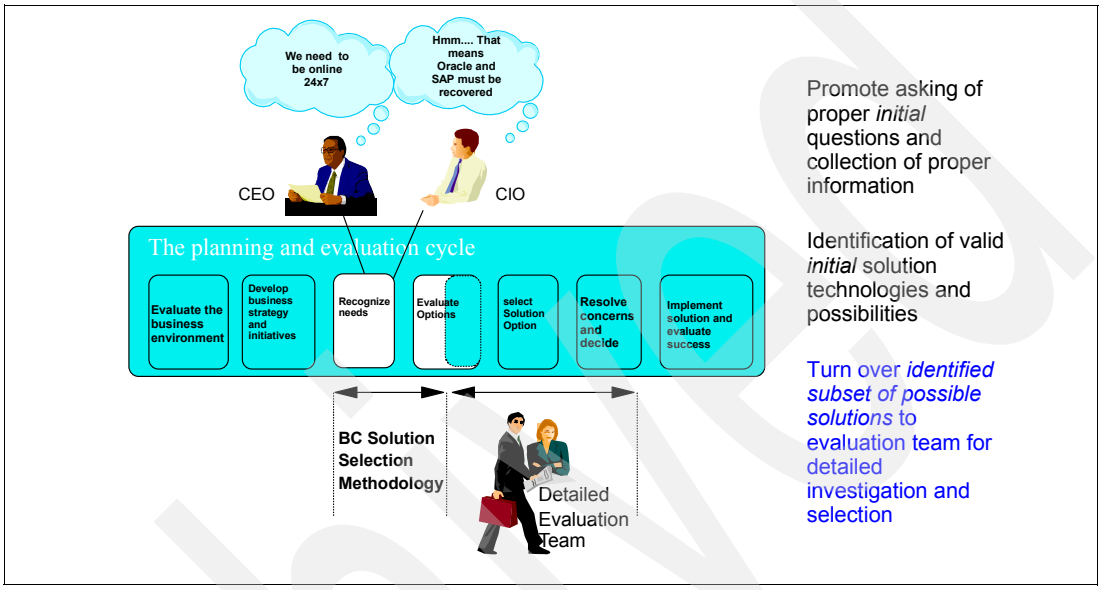

Figure 5-7 Intended use of Business Continuity Solution Selection Methodology

While the Business Continuity Solution Selection Methodology provides valuable insight into possible solutions, it:

- Cannot replace detailed solution recommendation configuration assistance.
- Cannot replace in-depth technical validation.
- Cannot replace detailed design and implementation skills and services.

Instead, the detailed evaluation team performs these functions.

The Business Continuity Solution Selection Methodology is not intended to be a perfect decision tree. Rather, it is a framework for organizing multiple Business Continuity technologies efficiently and identifying the proper possible solutions for any given customer set of requirements more quickly.

5.5.3 Principle: Asking requirements questions in a specific order

The Business Continuity Solution Selection Methodology is based on the principle of requesting and organizing necessary IT Business Continuity requirements in a *specific order*. This specific order is designed to minimize the information that we need to gather in order to arrive at a valid solution.

We therefore gather the requirements with using a specific starter set of questions. With these questions, we establish the basic IT environment, infrastructure, and desired recovery time parameters for the Business Continuity solution. Some of the questions require the business line to answer through their Risk and Business Impact Analysis. Other questions are for the Operations staff to answer from their knowledge of the IT infrastructure. (We

discussed the terms that are used in the questions in Chapter 3, “Business Continuity planning, processes, and execution” on page 43. Also, the termed are defined in Appendix B, “Terms and definitions” on page 357.)

The starter set of questions are:

1. What is/are the business processes and applications that need to be recovered?
2. On what IT platform or platforms does it run?
3. What is the desired RTO?
4. What is the distance between the recovery sites (if there is one)?
5. What is the form of connectivity or infrastructure transport that will be used to transport the data to the recovery site? How much bandwidth is that?
6. What are the specific IT hardware and software configurations that need to be recovered?
7. What is the desired level of recovery? (Planned / Unplanned / Transaction Integrity)
8. What is the RPO?
9. What is the amount of data that needs to be recovered?
10. Who will design the solution?
11. Who will implement the solution?

These are not all the possible questions, of course, but they have been carefully selected as a valid starter set of eleven questions. You can see additional questions in Appendix A, “Business Continuity Solution Selection Methodology matrixes” on page 329.

5.5.4 Tutorial: Using the Business Continuity Solution Selection Methodology

As stated previously, the methodology is used to identify candidate solutions after a requisite amount of Risk Assessment, Business Impact Analysis, and current environment assessment has been done. We discuss these steps in detail in Chapter 3, “Business Continuity planning, processes, and execution” on page 43 and illustrate an example of these steps in Chapter 6, “The Next Step Business Continuity workshop” on page 173.

From this preparatory, the desired and agreed upon RTO is the primary determining factor as to what Business Continuity candidate technologies are identified. Using these questions and the answers provided, we then use a process of elimination to identify the appropriate candidate solutions.

Figure 5-8 is an overview of this methodology.

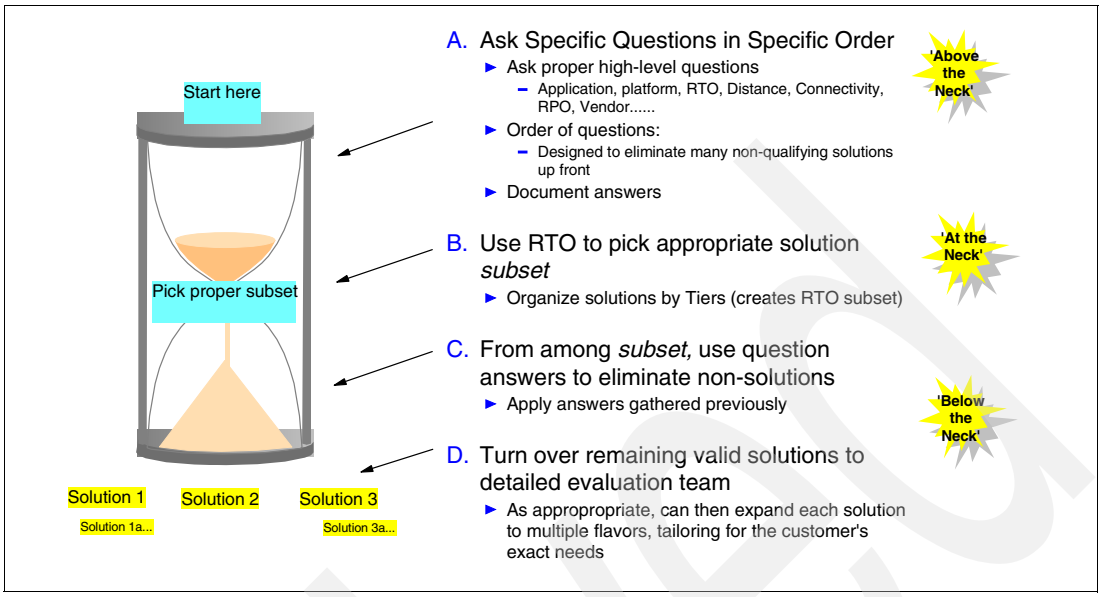

Figure 5-8 Overview of the Business Continuity Solution Selection Methodology

By segmenting the asking of questions into this hourglass concept and these three categories, it becomes possible to efficiently subset the nearly endless permutations of possible Business Continuity technology combinations and solutions into a manageable, methodical process.

Let us now step through the Business Continuity Solution Selection Methodology, step-by-step.

Step A: Ask specific questions in a specific order

In this first step, a series of Business Continuity IT requirements questions is asked in a specific order. If you have completed the Risk Assessment, Business Impact Analysis, and current environment assessment, the answers to these questions come from that work.

Remember that the order of the questions is by intention. The order of the questions is designed to eliminate non-solutions, even as you are performing the information gathering phase.

Figure 5-9 shows the questions that you need to know again.

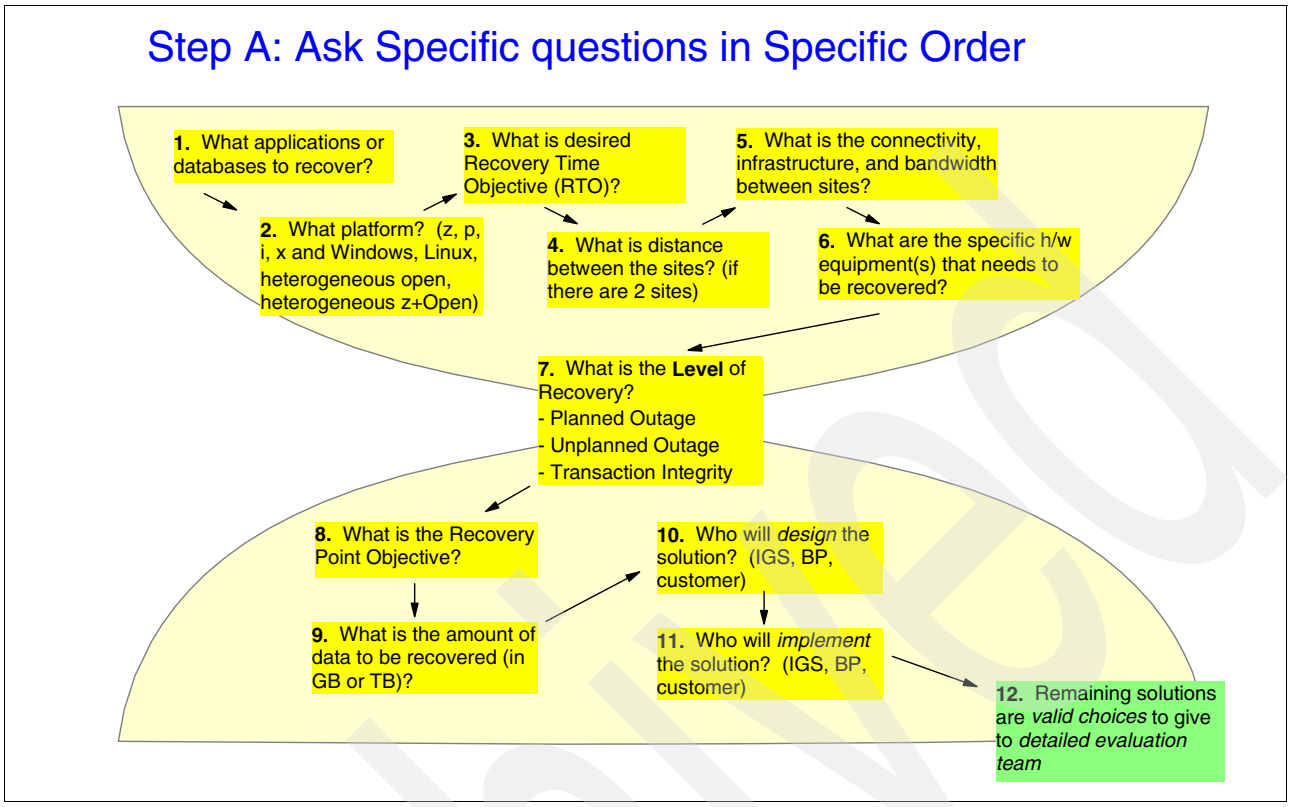

Figure 5-9 Step A: Gather answers to specific IT Business Continuity questions in specific order

The questions are designed to identify key basic business and IT requirements.

Important: It is essential that you answer these questions fully, because a lack of any of these answers means it is not possible to properly evaluate and identify the subset of solutions that you need to investigate. In this way, the methodology enforces the collection of proper business and infrastructure requirements before proceeding.

Before you continue, you must ensure that you have consensus on the answers to these questions from the enterprise's management, business lines, application staff, in addition to the IT operations staff.

Step B: Use RTO and Level of Recovery to identify candidate solutions

You now are ready to identify the preliminary candidate solutions. You use the answers from Question 3 (RTO) and Question 7 (Level of Recovery) to identify the first major subset.

Using only the Question 3 answer for RTO, you can identify an initial level of candidate solutions as we discussed in Chapter 4, "Tier levels of Business Continuity solutions" on page 137. For some circumstances, this might be sufficient, and you might have your identified candidate Business Continuity solutions at this point.

You can further refine the identified solution by using the identified RTO in combination with the answer to the Level of Recovery (Question 7). The Level of Recovery provides a major refinement clue for your candidate solutions. You define the Levels of Recovery as one of the following levels:

- **Planned outage:** The solution is required to only facilitate planned outages or data migrations. Unplanned outage recovery is not necessary.
- **Unplanned outage:** The solution is required, at the hardware and data integrity level, to facilitate unplanned outage recovery. It implies that planned outage support is also available in this solution. This level of recovery does not perform transaction integrity recovery at the application or database level.
- **Transaction integrity:** The solution is required to provide unplanned outage recovery at the application and database transaction integrity level. This level relies upon an underlying assumption that hardware level planned outage and unplanned outage support is also available.

Note: Each level builds on the previous level. In other words, a unplanned outage recovery also would have a planned outage capability. A transaction integrity recovery would also have an unplanned outage capability.

These three Levels of Recovery allow you to differentiate the fact that the Business Continuity technology and solutions that you use vary for planned outages, versus unplanned outages, versus transaction integrity.

Tip: For more information about the difference between unplanned outage and transaction integrity, see Chapter 8, “Planning for Business Continuity in a heterogeneous IT environment” on page 251.

Using the identified Level of Recovery (Question 7) in combination with the RTO (Question 3), you can identify a refined candidate set of solutions, provided in Table A-1 on page 331.)

Take the identified RTO (Question 3) answer and the Level of Recovery (Question 7) answers, and look into the Solution Matrix chart, Table A-1 on page 331.

Hint: Note that the RTO is directly mapped to a Business Continuity Tier. The Recovery Time Objective to Business Continuity Tier (Business Continuity Tier) mapping is provided as follows (note that these are general guidelines):

- Business Continuity Tier 7 RTO: Generally near continuous to 2 hours
- Business Continuity Tier 6 RTO: Generally 1 to 6 hours
- Business Continuity Tier 5 RTO: Generally 4 to 8 hours
- Business Continuity Tier 4 RTO: Generally 6-12 hours
- Business Continuity Tier 3 RTO: Generally 12-24 hours
- Business Continuity Tier 2, Business Continuity Tier1 RTO: Generally > 24 hours

In the Solution Matrix table (Table A-1 on page 331) identify the *intersection* of the RTO/Tier with the Level of Recovery. At the intersection, in the contents of the intersection cell, are the *initial candidate Business Continuity solutions for this particular RTO*, at this Level of Recovery.

An example of using this full solution matrix is shown for illustration purposes in Figure 5-10.

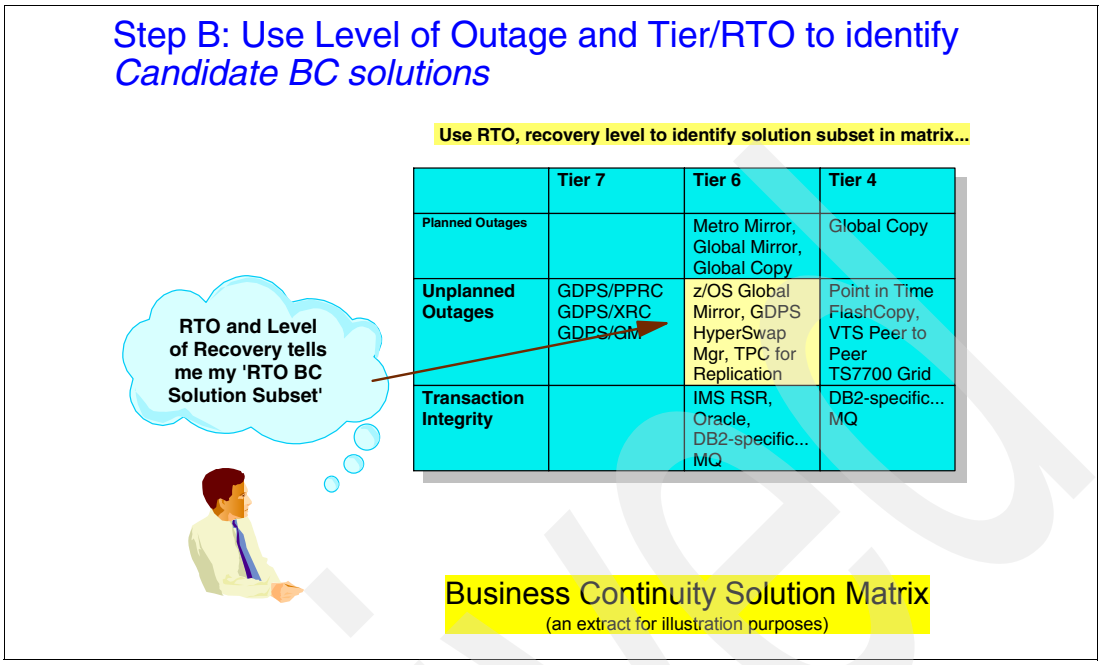

Figure 5-10 Step B: Full solution matrix

For example, in Figure 5-10, the identified preliminary candidate Business Continuity solutions are z/OS Global Mirror, GDPS HyperSwap Manager, and IBM TotalStorage Productivity Center for Replication.

Step C: Eliminate non-solutions

Now that you have identified the preliminary candidate Business Continuity solutions, you can *eliminate non-solutions* by applying the other answers gathered in Step A to the candidate solutions. For the solutions in this book, we supply a starter set of the *eliminate non-solutions* information in Appendix A, “Business Continuity Solution Selection Methodology matrixes” on page 329. As an example for illustration purposes, an extract from that table is shown in Figure 5-11.

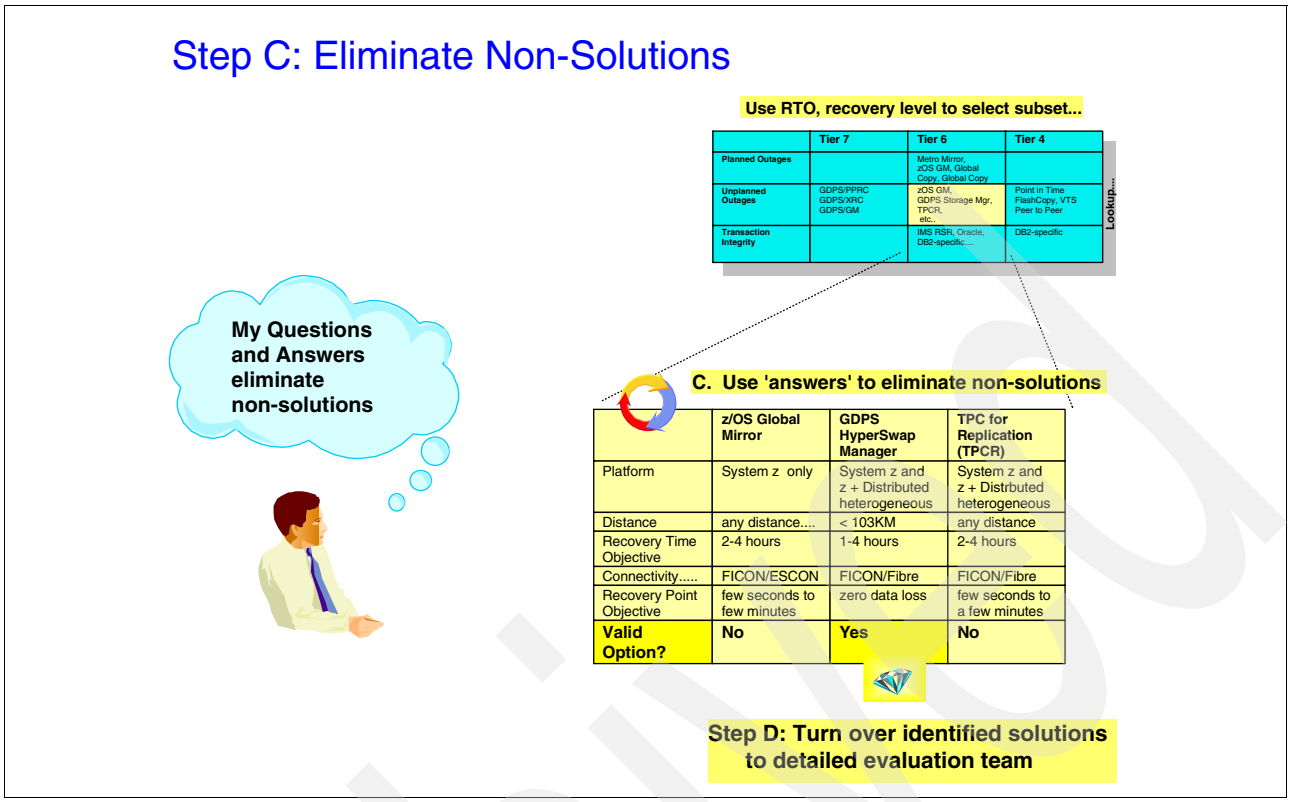

Figure 5-11 Step C: Eliminate non-solutions

By applying the answers from Step A on topics such as distance, non-support of platforms, and so on, those candidate solutions which do not apply are eliminated.

It will be normal to have multiple possible solutions after you complete Step C.

The Business Continuity candidate solution or solutions that remain after this pass through Step C are, therefore, your set of *initial candidate* Business Continuity solutions. At this point, you might need to pass these candidate technologies to the detailed evaluation team for further examination, refinement, and modification.

Step D: Turn over identified solutions to detailed evaluation team

Having identified a preliminary set of valid candidate Business Continuity solutions and technologies, you now turn over this set of candidate solutions to a skilled evaluation team that is made up of members who are qualified to contrast and compare the identified solutions in detail. The valid identified candidate solutions also dictates what mix of skills will be necessary on the evaluation team.

The evaluation team will in all likelihood need to further configure the candidate solutions into more detailed configurations to complete the evaluation. This is also normal. 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.

You should not expect this methodology to be a perfect decision tree. Rather, the intent is to provide an initial identification, in a repeatable, teachable manner, that can be performed by staff of varying skill levels, including relatively inexperienced staff.

At this point, the methodology is complete.

5.5.5 Value of the Business Continuity Solution Selection methodology

As simple as this methodology sounds, the process of quickly identifying proper candidate Business Continuity solutions for a given set of 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 can manage the preliminary evaluation phase more consistently and repeatedly and can be taught to others easily.

This methodology also supports our current best Business Continuity practices 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 re-run the methodology, and give the lower RTO Level of Recovery for those lower bands and applications, you would find the corresponding candidate solution technologies in the appropriate (lower) RTO solution subset cells.

5.5.6 Updating the methodology as technology advances

This methodology is also flexible. Because of the table-driven format, as technology changes, only the contents of the tiers chart change. The methodology itself need not change.

In particular, because Business Continuity technology is created or enhanced and results in an improvement of its tier of Business Continuity capability, this methodology simply:

- Adds the new technology to the appropriate RTO/Tier cell.
- Adds that solution as a column to the Eliminate Non-Solutions table.

In most cases the questions asked in either Step A or Step B do not need to change.

5.6 An example: Using the Business Continuity Solution Selection Methodology

To illustrate the use of the methodology in practice, here is an example.

5.6.1 Step A: Ask specific questions in a specific order

The first step in any Business Continuity solution evaluation is to gather the appropriate business and IT infrastructure requirements by working within your organization to reach agreement on the methodology starter questions.

For this example, let us suppose that the answers to the starter set of Business Continuity Solution Selection Methodology questions are as follows:

1. What is the application that needs to be recovered?

   Heterogeneous

2. What platform does it run on?

   System z

3. What is the desired RTO?

   3 hours, which is a Tier 6 solution

4. What is the distance between the recovery sites (if there is one)?

   35 KM

5. What is the form of connectivity or infrastructure transport that will be used to transport the data to the recovery site? How much bandwidth is that?

   Fibre Channel, DWDM, bandwidth = 50 MBps

6. What are the specific storage vendor hardware and software configurations that need to be recovered?

   IBM DS8000

7. What is the desired level of recovery? (Planned, Unplanned, Transaction Integrity)

   Unplanned Outage

8. What is the Recovery Point Objective?

   Near zero data loss

9. What is the amount of data that needs to be recovered?

   4 TB

10. Who will design the solution?

    to be determined

11. Who will implement the solution?

    to be determined

After you answer these questions, you can proceed to Step B.

5.6.2 Step B: Use level of outage and Tier/RTO to identify RTO solution subset

You now apply your Tier/RTO and Level of Recovery to your solution matrix. From questions 3 and 7 in 5.6.1 Step A, these were:

- Unplanned Level of Recovery
- Recovery Time Objective = 3 hours

Figure 5-12 shows a *simplified* version for illustration purposes. A full version of this table is in Appendix A, “Business Continuity Solution Selection Methodology matrixes” on page 329.

	7	6	5	4, 3	2, 1
RTO ===>	RTO Generally -near continuous to 2 hours	RTO Generally -1 to 6 hours	RTO Generally - 4 to 8 hours	RTO Generally Tier 4: 6-12 hours ; Tier 3: 12 - 24 hours	RTO Generally > 24 hours
Description:	Highly automated integrated h/w s/w failover	Storage and Server mirroring	S/W and database transaction integrity	Hot site, Disk PiT copy, Tivoli Storage Manager-DRM, fast tape	Backup software, tape
Planned Outage / Data Migrations - "Byte-Movers		Metro Mirror, Global Copy, Global Mirror, z/OS Global Mirror VTS Peer to Peer TS7700 Grid		FlashCopy, Global Copy VTS Peer to Peer, TSM, Tape	Tivoli Storage Manager, Tape
Unplanned Outage D/R, adds Data Integrity to byte-movers	GDPS/PPRC, GDPS/XRC, GDPS/GM, AIX HACMP-XD with Metro Mirror	Global Mirror, z/OS Global Mirror, GDPS HyperSwap Manager, TPC for Replication		Tier 4: VTS Peer to Peer, TS7700 Grid, FlashCopy, FlashCopy Migration Manager, Global Copy, FlashCopy, TSM, Tape	Tivoli Storage Manager, Tape
Database and Application Transaction Integrity - adds Transaction Integrity to Unplanned Outage Data Integrity	Database-level integration with clustered operating system		SAP, Oracle, DB2, SQL Server remote replication	Tier 3: MS SQL Server Database cluster with physical tape transport	

Figure 5-12 SystemStorage Business Continuity Solution matrix

By intersecting the Tier 6 RTO column with the Unplanned Outage row, you find that the preliminary candidate recommendations in you simplified table is:

- z/OS Global Mirror
- GDPS HyperSwap Manager C
- TotalStorage Productivity Center for Replication

5.6.3 Step C: Eliminate non-solutions

You now use the information that you gathered in Step A to eliminate non-solutions. Examine the Step C Eliminate Non-Solutions table for this Tier 6 Unplanned Outages, for which a starter set is supplied in Appendix A, “Business Continuity Solution Selection Methodology matrixes” on page 329. A *simplified* version of the eliminate non-solutions table, for the Tier 6 Unplanned Outage, is shown in Figure 5-13.

Solution	z/OS Global Mirror	GDPS HyperSwap Manager	TPC for Replication
Platform	System z	System z, Hetereogeneous including System z	System z, System p, System i, LINUX, Sun, HP, Windows, Heterogeneous (distributed)
Distance	Unlimited	< 100KM	< 300Km for Metro Mirror, Unlimited for Global Mirror
Connectivity	Fibre Channel, FICON	ESCON, FICON	FICON, Fibre Channel
Vendor (1)	IBM or Hitachi	PPRC-compliant storage from same vendor	DS8000, DS6000, ESS, SVC
Vendor (2)	IBM or Hitachi	PPRC-compliant storage from same vendor	IDS8000, DS6000, ESS, SVC
RPO	few seconds to few minutes	near zero	a few seconds to a few minutes
Amt of Data	any	any	any

Figure 5-13 Tier 6 Unplanned Outage eliminate non-solutions table

As you apply the different criteria sequentially from top to bottom, you find that:

1. Options in Figure 5-13, above support, IBM System z, Heterogeneous, and Distributed environments.
2. From a distance of 35 KM, all remaining solutions qualify.
3. From a connectivity standpoint of ESCON®, all remaining solutions qualify.
4. From a storage vendor hardware standpoint for site 1, all solutions qualify.
5. From a storage vendor hardware standpoint for site 2, all solutions qualify.
6. From a RPO standpoint of near zero, *only GDPS HyperSwap Manager qualifies*.

Therefore, you see that after applying the answers to the identified candidates and eliminating non-solutions, these is the valid preliminary candidate solution of GDPS HyperSwap Manager.

Note: The methodology can often result in more than one solution being possible. This is normal.

5.6.4 Step D: Turn over identified preliminary solutions to evaluation team

In the final step, you now give the solution (or solutions) that you identified to the detailed evaluation team. In all cases, whether you identify one or multiple possible solutions, the detailed evaluation team step is necessary to *validate* this preliminary set of identified solutions, as well as to accommodate a large variety of environment-specific considerations. As stated earlier, the methodology is not intended to be a perfect decision tree.

This completes the methodology example.

5.7 Summary

This methodology is meant as a framework and an organizational pattern for the efficient preliminary identification of proper Business Continuity solutions. This methodology is adaptable as the technology or environment changes, by updating the tables and questions used. It provides a consistent, teachable, repeatable method of locating the proper preliminary Business Continuity solutions.

This methodology is not meant as a substitute for Business Continuity skill and experience, nor is it possible for the methodology to be a perfect decision tree. While there clearly will be ambiguous circumstances (for which knowledgeable Business Continuity experts are required), the methodology still provides for the collection of the proper Business Continuity business requirements information.

In this way, the methodology provides an efficient process by which the initial preliminary Business Continuity solution selection can be consistently performed. In the end, this methodology can assist you in mentally organizing and using the information in this book, as well as navigating any Business Continuity technology evaluation process.

6

The Next Step Business Continuity workshop

Because of its broad nature, Business Continuity planning has often been viewed as complex and complicated, slow to get started, and less complete than it should be.

This chapter addresses these issues by providing a step-by-step, applied example of the Business Continuity planning, design, and solution methodologies that we describe in this IBM Redbook. The intent is to help you get started with Business Continuity planning for your own organization in an efficient, timely manner without requiring excessive preparation. We present the example in a workshop format, which you can use as a guideline when presenting and working within your organization or with your clients.

6.1 Objective and format of the workshop

This chapter provides a step-by-step, applied example of all of the concepts in this IBM Redbook so that you can apply those concepts quickly and efficiently within your own organization.

Business Continuity planning concepts are essentially the same for any organization. You can use this chapter as a template and foundation for performing an expanded Business Continuity planning project in a larger, more complex enterprise. The workshop format as documented here is also well suited for use in enterprises that are smaller or less complex. This chapter gives you a step-by-step cookbook to help define a valid *baseline* Business Continuity strategy, roadmap, and implementation plan for your organization. From there, you can expand and grow that strategy over time.

This chapter includes:

- Logistics, preparation, and setting proper expectations
- Overview of workshop methodology with recommended workshop steps
  - Collect information for prioritization
  - Vulnerability, risk assessment, scope
  - Define Business Continuity targets based on scope
  - Solution Option design and evaluation
  - Recommended IBM solutions and products
  - Recommended strategy and roadmap
- In Chapter 7, “Next Step Business Continuity workshop: Case Study” on page 233, we provide a complete case study that illustrates the practical application of this workshop to a real life client.

Note: This workshop and its steps and framework are just recommendations. Every circumstance is going to be a bit different. Feel free to add, delete, or rearrange the framework to suit your needs.

This chapter is a simplified version of the ideal Business Continuity planning process that we describe in Chapter 3, “Business Continuity planning, processes, and execution” on page 43. We cross-reference each of our steps to the relevant detailed documentation within that and other chapters in this book.

6.2 Workshop logistics and preparation

Here are some sample logistics that are necessary to prepare for conducting this workshop. If a more detailed version of this workshop is required, you can expand the time frames for these preparation milestones appropriately.

6.2.1 Workshop expectations, scope, and desired outputs

In order to get a good start on Business Continuity planning, we recommend that you not attempt to solve everything at the same time. Rather, our recommendation is that you narrow the scope of this workshop to just a few key business processes (one, two, or at most three). You can then focus on that achieveable scope, and establish a *baseline* Business Continuity strategy with skills and a methodology upon which you can then expand. Most importantly,

this approach can *teach* the requisite skills that are needed to repeat the methodology with an expanded scope.

To help reach this goal, we will:

- Provide fill-in-the-blank templates for important information, so that it is clear what information is to be gathered.
- Provide valid assumptions so that we can move forward with the planning process even if key information is not available.
- Provide templates for defining what should be done in achievable chunks, that is a phased approach over time. This phased approach is essential to accommodate realistic constraints on resource availability, skills, and manpower.

You should expect the preparation for this workshop to take about a week, on average, to discuss, identify, and collect the requested input information. The actual workshop itself is designed to last from one to two days.

The intended output is a presentation deck that documents the findings and requirements that came out of the workshop, restating necessary principles, and describing a derived baseline Business Continuity architecture. Associated products and solutions that fit the architecture are documented, as well as a roadmap and general implementation plan.

6.2.2 Desired participants

Obtaining the proper IT management representation at this workshop is essential.

Key client participants

Key IT decision makers and IT technical experts need to participate in this workshop to provide input on decisions or credible assumptions on which to base the recommendations. Depending on the scope of the workshop, the recommended IT client attendees are:

- CIO
- Data center manager
- Application development manager
- Database administrator
- Business development manager (for new offerings)

It is recommended to include Line of Business attendees as appropriate.

IBM or IBM Business Partner

Accompanying the client to this workshop should be the appropriate IBM or IBM Business Partner representatives from the account team, including as needed: database specialist, server and storage specialists, networking specialist, and IT architects.

6.2.3 Sample workshop objectives

Here are a set of starter objectives for this workshop, which should be customized to your specific needs:

- To have a high level understanding of the technical and financial implications of IT Business Continuity solutions, enabled upon a consolidated base of applications, databases, servers and storage in the data centers and remote locations.
- To understand the different technologies in developing cost-effective IT Business Continuity solutions, architecture and roadmap-server clustering, data replication with

consistency, infrastructure simplification and consolidation, virtualization, automation software, and their benefits specific to the enterprise.

- To give a good idea of the different options for IT Business Continuity for the key applications, databases, IT infrastructure.
- To provide a baseline strategy for achieving a reliable, repeatable, scalable, testable IT Business Continuity solution that will allow the business to manage the IT assets more effectively and efficiently, improve service levels, reduce operational costs, and be able to adapt and respond to risks and opportunities, in order to maintain continuous business operations, be a more trusted partner, and enable growth.

6.2.4 Sample workshop agenda for a one day, 4- to 6-hour workshop

Here is a starter agenda for a one day Next Step Business Continuity workshop:

1. Workshop introduction: participants, objectives, agenda

   Introduce the participants, confirm objectives, and agenda.

2. Infrastructure review

   Review the existing configuration and previously gathered information about key applications and databases. Information gathered includes: high level data center configurations, server and storage inventory, network, Service Level Agreements (availability), performance, data recovery objectives, capacity, current pains, strategy and growth scenarios, and more.

3. Overview of Business Continuity planning methods and technology

   Review key Business Continuity concepts to set the stage for the working session.

4. Working session for Business Continuity

   A facilitated, efficient working session to derive a scalable, reliable, and repeatable Business Continuity architecture and solution set.

   Our workshop methodology is a distilled version of Chapter 3, “Business Continuity planning, processes, and execution” on page 43. We define and narrow our scope to be high-level versions of:

   – Risk Assessment
   – Business Impact Analysis
   – Business Continuity program assessment, and
   – Business Continuity IT Strategy Design

   Included in this working session is the various information that will be *refined* during the course of the workshop: business objectives, risks, business impacts of risks, with associated discussions of what workloads and characteristics need to be recovered.

   We derive through the working session, a high-level Business Continuity architecture, identify candidate technologies and solutions, address financial, technical, and organizational issues, and document projected benefits, challenges, and a roadmap.

5. Recommendations, discussion, and action items

   We wrap up with discussions of Next Step recommendations and action items.

6.2.5 Preparing for the workshop (1 to 2 weeks prior)

One to two weeks prior to the workshop, you perform the following tasks:

- Confirm executive commitment and attendance.
- Have pre-workshop preparation conference calls to confirm workshop objectives, questions, agenda, outputs, date, location and participants, information needed to send to the workshop architect.
- Provide templates and discuss information to be gathered.

6.2.6 Preparing for the workshop (5 days prior)

Five days prior to the workshop, you need to confirm final collection of the required information.

As part of our workshop methodology, we recognize that it is rare that all relevant information is perfectly available. Therefore, we provide valid assumptions to help us to keep moving forward. Of course, the more accurate the information that is available, the more specific and relevant the recommendations and solution will be, and the more productive the workshop will be.

Information to be gathered

We provide templates for the information to be gathered. Reference the following figures to see these templates:

- Figure 6-13, “Step 1 - resource and business impact” on page 189
- Figure 6-17, “Step 2 - define vulnerability, risks in phases” on page 193
- Figure 6-21, “Step 3 - define decision criteria for the Business Continuity solution” on page 195
- Figure 6-57, “Step 4 - concept of building the solution step-by-step” on page 228

Tip: The questions that we include here are an extended list of the basic IT Business Continuity requirements questions that we detail in Chapter 5, “Business Continuity Solution Selection Methodology” on page 151.

The types of information that you need to gather include:

1. Business: How much data loss and what data loss will start to hurt bottom line? For what period of time?
2. Data Center configuration diagrams: servers, storage, networking and others.
3. Key application workload and databases description (applications, database, operating system, SLAs). Client production, internal production, development, test and pre-production, operations, new workloads. Which workload should be analyzed first?
4. Server, storage, application, database, network inventory-server type (model, CPU, RAM, quantity, average and peak utilization), storage type (model, raw, usable, allocation, file type, backup criteria, and so forth), growth rates of each.
5. Application and server/storage segmentation, priorities, desired recovery times of the above inventory.
6. What is considered the Continuous Availability segment and what is its Recovery Time Objective (RTO)?

7. What is the middle-level criticality segmentation (Rapid Data Recovery) and what is its RTO?
8. What is the application segmentation that only requires a basic level of recovery (Backup/Restore) and what is its RTO?
9. Server, storage, database, software, network costs—depreciation or amortization, maintenance, hosting and support (monthly or yearly), install or implementation (one time).
10. Software licensing charges (operating system, applications, databases)—per CPU, one time or yearly
11. Environment costs—hosting, space, power, cooling, fire protection, and so forth.
12. Operation costs—number of infrastructure personnel (server administrators), annual costs per full time equivalent, hosting.
13. Estimated business outage cost (an hour or a day) and financial acceptable criteria (such as 18 months return on investment, $*X* savings in *Y* months).
14. Management tools (application, database, server, storage, and network)
15. Current data center challenges, issues, next two-year business, outlook, applications and technical strategy.
16. List any objectives, issues and questions to be addressed during workshop.
17. Any other relevant information?

6.3 Next Step workshop methodology overview

In this section we review the methodology that we use in this Business Continuity planning workshop. We often find that one of the more difficult tasks in Business Continuity planning is being able to present and communicate the concepts, ideas, and logic flow, to the wide audience that must understand, participate, and reach consensus on the Business Continuity plan.

Therefore, we overview the methodology by providing a step-by-step set of visual aids, which are designed to give you a straightforward set of guidelines for successfully performing a Business Continuity planning session.

Our methodology overview provides:

- A sequential set of actions
- Each action in the methodology has an accompanying visual aid
- Each visual aid has as appropriate: supporting discussion points, questions to be asked and answered, necessary inputs and desired outputs, and documentation of relevant assumptions
- Each visual aid has cross-references to more information within this Redbook

Let us begin with Figure 6-1.

Figure 6-1 Welcome to the Next Step Business Continuity workshop

Welcome to the Next Step Business Continuity workshop. We provide a distilled, simplified step-by-step application of the concepts in this book.

Figure 6-2 shows the agenda.

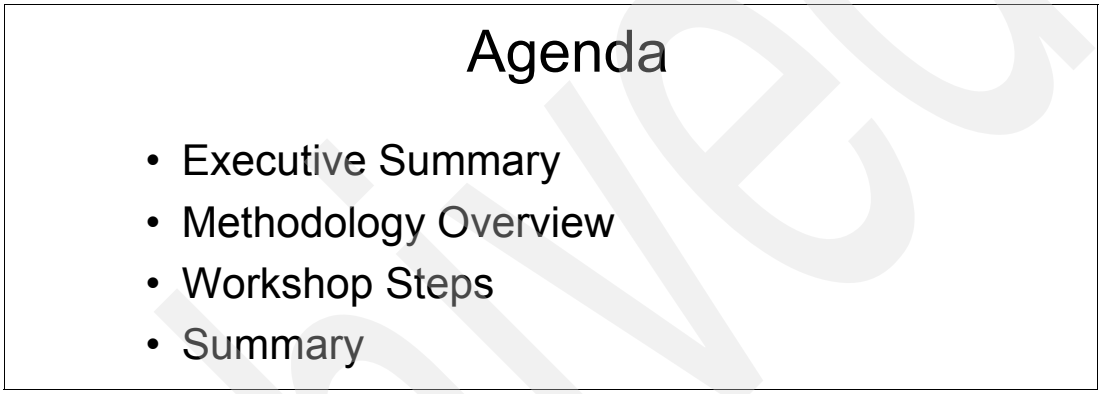

Figure 6-2 Agenda

6.3.1 Executive summary

In today's workshop, we discuss many topics, so we begin with a summary of the day's events.

We assume that the need for better Business Continuity is clear to all attendees, so we will not review those needs unless necessary. If it is necessary to discuss the value of Business Continuity, we provide discussion material for that in 6.5, "Appendix: Why Business Continuity" on page 223.

We will follow this outline for today's workshop:

1. We start with a Client Summary discussion—a key conversation by the sponsoring IT executives to the room, to set the proper expectations as to what we intend to do.
2. Then, the workshop facilitators provide an overview of the workshop's methodology. We discuss how the methodology leads us through determining key IT Business Continuity metrics, key IT *value* metrics, and Key Performance Indicators (KPI) for Business Continuity - these all will be discussed with both IT and with the Lines of Business.
3. These Business Continuity metrics will determine Business Continuity goals and drive design, budgets, return on investments and success criteria for the workshop. The desired service levels can include end-to-end systems availability, response time, disaster recovery objectives.

4. We will suggest a set of KPIs, which ideally need to be available and tracked on a monthly or quarterly basis. Examples of KPIs could include:
   - *Backup window*: What is the amount of data being backed up? Is our backup window staying steady or shrinking, and if so, how fast?
   - *Preparedness*: What is our current time to switch to the remote data center? Is it improving or getting worse
   - *Testing*: Do we have sufficiently frequent, affordable testing that can tell us the answer to the above two KPIs? What is the frequency of our testing? What is the success rate?
5. The methodology steps also guide us through assessing present attainment of these service levels and KPIs, to establish a base line for comparison and understanding of the challenges to meeting the target:
6. We assess the obstacles to meeting the targeted service levels and KPIs. Typical obstacles include lack of sufficient tools to monitor service levels and metrics, difficulty of diagnosing the source of problems, lack of expertise, and staff resources
7. The methodology will help us develop a baseline Business Continuity architecture that can best address the obstacles, specifying technologies and solutions which are becoming more affordable and feasible to implement every year
8. We will develop a project plan to pilot and implement the Business Continuity architecture and solutions, including monitoring, tracking progress and reporting to IT management and lines of business, to clearly demonstrate ongoing, quantifiable value to the business

6.3.2 Intended audience and scope for the workshop

Let us next review the methodology that we will use today. Figure 6-3 shows the intended audience and scope for this workshop.

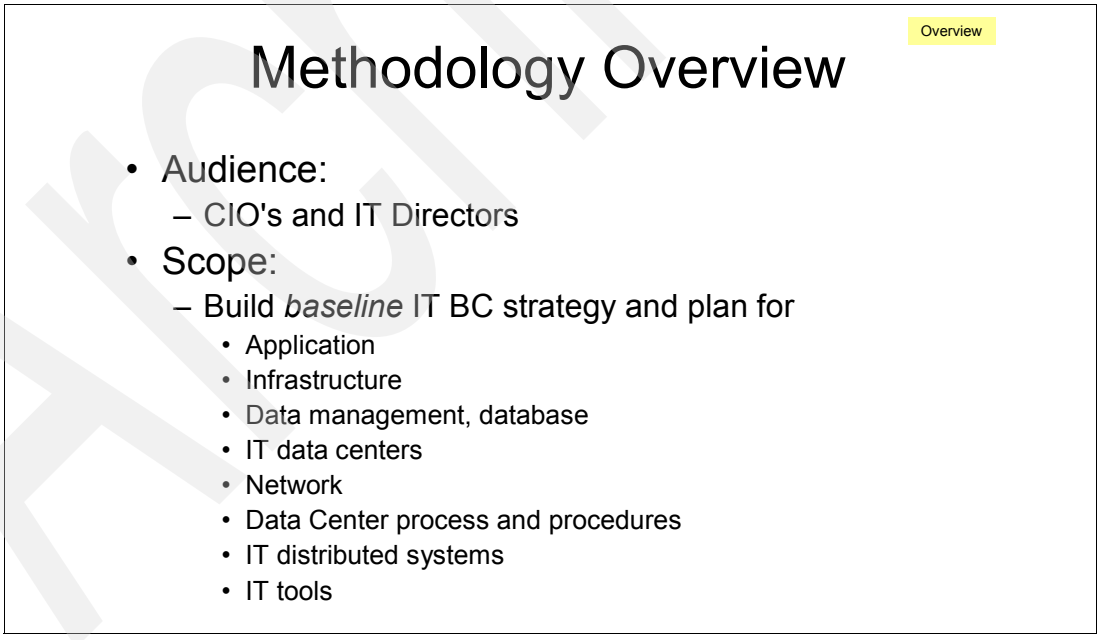

Figure 6-3 Workshop methodology overview - audience and scope

This workshop is aimed at the IT level. While true Business Continuity planning must also involve all other aspects of the business, this workshop is designed for the IT management and senior IT staff to understand and synthesize their roles and strategies as they pertain to the greater whole.

Note: The scope of this Business Continuity exercise is at the *business process* level. In today's business world, it is clear that it is the service to the lines of business, and the value that is offered to internal and external users as a unified whole, that is the primary value add of IT. Therefore, the scope of this workshop is on the necessary IT components needed to support the *entire* business process.

We provide templates and expertise to document and diagnose each of these factors listed under scope. You can also see why we set the initial expectation that we do not attempt to solve it all at the same time, and why we narrow the scope by prioritization.

Next, we discuss the objectives of our workshop methodology, shown in Figure 6-4.

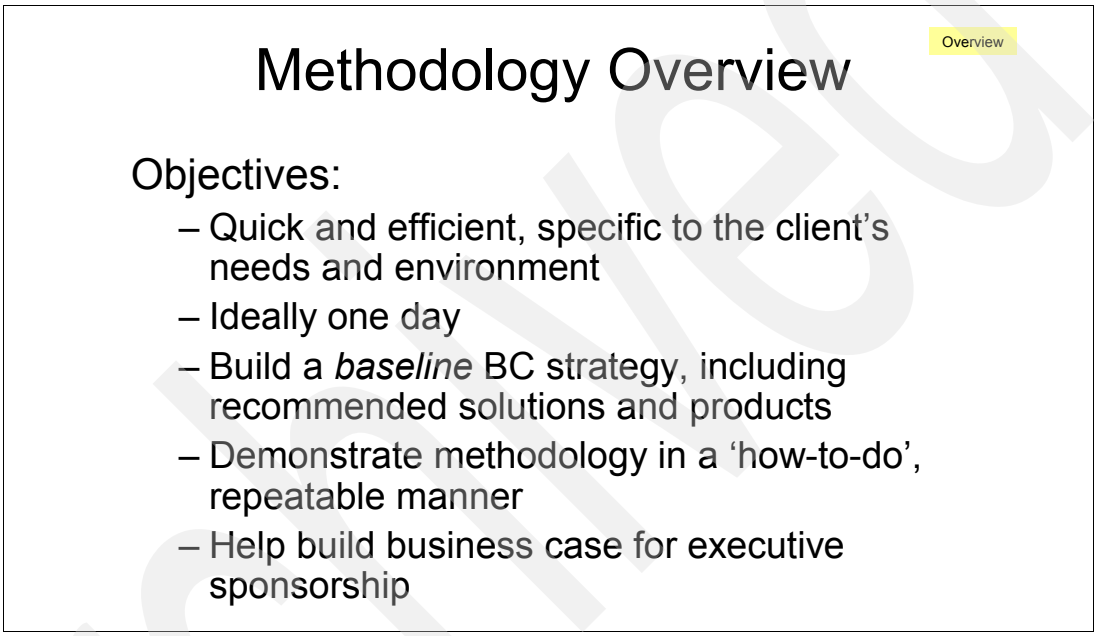

Figure 6-4 Workshop methodology overview

In summary, we expect that the ideal time span of this workshop is one day in many cases, perhaps two days at most in others. We provide templates to simplify the gathering of the relevant information. Where information is not available, we provide valid assumptions so that we can move forward with the planning process (Figure 6-5).

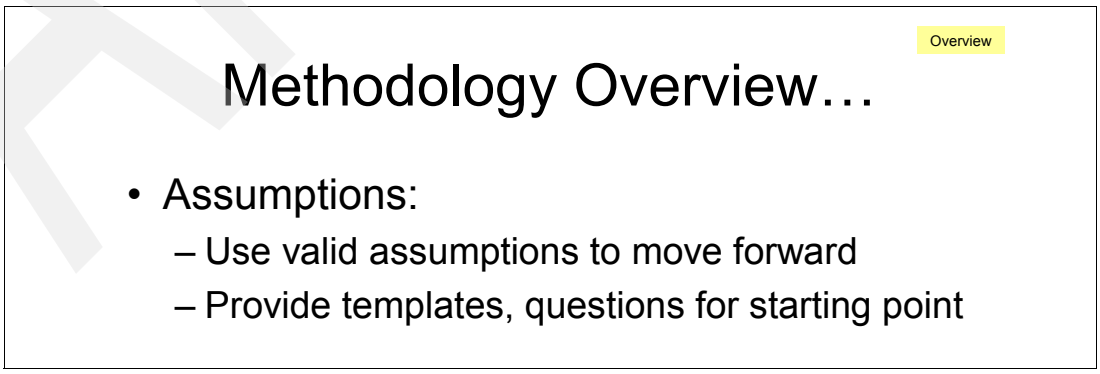

Figure 6-5 Workshop methodology overview - assumptions

Next, we discuss the intended outputs of our workshop and methodology, as shown in Figure 6-6.

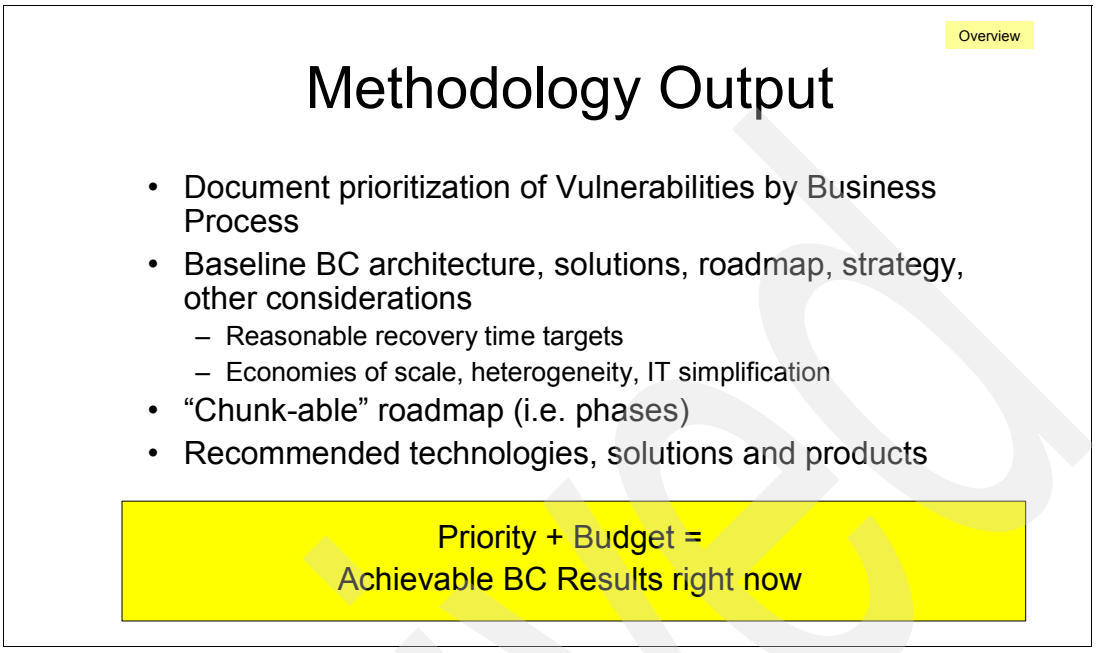

Figure 6-6 Workshop methodology overview - intended ouputs

To deliver the objectives stated previously, we include the following in our output:

- Regarding **priority of business processes and recovery**, an important documentation of the *Business Vulnerabilities* that we include in our scope. This is a distillation of aspects of the *Business Prioritization* portion of the ideal Business Continuity process that is described in detail in (3.4, "Business prioritization" on page 47).
- Regarding phases of a Business Continuity Solution, a baseline Business Continuity architecture, recommended solutions, roadmap, strategy, including:
  - Reasonable recovery time targets that can be achieved
  - Economies of scale, heterogeneity, IT simplification
  - Phased roadmap

    This is a further a distillation of the entire *Integration into IT* portion of the ideal Business Continuity process that we describe in detail in 3.5, "Integration into IT" on page 67.
- A specific set of recommendation on how to show benefits to the Line of Business. It is key that we can demonstrate how identified Business Continuity priorities can fit within existing budgets. This is a distillation and expansion of the entire *Manage* portion of the ideal Business Continuity process that we describe in detail in 3.6, "Manage" on page 125. This is an essential task, because we know that a truly effective, cost-justifiable Business Continuity strategy must be able to show continuing value in order to have joint ownership by both IT and the Line of Business.

In the next section are the actual workshop steps, in which we provide templates, guidelines, and discussion points. We provide cross-references for further information and reading within this book.

6.3.3 Performing the workshop

We begin by showing an overview of the workshop methodology steps. Figure 6-7 shows the key steps in the workshop.

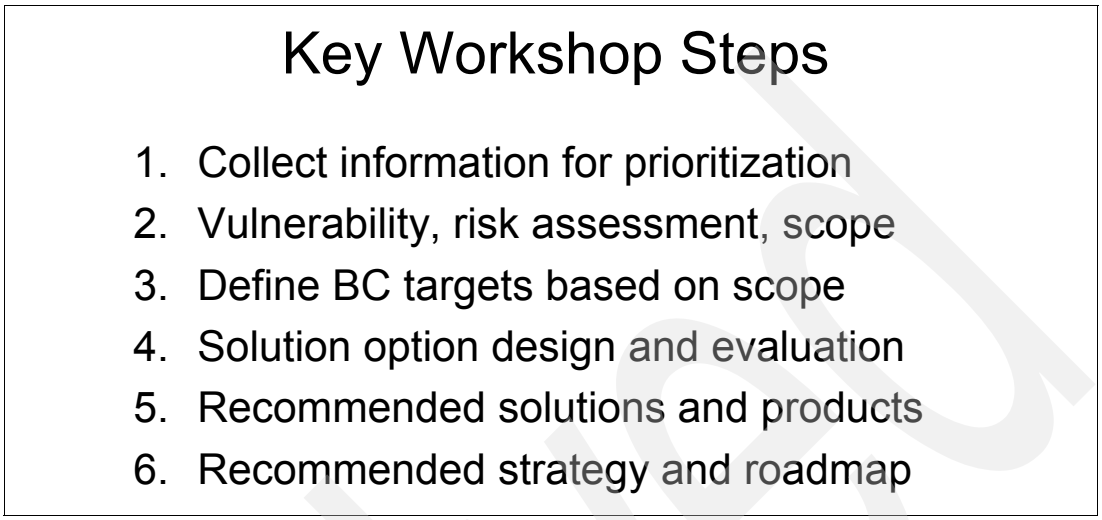

Figure 6-7 The key workshop steps

We document how each of our methodology steps map to the ideal Business Continuity planning process that we discuss in 3.3, “Ideal Business Continuity planning process” on page 46.

The first three steps are a distillation of the *Business Prioritization* portion of the ideal Business Continuity process, specifically *Risk Assessment*, *Business Impact Analysis*, and current *Program Assessment*:

1. Collect information for prioritization.
2. Vulnerability, risk assessment, scope.
3. Define Business Continuity targets based on scope.

See 3.4, “Business prioritization” on page 47 for detailed information about these steps.

The next three steps are a distillation of the *Integration into IT* portion of the Ideal Business Continuity process, specifically *Business Continuity Program Design*, *IT Strategy Design*, and *Implementation*:

4. Solution option design and evaluation.
5. Recommended IBM solutions and products.
6. Recommended strategy and roadmap.

See 3.5, “Integration into IT” on page 67 for detailed information about these steps.

Figure 6-8 is a diagram of the inputs, actions, and outputs of the key steps of the workshop.

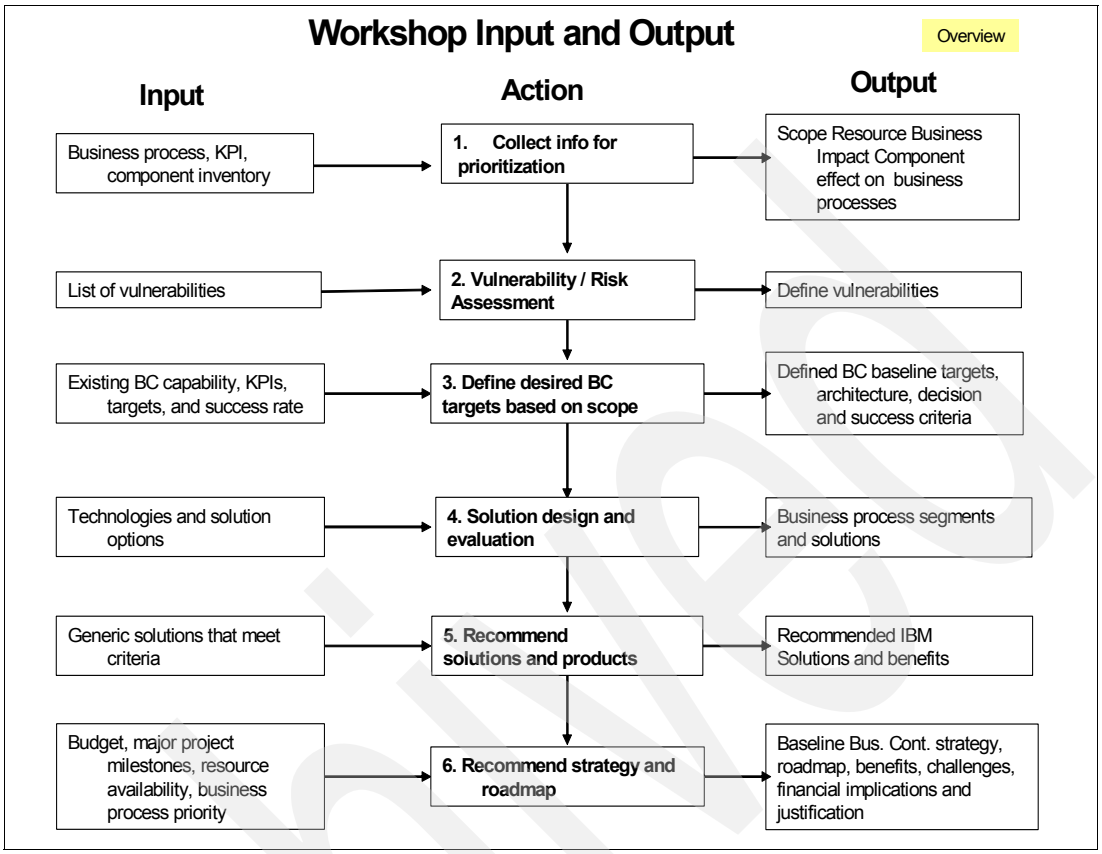

Figure 6-8 Workshop inputs, actions, and outputs

We recommend that you print a large version of this chart and place it on the wall to guide the workshop participants as to where they are in the process. Note also that each of these input and output documents are valuable communication tools for the enterprise.

Each document should be considered to be a living document, that is continually maintained, expanded, and made available to all interested parties in the enterprise. These documents, and the input/action/output flow that connects them, provides a actionable, achieveable baseline action plan for today's workshop, and for the organization over time as well.

The information in this book and in *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548 and *IBM System Storage Business Continuity Solutions Overview*, SG24-6684 are the reference material and foundations for this methodology.

Key reference chapters for each step are as follows:

1. Collect information for prioritization

   3.4.1, "Risk assessment" on page 47
2. Vulnerability, risk assessment, scope

   3.4.2, "Business impact analysis" on page 51
3. Define Business Continuity targets based on scope

   3.4.3, "Program assessment" on page 54

4. Solution Option Design and evaluation
   - 3.5.1, “Business Continuity program design” on page 67
   - 3.5.2, “IT strategy design” on page 89
5. Recommended IBM Solutions and Products
   - Chapter 4, “Tier levels of Business Continuity solutions” on page 137
   - Chapter 5, “Business Continuity Solution Selection Methodology” on page 151
   - *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548
   - *IBM System Storage Business Continuity Solutions Overview*, SG24-6684
6. Recommended strategy and roadmap

   In addition to the suggestions above for point 5, also see Chapter 2, “Industry Business Continuity trends and directions” on page 9.

Let us now begin with step 1.

6.3.4 Workshop Step 1 - Collect information for prioritization

Step 1 in the workshop is *collect information for prioritization* (Figure 6-9).

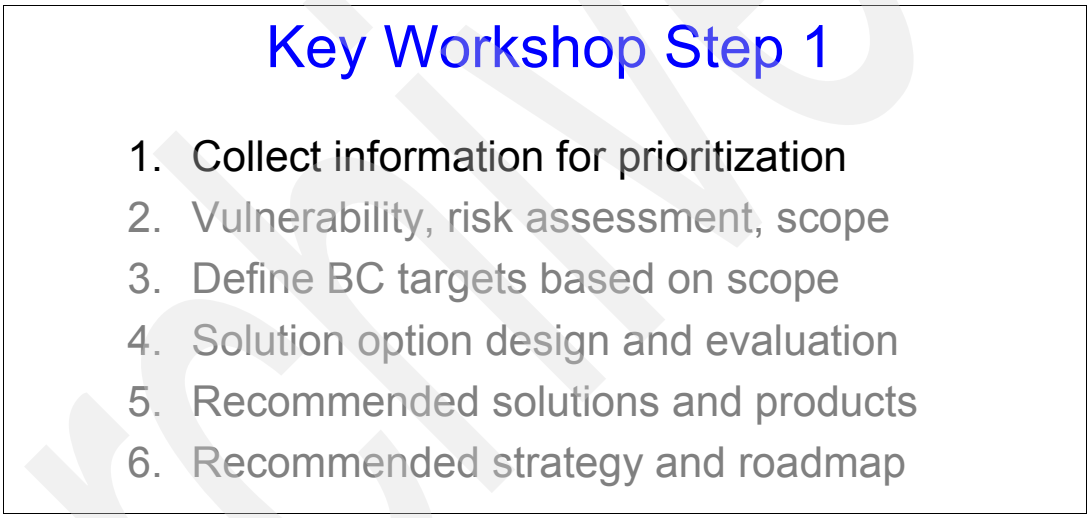

Figure 6-9 Workshop step 1 - collect information for prioritization

We begin by discussing key business processes (Figure 6-10). We suggest:

- Select a few key business processes to discuss and work on just those for this workshop
- Ideally three or less

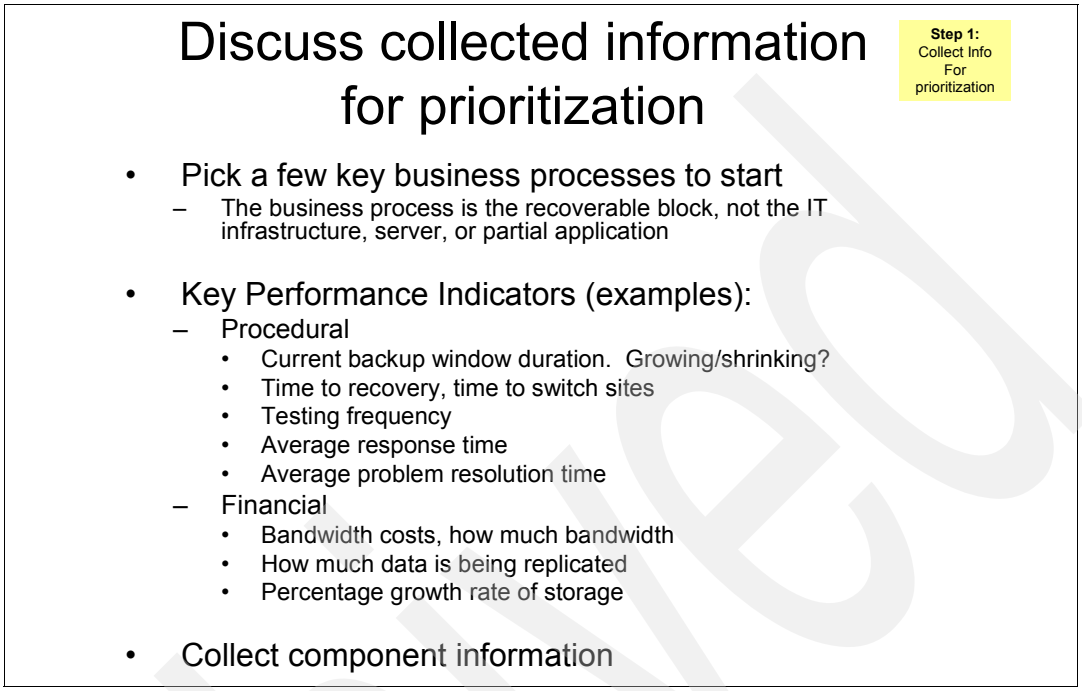

Figure 6-10 Step 1 - guidelines for 'collecting information for prioritization'

The value of narrowing the scope for this workshop to a few (preferably one, two, or at most three) key business processes, is that we can contain the nature of the workshop to an achievable level. After the client has become fluent in running our methodology with one business process, they can then repeat the methodology on other processes over time.

Note that we focus on the business process as the *recoverable unit*. We observe that today, the traditional approach of designing recovery at the individual server platform level or storage level is often insufficient. In today's business environment, the loss of just one key component will probably impact multiple steps of multiple business processes.

Tip: More information and a diagram on why the business process must be the focus, not just the technology, is in 2.1.1, "Shift of focus from technology to business processes" on page 10.

Next, we review the Business Continuity KPIs for these business processes.

Key Performance Indicators

To review briefly, KPIs are a business management concept that is particularly applicable to a Business Continuity project. We suggest some starter IT Business Continuity KPIs here; you can choose and modify these as you see fit.

KPIs reflect your client's IT Business Continuity requirements success factors. The Business Continuity solution KPIs must accurately portray the organization's Business Continuity objectives, they must accurately measure the keys to success, and they must be measurable and quantifiable.

When defining or suggesting Business Continuity solution Key Performance Indicators, consider the following:

- Should be considerations for the long-term
- Definition of what they are and how they are measured should not change often
- Must be accurately defined and measured
- Must have quantifiable targets and timeframes

Starter set of KPIs for IT Business Continuity

Figure 6-11 shows a template for metrics of a starter set of Business Continuity KPIs.

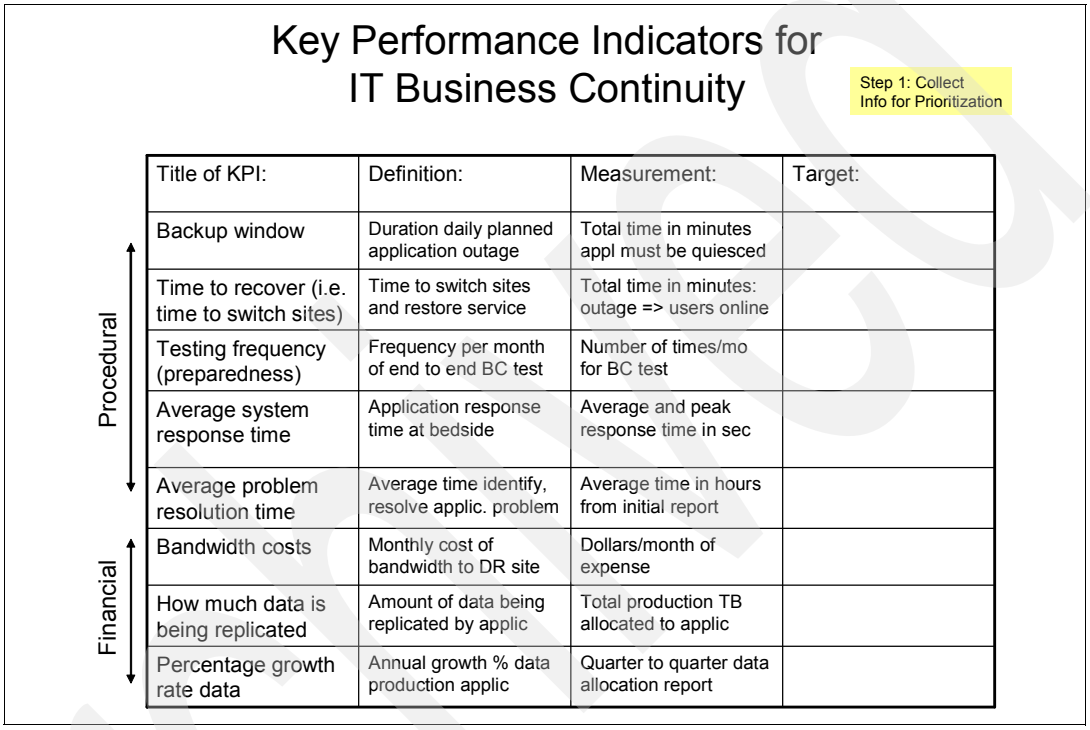

Title of KPI:	Definition:	Measurement:	Target:
Backup window	Duration daily planned application outage	Total time in minutes appl must be quiesced	
Time to recover (i.e. time to switch sites)	Time to switch sites and restore service	Total time in minutes: outage => users online	
Testing frequency (preparedness)	Frequency per month of end to end BC test	Number of times/mo for BC test	
Average system response time	Application response time at bedside	Average and peak response time in sec	
Average problem resolution time	Average time identify, resolve applic. problem	Average time in hours from initial report	
Bandwidth costs	Monthly cost of bandwidth to DR site	Dollars/month of expense	
How much data is being replicated	Amount of data being replicated by applic	Total production TB allocated to applic	
Percentage growth rate data	Annual growth % data production applic	Quarter to quarter data allocation report	

Figure 6-11 Step 1 - document current Key Performance Indicators for Business Continuity

The use of KPIs provides a key management tool. They give everyone in your organization a clear vision of what is important, and of what they will need to make happen.

