

Automating the Business of Business Recovery

A White Paper Discussion

May 2007



CONTENTS

EXECUTIVE SUMMARY 3

INTRODUCTION..... 4

DISASTER RECOVERY – WHAT DOES IT REALLY MEAN? 4

THE ROAD TO RECOVERY 4

IDENTIFICATION OF CRITICAL DATA 5

REPORT, MANAGE AND BACKUP 6

MANAGING DATA THROUGH FILTERS..... 7

MANAGING BACKUPS THROUGH AUTOMATION..... 7

RECOVERY 8

AUTOMATING THE RECOVERY PROCESS..... 8

BUILD CONFIDENCE..... 8

IN SUMMARY 9

EXECUTIVE SUMMARY

The concept of business recovery is nearly as old as mainframe servers, if not as old. Growth in applications, data center consolidations, business combinations, 24 by 7 operations, and many other dynamics have made the manual process of managing Business Recovery an almost humanly impossible task.

Over the years, new hardware and software technologies have come, gone and been improved. Yet, there has been little improvement regarding the succinctness of the business recovery process.

“Wait a minute,” you say, “There’s mirroring and replication of DASD!”

True, there are those technologies. But mirroring of DASD is only part of the business recovery process. In addition to the cost of the duplicated DASD, there are substantial telecommunications costs, and yet there most likely will still be critical tape datasets missing preventing a complete and successful business recovery.

Many enterprises have opted for an all-or-nothing approach to business recovery, i.e. mirroring everything regardless of cost, while others have opted for a subset of mirroring with replication managing the other datasets. In each case, the first and most critical function of business recovery, that of data identification, has been performed at a very high level or not at all, resulting in the replication of too much data (substantial meaningless data) or missing even one critical file necessary to continue the business processes at an alternate data center for the mission critical business application.

There are requirements for this level of recoverability as a result of governmental regulations throughout the world and the compliance guidelines for reporting and managing certain types of data. Out of self defense, many users have created a plethora of manually managed procedures and processes to restore the operations based on a recovery time objective and on a recovery point objective.

However, this does not address the basic issue – the process of business recovery and its subsets, such as data identification, have not kept up with the times. To this day, the bulk of the business recovery process requires a great deal of human resources, in the form of people, time and effort, which can be prone to error, and is costly on a single test basis (tests delayed for hours based on restarting) and on a recurring daily operational basis (excess offsite costs and capacity degradation).

In this document we discuss how OpenTech Systems’ business recovery solution, DR/Xpert, will simplify and reduce the human resource constraints normally related to this process.

INTRODUCTION

Ensuring business recovery has become a difficult and sometimes thankless task despite regulations and local catastrophes. Obtaining tools that aid in ensuring business recovery is one of the most important items with any strategy within this space. Knowing the criticality of your business applications and their associated data can help you make important decisions as to how the recovery strategy is to be implemented from a hardware and software perspective.

We will discuss specific areas necessary to ensuring business recovery and how much of these tasks can be automated using OpenTech Systems' solutions.

DISASTER RECOVERY – WHAT DOES IT REALLY MEAN?

Everyone talks about disaster recovery; they also talk about Business Continuity, Business Continuance or Business Recovery. They use the terms interchangeably. It gets confusing... To keep things simple, this document will work from the following definitions¹:

- **DISASTER RECOVERY (DR)** – Methods and procedures for returning a system, network or data center to full operation after a catastrophic interruption — including the recovery of lost data, and the use of alternative network channels if the primary channels are disconnected or malfunctioning.
- **RECOVERY POINT OBJECTIVE (RPO)** – A term used in disaster recovery and business continuity planning. The RPO defines what constitutes an acceptable loss of data — specifically, the required timeliness of the data that can be recovered using backups, journals or transaction logs.
- **RECOVERY TIME OBJECTIVE (RTO)** – Often called the "recovery window," the RTO defines how quickly information systems, services and processes must be made operational for disaster recovery purposes.

Disaster recovery is a subset of a business continuity plan; the terms continuity, continuance and resiliency, have similar connotations and will be treated as meaning the same thing throughout this document.

THE ROAD TO RECOVERY

The primary objective of disaster recovery is to have the capability to provide a process in which files, applications or operating systems can be recreated in a timely fashion, locally or at a remote location.

Presumably, what works locally should work remotely. However, the scale or scope of recovery tends to be very different when discussing local versus remote recovery. Most local

¹ "The Gartner Glossary of Information Technology Acronyms and Terms" Source: Gartner, Inc. 2004

recovery initiatives tend to involve a small group of data sets within the local operating complex. Many of the recovery initiatives occur locally, mostly due to human error. This could be as simple as accidental deletion of one or more files. Remote recovery implies that the local environment is compromised or no longer available for any number of reasons – due to some type of local or regional event. Regardless, the computing environment must be recreated (or recovered) to another location. The question then becomes, what should be recovered and when?