We suggest that you publicize and post the KPIs in places throughout your organization where all employees can have access to them: in the lunch room, on the walls of conference rooms, on your company's intranet, and even on the company's Web site for some of them. Show the target for each KPI, and show the progress toward that target. With good project management, your staff will be motivated and feel pride in reaching the KPI targets.

Tip: Additional information about KPIs is in "Key Performance Indicators" on page 55.

Next, we review collecting information about the IT environment and components that underpin the selected business process.

Component inventory information to be collected

During the logistics and preparation for this workshop, provide fill-in-the-blank templates to the client for organizing the component information to be gathered. These are the IT components that underpin the business processes that you have selected to study (Figure 6-12).

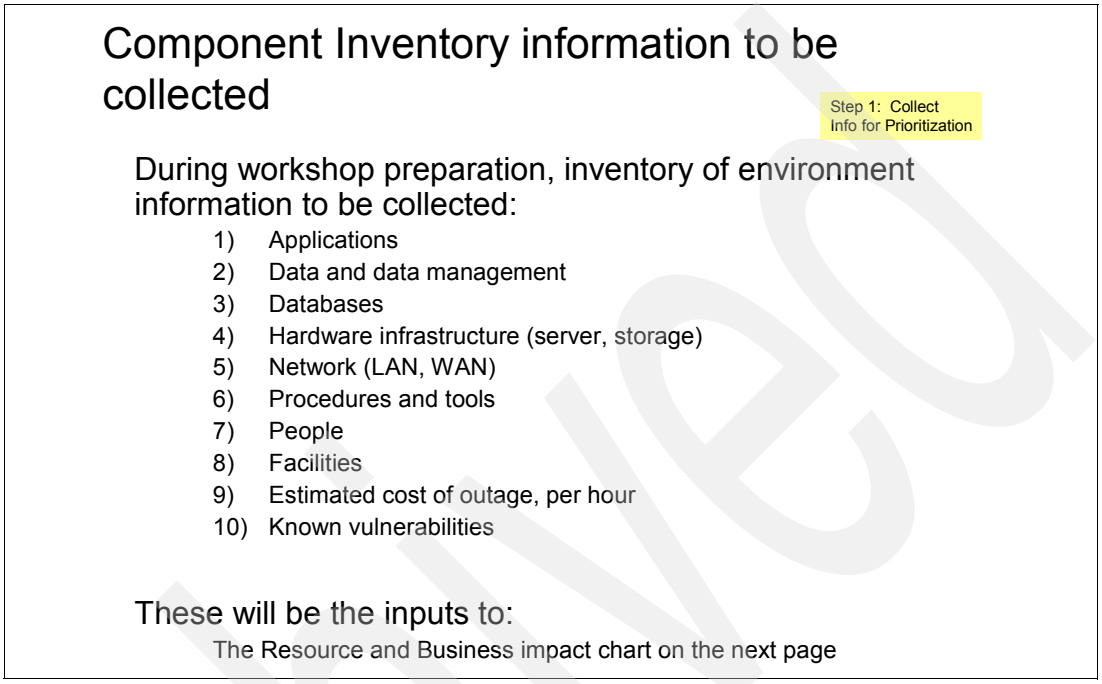

Figure 6-12 Step 1 - component inventory information to be collected

We include some important points about this component information here. A high-level description is sufficient. We drill down for more detail as appropriate during the workshop, using the expertise of the workshop participants. (We provide a template for documenting this information in Figure 6-13.)

- **Applications**: List the various applications that make up the business process. Diagram a basic flow chart of how they relate to each other.
- **Data and data management**: For these applications, describe at a high level, the current data allocation, data management, data backup policies, and the infrastructure for this data. amount of data in GB/TB. What are the change rates/day, how often do we back it up, what is the backup method, what are the data allocation policies (on what storage), and so forth.
- **Databases**: List the databases used to manage the data, and any relevant configuration facts. What facilities are used in that database software that might affect the recovery, and so forth.
- **Hardware infrastructure** (server, storage): List the IT hardware infrastructure on which the business processes, applications, data, and databases reside. Diagram the servers, storage, their interconnect diagram, any relevant model numbers, and so forth.
- **Network** (LAN, WAN): Describe and diagram the networking infrastructure that connects these business processes. Include line speeds/bandwidth, type of line, distances, telecom topologies, and so forth.
- **Procedures and tools**: Diagram the various procedures and the sequence they occur in this business process.

- **People**: What and who are the people that provide the skills and operate the business procedure. Where are they located, and what are the issues related to them?
- **Facilities**: What are the physical facilities and locations that make up this business process?
- **Estimated cost of outage, per hour**: Describe the estimated costs to the business of this business process is not available. as well as description of how those costs changes as length of outage progresses, and so forth.
- **Known vulnerabilities**: List of known vulnerabilities that are desired to be addressed.

At this point, the resulting KPIs, configuration diagrams, and inventory lists of the existing environment should be available and shared with all the participants. Ideally, this information would be posted on large flip charts or posters for visual referral during the subsequent steps of the workshop.

This information is input into the Risk Analysis and Business Impact Analysis portions of the workshop. If this information was not available prior to the workshop, the facilitator and participants need to take the time to diagram this information now, at a high level. The list of components shown in Figure 6-12 on page 188 are guidelines for the what is necessary to be listed in the configuration diagrams.

You can document this information in the provided template in Figure 6-13.

Step 1: Collect info for prioritization

Scope - Resource and Business Impact

Component	Business process 1	Business process 2	Business process 3
Applications			
Data and data management policies			
Databases			
Servers			
Storage			
Network (LAN, WAN)			
Procedures and tools			
People			
Facilities and physical location			
Known vulnerabilities			
Estimated cost of outage , per hour			
Business Impact			
Ranking			

Figure 6-13 Step 1 - resource and business impact

Tip: See Figure 7-2 on page 239 for a applied example of this template.

Ranking of business processes and components

With the key business processes identified and a sufficient description of the current environment available for the participants to see, you next complete the summarized template shown in Figure 6-14. Most importantly, you use this template to prioritize and *rank* the business processes, using the template from Figure 6-13 on page 189.

Note: This ranking becomes the *scope* for the workshop.

Component effect on business process

Figure 6-13 was organized by business process so that you can see, for each one, the included components. In Figure 6-14, for each of the components, the resulting diagram documents what business processes are affected should that component fail. You could perform the same analysis on the components of databases, servers, storage, and so forth. Note that to limit the scope of the workshop, this step is optional.

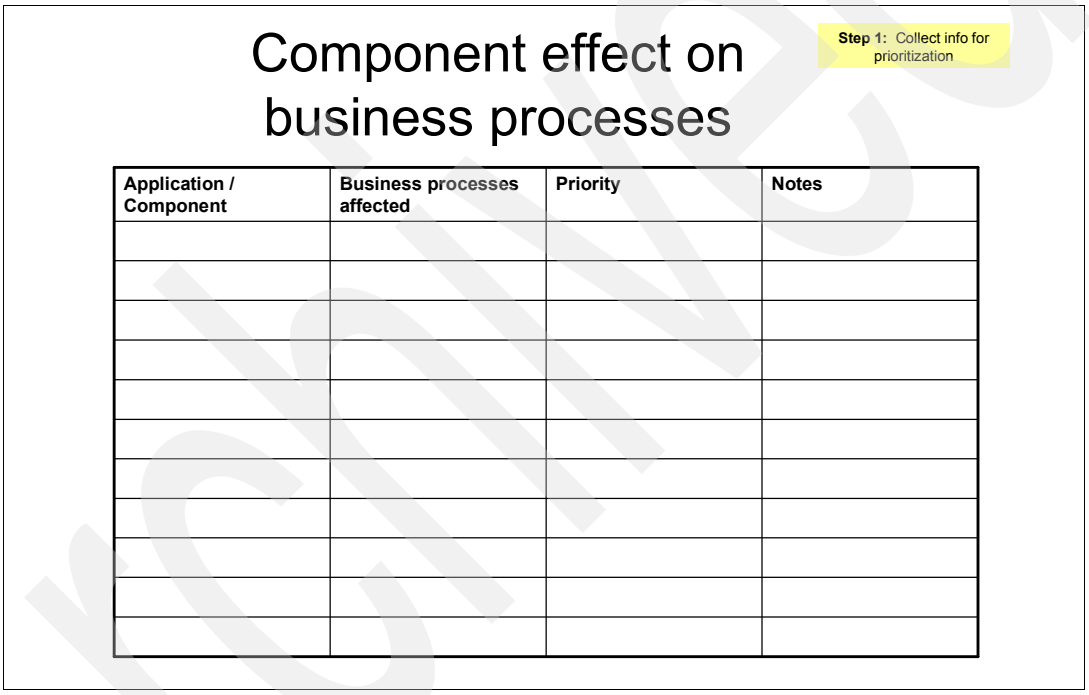

Figure 6-14 Step 1 optional - component effect on business processes

Tip: See Figure 7-3 on page 240 for a applied example of this template.

With these charts completed, you have now documented the scope of the current environment that you will study.

6.3.5 Workshop Step 2 - Vulnerability, risk assessment, and scope

You next proceed to assess the vulnerabilities and risks that you will choose to address, and further refine the scope of the workshop (Step 2 as shown in Figure 6-15).

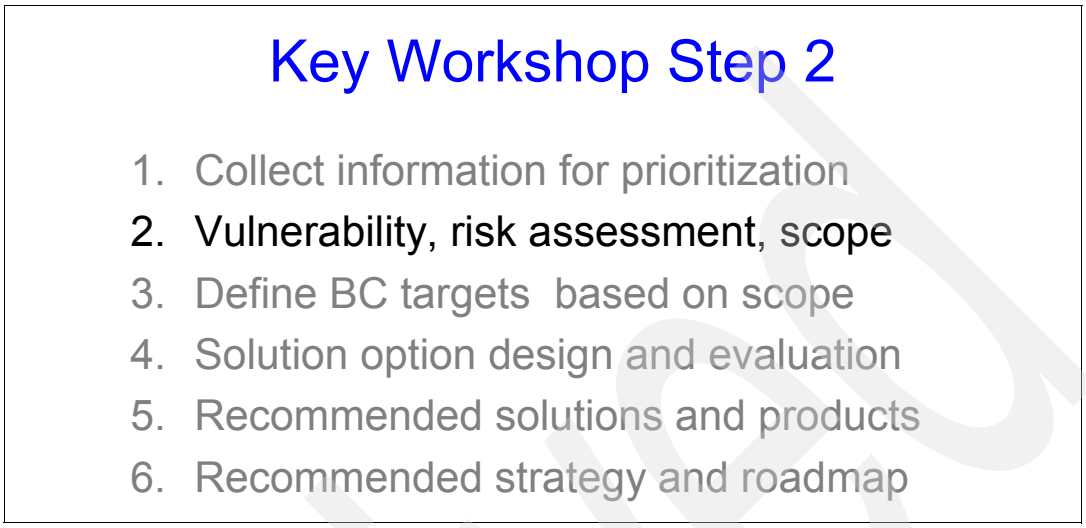

Figure 6-15 Step 2 - Vulnerability, risk assessment, and scope

As you have seen in our methodology and demonstrated again in this step, it is essential that we narrow the scope and prioritize, so that all parties will continue to have an achievable, meaningful experience of this workshop.

Tip: A more detailed discussion of risk assessment and the types of vulnerabilities, is in 3.4.1, “Risk assessment” on page 47.

For our workshop, we suggest that you use the template to diagram these three categories of vulnerabilities:

- Nature
- People
- Equipment

On the template, identify the kinds of nature, people, and IT equipment related vulnerabilities that you want to address. Continue as well with the likelihood ranking. Because you are working on a baseline strategy and a methodology, it is not necessary to cover every vulnerability but only the key high impact or high likelihood vulnerabilities.

To further narrow the scope, we also suggest that you do this just for the data center.

As always, after the methodology and skills are understood, you can then repeat the methodology on other business processes or part of the organization as time goes along.

Figure 6-16 is a template that you can use to obtain a simplified, high-level vulnerability and risk assessment.

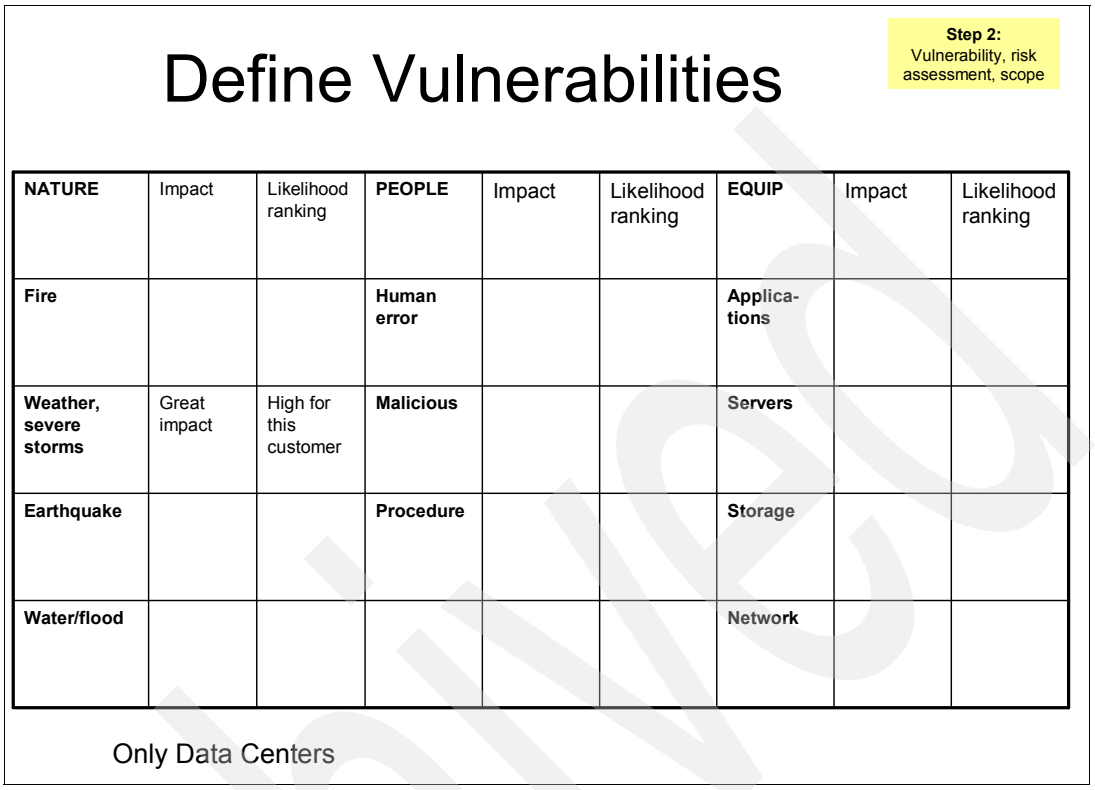

Figure 6-16 Step 2 - Define vulnerabilities - template with example vulnerabilities

As with the business processes, you then select the specific vulnerabilities that comprise an appropriate scope for the purposes of the workshop.

Tip: See Figure 7-4 on page 241 for an applied example of this template.

You can expect that the scope starts at an initial level and then expands in phases. This is a key message: You should address the risks and, thus, the associated scope in phases over time, as shown in Figure 6-17.

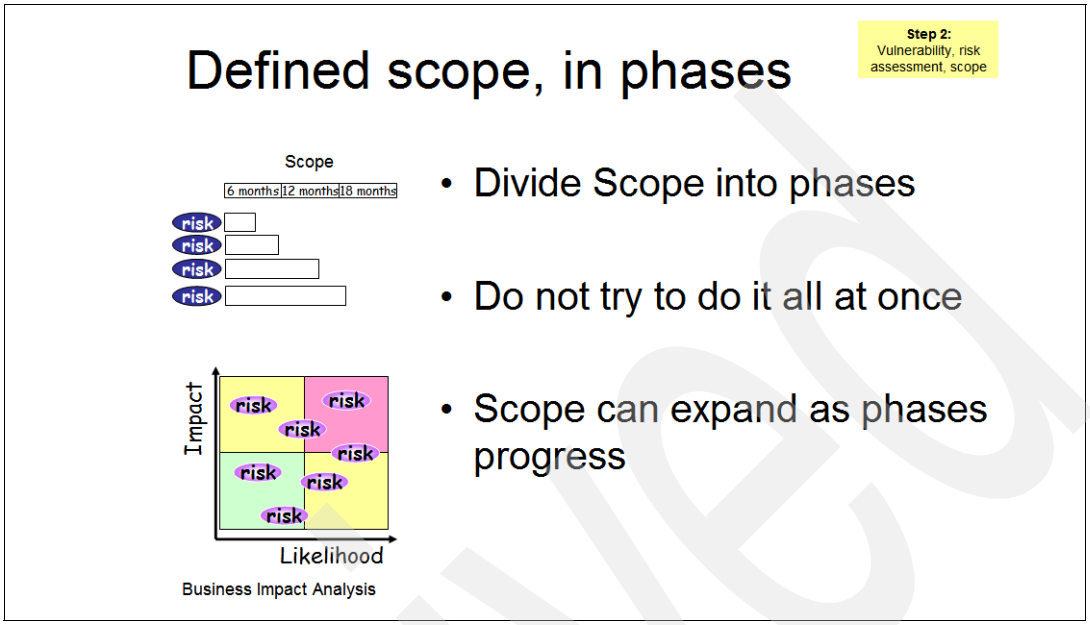

Figure 6-17 Step 2 - define vulnerability, risks in phases

We can also see in Figure 6-17 that the output of this work is a defined set of risks and vulnerabilities that we cover in the rest of the workshop.

The result of these first two steps is that you have arrived at a valid, reasonable *scope* for the remainder of our workshop, as shown in Figure 6-18.

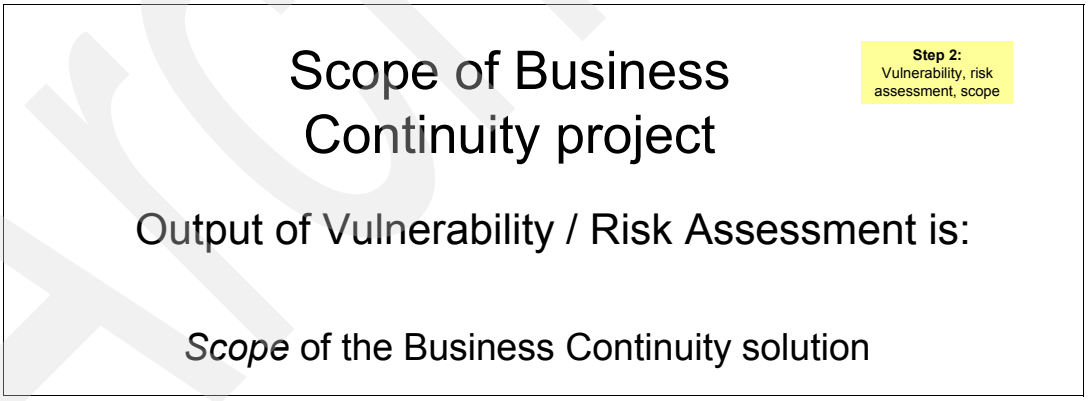

Figure 6-18 Step 2 - Scope of Business Continuity solution

6.3.6 Workshop Step 3 - Define Business Continuity targets based on scope

You next apply this defined scope onto the Business Process and IT environments that you documented in Step 1.

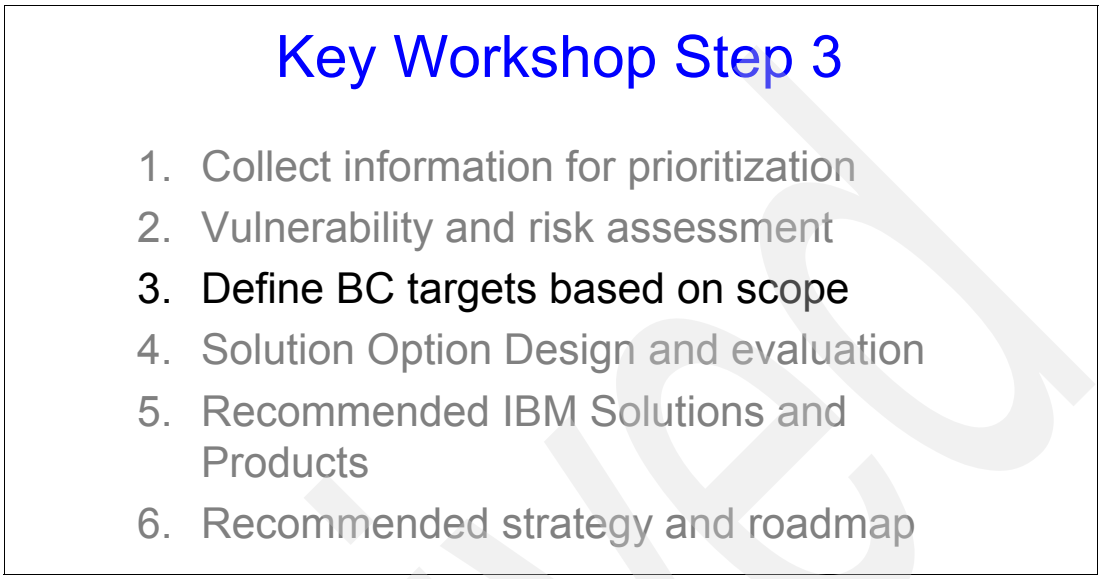

Figure 6-19 Step 3 - define Business Continuity targets based on scope

In the template in Figure 6-20, you take information that you gathered in Steps 1 and 2, and discuss an initial set of desired Business Continuity targets.

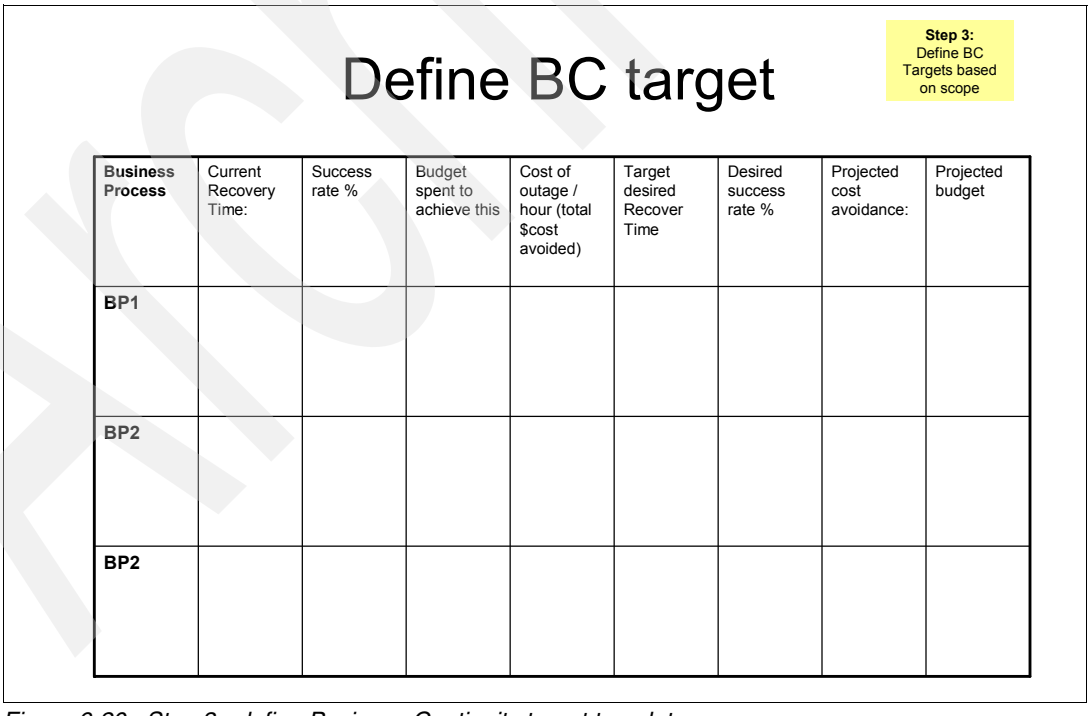

Figure 6-20 Step 3 - define Business Continuity target template

As we discuss the information in these columns and complete the template, some points on the success rate are:

- The percentage success rate of the current recovery time is directly proportional to the frequency of testing. Part of our strategy will be to address ways we can affordably increase the testing frequency, thus laying ground work for better, more provable and more audit-compliant success rates.
- List the desired target recovery time. As we discussed in (Chapter 5., “Business Continuity Solution Selection Methodology” on page 151), this is the *major* determinant of the Business Continuity technology that is selected. Use an initial desired time. You refine this in later steps as necessary. More importantly, you use this metric to determine the intermediate steps that would be necessary in pursuit of this RTO.

Tip: See Figure 7-5 on page 242 for an applied example of this template.

We next discuss a suggested decision criteria for the Business Continuity solution.

The normal criteria of budget, timeline, performance of course apply. We suggest, however, that an additional set of essential criteria should be met, in order for the business to best address the many coming trends as we describe in Chapter 2, “Industry Business Continuity trends and directions” on page 9.

Those additional criteria can be described as shown in Figure 6-21. Discuss these criteria and arrive at a consensus for decision criteria.

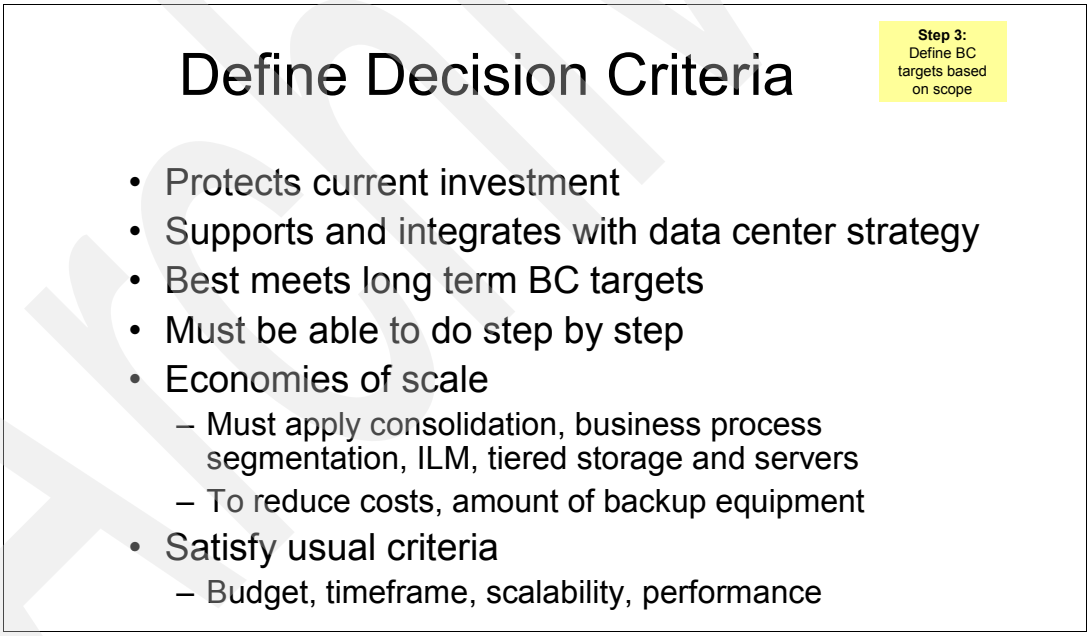

Figure 6-21 Step 3 - define decision criteria for the Business Continuity solution

In addition to the normal criteria of budget, time frame, performance, and so forth, for today’s business environment, we suggest the following additional criteria:

- **Protects current investment**: Resource and time constraints demand that today’s Business Continuity architecture cannot be a major restructure of the existing infrastructure.
- **Supports and integrates with IT data center strategy**: More importantly, as was discussed in detail in 2.2, “Justifying Business Continuity to the business” on page 12”, the

Business Continuity architecture can best be justified and adequately funded when it is an extension and an integral part of the larger IT Data Center strategy. It is a practical reality that a cost-justifiable Business Continuity architecture must be built upon a foundation of a appropriate amount of IT consolidation and standardization.

- **Best meets long-term Business Continuity targets**: The selected Business Continuity architecture must be able to reliably scale and flexibly expand to meet future Business Continuity targets, without major rework. The Business Continuity architecture that is put in place today is likely to continue to be the architecture that is used for a significant amount of time to come.
- **Must be able to do step-by-step**: We expect in this workshop to plan for an initial scope of Business Continuity coverage. with the intent that you can then expand that scope in phases over time. A good Business Continuity architecture will stage the necessary procedure and process changes, implementation changes, testing, and other aspects over time, building incrementally. This best matches today's budget and resource constraints. Building a phased approach also allows the business to accelerate or slow the implementation pace as the needs of the business and funding change over time.
- **Economy of scale**: For cost and affordability reasons, it is very important to apply techniques that can avoid duplicating the entire IT infrastructure at the backup data center.
  - By using techniques such as consolidation, business process segmentation by recovery criticality, Information Life Cycle Management for efficiently managing the data, tiered storage and tiered servers, and so forth.
  - The financial justification is easier, as we reduce the amount of data required, and manage that data in a more cost-effective fashion
  - These techniques also allow the Business Continuity architecture to reliably scale and extend its capabilities, potentially expanding on a larger scale across multiple Lines of Business. In today's business environment, it is essential to be able to support the Business Continuity objectives even in the event of acquisitions, mergers, and ever larger consolidations.
- **Satisfy usual criteria**: For budget, time frame, scalability, performance.

6.3.7 Workshop Step 4 - Solution option design and evaluation

With the decision criteria for the Business Continuity solution agreed upon, you can now proceed to the major part of this workshop, which is solution option design and evaluation (Figure 6-22).

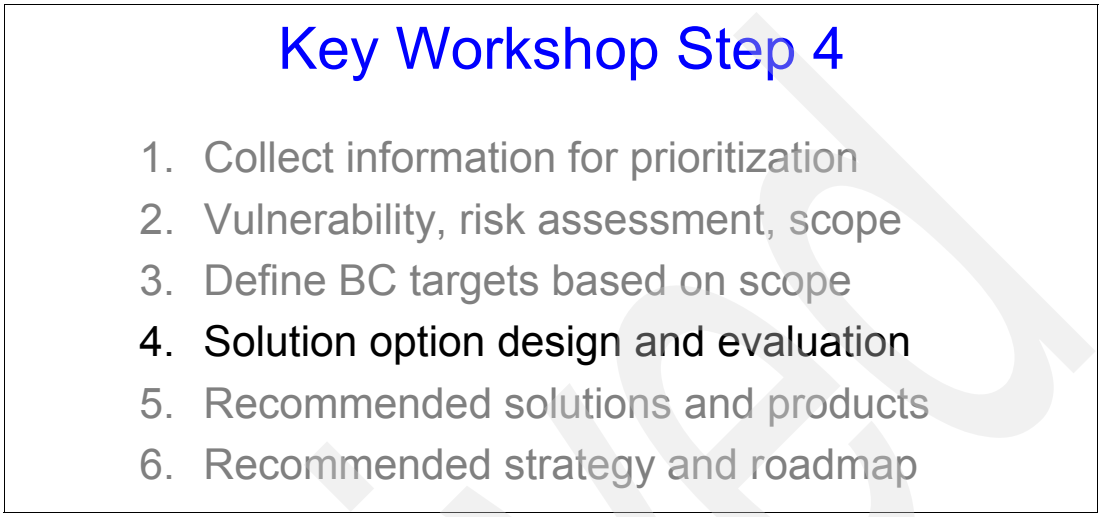

Figure 6-22 Step 4 - solution design and evaluation

You commence a design discussion, using the information that you have already developed. Begin by reviewing basic concepts and principles, as shown in Figure 6-23 and Figure 6-24.

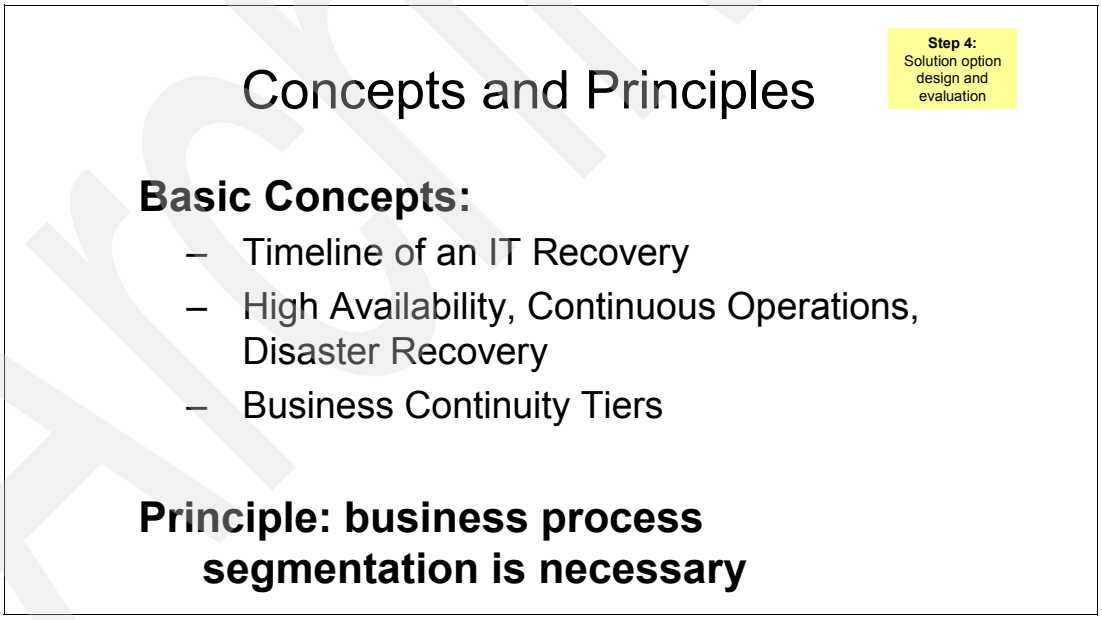

Figure 6-23 Concepts and principles - basics of Business Continuity and need for segmentation

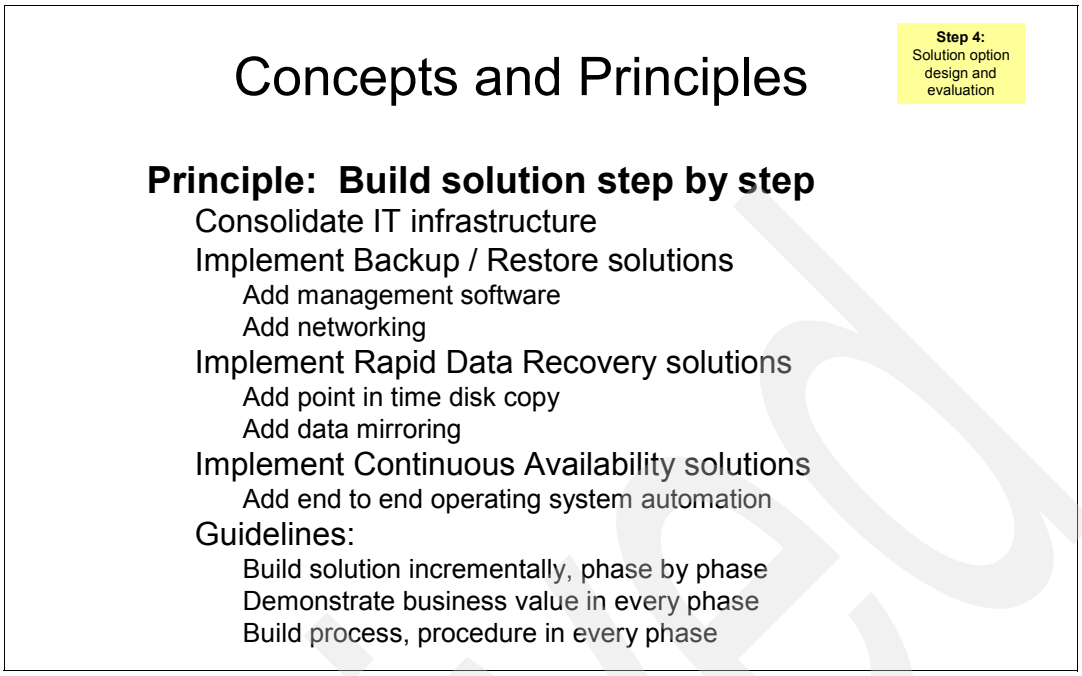

Figure 6-24 Concepts an principles - build solution step-by-step

Using the following charts as visual aids, describe, discuss, refine, and derive the appropriate solution options. The intent is to establish a good best practices knowledge base with the workshop participants as you start the solution options, design, and evaluation step.

As we discuss in “Key Performance Indicators” on page 186, have these key objectives available and posted in the room for all the participants to refer to as you review these concepts.

Concept: Timeline of an IT recovery

Review the concept of the *Timeline of an IT Recovery*, shown in Figure 6-25. This chart shows a *vision* of an end-to-end Business Continuity recovery, including the various people and processes, to assure that all participants are brought to the same base level of knowledge.

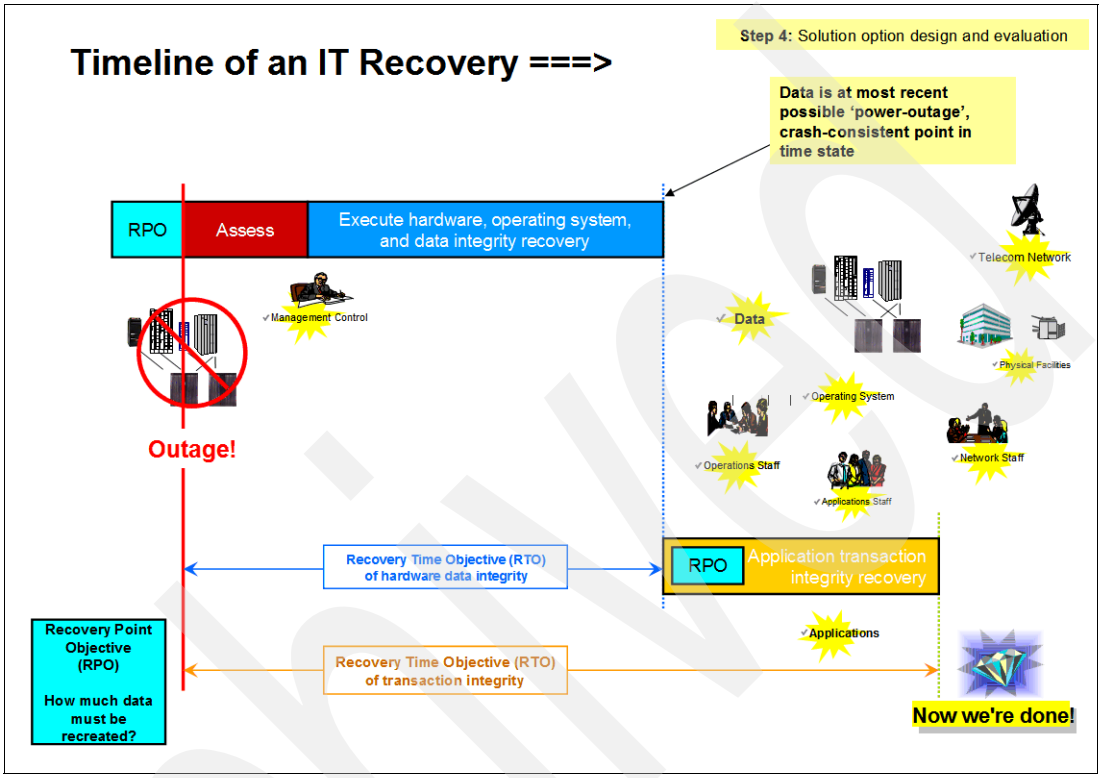

Figure 6-25 Step 4: concept - the timeline of an IT recovery

To briefly review, the concept here is that multiple processes, procedures, and actions make up the timeline of an IT recovery. Assessment is the first phase, recovery of hardware and operating systems is the next phase, followed by application-level recovery as the final phase.

Tip: The detailed discussion of the Timeline of an IT Recovery is in “Definition: RTO and RPO” on page 62.

Concept: High Availability, Continuous Operations, Disaster Recovery

Next, review the definitions of *High Availability*, *Continuous Operations*, and *Disaster Recovery*, shown in Figure 6-26. You design your Business Continuity solution to address these aspects.

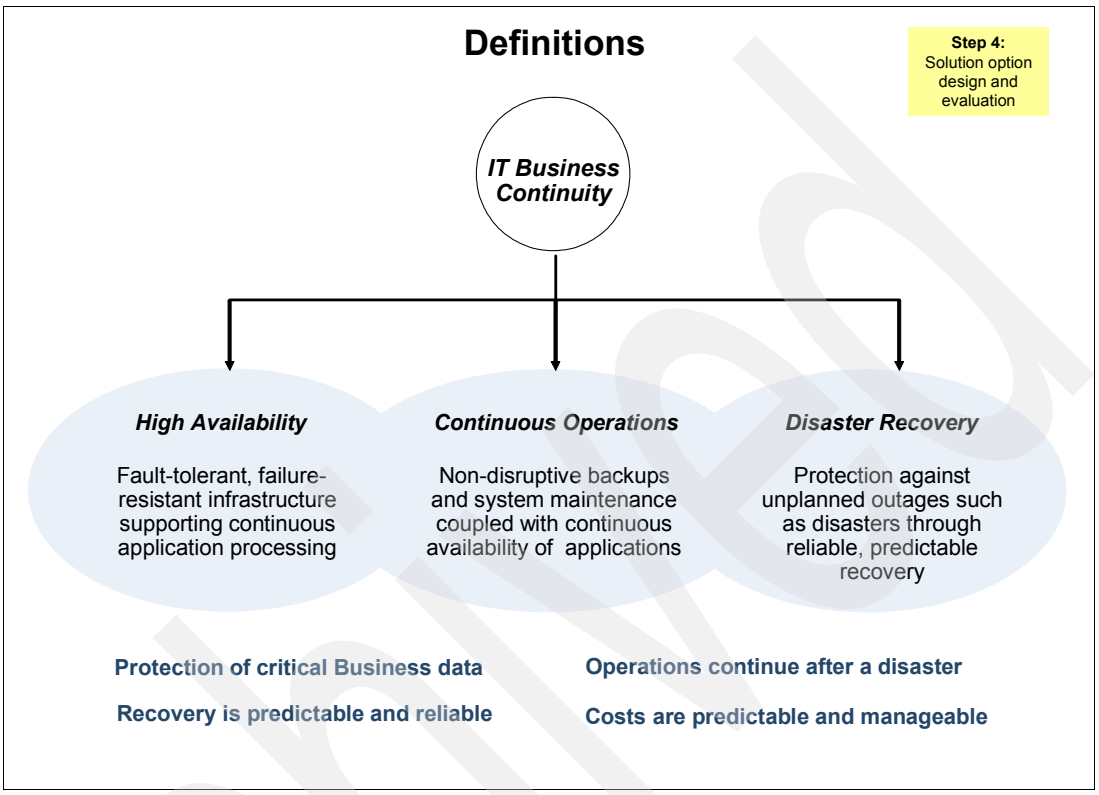

Figure 6-26 Step 4 - Definitions of High Availability, Continuous Operations, Disaster Recovery

To briefly review, the concept is that Business Continuity is made up of three overlapping areas: High Availability, Continuous Operations, and Disaster Recovery.

Tip: You can find the detailed discussion of these definitions in “IT High Availability design” on page 92.

Ideally, when you design the Business Continuity solution properly, the High Availability and Continuous Operations capabilities also provide Disaster Recovery. In this way, you design the solution to provide ongoing value in terms of High Availability and Continuous Operations, which is very important to be able to demonstrate for cost justification purposes.

Concept: Business Continuity tiers

Having defined what you mean by IT Business Continuity, you next review the important concept of the Business Continuity tiers, Figure 6-27.

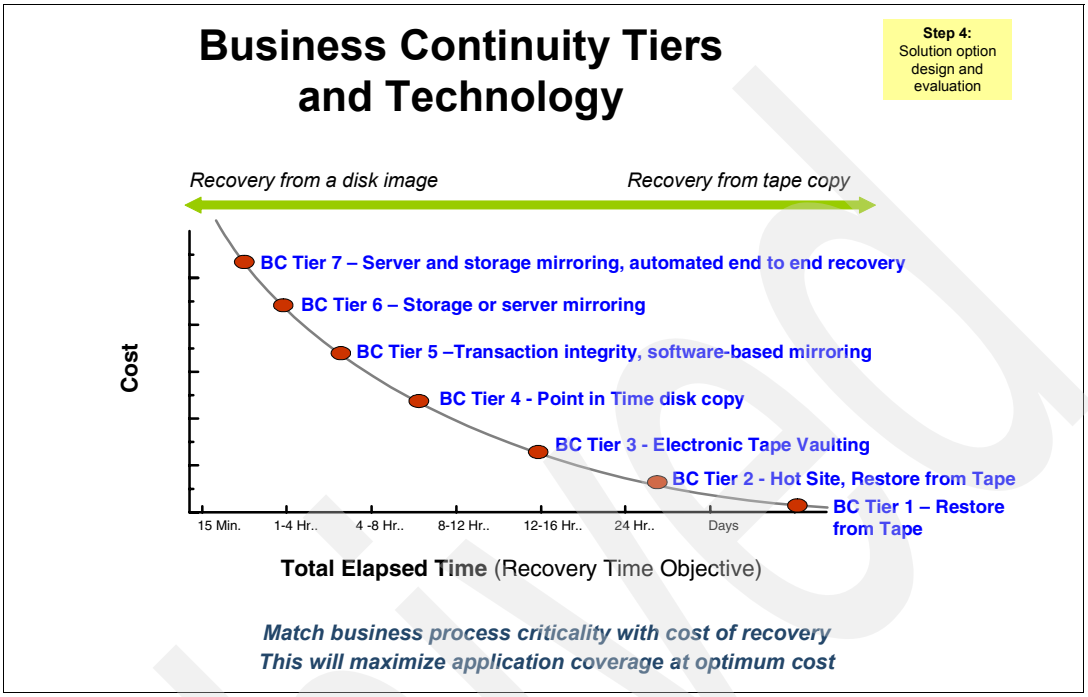

Figure 6-27 Step 4 - principle: Business Continuity tiers

To briefly review, the *Business Continuity tiers* concept allows you to categorize various Business Continuity technologies by their recovery time. In this way, you can more easily match your desired RTO time with the optimum set of technologies. The reason for multiple tiers is that as the RTO time decreases, the optimum Business Continuity technologies for RTO must change. For any given RTO, there are always a particular set of optimum price or performance Business Continuity technologies.

Tip: We discuss the concept of Business Continuity tiers in more detail in “IT Business Continuity tiers of technology” on page 94.

Principle: the need for business process segmentation

Next discuss the necessity of *business process segmentation*, as shown in Figure 6-28:

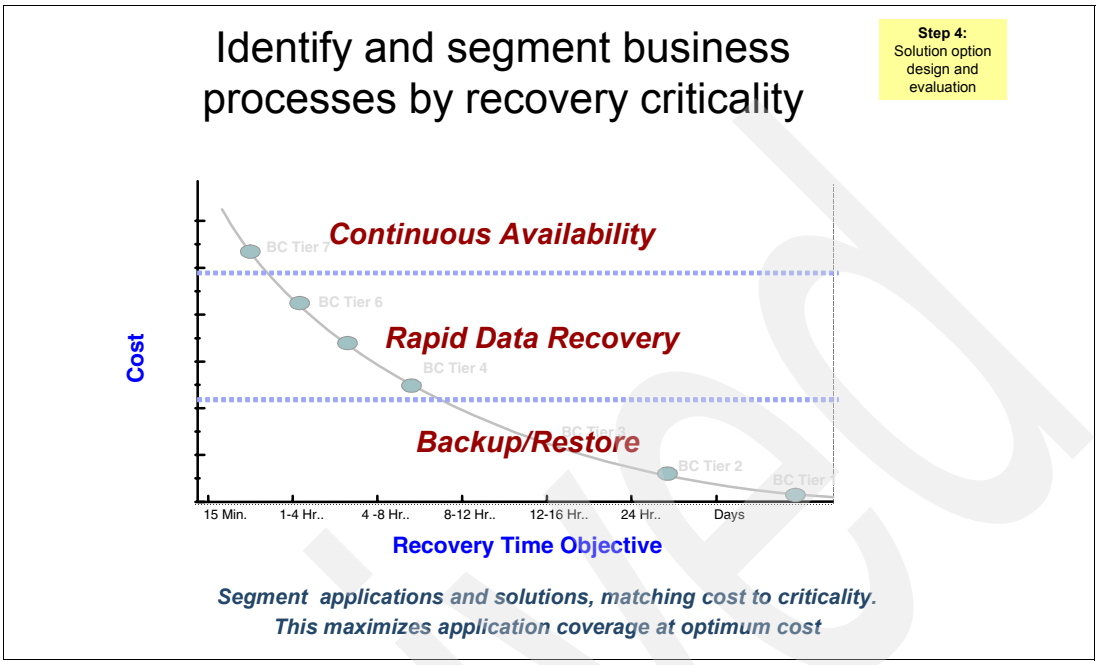

Figure 6-28 Step 4 - Principle: need for business process segmentation

To briefly review, you need to map the desired recovery time for a set of business processes onto the Business Continuity technology that is described in the Business Continuity tiers chart. To account for the fact that not all business processes need to be recovered at the same speed, you need to match recovery cost to recovery speed appropriately. The best way to do this is to architect a strategic segmentation of the business processes.

Tip: More information about the need for business process segmentation is in “Need for Business Process segmentation” on page 96.

You typically implement your solution from the bottom up, going one segment at a time. In other words, you build a good base of Backup/Restore capability, take a subset of those applications and add Rapid Data Recovery, and then take a subset of those applications and add Continuous Availability.

Template: Mapping business process segmentation to technology

Use the template in Figure 6-29 to map the identified critical business processes into this segmentation.

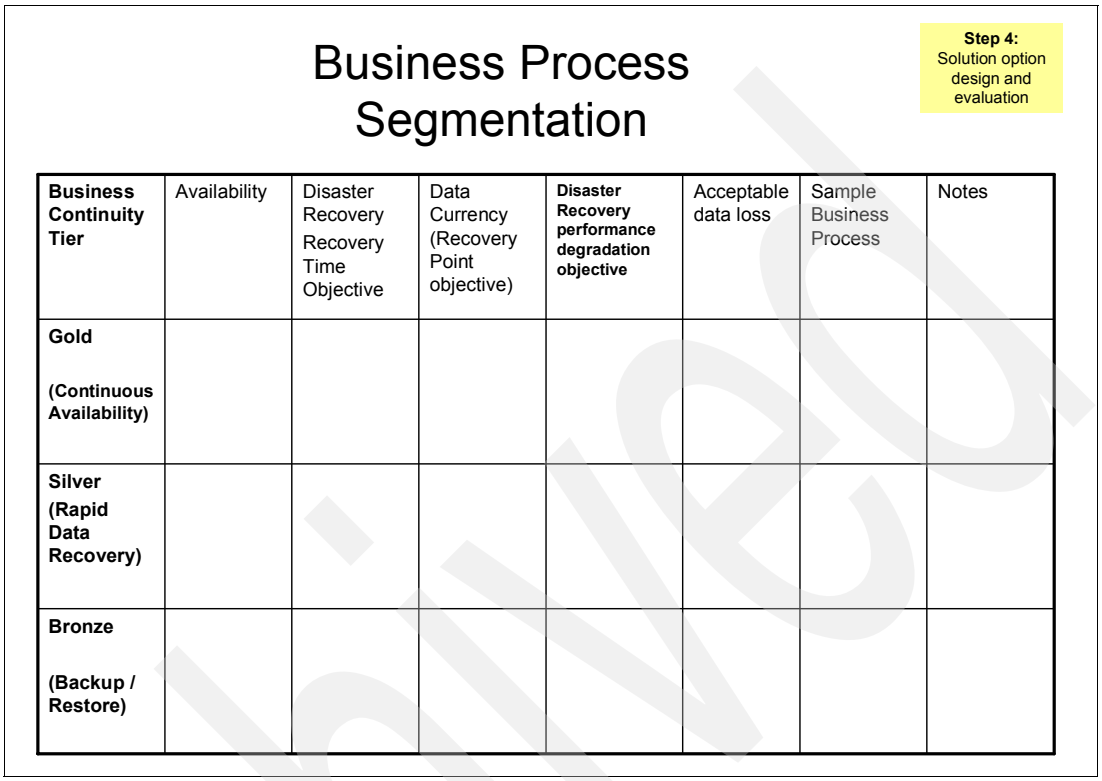

Business Process Segmentation

Step 4: Solution option design and evaluation

Business Continuity Tier	Availability	Disaster Recovery Recovery Time Objective	Data Currency (Recovery Point objective)	**Disaster Recovery performance degradation objective**	Acceptable data loss	Sample Business Process	Notes
Gold **(Continuous Availability)**							
Silver **(Rapid Data Recovery)**							
Bronze **(Backup / Restore)**							

Figure 6-29 Business Process segmentation template

Tip: See Figure 7-6 on page 243 for an applied example of this template.

The chart includes the key determinant information for the success of our workshop because it is the *key* input into our IT strategy design and Business Continuity solution selection methodology.

Foundation: IT consolidation and simplification

You can now briefly review the key concept that realizing IT consolidation and simplification is an essential foundation for the Business Continuity solutioning process. Tieing IT consolidation and simplification to the IT Business Continuity strategy is essential for both cost and business justification purposes (Figure 6-30).

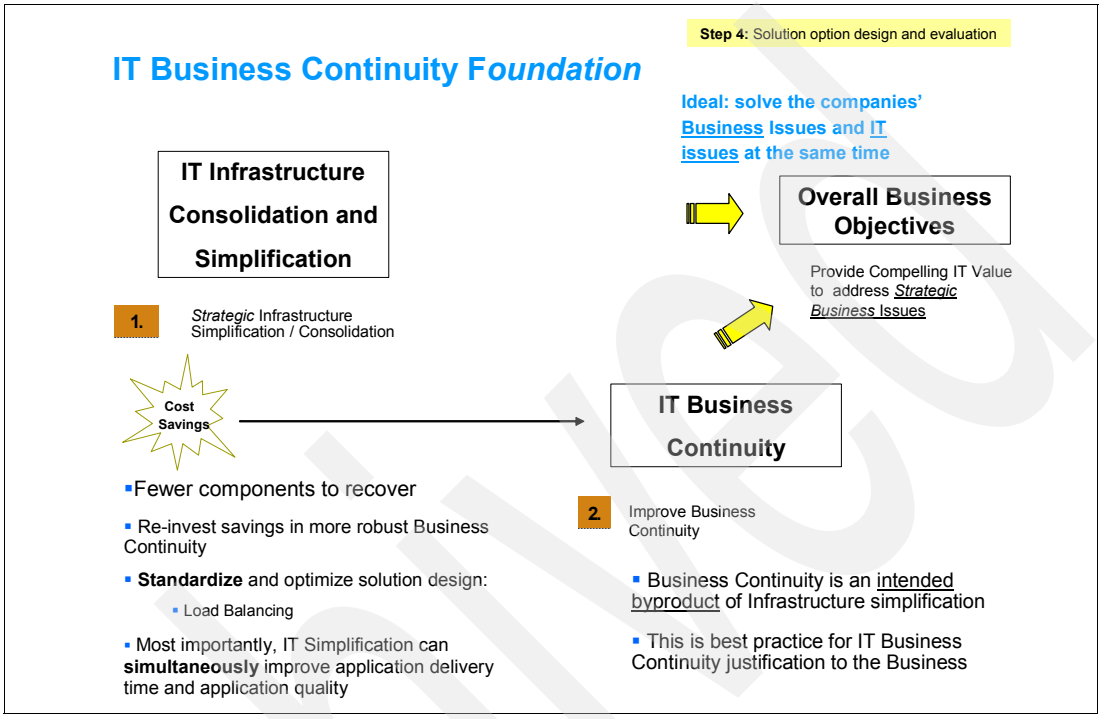

Figure 6-30 Step 4 - principle: consolidated infrastructure as a foundation

To briefly review, the concept is that the Business Continuity architecture should be built upon an adequately consolidated, standardized IT infrastructure. Such an infrastructure has fewer components to recover, and those components tend to be better managed. It is very likely that IT infrastructure simplification and consolidation projects are already in progress. If so, we should simply use and expand the consolidation strategy to introduce components and functions that produce the desired Business Continuity.

Tip: The detailed discussion of this concept is in 2.2.6, "Infrastructure simplification as a prerequisite to IT Business Continuity" on page 19, as well as "IT infrastructure simplification as a prerequisite to Business Continuity" on page 89.

Key IT techniques to be used in the IT consolidation and simplification include:

- Virtualization
- Information life cycle management
  - Tiered storage
  - Tiered servers
- Business process segmentation, that is mapping the criticality of the recovery requirements to the cost of the recovery solution

Work session to design Business Continuity architecture

You now commence a facilitated discussion of all of the factors involved and start to derive the appropriate Business Continuity architecture and solution. You can expect at this point that the skills of the IBM or IBM Business Partner staff or workshop leader (or leaders) are used to lead this working session.

As the working session continues, address common questions and provide valid assumptions that allow you to move ahead.

Begin in Figure 6-31 by defining the difference between *preventative* and *recovery* solutions.

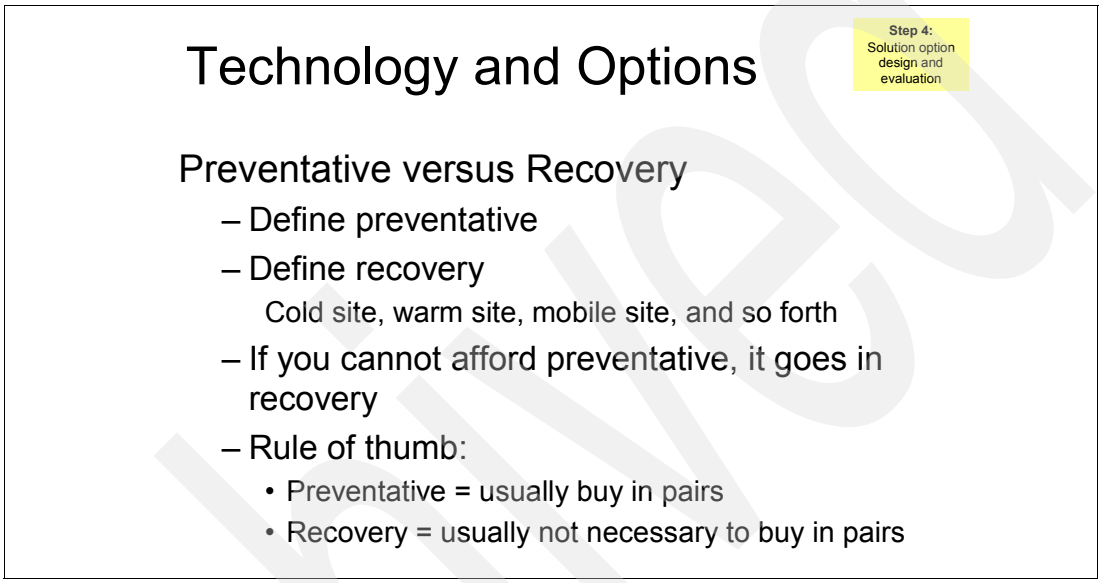

Figure 6-31 Preventative versus recovery

Basically, *preventative* capabilities are proactive high availability, continuous operations Business Continuity technologies. Examples are server clustering, storage mirroring, and so forth. Business processes that need very fast recovery typically are looking for this kind of technology option.

Recovery capabilities are the more traditional recoveries, often from tape, executed at remote sites that might be cold sites, warm sites, mobile sites, and so forth but executed after the fact.

Note: If preventative recovery is not affordable, then the only choice is to use recovery technologies and methods.

As a general rule of thumb:

- If the client uses recovery techniques, it is usually not necessary to plan to buy IT equipment in pairs.
- However, if the client uses preventative techniques, he should expect and budget to buy IT equipment in pairs.

Tip: You can find more information about preventative versus recovery and Risk Management in “Role of risk management” on page 70.

Assumptions and rules of thumb for moving forward

In this section, we document useful assumptions and rules of thumb to allow the Business Continuity planning process to proceed.

What equipment is likely to be needed, when?

Here we need to document the kinds of solutions and associated equipment which typically are needed at both the primary and backup data centers, and when. This is required because the client typically wants to know what equipment and budgeting expenditures are going to be needed and when.

Therefore, we give them a preview of what historically has been the typical progression of components and equipment, at what phase they are likely to arrive, and segment that information between the primary data center and the backup data center

Likely equipment requirements at primary data center

We start by examining the primary data center, as shown in Figure 6-32.

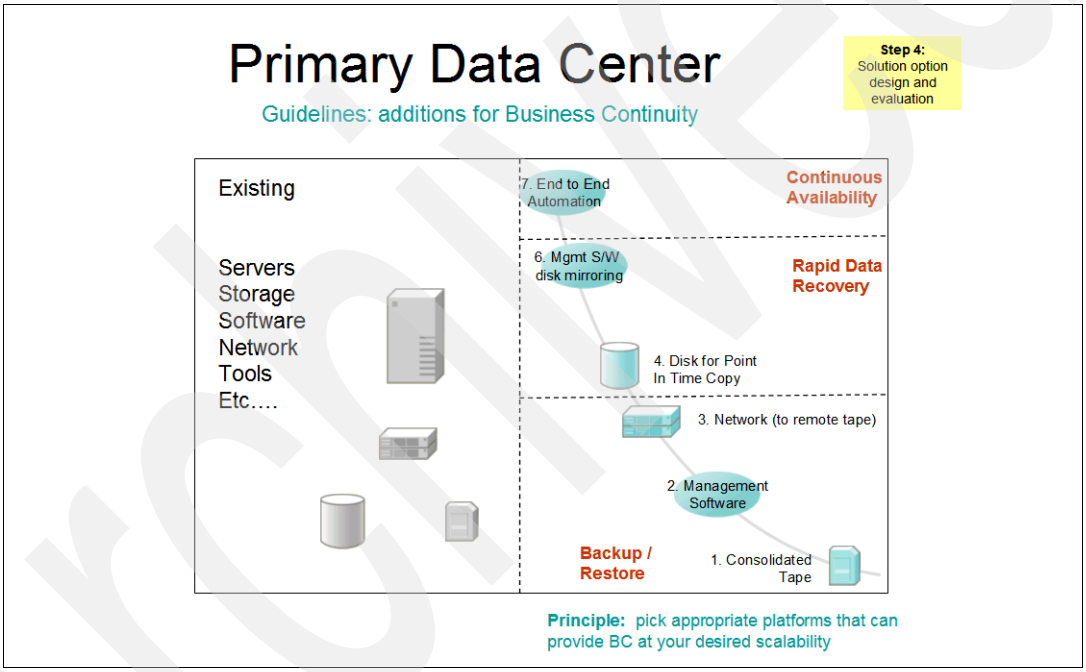

Figure 6-32 Step 4 - principle - build the solution step-by-step - primary data center

This chart shows the existing primary data center equipment, depicted on the left side of the graphic, continues to operate as it does today. On the right-hand side of the chart, we show what equipment is likely to arrive, by phase.

In keeping with our concept of the Business Continuity tiers and the concept of application integration, we will then build the solution in a series of steps, climbing upward through the Business Continuity tiers. We assume that a prerequisite amount of infrastructure consolidation and simplification is already under way.

We would start the first phase by building up the Backup/Restore foundation. We would then progress to adding Rapid Data Recovery only for those business processes and applications that require it. We would then add Continuous Availability only for those business processes and applications that require that level of protection.

Tip: A more detailed discussion of the building a Business Continuity solution in phases is also included in 6.6, “Appendix: Solution visual aids” on page 227.

Referring to Figure 6-32, the steps typically are:

1. Start by streamlining, consolidating, and simplifying the tape infrastructure, to provide a consistent tape backup and restore foundation.
2. Add management software to manage the consolidated tape, which begins the process of adding automation.
3. Add network bandwidth and network control tools, to prepare for the following steps where more rigorous data replication is to be done. We leverage remote tape vaulting as an initial production usage of the network, so that we have a relatively less mission-critical application in the network while we are learning to manage and control that network. Simultaneously, we can further improve the tape recovery speed as we have removed physical tape transport time.
4. Add point in time disk copy capabilities at the primary data center to provide near continuous application availability by avoiding application outages while making the backup. Add this point in time disk capability to the remote tape vaulting.
5. Segment out any recoveries that will be based on database or server software.
6. Add management software that can manage the addition of storage mirroring to the existing storage devices. We use the network that has been tested in previous Business Continuity Tiers 3 and 4.
7. Add continuous availability operating system integration upon all the foundations laid by each of the preceding steps.

We recommend that at each step, we document the incremental recovery time improvements, and communicate those improvements to higher management, thus demonstrating an ongoing value as the project progresses.

Figure 6-33 shows this typical progression of equipment in more detail.

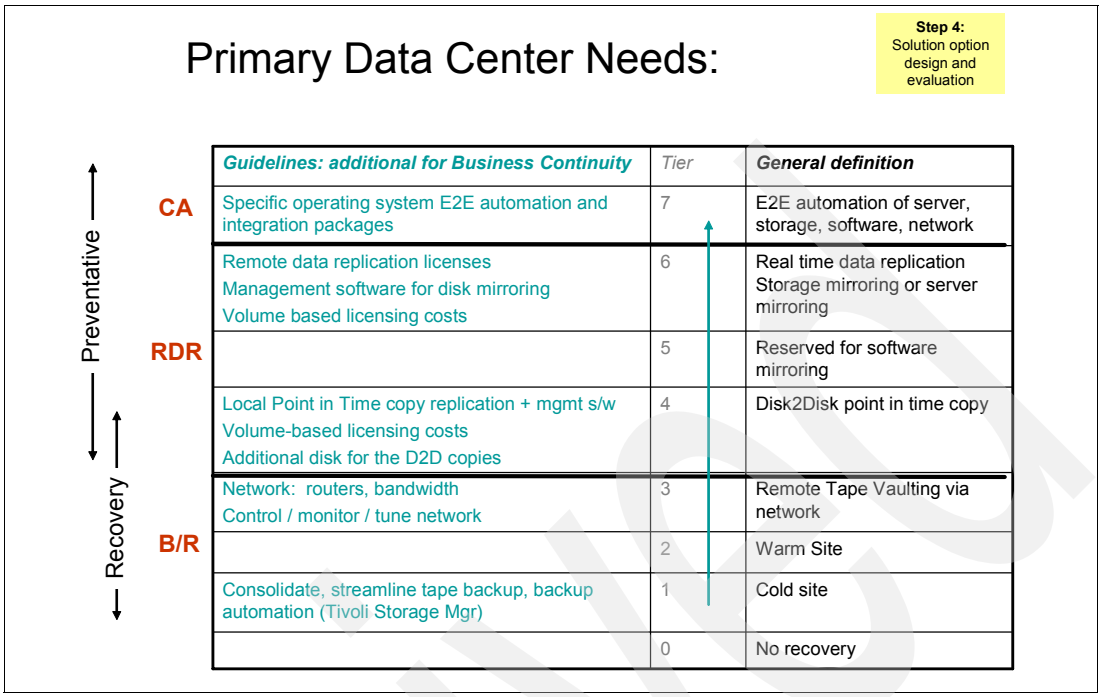

Figure 6-33 Step 4 - listing of primary data center equipment that will be added at each phase

With the information shown, you can give an expectation as to what kinds of investment will be required at the primary data center.

Likely equipment requirements at backup data center

You next show the likely progression of equipment that will be added at the remote data center in a step-by-step phased approach (Figure 6-34).

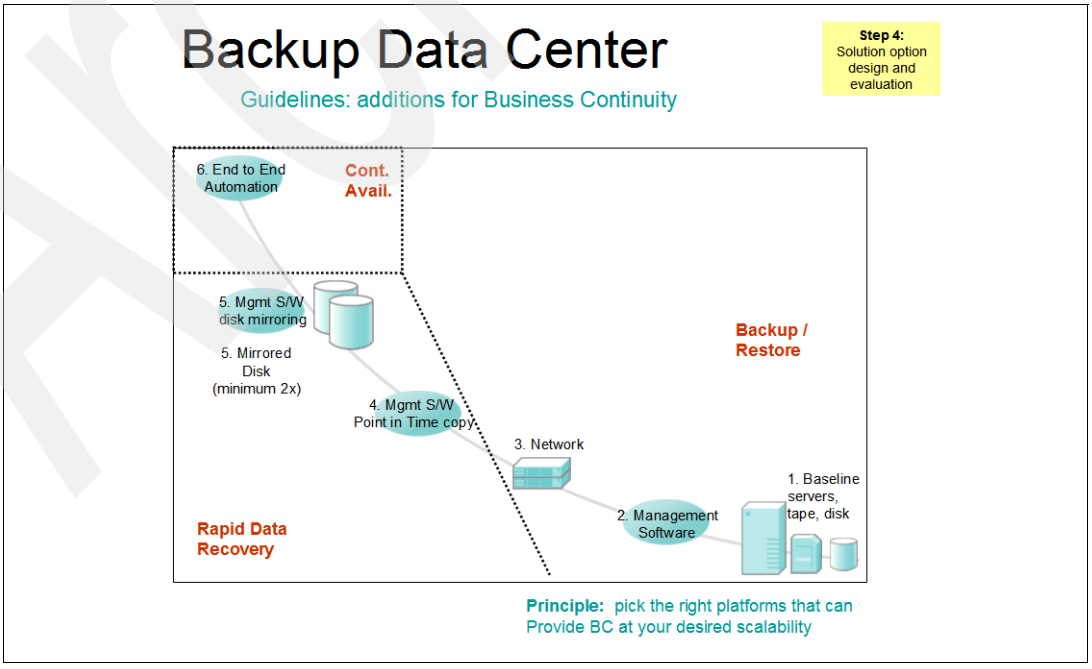

Figure 6-34 Step 4 - equipment that will be added at the remote data center

As you can see, at the backup data center we also follow the concept of building up the solution step-by-step.

For the remote data center, the equipment acquisition curve typically looks like:

1. Addition of a baseline amount of recovery servers, tape, and disk.
2. Add management software to manage the remote servers, tape, and disk, this begins the implementation of automation process.
3. Add network bandwidth and network control tools, as the receiving end of the same implementation at the primary data center. Implement the receiving end of the remote tape vaulting as an initial production usage of the network, simultaneously further improving the tape recovery speed.
4. Add receiving end management software at the remote site as necessary in conjunction with the primary site implementation of point in time copy, followed by remote vaulting.
5. Add receiving end storage for the disk replication. Add management software for management of the storage mirroring. We use the network that has been tested in previous Business Continuity Tiers 3 and 4.
6. Add receiving end continuous availability operating system integration at the remote site.

As with the primary site, for each step, we recommend documenting the incremental recovery time improvements gained by each step, and communicating it to higher management, thus demonstrating incremental improvements and ongoing value as the project progresses.

Figure 6-35 shows a guideline of the types of equipment that are typically needed at the various phase of the project, at the backup data center.

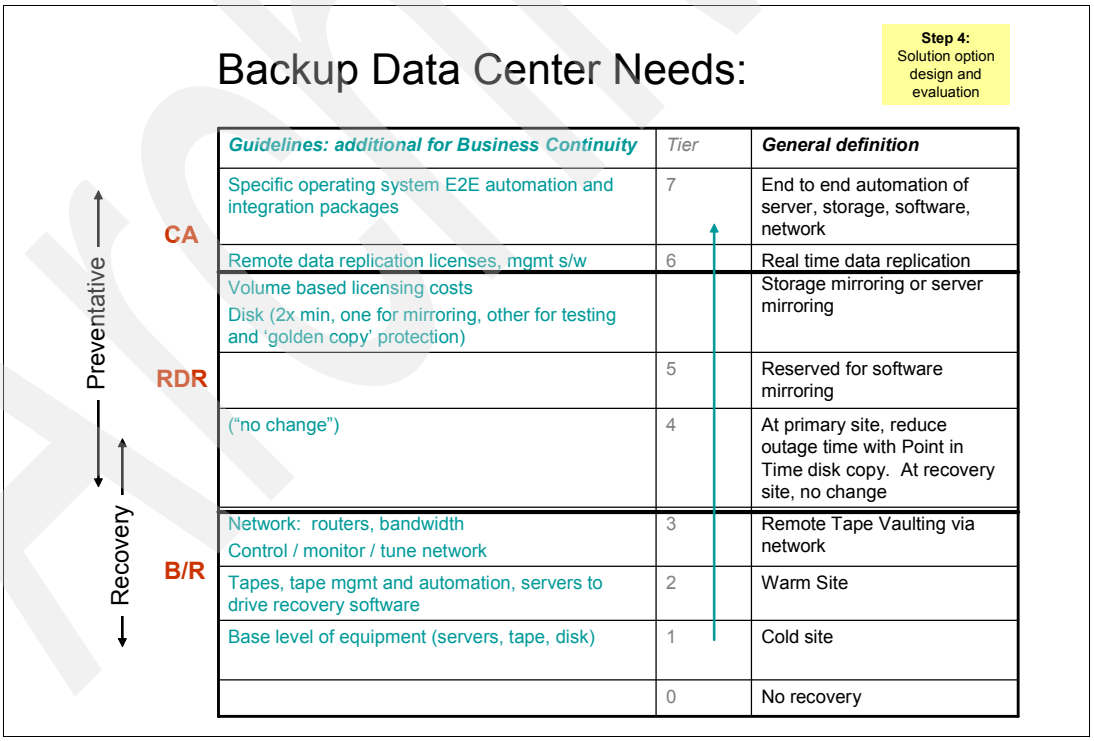

Figure 6-35 Step 4 - backup data center typical equipment needed

Telecom bandwidth requirements for data replication

At some point, the topic of bandwidth requirements, especially for disk replication, is likely to arise. Figure 6-36 shows a rule of thumb for this topic. You can use this to approximate the amount of bandwidth that is required.

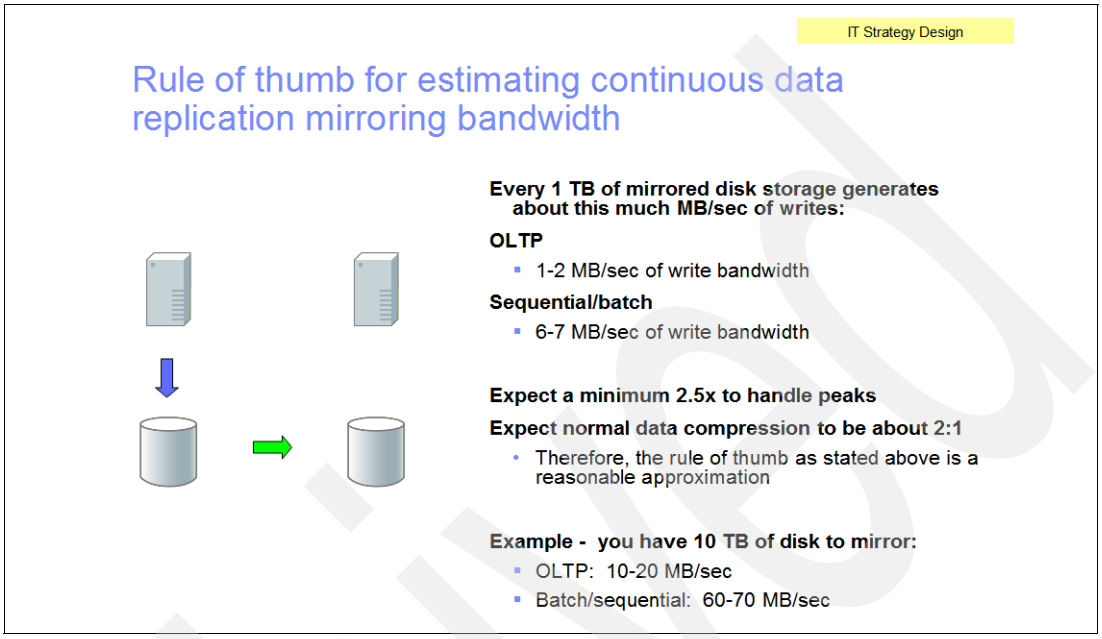

Figure 6-36 Rule of thumb for bandwidth

Tip: More information and the derivation of this rule for storage replication bandwidth is in “Telecom bandwidth capacity: Tips and a rule of thumb” on page 110.

Software replication, server replication, or storage replication?

Other discussions might arise about when to choose software replication, when to choose server replication, and when to choose storage replication. Figure 6-37 should assist in this decision.

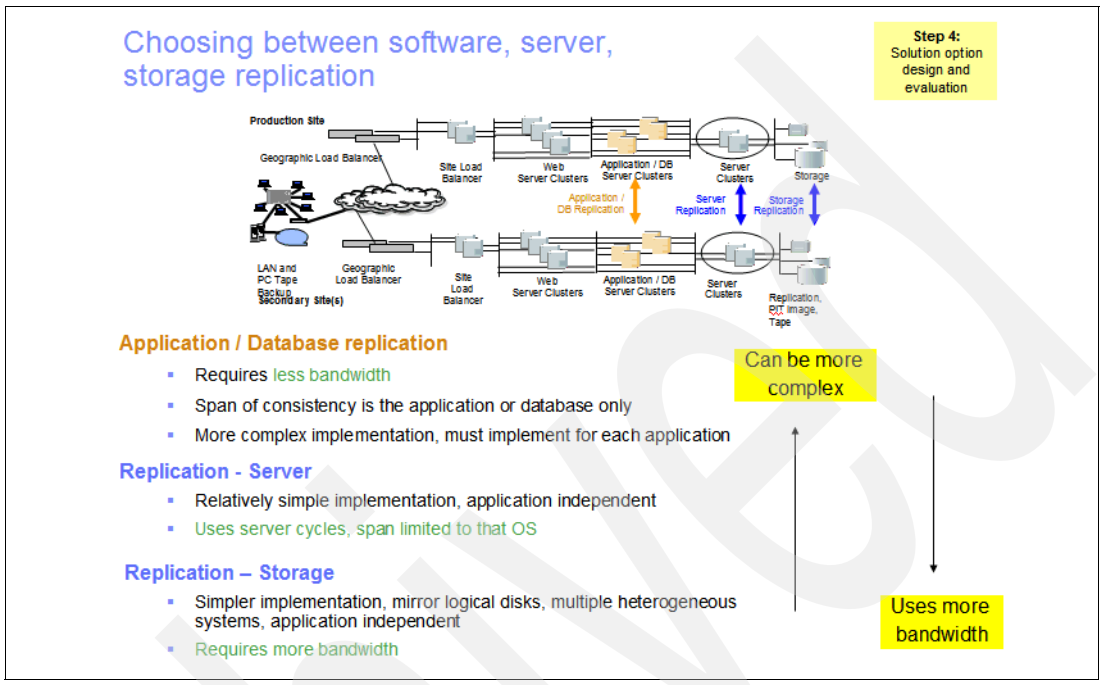

Figure 6-37 Considerations for choosing among software, server, or storage replication

Tip: More information about choosing between software replication, server replication, or storage replication is in “Data replication” on page 101.

Business Continuity Solution Selection Methodology

To provide assistance in the solutioning aspect of this workshop, see Chapter 5, “Business Continuity Solution Selection Methodology” on page 151. Figure 6-38 is an excerpt from the methodology. It lists the most important IT requirements questions for deciding upon the proper candidate solutions.

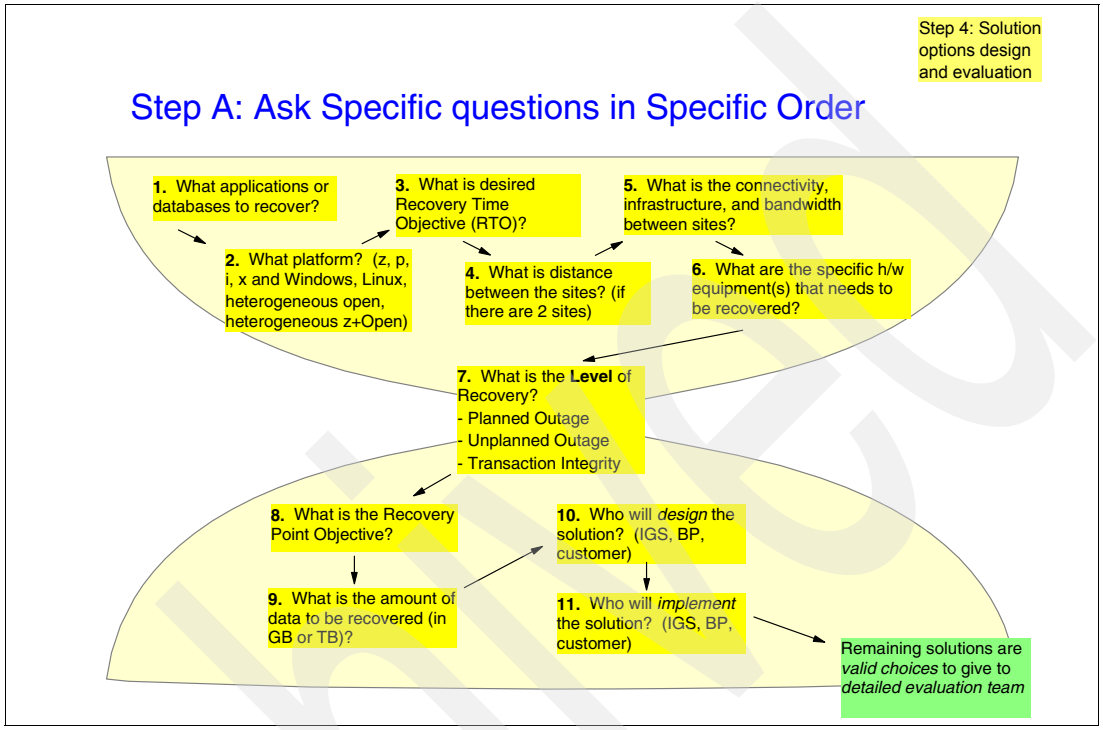

Figure 6-38 Solution Selection Methodology - the 11 IT Business Continuity requirements questions

Tip: More information about the use of these questions is in 5.5.3, “Principle: Asking requirements questions in a specific order” on page 160.

Summary

You should make the key points in Figure 6-39 during every Business Continuity workshop.

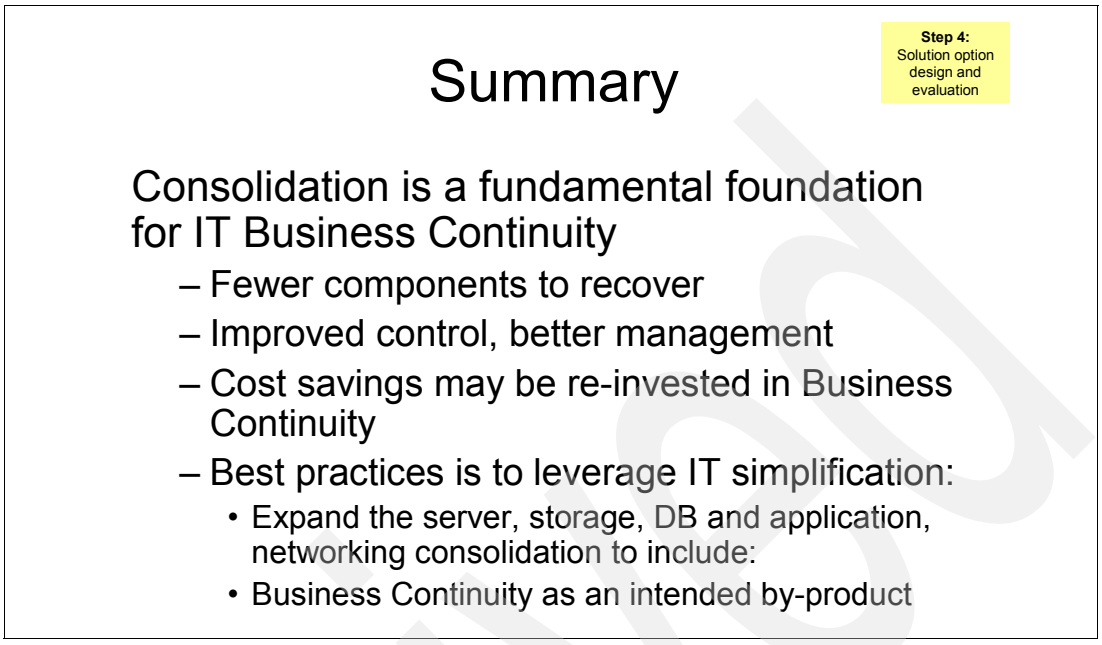

Figure 6-39 Summary - consolidation is a pre-requisite for successful IT Business Continuity

Tip: You can find more information about consolidation and infrastructure simplification as a prerequisite for IT Business Continuity in “IT infrastructure simplification as a prerequisite to Business Continuity” on page 89.

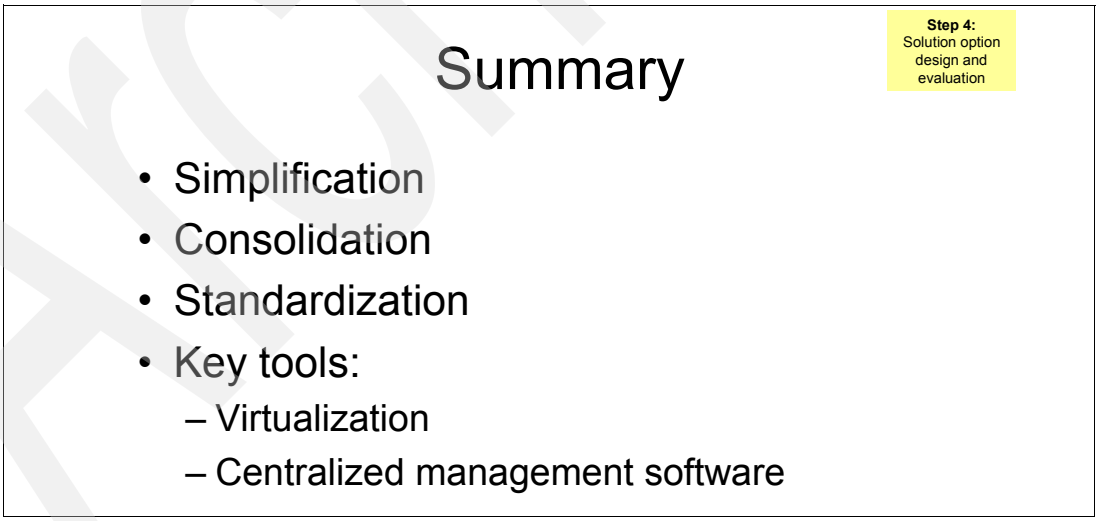

Figure 6-40 Summary - good tools for achieving consolidation include virtualization

Tip: You can find more information about how you can use virtualization and centralized management software in IT Business Continuity in *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548.

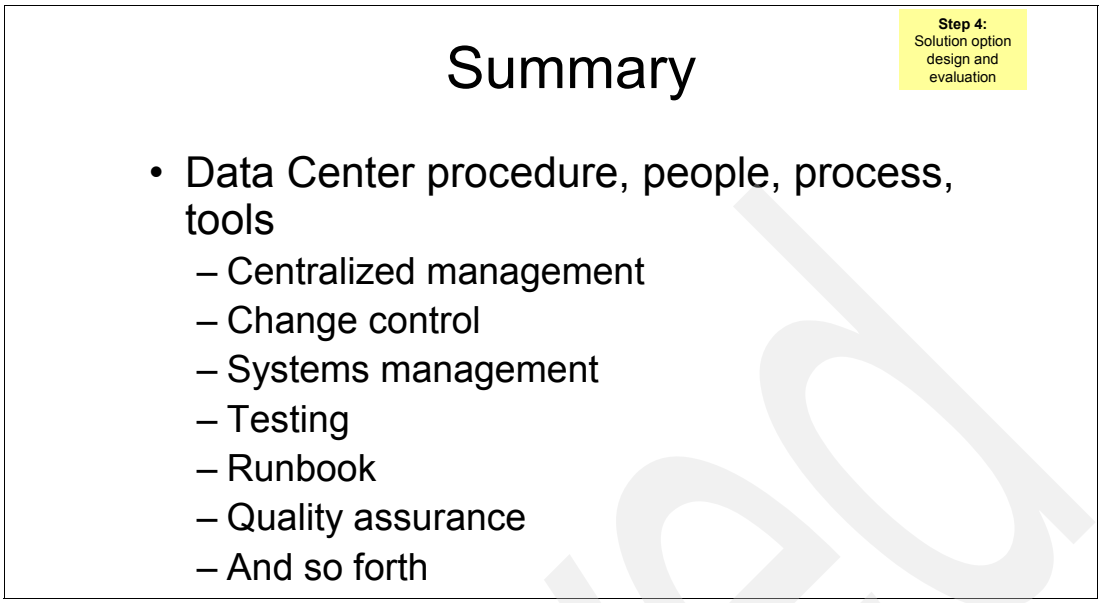

Figure 6-41 Summary - data center process and procedure

Tip: You can find more information about process and procedure in “Process and procedures” on page 68.

In summary, take the discussion, document it well, and apply the summary principles shown in Figure 6-42.

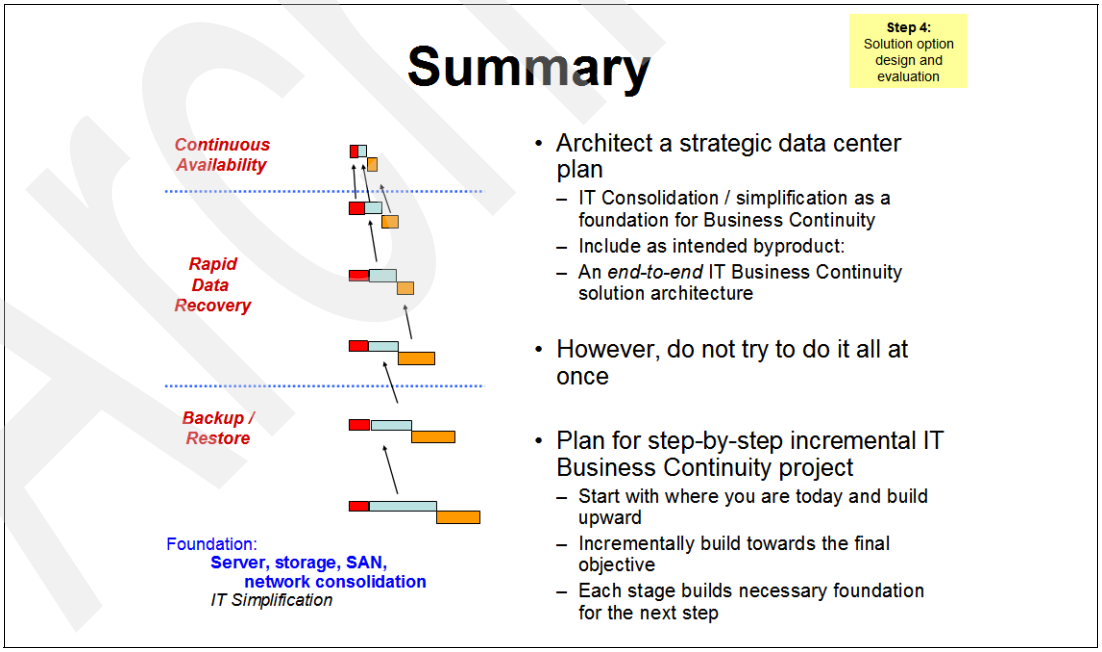

Figure 6-42 Summary - architect a solution and vision, plan for a step-by-step implementation

6.3.8 Workshop Step 5 - recommended IBM solutions and products

In the previous step, you had a long free form conversation regarding the various kinds of requirements, the need for consolidation and virtualization as a prerequisite, bandwidth, and defined and refined what would be the appropriate technology choices. With that discussion and using the expertise of the workshop leaders, you should at this point have arrived at a consensus of a good Business Continuity architecture and have a good idea of the kinds of specific technology options from which to choose (Figure 6-43).

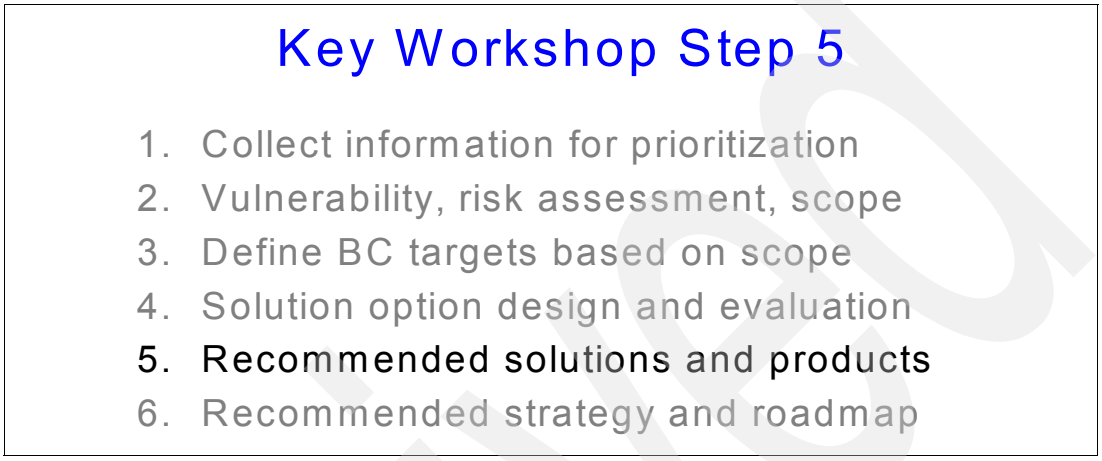

Figure 6-43 Step 5 - document recommended IBM solutions and products

The workshop facilitator and workshop subject matter experts would lead the participants to a joint consensus on the appropriate solutions that were the result of these conversations. Figure 6-44 is a template listing the main items that should be covered in the recommendations.

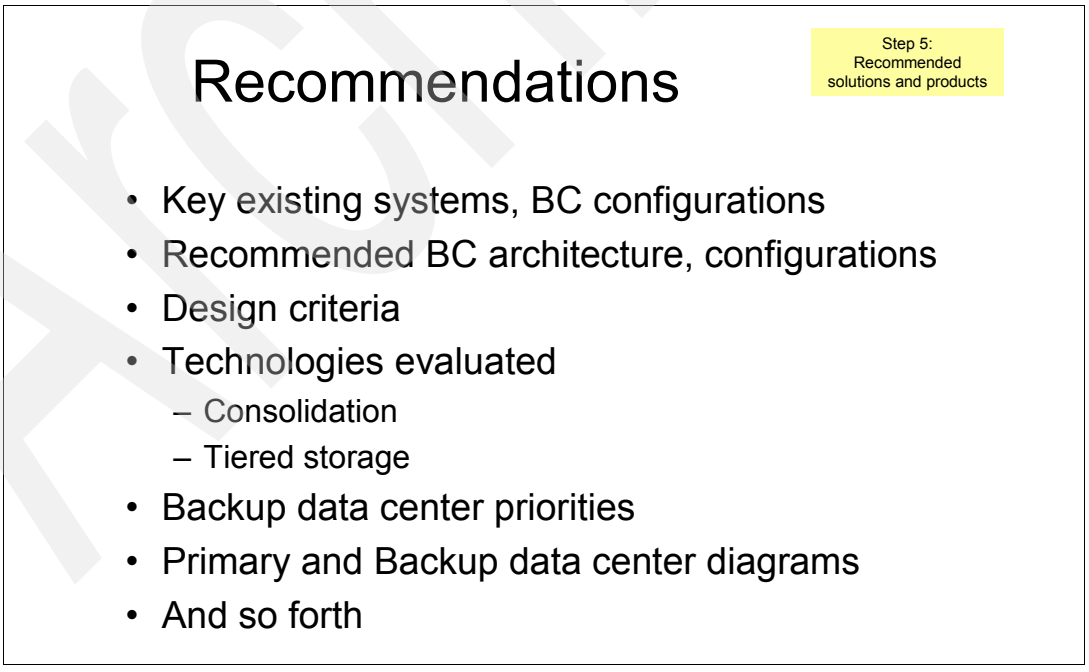

Figure 6-44 Step 5 - Solution recommendations template

Tip: We provide a detailed example of what these recommendations would look like in the case study in 7.4, “Recommended Business Continuity architecture and configurations” on page 245.

6.3.9 Workshop Step 6 - recommended strategy and roadmap

Finally, you end with the recommended strategy and roadmap, as shown in Figure 6-45.

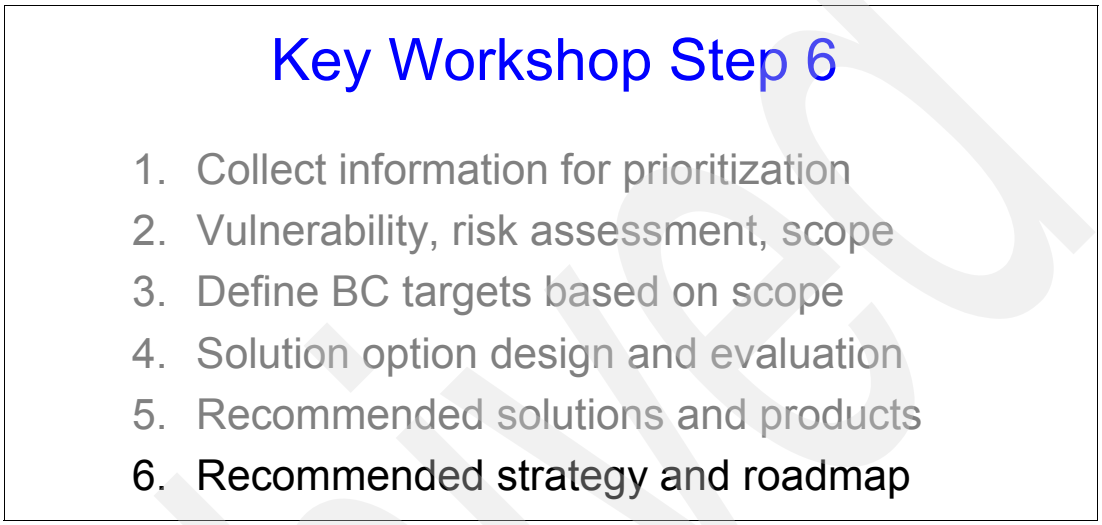

Figure 6-45 Step 6 - document a recommended strategy and roadmap for this client

List the key benefits and challenges to the IT organization for the proposed solution (Figure 6-46).

Figure 6-46 Step 6 - template for benefits and challenges

Tip: A completed Benefits and Challenges statement is in the case study in “Benefits” on page 246 and “Challenges” on page 247.

List the financial implications and justifications, as shown in Figure 6-47.

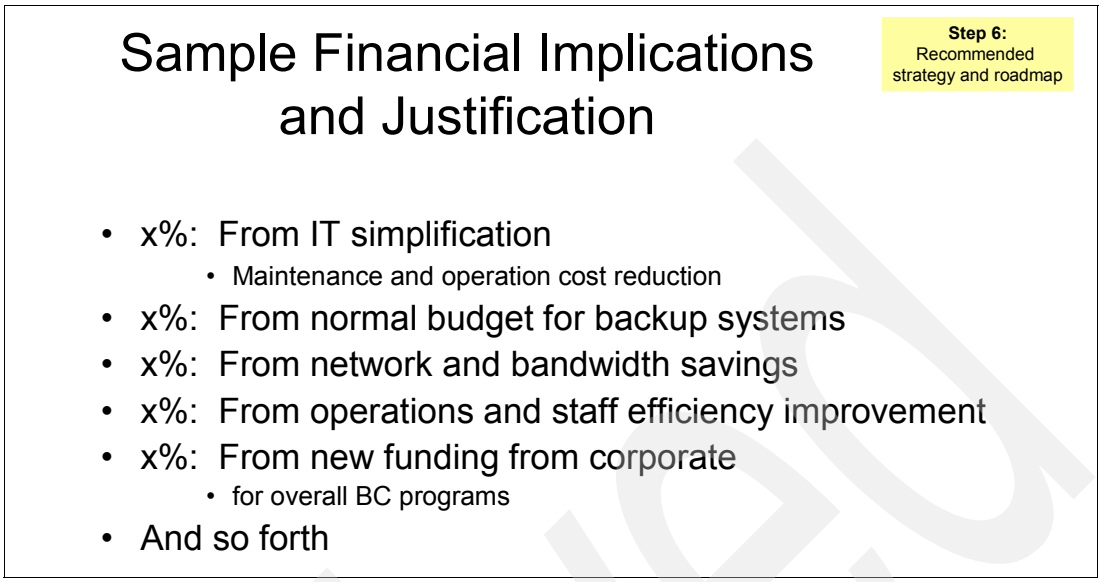

Figure 6-47 Step 6 - template for financial implications and justification

Tip: A completed Financial Implication and Justification statement is in the case study in 7.4.1, “Financial implications and justification” on page 247.

Provide an implementation plan.

Tip: A completed Implementation plan is in the case study in 7.4.2, “Implementation planning” on page 247.

Finally, provide a Next Steps list that is a roadmap of where the IT organization can and should go next. Figure 6-48 is a template.

Figure 6-48 Step 6 - Next Steps and roadmap template

Tip: A completed Next Step and Roadmap statement from the case study is in 7.4.3, “Next Steps and roadmap” on page 249.

As an executive summary, you can use Figure 6-49 as a visual template for laying out the high-level Next Steps and roadmap.

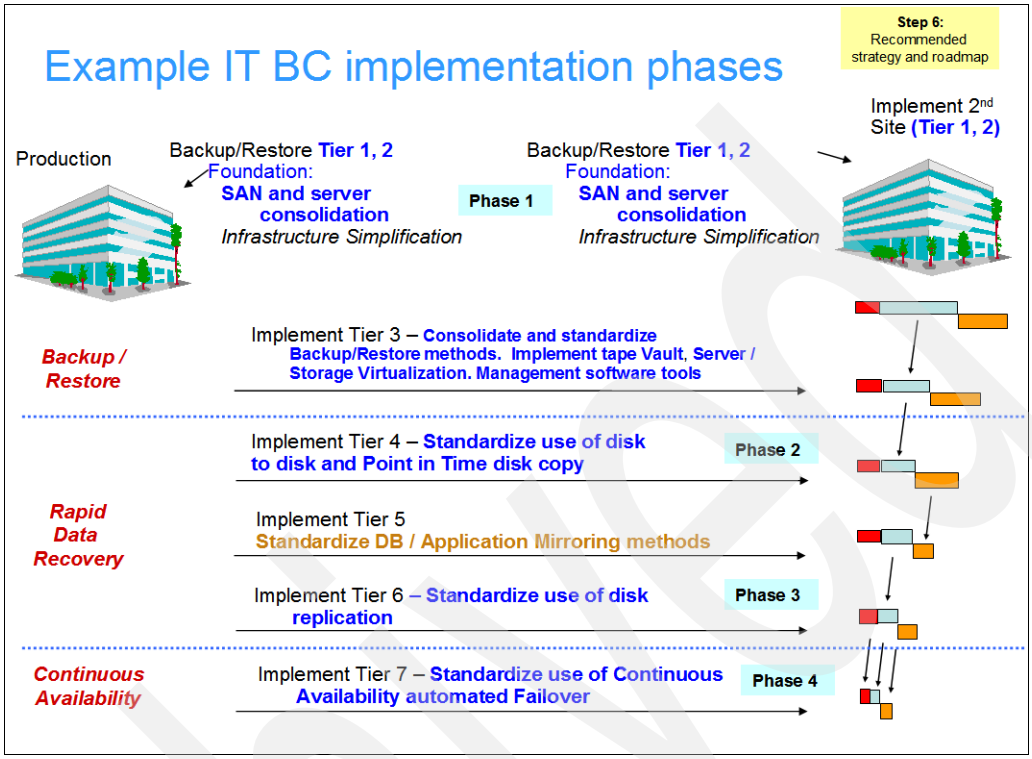

Figure 6-49 Step 6 - Executive summary Next Steps template

Finally, you need to make the points shown in Figure 6-50 at the end of each Business Continuity workshop.

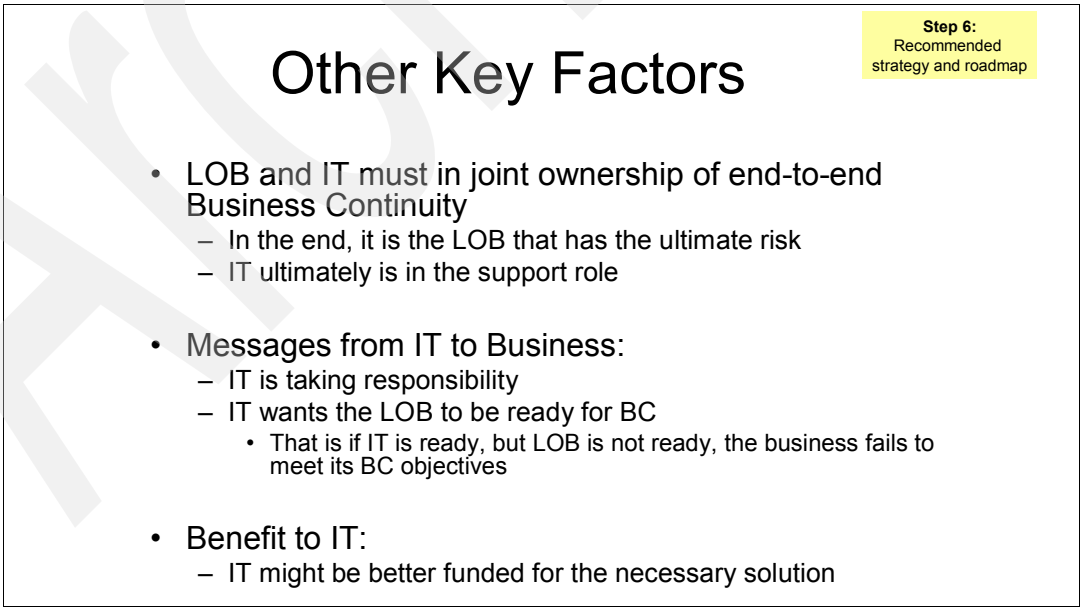

Figure 6-50 Step 6 - important recommendations to be included in any Business Continuity workshop

The Lines of Business and IT must be in joint ownership of the end-to-end Business Continuity project to achieve true Business Continuity.

This is especially important for IT, even though it is the Lines of Business that have the ultimate risk, and even though IT is in the support role, IT is often viewed as the reason for success or failure. Therefore, in order for IT to assure their own success, they must successfully have the LOB agree and participate in joint ownership. This is a very important message to IT.

Gaining such an agreement also can help IT to assure that funding for the IT aspects of the Business Continuity project is adequate.

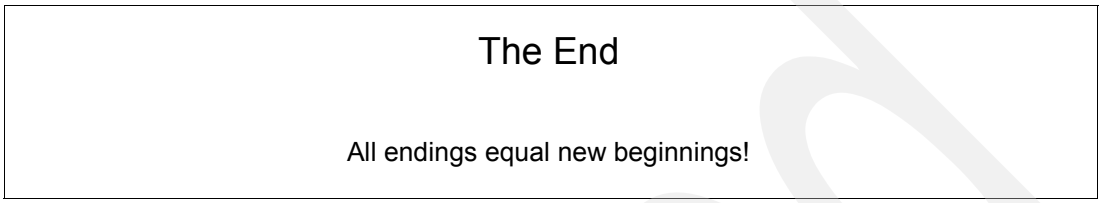

Figure 6-51 All endings equal new beginnings

We conclude our Next Step Business Continuity workshop methodology overview.

Tip: We recommend that next, you read and examine the supplied sample case study Chapter 7, “Next Step Business Continuity workshop: Case Study” on page 233, as an example of how this methodology can be applied.

We also provide the following supplemental information:

- 6.4, “Appendix: Sample Statement of Work for the workshop” on page 219
- 6.5, “Appendix: Why Business Continuity” on page 223
- 11.6, “Appendix: Solution visual aids” on page 227

6.4 Appendix: Sample Statement of Work for the workshop

This section presents a sample Statement of Work for the facilitators of this Next Step Business Continuity workshop.

Introduction

We are pleased to respond with this Statement of Work (SOW) to (*client name*)’s request to conduct a one day Next Step Business Continuity workshop for some key applications in the data centers for Business Continuity situations. The following are the objectives, scope and deliverable:

- To help (*client name*) CIO and IT department with a sample Business Continuity planning methodology in order to develop a Business Continuity strategy and plan for a few key applications, data and systems in the data centers.
- Based on the outcome of the sample assessment, to help (*client name*) make a decision on whether a comprehensive Business Continuity planning project should follow.
- To run some sample Business Continuity scenarios on two or three key applications or systems, using the Next Step Business Continuity methodology, to understand the common risks, business impacts, risk mitigation options and benefits.
- The deliverable will include assessment findings and recommendations in a presentation.
- The assessment will be conducted on site for one day. The report will be built off site and presented to the sponsoring executive two weeks afterwards.

Assessment scope, objectives, deliverables, and project initiation documentation are described herein.

Objectives and expectations

The section provides details on the objectives and expectations for the workshop.

Risk Assessment

The objective is to use a distilled version of the standard IBM IT Business Continuity planning methodology to understand the business and IT risks, impacts, costs, benefits for two or three key applications, data or systems in the two data centers to help (*client name*) IT team build a Business Continuity strategy and execution plan. The outcome should also help (*client name*) team to decide whether or not to conduct a comprehensive risk assessment and Business Continuity plan of the total system and all key applications.

Business Continuity and Disaster Recovery Plan

This workshop is the *first step* to build a comprehensive Business Continuity and recovery plan. The comprehensive plan is *not* in the scope of this effort.

Data Center Operations

Part of the Next Step Business Continuity workshop can include a high-level examination the existing data center operations such as run book, monitoring, help desk and other daily operations procedures understand the risk involved in a Business Continuity situation, based on the service level targets. Writing the new procedures will not be part of this effort.

Scope

The scope of this engagement will be the Business Continuity initial assessment for two or three key applications in the present data center environment and some operations procedures. The activities can include interviews with (*client name*) subject matter experts (SMEs) in related areas and examination of appropriate documentations. A summary of these findings, recommendations of best practices, technologies and development of architecture and deployment roadmaps can be part of the outputs. The writing of the new procedures will not be included in this effort.

Approach

IBM or IBM Business Partner will collect available (*client name*) documentation relevant to the effort. IBM or IBM Business Partner will use the Best Practices standards of its own operations as well industry standards as the basis for the assessment. Based on the review of the documentation and detailed discussions with the (*client name*) technical team, IBM or IBM Business Partner will create a report *Finding/Observations* and *Recommendations*. The *Findings/Observations* report will identify gaps in compliance with Best Practices and will identify those steps that must be taken to support the vision of the (*client name*) Business Continuity strategy and plan.

Sample Methodology used for Assessment

The sample methodology that is used for assessment is determined as appropriate from the charts presented in this chapter.

Deliverables

Deliverables from the workshop include:

- Findings and recommendations for a risk assessment for two or three key applications, data and operation procedures in Business Continuity situations in the primary and backup data centers
- The final report will be a presentation (30-50 slides) on the results of the Business Continuity scenario risk assessment of key applications - business risks, impacts, costs, possible risk mitigation options, investments and benefits, key Business Continuity performance indicators, present status of meeting these indicators and gaps, recommendations to close these gaps and recommended roadmap, architectures and solutions.

IBM or IBM Business Partner Responsibilities

IBM/IBM Business Partner are responsible for the following project management activities under this SOW, including but not limited to:

- Establishing and maintaining an appropriate framework for communications with the (*client name*) Technical Coordinator
- Working cooperatively with the (*client name*) Technical Coordinator to resolve Project issues that might arise.
- Providing regular status reports (through e-mail or other means) as appropriate.
- Preparing the materials that are specified in the *Deliverables* section, together with the IBM or IBM Business Partner Project team, and providing such materials to (*client name*).

(*client name*) Responsibilities

(*client name*) is responsible for the following project management activities under this SOW:

- Serving as the interface between the IBM or IBM Business Partner project team and all participating (client name) departments and provide interviewers as appropriate time and location.
- Working cooperatively with the IBM or IBM Business Partner Technical Coordinator to resolve any Project issues that arise and escalating any issues within (*client name*), if necessary.
- Monitoring and reporting ongoing and final project status to (*client name*) management on a regular basis.

Workshop logistics

Logistics of the workshop include:

- The IBM or IBM Business Partner assessment team will conduct in a one or two day contiguous effort for data collection and interviews on site at (*client name*) then follow up with a presentation of assessment and recommendations.
- All materials provided by (*client name*) during the preparation for the assessment will be treated as client sensitive information.

- Supply the following information to the IBM technical staff for review prior to coming on site:
  - Any existing Business Continuity, Disaster Recovery, and Risk Assessment related documentation
  - Disaster Recovery business impacts indicators – revenue, costs, customer satisfaction, governmental regulations and compliance, and so forth.
  - Key Performance Indicators that might be in place to monitor and manage the project
  - Any data center documentation (procedures, applications, service levels, and so forth.)
    - Data Center configuration diagrams.
    - Servers, storage, networking and others.
    - Application workload and databases description (applications, DB, OS, SLAs).
  - For the identified scope:
    - Application and server/storage segmentation, priorities, desired recovery times. For the above identified priority workloads:
    - What is considered the Continuous Availability segment and what is that Recovery Time Objective? (RTO = elapsed time to recovery)
  - Management tools in use:
    - Application
    - Data
    - Server and storage
    - Network

Charges and Payments

There will be no charge to (*client name*) for this effort.

Key participants

Key participants in the workshop include:

- For (*client name*)

  CIO, Data Center Manager, Database Administrator, Server Manager, Storage Manager, Network Manager, line of business manager (as appropriate)

- IBM or IBM Business Partner

  Business Continuity Architect

- Technical Coordinators

  Each Party will appoint a Technical Coordinator, who will be the person to whom all communications required by this SOW will be addressed, and who has the authority to act on behalf of their respective company in regards to all aspects of this SOW. Either Party can change its designated Technical Coordinator by providing at least five (5) days notice to the other Party.

- IBM Technical Coordinator and Business Continuity Architect
  - (*name*)
  - (*email ID*)
  - (*phone*)
- (client name) Technical Coordinator
  - (*name*)
  - (*email ID*)
  - (*phone*)

Authorization

If this proposed SOW is acceptable to (*client name*), then your agreement constitutes a Document of Understanding (DOU) between IBM, IBM Business Partner, and (*client name*)

(*client name*) Authorizing Manager

Street Address: _____
City, State and Zip: _____
Phone: _____
Fax: _____

6.5 Appendix: Why Business Continuity

In this section, we provide background information as to why Business Continuity planning has often been viewed as difficult and describe how the Next Step Business Continuity workshop is designed to address those issues. It is presented from the perspective of the executive and IT management.

This information is useful if you need to describe why Business Continuity is important and also to set an proper expectation of what the Next Step Business Continuity workshop is meant (and is not meant) to accomplish. Figure 6-52 provides an executive overview to the workshop.

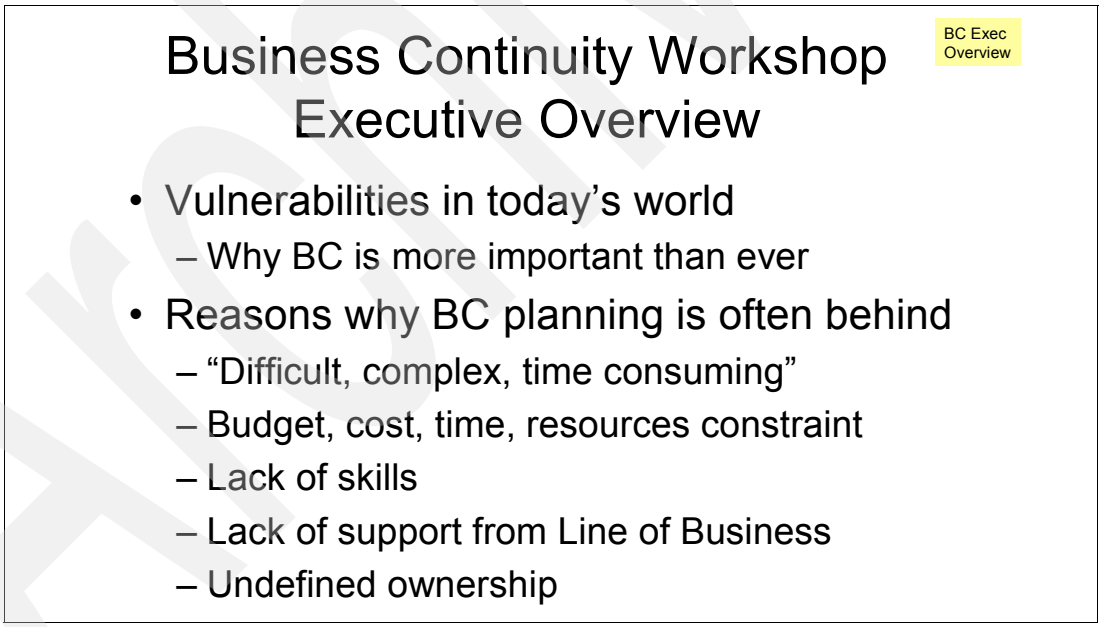

Figure 6-52 Executive summary and introduction to the workshop - 1

Begin by discussing the types of vulnerabilities that IT infrastructure is subject to in today's environment.

Vulnerabilities

It is no surprise to executives that Business Continuity is important today. Various increased threats, needs for higher uptime, needs to provide better service at a faster pace with high reliability, are all well discussed in the following chapters of this book:

- Chapter 2, “Industry Business Continuity trends and directions” on page 9
- Chapter 3, “Business Continuity planning, processes, and execution” on page 43

Reasons why Business Continuity planning is often behind

We next want to reassure executives that their past efforts in the area of Business Continuity have not necessarily been wasted. We also acknowledge the typical obstacles that cause Business Continuity planning to suffer and they are understandable (Figure 6-53).

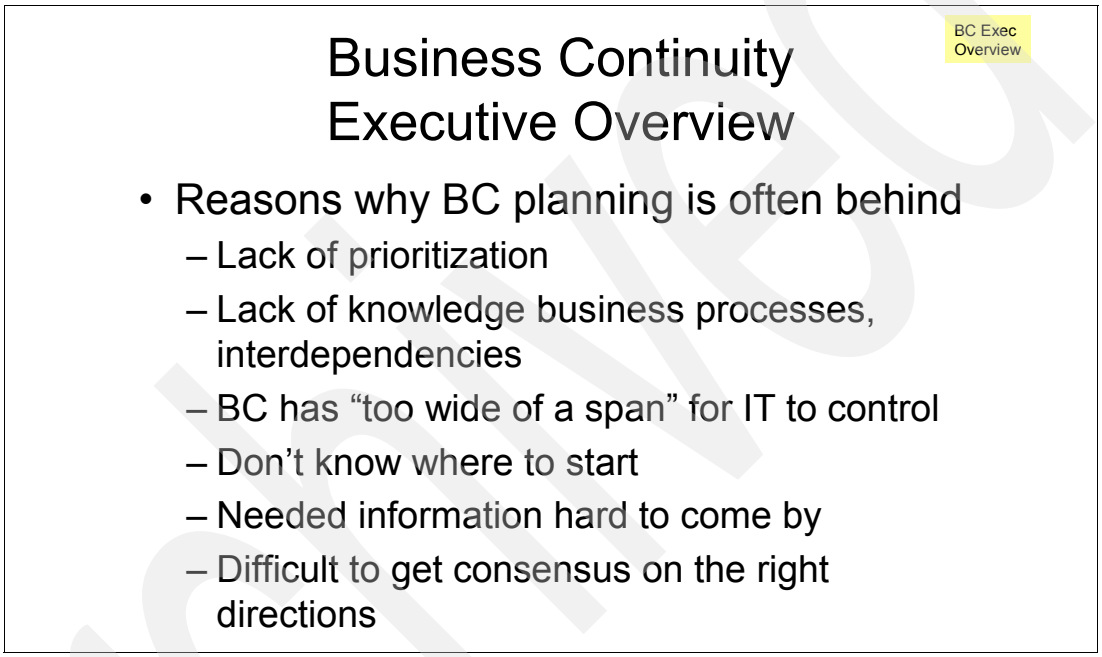

Figure 6-53 Executive summary - reasons why Business Continuity planning is often behind - 2

We list these obstacles specifically to point out what the Next Step Business Continuity workshop, supported by the information in this book, is designed address. The high-level baseline Business Continuity strategy that is the output of this workshop, considers and addresses not only technology issues but also the important organizational, budget, and resource issues.

We continue by listing additional typical obstacles.

For more information: Suggestions and further discussion for the obstacles listed in this chart are in Chapter 2, “Industry Business Continuity trends and directions” on page 9, as well as Chapter 3, “Business Continuity planning, processes, and execution” on page 43.

We next describe key aspects of our workshop methodology, which are designed to address the issues stated in the previous two charts (Figure 6-54).

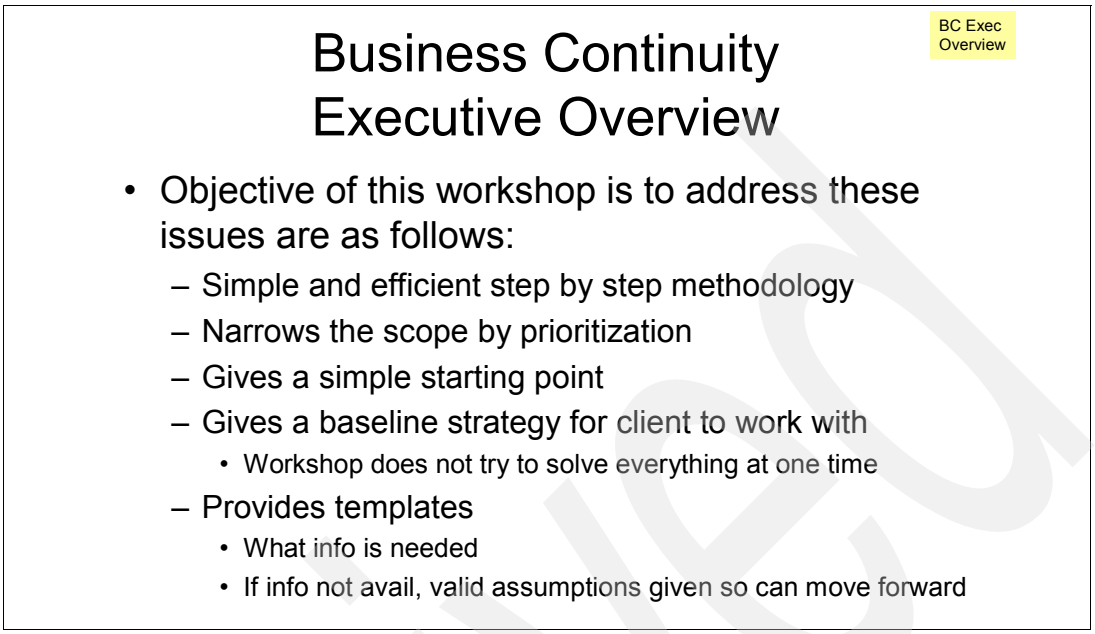

Figure 6-54 Executive summary - key workshop principles to address these issues

This workshop will:

- Use a proven step-by-step methodology, that is a distillation and simplification of the detailed Business Continuity planning process described in Chapter 3, “Business Continuity planning, processes, and execution” on page 43.
- Not try to solve everything at the same time. Rather, our intent is to narrow the scope of this workshop to just a few key business processes (one, two, or at most three). In this manner, we will establish a baseline Business Continuity strategy, skills, and methodology that he workshop audience can use, and that can be expanded upon.
- Provide fill-in-the-blank templates for important information, so that clarity exists on what information is to be gathered.

Important: We provide valid assumptions so that we can assure that we can move forward with the planning process even if key information is not available.

By narrowing this scope and using this methodology, the objectives are, as shown in Figure 6-55:

- Get started.
- Get skilled.
- Build a strategy, skill set, and methodology that is repeatable.
- Define what should be done in achievable chunks, that is a phased approach over time. This is essential to accommodate realistic constraints on resource availability, skills, and manpower.

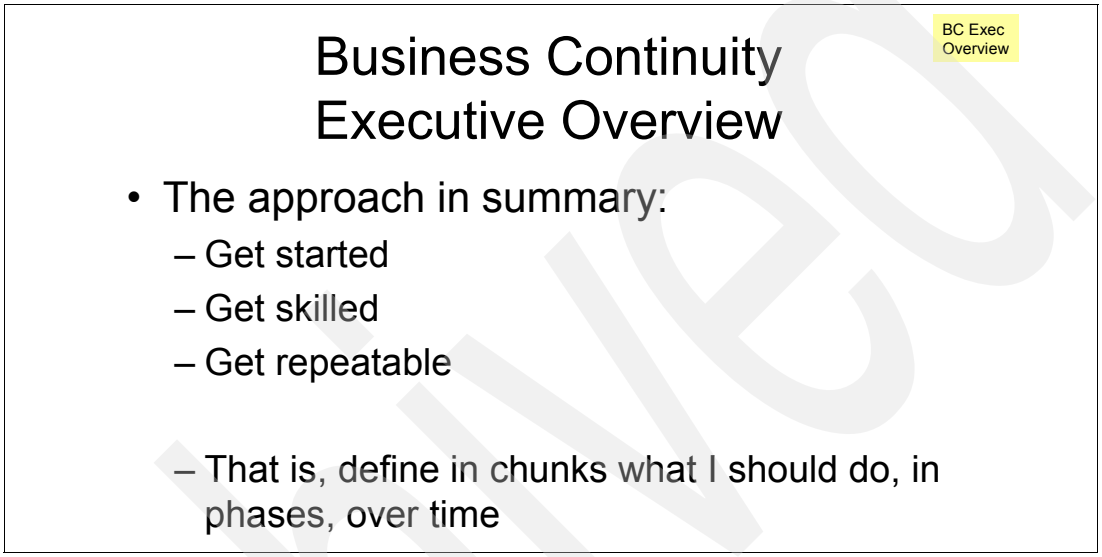

Figure 6-55 Executive summary - the approach

Figure 6-56 summarizes the executive overview.

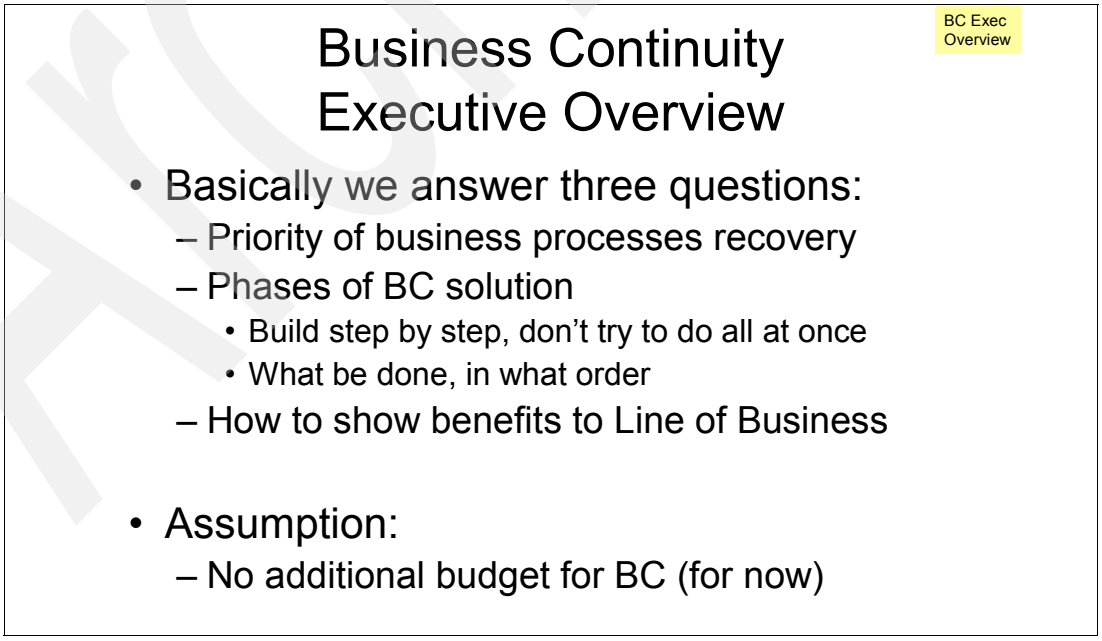

Figure 6-56 Executive summary - conclusion

You end the executive Business Continuity overview by summarizing three key questions that you aim to answer in the workshop:

- **Priority of business processes and recovery:** A distillation of the entire *Business Prioritization* portion of the ideal Business Continuity process that we describe in detail in 3.4, “Business prioritization” on page 47
- **Phases of a Business Continuity Solution:** A distillation of the entire *Integration into IT* portion of the ideal Business Continuity process that we describe in detail in 3.5, “Integration into IT” on page 67. (We again reiterate, that our workshop is not intended to do it all at once and that we will, as an output, derive what needs to be done next and in what order.)
- **How to show benefits to the Line of Business:** A distillation and an expansion of the entire “Manage” portion of the ideal Business Continuity process described in detail in 3.6, “Manage” on page 125. This is an essential task, as we know that a truly effective, cost-justifiable Business Continuity strategy must have joint ownership by both IT and the Line of Business.

We also make the following essential assumption:

- **No additional budget for now:** We assume in all stages of this planning workshop that no additional budget is available, which reflects the practical reality for the majority of our clients today.

This ends the Executive Overview.

6.6 Appendix: Solution visual aids

In this section, we provide some visual aids that can be useful as you discuss various aspects of solutioning.

Concept: a typical progression of a Business Continuity solution

The following set of charts give a visual explanation of how a typical IT Business Continuity solution project might progress, in phases.

You can use these charts and concepts to supplement the templates in the previous section. These discussion visual aids are intended to educate on the concepts of building the Business Continuity solution step-by-step.

To briefly review, we suggest to apply the following key Business Continuity solutioning principles:

1. Consolidate and simplify the IT infrastructure to provide a better managed, more cost-effective foundation upon which to build IT Business Continuity.
2. Use the business process Recovery Time Objective that has been defined and refined in this workshop, and apply your skill and leadership in solutioning to that RTO.
3. Build the end solution in phases. It is most affordable for a client not do it all at once.
4. In the process of building the plan in phases, each phase should build foundations for subsequent phases.

Example: To eventually implement end-to-end continuous availability, you first need to implement fast data replication. Prior to implementing data replication, you need to implement good data management policies and management software and probably need to uplevel the networking. Prior to implementing those prerequisites, you want to have a well-managed and consolidated tape backup and restore infrastructure.

An ideal Business Continuity solution would be built step-by-step, upwards from Backup/Restore, through Rapid Data Recovery, arriving at Continuous Availability.

Obviously, we rarely live in an ideal world. You should expect that you will need to overlap the building of the solutions for each of these business process segments.

Backup/Restore - Business Continuity Tiers 1, 2, 3

In this section, we provide visual aids for use during your solutioning discussion.

First, Figure 6-57 provides a visual for discussing solutioning for the Backup/Restore business process segment.

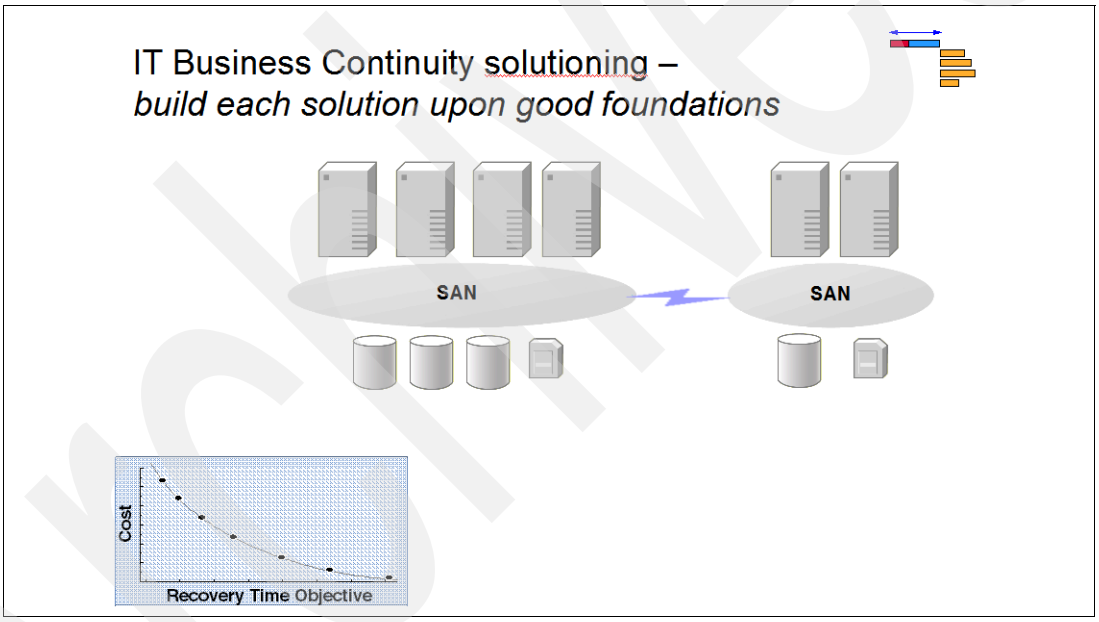

Figure 6-57 Step 4 - concept of building the solution step-by-step

Building an IT Business Continuity solution always depends on a good underlying foundations for the solution that you rae implementing. is to do it step-by-step, building upward from a well managed, consolidated foundation. We start as shown in Figure 6-58.

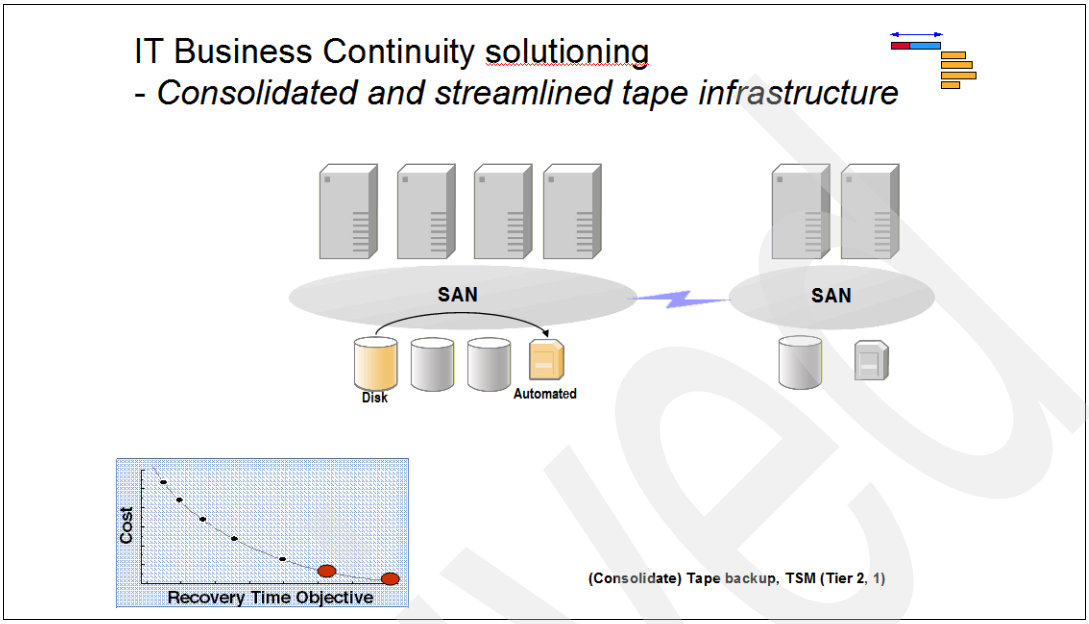

Figure 6-58 BC Tiers 1 and 2 - with Backup/Restore foundation and tape consolidation

As shown in Figure 6-58, at Business Continuity Tiers 1 and 2, we implement a streamlined, consolidated, and simplified tape infrastructure, to provide a consistent, cost-effective tape backup and restore foundation. We add storage management software to manage data and the consolidated tape, thus initiating process of exploiting automation.

A logical next step might be remote tape vaulting, thus saving physical transport time. As shown in Figure 6-59, at Business Continuity Tier 3 we show adding remote tape vaulting, and also adding network bandwidth and network control tools.

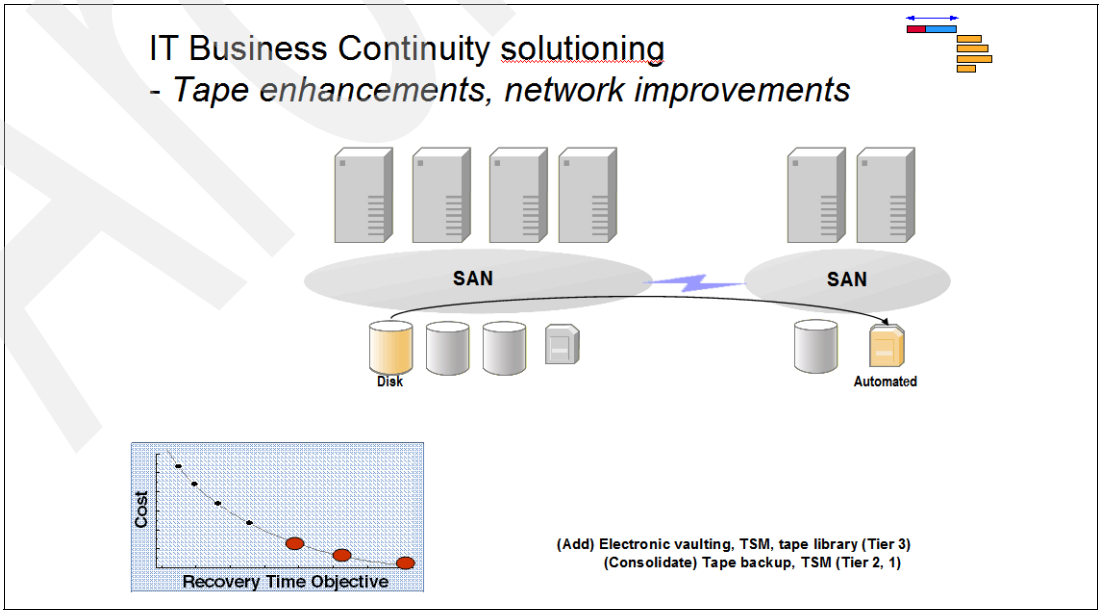

Figure 6-59 BC Tier 3 - add remote tape vaulting and appropriate management software

Adding networking at this earlier stage is an especially good strategy, if you are also preparing for subsequent Rapid Data Recovery solutions, where even more rigorous data replication is to be done. Rigorous data replication will require good network management and control.