In order to recover operating systems, mission critical business application(s) or data set(s), it is necessary to know what data is critical to the successful execution of that particular entity, identify any interdependencies, and recognize the appropriate Recovery Point Objective. Performing this task manually is not a viable option; business applications have grown to such complexity that it would take an inordinate amount of time and resources to accomplish the task and is guaranteed that something would be missed especially with data sharing across applications. It has been stated over the years that businesses backup too much data and businesses will miss a critical file. The only way to ensure that all critical data sets are identified is automation via software, which is where OpenTech Systems DR/Xpert, an automated recovery solution comes into the picture.

IDENTIFICATION OF CRITICAL DATA

The process of identifying what data is critical to one or more applications is not all that easy. More important is how the data is used: is it input? is it output? does it feed into other jobs or applications? is it a master file that is rarely updated? - and so on.

Why is the disposition of such data important? The disposition of a data set has a direct impact on whether or not the data is backed up. Dependent upon use, it may be that a data set must be backed up or only that JCL is available to allocate the file. Understanding the way data is used is as important as the identification.

In order to ensure that all data is classified as to being critical or non-critical, multiple sources of information must be used for analysis:

- **SYSTEM MANAGEMENT FACILITY (SMF)** – this facility collects and records a variety of information. For the purposes of business recovery, information related to job execution and data set disposition (created, updated, renamed or deleted), is extracted. Historical and current information is used so that annual, tri-annual, quarterly, monthly, weekly and daily jobs and data sets are identified. A great majority of the information is found through SMF; however, additional sources of information are needed in order to get the complete details of all application job flows.
- **JOB CONTROL LANGUAGE (JCL) LIBRARIES** – known as JCLLIBs or PROCLIBs, are scanned for more information. The information found in these libraries is used to identify data sets that may be listed but not opened, or steps that are only executed based on condition codes.

- **JOB SCHEDULING SYSTEMS** – contain information related to job names as well as when jobs are executed.

Once this information is obtained, DR/Xpert automatically analyzes this data to determine how the data is used and its disposition. This includes the following types of data sets:

- Need to exist (allocate only) in order to ensure successful application execution
- Dynamically allocated – internal to the executing program
- Pattern data sets – including, but not limited to, Generation Data Groups (GDGs)
- Concatenated data sets – not just the first data set, but all of them
- Read only data sets – how often does this type of file need to be backed up?
- Online versus batch data sets – determines when data sets are available for backup
- Migrated data sets – managing backups before migration
- Cross application intersections – when more than one application requires a data set.
- Knowing that such data sets need to be backed up is not sufficient. The key is to know when the data sets can be backed up to ensure consistency.

DR/Xpert aids in the understanding of these data set types. In addition, DR/Xpert will tell you why a data set (Tape or DASD) is deemed critical or not – at the touch of a button. Because DR/Xpert brings this understanding to your data, you now have the ability to determine why this data is to be backed up.

REPORT, MANAGE AND BACKUP

Implementing a DR solution takes time and effort. It's not necessarily a straight-forward process – applications are intertwined with one another, and data sets are accessed by more than one application. Getting it all straight is impossible without automation. Without reports there is no way to discern what data has been collected, and what to do with it.

The reports provided by DR/Xpert helps your staff proactively manage -

- Exception reporting that focuses immediate attention on problem issues, rather than wading through scores of reports showing what was successful
- What data should be backed up and when it should be backed up
- How many previous copies, or cycles, of data are kept
- Forecasts related to capacity requirements for business recovery at a remote site
- What data is stored in a remote vault

Plus, these reports also aid in achieving your compliance goals and initiatives by providing information related to backups, encryption, retention periods and vault management.

MANAGING DATA THROUGH FILTERS

DR/Xpert provides powerful filters to tailor what information should be included or excluded. These filters provide another benefit – they enable you to manage your data in other ways such as:

- Identification of application start and end jobs. Most DR solutions require this information to be researched and provided by you, the customer. DR/Xpert will perform this task automatically – all that is required is a basic job mask, DR/Xpert will automatically handle the rest.
- Support and management of tiered storage initiatives. DR/Xpert shows you what data is mirrored or not mirrored on disk. It also identifies critical data on virtual or physical tape. This allows you to ensure that those applications with short RPO/RTO objectives have their data sets resident on the volumes that are mirrored or allows you to redirect tape data to DASD to insure all data needed for recovery is mirrored and available off-site.
- Selective encryption of data. Analysts, such as Gartner and IDC, recommend that any tape media going off-site be encrypted in case of loss or theft.
- Consolidation of read only files. DR/Xpert automatically identifies and tracks these types of files so that they are only backed up when they have been updated. Over time the utilization of a tape may become inefficient, at which time DR/Xpert can be used to reduce the recovery volume set by automatically consolidating these files.