Using remote tape vaulting as an initial production usage of the network is an excellent way to initially place a relatively less mission-critical application in the network, while we are learning to manage and control that network. Simultaneously, we can further improve the tape recovery speed, by removing physical tape transport time.

Rapid Data Recovery Point in Time Copy - Business Continuity Tier 4

Figure 6-60 is a visual aid for use when discussing Point in Time disk copy solutions for the Rapid Data Recovery business process segment.

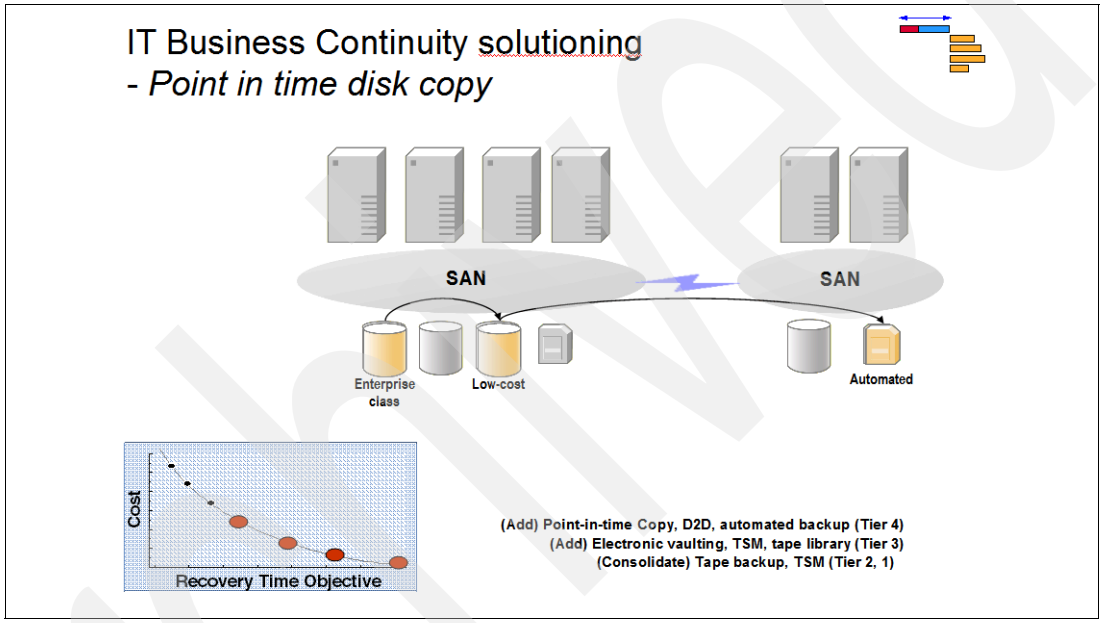

Figure 6-60 BC Tier 4 - add primary site point in time disk copy to the remote tape vaulting

We add point in time disk copy capabilities to provide near continuous *application* availability, by avoiding application outages while making the backup on disk.

This point in time disk capability would provide additional synergistic capability to any previous (or simultaneous) implementations of remote tape vaulting, tape library, and tape automation solutions.

Software-based Application / Database Rapid Data Recovery - Business Continuity Tier 5

As part of solutioning, you need to decide what portions of the Rapid Data Recovery data replication will be done by the application or database software or by server or storage replication. One method is not necessarily better than the other, and there are many specific cases where software or database replication might be the better method.

You can find more information about selecting between these methods of data replication in Chapter 3, “Business Continuity planning, processes, and execution” on page 43. Figure 6-61 is a solutioning visual aid chart on software and database replication and positioning it through the typical Recovery Time Objective achieved.

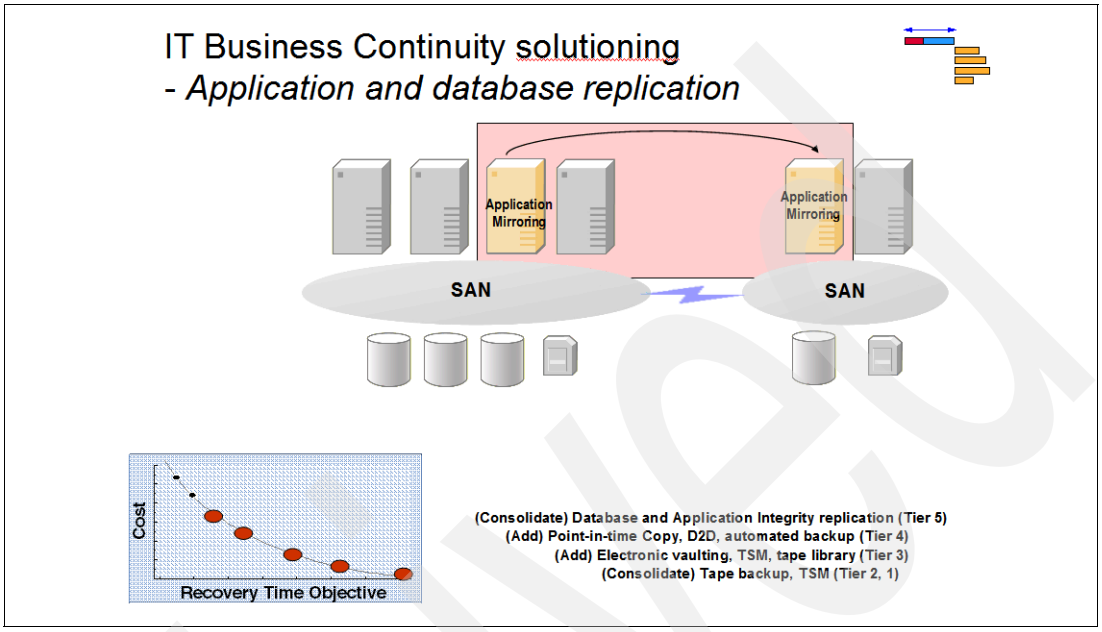

Figure 6-61 BC Tier 5 - segment out applications that use database and application level replication

Storage mirroring Rapid Data Recovery - Business Continuity Tier 6

Figure 6-62 is a solutioning visual aid chart for storage replication.

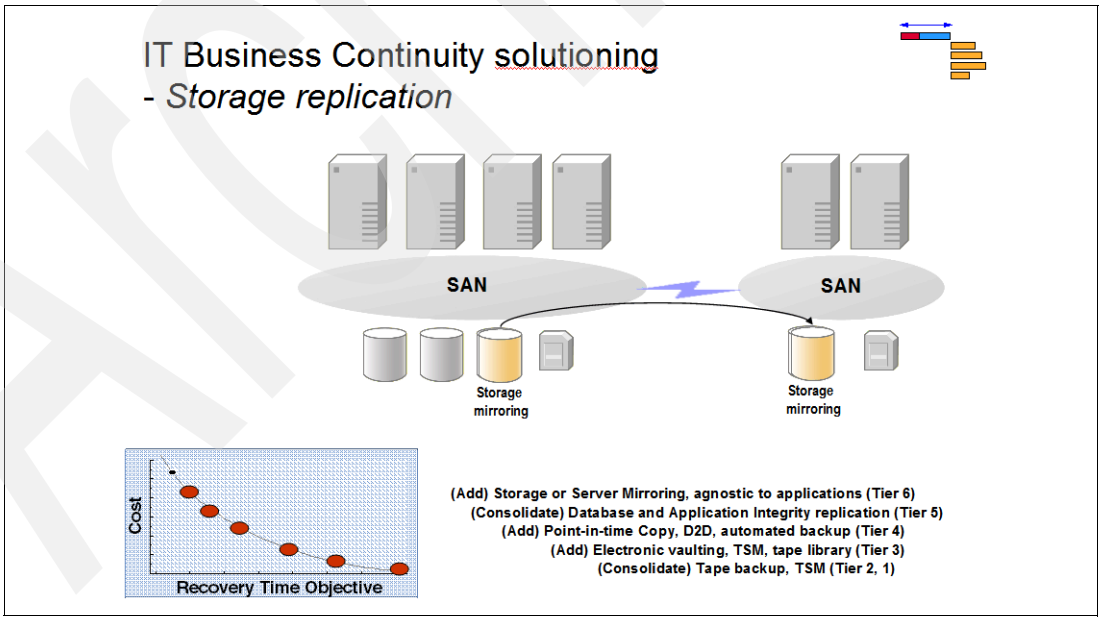

Figure 6-62 BC Tier 6 - add storage or server mirroring

At Business Continuity Tier 6, to the existing storage devices, we add storage mirroring along with management software for that storage mirroring. We can use the network that has been previously tested in Business Continuity Tiers 3 and 4.

Continuous Availability - Business Continuity Tier 7

At Business Continuity Tier 7, we add continuous availability operating system integration upon all the foundations laid by each of the preceding steps (Figure 6-63). Continuous Availability solutions require the previous Business Continuity Tiers to have been implemented; Continuous Availability solutions are fundamentally the automated single point of control for multiple Business Continuity technologies.

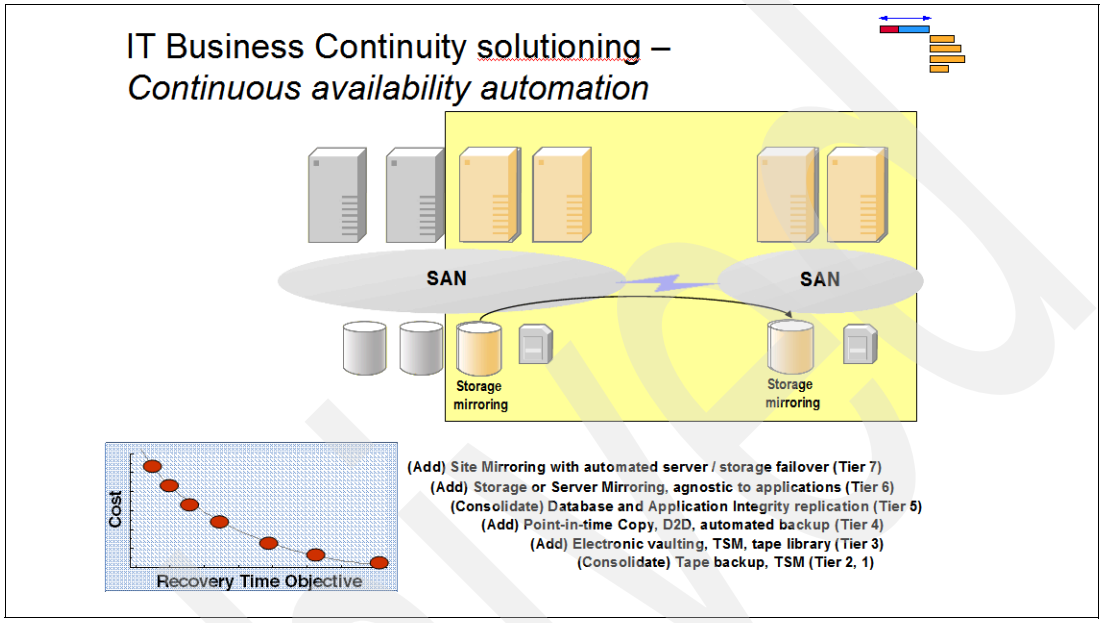

Figure 6-63 BC Tier 7 - add operating system integration for end-to-end automated failover

7

Next Step Business Continuity workshop: Case Study

In this chapter we provide an applied case study of the Next Step Business Continuity workshop that we described in Chapter 6, “The Next Step Business Continuity workshop” on page 173. The intent is show how the concepts of the Next Step Business Continuity workshop are applied in practice and to provide an example of what your own workshop could produce for you.

7.1 Introduction to the Case Study

In this chapter, we provide an applied case study which shows a practical example of executing all of the concepts described in the Next Step Business Continuity workshop. Our case study is for a large Medical Center client.

7.1.1 Motivation for the Next Step Business Continuity workshop

After learning about the catastrophic effects of a natural disaster on medical facilities in an industry conference, the senior management team of a university asked the CIO to prepare a Business Continuity plan for their medical center, with emphasis on risk assessment, budgets and project plan for the next three years.

Although there is some equipment redundancy and regular data backup for the medical facilities, there is no formal Business Continuity and disaster recovery plan in place. Each of the key systems is maintained and operated by different IT teams and is fairly independent from one another. They do share a common network infrastructure in the data center.

Because the mainframe system director has the longest tenure in the data center, he was asked to be the leader to produce the plan. IBM was requested to assist in the planning process for the IBM equipment. The customer team accepted the IBM recommendation of a workshop facilitated by Business Continuity specialists.

7.1.2 Client background

The large Medical Center has approximately 200 patient beds with five separate buildings on a 50 acre campus with over 1000 employees. It is located near the ocean and can be susceptible to floods and hurricanes. The executive management team has always had concerns about the staff readiness during a disaster. Learning about the effects of Hurricane Katrina on New Orleans hospitals and of their sufferings and liabilities, the executive team decided to develop a plan to address any situations.

Each of the departments in the hospital is responsible for their patients and medical staff in terms of facilities, equipment and IT applications. The IT department provides the IT infrastructure, the data center and daily operations which involve application and data backup. The department heads decide on the systems and information needed. It is not clear which hospital department is responsible for the overall Business Continuity planning.

The hospital is growing rapidly and their demand for IT applications and systems is growing along with the business. Similar to other hospitals, with increasing regulations and limited reimbursement, the hospital uses digital technologies, such as digital imaging and wireless networks, to improve patient care and staff efficiency. The applications and systems were selected by the individual department users. Consequently there is a huge amount of disparate IT equipment in the data center and other locations. Also, the volume of data collected is growing exponentially, making data backup very time consuming and impacting availability. There are business needs to connect the different department systems together for information sharing, sometimes in batch mode and at other times, real time.

7.1.3 Business challenges and issues

The Medical Center is facing the following challenges similar to others in the health industry:

- Reduce overall costs without sacrificing patient care quality
- Comply with government regulations to ensure reimbursement
- Keep medical staff efficient and satisfied with services

The hospital is investing heavily in technologies, such as digitization of patient care information and mobile diagnostics to improve patient care and lower costs. Each patient care or business process now involves several systems and infrastructure, making Business Continuity much more complex and requiring coordination among departments. At the same time, there is little budget allocated for cross systems projects, such as Business Continuity and disaster recovery as most of the system acquisitions are funded by the individual departments.

7.1.4 Technical challenges and issues

To keep patient care and business processes efficient and simple, the hospital is working on interconnecting their systems and facilitating data exchange among them using middleware software. From a Business Continuity perspective, the technical challenges are:

- Because each of the systems administrators is only responsible for their systems, cross system issues are being neglected. In this case, Business Continuity is only planned on an individual level and if disaster strikes, it is not clear if the hospital can continue to operate with increasing cross system interconnection and data exchange. Because the budgets are limited, Business Continuity requirements are usually low on the priority list.
- Every system administrator has their own management tools and Business Continuity design and plan. It is necessary to consolidate them to reduce complexity and cost. There are plans for a backup data center site. It has been estimated that it is financially impossible to fund redundant systems for each department in the backup site.

7.2 Workshop preparation

To ensure that they meet the customer expectations, the IBM team set up a call with the Hospital Business Continuity project team to discuss their requirements, expectations of the workshop, and outputs.

7.2.1 Client workshop expectations, scope, and desired outputs

Overall, the Medical Center team expressed the following expectations for the workshop:

- Review the Business Continuity planning methodology and medical industry best practices
- Review and evaluate different technology and solution options
- Develop recommendations on the solutions, IT architecture and overall Business Continuity plan
- The scope should be IT Business Continuity related issues, including the primary and backup data centers, all key systems, applications, data and tools
- The outputs should be recommendations of the technologies, solutions, tools, architecture, Business Continuity base plan and roadmap

The IBM team asked the Medical Center project team to send any relevant information:

- Existing Business Continuity plans and configurations
- Systems and applications details
- Data center configurations and diagrams, including networks, servers, storage and related equipment
- Service level targets
- Backup procedures and tools
- Key patient care and business processes and systems flow
- Any questions and issues to be addressed during the workshop

7.3 Results from the Next Step Business Continuity workshop

This section discusses the results from the actual day of the workshop.

7.3.1 Workshop agenda, desired participants, and information

The following agenda was accepted by the Project team:

- Team and workshop introduction – objectives, scope, output (15 minutes – IBM leading)
- Business Continuity planning methodology discussion (30 minutes – IBM leading)
- Existing Business Continuity plan and systems review (30 minutes – Hospital team leading)
- Risk assessment and prioritization with key patient care and business processes (60 minutes – IBM leading)
- Business Continuity objectives and success criteria (30 minutes – IBM leading)
- Existing Business Continuity systems and applications readiness assessments (60 minutes – IBM leading)
- Solution and architecture vision (30 minutes – IBM leading)
- Technologies and solution option evaluation and architecture discussion (60 minutes – IBM leading)
- Recommendations, roadmap and action plans (60 minutes – IBM leading)

It was recommended that key decision makers and information providers should attend the six hour session. Critical information should be gathered before the workshop. When information is not readily available, IBM experts will supply assumptions based on industry best practices and experiences to move the workshop forward. The decision makers and information providers should evaluate these assumptions and adjust according to their environments and needs. The assumptions can be adjusted over time.

7.3.2 Assumptions for this Next Step Business Continuity workshop

Before the workshop and after the collection of the individual system Business Continuity plans, the Medical Center came to these conclusions:

- The budget and configurations of the individual systems are in place and ready for the Business Continuity plan.
- Definitions and scopes of the individual systems vary. Some “good” Business Continuity practices, skills can be leveraged across the teams. Pooling of resources will accelerate implementation, success and lower costs.
- Because there is interdependency among the systems and data, the Business Continuity of a patient and business process depends on the availability of several systems.

At the beginning of the workshop, these scope and objectives were agreed upon:

- The “new” Business Continuity solution and architecture should address the patient care and business processes and should not stop at the individual system level.
- Business Continuity service tiers, standards and policies should be in place to simplify the design, management and recovery process and also minimize the overall costs. Each patient care and business process should fit into these tiers, standards and policies based on business priorities and budgets.
- All solutions should conform to a common set of architecture and technologies for easy deployment and management.
- The recommendation of the workshop can be the base standards, policies and architecture for the overall Business Continuity plan.
- Building the Business Continuity plan for the hospital will take some time and the implementation should be built in phases based on priority.

With these base assumptions, the steps of the workshop conducted as follows:

7.3.3 Collect information of the key patient care and business processes for prioritization and Business Continuity plan criteria

There were over 50 patient care and business processes considered as important in the Medical Center. The three processes, shown in Figure 7-2, were discussed in the workshop. These are patient records, radiology, and e-mail. Other processes will be evaluated at a later time with the Line of Business (LOB) teams.

7.3.4 Key business process and related IT components

Key Performance Indicators (KPIs) were not available or fully defined yet by this client. A partially filled in template, using the suggested starter set KPIs, was used, as shown in Figure 7-1.

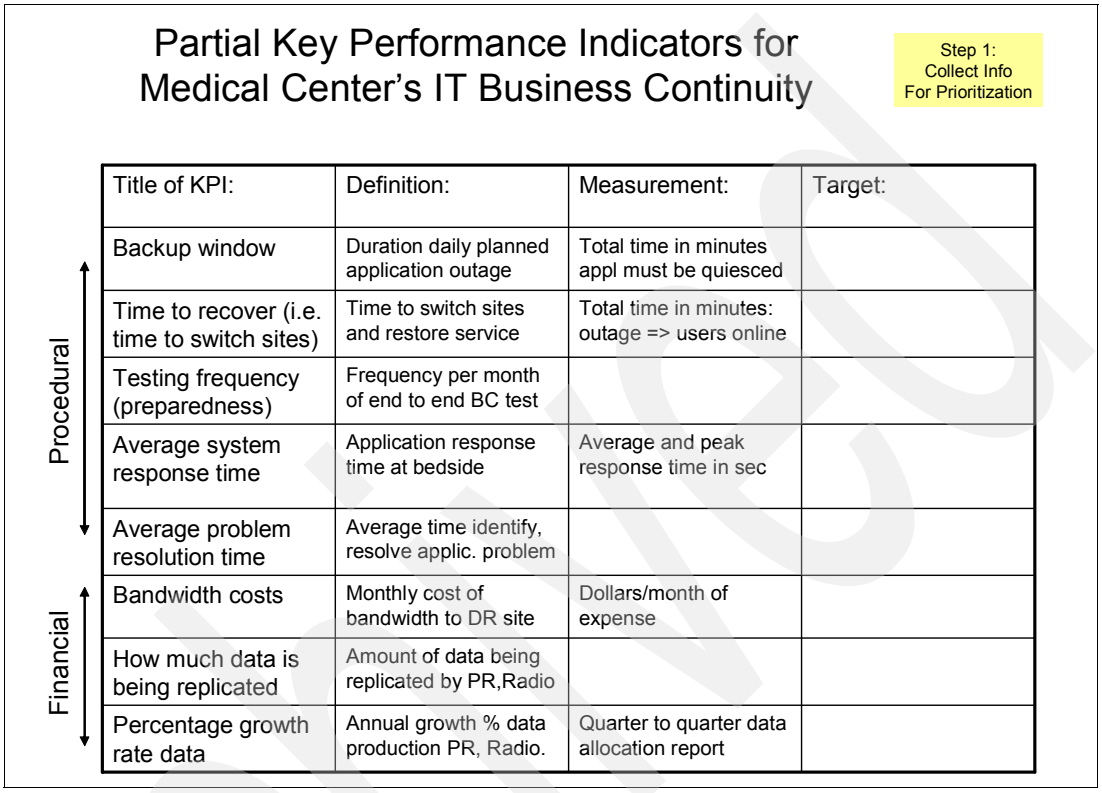

Title of KPI:	Definition:	Measurement:	Target:
Backup window	Duration daily planned application outage	Total time in minutes appl must be quiesced	
Time to recover (i.e. time to switch sites)	Time to switch sites and restore service	Total time in minutes: outage => users online	
Testing frequency (preparedness)	Frequency per month of end to end BC test		
Average system response time	Application response time at bedside	Average and peak response time in sec	
Average problem resolution time	Average time identify, resolve applic. problem		
Bandwidth costs	Monthly cost of bandwidth to DR site	Dollars/month of expense	
How much data is being replicated	Amount of data being replicated by PR,Radio		
Percentage growth rate data	Annual growth % data production PR, Radio.	Quarter to quarter data allocation report	

Figure 7-1 Medical center - partial filled in Key Performance Indicators

Figure 7-2 is the filled-in template for Scope - Resource and Business Impact.

Scope - Resource and Business Impact

Step 1:
Collect Info
For Prioritization

Component	Patient Records	Radiology	E-mail
Applications	PR1	EPIC1, PR1	MSE1, MSE2
Data and data management policies	Backup daily at night by TSM	Backup daily by Veritas	Backup daily by MSM
Databases	DB2	Oracle, DB2	MSSQL1
Servers	Mainframe1	UNIX1, Mainframe1	MS222, MS223, MS224
Storage	33801, 33802	SUN101	MSS1, MSS2
Network (LAN, WAN)	SLAN1 and SLAN2	SLAN2 and SLAN3	LAN2, LAN3
Procedures and tools	Data Center M1	Data Center M1, M2	Data Center M3
People	J1, J2	M1, J1	K1, K2
Facilities and physical location	DC1	DC1	DC1
Known vulnerabilities	Flood	Flood	Flood, servers
Estimated cost of outage , per hour	$200K*	$150K*	$20K*
Business Impact	High	High	Medium
Ranking	1	2	3

Figure 7-2 Medical Center Business Continuity Scope - Resources and Business Impact

Note: The values were calculated based on Loss of Business and lost productivity of the personnel (such as 500 staff losing 30 minutes or 5% productivity for the day) – their daily costs times 5%.

Note: Many of the cells in Figure 7-2 point to other documentation that is prepared either by the customer or the organization that conducts the workshop. For example, the entries in the Procedures and Tools row will point to appropriate system documentation that the client provides.

7.3.5 Key application and IT components

Figure 7-3 provides the key application and IT components impacts on key process to provide priority for the roadmap.

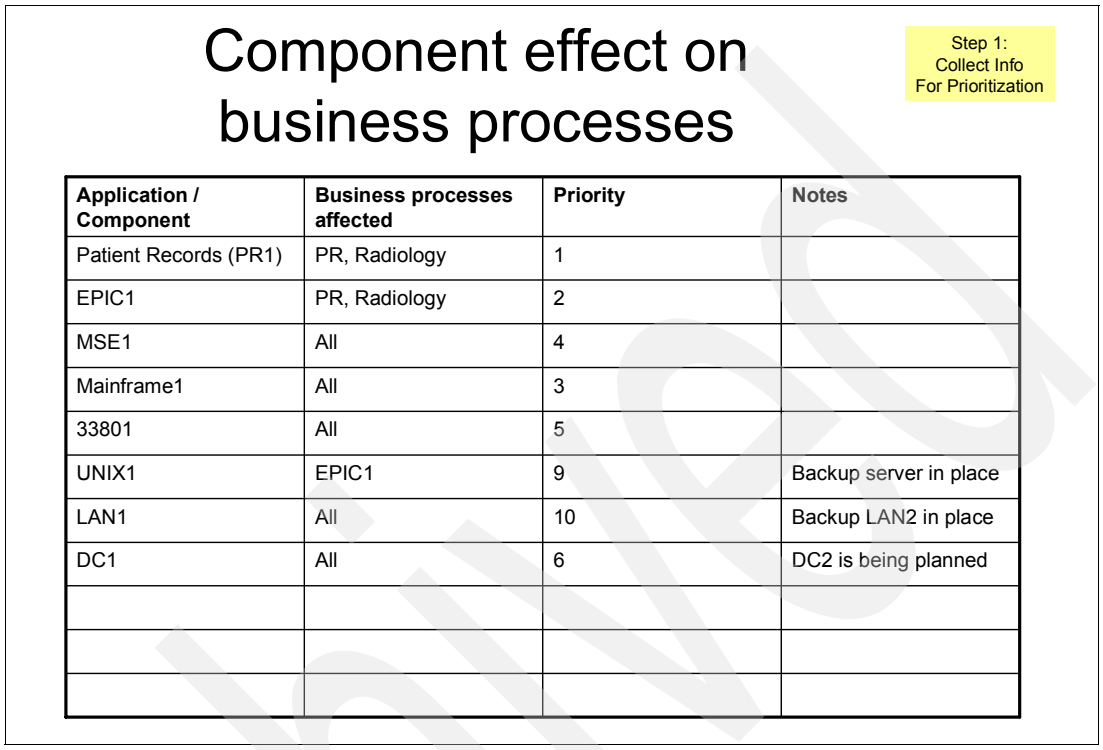

Application / Component	Business processes affected	Priority	Notes
Patient Records (PR1)	PR, Radiology	1	
EPIC1	PR, Radiology	2	
MSE1	All	4	
Mainframe1	All	3	
33801	All	5	
UNIX1	EPIC1	9	Backup server in place
LAN1	All	10	Backup LAN2 in place
DC1	All	6	DC2 is being planned

Figure 7-3 Medical Center case study- IT applications and components effect on business processes

We take the information in these two sections and proceed to Risk Impact and Business Impact Analysis.

7.3.6 Risk Impact and Business Impact Analysis Priority

Figure 7-4 is the completed template for our Risk Impact and Business Impact Analysis. We used three levels—high, medium and low, indicating the relative priority.

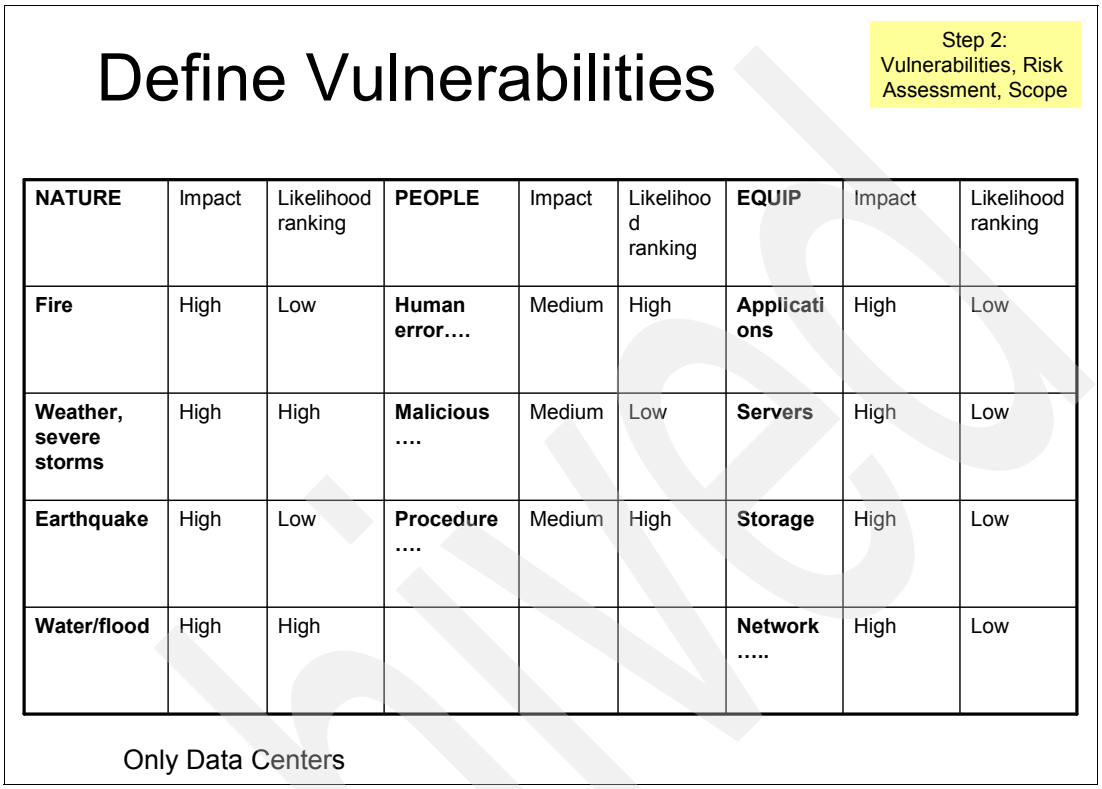

Define Vulnerabilities

Step 2: Vulnerabilities, Risk Assessment, Scope

NATURE	Impact	Likelihood ranking	PEOPLE	Impact	Likelihood ranking	EQUIP	Impact	Likelihood ranking
Fire	High	Low	**Human error….**	Medium	High	**Applications**	High	Low
Weather, severe storms	High	High	**Malicious ….**	Medium	Low	**Servers**	High	Low
Earthquake	High	Low	**Procedure ….**	Medium	High	**Storage**	High	Low
Water/flood	High	High				**Network …..**	High	Low

Only Data Centers

Figure 7-4 Medical Center case study- define vulnerabilities

The conclusion in terms of priority and immediate focus, is weather, floods, and procedures.

7.3.7 Key Business Process Business Continuity targets – current assessment and desired targets

As a result, we defined the key Business Process Business Continuity targets shown in Figure 7-5.

Define BC target

Step 3: Define BC targets based on Scope

Business Process	Current Recovery Time:	Success rate %	Budget spent to achieve this	Cost of outage / hour (total $cost avoided)	Target desired Recovey Time	Desired success rate %	Projected cost avoidance:	Projected budget
Patient Records	4 hours	???	$100K	$200K	30 min	100%	???	???
Radiology	8 hours	???	$40K	$150K	15 min	100%	???	???
E-mail	30 min	good	$10K	$20K	15 min	100%	???	???

Figure 7-5 Medical Center case study - defined Business Continuity targets

7.3.8 Conclusions

Here are the resulting conclusions from the previous documentation work:

- **High priority business processes:**
  - Patient Records
  - Radiology
  - Email
- **High priority applications, systems and IT components**
  - Patient Records 1 (PR1)
  - Mainframe1
  - EPIC1
- **Highest vulnerability and risk**
  - Hurricane
  - Flood
  - Data center procedures and human errors

7.3.9 Defined Business Continuity service targets and priorities by business process

Figure 7-6 shows the defined standards and policies into which other business processes fit.

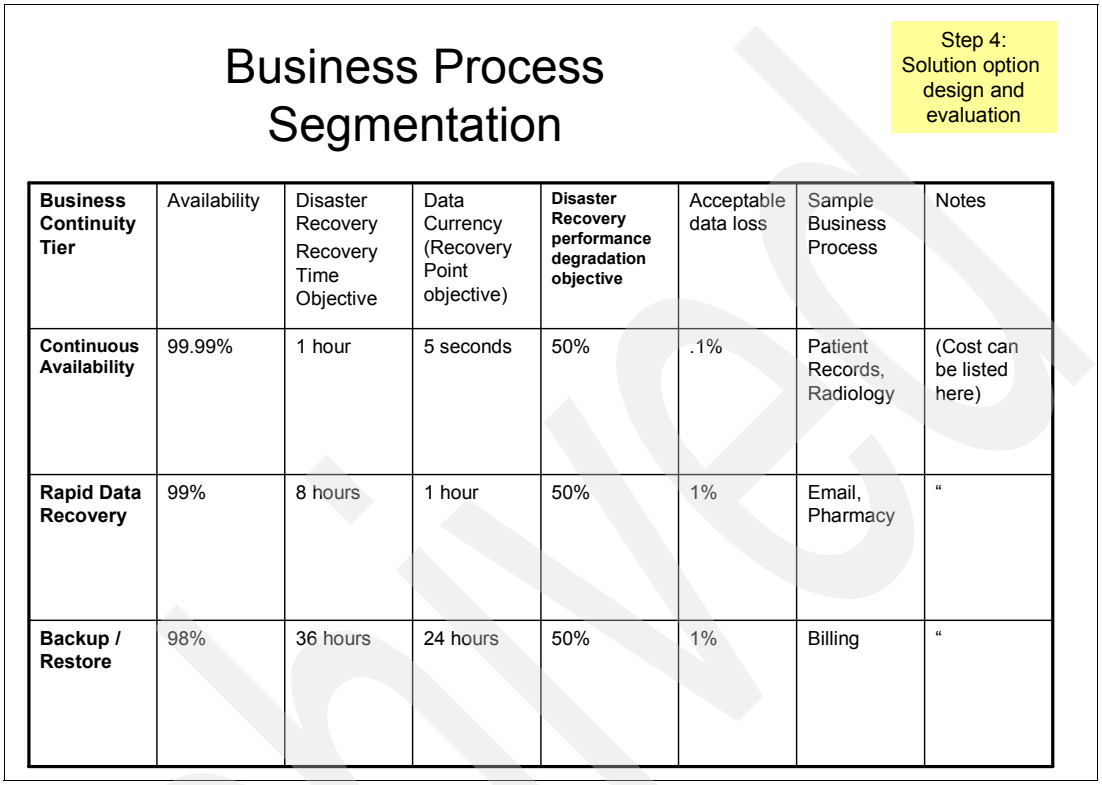

Business Continuity Tier	Availability	Disaster Recovery Recovery Time Objective	Data Currency (Recovery Point objective)	**Disaster Recovery performance degradation objective**	Acceptable data loss	Sample Business Process	Notes
Continuous Availability	99.99%	1 hour	5 seconds	50%	.1%	Patient Records, Radiology	(Cost can be listed here)
Rapid Data Recovery	99%	8 hours	1 hour	50%	1%	Email, Pharmacy	"
Backup / Restore	98%	36 hours	24 hours	50%	1%	Billing	"

Figure 7-6 Medical center case study - business process segmentation and solution RTO criteria

This key information now allowed us to proceed with designing a baseline Business Continuity architecture and solution design.

7.3.10 Defined baseline Business Continuity and IT architecture and design criteria

Using the business process segmentation described in Figure 7-6, we defined the following Business Continuity solution and architecture criteria:

- Backup Data Center is a must to address the flooding and weather vulnerability.
- Tiered storage should be a major component for standardization and simplification.
- Hardware virtualization should be considered if the applications are supported and service levels are acceptable.
- Operations and management tools need to be standardized and centralized if possible to minimize skills and personnel requirement during recovery. Increasing automation will also reduce complexity, human errors, and improve recovery time.
- Data center and Business Continuity procedures would need to be reviewed and rewritten for a comprehensive plan for Disaster Recovery and coordination with LOB activities.
- Budgetary considerations.

- Sharing hardware is acceptable at the backup site to lower overall cost of ownership but must be logically separate.
- The networking requirements (technology used and bandwidth requirements) can be a huge expenses depending on the distance, timeliness and volume of the data transmitted between the two data centers.

Other considerations

We also identified that the individual systems will continue to be independent but will consider consolidating some common components, such as DNS, HTTS server, load balancers in future design.

7.3.11 Key existing systems and Business Continuity configurations

We documented the following key systems in the existing data center, shown in Figure 7-7:

- 1 IBM System z mainframe
- 1 EMC 8530 attached to mainframe
- 2 IBM System p servers
- 2 HP UNIX servers
- 4 HP Alpha servers
- 110+ Intel® Window servers (various vendors) with close to 40 TB of storage from different vendors

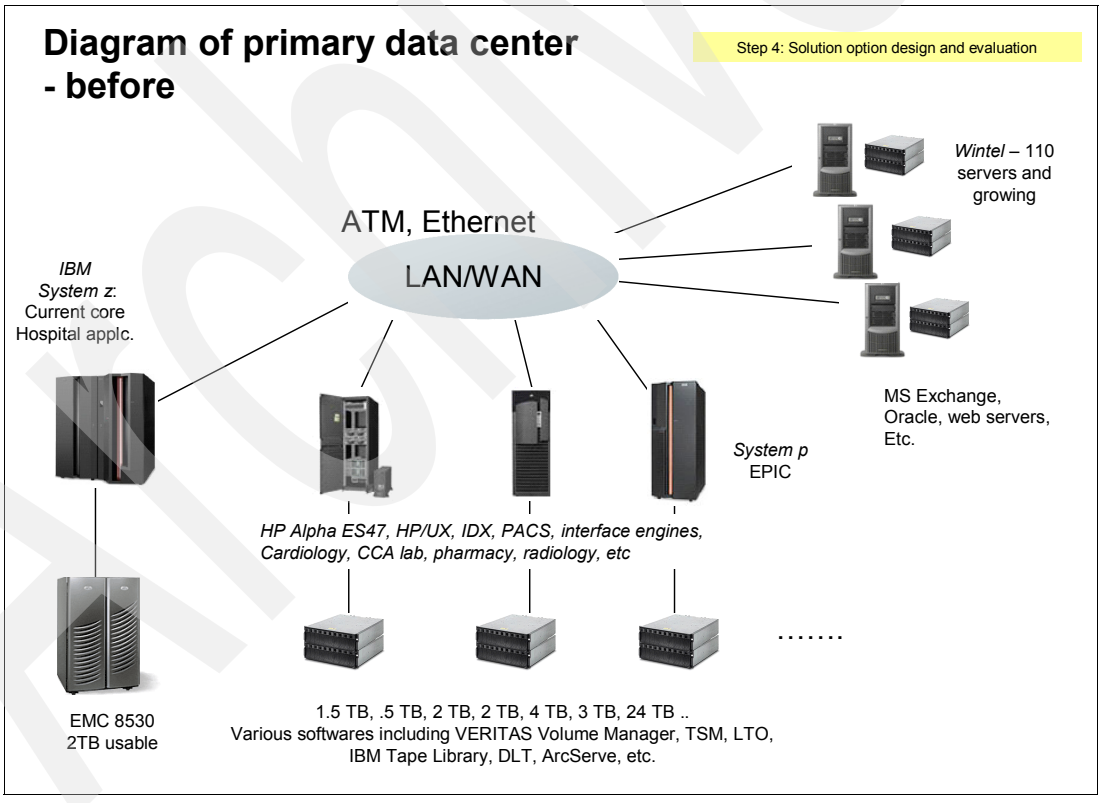

Figure 7-7 Medical center case study - key existing systems and Business Continuity configurations

7.4 Recommended Business Continuity architecture and configurations

Here is the recommended solution.

Design criteria

- Supports most existing IT investment and applications and leverage existing skills and expertise
- Drives the three Business Continuity tiers or segments – Continuous Availability, Rapid Data Recovery and Backup/Restore
- Fits in existing IT strategy and budgets

Technologies evaluated

- Virtualization – server (VMWare, MS Virtual Server) and storage (IBM SVC)

Consolidation

- Server pooling using virtualization

Tiered storage approach

Recommended using Information Life Cycle Management (ILM) tool such as TotalStorage Productivity Center product for automation through the Tiers using policies:

- Storage Tier 1 – High performance FC drives
- Storage Tier 2 – Medium performance FC drives
- Storage Tier 3 – SATA drives
- Storage Tier 4 – Tape Library

Backup Data Center priorities based on the risk assessment:

- Similar Mainframe installed at the backup data center using some of the main site mainframe upgrade budgets and sharing loads between the two. The backup mainframe is a hot site.
- Because the current EMC storage system required an upgrade any way, it was decided to move to an IBM storage (DS8100) and have a similar architecture product (at a lower costs) DS6800 at the back up site to take advantage of the lower costs and the data traffic transmission savings in terms of network bandwidth.
- For existing systems which already have a redundant system for backup such as HPUX and Alpha, SUN and IBM UNIX servers, the redundant backup systems would be moved to the backup site.
- For major Intel servers, some of the redundant backup servers will also move to the backup sites.
- For the Intel servers without redundant backups, there would be 4-8 ways Intel servers running VMWare to support a virtual redundant environment plus ad hoc request for additional capacity.
- For the storage, virtualization using IBM SVC and TotalStorage Productivity Center will support both the main and backup sites to improve data replication and lower network costs to meet the service levels and tier architecture.

Figure 7-8 shows the recommended solution for the backup data center.

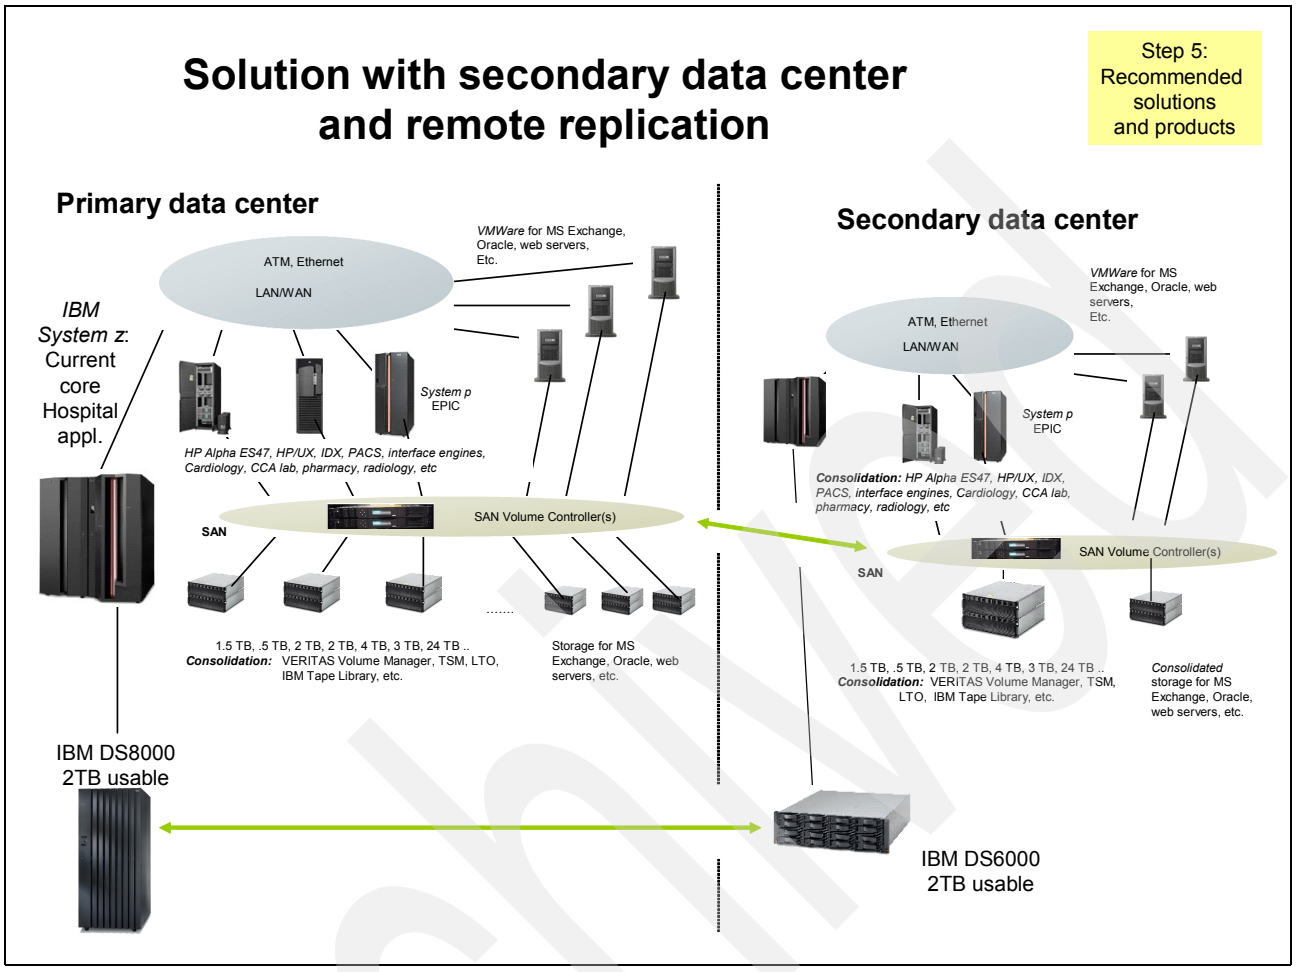

Figure 7-8 Medical center case study - recommended solution for backup data center

Benefits

- Maximize existing IT investment and skills on the UNIX systems, meeting Business Continuity and budget objectives
- Consolidate the appropriate components, such as storage and network transmission to reduce costs and simplify operations
- Leverage normal systems upgrade and refresh costs to support Business Continuity objectives without huge additional investment
- Improve existing asset utilization and operation by automation to lower overall cost of ownership
- Meet Business Continuity financial targets (minimum additional outlay) by cost savings through improved efficiency of existing operations (network and data center), and normal upgrade budgets
- Improve existing IT operations and lower overall costs of ownership

Challenges

- Locating a backup data center that is close enough to minimize network costs and far enough to avoid weather and flood impacts. It was found five miles away on higher ground.
- Seeking LOB support on a physically sharing IT environment over security and funding concerns. It took some explanation and benefits analysis and it was acceptable for some users.
- Analysis and prioritization of the rest of patient care and business processes took a lot of time because of various reasons, mostly politically (few department would put their processes in lower priority). The base architecture and design support all these processes, it was a matter where they would go into the Business Continuity tiers.

7.4.1 Financial implications and justification

Overall, the financial funding for the Business Continuity solutions came from the following sources:

- 20% from the costs savings of the simplification of the existing data center hardware and server and storage consolidation (maintenance and operation cost reduction)
- 40% from the usual acquisition of the redundant systems (as part of the usual practice of buying backup systems)
- 10% from network equipment and bandwidth savings (not having to pay additional assets and services)
- 10% from operations and staff efficiency improvement (not having to hire additional staff and tools)
- 20% from new funding from corporate for overall Business Continuity programs

The Lines of Business and IT management agreed on these justifications and the associated financial commitment, as the hospital would improve their IT Business Continuity readiness by demonstrating the ability to continue operations of the key patient care and business processes in two hours, after the primary data center is out of action.

7.4.2 Implementation planning

Here are the steps the IT project team took for the IT Business Continuity implementation plan:

- Simplify, consolidate, standardize and centralize infrastructure

  The less number of servers, storage and network equipment footprints are, the less number of application instances and operating systems to be supported, the less complexity and easier to backup and managed; technologies, such as server and storage virtualization, clustering and centralization (such as SAN and NAS) are important to support the service level goals; it will help shorten backup window and application maintenance downtime

- Evaluate existing data center Business Continuity related procedures and tools

  The existing procedures were fairly fragmented under each system and NOT processes. For the Business Continuity plan, new procedures were developed to reflect business process recovery to support the new solutions and tools.

- Evaluate existing personnel training and skills to support new plan

  The new Business Continuity plan was based from a total system standpoint and new processes, tools training was provided for the personnel.

- Evaluate existing hardware and tools to fit into new architecture

  At a later stage, for simplification and cost savings, the existing hardware and tools were evaluated to see if they could be incorporated or eliminated.

Figure 7-9 shows some proposed Key Performance Indicators for this client.

Suggested Key Performance Indicators for Medical Center's IT Business Continuity

Step 6: Recommended Strategy and Roadmap

	Title of KPI:	Definition:	Measurement:	Target:
Procedural	Backup window	Duration daily planned application outage= 20 min	Total time in minutes appl must be quiesced	Reduce to from 20 min to 1 min by 1Q2007
	Time to recover (i.e. time to switch sites)	Time to switch sites and restore service	Total time in minutes: outage => users online	Reduce to 30 minutes by YE2007
	Testing frequency (preparedness)	Frequency per month of end to end BC test	Number of times/mo for BC test	Improve from 2x/yr to 1x per month
	Average system response time	Application response time at bedside	Average and peak response time in sec	In DR mode, maintain no more than 40% increase in resp. time
	Average problem resolution time	Average time identify, resolve applic. problem	Average time in hours from initial report	Reduce average time to 2 hours by YE2007
Financial	Bandwidth costs	Monthly cost of bandwidth to DR site	Dollars/month of expense	Keep % annual expense growth < 20%
	How much data is being replicated	Amount of data being replicated by PR,Radio	Total production TB allocated to PR, Radio.	Maintain at 20% of total production TB
	Percentage growth rate data	Annual growth % data production PR, Radio.	Quarter to quarter data allocation report	% growth = 10% less than patient growth %

Figure 7-9 Medical center - key performance indicators

7.4.3 Next Steps and roadmap

Figure 7-10 shows the recommended Next Steps for the client. This plan also represents major milestones for the IT Business Continuity project to be executed over the next 18 months.

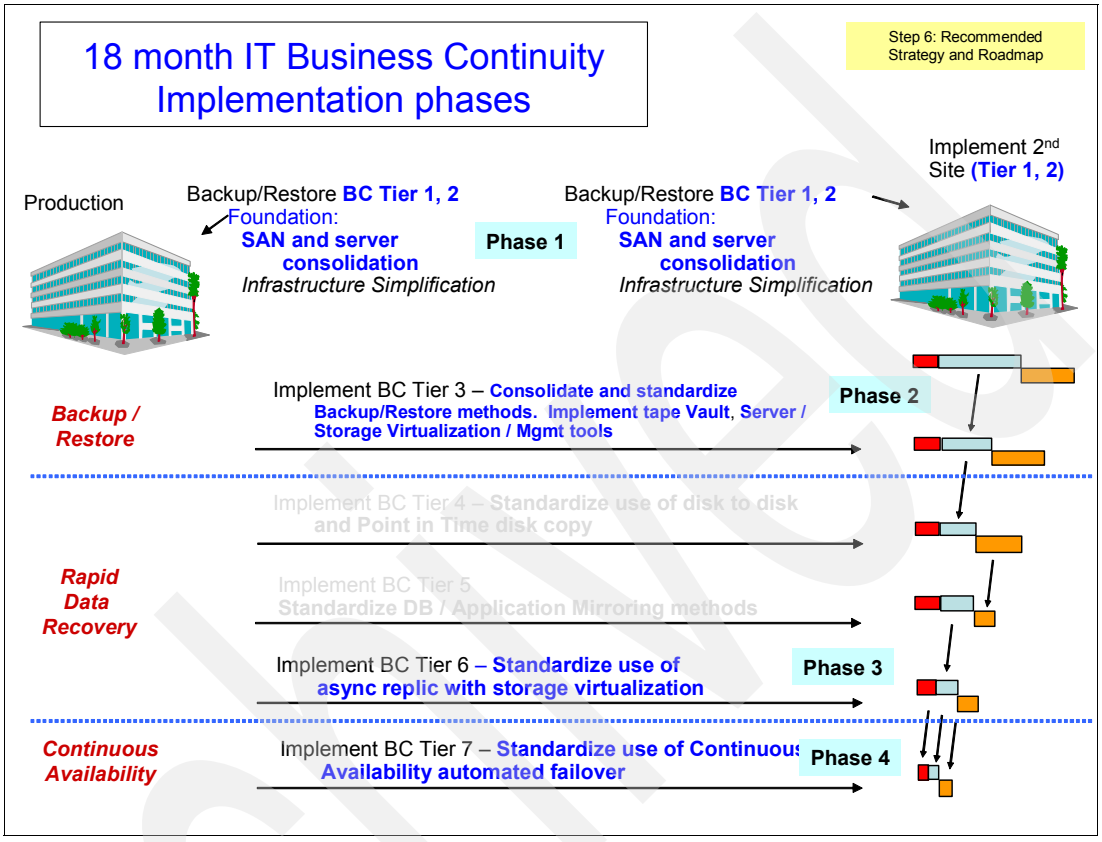

Figure 7-10 Medical center case study - high level eighteen month roadmap

- **Phase 1:**
  - Risk and business processes priority assessment
  - Architecture and solution design
  - Backup data center decision
  - Business Continuity plans, budgets and solution approval
  - Existing data center consolidation – applications, hardware, software, tools and procedure
  - Backup data center final design and configuration based on priority
- **Phase 2**
  - Mainframe and mainframe storage upgrade and migration (both main and backup site)
  - HP and UNIX backup systems move over to backup site
  - Storage virtualization and simplification project design
  - Implement storage virtualization at main site and backup site
  - Virtualization support for Intel servers at backup site
  - Overall Business Continuity strategy and planning with the Line of business
- **Phase 3**
  - Testing and dry run for disaster scenarios, including storage virtualization replication
  - Monitoring of Business Continuity readiness

- **Phase 4**
  - Future end-to-end automated failover/failback

The immediate next strategic Business Continuity project was to work with the Lines of Business to develop a hospital wide Business Continuity plan, for a scenario in which the whole hospital is out of action due to floods and hurricane

7.5 Case study summary

The client considered the Next Step Business Continuity workshop to be a success. They felt that they had been able to bring together the various departments, skills, and requirements together, and create a valid baseline Business Continuity architecture and solution.

The project team realized that Business Continuity planning is a continued process. IT is part of the overall design of application, business process and IT operations. For Business Continuity to be successful, the different departments in the organization need to work together.

8

Planning for Business Continuity in a heterogeneous IT environment

Today's business world involves acquisitions, mergers, and a very fast pace of change in an on demand world. As customer environments continue to proliferate, a major future trend in customer Business Continuity requirements is heterogeneous platform Business Continuity.

This chapter examines the concepts of effectively planning Business Continuity for a heterogeneous platform environment and also details current IBM System Storage heterogeneous Business Continuity solutions and their concepts.

8.1 The objective for heterogeneous Business Continuity

Let us begin by defining the objective for unplanned outage Business Continuity in a heterogeneous environment. In any unplanned outage, Business Continuity is about achieving resumption of service at the *application and user* level.

Resumption of service means that the applications staff must be able to perform a fast, repeatable, and of consistent duration *application and database restart*. This requires the Business Continuity solution to assure the hardware recovery, data integrity and data consistency.

What must be avoided is a situation in which the application data does not have logical integrity. If the database reports that indexes or data might be corrupted, the application database (and database personnel) must perform a time-consuming and of unpredictable duration *application and database recovery,* see Figure 8-1.

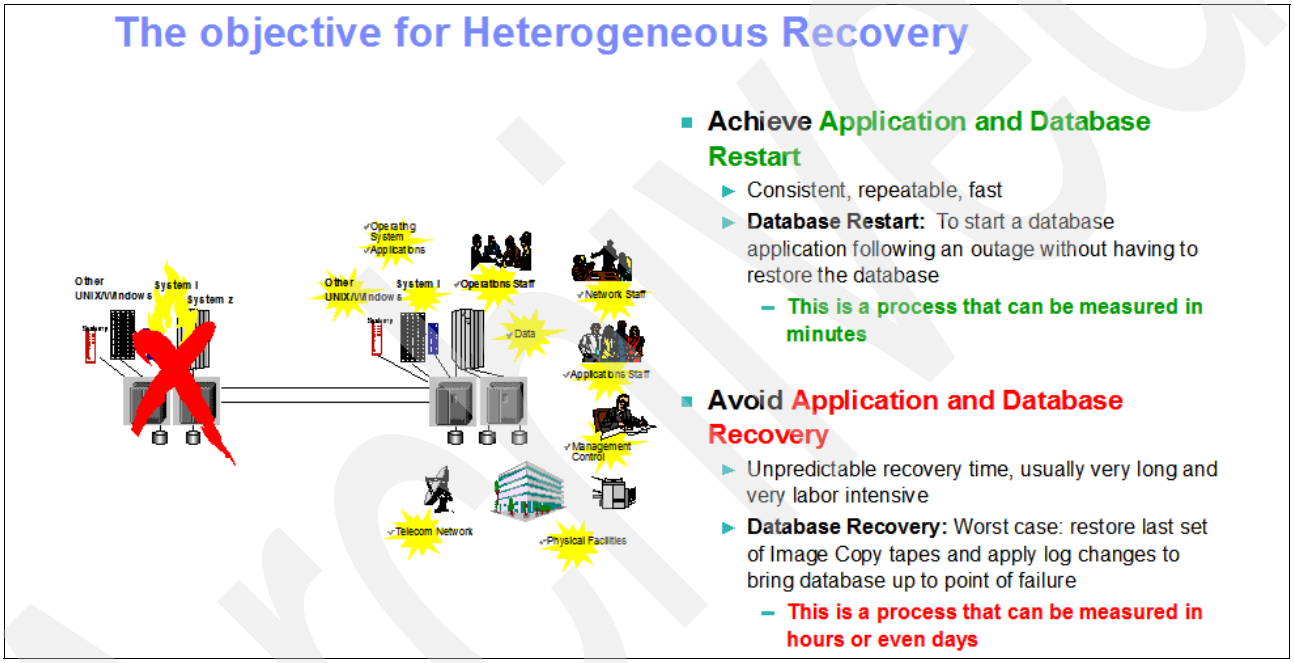

Figure 8-1 Objective of heterogeneous Business Continuity

To provide for an application or database restart in a heterogeneous environment, first review the timeline of an IT recovery and then see how that timeline is updated and managed in a heterogeneous environment.

8.1.1 Timeline of an IT Business Continuity recovery

Figure 8-2 is a diagram of an IT recovery. Time proceeds from the left to the right.

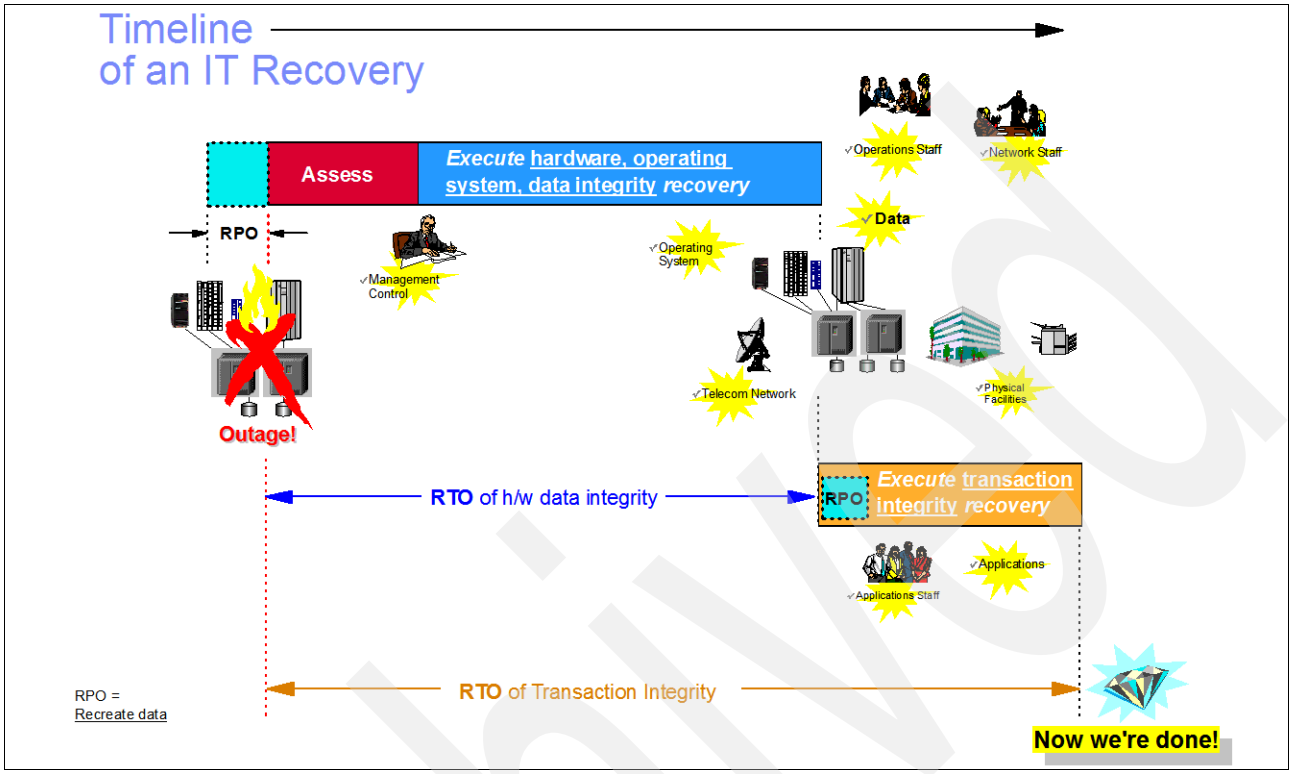

Figure 8-2 Timeline of an IT recovery

Note: In order to maintain zero data loss to a transactional level, there must be some means to recreate lost data. This means might be based on the database application (logs mirrored synchronously) or through manual re-entry of lost transactions. Without some mechanism to bring the transactional integrity to zero data loss, however, any transactions that are issued between the end of mirroring and halting of the applications will be lost.

Here are some things to note on the timeline of an IT recovery in Figure 8-2:

- After the outage occurs, the first step is for *management to assess* the outage (this incurs elapsed time). Because there is a significant capital cost to declaring a disaster and executing a recovery, management needs to be sure that the situation warrants committing their organization to that expense. After management has decided to declare a disaster, then they activate the Business Continuity plan.
- The first stage of the Business Continuity plan is to recover the hardware, operating systems, and the data itself. Operations, networking, telecommunications, physical facilities, and associated staff are involved.
- At the end of this stage, the operating systems and the data are recovered. The data ideally is accurate and data consistent to a point-in-time prior to the outage. The time duration to this point is the Recovery Time Objective (RTO) of hardware data integrity.

- However, we are not recovered from a user standpoint. The servers and storage are only capable of accurate *byte movement* (storage) and *proper data write sequencing* (servers and server clustering). Hardware data integrity is not the same as database integrity. The storage and servers cannot know what the logical database relationship is between multiple data blocks in the database. Therefore, after the first stage is complete, the transaction integrity recovery must next be performed by the applications staff, on the application and database.
- The applications staff performs transaction integrity recovery. Hopefully, this is a database restart and *not* a database recovery. This process will back out incomplete logical units of work, and restore the database to logical integrity as of the most recent time possible.
- When the transaction integrity recovery (rollback or roll forward) is complete, we now have the application and databases ready for user access. This duration is the RTO of transaction integrity.
- The difference in elapsed time between RTO of hardware data integrity, and RTO of transaction integrity. When discussing the RTO, it is important to distinguish which of these two is being referred to. Operations and application staff can have differing perceptions of RTO depending on whether the RTO is assumed to be at the hardware recovery level, or at the application recovery level. The fact that there are these two different RTOs is essential to understanding and planning for heterogeneous Business Continuity.
- Observe how Recovery Point Objective (RPO) is depicted in Figure 8-2. RPO (which is the amount of data recreation required prior to the point of outage) is shown as the time offset before the outage occurred. Note that the RPO data recreation happens in the transaction integrity recovery stage. RPO data recreation cannot be done in the hardware and operating system recovery, as the server and storage components do not have knowledge of the logical relationships between multiple application and database blocks of data.

8.1.2 Today's ever-proliferating applications, servers, and platforms

Simpler environments (that is, with a common architectural platform with common operating system, application software, and database) lend themselves to simpler Business Continuity procedures. In today's environments, however, there many reasons why enterprises choose not to have a common architecture, including acquisitions, mergers, different applications requirements, pricing, an ever changing expansion of technology, and varying relationships with vendors. These factors often result in the situation shown in Figure 8-3.

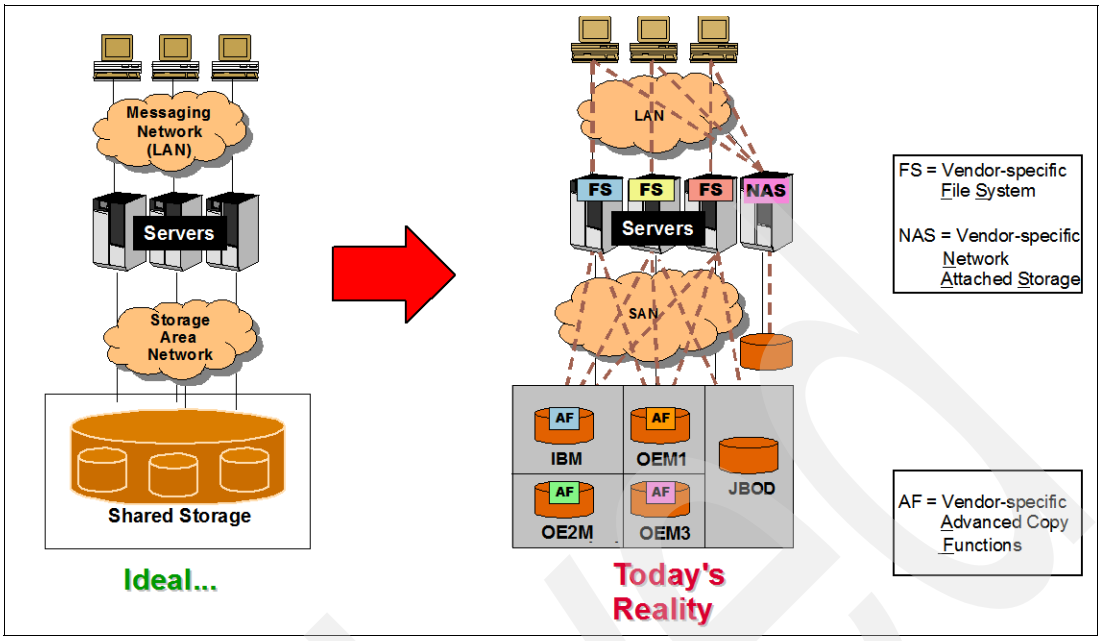

Figure 8-3 Today's ever-proliferating applications, servers, and platforms

The challenge of successfully performing Business Continuity in the environment on the right environment is made worse by the following difficulties:

- Coordinating multiple platform recovery, both technically and politically.
- Managing dynamic application changes that are happening at a very fast pace in all of the environments.
- Managing dynamic IT environment changes.
- Managing data and storage allocation, which can be spread across multiple departments, staffs, and platforms.

Let us now see what happens when we apply the timeline of IT recovery to a heterogeneous platform Business Continuity scenario and how the presumption of a common operating system platform recovery breaks down in that scenario.

Multiple platform management complexity

In a a heterogeneous platform Business Continuity scenario, it is generally *not* desirable to attempt to somehow coordinate multiple traditional *common platform time lines of IT recovery*. To do so invites complexity, as shown in Figure 8-4.

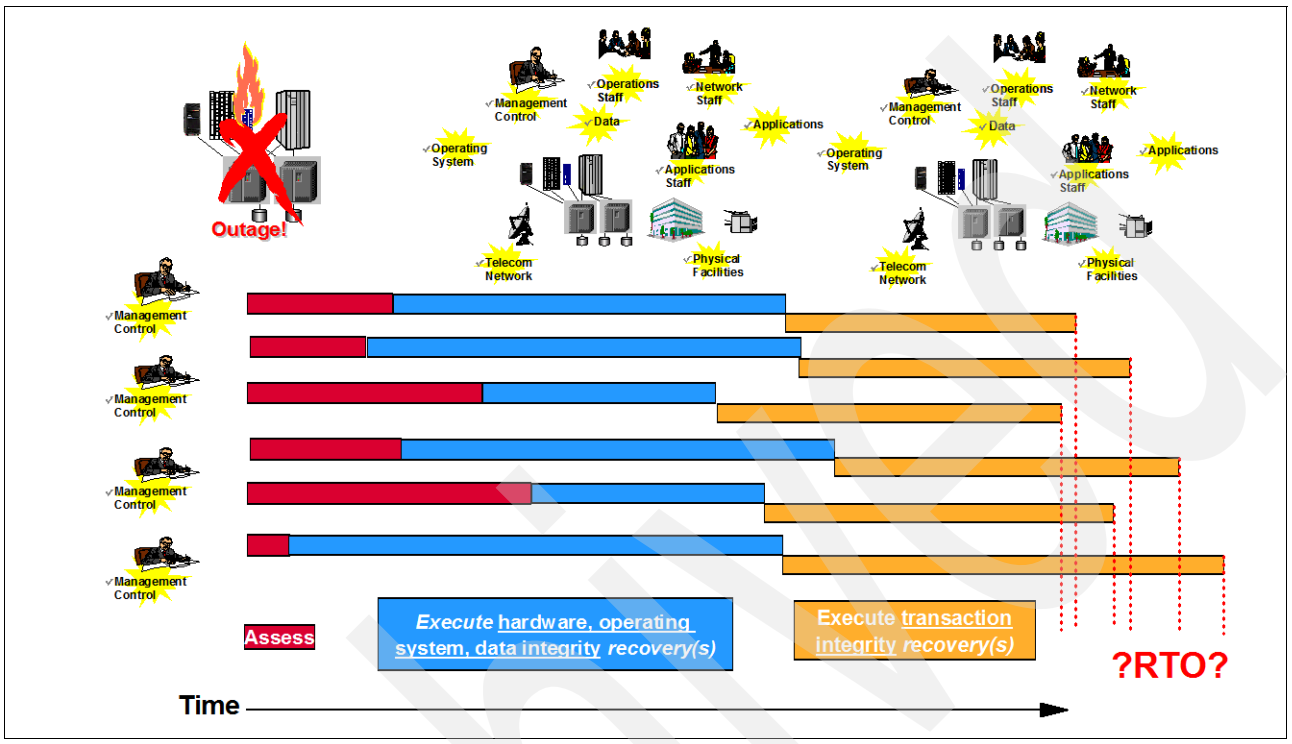

Figure 8-4 Undesirable: Multiple platform Business Continuity complexity

This scenario, with multiple management controls and multiple instances of staff doing potentially duplicate work, is very complex to manage. Thus, it will have significant difficulty in coordination and will have a lower chance of success.

From a business standpoint, this difficulty in managing the platform imposes severe limitations on the enterprise's ability to continue business success and growth, because the associated applications and IT infrastructure expansion becomes difficult or impossible to provide for effective Business Continuity from an enterprise standpoint.

8.2 Solutions for heterogeneous platform Business Continuity

With the application of new technology, the timeline of IT recovery changes for an effective heterogeneous Business Continuity solution. There are two common viable heterogeneous Business Continuity possibilities with today's technology:

1. Multiple applications, multiple platforms (application agnostic)
2. One application, multiple platforms (transaction integrity)

We examine each of these possibilities in detail.

8.2.1 Heterogeneous IT recovery #1: Multiple applications, multiple platforms

The first heterogeneous IT Business Continuity configuration possibility is multiple applications, running on multiple platforms. In this configuration, recovery is provided at the *hardware and data integrity level* recovery across multiple operating systems and platforms. This solution configuration does not attempt to perform transaction integrity recovery, which is the individual platform's responsibility.

Figure 8-5 illustrates the fundamental concept of this solution.

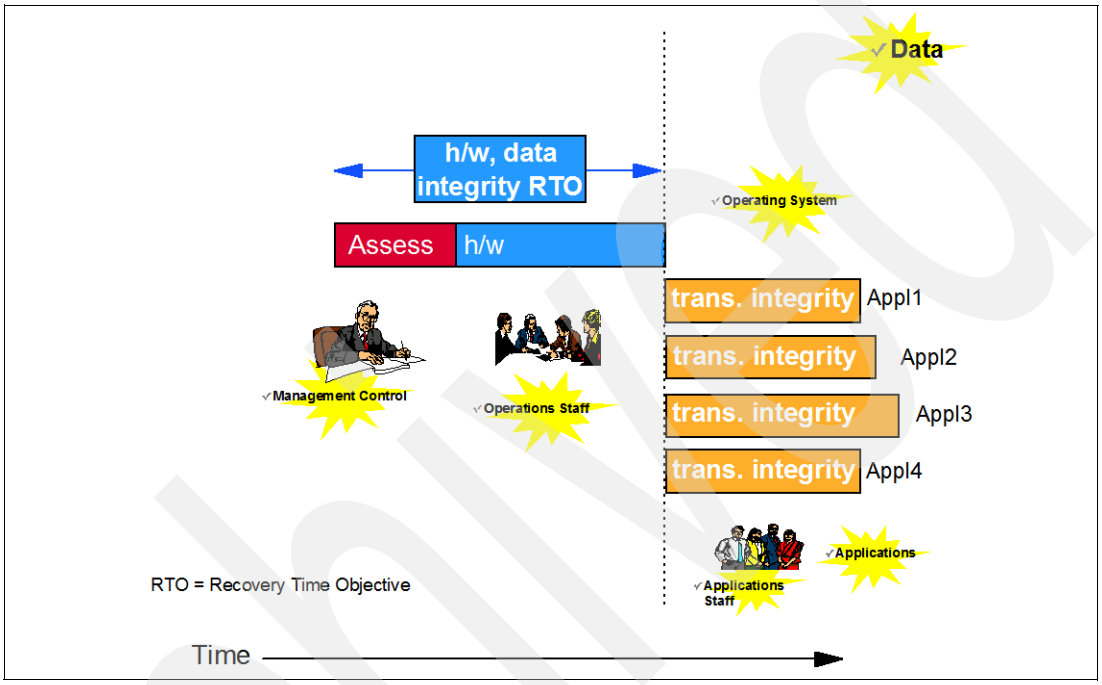

Figure 8-5 Heterogeneous recovery possibility #1: Multiple applications, multiple platforms

Note the following for this timeline of heterogeneous IT recovery, as compared with the unmanageable configuration in Figure 8-4 on page 256:

- The major enhancement of these solutions is the provision of a *consolidated* hardware, operating systems, and data recovery into a single coordinated recovery.
- These solutions provide the functionality for the data to be data consistent at a *single point-in-time across multiple platforms* in the event of an outage, which is represented by the single vertical line that represents all the data to be consistent as of that point-in-time. The data at the recovery site has the exact data image as though an instantaneous power outage had occurred at the primary site. Therefore, each platform, in starting its recovery, starts from the same point-in-time.
- These solutions do not attempt to perform transaction integrity, which is an advantage because with this limit of scope, the solutions are *application agnostic*. This configuration insulates against *difficult to manage* data and *difficult to manage* applications. Because these solutions are not designed to interlock with the application software or databases, there is no need to install and maintain client code on each of the platforms.
- After the hardware, operating system, and data recovery stage is complete and has the consolidated single point-in-time data available, each individual application or platform, then will perform their own individual restart and individual transaction integrity recovery, just as they do today.

- The application or platform restart and transaction recovery procedures are unchanged from the exact same procedures they would use today to recover from an instantaneous power outage.
- This configuration also has the advantage that there is not an architected necessity for the multiple platforms to interlock their recovery with each other. This architecture makes for very good scalability, because all departments in recovery can proceed in parallel.
- When all individual recoveries are complete, the result is the best possible time to recovery for a coordinated enterprise collection of heterogeneous applications or platforms.
- This configuration can be used as foundation for further application level transaction integrity.

Described in this book are various IBM server and System Storage integrated Business Continuity solutions that are available today to provide heterogeneous IT Business Continuity at the hardware, operating system, and data integrity level. These solutions include, but are not limited to:

- GDPS/PPRC Open LUN Management
- GDPS/Global Mirror Open LUN Management
- GDPS HyperSwap Manager with Open LUN Management
- TotalStorage Productivity Center for Replication with Metro Mirror or Global Mirror on DS6000, DS8000, or SVC (Metro Mirror only)
- SAN Volume Controller Metro Mirror or Global Mirror
- Global Mirror on DS6000, DS8000, and SVC

We give some more information about these solutions following in this section and more detailed information in *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548.

How consolidated hardware level data recovery is achieved

In order to recover applications across one or many platforms, it is necessary to ensure that all of the data is received in the recovery site. While this is often assumed as fact, it can be a dangerous assumption to make and understanding the Consistency Group functions, and their various implementations, is key to this.

Choosing the appropriate platform for data mirroring

When planning out a Rapid Recovery environment for Business Continuity, one common question deals with how to mirror the data. Even within the broader categories of synchronous versus asynchronous mirroring, there is still the decision to be made as to whether the mirroring will be handled at the hardware block level or at the transaction level through the database application itself.

While this is never a simple decision, we have assembled a list of questions to consider in guiding you toward the appropriate decision:

1. What needs to be recovered?
   - Which applications need to be recovered?
   - How are the applications segmented (Backup and Restore, Rapid Recovery, Continuous availability)?
   - What level of recovery is needed? Transaction? Block level?
   - On what platforms do the applications reside (server, operating system, storage)?
   - Is there non-database data to be recovered?

2. What is the existing skill base?

   Do your employees have skills that are geared towards doing mirroring and recovery on any particular storage or application platform?

3. What are the internal standards?

   Which DBMS, operating system, server, and storage are within your business's list of standard vendors?

The following sections provide a detailed discussion of these questions.

What needs to be recovered?

This is the first question that really needs to be understood. Before dealing with more abstract concepts along the lines of recovery points or tiers, the basic understanding of which applications need to be recovered is necessary.

Moreover, it is critical to understand how they are segmented. That is to say that some applications can tolerate a longer outage than others. The most critical applications can go into the category of *continuous availability* while those applications that are needed with the least frequency and which can tolerate more potential data loss would go into the category of *backup and restore*. Rapid data recovery falls between the other two segments.

Understanding the hardware infrastructure of where the applications reside helps further refine the selection criteria.

Finally, it is important to consider whether all of the critical data is going to be found within databases. In many cases environments will have data that exists outside of any of their databases and, as such, will need some other form of replication. Examples of such files could be flat files and image data (such as X-rays for hospitals). Again, it is important to understand how this data is segmented. If it is critical that these files are maintained as current as possible, then you will likely want to segment it as continuous availability or rapid recovery and thus consider block level or logical mirroring. If it is segmented as backup and restore, however, then there might be better alternatives.

What is your existing skill base?

If your administrators have skills in place for a specific type of mirroring then going with that form of mirroring might be less disruptive to your environment in the short term. You should still consider alternatives; however the simplicity of implementation is certainly something to consider.

What are your current internal standards?

Further narrowing the field of considerations are any internal standards for technology purchasing.

Rolling disasters

Maintaining data consistency is important because disasters do not occur in a single instant. An explosion travels at the speed of sound, a power failure travels at the speed of light, and batteries designed to ride out power fluctuations fail at different times. The point is that a single disaster will not affect an entire data center simultaneously. Thus, data centers are exposed to what we refer to as *rolling disasters*.

Take, for example, the situation where a fire starts in the data center. As the fire spreads, perhaps the adapters or the connectivity that is responsible for mirroring some of the data is damaged. If the storage system is able to continue operating, some transactions can continue mirroring while others cannot. This situation becomes a serious issue where dependent writes are concerned.

Note: It is important to note that one very big reason that you want to ensure that data is consistent to the point of the first failure is that there are times when the data can have inconsistencies but the database restart will still process correctly. This situation is a very dangerous situation to be in, because it is possible that you will not notice the corruption for some time and that the recovery actions that need to take place will be even more difficult when you notice the corruption.

For example, Figure 8-6, shows a database transaction that consists of a log write with an intent to update the data, the update to the data, and the log write when the update completes. If the data update cannot mirror, but the log continues to do so, then the data in the recovery site is inconsistent. When the disaster is declared, the database might not restart in the recovery site because, although it has log updates reflecting all of the entries that were written in the production site, it does not have any of the corresponding data. If a restart is successful, the condition cannot be detected. As a result, the database must be recovered manually. Depending on the amount of time between the failure and the halting of the applications, the database recovery could take days or might be impossible.

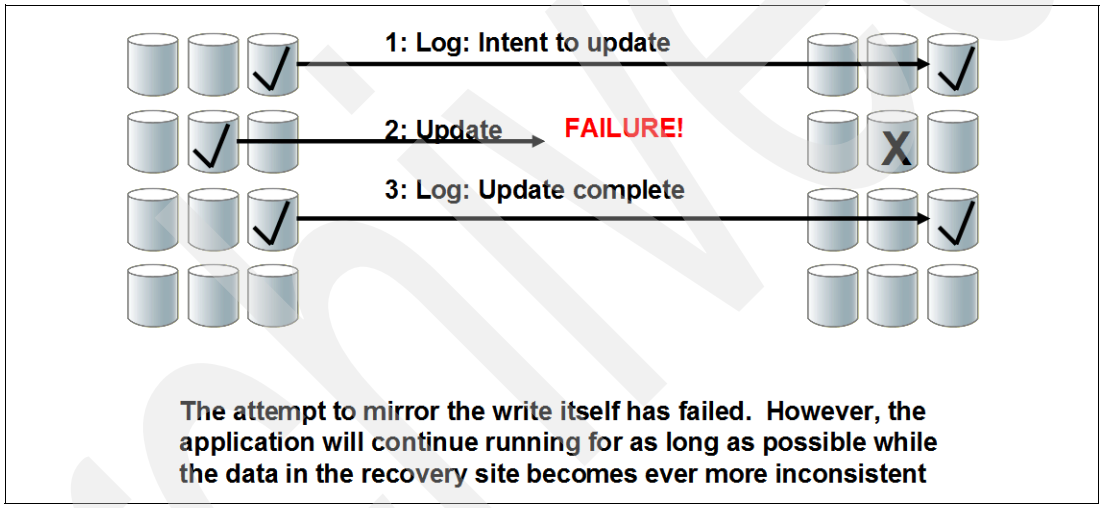

Figure 8-6 The impact of a rolling disaster creating inconsistent data as some pieces fail before others

To avoid these situations we must, at a minimum, make use of *Consistency Groups*, but consistency groups have different implementations depending on the hardware used and whether the mirroring is done synchronously or asynchronously.

Note: All data that needs to be held to the same point in time of consistency will nee to be replicated through the same type of mirroring technology. That is to say that all of the data in question needs to be mirrored through DS6000 or DS8000 Metro Mirror, through SVC Global Mirror, or through the platform of choice. If you choose to mirror through SVC and DS8000 and DS4000 independently (as opposed to, for example, mirroring all data through the SVC), you end up with multiple points of consistency, even if you have control software running on top of the environment.

Metro Mirror

Metro Mirror is the IBM terminology for synchronous remote copy between two storage controllers. Synchronous remote copy technology is a form of mirroring where each write is sent to the target disk in order and the write completion is not acknowledged to the server until the write to the secondary completes and acknowledges back to the primary. Metro Mirror is available on three different IBM System Storage products:

- DS6000/DS8000
- DS4000
- SVC

In this case, we deal with DS6000/DS8000 and SVC. The DS4000 does not have a consistency group function in its implementation of Metro Mirror but instead uses Global Mirror asynchronous copy in situations where consistency of dependent writes is a concern.

Maintaining data consistency in synchronous environments is complex and yet is of the highest importance because synchronous remote copy sends its data on an ongoing basis as each transaction comes in. It is a common belief that because each transaction must be written to both disk systems before the next transaction begins, that the data will inherently be consistent. This is a dangerous line of thought because, as illustrated in Figure 8-6 on page 260, the mirror continues trying to mirror although some segments of data cannot be mirrored.

DS6000 and DS8000 Metro Mirror

In the DS6000 and DS8000 environment, the first line of defense against this inconsistency is the creation of Consistency Groups. In these disk systems a consistency group is a set of volumes which are typically associated with a single application.

In Figure 8-7, the volumes that are outlined are in the consistency group. Should any of the volumes in the consistency group be unable to mirror its write (as is the case in Figure 8-7), that volume triggers an Extended Long Busy (ELB) and the writes are held. It is important to be aware, however, that this is not enough to protect against a rolling disaster in and of itself.

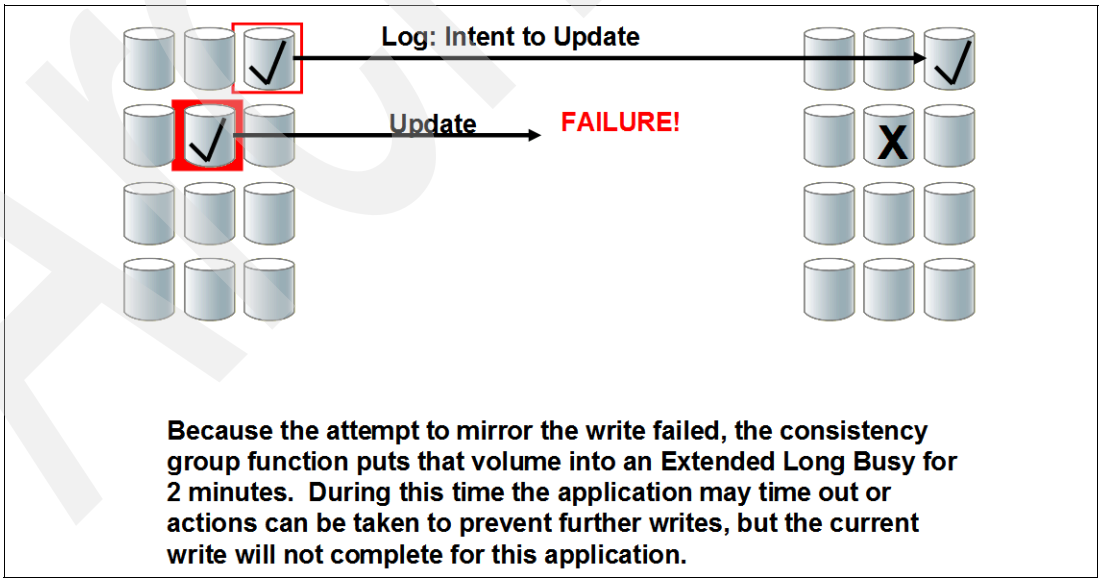

Figure 8-7 The Consistency Group function puts a volume in to ELB

There are two reasons that the Consistency Group function fails to provide complete protection in this case:

- As shown in Figure 8-8, only the specific volume that failed to write is protected. As such, even if a disk system had one large Consistency Group, independent writes continue from other applications to other volumes in the same disk system or to other attached storage. Without this ELB, at a minimum, the consistency and recoverability of the dependent applications can be called into question.

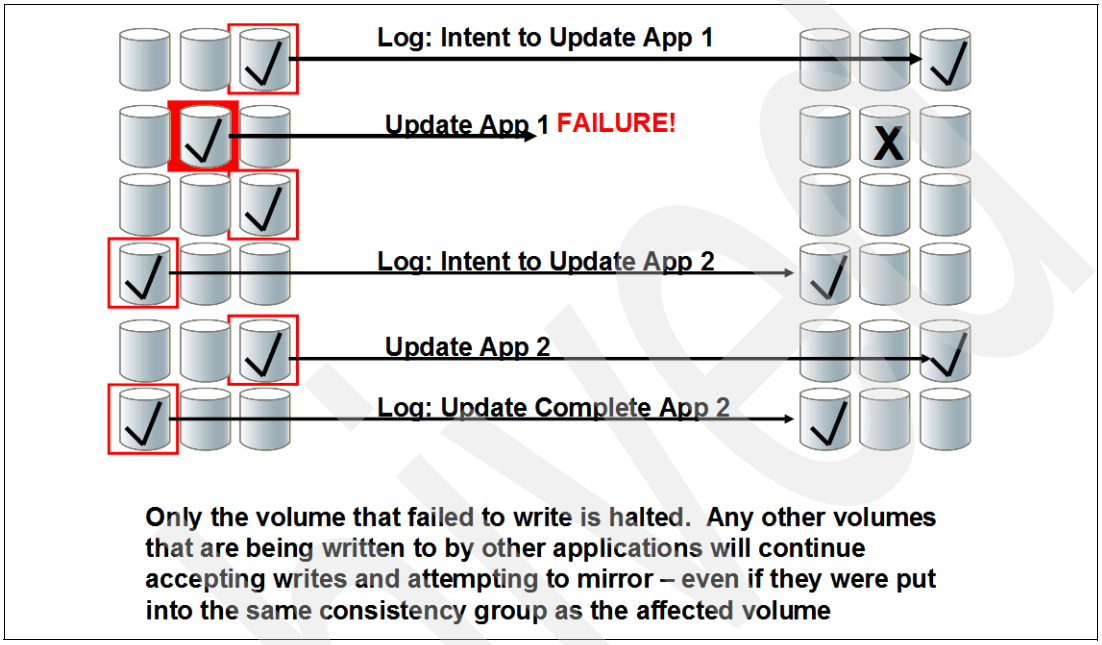

Figure 8-8 The Consistency Group protects application 1, but application 2 continues without pause

- Even within a Consistency Group, the consistency is not protected forever. The failed write within a Consistency Group triggers a timer as it sends the volume into the ELB state. This timer lasts for two minutes (by default, extendible up to 18 hours), during which you can take actions to protect consistency.

  At the end of the ELB timer period, the ELB is released. If the application has not timed out, it resumes writes and the storage controller tracks data changes in a bitmap, the assumption being that when the problem is fixed the changes are sent to the target disk. Other volumes which had previously been held from updating but which are still capable of mirroring will do so when the ELB is released.

Because of these two issues, it is highly recommended that any implementation of Metro Mirror on the DS6000 or DS8000 include control software which makes use of a *freeze function*.

Note: Control software is a server based program that monitors all of the replication functions on an ongoing basis and reacts according to policies set by the administrator.

Using the freeze function

A freeze function is a function that reacts to a Consistency Group failure. Rather than depending solely on the affected volume that remains in ELB, the freeze function quickly suspends all mirroring between all volume pairs within the Consistency Group that is protected by the Freeze. As a result, if you place all mirrored pairs into the same Consistency

Group, as shown in Figure 8-9, the consistency of dependent writes is protected on all volumes, LSSs, and disk systems.

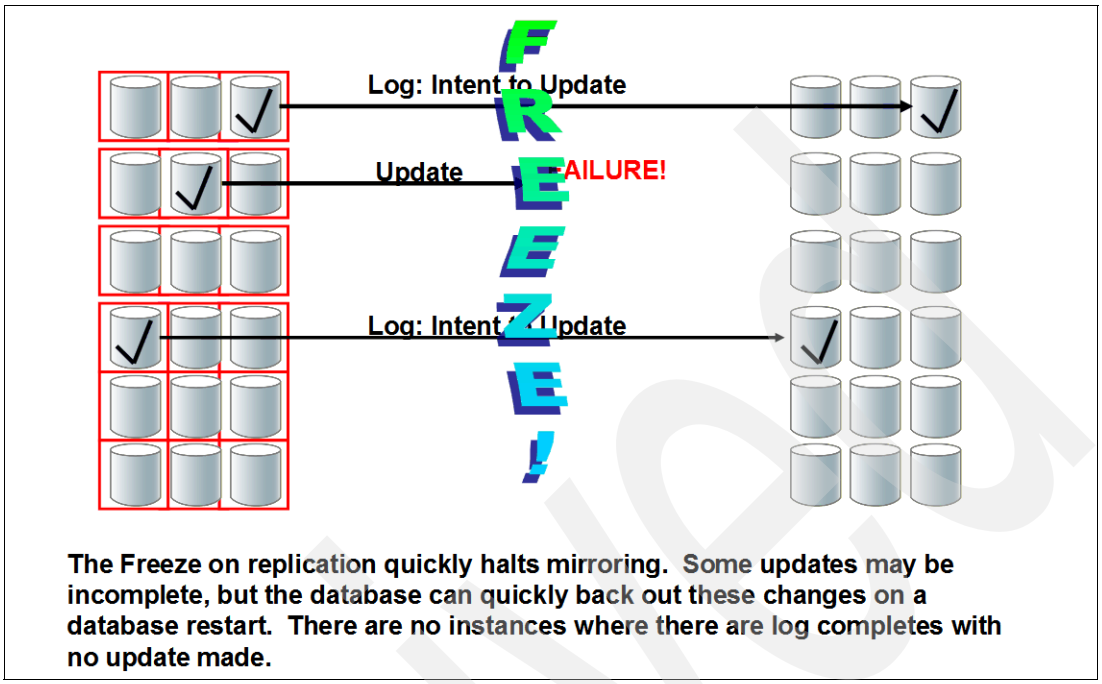

Figure 8-9 Freeze stops all replication at the first replication failure and both applications are protected

In IBM products and services, freeze functions typically come in two forms:

- **Freeze and run** (sometimes called *freeze and thaw*) indicates that when a Consistency Group is triggered, the mirror is suspended but I/Os continue to the production disk system. In this case, data consistency is protected but all further writes are lost if the triggering event is an actual disaster.
- **Freeze and stop** indicates that the mirror is halted and all I/Os to the production volumes likewise halt. Although in-flight transactions are lost, no further writes are issued, so no other data is lost.

Of the two forms, freeze and run is far more common because non-disasters, such as construction crews accidentally cutting a fiber conduit, are far more common as causes of a halted mirror than actual disasters that would bring down the data center itself. As such, most users are willing to accept that they will lose the transactions processed between the component failure and end of problem determination, as long as they know that data is consistent and restartable in the recovery location.

The alternative provides few transactions to recreate but all applications must then be restarted and, in a non-disaster, the business could be disrupted by the application downtime.

Freeze functions are included with FlashCopy and as features in GDPS/PPRC, TotalStorage Productivity Center for Replication, and PPRC Migration manager as supported products and services from IBM. Freeze is also a command that can be issued to the DS6000 and DS8000 in custom scripting, reacting to SNMP messages resulting from the triggering of a consistency group.

CRIT(YES) and CRIT(NO)

CRIT(YES) is setting in Metro Mirror which pre-dates the existence of Consistency Groups or Control Software with integrated freeze functions. Use extreme caution when using the CRIT(YES) setting. In the event that any volume within the Metro Mirror configuration fails to mirror, all I/Os to that disk system immediately halt when set to `Heavy`. Even when CRIT(YES) is set to be `Light`, it halts all I/Os to that volume.

Obviously, there is substantial risk involved in this setting, including:

- The odds that a non-disaster event has caused the failure substantially outweighs the odds of a disaster having occurred in most circumstances.
- For most administrators, the prospect of having applications come to a grinding halt every time a construction accident cuts through a fiber conduit is daunting.
- Unlike Control Software issuing a *Freeze and Stop*, which can span the entire environment if that is what is desired, CRIT(YES) affects only that disk system. So, other applications can continue running and their volumes will not be synchronized with the ones that shut down.

As a result of these impacts, while CRIT(YES) protects the consistency of some data, it is *not* recommended.

CRIT(NO) is the default setting for CRIT in Metro Mirror, but it is not a substitute for Consistency Groups. Specifying CRIT(NO) is the appropriate specification when defining Consistency Groups, which, allow you to protect data without creating the immediate and permanent halt to applications that CRIT(YES) delivers.

CRIT(NO) is the opposite of CRIT(YES) in terms of impact. When CRIT(NO) is specified, the primary volume returns an error condition to the I/O writes, even if the secondary volume cannot be successfully updated. Metro Mirror is aware of the issue and can handle it properly using the Consistency Group function, yet it allows the application to continue to update the primary volumes. This combination of CRIT(NO) and Consistency Group allows production application workloads to continue in the event of momentary, transitory conditions that affect only the remote links or remote site.

In addition, in terms of allowing you to have a crash consistent copy, using CRIT(NO) and Consistency Groups is the appropriate specification when using Control Software with a freeze function as we discussed earlier.

Specifying CRIT(NO) without the use of Consistency Groups is not recommended, as such a specification would not prevent a rolling disaster. Specifying CRIT(NO) without the use of Consistency Groups does not halt the mirroring of the block level data in a crash-consistent, power-outage consistent manner. Do note, however, that specifying CRIT(NO) is the appropriate prerequisite specification for proper setup of Consistency Groups. Consistency Groups (with freeze) are the recommended method to allow you to protect data without the negative impacts of CRIT(YES).

SAN Volume Controller

Writes in the SAN Volume Controller (SVC) version of Metro Mirror occur in much the same way as in the DS6000 and DS8000. However, the Consistency Group function is implemented quite differently. While consistency groups on the DS6000, DS8000, and SVC can span multiple Logical SubSystems (LSSs) across multiple storage controllers (with freeze and run control), an SVC consistency group can have up to 1024 volumes.

The second difference is that the SVC consistency group failure does not trigger an ELB on a single volume. Instead, a freeze and run function is integrated into the SVC. If any of the volumes in the consistency group are unable to mirror, the SVC will halt mirroring across all of its links and attached disk systems but will allow I/Os to continue and track changes in a bitmap-like concept called *grains*.

So, while TotalStorage Productivity Center for Replication is of value in a SVC Metro Mirror environment for purposes of management, it is not strictly necessary from a data consistency perspective.

Global Mirror

By its nature, Asynchronous Remote Copy maintains consistency. This consistency is necessary because data can be sent periodically in groups of writes (as opposed to individually in order) or can be sent out of order, depending on the mirroring technology used. The exception to this being Global Copy or other non-synchronous writes, which do not necessarily attempt to maintain consistency.

That said, it is still important to understand how consistency is maintained in these environments and that will vary based on the specific form of remote copy in use.

All IBM asynchronous remote copy technologies are referred to as *Global Mirror*. Each will be described individually in this section.

DS6000/DS8000

In the DS6000 and DS8000, Global Mirror maintains consistency by way of a series of bitmaps. Initially all new writes will be indicated in a bitmap labeled *out of sync*, and Global Mirror will scan through this bitmap and look for changed tracks of data. As it finds these changes, they are sent across a high performance link to the secondary disk systems.

None of this data is consistent during this phase because it is being sent based on the scan of the bitmap finding a track that has been updated, not based on the order in which data has been written. As such, the consistency group function in this technology is arranged in a process of switching bitmaps.

Periodically (at times determined by the users or by Global Mirror), a consistency group will form. In this case, a consistency group consists of data sent to the recovery site by all volumes being mirrored in a single global Mirror session, no matter how many DS6000s or DS8000s it spans.

Note: In DS6000 and DS8000 Global Mirror, consistency is maintained across all disk systems in the Global Mirror session without requiring external software to manage this. While control software can be useful for management of the Global Mirror session, it is not needed to ensure data consistency across multiple disk systems.

This consistency group formation process occurs when Global Mirror stops updating the out of sync bitmap and begins updating a new *change recording* bitmap as writes come in. This, again, is coordinated across all disk systems in the Global Mirror session. The toggle of the bitmap occurs during a dynamic *data freeze* during which all I/O is held in a brief ELB (design objective for this freeze is 2 to 3 ms).

The disk system which has been labeled *Master* coordinates the switch with all subordinates by sending commands and receiving responses across the SAN. They only switch to the change recording bitmap when they can all switch, to make certain that the out of sync bitmaps are all consistent to a single point in time. When this happens the writes that had been noted in the out of sync bitmap will continue to drain to the secondary disk. By making

this switch, consistency is preserved to a single given point in time without impacting the production write activity.

When the updates indicated in the out of sync bitmap finish being written to the target disk, the data is consistent to the point in time when the bitmap switch occurred. In order to preserve that consistency the next thing that happens is the issuance of an in-band FlashCopy command. By doing so, the consistent data is preserved and the process can begin again with the change recording bitmap becoming the new out of sync bitmap.

Important: Because a target FlashCopy is performed, twice the disk capacity is required at the Global Mirror target site.

It is important to note that the consistency group will only form if all of the disk systems can participate. If any cannot, Global Mirror will back off and wait for a cycle before it attempts to form a consistency group again.

DS4000

The DS4000 does Global Mirror as a subset of its Enhanced Remote Mirror (ERM) technology, which also includes Metro Mirror and Global Copy. Although it is also referred to Global Mirror, the mechanism for how asynchronous remote copy is performed in the DS4000 is completely different than that in the DS6000 and DS8000.

In this case, the DS4000 maintains a separate LUN in each side of the mirrored pair known as the *Mirror Repository Logical Drive*. This drive maintains a queue of writes that have been made to the primary disk system. These writes are then sent, in order, to the secondary as the disk system is able while writes continue. Because these writes are sent in a First in First Out sequence, the data in the secondary disk system is always consistent and it never goes through a period in which the data is unusable.

As the DS4000 version of Metro Mirror does not support consistency groups, Global Mirror should be used in any DS4000 environment where there are concerns about maintaining dependent write sequencing.

SVC

As is the case with Metro Mirror, SVC supports a version of Global Mirror, but this is its own form of Global Mirror. In the SVC version of Global Mirror, we issue sequence numbers in order to maintain the proper order in which to apply them in the recovery site. These sequence numbers are used based on a few criteria:

1. Sequence numbers are applied in such a way that if B is dependent on A then B will be given a higher sequence number.
2. Those sequence numbers are applied from lowest to highest, so that if B has a higher sequence number than A, B will be applied after A.
3. If writes are not dependent on each other, they can be given the same sequence number.

As writes are received in the secondary SVC, they are held in cache and only applied when the proper order is received. By doing so we are assured that the data is applied in a restartable format, that is to say that if we have writes 1, 2, and 3 if the remote SVC receives writes 1 and 3 it will hold them in cache until it receives 2.

Writes that do not require a specific sequence, such as those in separate applications, can be given identical sequence numbers, so multiple 1s can be applied simultaneously, for instance.

z/OS Global Mirror

z/OS Global Mirror (formerly known as XRC) uses a unique method of maintaining consistency groups. In this particular case, there is activity occurring on both the primary and secondary side of the environment.

First, each write to the primary disk system is time-stamped from a common time source (that is, Server Time Protocol). This gives an exact sequence of how to apply them when they reach the recovery site. the time stamps are kept in a portion of cache called the *Side File*.

The second half of this comes from the recovery site.

While Global Mirror “pushes” data from the primary to the secondary Disk Systems, z/OS Global Mirror instead uses a an active z/OS LPAR (or coupled LPARs) in the recovery site to “pull” the data across with System Data Movers (SDMs).

The SDM is a function within DFSMSdfp™, so any z/OS customer already has it without needing to buy additional software. The System Data Mover examines the contents of side files on all LSSs on all Disk Systems in the environment. It looks at the time stamps and selects the minimum of the maximum time stamps. All updates that are time stamped before this selected time are included in the currently formed Consistency Group and then that Consistency Group is journalled and applied to the target disks. Typically, the SDM forms several Consistency Groups per second and the entire Consistency Group must be applied for it to be committed.

In terms of scalability in consistency, z/OS Global Mirror is unmatched. While one SDM can handle somewhere between 1000 and 2000 volumes, it is possible to make SDMs work together to scale far beyond that. 13 SDMs can be *clustered* together in a single LPAR. These clustered SDMs can then be *coupled* with up to 14 clusters participating. As a result, z/OS Global Mirror is capable of handling even the largest z/OS environments with a single point of consistency.

8.2.2 Heterogeneous recovery #2: One application, multiple platforms

The second of the two heterogeneous Business Continuity solution possibilities is one application, multiple platforms. In this configuration, there is the capability to maintain transaction integrity across multiple platforms. Figure 8-10 illustrates this configuration.

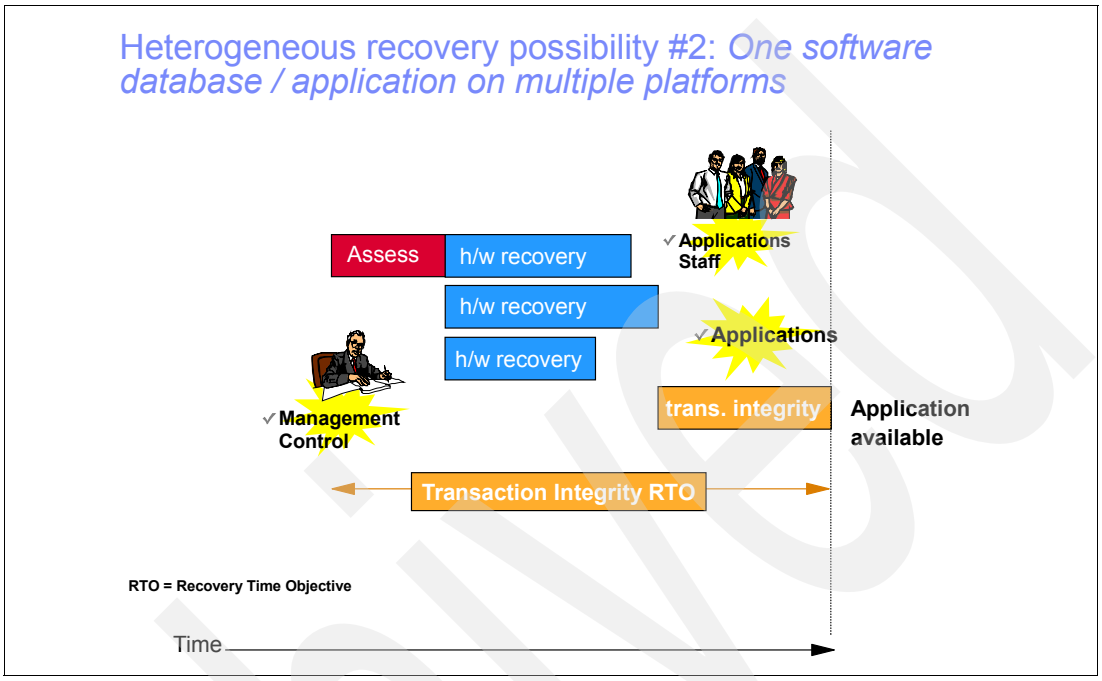

Figure 8-10 Heterogeneous recovery possibility #2: One application, multiple platforms

As we have discussed, transaction integrity cannot be performed at the hardware and operating system level, It can only be performed by the actual application and database. Therefore, the prerequisite functionality for this configuration is that the cross-platform functionality must be imbedded in the application and database software.

Transaction integrity across multiple platforms can be accomplished only when the software has this embedded functionality to interlock all the instances and platforms. It is not possible with current technology to provide transaction integrity across multiple applications, if the applications themselves do not have the supporting interlock functionality.

Some specific examples of this kind of multiple platform application or database software are (this is not an all-inclusive list):

- SAP Advanced Infrastructure Solutions
- DB2
- MQSeries

The one application, multiple platform heterogeneous Business Continuity configuration is implemented using the application-specific software skills, design, and implementation. High application awareness is required. In return for the application awareness and skills required to implement, the one software, multiple platform configuration can provide Business Continuity transaction integrity.

Because the heterogeneous recovery is based in the software's capability, a trade-off is that the transaction integrity recovery cannot extend past that application software's boundaries.

8.3 Comparison and decision tree

Let us compare and summarize the two viable heterogeneous Business Continuity configurations, as shown in Figure 8-11.

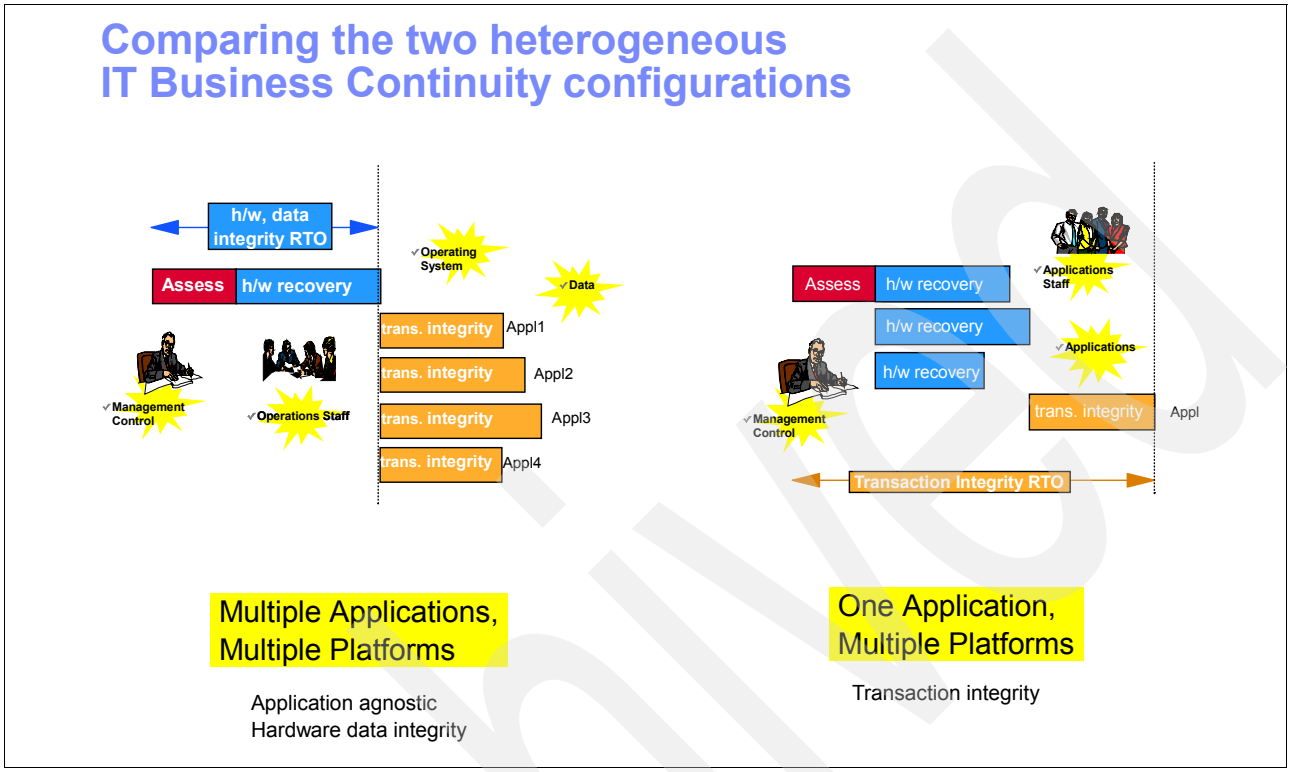

Figure 8-11 Comparing the two heterogeneous Business Continuity configurations

It is a business trade-off that determines which of the two configurations, or a blend of the two, might be appropriate for any given heterogeneous IT Business Continuity environment:

- **Configuration #1:** Multiple applications and multiple platforms, offers:
  - Insulation from hard to define / hard to manage applications or data.
  - Can be used as the foundation for application level transaction integrity.
  - Trade-off: By itself, it does not provide transaction integrity.
- **Configuration #2:** One software and multiple platforms, offers:
  - Only available when application-level recovery capabilities are present.
  - Transaction integrity can be provided at the application layer.
  - Can and should layer on top of hardware data integrity recovery.
  - Trade-off: Cannot extend transaction integrity past that application software's capabilities.

A database restart does assume that hardware and operating data integrity are provided as an underlying foundation for the transaction integrity Business Continuity. The first configuration, hardware and operating system level heterogeneous Business Continuity, can be used as an underlying foundation for the second configuration, heterogeneous one application, multiple platform Business Continuity. The two configurations are not mutually exclusive.

8.4 The value of control software in a heterogeneous environment

While we have discussed the value of control software in terms of its impact on the task of maintaining consistent data, it is important to be aware of its other impact. Control Software, as the name indicates, provides a single point of control for replication functions. By using a product or service such as TotalStorage Productivity Center for Replication or GDPS (PPRC, PPRC HyperSwap Manager, or Global Mirror), all replication functions come through one interface.

Rather than requiring an administrator to maintain separate replication functions across multiple platforms and using multiple interfaces, we can have that same administrator use a single application in order to accomplish all tasks in terms of adding, removing, suspending, or deleting pairs.

8.4.1 TotalStorage Productivity Center for Replication

TotalStorage Productivity Center for Replication is an application available through IBM System Storage and provides consistency of data and a GUI for management of replication functions on the ESS800, DS6000, SAN Volume Controller, and DS8000s. TotalStorage Productivity Center for Replication can maintain consistency of any mirrored data, open or mainframe, and runs on UNIX, Linux, or Windows. For more information about this product, see *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548.

8.4.2 GDPS

The GDPS solution include Open LUN Management and GDPS Multiplatform Resiliency for System z.

Open LUN Management

While primarily viewed as a mainframe centric solution for Continuous Availability, some forms of GDPS (those which use the RMC feature code, either Metro Mirror or Global Mirror), can also serve as a control point for managing the remote copy of open systems storage.

In order to accomplish this, the disk system used for mirroring the open systems data must contain a volume of Count Key Data (CKD). This allows for FICON® connectivity and visibility into the Disk System. It should be noted that GDPS will not recover open systems servers or applications, but only serves as a control point for the mirroring functions themselves.

GDPS Multiplatform Resiliency for System z

While most non-z/OS platforms are limited to maintenance of consistency and replication functions in terms of GDPS support, some Linux Distributions for System z can benefit from full GDPS protection. This means that GDPS, when enhanced with Multiplatform Resiliency for System z, will provide support for reconfigurations of LPARs and even support for HyperSwap.

This feature is available only for GDPS/PPRC.

Further information about GDPS is available in *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548

8.5 Summary

As customer environments continue to proliferate, a major coming trend in customer Business Continuity requirements is heterogeneous platform Business Continuity. The two levels of available heterogeneous Business Continuity solutions are:

1. Multiple applications, multiple platforms
   - Is application agnostic
   - Provides hardware and data integrity
2. Single application or database, multiple platforms
   - Provides transaction Integrity
   - Depends on the application to interlock the heterogeneous platforms

At the current time with the current technology, there are not consolidated heterogeneous Business Continuity solutions outside of these two possible configurations.

In order to ensure that data is in a state where the database can be *restarted*, it is important to avoid the *rolling disaster* scenario:

- A rolling disaster occurs when failures within a data center are spread out over time, as opposed to occurring at one instant.
- Consistency Groups can help with rolling disaster avoidance by protecting the consistency of the recovery data, but are not comprehensive protection in the DS6000 and DS8000 configurations.
- Control Software, such as GDPS or TotalStorage Productivity Center for Replication, can assist with ensuring that the data is consistent through a freeze while also providing a management platform for better control of the mirror in a heterogeneous environment.
- Asynchronous replication maintains internally consistent data but can still benefit from control software's management benefits.

Transaction integrity heterogeneous solutions are dependent on software functionality. Software that can support this functionality include (the list is not all inclusive):

- SAP
- DB2
- WebSphere
- MQseries

The decision to go with application based mirroring or hardware based block level mirroring is complex and includes many areas of consideration.

9

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.

9.1 Small and medium sized business overview

SMBs play a very important role in the worldwide economy. As an example, according to U.S. 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 U.S., SMB represents nearly 99.7% of all employers, responsible for nearly three quarters of the net new jobs created to the U.S. economy in the last three years. SMB accounts for over half of the U.S. private work force and drives over 40% of private sales.

While SMB statistics 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.

9.1.1 SMB company profiles and Business Continuity needs

Recent natural disasters and terrorist threats have put Business Continuity as a top priority for enterprises worldwide, creating a sense of urgency within SMB companies to gear up their Business Continuity capabilities because 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. The key SMB Business Continuity drives are:

- 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 U.S. 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 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 Business Continuity planning usually is an afterthought. With increasing pressure from the lines of business and Business Continuity technologies and solutions becoming more affordable and simple, SMB IT management is moving Business Continuity projects forward. In some cases, Business Continuity can be used 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 up-time, planned outage windows are smaller by the day, affecting the ability to back up systems, applications, and data adequately. Some data might not be backed up at all, exposing the business to liabilities and losses.
- Inefficient tools—because 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—Business Continuity 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. Business Continuity 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 Business Continuity implementation.

In this chapter, we answer the following questions related to the successful planning, evaluation, and implementation of Business Continuity solutions for SMB companies:

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

9.1.2 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 might 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.

9.1.3 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 in SMB IT staff.

9.2 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 Business Continuity 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.

9.2.1 Major SMB Business Continuity design components

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

- Prevention Services
- Recovery Services

Because budget and value are the decision criteria for SMB companies, recovery services are usually the starting points for Business Continuity. As prevention services are becoming more affordable, usually Business Continuity 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. The aspects of Business Continuity that fall under prevention services are:

- **High Availability**—builds 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—**minimizes data center operation impacts on up time, including 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. You can find more information about Continuous Availability solutions 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. The aspects of Business Continuity that fall under prevention services are:

- **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 dictate the degrees of prevention and recovery services that are required and the budgets that support them. Usually it is a decision on risks—a balance between the avoidance costs and Business Continuity solutions investments.

9.2.2 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 forth. 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 Business Continuity measure necessity, design and budgets.

9.3 Successful SMB Business Continuity planning and implementation

We recommended the following planning steps, especially for the SMB enterprise:

1. Conduct a risk assessment to develop a set of Business Continuity service targets and IT metrics for key business processes with lines of business.

   The assessment results should determine Business Continuity 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 a comparison of and an understanding of the challenges to meet these 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.

   Business Continuity 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 Business Continuity strategy and plan.

   It is important that the IT Business Continuity plan coordinates with the overall business plan.

9.3.1 SMB Business Continuity implementation steps

We recommended the following 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 Business Continuity 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

9.3.2 SMB Business Continuity affordability

There are two major factors in assessing affordability:

- How much the business can afford to lose
- How much the business can afford to pay

Basically, this is a risk and investment assessment. It is somewhat similar to a home owner's insurance. Although Business Continuity 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

To calculate affordability, especially for the SMB enterprise, you need to determine your *recovery objective* by asking the following questions:

- 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?

You also need to determine you *budget objective* by asking the following questions:

- 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).

9.4 SMB Business Continuity solution components

The table in this section lists the components that typically make up a cost-effective SMB Business Continuity solution, at differing levels of recovery. While your results will vary according to your specific requirements, this table can provide a good beginning guideline. You can use it to build a chart for your enterprise.

Figure 9-1 shows the typical SMB Business Continuity 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 9-1 SMB typical Business Continuity solution components

The definitions of these terms, the tiers, and the various types of solution components is in Appendix B, “Terms and definitions” on page 357 and Chapter 4, “Tier levels of Business Continuity solutions” on page 137.

9.4.1 Typical SMB Business Continuity solutions: Performance and downtime

This section shows the performance and downtime characteristics of typical Business Continuity solutions in the SMB environment. While your results will vary according to your specific requirements, the table in this section can provide a good beginning guideline to what you might expect at differing tier levels of recovery.

You can use the table in Figure 9-2 to build a chart for your enterprise's Business Continuity solution. Components for typical SMB Business Continuity solutions are described in their respective chapters. The definitions of these terms, the tiers, and the various types of solution components are in Chapter 4, "Tier levels of Business Continuity solutions" on page 137.

Typical BC Solution by Tier	Performance in event of unplanned outage	Downtime (typical)	**Typical Solution Components** Components are cumulative, each solution has as pre-req, 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, 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 available at all	> 72 hours	

Figure 9-2 Typical SMB Business Continuity solution - performance and downtime characteristics

The components shown in Figure 9-2 are not the only components or products that can be part of the solution. These are meant as general guidelines. The products shown are typical and can 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.

9.5 Summary

Business Continuity is, and will be, a key requirement for SMB to conduct business. Solutions and technologies continue to improve and 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.

10

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 that provides 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 Disaster Recovery site that is separated by unlimited distances to provide Disaster Recovery protection against regional events.

This chapter provides basic information regarding networking options that you need to consider prior to designing two-site or three-site network infrastructures. Choosing an appropriate networking infrastructure is an often underestimated but a 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 book to provide detailed guidance but a network expert should be an early and constant part of the Business Continuity solution design team.

10.1 Network topologies

Remote storage and remote backup are key components in either a continuous availability, rapid data recovery, or backup/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 if selected.

Figure 10-1 shows two-site and three-site recovery (*cascaded*) topologies which 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.

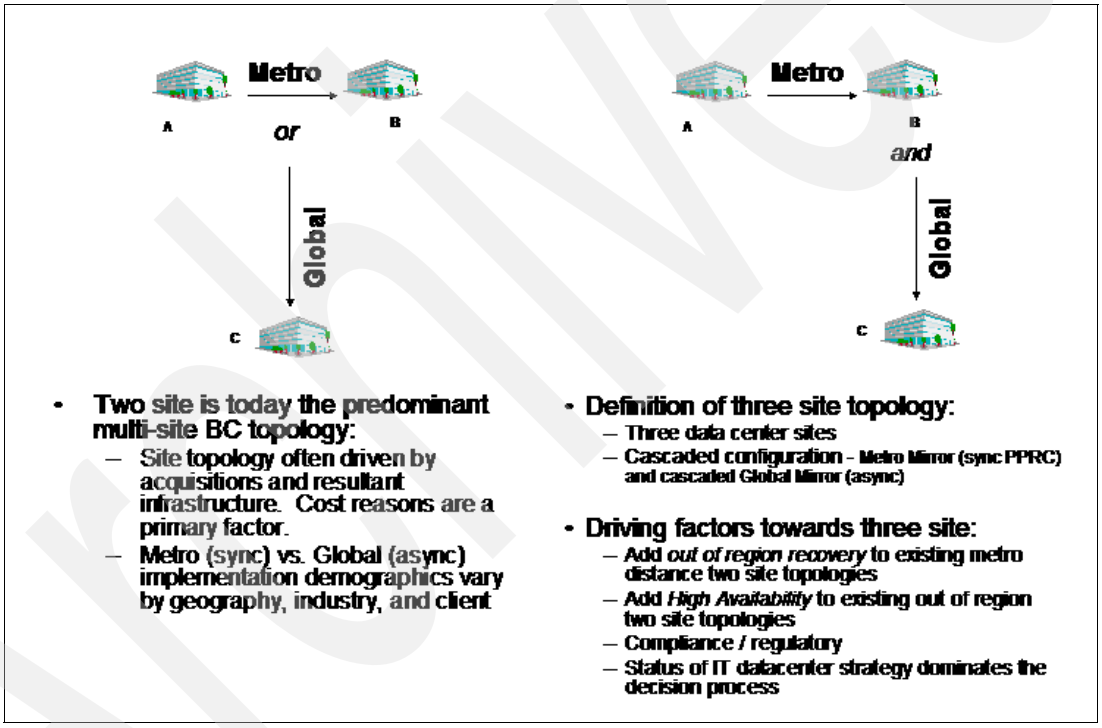

Figure 10-1 Two-site and three-site topologies

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 questions:

- Distance. How far away is the remote site?
- Bandwidth. How much data is transported in what time frame?
- Recovery Point Objective (RPO). How much data can you afford to lose?
- Recovery Time Objective (RTO). How long can you afford to be without your systems?
- Network Recovery Objective (NRO). 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.

10.2 Fiber transport

This section defines the two major options in fiber transport technology: dedicated fiber and Synchronous Optical NETwork (SONET).

Managed Services Service Level Agreements (SLAs): Telecommunication and Network Connectivity providers will often be interested in offering managed services for dedicated fiber links in addition to SONET rings which are managed services by their very nature. In this case, dedicated fiber is not purchased outright but leased, although the performance benefits of having dedicated fiber strands remain.

Under these arrangements there will be SLAs set forth as part of the service contract. These are extremely important within disaster recovery solutions because they set forth the requirements under which the network provider will handle availability and the penalties associated with not abiding by those requirements. When negotiating SLAs make sure that RTO, RPO, and NRO (see Appendix B, "Terms and definitions" on page 357) form the basis of discussion, because they are dependent on the managed service provider maintaining network connectivity.

10.2.1 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.

10.2.2 SONET

Synchronous Optical NETwork or 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 10-2 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.

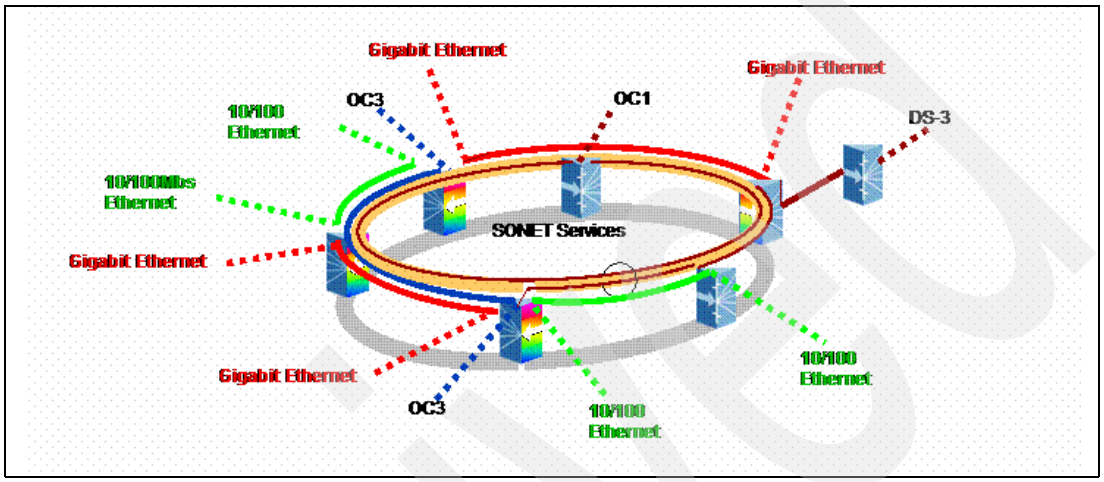

Figure 10-2 SONET

10.2.3 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 or its capacity to move data as measured in millions of bits per second (Mbps) or 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.

Table 10-1 shows a comparison of the various transport links that are available.

Table 10-1 Data transport link bandwidth and speed

Link Type	Bandwidth	Speed
Copper Telco Links		
T1/DS1	1.5344	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
Fibre Channel/FICON	1	Gbps
	2	Gbps
Ethernet Links		
10Base-T	10	Mbps
100Base-T	100	Mbps
Gigabit Ethernet	1	Gbps

It is also useful to compare some of the relative line speeds as shown in Table 10-2.

Table 10-2 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 might 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 can be unacceptable. Latency is a particularly difficult challenge because, unlike bandwidth, spending more money for higher speeds will not reduce latency.

10.2.4 Technology selection

Figure 10-3 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.

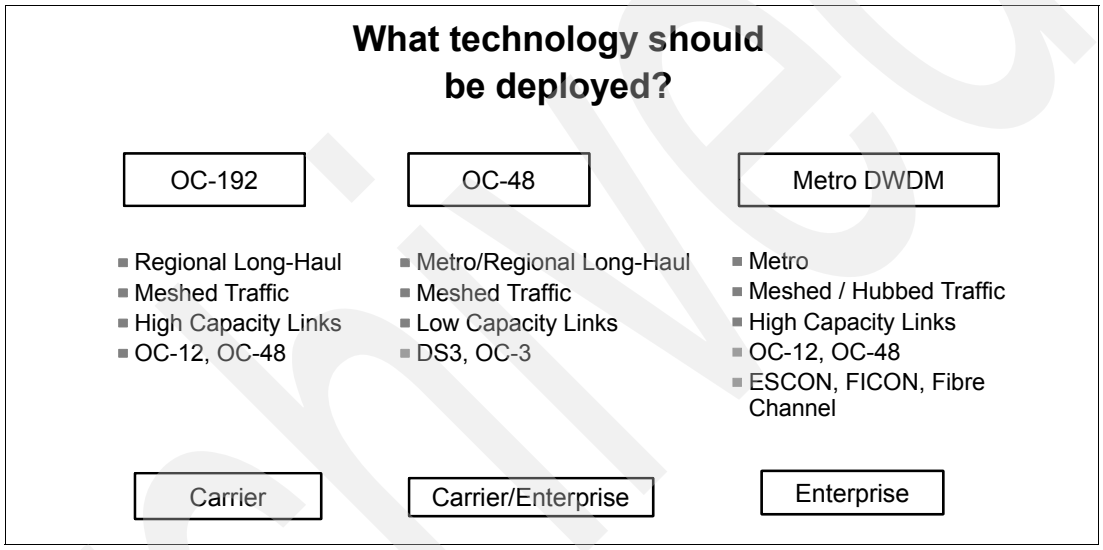

Figure 10-3 Network transport technology selection

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 don't 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).

10.3 Wavelength Division Multiplexing

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 10-4 demonstrates the flow as the channels connect to the WDM and enter the fiber strand.

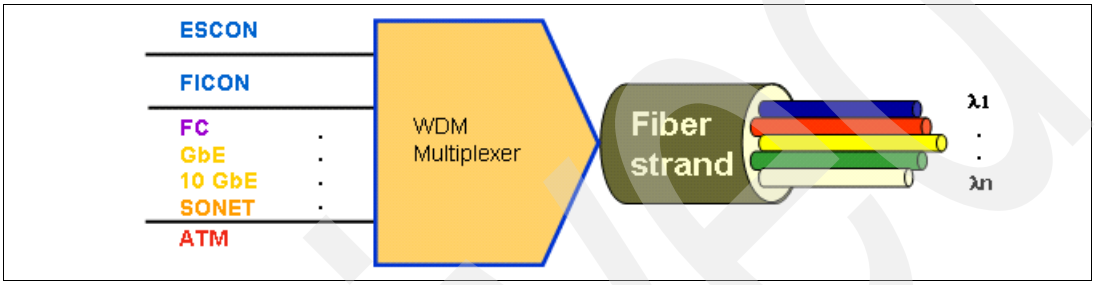

Figure 10-4 WDM conceptual flow

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 determine whether it is supported in your configuration.

10.3.1 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 through 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.

10.3.2 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

Coarse Wavelength Division Multiplexing (CWDM) spreads out the light waves 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

The most common form of multiplexing, Wavelength Division 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 1 pair of fiber, and this can be multiplied through technologies such as subrate multiplexing.

10.3.3 Subrate multiplexing

Current standards in WDM technology allow devices to split the light within one pair of fiber strands into 8 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 2 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, might not be supported due to the minor variances in timing. Other channel types are far more tolerant.

10.4 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, might make better financial sense than using dedicated fiber. Figure 10-5 shows how channel extension can be used to transfer large amounts of data between data centers.

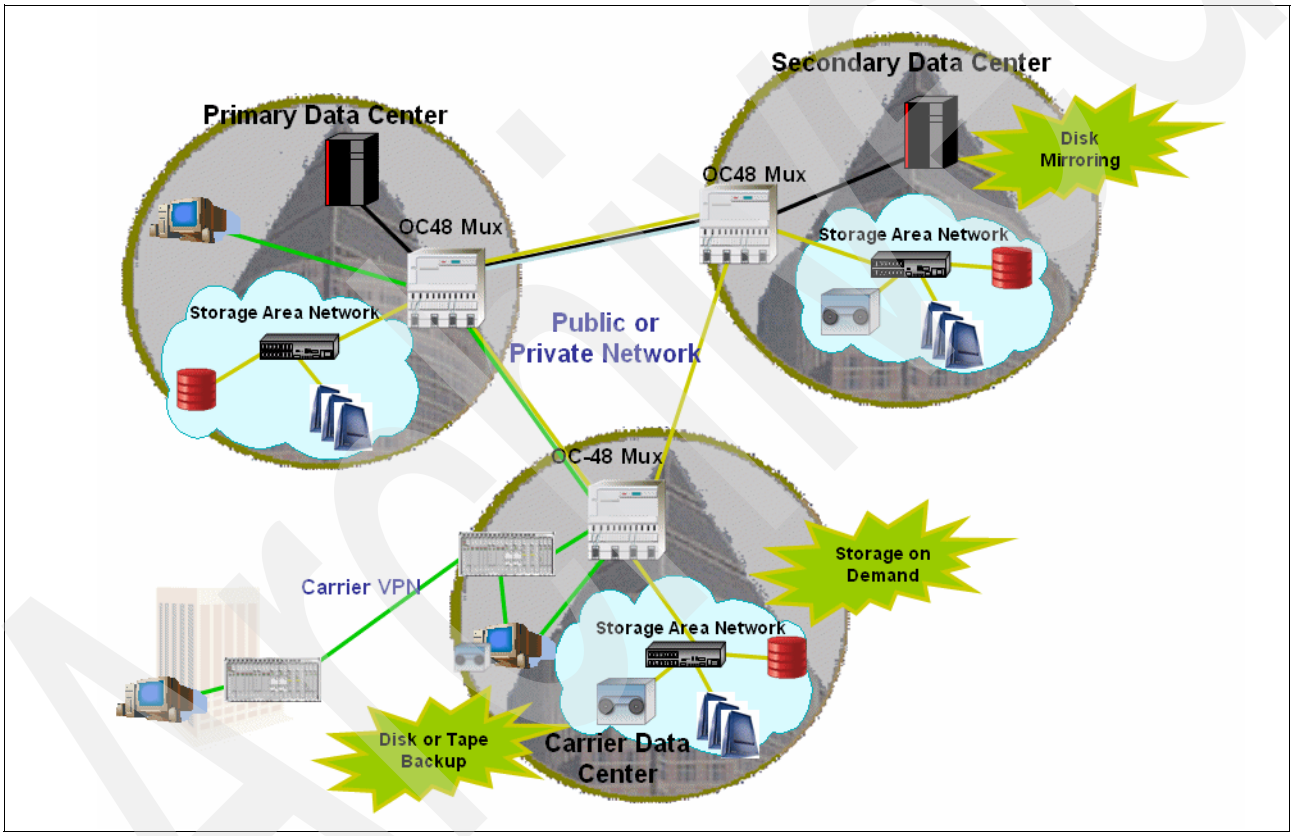

Figure 10-5 Example of channel extension implementation using SONET

10.4.1 Methods of channel extension

Because SONET is used over such long distances, special methods must be used to ensure that devices don't 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 which 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 might prove to be more cost effective than other channel extension technologies.

There is one word of caution, however, when investigating Fibre Channel over IP. Large disk mirroring implementations can take up a sizable amount of the available bandwidth and their use can 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.

10.5 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 can 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 might not always be a workable solution, especially in solutions that transport over continental distances. In these cases, an alternative might 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.

10.6 Bandwidth sizing

When designing a storage replication solution across sites, sooner or later one fundamental question will be asked: What bandwidth must be put in place to satisfy the remote write requests while maintaining adequate performance on the local production site? This section addresses these problems.

10.6.1 Concepts

We start by illustrating some fundamental concepts and seeing how these concepts are correlated.

RPO

RPO defines the amount of data that can be lost in the case of a disaster. It is used as a design factor: For example, the customer states that he can accept to lose a maximum of three hours worth of data updates in a disaster. In asynchronous replication mode the RPO can vary with the application write activity and the available bandwidth for replication.

Synchronous and asynchronous

Assume that we have two storage systems located in two locations separated by a certain distance, as shown in Figure 10-6.

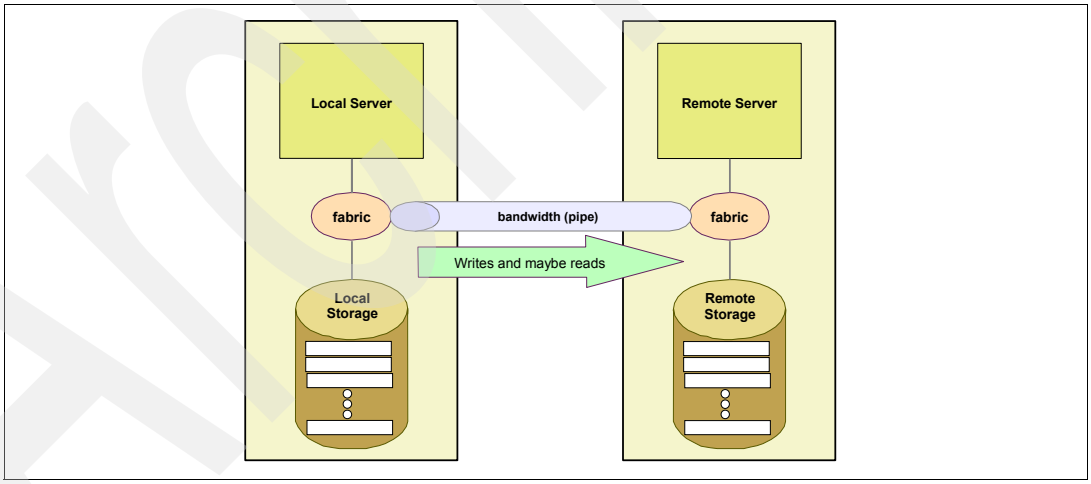

Figure 10-6 Replicated storage in two sites

The objective is to write updates in both sites so as to have a valid copy of data in the event of a disaster. Writes need to be replicated from the local site, where we assume our application is running, to the remote site. Replication operations can be either synchronous or asynchronous:

- *Synchronous* is when updates (writes) are written to both the local and remote site and then, when both write operations are complete, the application is notified that the write completed successfully. Synchronous replication is usually very sensitive to both

bandwidth and latency. These factors can negatively influence I/O and application response times.

- *Asynchronous* is when updates (writes) are written to the local site and then the application is notified that the write completed successfully. The write operation is buffered locally and sent to the remote site at a later time, depending on the bandwidth available. Asynchronous replication is impacted much less by latency; a reasonable bandwidth must be guaranteed so that the data gets replicated in a timely manner, to guarantee an agreed RPO.

Bandwidth

In storage terms, bandwidth measures the amount of data that can be sent in a specified amount of time. Networking link bandwidth is usually measured in bits and multiples of bits. We speak of kilobits and megabits, respectively thousands or millions of bits. Seconds are usually used for time measurement so we can speak of Kbps or Mbps.

Applications issue read and write requests to storage devices and these requests are satisfied at a certain speed commonly called the data rate. Usually disk and tape device data rates are measured in bytes per unit of time and not in bits. One million bytes per second will be expressed as 1 MBps. Current technology storage device LUNs or volumes can manage sequential sustained data rates in the order of up to 80-90 MBps. In other terms, an application can be writing to disk at 80 MBps. Assuming a conversion ratio of 1 Megabyte to 10 Megabits,. we can have a data rate of 800 Megabits. It is always useful to check and make sure that one correctly correlates MBps to Mbps.

Latency and the speed of light

Bandwidth is not the only limiting factor. Latency can have a large impact on application performance. Latency is the time that is required to send a command to a remote server or storage device and receive an answer. The longer the time, the greater the performance impact. Latency depends on the speed of light (c): c = 3 x 10^8m/second. (vacuum) = 3.3 µsec/km (µsec represents microseconds, one millionth of a second). The speed through glass fiber is around 2/3 of c, around 5 µsec/km. These values appear very small.

Now a SCSI write over Fibre Channel requires two round trips per I/O operation, we have 2 (round trips) x 2 (operations) x 5 µsec/km = 20 µsec/km. At 50 km we have an additional latency of 20 µsec/km x 50 km = 1000 µsec = 1 msec (msec represents millisecond). Each SCSI I/O has 1 msec of additional service time. At 100 km it becomes 2 msec additional service time.

We can see that latency can negatively impact applications with high I/O rates, short IOs and fast response times, and these are usually transactional database applications.

Protocols and latency implications

Discussing latency we have seen that SCSI write operations over Fibre Channel require two round trips per I/O operation. The number of round trips is protocol dependent. Some protocols require more round trips than others. In the previous section we have seen how this can impact latency and response time.

- The ESCON protocol requires approximately 6 protocol exchanges per I/O operation, this means a signal latency of approximately 60 µs/km round trip. For sites that are separated by 10 km, this equates to a latency of 600 microseconds, or 0.6 milliseconds. For sites that are separated by 100 km, this equates to a latency of 6000 microseconds, or 6 milliseconds.
- The ESCON PPRC implementation on the Enterprise Storage Server® (ESS) typically requires 3 protocol exchanges per remote copy I/O operation. Users will typically

experience propagation delays of about 1 ms per 30 km of site separation distance. For a 10 km site separation, this equates to a latency of about 0.3 milliseconds. For a 100 km site separation, this equates to a latency of about 3 milliseconds.