MANAGING BACKUPS THROUGH AUTOMATION

DR/Xpert uses the gathered data to complete the next phase of the business recovery process. The core design of DR/Xpert ensures that your critical data is being backed up –automatically. This includes the automatic generation of JCL – for both the backup and recovery process, with or without encryption, thus removing the possibility of error via human intervention. DR/Xpert validates backup processing and will dynamically retry the backup in the event of a problem.

Not all critical data is resident on DASD. In some cases, it is resident on virtual or physical tape. This data cannot be ignored by the identification and backup processes. DR/Xpert can automatically include tape data.

Because of these automated processes, your staff no longer has to devote large amounts of time to ensure that the backup and recovery JCL is correct and error-free, nor do they have to manage the updating of such JCL when there are changes. DR/Xpert monitors changes so that new procedures, jobs or data sets are automatically detected and if necessary, adjusted when backups occur, without intervention.

DR/Xpert tracks all business recovery backups and dynamically maintains its own database and backup history. By providing this functionality, your staff no longer has to take the time to monitor this process.

RECOVERY

This is what the “data identification and backup processes” are all about. It’s all well and good to identify and backup the data necessary to continue the business operations, but if no one is available that understands the order of recovery or how to accomplish the recovery, the resources, time and effort are for naught!

The goal of all recovery methodologies should be to backup only what is needed to ensure successful business operations. Unfortunately, that is not always the case. These same methodologies should also provide you with the ability to take advantage of physical media investments, such as high-density tapes.

Asked what should be the first thing to recover, most would say either “Everything” or “Our online systems, because that’s what makes the money for the business.” Both answers are correct however, one thing has been left out. The operating system(s) is what really needs to be recovered first. While you might say “that’s obvious”, it is important to remember if only for recovery planning.

AUTOMATING THE RECOVERY PROCESS

As part of the recovery process, it is important to remember to include your recovery software. The information used to identify and backup your critical data is contained within the software databases and is necessary for recovery. Failure to include this important recovery investment may hamper your recovery process.

DR/Xpert, once recovered, provides automated recovery of business applications and data sets by executing the recovery job streams created during the backup process. If there are differences at the recovery site, the product has been designed to allow for changes onsite with automatic regeneration of recovery JCL.

As part of this process, tracking of recovery jobs occurs so that the databases are automatically updated. Once this process is completed, DR/Xpert continues to automatically manage the identification and backups at the recovery site. This ensures that all production applications can be automatically recovered when the home site is available.

BUILD CONFIDENCE

Because DR/Xpert is a dynamic and automated solution, it allows you to not only protect your critical data, but to implement a repeatable process that ensures business recovery.

By automating your recovery process, you reap the following benefits:

- Centralized identification, backup and recovery management
- A consistent accurate, predictable, and repeatable process
- Reduction of previously manually intensive processes through automation

IN SUMMARY

DR is a process, that once started is never ending. It does require 'flexibility' to support the dynamics of today's changing applications - it cannot be ignored. However, it does not have to be resource intensive or prone to human error when other automated solutions are available.

DR/Xpert is an automated solution that frees you and your staff from the drudgery of using a product or home grown process – and does not degrade the performance or reliability of the DR process.

DR/Xpert is the next generation business recovery solution that -

- Performs dynamic discovery of critical data files and applications
 - Automatically works with scheduling packages to ensure backups are scheduled throughout the day thereby eliminating the “backup window” and allowing backups to occur “anytime”
 - Reports on exceptions (as well as things that worked) so that your staff can focus on problem issues and not on what worked correctly.
 - Works with standard, third-party, job, tape and disk management systems
 - Is completely self-contained and does not require additional third party software
 - Aids in compliance by reporting that critical data is protected
 - Reduce costs through the management of resources such as data with high-availability requirements and high-density tapes
-

About OpenTech Systems, Inc.

OpenTech Systems, Inc. specializes in Disaster Recovery, Data Availability, and Data Security solutions for IBM mainframe-centric data centers. OpenTech Systems is the vendor of choice for many leading storage suppliers, IT outsourcers, and Fortune 1000 companies focused on ensuring business continuity, increasing operating efficiency and improving data security.

For further information, visit www.opentechsystems.com.

North America – Worldwide
Headquarters
OpenTech Systems, Inc.
405 State Hwy, 121 Bypass
Building C, Suite 130
Lewisville, TX 75067-4182
+1 469 635 1500
+1 469 635 1507 (fax)

Europe
OpenTech Systems, Ltd.
3000 Hillswood Drive
Hillswood Business Park
Chertsey, Surrey, KT16 0RS, UK
+44 (0) 1932 895 205
+44 (0) 1932 895 501 (fax)

Germany
OpenTech Systems, Ltd
Deutschland
Lyoner Strasse 26
60528 Frankfurt am Main
+49 69 6 77 33 415
+49 69 6 77 33 200 (fax)