- The FICON protocol requires only 1 protocol exchange per I/O operation, which in turn results in a signal latency of approximately 10 microseconds per km site separation. round trip. For a 100 km site separation, this equates to a latency of about 1 millisecond.
- The Fibre Channel Metro Mirror implementation on the DS8000 requires 1 protocol exchange per remote copy I/O operation.

WANs are usually characterized by high latencies and limited (in other words costly) bandwidths compared to storage, SANs and LANs. While investments in networking infrastructure combined with greater bandwidth demand can result in cost-effective bandwidths over the wide area, the relatively higher latencies of wide area communications are a problem that is likely to persist for the foreseeable future. Fundamental limitations placed by the speed of light and technological limitations placed by the overhead of routing and switching together make the latency of wide area access at least an order of magnitude worse than storage area network latencies. The latency problem is dictated by the speed of light and by the routing and switching overhead when intermediate routers, switches and bridges are used between the communicating endpoints. In practice, it has been observed that the bulk of the WAN latency is attributed to the router delays and not to the speed of light, so the calculations shown above are optimistic.

10.6.2 Replication implementations

There are various implementations of remote mirroring. These can be implemented either in software or in hardware.

LVM mirroring

Many operating systems rely on an underlying logical volume manager (LVM) software layer. Most LVMs offer storage mirroring capabilities across different physical storage volumes or LUNs, and if you place a second set of LUNs onto a remote storage device you have a remote mirroring solution. This approach has advantages and disadvantages.

The advantages include the transparency of failover: If one of the two mirror volumes is lost the application will not stop or crash; in AIX, LVM will log an error to the AIX error log and continue with the remaining volume.

Among the disadvantages are management complexities. Each system has its own LVM to manage. In most LVMs, both writes and reads can be sent to both the local and the remote storage device. So the network transport between sites might be required to carry both write and read traffic.

Note: Depending on the LVM, it is possible to read from the local site and send only writes to the remote and local site. AIX LVM is one example of this capability.

A second, small, disadvantage is that we use server CPU cycles to perform the data replication, although the amount of CPU cycles used for I/O normally is very small: 1-2% of the system.

Storage system mirroring

Most IBM System Storage systems have a remote mirroring feature. These storage systems mirror only write operations between the local and remote storage systems, reads are

satisfied by the local system. The synchronous implementation is called Metro Mirror and the asynchronous implementation is called Global Mirror.

The advantages of off loading replication to the storage systems are multiple. We have a single point of control and do not use host server CPU cycles to replicate the data.

The disadvantage is that this approach is not transparent to the host server and application. In the event of loss of the primary storage system we switch to the secondary system and this is not transparent to the application.

Software and application replication

Another form of replication is what we call software application replication. In this case the replication is performed by software in the production system, using TCP/IP as a network transport. Application examples that offer replication solutions are DB2, Domino® and Oracle. This category also encompasses file replication solutions such as NSI Doubletake, VERITAS Volume Replicator, and VERITAS Storage Replicator.

DB2 UDB offers a High Availability Disaster Recovery (HADR) feature. HADR is a replication that takes place at the database level. The HADR requires two active systems: A primary and a standby. All database transactions take place at the primary system. The transaction log entries are continuously shipped to the standby machine through TCP/IP. The standby system receives and stores log entries from the primary system and applies the transactions to the database. If the primary fails, the standby can take over the transactional workload and become the new primary machine.

Oracle offers a similar function called Oracle Data Guard. Data Guard performs synchronous or asynchronous log shipping from the production data base to one or more standby databases.

Database replication using database tools is generally more bandwidth efficient than using other forms of replication. To understand why, consider the following two examples:

1. Database update write activity to a file system or volume requires three write operations:
   a. Write before image to the database log
   b. Update the database
   c. Write after image to the database log
2. Database update to a remote system using database tools and log shipping requires one write operation, which is write the update to the remote system.

Other software products perform file system or volume based replication. These products usually intercept write IOs and propagate them to the remote sites.

10.6.3 Measuring workload I/O characteristics

To size the bandwidth required to replicate data between two different locations we need to know the amount of data that is written to disk: The write profile. We need to obtain this information for all the servers that will be involved in replication and must add it together to get the I/O profile of multiple servers.

The write activity measurements should span a reasonable amount of time, days or weeks, and should be gathered in busy periods such as month or year end. A reasonable long term sampling interval is 5 to 15 minutes, and a more detailed analysis can be performed for the peak hour.

There are many measures of I/O activity. For bandwidth sizing one measure is essential: The number of bytes per second that is written to disk. This can be calculated as the average I/O size multiplied by the number of IOs. Different platforms have different measurement tools and return different types of measurements.

z/OS tools

Resource Measurement Facility (RMF™) provides performance information for the DS6000, DS8000, SVC, ESS, and other disk systems for z/OS. RMF Device Activity reports account for all activity to a base address and all its associated alias addresses. Activity on alias addresses is not reported separately, but RMF will report the number of Parallel Access Volumes (PAV) addresses (or in RMF terms, exposures) that have been used by a device, and whether the number of exposures has changed during the reporting interval.

RMF cache statistics are collected and reported by a logical control unit (LCU). So a fully configured ESS would produce 16 sets of cache data. One report only reflects the status of one logical system. In other words, if you want to check out the status of the whole cache, you have to check out all of the 16 reports.

With any RMF reports, the I/O information reported is from one z/OS system only. If you are sharing I/O across multiple systems, you will need to review RMF reports from each of the sharing systems in order to see the complete I/O activity to the LCUs and devices. This is not true for the Cache reports, however; because cache data is obtained from the control unit, it does include I/O activity from all sharing systems. Contention from a sharing system will generally be seen as increased pending (PEND) and disconnect (DISC) times in the Device Activity report.

The RMF FICON Channel Path Activity report shows the MBps read and write data rates for each channel. You must perform aggregation of the data for all the channels attached to the storage systems you want to replicate. Another way of obtaining and summarizing RMF performance data is to use the RMF Magic[1] tool.

RMF Magic is a new tool developed by Intellimagic designed to analyze the I/O workload volume and performance of disk systems attached to one or multiple z/OS or OS/390® system images. The program processes RMF records; an extended period of up to a month can be analyzed in a single study.

RMF Magic creates both a consolidated and comprehensive view of all I/O activity in a z/OS environment. The tool can be used for disk system-related capacity planning and provides special support for sizing Extended Remote Copy (XRC) or Peer-to-Peer Remote Copy (PPRC). It can also be used to debug performance anomalies, in particular those that require an analysis of performance behavior over an extended period of time.

Data presentation is done through Excel® data sheets and charts, allowing the analyst to take advantage of Excel's rich user interface and data formatting capabilities and to adjust the output format and detail as required for the subject under study.

Except for the data collection, every step of an RMF Magic analysis can be done on a PC, made possible by a better than 10:1 reduction of the data volume. This facilitates a process of incremental discovery where each step in the analysis can prompt the need for another view on the source data.

Additional information about RMF Magic is available at:

http://www-304.ibm.com/jct09002c/gsdod/solutiondetails.do?solution=12118&expand=true&lc=en

[1] RMF Magic is a product developed by IntelliMagic and is available from IntelliMagic.

For further information about techniques to gather host system I/O performance data, refer to:

- *IBM TotalStorage Enterprise Storage Server Model 800 Performance Monitoring and Tuning*, SG24-6422
- *IBM TotalStorage DS6000 Series: Performance Monitoring and Tuning*, SG24-7145
- *IBM TotalStorage DS8000 Series: Performance Monitoring and Tuning*, SG24-7146

TotalStorage Productivity Center

Sometimes it might not be easy to collect the information at the server level. You can then choose to evaluate alternatives. Figure 10-7 illustrates the different types of performance data that TotalStorage Productivity Center for Disk can provide when monitoring SMIS-capable storage devices.

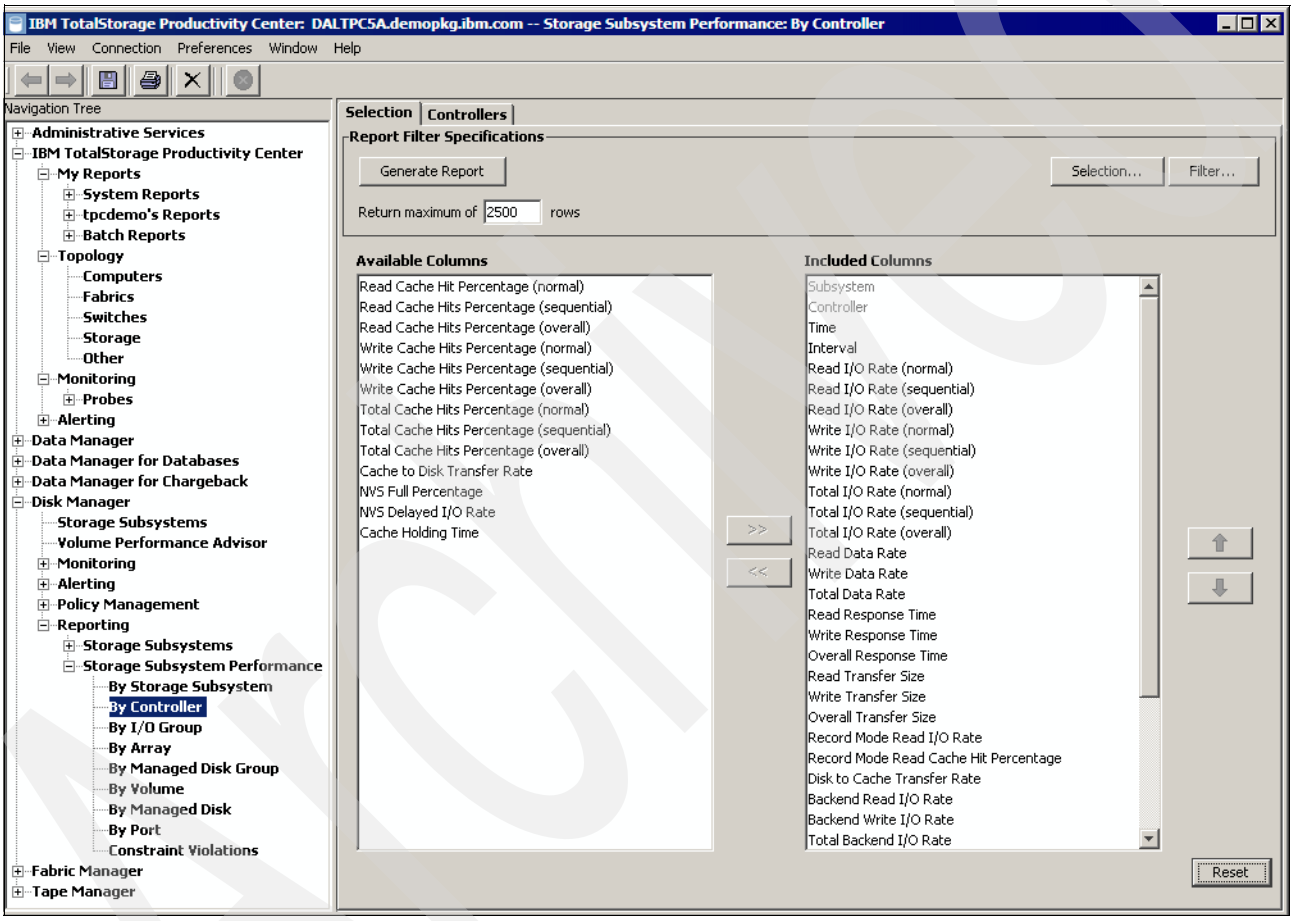

Figure 10-7 TotalStorage Productivity Center for Data disk performance reports

TotalStorage Productivity Center for Disk provides gauges to track real-time performance so that the administrator can monitor performance metrics across different storage systems from a single console. Alerts can also be set to automate appropriate actions based on the exceeding of performance thresholds.

These storage system metrics can help estimate bandwidth requirements for channel extension.

AIX tools

The **iostat** command displays I/O statistics for the server. The **iostat** command gives the read and write kilobytes per second and the transactions per second. Unfortunately it shows

the data by individual disk rather than aggregated, so summarization must be performed manually. Also hdisks could be multiple exposures of the same disk.

A better way to collect AIX write data is to use **nmon** and nmon analyzer. These tools are freely available on the IBM Web site at:

http://www.ibm.com/developerworks/aix/library/au-analyze_aix
http://www-941.haw.ibm.com/collaboration/wiki/display/WikiPtype/nmon
http://www-106.ibm.com/developerworks/eserver/articles/nmon_analyser/index.html

The nmon tool offers a long term collection capability that is documented in the nmon analyzer manual. You collect one day's worth of data by starting the command **nmon -x**. The command produces an output file that is transferred to a personal computer and processed by the nmon analyzer worksheet. The resulting worksheet includes a tab called DISK_SUMM, shown in Figure 10-8, with the write profile.

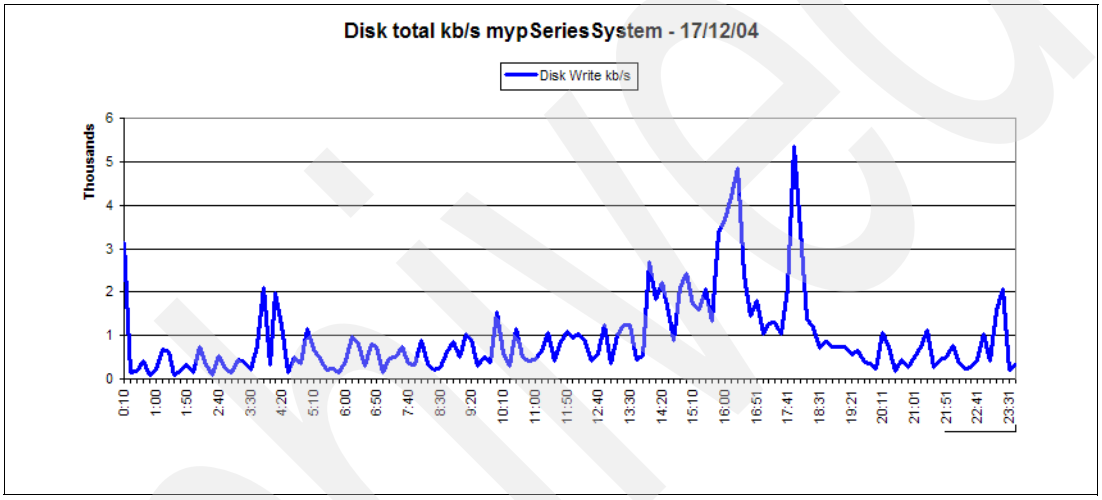

Figure 10-8 I/O write profile

This information is input to the bandwidth sizing exercise. You also need to know the data collection interval, which is 600 seconds in this example.

Linux tools

The **iostat** utility is included in many Linux distributions. The **iostat** command is also part of the Sysstat set of utilities, which is available from:

http://perso.wanadoo.fr/sebastien.godard/

The **iostat** command lets you see a report about the activities of the disk system of the server. The report has two parts: CPU utilization and device (disk) utilization. You can use **iostat -x** to collect detailed I/O statistics including read and write kilobytes per second.

The **nmon** tool is also available for Linux SUSE and Red Hat on Intel and IBM System p platforms. You can use this tool as discussed in "AIX tools" on page 296.

For additional information about collecting Linux I/O performance data see *Tuning IBM System x Servers for Performance*, SG24-5287.

Windows tools

Windows offers the performance monitoring tool perfmon. To collect data over a period of time you can use Microsoft Windows perfmon counter logs. Use the following procedure to collect the write activity data:

1. Start perfmon by selecting **Start** → **Run** → **perfmon**.
2. Select **Counter Logs** from the left pane.
3. Right-click in the right pane.
4. Select **New Log Settings**.
5. Choose a name for the log. A new window opens.
6. In the General tab, select **Add Counters** and select the following counters:
   - PhysicalDisk - Disk Write Bytes/second for all instances
   - PhysicalDisk - Disk Writes/second for all instances
7. Select an interval for the data collection. A 10 or 15 minute interval is adequate for a long running collection.
8. Select **OK** to save the customization.
9. The new counter log appears in the right hand pane. Red indicates it is stopped, and green indicates that it is running and collecting data.
10. To start data collection select the counter log by right-clicking and pressing **Start**.
11. Make a note of the log file name as it will be needed later.
12. The data collection runs until you select the task with the name that you assigned in step 5 by right-clicking and pressing **Stop**.

After the data has been collected, you need to analyze it. We will show how to use the log file that we collected and how to save the information into a Microsoft Excel worksheet:

1. Start perfmon by selecting **Start** → **Run** → **perfmon**.
2. Select the right window pane and press Ctrl+E to reset counters.
3. Select the right window pane and press Ctrl+L to load log data from file.
4. Select **Source - Data source - Log files** and click **Add**.
5. Select the log file name that you recorded in the previous step and click **Open**.
6. Select the Data tab click **Add**.
7. Select the following counters from the list:
   - PhysicalDisk - Disk Write Bytes/second for all instances
   - PhysicalDisk - Disk Writes/second for all instances
8. Click **OK**. A chart with the data you collected opens.
9. Right-click in the chart area and select **Save Data As**.
10. Choose a file name and in the *Save as type* field choose **csv** and click **Save**.

The data in the .csv file can be loaded into a spreadsheet such as Excel. Table 10-3 shows a sample of the data.

Table 10-3 Sample Windows write profile

(PDH-CSV 4.0) (Pacific Standard Time)(480)	\\TARELLA\PhysicalDisk (_Total) \Disk Write Bytes/s	\\TARELLA\PhysicalDisk (_Total) \Disk Writes/s
22:08:45	48678	3
22:09:00	49116	7
22:09:15	14744	2
22:09:30	2730	0
22:09:45	8	0
22:10:00	6894	1
22:10:15	6007	1

System i tools

System i also offers I/O interval reports. These interval reports provide disk utilization details. You can consolidate the information for the peak 15 minute intervals where throughput and response time are relevant for your applications and get the peak I/Os per second, KB per I/O, and also the read-to-write ratio. With this information, you can start to estimate more accurately the bandwidth requirements that will adequately support your IBM System i workload demands. For further information about techniques to gather host system I/O performance data, refer to:

- *IBM TotalStorage Enterprise Storage Server Model 800 Performance Monitoring and Tuning Guide*, SG24-6422
- *IBM TotalStorage DS6000 Series: Performance Monitoring and Tuning*, SG24-7145
- *IBM TotalStorage DS8000 Series: Performance Monitoring and Tuning*, SG24-7146

10.6.4 Determine the bandwidth

In the preceding section, we discussed how to collect the performance data to determine write activity. These writes need to be replicated to the remote system.

Using the I/O write profile to determine required bandwidth

If we have collected write activity from multiple systems and platforms we will first need to summarize it to get a profile that aggregates write activity from the multiple systems. We need to sum the write performance data, interval by interval, for all the systems we want to replicate.

Data summarization across systems and platforms will be easier if all system clocks are aligned and the data collection interval is the same for all systems and all intervals start at the same time, say on the hour.

There is no specific tool or method to collect this data across all servers and platforms. The task will be manual and can be accomplished by using simple scripts and spreadsheet programs. The end result that is needed is the average and peak write data rate.

Synchronous or asynchronous

For synchronous replication we need to determine the peak write rate and choose a network bandwidth that is greater than this peak. The reason for this choice is to not impact the performance of the production applications. If we choose anything less than the peak write rate, the primary application can incur delays due to writes at the remote site. These delays are often felt more by batch applications because they are more write intensive. A second factor to consider is additional latency for the I/O operations and this translates into higher response times as illustrated in 10.6.1, “Concepts” on page 291 under “Latency and the speed of light”.

Asynchronous replication is more complex. In general it supports adequate operations with less bandwidth. The penalty is that RPO will be greater than zero and will depend on the bandwidth and the arrival of writes. Also, depending on the replication implementation, data writes might need to be stored in a temporary storage area or log and the amount of log space must be determined.

Tools

There are various IBM internal tools available to estimate RPO for asynchronous replication solutions. These tools will be normally used by your IBM replication specialist.

In an asynchronous replication implementation a model can be built to determine RPO and log size when we have the application write I/O profile and the bandwidth. The relationship can be established between write I/Os, bandwidth and RPO and log size. We will illustrate a conceptual model and show some examples of results using this model.

Assume we have an arrival rate of x units of data per interval, for example *x* MBps average, and a bandwidth or pipe between sites of *y* MBps. We perform the following calculation for each interval:

1. Interval 1: If x1<y all data can be sent in interval.
2. If x1>y only part of the data (y) can be sent and we have x1-y=r1 (remainder) left over. The remainder needs to be sent in the next interval, interval 2.
3. Interval 2: we calculate how much data to send: x2+r1=x2’.
4. If x2’<y it gets all sent.
5. If x2’>y then we only send y and send the remainder (x2’-y=r2) to the next interval, interval 3.

This process is repeated for all intervals.

The calculation we perform is not mathematically foolproof, for example, if we perform no writes in 90% of the interval and all writes in the last 10% of an interval, then the assumption that the average data rate (*x* MBps) is less than pipe size (*y*) no longer holds true. On average we have an acceptable approximation.

What do we obtain? The remainder we determine in each interval (*r*) represents the RPO, data that has not yet been sent to the remote site (calculated as [*r* MBps / bandwidth MBps] * intervalSec). The maximum remainder is the maximum RPO. Also if we multiply the maximum remainder (*r*) MBps by the interval length we get the maximum log space required.

A sample worksheet is shown in Figure 10-9. Given an I/O profile (Write kb/int) and a bandwidth or pipe size of 30 Mbits we calculate a max RPO of under 14 minutes and 2 GB of local log space.

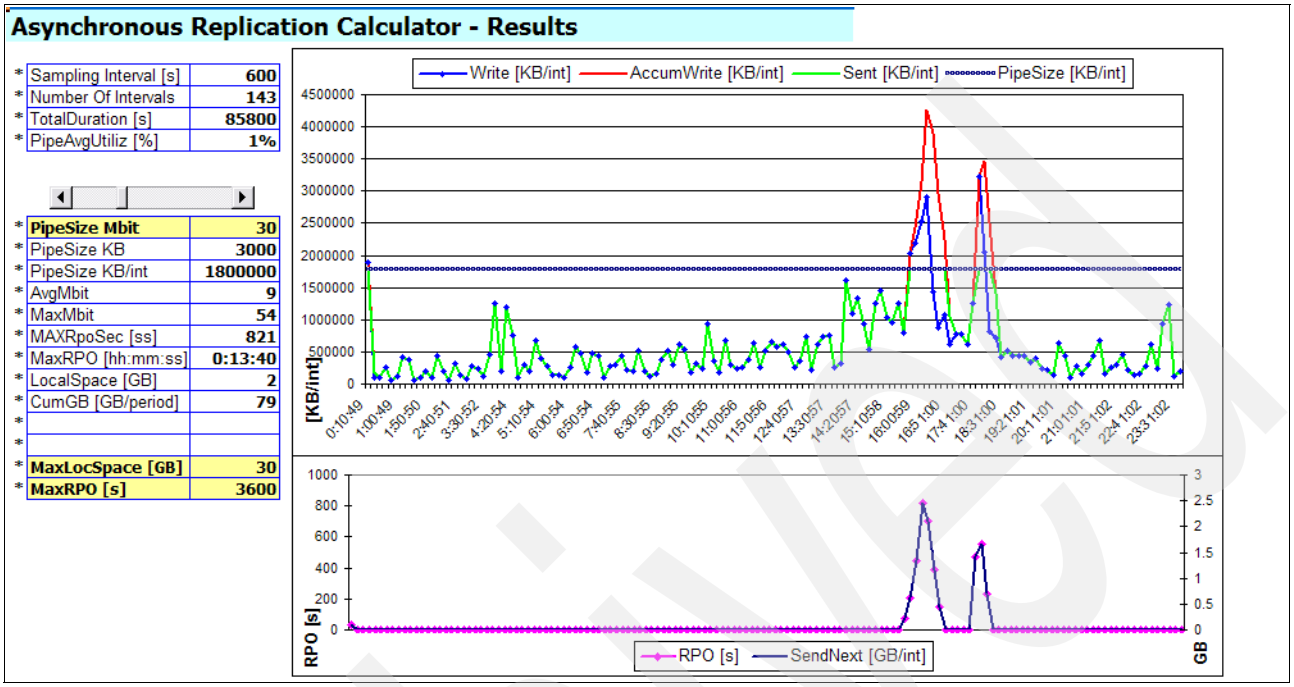

*Figure 10-9 Asynchronous replication with 30 Mbit bandwidth - **- original is in RPOCalc.XIs if want to change***

We can compare the above results with the behavior calculated in Figure 10-10. We changed the bandwidth to 5 Mbits. In this case we have a maximum RPO of over 20 hours in a 24 hour interval. The value AccumWrite, representing data to be replicated that has yet to be sent to the remote site, is growing in an uncontrolled way. It is extremely unlikely that a 5 Mbits pipe can satisfy this workload.

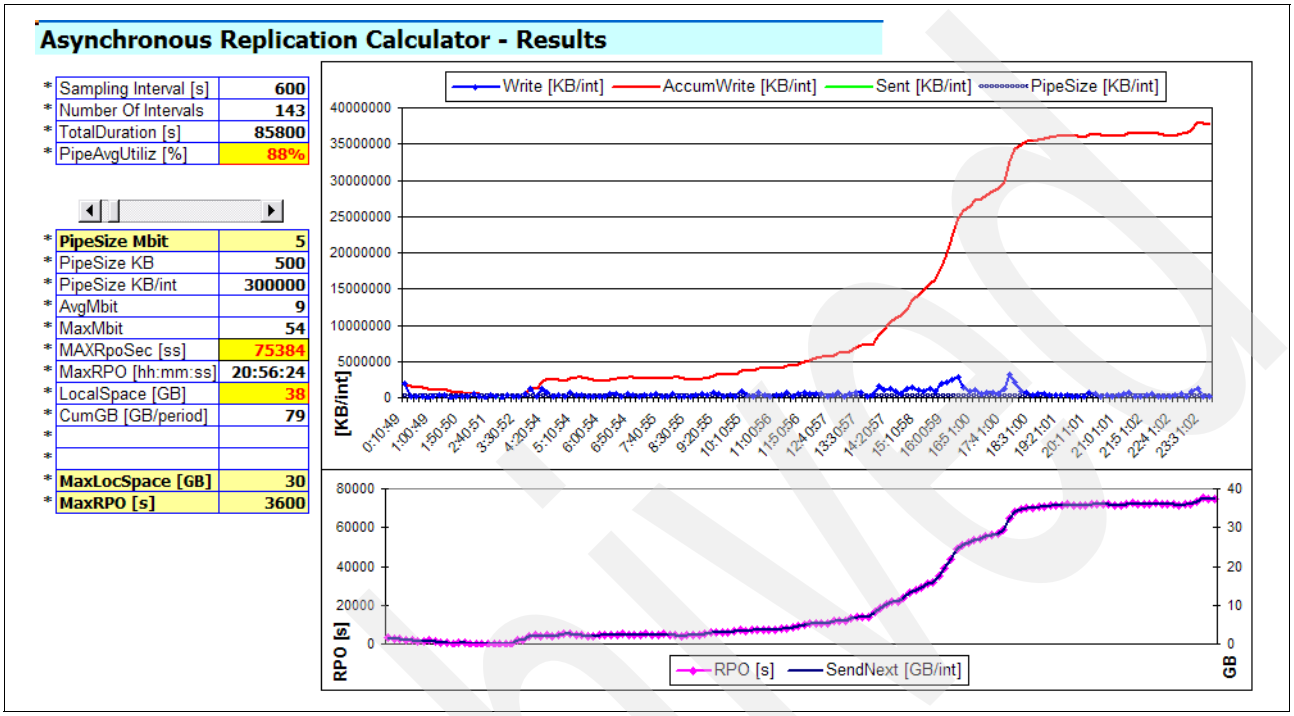

Figure 10-10 Asynchronous replication with 5 Mbit bandwidth

The model also calculates RPO and local write log space and plots these values over time. It can be used to evaluate what-if scenarios and quickly and intuitively determine an adequate bandwidth.

Rules of thumb

If no information is available, we can give a first rough estimate of writes using rules of thumb based on industry averages. Note that your installation might be completely different.

Computer Measurement Group[2] is a worldwide association of IT system performance professionals. CMG user groups meet several times a year, worldwide, to exchange IT performance information. Over the past five years, CMG has documented publicly that a reasonable worldwide average disk performance measurement is that every 1 GB of disk data produces, on average, a little less than 1 I/O per second per GB. This estimate is reasonable for both mainframe and open environments. This metric is called access density.

We can use this access density fact to derive a straightforward Rule of Thumb (ROT) for estimating the amount of production I/Os that are generated by a given amount of disk GB or TB of data. Using the access density above we can estimate the amount of write data per second that a given amount of data will produce.

[2] http://cmg.org

There are two rules of thumb using an average, conservative, access density of 1 I/O per second per GB are:

- 1 TB of production OLTP disk storage generates about 1 MBps of write data.
- 1 TB of production BATCH disk storage generates about 6.75 MBps of write data.

The derivation of the two rules of thumb is outside of the scope of this discussion. Worldwide, CMG white papers have shown that these estimations are valid. Using this information, we can calculate and extrapolate the number of transactions or batch updates that might be represented by a given amount of production disk mirrored storage.

11

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 book, but we refer to them when it is appropriate.

This chapter first explains the methods on how to make systems highly available. Then it introduces the technologies and functions that are suitable for disaster recovery and fault tolerant solutions for databases.

11.1 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. In Chapter 3, “Business Continuity planning, processes, and execution” on page 43, we explain the basic disciplines of Disaster Recovery and the tasks which must be in place.

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 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 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 use technologies that provide Continuous Availability provide the necessary IT infrastructure as building blocks for business continuance depending on the availability requirements, as shown in Figure 11-1.

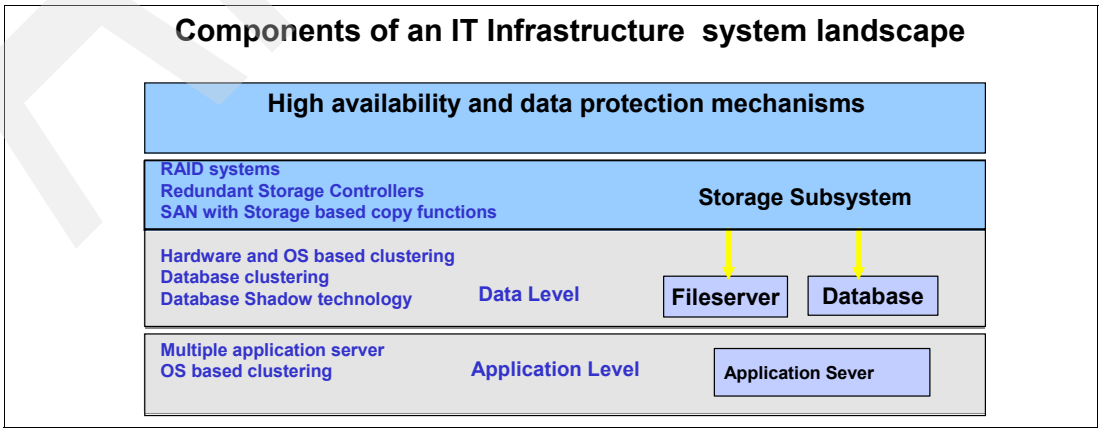

Figure 11-1 Components of IT infrastructure system landscape

11.1.1 Selecting IBM Server and System Storage solutions for High Availability

IBM Server and System Storage products are built and proven for availability, scalability and reliability. In combination with IBM system management products, such as IBM Director or Tivoli Enterprise™, a system (or a system's component) automatically notifies an administrator prior to a failure of a suspicious component which is likely to fail, because a built-in hardware threshold for error correction has been exceeded.

Hardware redundancy prevents downtime caused by hardware failures by detecting a failing component before it actually fails and bypassing a failure when it does occur. Most of the IBM servers provide predictive failure analyses (PFA) for their most critical components such as processors, memory, disks, voltage regulators, fans, adapters and power supplies. Therefore a problem on a single component can be detected and proactively (hot) replaced without any system outage or maintenance window. This is part of the strategy from IBM of autonomic computing. For more information, see *Getting Started with IBM Tivoli Monitoring 6.1 on Distributed Environments*, SG24-7143 and *Automated Distribution and Self-Healing with IBM Tivoli Configuration Manager V 4.2*, SG24-6620.

Key terms that are related to IBM Server and System Storage solutions for High Availability include:

- **Redundant power supplies**

  Redundant server and disk array power supplies provide a secondary power supply if the primary power supply fails.

- **Redundant fans**

  Redundant fans ensure that sufficient cooling exists inside of the server if a cooling fan fails.

- **Redundant storage system**

  A redundant storage system provides protection against the failure of a single disk drive or controller.

- **Redundant memory**

  Redundant memory provides memory if a memory bank fails.

- **IBM Chip Kill memory**

  Error-correcting code (ECC) memory detects and corrects single-bit errors and takes the memory offline if a double-bit error occurs. IBM Chipkill™ Memory enables increased availability by detecting and correcting multiple-bit memory DIMM errors.

- **Redundant Bit Steering**

  Redundant Bit Steering is similar to the hot-spare of a disk array which utilizes unused bits in each memory DIMM (hot spare bits). It can utilize these spare bits in the event of a failure in a DIMM. This dramatically increases the memory reliability, especially in multi-gigabyte main memory capacities in combination with in-memory databases.

- **Redundant network interface cards**

  Using redundant network interface cards (NICs) ensures that clients can connect to the data center if a NIC or a network connection fails. In addition to redundancy some network adapters provide load balancing for better performance as well as for High Availability.

- **Power-on monitoring**

  When the server is initially turned on, during the power on self test (POST) the server detects startup failure conditions, such as abnormal temperature conditions or a failed fan.

- **Lock-stepped processors**

  Lock-stepped processors are two processors that execute the same instruction stream and that crosscheck each other as with MARATHON Technologies. For more information, see their Web site:

  http://www.marathontechnologies.com

Parity checking on internal buses using IBM server hardware minimizes the likelihood of hardware failures and uses monitoring and redundancy to limit the effect of a failure on data availability.

In Chapter 4, “Tier levels of Business Continuity solutions” on page 137, we explain the hierarchical dependencies of the system layer architecture. When we refer to the system dependencies with databases, you need to consider the components of the system layer architecture for the requirements of High Availability within a system, as shown in Figure 11-1 on page 306.

11.2 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, in addition to 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
- Geographically dispersed clusters

In the following sections, we describe the basic cluster technologies and how you can deploy them as highly available building blocks in disaster tolerant solutions in conjunction with databases.

11.2.1 Shared nothing clusters

A s*hared nothing cluster*, shown in Figure 11-2, is defined as a system of independent servers (two or more) with access to a common storage system but operating 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.

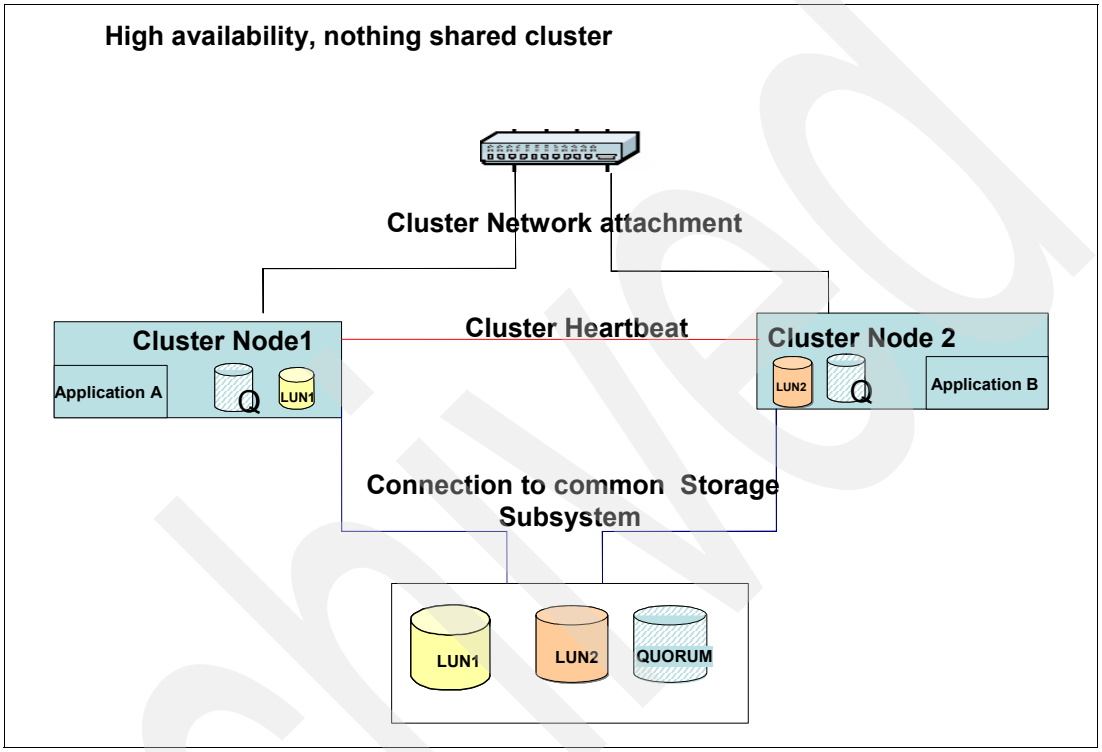

Figure 11-2 High Availability, shared nothing cluster

The nodes can be configured in two ways: either *active/active* or *active/passive*. In 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 failover to the surviving node, as shown in Figure 11-3. The cluster heartbeat is a dedicated network which 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 8 nodes. IBM AIX HACMP is another example of a nothing shared cluster.

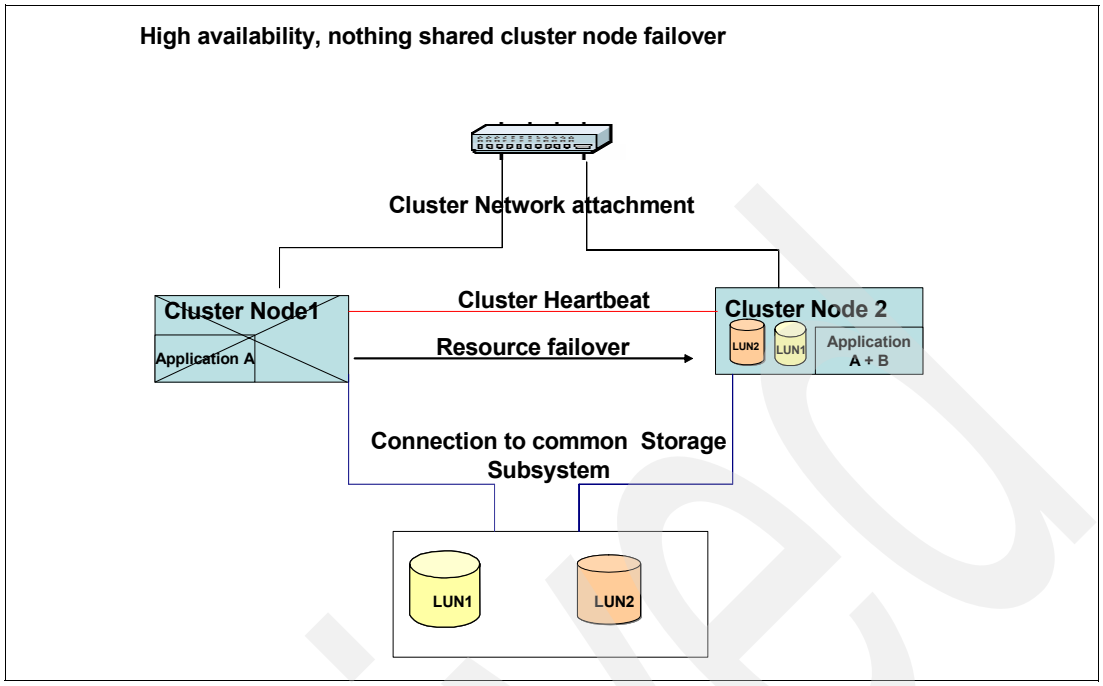

Figure 11-3 Failover within a nothing shared cluster

11.2.2 Common shared cluster

Physically, a common shared cluster can consist of the same or similar building blocks as the nothing shared 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 which enables data locking.

Figure 11-4 is a common shared cluster schematic.

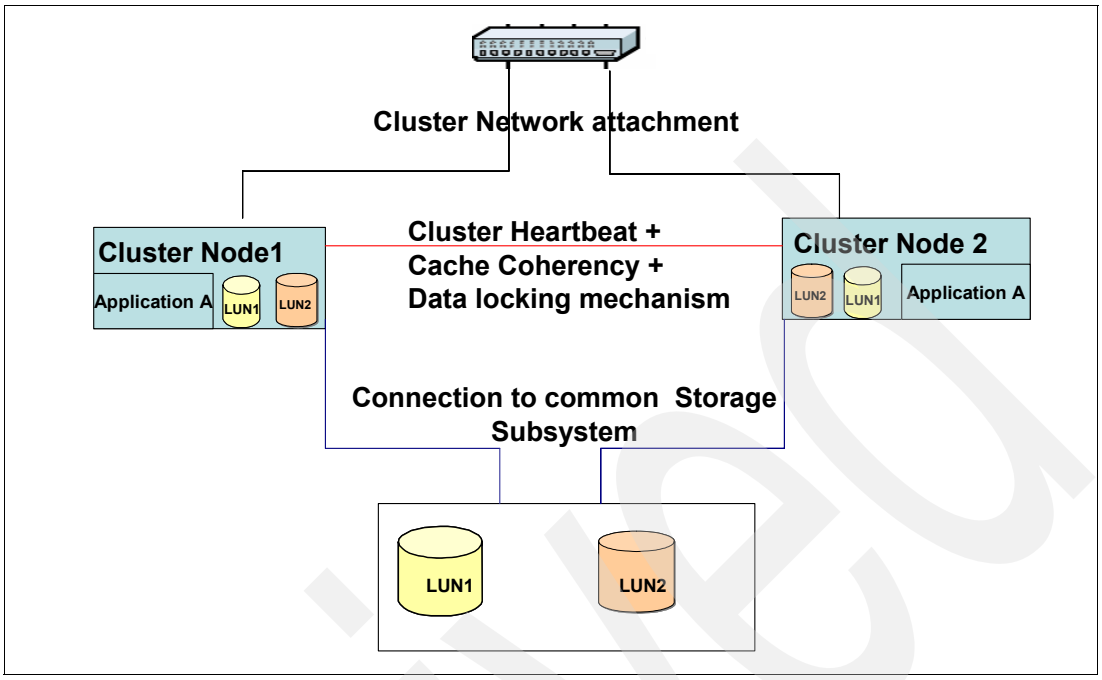

Figure 11-4 Highly available and scalable common shared cluster

11.2.3 Shared nothing application cluster

In a shared nothing application cluster, shown in Figure 11-5, the users are linked to a farm of independent cluster nodes which act as one single system. Each node has its own data and the load between the nodes is distributed and coordinated by the application's 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 up to the crash of the individual node might be lost depending on the application's behavior. An example for an application server farm is CITRIX or IBM WebSphere.

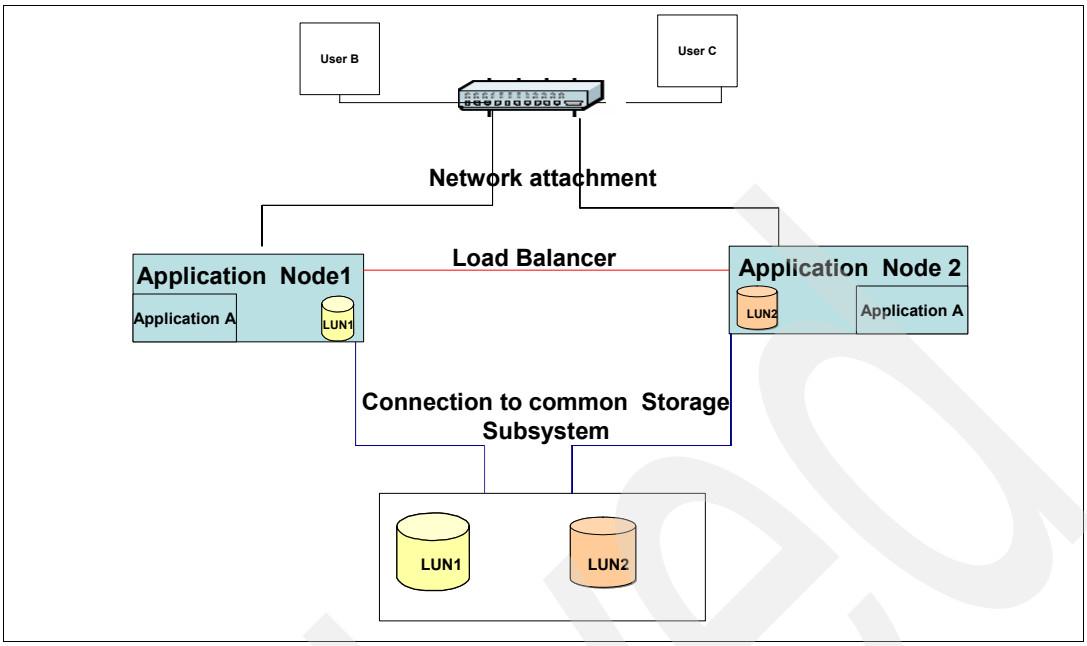

Figure 11-5 Nothing shared application cluster

11.2.4 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 systems across distances.They do not protect the data from logical or human errors.

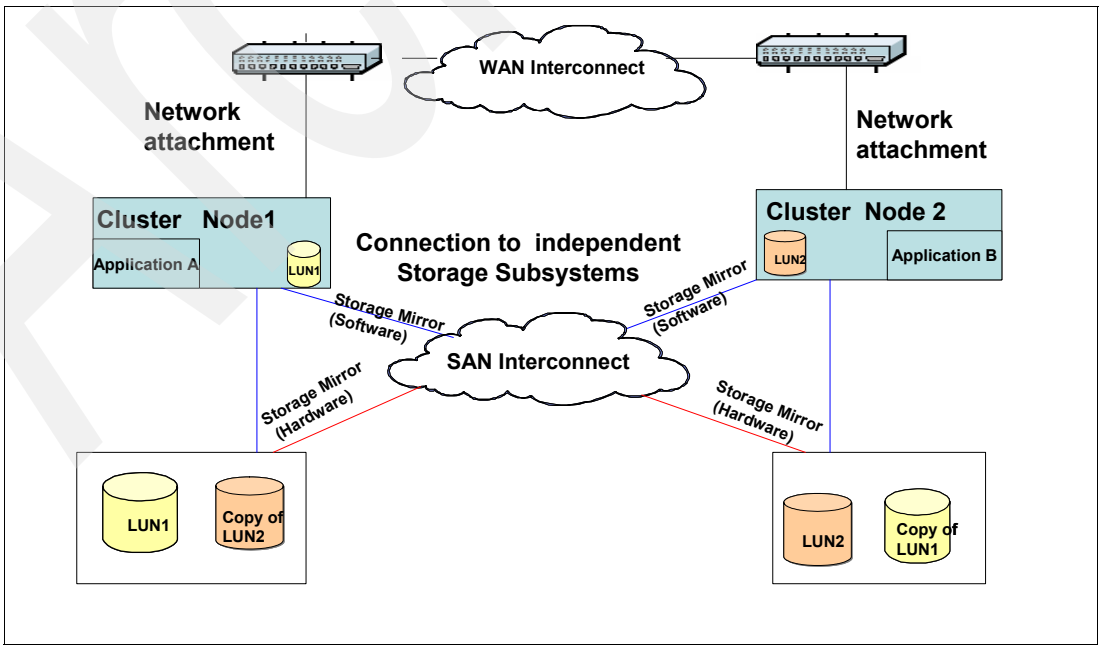

Figure 11-6 Geographically dispersed High Availability cluster

As indicated in Figure 11-6, there are two general possibilities within geographically dispersed clusters for synchronization and data mirroring of the storage systems between the sites:

- Host based mirroring, such as with AIX Logical Volume Manager
- Storage based mirroring, such as
  - Metro Mirror for DS6000, DS8000, ESS, SVC
  - Enhanced Remote Mirroring (ERM) with DS4000
  - SyncMirror® for N series

Either method has its advantages and disadvantages.

Examples of geographically dispersed High Availability clusters include GDPS/PPRC and Copy Services for System i. You can find more information about these technologies in the Continuous Availability chapter of *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548.

11.2.5 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 an 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 in combination with the appropriate technologies are a key component within the disaster protection and data recovery methodology. With databases two general methods of backup technologies are to be considered:

- *Offline backup* means to shutdown 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 the online backup methods, we take a look at the database level and how the database components interact with each other. To understand how an online backup works, we look closer at how databases work. A transaction-oriented online database of any kind consists of three major disk storage components:

- **Data:** The data is organized in flat files, tables, or tablespaces depending on the database.
- **Log files**: The log files are the files which keep the records of all the transactions which have been written to the database. As log files fill up, they switch the processed transactions to archive logs as shown in Figure 11-8 on page 314.
- **Archived log files:** The log files which contain already-processed transactions are copied to the archive log before the logs are overwritten with new transactions.

Figure 11-7 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 is written consistently in the event of a server crash.

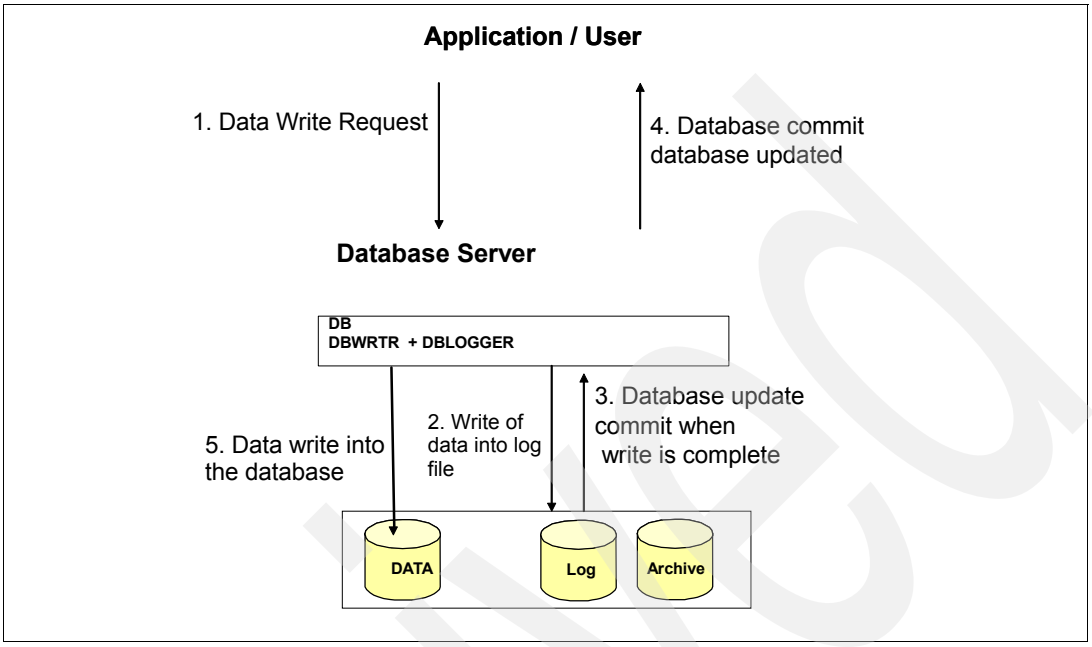

Figure 11-7 Simplified transaction based database write

After a log file has filled up with transactions, a switch to a new log file occurs. In the mean time the transactions from the old log file are copied to the archive logs, as shown in Figure 11-8.

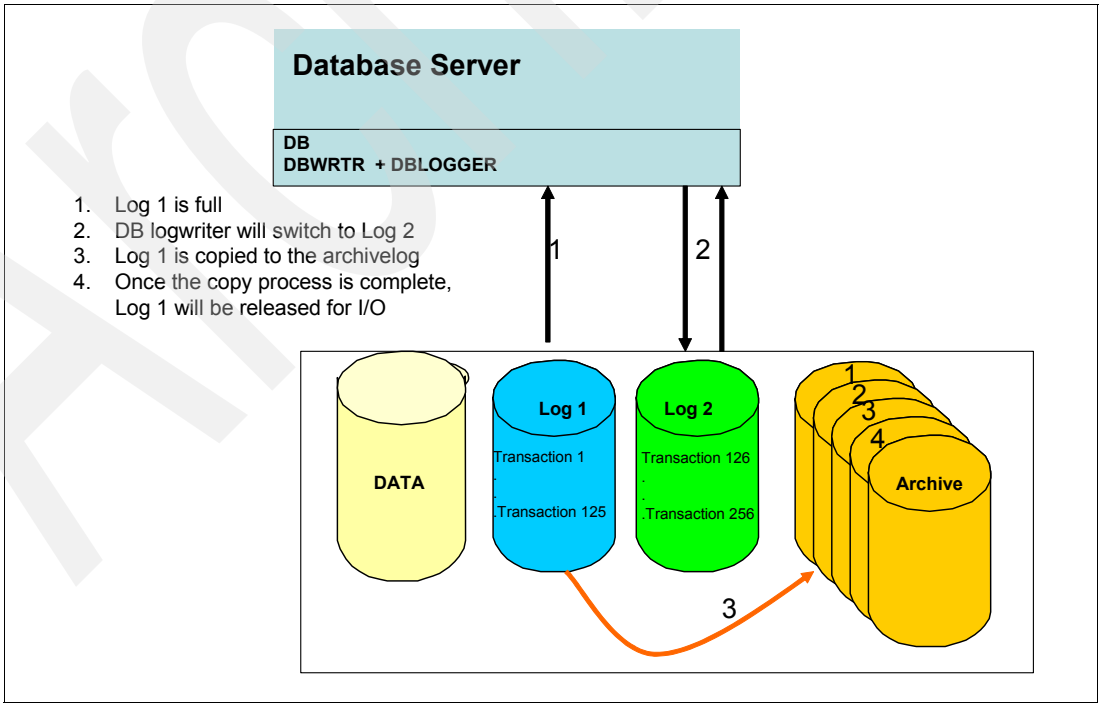

Figure 11-8 Simplistic database log file switch

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

In addition to the physical tape speed, 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. With database recovery, you need to consider the following time windows for *recovery time*:

- **Recovery Time (R*t*)**: Time window from physical restore to data availability
- **SetupTime (S*t*)**: Administrative time to prepare the restore
- **Physical Recovery Time (PR*t*)**: 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 (R*t*) is:

```
Rt = St + PRt + nT / nTs
```

Important: To recover a database to a certain point-in-time, you must have the logfiles + the database + the archive files in sync! A database recovery always depends on the backup status of the data as well as the sequential availability of the archive and logfiles. Without the archive and logfiles a recovery might be impossible.

Example

Figure 11-9 shows how a simplified point-in-time database recovery works.

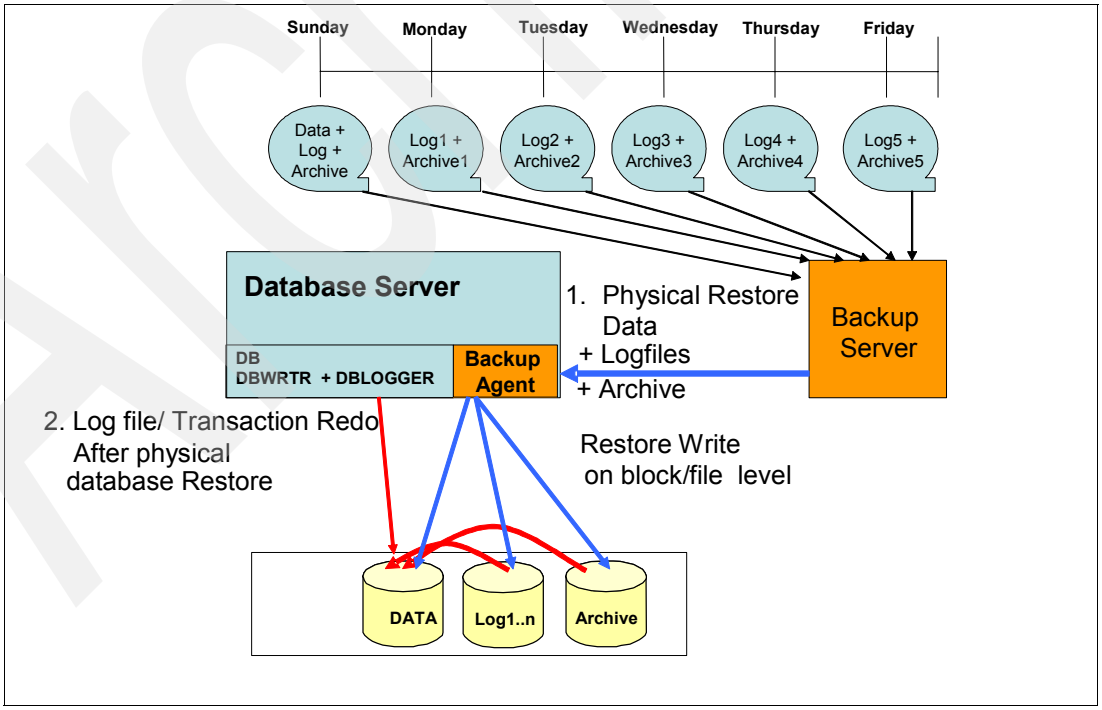

Figure 11-9 Simplistic database restore - point in time

In this example, we consider a disaster on Friday, where we have to recover from tape. If the backup policy in this scenario was to do a full database backup consisting of the database, the log files and the archives every Sunday, then we need to keep backups of the log files and archives over 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 which 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.

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 (LOB) or physical online backup.

In *LOB*, shown in Figure 11-10, the database and the backup software communicate through backup agents and APIs. This type of backup implies a close integration of the backup software to the database. Therefore all components, the database version, the backup software, the online agents, and so on, must be compatible with each other.

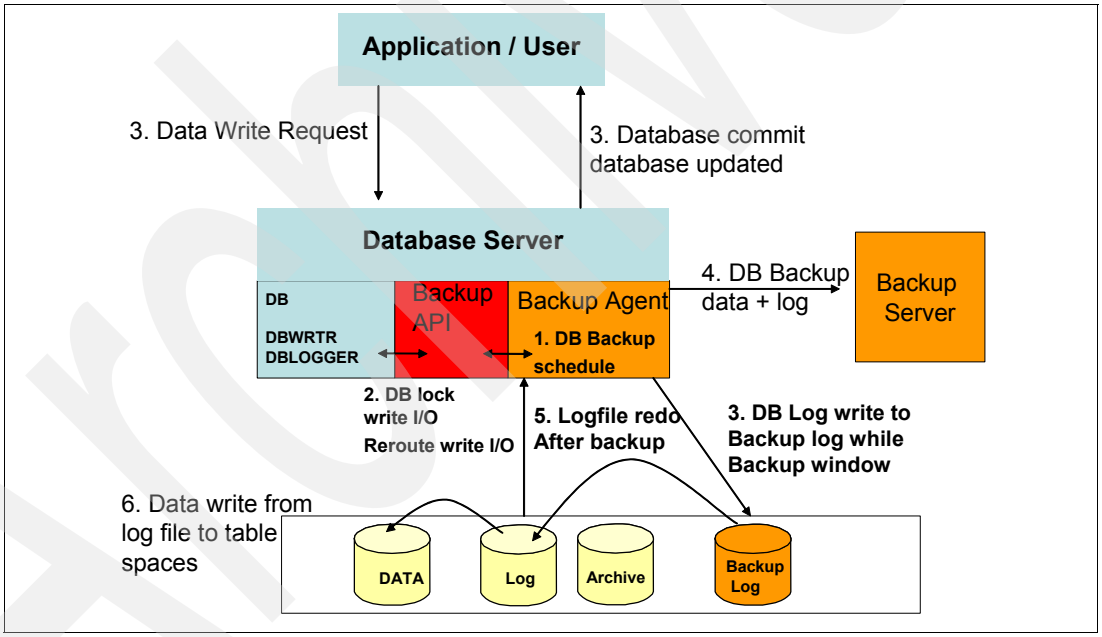

Figure 11-10 LOB

The backup agent triggers the backup. While the online backup is running, all write activity against the database is rerouted to the dedicated backup log. This guarantees 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.

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 is dependent on the size of the database and the backup technology in place. For information

about appropriate tape technologies, see *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548.

Physical online backups are considered as a physical image copy of the database and its related log files but executed on the storage (*the physical*) level. This type of backup 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 buffer pools 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. In 11.2.6, “Remote storage mirroring” on page 318 we discuss the details of storage based remote copy.

As shown in Figure 11-11, the copy process itself is executed on the storage layer versus the logical online backup where the backup is executed on the *server layer.*

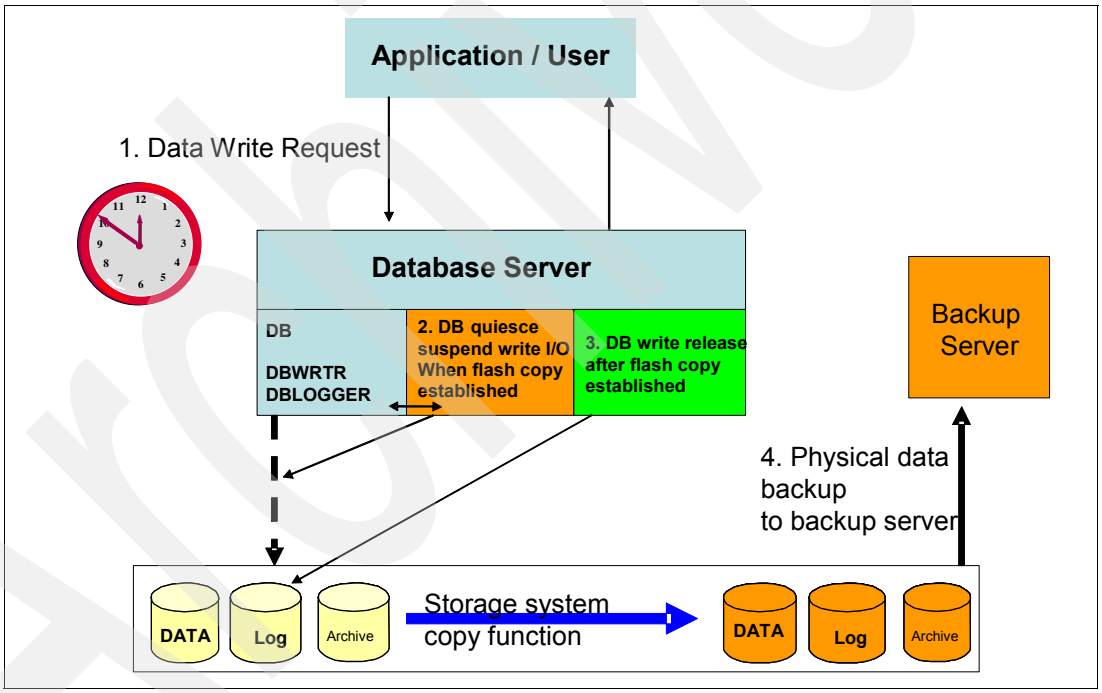

Figure 11-11 Physical database copy using storage copy function

After the FlashCopy process has been initiated, which takes just a couple of seconds, the actual copy process of the data to the FlashCopy volume is executed in the background. This background operation allows the database to resume write access. Figure 11-12 illustrates the copy on write mechanism of FlashCopy.

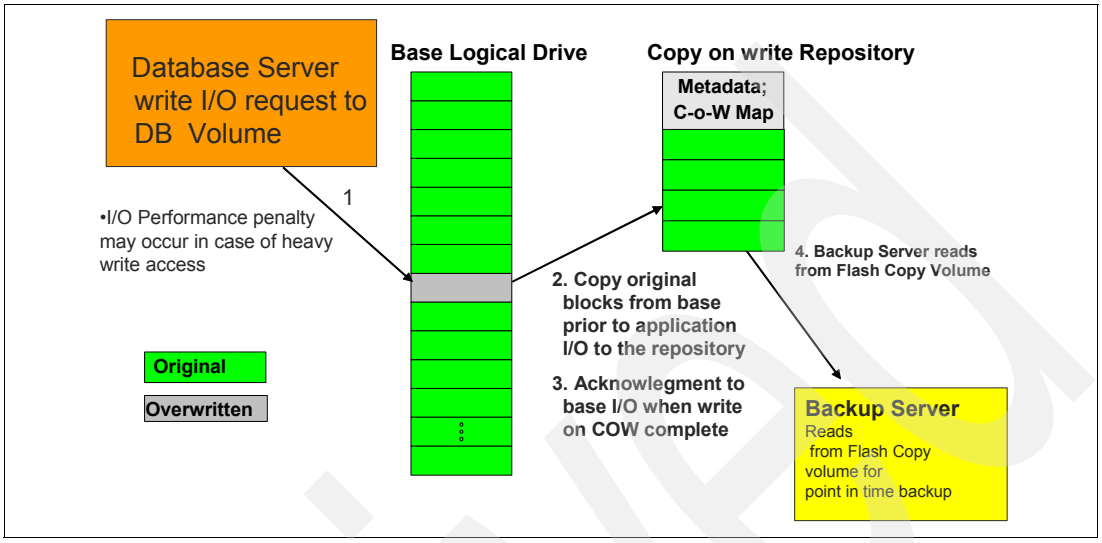

Figure 11-12 Simplistic FlashCopy operation

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, you can mount them and dump them. For more information, see *IBM System Storage DS8000 Series: Copy Services in Open Environments*, SG24-6788.

11.2.6 Remote storage mirroring

Remote storage mirroring is a *storage* based method for data mirroring between two independent storage systems. It offers basic functionality for real dispersed clusters and remote online data protection on storage systems. However, as with the cluster technologies, it does not provide protection for any kind for logical or human errors. Data mirroring between two sites can be achieved on two levels:

- Logical storage mirroring
- Physical storage mirroring

Logical storage mirror (logical volume mirror)

By using logical mirroring (or *host based mirroring* as it often called) the data mirror between storage systems is achieved by additional server based software or with built in capabilities of the database itself which synchronizes the writes between two storage systems.

Assume we have two storage systems, each with one LUN which is protected on the storage system itself by RAID and other HA components. By using the logical volume mirror both LUNs in the two storage systems will be reflected as one disk to the server operating system. Thus the mirror between the systems is synchronous. In case one storage system fails, the remaining storage system will continue without any interruption to the application (the database). By using LVM technology the database can be *physically* protected with more then one storage system and also N+1 servers.

The advantage of using LVM is the fact that the operating system or clustering software is constantly aware of both disks. Should one volume become unavailable for any reason, there

is no failover period. Instead the Operating system or clustering software will continue functioning with one disk instead of two, as shown in Figure 11-13.

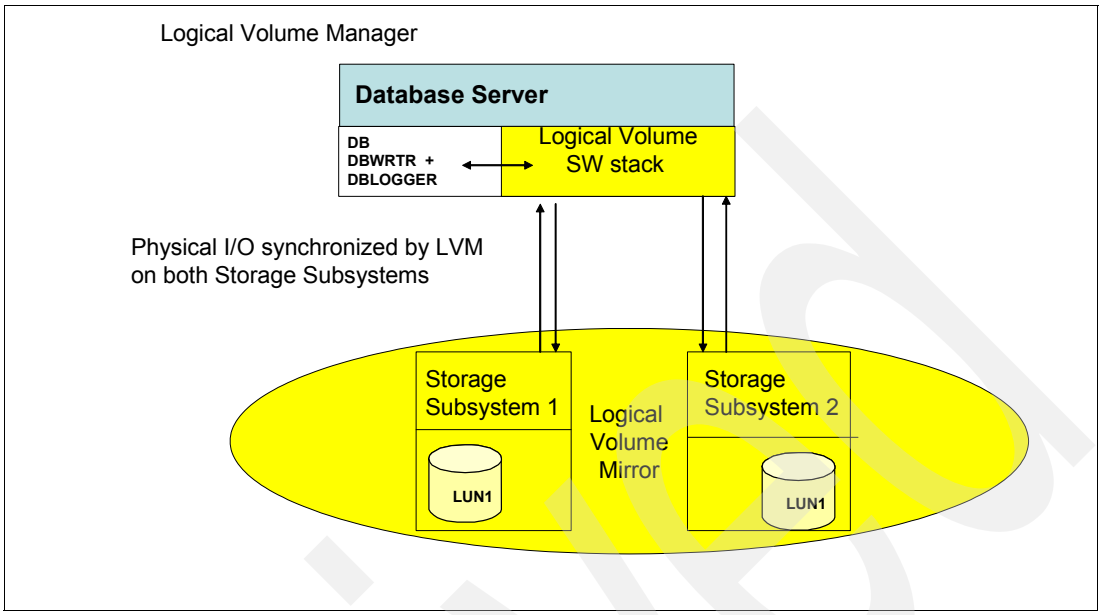

Figure 11-13 Logical volume mirroring

Block level hardware mirroring (controller based mirroring)

Unlike logical volume mirroring discussed previously, the controller based mirror is executed and maintained exclusively at the storage system level. Implementations of the physical storage mirror functionality are Metro Mirror and Global Mirror, including Enhanced Remote Mirroring (ERM) for the DS4000. The advantage of this technology is derived from the fact that the mirroring is done completely at the storage controller level, without regard to the application or server environment. This enables us to use block level replication as a single point of control for all mirroring in the environment and can represent a major simplification to the infrastructure of the Business Continuity solution.

The limitations on block level mirroring will depend on the type of disk system and specific form of mirroring in use. In the disk-related chapters in *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548, we discuss the individual implementations of the storage mirroring functionality on those platforms. Additionally, depending on the operating systems in the environment, Control Software might be available to enhance the availability of the data and control the recovery of the server environment. Those technologies are described in the Continuous Availability chapter in *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548.

Databases on their own do not actively take advantage of a physical mirror. In the event of a storage system failure on the production site, the surviving storage system must be made available to the (cold standby) database server. In case of data corruption of the database, or any kind of logical error also on the second storage server, it would be necessary to recover from tape or a previous point-in-time copy. Figure 11-14 shows the operation of Metro Mirror.

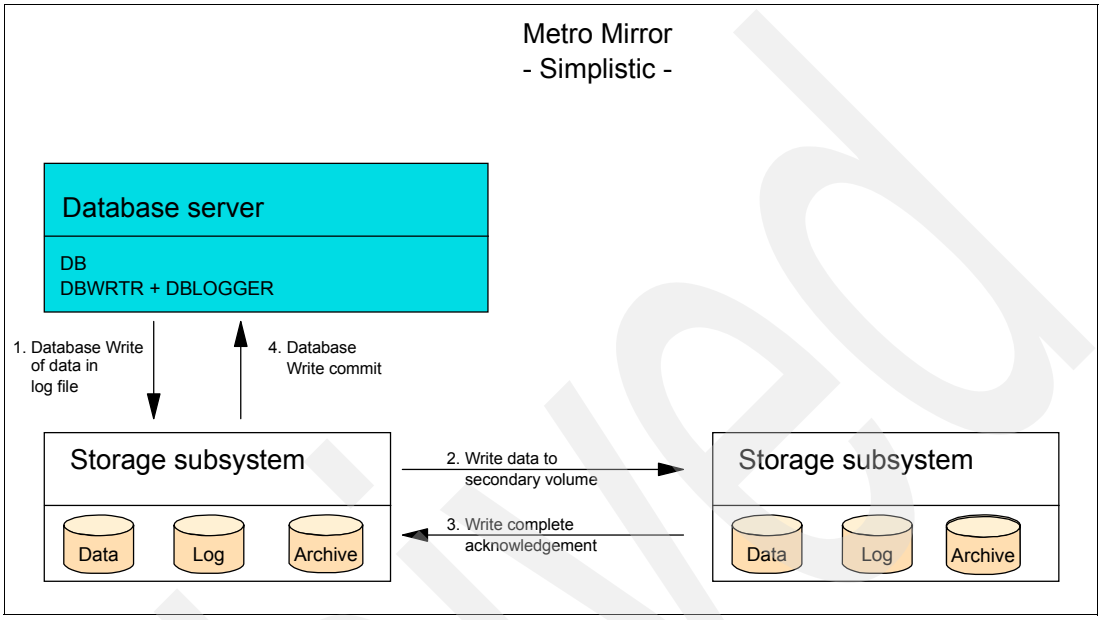

Figure 11-14 Metro Mirror replication

11.2.7 General Parallel File System

The General Parallel File System™ (GPFS™), which can be striped over multiple servers and thousands of disks, is designed for scalability and High Availability. The general design of GPFS is shown in Figure 11-15 on page 320, which represents an example with Linux. Other operating systems, such as AIX, also provide GPFS capabilities.

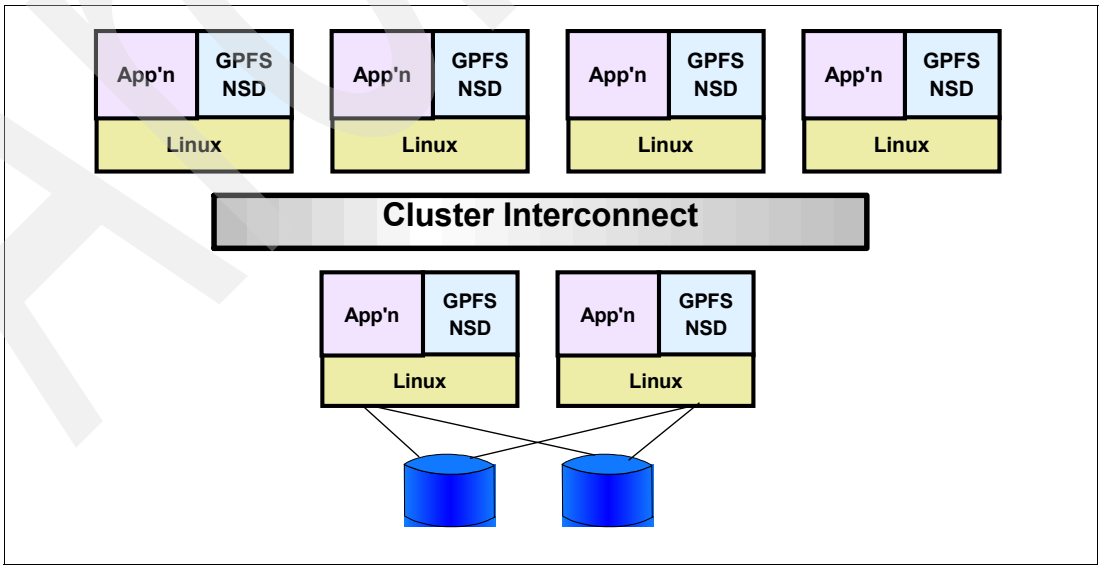

Figure 11-15 GPFS file system

As you can see, all applications using GPFS have access to the same data. However, while GPFS is able to make data highly available within a local environment, it is not suited to metropolitan or global mirroring of data.

11.2.8 Shadow databases

Shadow databases are physical and logical copies or replicas of a database dispersed over two or more sites. The management of the updates to multiple copies of the database is done by the database management system itself not the server operating system or the storage system. This provides a greater protection from physical or logical errors than can be provided by operating-system-based or storage-system-based solutions, but shadow databases do not protect from human errors or erroneous transactions. On their own, shadow databases are just another building block for the IBM Business Continuity Solution Selection methodology.

You can set up shadow databases in two ways:

- Synchronous (no data loss) database shadow (also called distributed relational database)
- Asynchronous (data loss) database shadow (also called warm standby or log-shipping approaches)

In this section we discuss both methods in detail.

Synchronous database shadow (no data loss)

With synchronous database shadow technology the control of data integrity, transaction logic and I/O behavior is managed by the *database management software*, itself. This is very different from the previous methods we discussed along with the *logical* or *physical* storage mirroring possibilities.

Synchronous database shadows provide online availability at the database level. Instead of handling database updates at the *storage* layer using SAN connectivity, the mirror of the shadow database is updated by sending the transaction from the production database server to the shadow database server through the LAN or WAN. When using this technology, the servers only need connectivity to their respective local storage.

Shadow databases can be implemented on independent hardware and are only logically (on the database level) linked together as displayed in Figure 11-16.

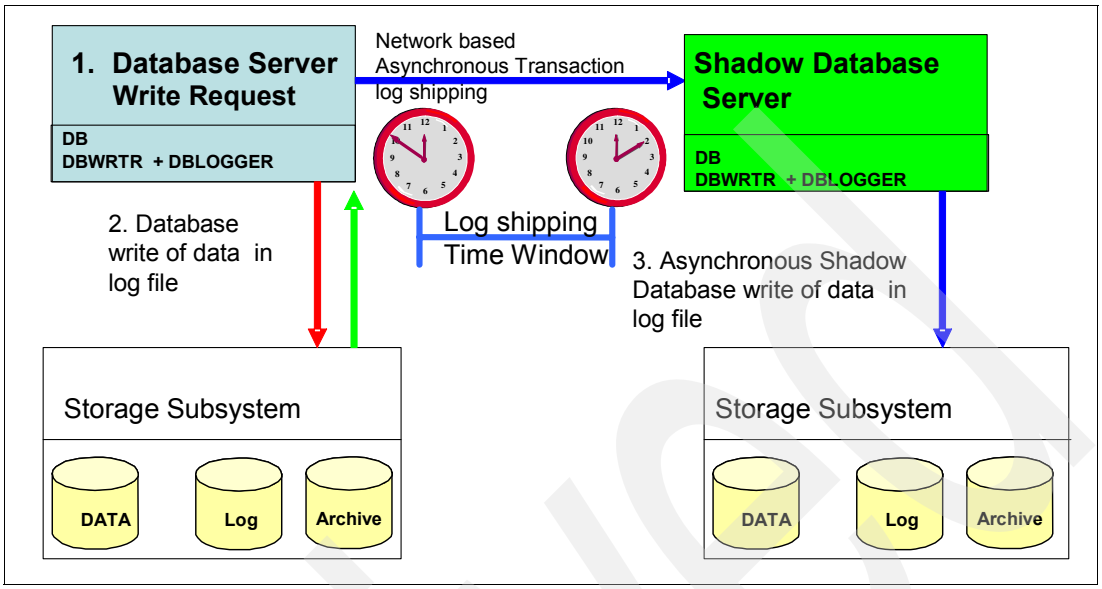

Figure 11-16 Synchronous shadow database

Asynchronous shadow database (warm-standby database)

Many relational database offerings provide a capability for creating a duplicate standby database at another location. In the event of a disaster that destroys access to the primary database, application processing can be switched to an alternate site, accessing the standby or secondary database. The standby database can be updated frequently, to keep it close in currency to the primary database. The goal is to switch over to the standby database and restore the production environment in the least amount of time with as little effort possible.

To establish this environment, you first backup the entire primary database. Then you restore this image on a different, but identical, server. You also might need to make some configuration and network definition changes, if the primary and standby servers will be connected to your network at the same time.

As transactions are processed at the primary database site, these transactions are captured in the database transaction logs, also called *re-do* logs. These logs are shipped to the standby site using TCP/IP.

When these transaction logs are applied to the standby database instance, it brings the standby database up to currency as of the last transaction captured in each log. Of course, transactions continue to be processed and logged at the primary site, so the standby site is always slightly behind the primary site in transaction currency because transactions are only applied at the standby site when the primary transaction log is closed and shipped to the standby site.

Some shadow database implementations offer a more granular approach, they do not wait for a log to be closed but intercept log write operations and send the information immediately to the remote site. This allows you to have more current data at the secondary site.

You can take periodic full backups of the database from the standby site copy. The advantage of this is that the primary database server instance does not have to be shut down or quiesced to do this.

In Figure 11-17 we consider a time window of 20 minutes. This will provide the ability to recover the database with a specified Recovery Point Objective of, for example, 20 minutes. In our previous synchronous shadow database example all transactions within the 20 minute time window, which are not synchronized, would have been lost in case of a disaster on the production site. Therefore the desired synchronization window depends, as already mentioned, on the business requirements and the physical network speed.

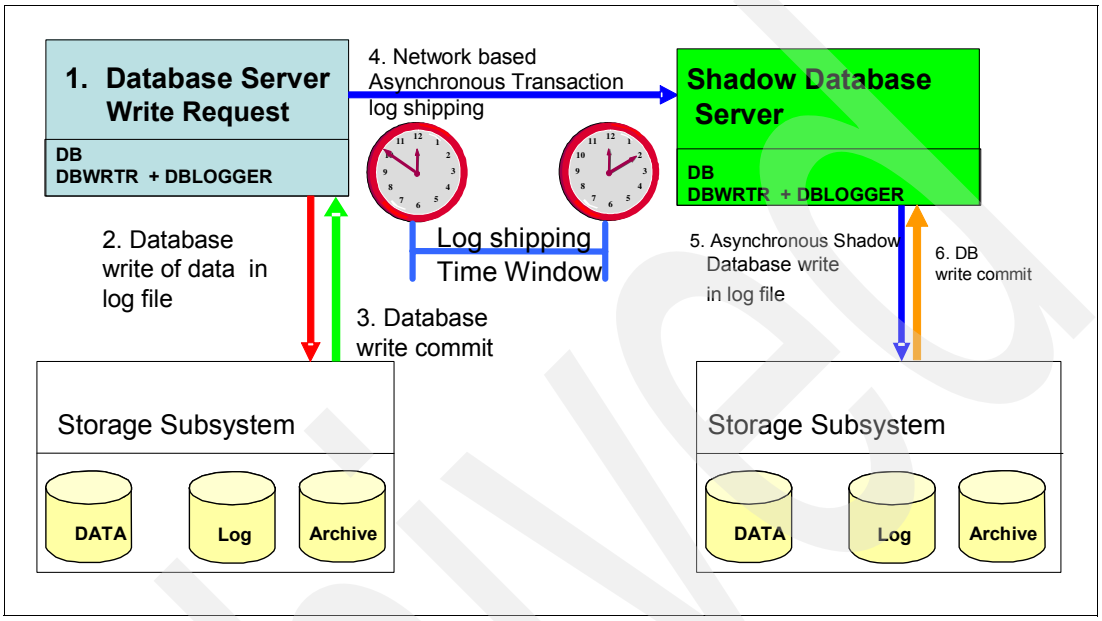

Figure 11-17 Asynchronous shadow database

This approach can also help protect from the consequences of corruption of the primary database. If the primary database is corrupted, or if an important table space is dropped (deleted) in error, you can protect yourself from propagating this error to the standby site merely by not applying that log to the standby. The database administrators can restore the table space using the database manager's facilities to move the uncorrupted table space from the standby image back onto the primary database.

There are some types of changes that can be applied to the primary database that will not be propagated to the standby database by merely applying the transaction logs. An example of this is if the database administrator increases the storage capacity of the primary database. There are other processes that can be done on the primary database that will cause the standby database to become out of date, and the standby will have to be recreated from a copy of the primary. An example of this is if the logs at the primary are reset. Overall, the log-shipping, or stand-by approach can provide a disaster recovery copy of the database, but there is some administrator attention that is needed to ensure that the stand-by copy is valid.

Conclusion

Any kind of online database backup avoids interrupting production without shutting the database down. It also avoids 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 remaining sections in this chapter describe examples of how you can achieve database security using certain technologies and methods. They also describe the advantages and disadvantages of each solution.

However, these scenarios are just examples. It is not the intention and purpose of this book to cover all possible combinations of databases, operating systems and hardware platforms. Therefore we always recommend to take into consideration 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.

11.3 Databases

There are a number of factors to consider when planning database storage requirements. We can only present some generalizations here. Your database consultant can provide more detailed recommendations.

11.3.1 Storage planning and database preparation

One important step in the planning process for a database installation is the definition of the database physical layout on the storage. The layout is very much depends on the nature of the database and the performance requirements.

We differentiate between the following database workload types:

- **Business intelligence and data warehouse typical characteristics**
  - Decision support
  - Huge amounts of data, long and complex SQL queries
  - Bandwidth intensive (MBps)
  - Read intensive
  - Small block (4 KB to 8 KB) sequential reads
- **Enterprise Resource planning typical characteristics**
  - Small block (4 K to 8 K) random reads
  - Online Transaction Processing
  - Large number of users and transactions, short response times
  - Mix of reads and writes (usually about 70% reads)
  - Mainly random access
- **Multimedia systems planning typical characteristics**
  - Bandwidth intensive (MBps)
  - Read intensive
  - Large block (64 KB to 128 KB) sequential reads

11.3.2 General recommendations for database storage layout

Traditionally, database vendors have given recommendations on the placement of database data files on the underlying storage devices. Some of the recommendations for databases and their layout in storage include:

- Plan the number of storage systems based on total expected throughput requirements.
  - Transactions/hour.
  - Type of database workload (OLAP, OLTP, BI).
- Typical database 70% reads 30% writes for most OLTP systems.
  - Use RAID-1 (or RAID-10) for log files and archives while the database files are installed on RAID-5.

- High performance requirements for database I/O:
  - Use RAID-10 for log files and data.
  - Use more smaller disk drives.
  - Use high performance disks such as the 15k drives.
- Data placement / data layout considerations to maximize performance and avoid bottlenecks. For example:
  - Spread active data over as many storage components as possible
  - Separate database files and redo logs so they are placed on different storage components.

While many of the recommendations remain valid, the applicability of the RAID recommendations is questionable when we use enterprise class storage systems. These systems often offer advanced functionality that masks the performance penalties of specific RAID implementations. For example RAID5 is more than adequate for most kinds of database data while RAID10 performs particularly well in presence of high rates of random write I/Os.

Database relevant RAID basics

The number of read and write operations needed for RAID-1 and RAID-5 are:

- RAID-1 write = 2 * Write
- RAID-5 write = 2 * Read + 2 * Write = RAID-5 write penalty

11.3.3 Database tuning considerations

In many situations where a database system does not perform as required, it might seem easier to start the tuning at the storage level. However, you need to consider the following steps before tuning the storage:

1. Analyze and tune the global business rules.
2. Analyze and tune the data design.
3. Analyze and tune the application design or recompile the application with server platform optimized functions (such as, multi threaded versus single threaded).
4. Analyze and tune the logical database layout (tables, indexes, schemas).
5. Analyze and tune the database operations.
6. Analyze and tune or redistribute the data access paths.
7. Analyze and tune DB memory allocation.
8. Analyze and tune DB server infrastructure (avoid bottlenecks).
9. Analyze and tune I/O and physical structure (this is where Storage is considered!).
10. Analyze and tune Resource Contention.
11. Analyze and tune the underlying platform(s).

11.4 Benefits of database, storage, and logical mirror functions

Hardware and software based mirroring both have their own separate strengths. We discussed some criteria in “Data replication” on page 101. However, this section also includes a summary list of the benefits.

Database application mirroring

Database application mirroring has the following features:

- **Application aware:** In database mirroring, such as HADR with DB2, the application is aware of what data has been written to the recovery location. This can help avoid rolling disasters and maintain *transactional integrity* in one step.
- **Low-bandwidth:** Application level mirroring typically minimizes bandwidth utilization by sending smaller chunks (versus block level) of data over IP.
- **Warm standby:** Database applications can have their secondary prepared to take over. Should an outage require recovery in the secondary site, it is more capable of quickly restoring operations at the transaction integrity level.

Server driven logical volume mirroring

Logical volume (operating system) mirroring has the following features:

- **No failover:** Logical volumes do not failover in the typical sense but, rather, the server or clustering application is constantly aware of all disk that it is attached to. Should one half of the volume group fail, it will simply continue accessing the other half
- **Application independent:** Logical volume mirroring can be done with any data associated with the platform that it is being managed through, rather than data that is associated with a specific application.

Storage Block Level Mirroring

Storage level mirroring has the following features:

- **Application independent:** Hardware mirroring can be done across all data and applications, regardless of the type of server or application.
- **Centralized management:** Management is simplified because there is only one level to deal with. Rather than handling each server or application as a separate entity to be managed, the storage controller can be the control point for all data being mirrored.

11.5 Summary

Remember, protecting the database and the data is only one link in the availability chain of the business process, as we have already stressed, and does not provide business continuance on its own. Certain backup and recovery methods for which you also have to plan in conjunction with the data protection of the databases itself are discussed in the Storage Management software chapter in *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548.

In this chapter, we discussed the advantages to each of the types of mirroring available and that fact that any or more than one could be applicable to a given customer environment.

Throughout the database planning process IBM consulting teams can help and support you, to build and propose the best solution as a whole, exactly matching your needs based on your business requirements.

You can find a wide selection of documents on High Availability solutions for applications and databases connected to IBM system storage on the IBM Web site. You can adapt these solutions easily to the DS4000, DS6000, and DS8000 series. The following link points to IBM System Storage white papers:

http://www.ibm.com/servers/storage/disk/ess/whitepapers.html

A

Business Continuity Solution Selection Methodology matrixes

This appendix provides the following tools for use with the Business Continuity Solution Selection Methodology that we describe in Chapter 5, “Business Continuity Solution Selection Methodology” on page 151.

- “Starter set of business requirement questions” on page 330 is the starter set of Business Requirements questions for the Methodology.
- “Business Continuity Solution Matrix” on page 331 is the Solution Matrix that organizes the collection of System Storage solutions in this book into tiers.
- “Eliminate non-solutions matrixes” on page 333 are the Eliminate Non-Solutions tables, one for each cell in the Solution Matrix.

We include additional questions that are useful for building business justification for the Business Continuity solution, as well as further information needed by the detailed evaluation team, in “Additional business requirements questions” on page 349.

Starter set of business requirement questions

This section includes a suggested starter set of Business Continuity requirements questions and answers that you need to obtain prior to entering into the methodology. These questions are designed to elicit enough basic information to start the process. (You can find a detailed list of additional questions in “Additional business requirements questions” on page 349.)

Some of these questions require the business line to answer through a Risk and Business Impact Analysis. Other questions are for the operations staff to answer from their knowledge of the IT infrastructure.

The terms used in the questions are defined in Appendix B, “Terms and definitions” on page 357.

The starter set of business requirements questions are:

1. What is the application (or applications) that needs to be recovered?
2. On what platform or platforms does it run?
3. What is the desired Recovery Time Objective?
4. What is the distance between the recovery sites (if there is one)?
5. What is the form of connectivity or infrastructure transport that will be used to transport the data to the recovery site? How much bandwidth is that?
6. What are the specific hardware and software configurations that need to be recovered?
7. What is the Recovery Point Objective?
8. What is the amount of data that needs to be recovered?
9. What is the desired level of recovery? (Planned, Unplanned, Transaction Integrity)
10. Who will design the solution?
11. Who will implement the solution?

Business Continuity Solution Matrix

Table A-1 shows the full Solution Matrix for use with the System Storage solutions in this book.

Table A-1 Business Continuity Solution Matrix

	Tier 7	Tier 6	Tier 5	Tier 4, 3	Tier 2, 1
RTO ===>	**RTO: Generally near continuous to 2 hours**	**RTO: Generally 1 to 6 hours**	**RTO: Generally 4 to 8 hours**	**General RTO: Tier 4: 6-12 hrs Tier 3: 12-24 hrs**	**RTO: Generally > 24 hours**
Description:	**Highly automated integrated hardware and software failover**	**Storage and Server mirroring**	**Software, application, and database transaction integrity**	**Hotsite, disk PiT copy, database journaling and forwarding, comprehensive backup s/w, fast tape, electronic vaulting**	**Backup software, physical transport of tape**
Level of Recovery: **Planned Outages, data migrations "Byte Movers"**	►GDPS/PPRC ►GDPS/PPRC with HyperSwap ►GDPS/XRC ►AIX HACMP/XD with Metro Mirror	► Metro Mirror ► Global Copy ► z/OS Global Mirror ► FlashCopy ► PCR	► Software ► Application ► Database level facilities	► FlashCopy ► Global Copy- Tivoli Storage Manager - DRM- Tape	► Tivoli Storage Manager ► Tape
Level of Recovery: **Unplanned Outage, adds data integrity to "Byte Movers"**	► GDPS/PPRC ► GDPS/PPRC with HyperSwap ► GDPS/XRC ► AIX HACMP/XD with Metro Mirror	► z/OS Global Mirror ► GDPS HyperSwap Manager ► TPCR ► PPRC Migration Manager ► AIX LVM ► AIX HACMP/XD	► Software ► Application ► Database level facilities	**Tier 4**: ► PtP VTS ► TS7700 Grid ► FlashCopy ► FlashCopy Manager ► Global Copy ► N series Snap Mirror **Tier 3:** ► FlashCopy ► Tivoli Storage Manager ► Tape	► Tivoli Storage Manager ► Tape

	Tier 7	Tier 6	Tier 5	Tier 4, 3	Tier 2, 1
Level of Recovery: **Transaction Integrity, adds transaction integrity hardware "Unplanned Outage" data integrity**	Database level recovery on top of any Tier 7 Unplanned Outage recovery. *Examples*: DB2 with GDPS or AIX HACMP/XD, SAP and DB2 remote replication, shadow database with forward recovery, split mirror, and so forth.	Database level recovery on top of any Tier 6 Unplanned Outage recovery. *Examples:* DB2 UDB with AIX HACMP/XD,SAP and DB2 remote replication, shadow database with forward recovery, split mirror, and so forth.	► Remote Replication with DB2, Oracle, SQL Server, and so forth. ► Software, application, and database level facilities. *Examples:* shadow database with forward recovery, Split mirror, and so forth. See Chapter 11, "High Availability clusters and database applications" on page 305.	► Tier 4: Database-level journal file forwarding and remote application ► Tier 3: Database-level recovery with physical tape transport	Database-level recovery with physical tape transport
Tolerance to Outage	Low tolerance to outage	Low tolerance to outage	Low tolerance to outage	Somewhat tolerant to outage	Very tolerant to outage

While PtP VTS is still a valid name for existing equipment, TS7700 Grid is the new name for the IBM TS7700. Both these terms address a similar function, on two different technologies

Notes on the Solution Matrix cells

As a general comment, recall that the Business Continuity Solution Selection Methodology is not intended to be a perfect decision tree, and the boundaries and contents of the cells are of necessity giving general guideline suggestions rather than attempting to be all inclusive.

The methodology allows room for product and Business Continuity experts to add their expertise to the evaluation process after an initial preliminary set of candidate solutions is identified.

The intent of the methodology is to provide a framework for efficiently organizing multiple Business Continuity technologies, and more quickly identifying the proper possible solutions for any given customer set of requirements.

Tiers 7, 6, and 5 Transaction Integrity

The solutions for transaction integrity are specific to the database and application software being used. Because of this, the list of possible solutions is very broad, and it is not feasible to be all-inclusive. You should involve a software specialist skilled in the application and database set that you are using for detailed evaluation of transaction integrity recovery specific to your database and application. We do show common examples of solutions in these cells in the matrix.

Chapter 11, “High Availability clusters and database applications” on page 305 gives an extensive discussion on specific solutions and guidance relating to database and application level transaction integrity recovery.

Eliminate non-solutions matrixes

You use the tables in this section to eliminate non-solutions. There is one matrix set for each cell in the solution matrix. Note that for some solutions, greater distances can be supported by special request. See your IBM representative for details. However, for synchronous replication in particular, the specific workload and response time requirements can reduce the distance that can be spanned.

Tier 7 Planned Outage

Table A-2 is the matrix for Tier 7 and Planned Outage.

Table A-2 Tier 7 and Planned Outage matrix

Solution	GDPS/PPRC	GDPS/PPRC with HyperSwap	HACMP/XD
Platform	System z System z + Distributed	System z System z + Distributed	AIX
Distance	<100km	<100km	<100km
Connectivity	FICON Fibre Channel	FICON Fibre Channel	Fibre Channel TCP/IP
Supported disk storage (primary site)	PPRC- compliant disk systems	PPRC HyperSwap capable disk systems	DS8000 DS6000 DS4000 SVC ESS
Supported disk storage vendor (secondary site)	PPRC-compliant disk systems	PPRC HyperSwap capable disk systems	DS8000 DS6000 DS4000 SVC ESS N series
Recovery Point Objective	near zero	near zero	near zero
Amount of Data	any	any	any
Other notes			

Tier 7 Unplanned Outage

Table A-3 is the matrix for Tier 7 and Unplanned Outage.

Table A-3 Tier 7 and Unplanned Outage matrix

Solution	GDPS/PPRC	GDPS/PPRC with HyperSwap	HACMP/XD
Platform	System z System z + Distributed	System z System z + Distributed	AIX
Distance	<100km	<100km	<100km
Connectivity	FICON Fibre Channel	FICON Fibre Channel	Fibre Channel TCP/IP
Supported disk storage (primary site)	PPRC-compliant disk systems	PPRC HyperSwap compliant disk systems	DS8000 DS6000 DS4000 SVC ESS
Supported disk storage (secondary site)	PPRC-compliant disk systems	PPRC HyperSwap compliant disk systems	DS8000 DS6000 DS4000 SVC ESS N series
Recovery Point Objective	near zero	near zero	near zero
Amount of Data	any	any	any
Other notes	Delivered as IBM Global Services Offering	Delivered as IBM Global Services Offering	

Tier 7 Transaction Integrity

This matrix is for Tier 7 and Transaction Integrity. For readability, the multiple columns for this matrix are separated into multiple parts, Parts 1 and 2, in Table A-4 and Table A-5.

The solutions for transaction integrity are specific to the database and application software being used. Because of this, the list of possible solutions is very broad, and it is not feasible to be all-inclusive. You should involve a software specialist skilled in the application and database set that you are using for detailed evaluation of transaction integrity recovery specific to your database and application. We do show common examples of solutions in these cells in the matrix.

Chapter 11, “High Availability clusters and database applications” on page 305 includes an extensive discussion on specific solutions and guidance relating to database and application level transaction integrity recovery.

Table A-4 Tier 7 and Transaction Integrity, Part 1

Solution	Database transaction recovery layered on GDPS/PPRC	Database transaction recovery layered on GDPS/PPRC with HyperSwap	Database transaction recovery layered on AIX HACMP/XD with Metro Mirror
Platform	System z System z + Distributed	System z	System p
Distance	<100km	<100km	<100 km
Connectivity	FICON Fibre Channel	FICON, Fibre Channel	Fibre Channel TCP/IP
Supported disk storage (primary site)	PPRC-compliant disk systems	PPRC HyperSwap compliant disk systems	DS8000 DS6000 DS4000 ESS SAN Volume Controller
Supported disk storage (secondary site)	PPRC-compliant disk systems	PPRC HyperSwap compliant disk systems	DS8000 DS6000 DS4000 SAN Volume Controller ESS
Recovery Point Objective	near zero	near zero	near zero
Amount of Data	any	any	any
Other notes	Delivered as IBM Global Services Offering	Delivered as IBM Global Services Offering	

Table A-5 Tier 7 and Transaction Integrity, Part 2

Solution	Shadow database with forward recovery	Split Mirror Database with Metro Mirror
Platform	any	any
Distance	any	<300km
Connectivity	any	FICON Fibre Channel
Supported disk storage (primary site)	any	PPRC-compliant disk systems
Supported disk storage (secondary site)	any	PPRC-compliant disk systems
Recovery Point Objective	depending on the log shipping mechanism, loss of only few transactions possible	near zero
Amount of Data	any	any
Other notes		

Tier 6 Planned Outage

For readability, the multiple columns for this matrix are separated into multiple parts, Parts 1 and 2, Table A-6 and Table A-7.

Table A-6 Tier 6 Planned Outage, Part 1

Solution	Metro Mirror	Global Mirror	z/OS Global Mirror (XRC)	MetroMirror Global Mirror
Platform	any	any	System z	Distributed
Distance	<300 km for DS8000 DS6000 ESS <100km for other disk systems	any	any	<100Km DS8000 DS6000 DS4000 SVC N series ESS
Connectivity	FICON Fibre Channel	FICON Fibre Channel	FICON	Fibre Channel
Supported disk storage (primary site)	Metro Mirror-capable disk systems	DS8000 DS6000 ESS	z/OS Global Mirror capable disk systems	Metro Mirror-capable disk systems
Supported disk storage (secondary site)	Metro Mirror-capable disk systems	DS8000 DS6000 ESS	any z/OS supported disk systems	Metro Mirror-capable disk systems
Recovery Point Objective	near zero	few seconds to few minutes	few seconds to few minutes	near zero to a few minutes
Amount of Data	any	any	any	any
Other notes				

Table A-7 Tier 6 Planned Outage, Part 2

Solution	PtP VTS Synchronous TS7700 Grid	PtP VTS Asynch. TS7700 Grid
Platform	System z	System z
Distance	<250km	any
Connectivity	FICON Fibre Channel ESCON	FICON Fibre Channel ESCON
Supported disk storage (primary site)	any System z supported disk systems	any System z supported disk systems
Supported disk storage (secondary site)	any System z supported disk systems	any System z supported disk systems
Recovery Point Objective	near zero	few seconds to few minutes, minutes to hours, (defined by user policy)
Amount of Data	any	any
Other notes		

Tier 6 Unplanned Outage

For readability, the multiple columns for this matrix are separated into multiple parts, Parts 1 and 2, Table A-8 and Table A-9.

Table A-8 Tier 6 Unplanned Outage, Part 1

Solution	z/OS Global Mirror (XRC)	GDPS HyperSwap Manager with Metro Mirror	Rapid Data Recovery with TPC for Replication	PPRC Migration Manager
Platform	System z	System z	System z or Distributed	System z
Distance	any	<100km	>300 km	any
Connectivity	FICON ESCON	Fibre Channel TCP/IP	Fibre Channel TCP/IP	FICON Fibre Channel ESCON
Supported storage (primary site)	z/OS Global Mirror capable storage	PPRC-compliant storage	DS8000 DS6000 ESS	DS8000 DS6000 ESS
Supported storage (secondary site)	any z/OS supported storage	PPRC-compliant storage	DS8000 DS6000 ESS	DS8000 DS6000 ESS
Recovery Point Objective	few seconds to few minutes	near zero	near zero	near zero
Amount of Data	any	any	any	any
Other notes		Delivered as IBM Global Services Offering	Delivered as IBM Global Services Offering	Lower cost special bid offering, delivered through IBM Storage Services

Table A-9 Tier 6 Unplanned Outage, Part 2

Solution	Global Mirror on DS8000, DS6000, ESS	Global Mirror Metro Mirror on SVC/ DS4000	AIX LVM	AIX HACMP without Metro Mirror	AIX HACMP/XD with Metro Mirror
Platform	System z, Distributed, z+Distributed	Distributed	System p	System p	System p
Distance	any	any	10 km	any	metropolitan distances
Connectivity	Fibre Channel, channel extender	Fibre Channel, channel extender	Fibre Channel	Fibre Channel, TCP/IP	Fibre Channel, TCP/IP

Solution	Global Mirror on DS8000, DS6000, ESS	Global Mirror Metro Mirror on SVC/ DS4000	AIX LVM	AIX HACMP without Metro Mirror	AIX HACMP/XD with Metro Mirror
Supported storage (primary site)	DS8000 DS6000 ESS	With SVC - Any	any storage supported by AIX	any storage supported by AIX	DS8000 DS6000 SAN Volume Controller ESS N series
Supported storage (secondary site)	DS8000 DS6000 ESS	With SVC- Any	any storage supported by AIX	any storage supported by AIX	DS8000 DS6000 SAN Volume Controller ESS N series
Recovery Point Objective	3-5 seconds, assuming sufficient bandwidth and distance latency	bandwidth dependant	near zero, few seconds to few minutes	near zero, few seconds to few minutes	near zero, few seconds to few minutes
Amount of Data	up to 17 primary boxes with RPQ	DS4000 - 64LUNs SVC - 4096 LUNs	any	any	any
Other notes					

Tier 6 Transaction Integrity

The solutions for transaction integrity are specific to the database and application software being used. Because of this, the list of possible solutions is very broad, and it is not feasible to be all-inclusive. You should involve a software specialist skilled in the application and database set that you are using for detailed evaluation of transaction integrity recovery specific to your database and application. We do show common examples of solutions in these cells in the matrix.

Chapter 11, “High Availability clusters and database applications” on page 305 includes an extensive discussion on specific solutions and guidance relating to database and application level transaction integrity recovery.

Table A-10 shows examples for Tier 6 and Transaction Integrity.

Table A-10 Tier 6 and Transaction Integrity

Solution	- Database level transaction recovery on top of any Tier 6 Unplanned Outage recovery	Shadow database with forward recovery	Split Mirror Database with Metro Mirror
Platform	database and application specific	any	any
Distance	database and application specific	any	<300Km on DS8000, DS6000, ESS <100Km on DS4000, SVC
Connectivity	TCP/IP	any	Fibre Channel FICON
Supported disk storage (primary site)	any	any	PPRC-compliant disk subsystem
Supported disk storage (secondary site)	any	any	PPRC-compliant disk subsystem
Recovery Point Objective	near zero, few seconds to few minutes, minutes to hours (dependent on specific database, application, and hardware)	depending on the log shipping mechanism, loss of only few transactions possible	near zero
Amount of Data	any	any	any
Other notes			

Tier 5 Planned Outage

The solutions in Tier 5 are specifically defined as database and application *software* functionalities for planned, unplanned, and transaction integrity recovery. These solutions are dependent on each individual software's capabilities.

The scope of this book is to focus on hardware and operating system level System Storage Business Continuity solutions. You should involve a software specialist skilled in the application and database set that you are using.

However, as a general statement, robust databases have integrated software functionalities to enhance and minimize planned outages. See Chapter 11, "High Availability clusters and database applications" on page 305 for additional suggestions and guidance.

Table A-11 Tier 5 Planned Outage

Solution	Software, application, and database level facilities
Platform	software, application, and database specific
Distance	software, application, and database specific
Connectivity	Fibre Channel TCP/IP
Supported disk storage (primary site)	any
Supported disk storage (secondary site)	any
Recovery Point Objective	software, application, and database specific, typically defined by software policy
Amount of Data	any
Other notes	

Tier 5 Unplanned Outage

The solutions in Tier 5 are specifically defined as database and application *software* functionalities for unplanned and transaction integrity recovery. These solutions are dependent on each individual software's capabilities.

The scope of this book is to focus on hardware and operating system level System Storage Business Continuity solutions. You should involve a software specialist skilled in the application and database set that you are using.

However, as a general statement, robust databases have integrated software functionalities to do unplanned outage recovery. See Chapter 11, "High Availability clusters and database applications" on page 305 for additional suggestions and guidance.

Table A-12 Tier 5 Unplanned Outage

Solution	Software, application, and database level facilities
Platform	software, application, and database specific
Distance	software, application, and database specific
Connectivity	Fibre Channel TCP/IP
Supported disk storage (primary site)	any
Supported disk storage (secondary site)	any
Recovery Point Objective	software, application, and database specific, typically defined by software policy
Amount of Data	any
Other notes	

Tier 5 Transaction Integrity

The solutions in Tier 5 are specifically defined as database and application *software* functionalities for transaction integrity recovery. These solutions are dependent on each individual software's capabilities.

The scope of this book is to focus on hardware and operating system level System Storage Business Continuity solutions. You should involve a software specialist skilled in the application and database set that you are using.

However, as a general statement, robust databases have integrated software functionalities to do remote replication. To maintain transaction integrity, the database functionality must be integrated into whatever replication architecture is being used. See Chapter 11, "High Availability clusters and database applications" on page 305 for additional suggestions and guidance.

Table A-13 Tier 5 Transaction Integrity

Solution	**- Remote Replication transaction integrity with DB2, Oracle, SQL Server, and so forth.**
Platform	software, application, and database specific
Distance	software, application, and database specific
Connectivity	Fibre Channel TCP/IP
Supported disk storage (primary site)	any
Supported disk storage (secondary site)	any
Recovery Point Objective	- near zero, few seconds to few minutes, minutes to hours - dependent on user policy
Amount of Data	any
Other notes	

Tier 4, 3 Planned Outage

Table A-14 is for Tier 4, 3 and Planned Outage.

Table A-14 Tier 4, 3 and Planned Outage matrix

Solution	Point in Time Copy	Global Copy (PPRC-XD)	PtP VTS	Tivoli Storage Manager - Disaster Recovery Manager	TAPE
Platform	any	any	System z	Distributed	any
Distance	any	any	any	any	any
Connectivity		FICON Fibre Channel	FICON Fibre Channel		FICON Fibre Channel
Supported disk storage (primary site)	DS8000 DS6000 DS4000 SVC N series ESS	DS8000 DS6000 DS4000 SVC N series ESS	any	any	any
Supported disk storage (secondary site)	DS8000 DS6000 DS4000 SVC N series ESS	DS8000 DS6000 DS4000 SVC N series ESS	any	any	any
Recovery Point Objective	minutes to hours	minutes to hours	minutes to hours	minutes to hours	hours
Amount of Data	any	any	any	any	any
Other notes					

Tier 4 Unplanned Outage

Table A-15 is for Tier 4 and Unplanned Outage.

Table A-15 Tier 4 and Unplanned Outage matrix

Solution	PtP VTS Asynchronous TS7700 Grid	FlashCopy	FlashCopy Manager	Global Copy (PPRC-XD)
Platform	System z	any	System z	any
Distance	any	any	any	any
Connectivity	FICON ESCON	n/a	n/a	Fibre Channel FICCON
Supported disk storage (primary site)	any	any	any	DS8000 DS6000 DS4000 SVC N series ESS
Supported disk storage (secondary site)	any	any	any	DS8000 DS6000 DS4000 SVC N series ESS
Recovery Point Objective	minutes to hours	minutes to hours	minutes to hours	minutes to hours
Amount of Data	any	any	any	any
Other notes			IBM Storage Services Offering	

Tier 4 Transaction Integrity

Table A-16 is for Tier 4 and Transaction Integrity.

Table A-16 Tier 4 and Transaction Integrity matrix

Solution	Database-level journal file forwarding and remote application
Platform	software, application, and database specific
Distance	software, application, and database specific
Connectivity	Fibre Channel TCP/IP
Supported disk storage (primary site)	any
Supported disk storage vendor (secondary site)	any
Recovery Point Objective	- minutes to hours - dependent on user policy
Amount of Data	any
Other notes	

Tier 3 Unplanned Outage

Table A-17 is for Tier 3 and Unplanned Outage.

Table A-17 Tier 3 and Unplanned Outage matrix

Solution	FlashCopy	Tivoli Storage Manager	Tape
Platform	any	any	any
Distance	any	any	any
Connectivity	n/a	TCP/IP	n/a
Supported disk storage (primary site)	FlashCopy capable storage	any	any
Supported disk storage (secondary site)	FlashCopy capable storage	any	any
Recovery Point Objective	minutes to hours	minutes to hours	minutes to hours
Amount of Data	any	any	any
Other notes			

Tier 3 Transaction Integrity

Table A-18 is for Tier 3 and Transaction Integrity.

Table A-18 Tier 3 and Transaction Integrity matrix

Solution	**Database-level recovery using electronic tape vaulting**		
Platform	software, application, and database specific		
Distance	software, application, and database specific		
Connectivity	Fibre Channel TCP/IP		
Supported disk storage (primary site)	any		
Supported disk storage (secondary site)	any		
Recovery Point Objective	- minutes to hours - dependent on user policy		
Amount of Data	any		
Other notes			

Tier 2, 1 Planned Outage

Table A-19 is for Tier 2, 1 and Planned Outage.

Table A-19 Tier 2, 1 and Planned Outage matrix

Solution	**Tivoli Storage Manager**	**Tape**	
Platform	software, application, and database specific		
Distance	software, application, and database specific		
Connectivity	Fibre Channel TCP/IP		
Supported disk storage (primary site)	any		

Solution	Tivoli Storage Manager	Tape	
Supported disk storage (secondary site)	any		
Recovery Point Objective	- minutes to hours - dependent on user policy		
Amount of Data	any		
Other notes			

Tier 2, 1 Unplanned Outage

Table A-20 is for Tier 2, 1 and Unplanned Outage.

Table A-20 Tier 2, 1 and Unplanned Outage matrix

Solution	Tivoli Storage Manager	Tape	
Platform	software, application, and database specific		
Distance	software, application, and database specific		
Connectivity	Fibre Channel TCP/IP		
Supported disk storage (primary site)	any		
Supported disk storage (secondary site)	any		
Recovery Point Objective	- minutes to hours - dependent on user policy		
Amount of Data	any		
Other notes			

Tier 2, 1 Transaction Integrity

Table A-21 is for Tier 2, 1 and Transaction Integrity.

Table A-21 Tier 2, 1 and Transaction Integrity matrix

Solution	**Database-level recovery using physical tape transport**		
Platform	software, application, and database specific		
Distance	software, application, and database specific		
Connectivity	n/a		
Supported disk storage (primary site)	any		
Supported disk storage (secondary site)	any		
Recovery Point Objective	► hours to days hours ► dependent on user policy		
Amount of Data	any		
Other notes			

Additional business requirements questions

This section includes a list of additional business requirement questions that you need to answer prior to entering the Business Continuity Solution Selection Methodology.

Justifying Business Continuity to the business

Because Business Continuity solutions are by their very nature *insurance*, the following questions can help identify the ongoing daily payback value of a proposed Business Continuity solution.

You can partially or fully justify the requested investment for Business Continuity to senior management by quantifying the values below, and then portraying the proposed Business Continuity solution cost as only a portion of the anticipated ongoing daily benefits, as identified by these questions.

Tangible compelling IT values

Tangible compelling IT values include:

- Savings due to Planned Outage Reductions
  - Benefits and savings, revenue increases, due to business being able to operate without the planned outage
  - Benefits in personnel productivity
  - Savings in removing overtime compensation, overtime savings for planned outages

- Savings due to Better testability / maintainability of recovery solution
  - Lowered cost for every test
    - Savings due to lowered system resource impact for test
    - Savings due to better control of testing
    - Savings due to better reliability in testing
    - Savings due to better speed in test completion
  - Savings in maintainability costs, because GDPS insulates the Business Continuity scheme:
    - From application changes
    - From hard to define / hard to manage applications
    - From hard to manage / inability to manage data
- Benefits of absolute confidence in switch or cutover
  - Better information flow to decision team due to automation messaging of status.
  - Lowered cost of maintaining solution.
  - Increased accuracy of test and switch due to automation.
- Savings due to personnel cost reductions
  - Savings due to reduced labor and cost of a custom roll-your-own implementation.
  - Savings due to reduces labor costs due to automation.
  - Automation reduces salary/overtime cost of personnel to perform a recovery.
  - Pre-install and post-install difference in amount of staff required for change windows, unplanned outages, and practice and execute DR.
  - Savings because less costly / more available 'B' and 'C' personnel team can perform planned / unplanned outage recovery.
  - Lowered skill requirements for operations or recovery team.
  - Survivability without requiring key personnel.
- Benefit of providing DR after large Storage or Server Consolidation

  Benefits of providing efficient, trustable Recovery of large consolidated data center of servers / storage.
- Cost Savings of bringing DR In-house versus Out-sourced Service Provider
  - For same expenditure: Better recoverability, removal of dependencies on other Service Provider clients, no expiration time limit in recovery center.
  - Savings due to removal of out-sourced recovery center for equivalent functionality.

Tangible compelling business values

Tangible compelling business values include:

- Strategic and Competitive Advantage
  - 24x7 Internet customer availability required on new applications
  - Worldwide customer availability required on new applications
  - Meet mandatory regulatory requirements
  - Avoidance of large monetary impact to business of a disaster (customer $K/Hr.)
  - Exploit existing investment in installed equipment
  - Future regulation compliance in affordable, strategic approach
- Confidence
  - Regulatory agency confidence
  - Shareholder confidence
  - Financial Markets confidence
  - Senior management confidence and trust in the recovery
  - Maintenance of brand image
  - Willingness to use the recovery or switch because of the switch

- Tactical
  - Employee idling labor cost
  - 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
  - Interest value on deferred billings
  - Penalty clauses invoked for late delivery and failure to meet Service Levels
  - Loss of interest on overnight balances; cost of interest on lost cash flow
  - Delays in customer accounting, accounts receivable and billing/invoicing
  - Additional cost of working; administrative costs; travel and subsistence and so forth

Intangible compelling values - IT

Intangible compelling values for IT include:

- Value of DR Strategy that resolves failed previous DR schemes
- Personnel
  - Savings due to reduced number of storage administrators required per TB of disk storage
  - Recruitment costs for new staff on staff turnover
  - Training / retraining costs for staff
- Confidence in recoverability because of:
  - More frequent tests
  - Success of tests
- Planned Outage Reductions creates new options in testing or site maintenance
  - Confidence and accuracy value due to more frequent testing
  - Savings due to less expensive cost for testing
  - High confidence in switch
  - Value of prior and post 'Planned outage minutes / year'
  - Business impact of a planned outages / year (planned outage customer cost * planned outage minutes)
- Testing
  - Assuring successful recovery through increased frequency of testing
  - Catching errors in recovery through increased frequency of testing
  - Repeatability
- Automation value
  - Repeatability
  - Trustability

Intangible compelling values - business

Intangible compelling business values include:

- Unplanned Outage Revenue Loss Avoidances
  - Lost revenue
  - Loss of cash flow
  - Loss of customers (lifetime value of each) and market share
  - Loss of profits
- Unplanned Outage Cost Avoidances - IT
  - Cost of replacement of buildings and plant
  - Cost of replacing equipment
  - Cost of replacing software

- Unplanned Outage Business Impacts
  - Brand image recovery
  - Fines and penalties for noncompliance
  - Liability claims
  - Additional cost of advertising, PR and marketing to reassure customers and prospects to retain market share
  - Loss of share value
  - Loss of control over debtors
  - Loss of credit control and increased bad debt.
  - Delayed achievement of benefits of profits from new projects or products
  - Loss of revenue for service contracts from failure to provide service or meet service levels
  - Lost ability to respond to contract opportunities
  - Penalties from failure to produce annual accounts or produce timely tax payments
  - Where company share value underpins loan facilities, share prices could drop and loans be called in or be re-rated at higher interest levels.
  - Additional cost of credit through reduced credit rating

Business requirements questions for detailed evaluation team

The detailed evaluation team needs to address the following list of questions and answers in the course of quantifying, justifying, and designing the Business Continuity solution. Some questions are business in nature, others are IT or infrastructure in nature. They are the expanded super-set from which the basic Starter Set Business Requirements questions in “Starter set of business requirement questions” on page 330 are derived. We provide them here so that you can have a guideline for the types of information that will need to be gathered and analyzed by the detailed evaluation teams, to finalize an in-depth recommendation.

Business Profile

1. What is the customer business/industry?
2. What is the compelling reason for the customer to act at this time?
3. Who is the sponsor within the organization?
4. What is the budget that is allocated for this project?
5. When do they expect to have this implemented?
6. What are our chances of winning this business?
7. What are your goals that you feel are important for a successful project?
8. Which business sponsors do we need to engage with to properly determine the critical success factors for the project?
9. Will funding come from these mission critical business sponsors or from within the previously constructed I/T budget? Are funds allocated?
10. Have you designed an I/T recovery program, which incorporates various “speeds” of recovery in the event of interruption? Who is your current business continuance provider? Current contract expiration date?
11. Does your recovery plan take into account any acceptable level of transaction data loss and data unrecoverable? Explain.
12. What would the financial impact be on the interruption to your company due to some unexpected, unplanned catastrophic event?
13. Which business processes require an advanced level of recoverability in the event of an unplanned medium to a long-term interruption of I/T services?

14. Do you backup all of your company's critical data on a regular basis? Frequency? If a declared disaster occurred, would you be ready and able to restore your company's critical data to the Point of Failure?
15. Are critical applications replicated offsite in case of disaster? Can you access the site quickly with your staff in the time you have established?
16. If you don't have a business continuance program in place, what is the motivating factor associated with this change in strategy? Why are you interested in doing this now?
17. What is your current yearly cost associated with business continuance? If internal, approximate cost.
18. What is your current time frame for the business continuance project?
19. What type of disk do you currently use? Manufacturer? Total capacity? Mixed environment? Utility S/W? Upgrade plans? Explain.
20. Can you supply a total inventory list of all current server hardware?
21. Is your company a current IBM Hot Site customer?

Business Continuity Planning and Infrastructure Sizing

22. Has a sizing exercise been done? (Disk Magic for Metro Mirror or z/OS Global Mirror sizing)?
23. Is the implementation for data migration or Business Continuity?
24. If the implementation is for data migration, is the plan to minimize the amount of time in duplex by using the hardware bit maps?
25. Has the customer been made aware of the various Business Continuity documents and tools?
26. Will IBM Global Services or a Business Partner be involved in the implementation?
27. Has FlashCopy or some other point-in-time copy been considered?
28. What is the current Business Continuity objective (RTO or RPO)?
29. Recovery Time Objective (RTO)? What is your desired elapsed time objective from time of disaster until time of full recovery and accessibility to end-users? (Includes database recovery time), or,
30. Recovery Point Objective (RPO)? At the time that the RPO is complete, how much data is possible to recreate? (Measured in terms of seconds, minutes, and hours.)
31. How long does it take to Initial Program Load (IPL) the system following an unplanned system failure?
32. What is the critical application restart time after a system failure (after the system is IPL'ed)?
33. What is the planned system restart time under normal conditions (length of time to bring up your system)?
34. What is the planned system shutdown time (length of time to stop all applications and the system)?

35. What platforms are required to be recovered?
    - Z/OS
    - IBM System z
    - OS/390
    - VM
    - VSE
    - TPF
    - LINUX/390
    - UNIX
    - IBM System p
    - RS/6000®
    - AIX (non-clustered or clustered?)
    - SUN Solaris™ (non-clustered or clustered?)
    - HP-UX (non-clustered or clustered?)
    - IBM System i (AS/400®, OS/400®)
    - IBM System x™ and BladeCenter®
    - Windows 2003/XP
    - LINUX
    - Other
36. What storage mirroring technologies will be used (Metro Mirror, Global Mirror, Global Copy and so forth)?
37. Are Coupling Facilities being used?
38. How many/Type/Model/Vendor?
39. Are facilities to handle data integrity included?
40. Are there adequate resources for managing Internet security and intrusion, with ongoing monitoring and management?
41. Is the IT recovery strategy in line with the business objectives? Does the business or IT operations hinge on the availability of an individual person's skills?

Primary Side Hardware

42. How many Primary Control Unit's will be installed?
43. What is the vendor?
44. How many volumes/LUNs are expected to be recovered?
45. What processors are installed?
46. How many/Type/Model/Vendor?
47. Are there tape drives involved in this proposal (if so, describe)?

Secondary Side Hardware

48. Is the secondary site customer owned or are you using a Business Recovery Center? If so, which one? Customer Owned?
49. How many secondary Control Unit's will be installed?
50. What is the vendor?
51. How many volumes are expected?
52. What processors are installed?
53. How many/Type/Model/Vendor?
54. Are there tape drives (if so, describe)?

Performance

55. Has a bandwidth analysis been performed by collecting and analyzing data on the production applications?
56. What percentage of the workload is required to be mirrored?
57. What is the method of automation to be used (GDPS, other)?
58. Is cross platform data consistency required?
59. What platforms?
60. What level of consistency?

Connectivity

61. What is the distance to the remote site (miles or kilometers)?
62. What is the infrastructure to the remote site (Dark Fiber, Fibre provider / DWDM, Telecom line - what speed and flavor, T1 - 128KBytes/sec., T3 - 5 Mbytes/sec, OC3 - 19 Mbytes/sec, IP)?
63. Will channel extenders be utilized? If so, which channel extender vendor is preferred?
64. What is the write update rate (MBps, operations per second, how does it vary by time of day/month)?

B

Terms and definitions

This chapter addresses and defines some of the common terms that are used when determining or planning disaster recovery solutions and other terms that we use in this book.

Terms

This section provides a brief explanation of the various terms used in a disaster recovery discussion.

Automation

Software that recognizes and enables Disaster Recovery and minimizes the dependency on human intervention. Automation works directly with the devices or applications in a disaster recovery environment to ensure consistency and restore production.

Business Continuity

Business Continuity is a management process that relies on people, facilities, business processes, infrastructure, and applications to sustain operations at all times and under any circumstances.

Business Continuity Plan (BCP)

The BCP focuses on sustaining an organization's business functions during and after a disruption. A BCP can be written for a specific business process or can address all key business processes. The IT infrastructure addressed by the BCP is only based on its support for business processes or a subset.

Back-up Window Objective (BWO)

The BWO addresses the backup side of the disaster recovery solution. It represents how long operations can be halted while building/creating the backup data needed for the recovery. This parameter combined with the Network Recovery Objective (NRO), Recovery Point Objective (RPO), and Recovery Time Objective (RTO) are the main fundamentals of the disaster recovery solution.

Consistent data

Consistent data is mirrored or point-in-time data that can be verified to have completed all aspects of the copy process's last transaction. In a large environment using multiple storage systems or heterogeneous servers, consistent data requires that all copy processes stop at the same point-in-time. This does not necessarily imply zero data loss but does imply that such data could be used for a database restart.

Clustering

Clustering is a set of server images working together sharing the same resources. A good example of this in the mainframe environment is the Parallel Sysplex where an application can be spread through different images for Continuous Availability. In the open world this refers more often to a configuration with a hot stand-by server ready to take over the failing operating server.

Database recovery

Database recovery is the process of restoring the database from backups and using log files to update to the last consistent transaction. Depending on how recent the last backup was taken this could be a process measured in hours or days.

Database restart

Database restart is a process by which the database application is activated in the recovery facility and works as normal. To perform a database restart it is necessary to have data that is consistent.

Data loss

Data which was corrupted, in-flight, or inconsistent at the time of the disaster is unusable and, thus, lost until/if it can be recreated.

Data mirroring

Data mirroring is the process of replicating data between pairs of disk or tape subsystems. Data Mirroring can be based in software (such as DB2 using a two-phase commit), Operating Systems (such as AIX Logical Volume Mirroring), in the disk subsystem (such as Metro Mirror), or in a combination of two or more (such as Extended Remote Copy).

The method of data mirroring can be either synchronous or asynchronous.

- *Synchronous* data mirroring requires that a strict write order sequence be followed. Each update to the primary subsystems must also be updated in the secondary subsystems before another transaction can process. This results in near perfect data currency but can result in some lag time or latency between transactions.
- *Asynchronous* data mirroring allows additional transactions to process without ensuring receipt in the secondary subsystems. The lag in data currency could be seconds, minutes or hours, depending on the specific technology used, but there is usually low or no impact on transaction processing.

Disaster Recovery Plan (DRP)

The DRP applies to major, usually catastrophic, events that deny access to the normal facility for an extended period. DRP refers to an IT-focused plan designed to restore operability of the target system, application, or computer facility at an alternate site after an emergency. A DRP does not usually address minor disruptions that do not require relocation.

Enterprise Remote Copy Management Facility (eRCMF)

eRCMF is a service offering from IBM Global Services. enterprise Remote Copy Management Facility is a multi-site disaster recovery solution which manages data on volumes for Fixed Block (open systems hosts).

Extended Remote Copy (XRC)

XRC is a combined hardware and z/OS software asynchronous remote copy solution. All critical data is mirrored between the primary and secondary sites

Failback

Failback is the act of returning production to its original location after a disaster. Also a command under PPRC on ESS/DS6000/DS8000 (see in *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548).

Failover

Failover is the act of switching production to a back-up facility. Also a command under PPRC on ESS/DS6000/DS8000 (see *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548).

Fiber

Fiber is the term used to describe strands of cable made up of glass *threads* used to focus and transport light waves from one point to another.

Fibre
Fibre is the version of the word *fiber* used to indicate the difference between physical Fiber Optical Cable and the Fibre Channel Protocol (FCP).

FlashCopy
FlashCopy makes a single point-in-time copy of a LUN. This is also known as a *time-zero copy*. Point-in-time copy functions give you an instantaneous copy, or view of what the original data looked like at a specific point-in-time.

Freeze
Freeze is a process that halts I/Os to the primary Logical Subsystem of a mirrored pair. This results in data that is consistent to the point-in-time when the freeze occurred. When combined with automation this can prevent rolling disasters and ensure that recovery site databases can be restarted.

From an application perspective, the storage system is seen as being in the busy state.

Fuzzy copy
A fuzzy copy is mirrored data which has been received in the remote location but which is not verified to be consistent. PPRC-XD creates a fuzzy copy to quickly move tracks of data from one disk subsystem to another and is fully detailed in *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548.

Geographically Dispersed Open Clusters (GDOC)
GDOC is a multivendor solution designed to protect the availability of critical applications that run on UNIX, Windows or Linux servers. GDOC 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.

Geographically Dispersed Parallel Sysplex (GDPS)
GDPS is a multi-site application availability solution which provides the capability to manage the remote copy configuration and storage systems. It automates Parallel Sysplex operational tasks and does failure recovery from a single point of control, thereby improving application availability.

General Parallel File System (GPFS)
Available for both AIX and Linux Clusters, GPFS provides a cluster-wide file system allowing users shared access to files spanning multiple disk drives. GPFS is based on a shared disk model, providing lower overhead access to disks not directly attached to the application nodes, and using a distributed protocol to provide data coherence for access from any node.

Global Copy
Global Copy is an asynchronous remote copy function for z/OS and open systems for longer distances than are possible with Metro Mirror. With Global Copy, write operations complete on the primary storage system before they are received by the secondary storage system.

Global Mirror
Global Mirror provides a two-site extended distance remote mirroring function for z/OS and open systems servers. With Global Mirror, the data that the host writes to the storage unit at the local site is asynchronously shadowed to the storage unit at the remote site.

Global Mirror Utilities (GMU)

The GMU is a tool that can be used only to manage Global Mirror and not other Copy Services. It is distributed on a CD together with installation and user guide documentation with every ESS running Licensed Code (LIC) 2.4.0 or higher. It is available for ESS while DS6000 and DS8000 is planned for the 2nd half of 2005.

High Availability

High Availability is ensuring the resiliency of a device by removing single points-of-failure (SPOF). This reduces the chance of a disaster occurring due to a device failure, but it does not constitute Disaster Recovery, because it does not involve restoration of systems that have failed.

HACMP/XD (High Availability Cluster Multiprocessing XD)

HACMP allows continuous access to data and applications typically through component redundancy and failover in mission-critical environments. HACMP/XD (Extended Distance) manages failover to back up resources at remote sites.

Levels of recovery

There are three levels of recovery and each level builds upon the previous level:

- *Planned Outage* provides facilities for planned outages only. Typical usages include: application quiesces for backup, planned site switches, data migration, and relocation. Unplanned Outage recovery is not provided at this level, in other words, this level does not provide unplanned outage data integrity.
- *Unplanned Outage* data consistency and integrity is provided at the hardware and operating system level. Unplanned Outage level implies that Planned Outage support is also available. This level of recovery does not perform transaction integrity recovery at the application or database level.
- *Transaction Integrity* provides unplanned outage recovery at the application and database transaction integrity level. Solutions at this level typically rely upon an underlying level of Unplanned Outage support to be available.

Note that disaster recovery technology and solutions used for a given environment will vary depending on whether the level is Planned Outage, Unplanned Outage, or Transaction Integrity. More discussion on the levels of recovery is in Chapter 8, “Planning for Business Continuity in a heterogeneous IT environment” on page 251.

HyperSwap Manager

The HyperSwap Manager is a tool for GDPS/PPRC. It is designed to extend the availability attributes of Parallel Sysplex redundancy in a single site to disk subsystems. It provides the ability to transparently switch primary PPRC disk subsystems with the secondary PPRC disk subsystems for a planned or unplanned reconfiguration.

Metro Mirror

Metro Mirror is a remote data-mirroring technique for all supported servers, including z/OS and open systems. It is designed to constantly maintain an up-to-date copy of the local application data at a remote site which is within the metropolitan area (typically up to 300 km away using DWDM).

Metro/Global Copy

A three site solution, a combination of Global Mirror and Global Copy, called Metro/Global Copy is available on the ESS 750 and ESS 800. It is a three site approach and it was previously called Asynchronous Cascading PPRC. You first copy your data synchronously to an intermediate site and from there you go asynchronously to a more distant site

Network Recovery Objective (NRO)

The NRO is the requirement to establish network communications with the backup systems in a disaster recovery environment. This objective is expressed in terms of the amount of time that a business can afford to be without network communications. For example, a business that must have network communication established within two hours has a NRO of two hours.

Point-in-time Copy (PIT)

A point-in-time copy is an image that reflects selected data within a disk subsystem at a particular instance in time. PIT copies are used for backups, availability, and disaster recovery solutions. A good example of this is the FlashCopy feature on the IBM System Storage DS8000. Such images can be made from a single subsystem or across multiple subsystems.

Pick-up Truck Access Method (PTAM)

PTAM is the shipping of backup tapes from production data centers or to backup facilities by ground transportation. An often slow and unpredictable and also inexpensive method of data transportation.

Production (or primary) data center

The production data center is the data center in which production workload is processed. This can be paired with a Recovery data center to create two of the basic components of a disaster recovery solution. It is important to note that in some cases production workload or storage is based in both locations. In this case the location that is the production data center for one application can be the recovery data center for another application.

Remote copy

See *data mirroring*.

Recovery (or secondary) data center

A recovery data center is a data center facility which has been established primarily to hold data and systems that will be used to recover in the event of a disaster. Recovery data centers can be and often are used for other purposes such as testing or development during non-disaster situations.

Recovery Point Objective (RPO)

The RPO is the requirement for currency of data. Also expressed in the amount of data that could acceptably be recreated post disaster. For example a business that believes it could acceptably afford to create or lose 5 minutes worth of data has a RPO of 5 minutes.

Recovery Time Objective (RTO)

The RTO is the requirement for restoration of systems and applications, expressed in terms of how long a business can afford to have systems and applications down after a disaster. For example, a business that believes that it could afford to be without systems for 8 hours has a RTO of 8 hours.

Rolling disaster

A rolling disaster occurs when failures strike disk subsystems at different intervals. An explosion that damages disk subsystem A immediately and then strikes subsystem B three milliseconds later could result in inconsistencies between A and B. The failures *roll* across the data center, therefore a rolling disaster.

Single Point-of-Failure (SPOF)

An SPOF is any one component of a device or plan, the disabling of which, would result in that device or plan no longer functioning.

Tiers

The seven categories of disaster recovery solutions defined by SHARE are known as *tiers*. They range from 1 (off site backup with no recovery data center) to 7 (highly automated recovery with zero or little data loss) depending on the speed and value of the solution. See 4.2, “A breakdown of the seven tiers” on page 139 for further discussions on tiers.

Virtualization

Virtualization is the use of technology to change the way that storage is viewed. Multiple storage systems can be viewed and managed as a single level or as separate pools of data instead of being viewed individually. Such technology is available through devices such as the SAN Volume Controller and through the IBM System i.

Virtual Disk Service (VDS)

VDS is a component of the Microsoft Windows Server® 2003 infrastructure that simplifies disk management by enabling the administrator to use a standardized storage management user interface to manage multi vendor storage arrays and direct attached storage.

Volume Shadow Copy Service (VSS)

VSS is a component of the Microsoft Windows Server 2003 infrastructure that enables your storage array to interact with third-party applications that use the VSS Application Programming Interface (API). Volume Shadow Copy Service is able to create shadow copies through its coordination with business applications, backup applications, and storage hardware.

C

Services and planning

In this appendix we discuss the following service related topics:

- The services required and who can provide them
- IBM Global 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 that we illustrate in this book rely on interactions between multiple components that can come from multiple vendors. The components can 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
- Application software automation and functions 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 customers, 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 to choose IBM Global Services or an IBM Business Partner? This question is not easily answered. It depends on 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 customer as whether to involve IBM Global Services or an IBM Business Partner.

IBM Global 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. 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* or *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 C-1.

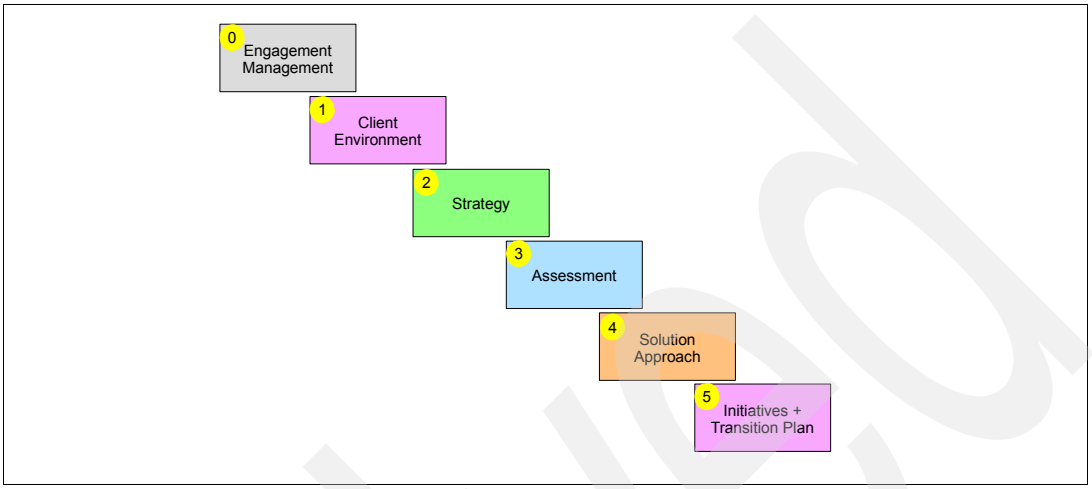

Figure C-1 IBM Global Services Method for resiliency and storage related services

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, 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 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 Step 4 and 5 define the path to reach the goal of where we want to go.

The approach illustrated above applies to more comprehensive consulting type of engagements. Another way of defining and classifying services is illustrated in Figure C-2. This is a more solution life cycle oriented classification.

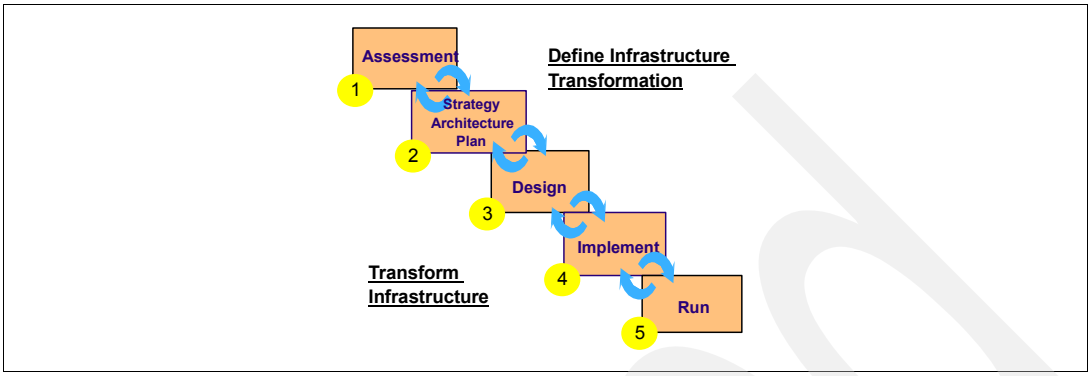

Figure C-2 IBM Global Services classification

Steps 1 and 2 are the steps where you decide what business needs you have to address. After these needs are understood and documented, we can proceed to define a target storage environment architecture. This phase can 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 perform the transformation itself:

- *Step 3 - Design* defines the detailed physical and logical design of different components of the architecture. Step 3 can 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* address storage management activities. IBM Global Services Managed Storage Services offer customers the opportunity of out tasking storage management related activities to IBM Global Services.

You can 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 might 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 across major initiatives to solve specific customer problems or sets of problems. Figure C-3 shows the On Demand initiative alignment. The On Demand framework addresses all customer major focus areas: Virtualization, Automation and Integration.

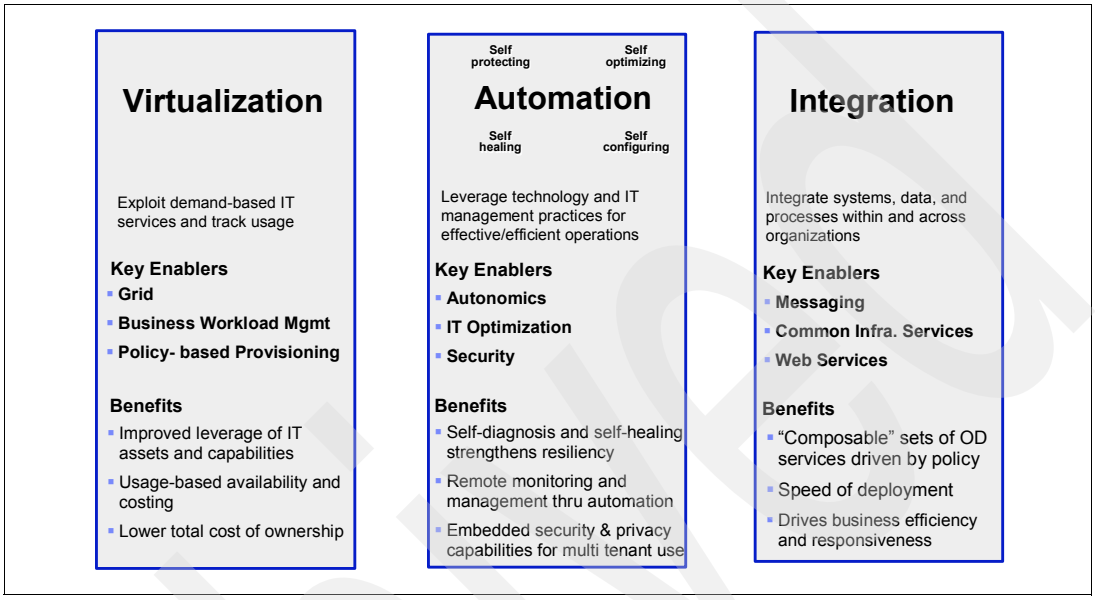

Figure C-3 On Demand initiatives

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 C-4.

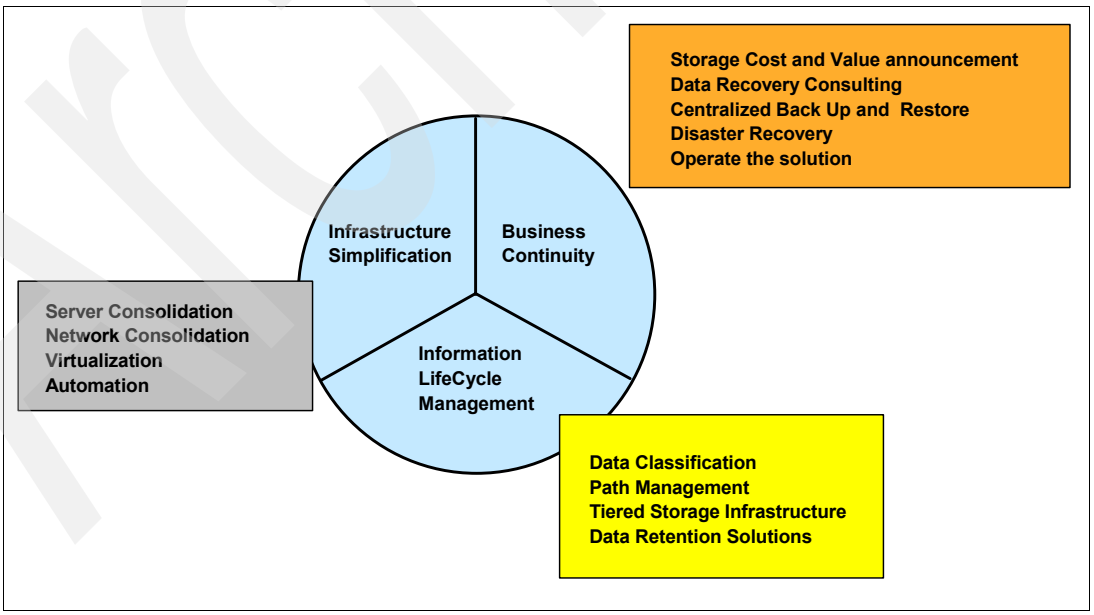

Figure C-4 IBM Global Services solution classification for System Storage

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 C-5. IBM Global Services offers a vast portfolio of services that can be tailored to meet virtually any customer requirement.

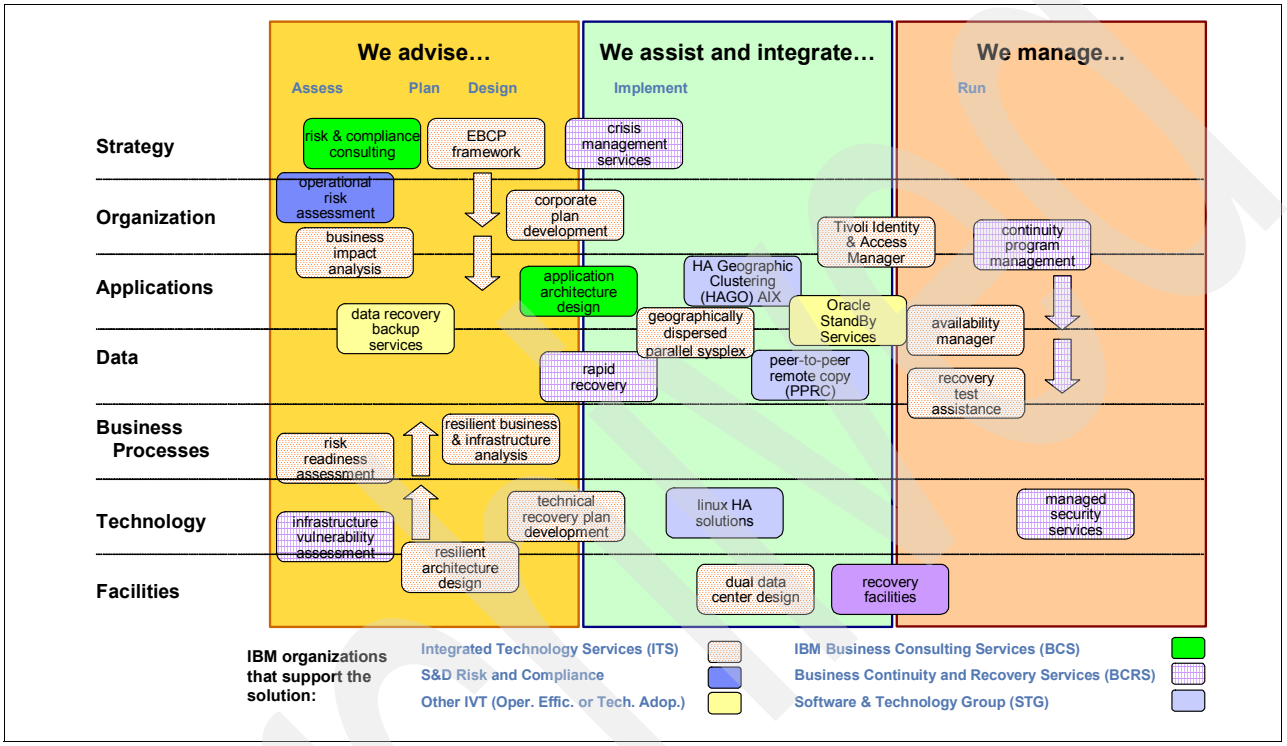

Figure C-5 IBM Global Services Sample Business Resiliency services

Figure C-5 illustrates a suggested set of services that you could use to build and manage a resilient infrastructure.

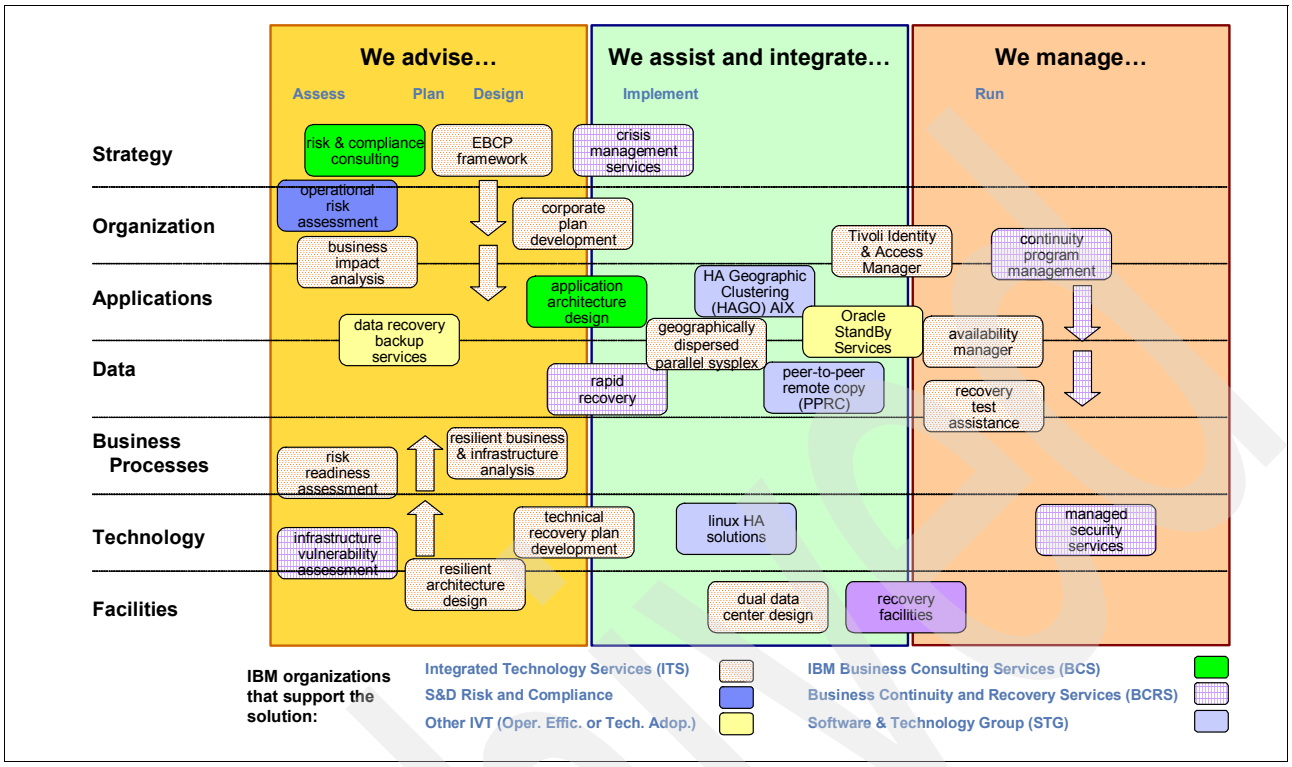

Figure C-6 Possible IBM Global Services for a resilient infrastructure

For more information about IBM Global Services see one of the following Web sites or contact your IBM representative or IBM Business Partner:

`http://www.ibm.com/services/us/index.wss/gen_it`

For storage services see:

`http://www.ibm.com/services/us/index.wss/it_services/its/a1000416`

For Business Continuity and resiliency Services see:

`http://www.ibm.com/services/us/index.wss/it_services/bcrs/a1000411`

For infrastructure and systems management services see:

`http://www.ibm.com/services/us/index.wss/it_services/its/a1000415`

For technical support services see:

`http://www.ibm.com/services/us/index.wss/dt/its/a1001327?cntxt=a1000416`

In the next sections, we give an overview of three sample storage services. These are just examples, for a complete list and more in depth information refer to the Web sites that we listed previously.

IBM Managed Hosting storage and backup services

IBM Managed Hosting storage and backup services are designed to provide managed storage capacity and other cost-effective solutions to help increase your data availability and protect your critical information from accidental loss or destruction. Flexible backup and restoration services, expert technical assistance and managed storage capacity can help address the rapidly growing data volumes associated with your on demand business.

Highlights of IBM Managed Hosting storage and backup services include:

- Offer on demand access to scalable storage capacity, helping you accommodate and plan for growth
- Support next-generation storage-intensive applications and expansion with scalable tape library options
- Provide off-site storage to facilitate Disaster Recovery
- Offer a security-rich storage environment to assure the integrity of business-critical data
- Provide access to specialized technical skills

IBM Managed Hosting services offer the following benefits:

- SAN managed storage services

  Managing local data stores can slow system performance and impact database performance. With SAN (storage area network) managed storage services, you can store large volumes of data in a security-enhanced SAN environment provided and managed by IBM. SAN managed storage services are designed to allow you to purchase reliable, high-performance storage capacity for a flat monthly charge for each shared or dedicated Fibre Channel access connection to the SAN environment, plus a monthly charge per gigabyte of allocated disk space. Data can be stored at the same IBM e-business Hosting™ Center where your servers are located or at multiple e-business Hosting Centers around the world.

- Standard backup and restore

  With the standard backup and restore service for hosting environments, you gain reliable, flexible, security-rich, network-based backup (through IBM Tivoli Storage Manager) at a very affordable price. That's because the costs for hardware and management tools, such as servers, storage pools and tape libraries, are shared among all clients of the service. At the same time, IBM employs security features and management mechanisms at the network, backup server and storage device layers to help ensure separation and protection of your data. You maintain control over the actual backup process itself. If extra storage or data transfer is necessary, you can initiate ad hoc backups at any time. Billing is based on the amount of data transferred to backup or from restore. You pay for the services you use.

- Dedicated backup and restore

  With a dedicated storage management server to handle tape backup and restoration functions, you gain a higher degree of flexibility and schedule control than the standard backup and restore service allows, especially if you have large volumes of data. You can choose between lower-cost shared tape library resources or optional dedicated resources.

- Online hot database backup

  When you have mission-critical data and cannot tolerate the downtime associated with offline backup and restoration services, the online hot database backup service is designed to allow you to back up selected databases without interrupting business transactions. The 24x7 capability can be added to our standard backup and restore service. You can select the databases and frequency of the backups. IBM DB2, Informix®,

Oracle and Microsoft SQL Server databases are supported. Billing consists of a one-time setup charge per database, plus the ongoing charges associated with the standard backup and restore service.

- High performance backup and restore

  As multi-terabyte databases become common and available backup windows continue to shrink, the issue of backing up large databases in the shortest possible time has become increasingly urgent. High performance backup and restore provides backup and restoration of large amounts of data in significantly less time than standard services by moving data transfers from the IP communications network to a data network or SAN. These fiber connections provide throughput superior to LAN-based IP connections because they support large frame size and high-bandwidth transfers. High-throughput data transfer means shorter backup and restoration intervals, which translates to higher application availability and potentially lower costs.

- Offsite data storage

  You can increase the level of data protection provided by our standard backup and restore service by adding offsite data storage. Duplicate copies of your backup tapes are automatically created by Tivoli Storage Manager and transferred daily to a security-rich offsite storage facility. Should recovery be impossible using onsite backup media, quick access to these vaulted copies can help limit business disruptions. Billing is per gigabyte stored in the offsite facility, with separate charges for each emergency media recall.

- Tape library partition

  You can acquire terabytes of cost-effective, scalable tape storage with tape library partition services, designed to provide dedicated access to a logical library partition installed in the shared backup infrastructure at the IBM e-business Hosting Center. The partition includes one or more tape drives plus tape library slots dedicated to your use, with Fibre Channel connections for high-speed data access.

- Storage assistance

  Effective backup and restoration strategies are crucial to maintaining the availability of Web-based business processes. Storage assistance offers on-call consulting support from skilled IBM personnel who can help you implement a sound backup and restoration strategy for your hosting environments and provide onsite support for operational tasks performed at the IBM e-business Hosting Centers, including restoration of data files and review of daily server logs.

- Hosting services are designed for on demand business

  IBM e-business hosting offers a flexible array of scalable services that grow with your business and enable you to respond to a rapidly changing marketplace dynamically. The services are modular so you can choose what you need, when you need it. As your business requirements change or accelerate, IBM will work with you to build a services solution that is capable of addressing them.

Resilient business and infrastructure assessment

An important aspect of on demand infrastructure is resiliency—having the flexibility to respond rapidly in the event of changes and threats—be they computer viruses, earthquakes, or sudden spikes in demand for IT resources.

Highlights of resilient business infrastructure include:

- Examines the many working layers required for an optimal infrastructure implementation, including strategy, organization, business and IT.
- Compares your environment to other similar companies in your industry.

- Provides a definition of potential threats or disruptions, prioritizes the level and impact of risk.
- Recommends areas of your business that need to be addressed to meet business goals.
- Incorporates business and IT components into a single, security-enhanced and competitive resiliency strategy.

A resilient business and infrastructure assessment helps you evaluate the ability of your infrastructure to:

- Provide a data and application environment that consists of systematic methods, processes and technologies and that is designed to be security-rich, agile, available and recoverable.
- Enable changing business models and strategies and link business strategy to risk tolerance and IT readiness levels.

IBM resilient business and infrastructure solutions

Today's business requires resiliency to be imbedded into the fabric of the business infrastructure. A resilient business infrastructure provides a security-rich, agile, available and recoverable environment that can handle planned and unplanned events, and positive and negative impacts to business. Survival is dependent upon the rapid response of both critical business processes and the supporting infrastructure. IBM resilient business and infrastructure solutions include:

- Identifies business processes and elements that are critical to a resilient, agile enterprise.
- Determines risks (stresses and demands) to your business.
- Evaluates your current business and technology infrastructure using a six layer framework.
- Creates vulnerability, responsiveness and prevention indices.
- Prepares next steps and recommendations for mitigating risks.

In your effort to achieve a resilient, agile enterprise, we can help you understand the linkage required between business processes and the technologies that enable them. We take a broad view of your business and use a structured approach to assess the resiliency of your enterprise across six key areas:

- Strategy
- Data and applications
- Technology
- Business and IT processes
- Organization
- Facilities and security

We then evaluate how well the IT functions support the needs of the business functions. Strengths and weaknesses are identified and reflected in sophisticated indices, including a comparison of your company's resiliency to other similar organizations. We can help you address any weaknesses, utilizing our wide range of expertise and service offerings.

Other storage services

In this book we illustrate many storage services in the various solution chapters. For additional information about the GDPS and HACMP/XD, see the Continuous Availability chapter in *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548.

Network Consulting and Integration services

IBM Global Services offers the broadest range of services offerings in the industry for Networking solutions. Their services capabilities, skills depth and breadth, and the IBM reach and range are simply unmatched in the networking services industry.

There are two basic dimensions to their networking services offerings:

- The Network Life cycle dimension
- The Network Technology dimension (for example, routing and switching, VPN/Security, wireless, Optical Networking, and so forth)

The ITS networking services offering portfolio covers both dimensions. You can find comprehensive information about IBM Network Services at:

http://www.ibm.com/services/us/index.wss/az#N

Optical/Storage Networking

Here are some of the service offerings in the area of optical and storage networking, Network deployment and cabling.

IBM Network Consulting for Optical Networks

This offering provides a set of supporting materials (sales positioning presentations, SOW, technique papers, spec sheet and Webcast) for optical networking consulting engagements. This offering enhancement extends the 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 customer 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, specifications sheet Webcast) 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 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, and 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 Webcast 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 give each branch direct, full-speed connections to centralized storage through Dense Wavelength Division Multiplexing (DWDM) technology.

This offering provides supporting materials (spec sheet, sample SOW and 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 customers 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, specifications sheet, and Webcast) 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.

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 IC 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 can include, procuring hardware, configuration, installation, cabling, site services, site surveys, and so forth. for network upgrades, enhancements or new installations at multiple sites both within a single country and multi-nationally.

IBM Network Consulting, 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 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 can be used to implement the design.

The guide contains suggestions 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.

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 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.

D

Networking terminology tutorial

Networking is one of the most important technology components of any IT infrastructure. This appendix is a brief tutorial on the terminology and concepts of networking. It is intended to:

- Provide a technical supplement to Chapter 10, “Networking and inter-site connectivity options” on page 281.
- Define and position the large number of networking technologies relative to each other, using the Open Systems Interconnect architecture to unify the technologies.
- Promote a better understanding of many commonly used networking concepts so that you can better evaluate and select the appropriate networking components for your Business Continuity solution.

Open Systems Interconnect network model

To understand modern IT networking, we need to understand the architecture upon which all of today's networking technology is founded and built. That model is known as the Open Systems Interconnection model or *OSI* for short. Understanding the OSI architecture gives us the proper framework to understand the myriad of networking acronyms.

Figure D-1 illustrates an all-in-one overview of the OSI model. We go through this chart, level-by-level. When we finish, you will be able to define many common networking technologies, to position them in relation to each other, and to understand the basic role that each can play in your final networking design.

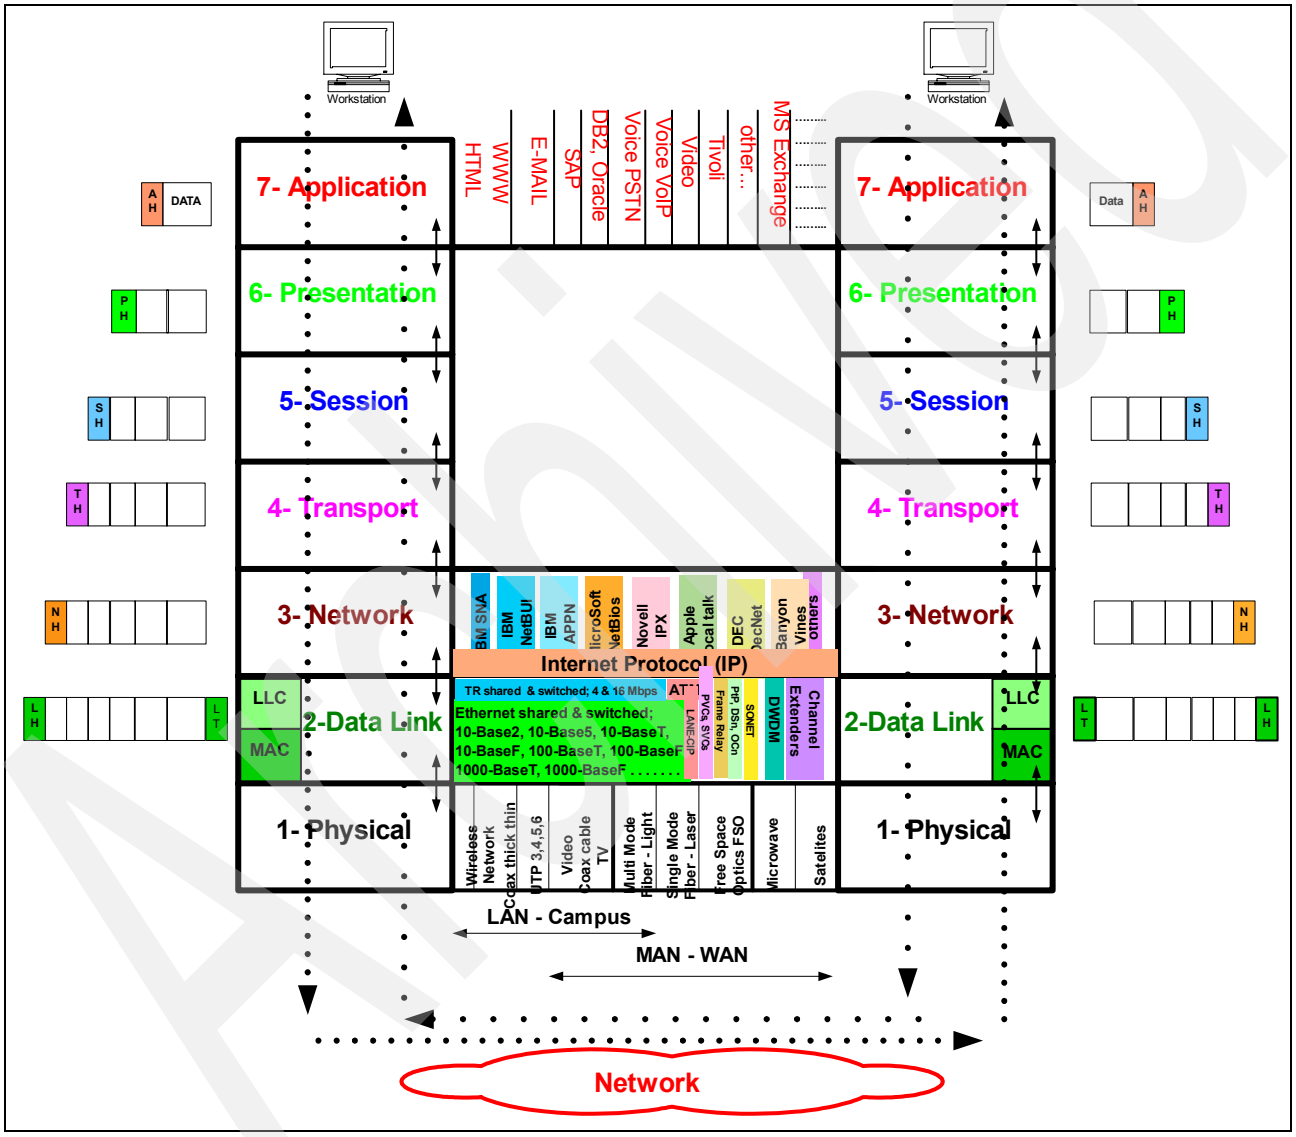

Figure D-1 OSI network model overview

There are seven layers, and each of the different layers represents a specific functionality that allow Layer 7 applications (such as e-mail, Lotus® Notes®, Web servers, and Web browsers) to connect to each other. This chart represents the details of what is required to make that happen.

There are headers shown on the left and right hand side of the OSI chart. The application's data is progressively encapsulated with headers as the data is passed from layer to layer, in preparation for transmission to the other site. At the other end, headers are progressively removed as the data is passed up towards the receiving application.

The header abbreviations are:

- AH - application header
- PH - presentation header
- SH - session header
- TH - transport header
- NH - network header
- DH - data link header

Let us now zoom into this chart and examine the OSI layers one-by-one.

OSI layer 1 (physical layer)

Figure D-2 shows the OSI layer 1 (the physical layer), showing the types of physical connections that can reside within it.

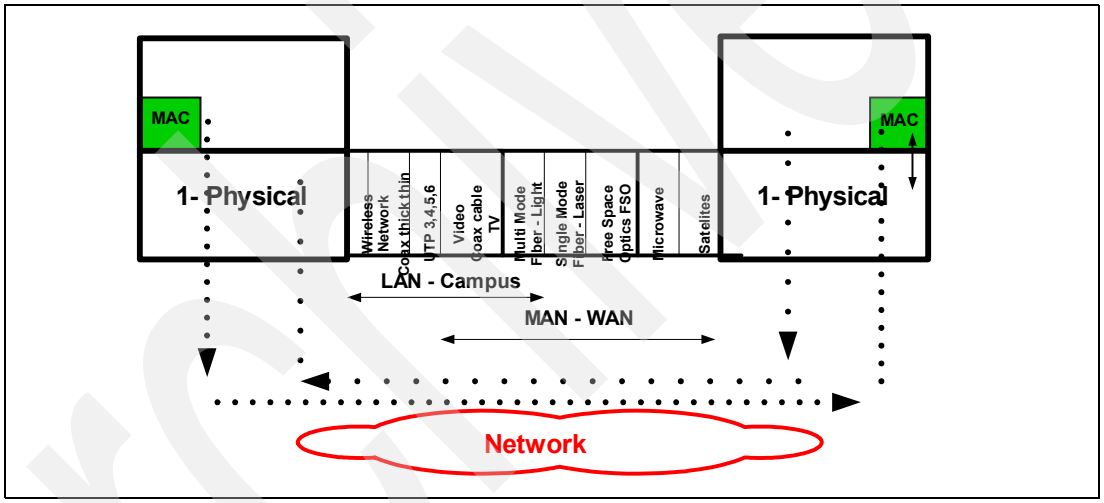

Figure D-2 OSI layer 1 - physical layer

This layer is the physical connection layer. As shown in Figure D-2, physical connections between two ends of the network can consist of (but are not limited to):

- Wireless, including variants of today's wireless LAN technology and cell phone technology.
- Network thick or thin coax.
- Unshielded Twisted Pair (UTP), which comes in various levels (3, 4, 5, 6), each of which describes a certain level of drive distance, bandwidth capability, and resistance to electromagnetic interference.
- Video and cable, television coax.
- Fibre optic cable, the current strategic physical connection networking technology. Fibre optic cable can be either multi-mode or single mode cable. These modes are defined in Figure D-8 on page 389.

- Free Space Optics, wireless or microwave technologies to handle one of the biggest problems in using the strategic fiber optic capability, which can be: “how do I get fiber run *the last mile*” from the network provider's closest connection point, to my particular data center?'
- Microwave and Satellites - wireless technologies used for networking communication especially at very long distances requiring no physical interconnect.

As you can see, the LAN-Campus arrow shows that in the LAN-Campus environment, the networking Layer 1 physical connection technologies are usually one or more of the following: Wireless, Network coax, Unshielded Twisted Pair, Video coax, and Multi-mode fiber optic.

Conversely, in the Metropolitan Area Network or the Wide Area Network (that is, longer distance than LAN-Campus) the Layer 1 physical interconnection technologies are typically: UTP, video coax, multi-mode and single-mode fiber optic, free space optics, microwave, and satellites.

To begin any networking connection, the first requirement is that a layer 1 connection must be established using one or more of these technologies.

OSI layer 2 (data link layer)

Having established Layer 1 (the physical connection), which at this point is waiting for something to travel upon it (whether that be electrical signals, light, or radio), the next step is to establish the proper layer 2 (data link connection).

Figure D-3 shows the OSI data link layer 2, showing physical protocols that reside within it.

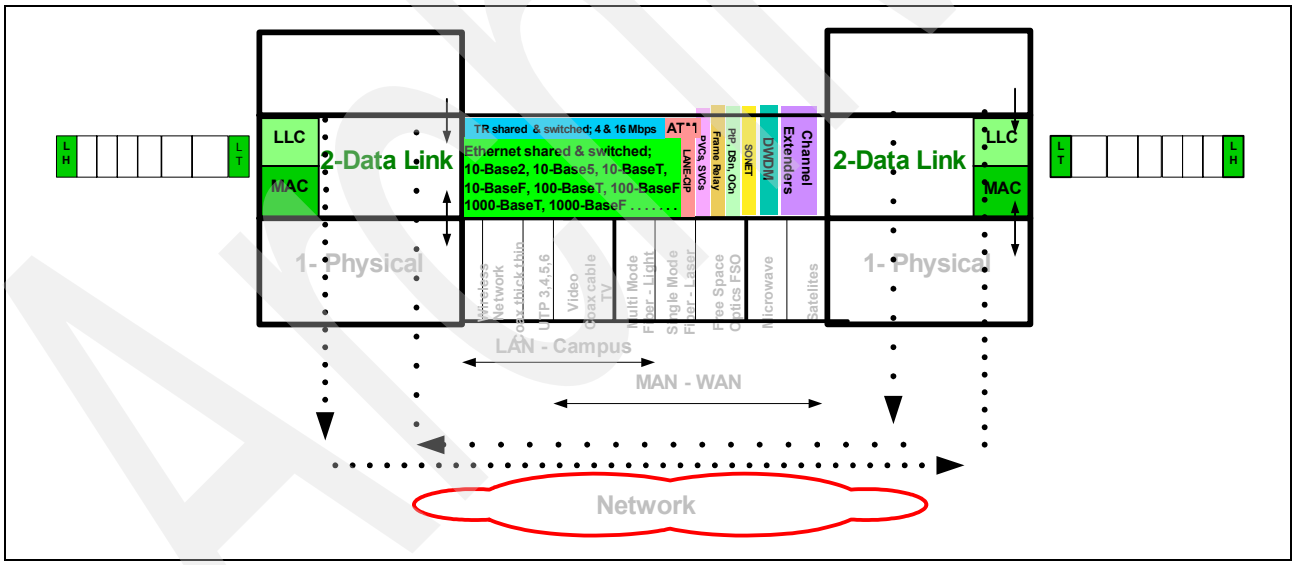

Figure D-3 OSI layer 2 - data link layer

You might think of the data link connection as the physical protocol that allows a ring or a network of many physical connections to be managed, load-balanced, configured and reconfigured (dynamically in many cases). This layer specifically refers to managing, load-balancing, and configuring at the physical connection level, not at the logical software connection level.

Tip: The standard networking speed nomenclature is that a lowercase *b* stands for *bits* and an uppercase *B* stands for *Bytes*. Thus, a line which runs at:

- 1.544 Mbps is a 1.544 Mbps line, in other words, approximately 200 Kbps
- 100 BaseT, a 100 Mbps LAN, is approximately 10 MBps

Always make sure you note the difference. Otherwise, your conversation could be off by a factor of 8 or 10, in terms of the speeds that you are discussing.

In OSI Layer 2 data link layer, the various technologies in this layer include (but are not limited to):

- Ethernet (shared and switched, with various speeds including10-BaseT (10 Mbps or approximately 1 MBps raw), 100 BaseT (100 Mbps or approximately 10 MBps raw), 1000 BaseT (100 Mbps or approximately 10 MBps raw, commonly known as *Gbit Ethernet*).

  Layer 2 note: In the LAN/campus environment, Ethernet has become the standard data link protocol that is used.

- Token Ring, at either 4 Mbps (.5 Mbps) or 16 Mbps (2 MBps).
- Asynchronous Transfer Mode (ATM): A method of transmitting data over this data line with asynchronous characteristics (and reassembling the data in proper sequence at the receiving end).
- Private Virtual Circuits and Sxxx Virtual Circuits (PVC and SVC): A method of negotiating and establishing a private, secure connection over this data link even though we are physically sharing Layer 1 physical links.
- Frame relay: Another technology designed to a method of transmitting data over this data line with asynchronous characteristics (and reassembling the data in proper sequence at the receiving end).
- PtP DSn OCn: Point to point connections, DSn, OCn (Optical Cable) data link connections. Typically used by telcos to describe some aspect of the type of connection and the raw speed potential in the case of DSn, Ocn.
  - OC-3 denotes a raw speed capability of approximately 19 MBps.
  - OC-12 denotes a raw speed that is 4x the capability of OC-3, in other words, approximately 78 MBps.
- SONET: A telco architecture that manages large (usually fiber optic) networks, and allows powerful problem determination, performance monitoring, and reconfiguration capabilities. In and of itself, SONET does not imply any particular speed.

  These are raw speeds, and the effective transfer rate due to overhead from a Layer 3 and above functions must be deducted from this raw speed.
- Dense Wave Division Multiplexing (DWDM): A method of multiplexing multiple channels of fiber optic based protocols (such as ESCON, Fibre Channel, FICON, Gbit Ethernet) onto on physical cable, by assigning different wavelengths of light (that is, *colors*) to each channel; then fanning it back out at the receiving end. Major players in the enterprise class DWDM marketplace are: Nortel Networks, Cisco (ONS 15540), and Lucent.

  Dense Wave Division Multiplexors (DWDM) are data link Layer 2 tools. Thus, the typical DWDM machine does not perform any switching, routing or protocol conversion.

- Channel Extenders: These are devices which take a data center protocol (typically Fibre Channel, ESCON, FICON, or Fibre Channel) and convert it to Wide Area Network protocols for transmission to its companion machine at the remote site, where the signal is then converted the signal back to its original form.

  For SAN Distance extension (that is, Fibre Channel), there are many vendors and Business Partners. For more information, see the System Storage Proven™ Web page at:

  http://www-03.ibm.com/systems/storage/proven/index.html

- Media Access Control (MAC): A data link functionality which defines how the computer in question participates or is identified in the network from a hardware standpoint; the best known example of a MAC address is the firmware unique hardware address.
- Logical Link Control (LLC): A data link functionality which defines how the computer in question participates or is identified in the network from a network address standpoint. The most prevalent LLC protocol you have probably seen labeled as IEEE 802.2.

OSI layer 3 (network layer)

Figure D-4 shows the OSI data link layer 3, showing logical network protocols that reside within it.

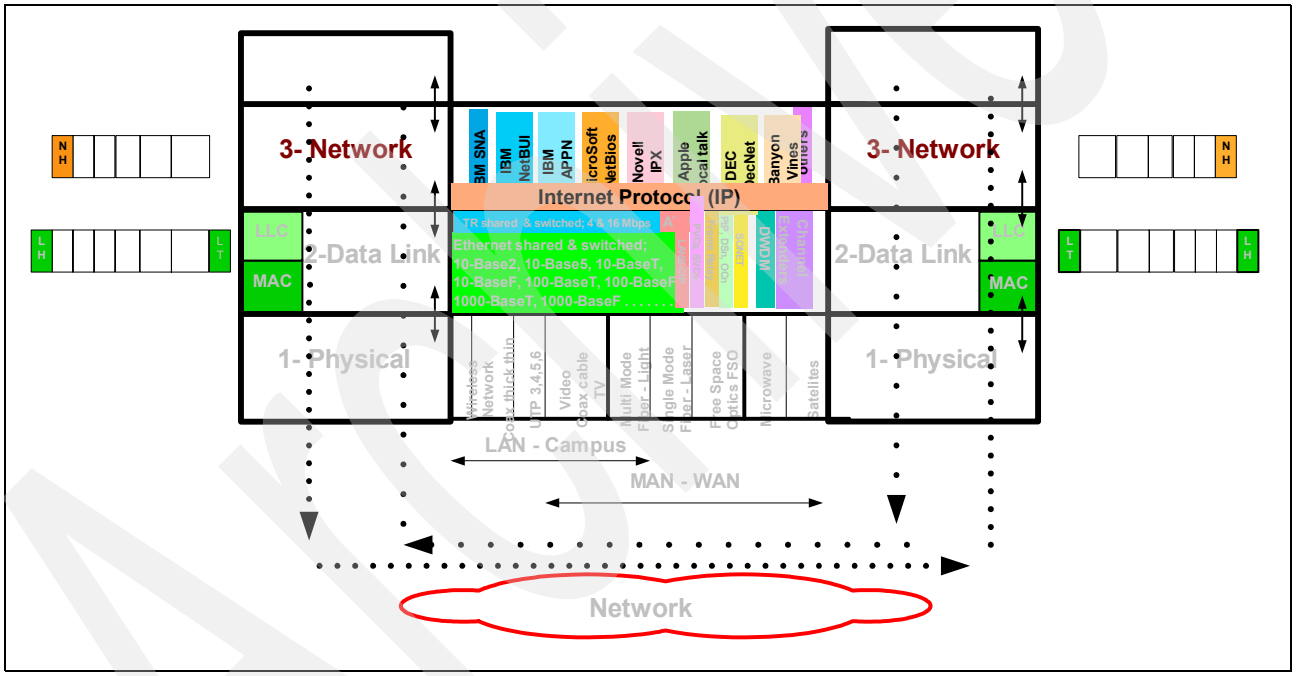

Figure D-4 OSI layer 3 - network layer

You can think of the OSI layer 3 network layer as being the actual logical *language* that the two communicating computers use to talk to each other. Let's illustrate with an example using a telephone call as an analogy (this is only an analogy for explanation purposes):

- Layer 2 Data Link: You are in San Jose, California, U.S., you pick up a telephone, you get dial tone, you dial Kuwait, and someone answers. This is an analogous example of a successful OSI layer 2 data link connection using the telco's OSI layer 1 physical connection.
- Layer 3 Network Link: Now, you start speaking English. The other person starts speaking Arabic. Neither of you understand each other. This is a failed layer 3 network link. If either of you switches to a language you both understand, that is a successful layer 3 network link.

With that as an example, the layer 3 network link logical protocols include (but are not limited to):

- IBM SNA - IBM Systems Network Architecture logical protocol, heavily used in the past for example in 3270/VTAM/NCP networks
- IBM NetBEUI - an IBM PC LAN logical protocol
- IBM APPN: IBM Advanced Program to Program Networking - an older IBM logical protocol for program to program communication
- MicroSoft NetBIOS: MicroSoft LAN networking logical protocol
- Novell IPX: Novell LAN networking logical protocol
- Apple LocalTalk: Apple LAN networking logical protocol
- DEC DecNet: DEC LAN and inter-DEC computer logical protocol
- Banyan Vines: Banyan LAN logical protocol
- IP - Internet Protocol is the best known logical protocol today; it is the logical protocol upon which TCP/IP and the Internet uses to connect the world with each other

Layer 3 note: IP is not a hard requirement. However, IP is shown as overlaying all protocols, because it has become the *de facto world standard* in the campus and network.

Often the other logical protocols that are listed are encapsulated within IP message frames, to exploit the position of the IP as the de facto world standard.

Note: While each OSI functionality layer is distinctly identified, any individual network component is likely to have some overlap of functionality across layers.

For example, a *layer 3 switch* is probably going to have aspects of layer 2 or layer 4 in it also. Understanding the architectural functions of a layer allows you to decode what a vendor or technician really means when they loosely say *layer 2 switch*, *router*, *bridge*, or *hub*.

OSI Layer 4, 5, 6, 7 - Transport, session, presentation, application layers

The additional layers are transport, session, presentation, and application and are shown in Figure D-5.

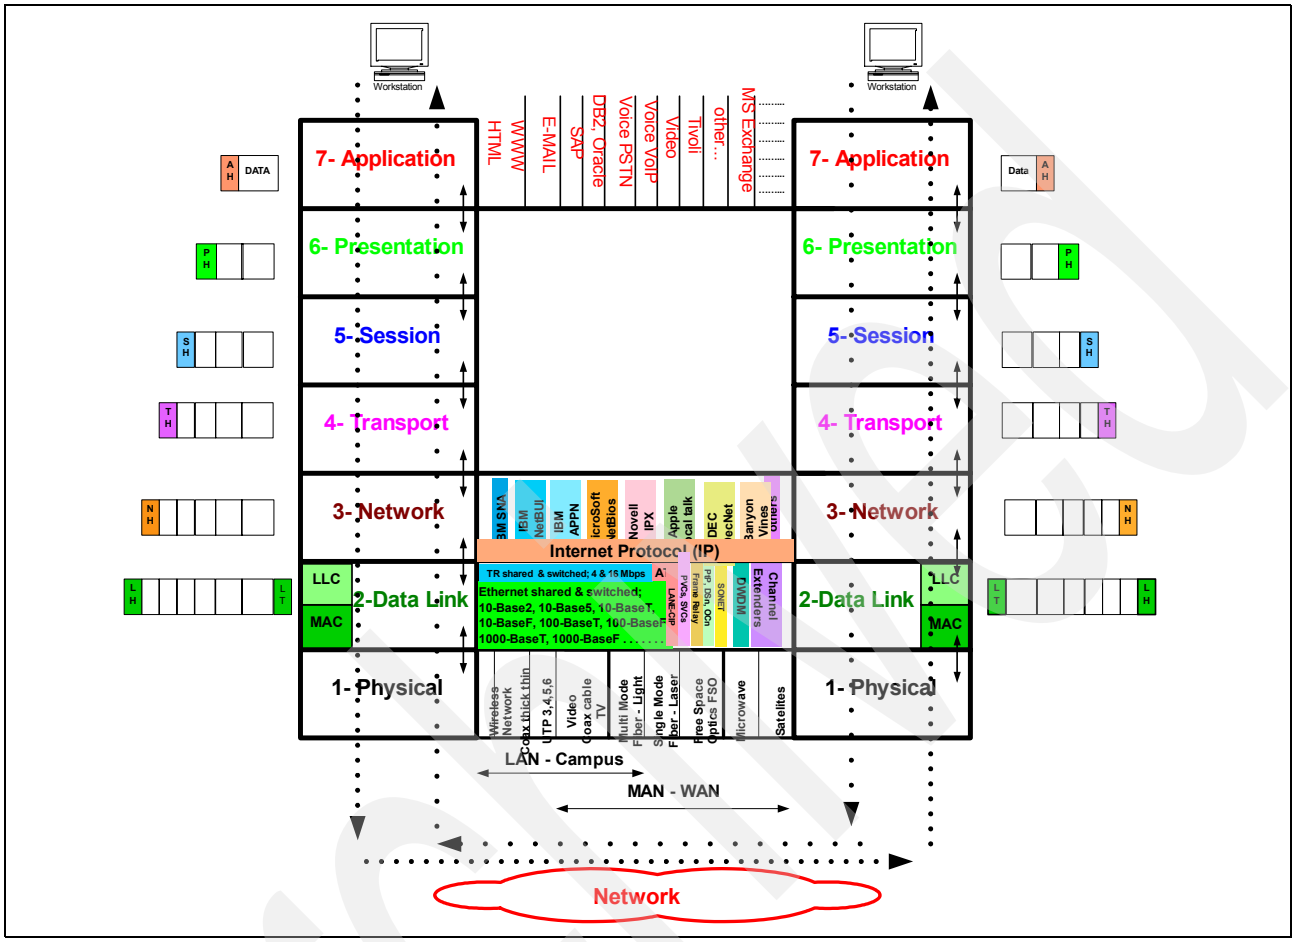

Figure D-5 OSI layers 4, 5, 6, 7 - Transport, session, presentation, and application

The remainder of this chapter focuses on what happens at layers 1, 2, and 3. You can safely assume that if you can connect at the OSI layer 3 level, then from there the software can provide the additional logic to transverse the other layers.

OSI layer 4 (transport layer general comment)

Transmission Control Protocol (the TCP portion of TCP/IP) is in the OSI layer 4 transport layer. In this layer, TCP/IP data packets undergo:

- Flow and congestion control
- Reassembly of received -out-of-order packets
- CRC error checking

Interfacing different networks together

With the OSI model now understood, we turn next to the need and methods of *interfacing* between two networks. There are four basic types of network interface devices (Figure D-6):

- Repeaters
- Hubs / Bridges
- Switches
- Routers

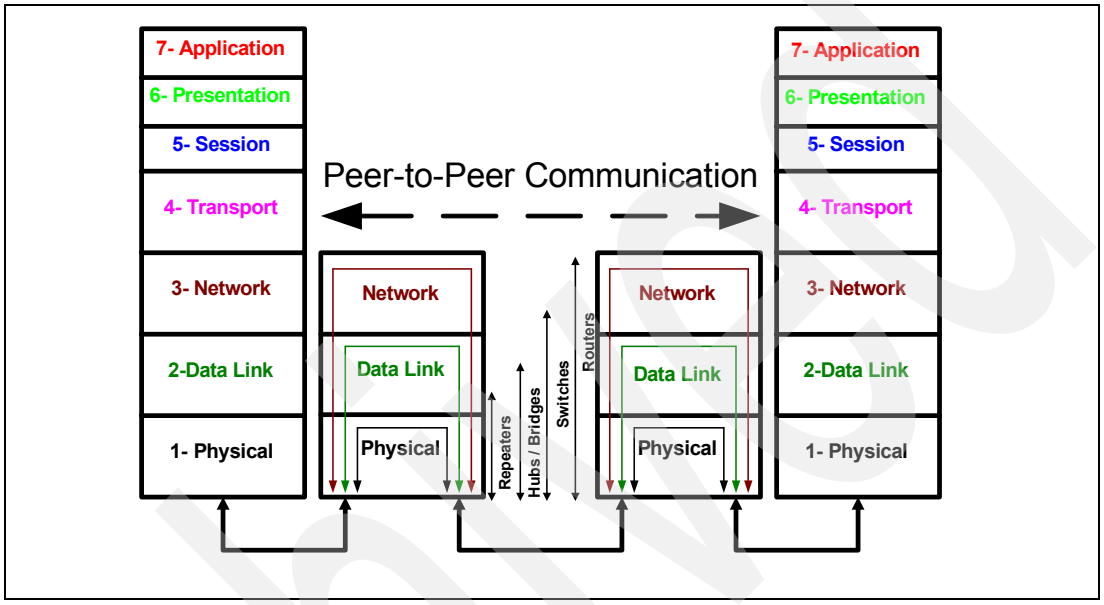

Figure D-6 Interfacing networks using repeaters, hubs/bridges, switches, and routers

As you can see from the vertical arrows, the determinant of which device to be used depends on what layer of the OSI model is to be interfaced.

- *Repeaters* are used when highest level of network interconnection is fundamentally at the OSI layer 1 physical connection level.
- *Hubs and Bridges* are used when highest level of network interconnection is fundamentally at the OSI layer 2 data link level.
- *Switches* are used when highest level of network interconnection is fundamentally at the more basic levels of the OSI layer 3 network level.
- *Routers* are used when highest level of network interconnection is fundamentally deep within the OSI layer 3 network level.

Obviously, there is no hard and fast line between these four types of interconnection devices. Vendors all have varying levels of technology, and that technology can and does span layer interconnections.

In Figure D-7 is another example of where repeaters, hubs and bridges, switches, and routers would typically be used.

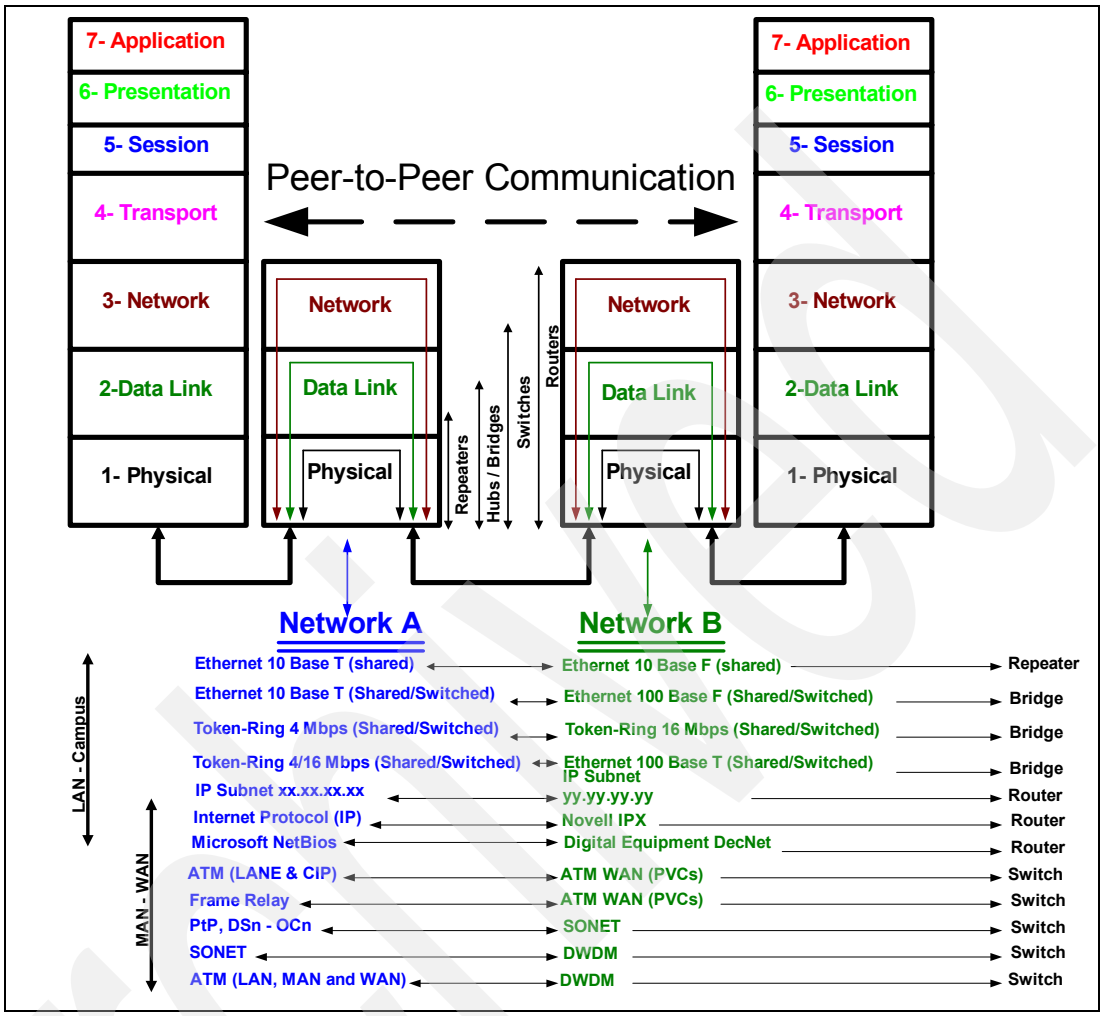

Figure D-7 Another example of where repeaters, hubs/bridges, switches, and routers would be used

By understanding the OSI model one is able to ask intelligent questions of the vendor's implementation in order to determine exactly what layer, and to what level of interfacing that vendor product is doing.

As a consequence, the following terms are unfortunately used somewhat loosely:

- Layer 2 Switch
- Layer 3 Switch
- Layer n Switch
- Gateways
- Bridging and Routing
- Hub / Switch
- Shared / Switched
- Full and Half Duplex
- Multiplexors
- Extenders
- Directors
- Internetworking
- Interworking

Network interfacing is very flexible. It is basically possible to connect almost any network to almost any other network. The variables that determine if this is possible or not are money, funding, and cost justification; what level of interface gear is available, how much does it cost, and is that cost justifiable?

Fiber optic cables - used in OSI layer 1

At the OSI layer 1, there are multiple common different types of physical layers, including UTP (unshielded twisted pair), coax, and so forth. Let us investigate a strategically important type of physical connection: the fiber optic cable.

Fiber optic cable is at the OSI layer 1 physical layer. The data link layer, network layer, and transport layers are on top of the physical layer. Therefore, multiple different types of protocols in these other layers can traverse a physical layer fiber optic cable at the same time.

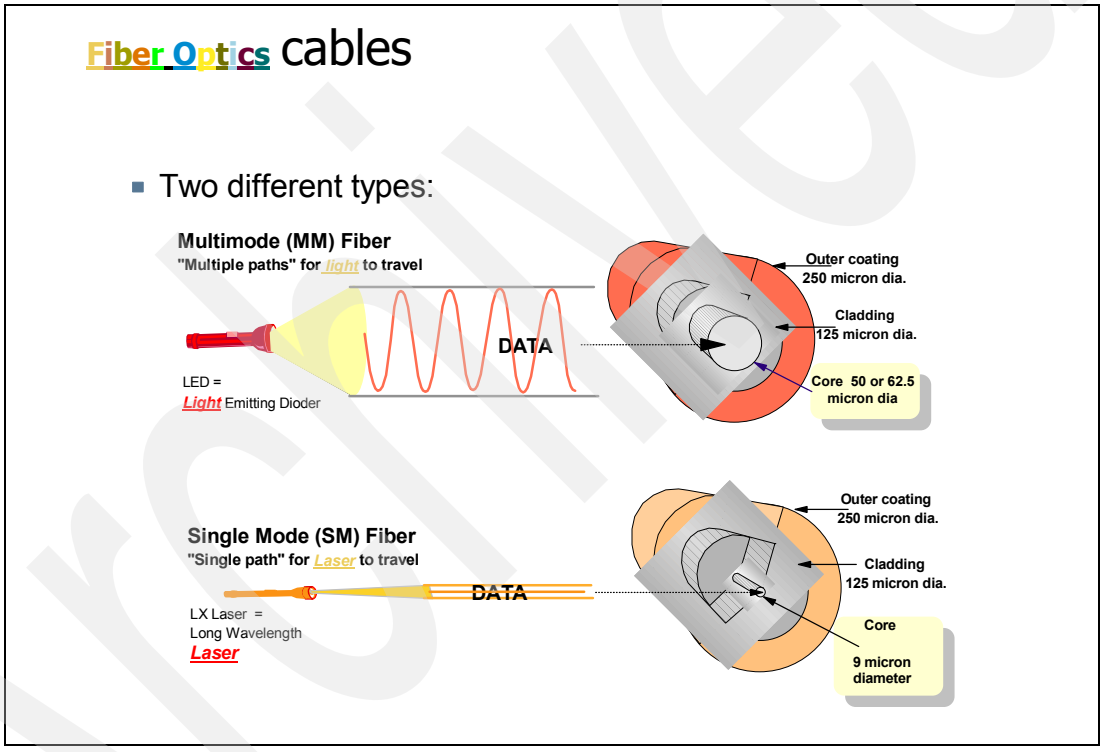

Figure D-8 Fiber optic cables - single-mode and multi-mode

There are two major types of fiber optic cable:

- Multi-mode fiber: 50 micron, typically used within the data center or short haul distances (on Fibre Channel SANs or Gbit Ethernet, typical max distance is about 500 meters). Uses shortwave laser, which can transmit 300m at 2Gbps.
- Single-mode fiber: 9 micron, more expensive per foot/installation than multi-mode, and typically used for longer distance transmissions. Uses longwave laser, which can transmit 10km at 2Gbps.

Note that *two* strands per fiber optic link will be required; one to transmit in each direction.

The strengths of using fiber optic cable

For large capacity data transmissions, from a strategic standpoint, fiber optic cable is the most attractive type of OSI layer 1 physical connection, as shown in Figure D-9.

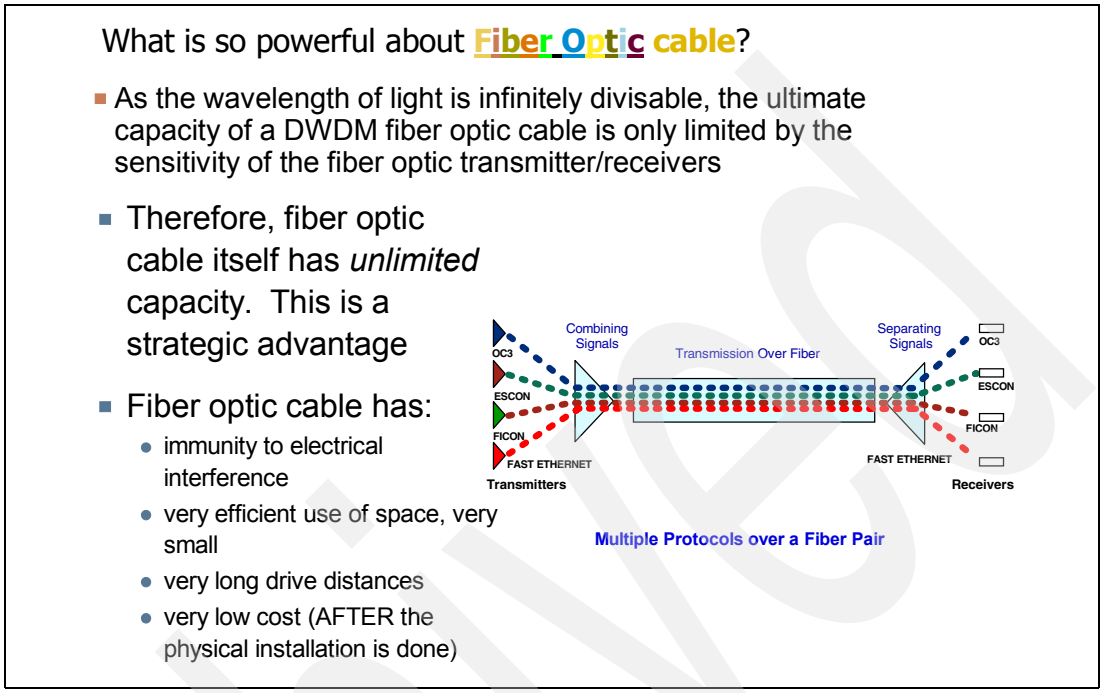

Figure D-9 What is powerful and strategic about fiber optic cable

Fiber optic cable offers strategically infinite bandwidth and performance, and has a very low cost of maintenance (after it has been installed).

As described earlier in 10.3, “Wavelength Division Multiplexing” on page 287, Coarse Wave Division Multiplexing (WDM) and Dense Wave Division Multiplexing (DWDM) further exploit fiber optic technology to:

- Assign a input fiber optic channel to a particular wavelength of light (think of it as a 'color')
- Transmit those 'colors' to the other end
- Separate the 'colors' into each individual output channel

Because the wavelength of light is infinitely divisible (gated only by the granularity and sensitivity of the transmitters/receivers), after the fiber optic cable is installed, there is unlimited bandwidth available over time using that same cable. No other known current physical layer media has this powerful potential to the same degree.

There are other advantages which make dark fiber the fundamental physical layer networking media of choice in today's world and for a very long time to come. They are:

- Immunity to electrical interference
- Very efficient use of space, very small
- Very long drive distances
- Very low cost (AFTER the physical installation is done)

For a very complete dissertation on optical fiber, see the IBM Redbook *Understanding Optical Communications*, SG24-5230, which is available online at:

`http://www.redbooks.ibm.com/redbooks/pdfs/sg245230.pdf`

Even though this book was published several years ago, the information is still relevant.

On the long distance interconnect, the telecom companies have been building a fiber optic infrastructure for over 20 years. They have their own internal telecom technologies (SONET is one example) to capture, encapsulate, transmit, problem-determine, configure, and manage their environments.

Other general comments about networking

We close with a few final comments.

The last mile issue

One of today's most important issues with dark fiber is whether the dark fiber or network access is actually available at a specific customer location. To be ultimately usable to each specific customer, the telecom provider must have a cost-effective means of getting that fiber or network, *the last mile* to the specific location that the customer wants.

This is not a trivial task, and should be investigated early in the cost estimation process. The last mile issue is one of the critical elements in a cost-justifiable network design.

Basic network design concepts

Just as in Business Continuity, the basic concept is to divide the enterprise into tiers. At the top of the hierarchy tier are the major central data centers which act as the major hubs. There is a next lower tier of data/regional centers, followed by the local sites.

Dark fiber strand pricing and configuration

You can expect dark fiber strand pricing to be quoted in dollars or euros per strand per mile per month. While costs vary, and clearly are decreasing depending on the geography and competition, the cost for dark fiber strands is still usually substantial. Therefore, to avoid unpleasant sticker shock for dark fiber infrastructure, bids to determine your cost should be submitted very early in the infrastructure design process.

Be aware that you need two dark fiber strands (one for transmit, the other for receive).

In many circumstances there will be a need for at least two geographically separate paths to the other site, so there is another factor causing the cost to effectively double again.

Although pricing for dark fiber varies by geography, and the cost is clearly decreasing over time, you should still expect that the pricing of dark fiber strands is not small. As an example in 2002, prices per dark fiber in the U.S.A could cost about US$300 to $400 per strand per mile per month.

Telecom providers are not the only organizations that have their own dark fiber. Utilities and transportation companies (because they already have right of way) often already have their own dark fiber infrastructures.

Summary

Networking is a sophisticated technology with its own terminology, architecture, and products. In today's world requires more and more often, that we know something about other technologies in order to evaluate, select, and maintain our complex IT infrastructures.

Related publications

We consider the publications that we list in this section particularly suitable for a more detailed discussion of the topics that we cover in this IBM Redbook.

IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 396. Note that some of these documents that we reference here might be available in softcopy only.

General

- *IBM System Storage Solutions Handbook*, SG24-5250
- *IBM System Storage Business Continuity: Part 2 Solutions Guide, SG24-6548*

Software

- *IBM System Storage SAN Volume Controller*, SG24-6423
- *Disaster Recovery Using HAGEO and GeoRM*, SG24-2018

Disk

- *IBM System Storage DS8000 Series: Architecture and Implementation*, SG24-6786
- *IBM System Storage DS6000 Series: Copy Services with IBM System z*, SG24-6782
- *IBM System Storage DS6000 Series: Architecture and Implementation*, SG24-6781
- *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
- *IBM TotalStorage Enterprise Storage Server Model 800*, SG24-6424
- *DS4000 Best Practices and Performance Tuning Guide, SG24-6363*
- *IBM System Storage DS4000 Series and Storage Manager*, SG24-7010
- *Implementing Linux with IBM Disk Storage, SG24-6261*
- *The IBM System Storage N Series*, SG24-7129

Tape

- *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632
- *Implementing IBM Tape in UNIX Systems*, SG24-6502
- *IBM Tape Solutions for Storage Area Networks and FICON*, SG24-5474
- *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229
- *IBM TotalStorage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide*, SG24-6115
- *IBM System Storage Tape Library Guide for Open Systems*, SG24-5946

SAN

- *Introduction to Storage Area Networks, SG24-5470*
- *Designing an IBM Storage Area Network*, SG24-5758
- *IBM System Storage: Implementing an 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 McDATA Portfolio*, SG24-6149
- *IBM SAN Survival Guide Featuring the IBM 3534 and 2109*, SG24-6127

Tivoli

- *IBM Tivoli Storage Manager Implementation Guide*, SG24-5416
- *Tivoli Storage Manager Version 5.1 Technical Guide*, SG24-6554
- *IBM Tivoli Workload Scheduler Version 8.2: New Features and Best Practices*, SG24-6628
- *Disaster Recovery Strategies with Tivoli Storage Management*, SG24-6844
- *IBM Tivoli Storage Management Concepts*, SG24-4877
- *Deploying the Tivoli Storage Manager Client in a Windows 2000 Environment*, SG24-6141
- *IBM Tivoli Storage Manager: Bare Machine Recovery for AIX with SYSBACK, REDP-3705*
- *IBM Tivoli Storage Manager Version 5.3 Technical Guide*, SG24-6638
- *Get More Out of Your SAN with IBM Tivoli Storage Manager*, SG24-6687
- *Using IBM Tivoli Storage Manager to Back Up Microsoft Exchange with VSS*, SG24-7373
- *IBM Tivoli Storage Manager for Advanced Copy Services*, SG24-7474

Online resources

These Web sites and URLs are also relevant as further information sources:

- GDPS:
  - http://www.ibm.com/servers/eserver/zseries/gdps/
  - http://www.ibm.com/services/us/index.wss/so/its/a1000189
  - http://www.ibm.com/common/ssi/fcgi-bin/ssialias?infotype=an&subtype=ca&appname=Demonstration&htmlfid=897/ENUS205-035
  - http://www.ibm.com/common/ssi/fcgi-bin/ssialias?infotype=an&subtype=ca&appname=Demonstration&htmlfid=897/ENUS305-015
  - http://www.ibm.com/servers/storage/disk/
- GDOC:

  http://www-1.ibm.com/servers/eserver/pseries/ha/disaster_tech.html
- SAN Volume Controller

  http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002442#_Inter_Cluster
- IBM TotalStorage DS8000 interoperability matrix:

  http://www-1.ibm.com/servers/storage/disk/ds8000/interop.html
- IBM TotalStorage DS6000 interoperability matrix:

  http://www-1.ibm.com/servers/storage/disk/ds6000/interop.html

- High availability:
  - 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
- IBM services:

  http://www.ibm.com/services/storage
- TDM:

  http://www-1.ibm.com/servers/storage/services/featured/microsoft_application_environment.html
- Tivoli:
  - http://www.ibm.com/software/tivoli/products/storage-mgr/
  - http://www.ibm.com/software/tivoli/solutions/disaster/
  - http://www.ibm.com/software/tivoli/products/storage-mgr-sysback/
  - http://www.ibm.com/software/tivoli/products/storage-mgr-erp/
  - http://www.ibm.com/software/tivoli/products/storage-mgr-hardware/
  - http://www.ibm.com/services/storage
  - http://www.ibm.com/software/tivoli/products/storage-mgr-mail/
  - http://www.ibm.com/software/tivoli/products/storage-mgr-db/
- IBM servers:
  - http://www.ibm.com/servers/eserver/pseries/
  - http://www.ibm.com/servers/eserver/xseries/
- IBM tape:

  http://www.ibm.com/servers/storage/tape/3583/
- SAP:
  - http://www-1.ibm.com/servers/storage/solutions/sap/index.html
  - http://www50.sap.com/softwarepartnerdir/
- GDPS:
  - http://www.ibm.com/servers/eserver/zseries/zos/news.html
  - http://www.ibm.com/servers/eserver/zseries/gdps/
- IBM System i

  http://www.iSeries.ibm.com/ha
- IBM & Cisco Alliance

  http://www-1.ibm.com/services/alliances/cisco/
- Brocade

  http://www.brocade.com
- Disaster Recovery Institute Canada

  http://www.dri.ca
- Disaster Recovery Institute International

  http://www.drii.org
- Disaster Recovery Journal

  http://www.drj.com
- McData

  http://www.mcdata.com

- Microsoft
  - http://www.microsoft.com
  - http://www.microsoft.com/sql/techinfo/planning/default.asp
- Microsoft VDS

  http://www.microsoft.com/windowsserversystem/storage/technologies/shadowcopy/stormgtusingvdsvss.msp
- VERITAS:

  http://www.veritas.com

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications, and Additional materials, as well as order hardcopy 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

D

E

F

G

H

I

L

M

N

O

P

R

S

T

IBM System Storage Business Continuity: Part 1 Planning Guide

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

IBM System Storage Business Continuity: Part 1 Planning Guide

Describes current trends and strategies for Business Continuity

Explains how to select an appropriate Business Continuity solution

Presents a step-by-step Business Continuity planning workshop

A disruption to your critical business processes could leave the entire business exposed. Today's organizations face ever-escalating customer demands and expectations. There is no room for downtime. You need to provide your customers with continuous service because your customers have a lot of choices. Your competitors are standing ready to take your place. As you work hard to grow your business, you face the challenge of keeping your business running without a glitch. To remain competitive, you need a resilient IT infrastructure.

This IBM Redbook introduces the importance of Business Continuity in today's IT environments. It provides a comprehensive guide to planning for IT Business Continuity and can help you design and select an IT Business Continuity solution that is right for your business environment.

We discuss the concepts, procedures, and solution selection for Business Continuity in detail, including the essential set of IT Business Continuity requirements that you need to identify a solution. We also present a rigorous Business Continuity Solution Selection Methodology that includes a sample Business Continuity workshop with step-by-step instructions in defining requirements. This IBM Redbook is meant as a central resource book for IT Business Continuity planning and design and is intended for anyone who wants to learn about Business Continuity trends and strategies.

The companion title to this IBM Redbook, *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548, describes detailed product solutions in the System Storage Resiliency Portfolio.

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

SG24-6547-03 ISBN 0738489700