

# LINKING DISASTER RECOVERY TIME OBJECTIVES TO BUSINESS AND COMPLIANCE REQUIREMENTS

White Paper  
August 2006

## **Abstract**

Too often disaster recovery objectives focus on systems rather than business needs. This paper provides guidelines for basing recovery time objectives and recovery point objectives on real-world business processes.

## Table of Contents

Linking disaster recovery time objectives to business requirements	
Executive summary .....	1
Introduction .....	2
Preparing for the worst .....	3
Using business objectives to define recovery objectives .....	4
Assigning RTO/RPO tiers .....	5
Finding RTOs and RPOs .....	6
Conclusion .....	7
Contact Information .....	7

## Chapter 1

# Executive summary

IT disaster recovery managers tend to base recovery time objectives (RTOs) and recovery point objectives (RPOs) on system requirements rather than business needs. However, a process is needed for linking RTOs and RPOs to critical business processes.

In practice, this linkage is the responsibility of the business owner. It begins with the assignment of tiers of criticality to each business process on which recovery processes are then based. Only the business owner can assess the risk that losing key systems poses to certain business processes and the overall health of the organization. If compliance issues such as HIPAA, Sarbanes-Oxley, GLB, COOP, or Basel II are a factor which impacts business processes within that organization, the risks become even more complicated and intricate with application and business process interdependency.

Balancing risk with the costs of various recovery strategies is an iterative process. Business owners must coordinate closely with business continuity managers, senior level management, and legal and risk departments to establish the key factors of Security, Storage, Data Archiving and Retrieval, and Restoration of Systems to meet the Enterprise RTO/RPOs. During cost/risk discussions, there are questions that need to be asked to help ensure the linkage of recovery to business needs. This paper introduces those questions as well as provides a sample decision input matrix to help ensure your recovery processes are based on real-world business drivers.

RTO is the time to recover the system after a disaster, while RPO reflects the time allowed to elapse since the last backup prior to a disaster.

Don't fall into the trap of thinking that business processes are secure just because your organization has system recovery time objectives

## Chapter 2

# Introduction

If you are an IT disaster recovery (DR) manager, you most likely have DR plans in place that provide for recovery of critical business applications and the systems to support those applications. Consider the following questions to examine if your current plans meet business and compliance requirements:

- What processes and criteria did you use to determine the recovery time objectives and recovery point objectives for those systems?
- Did you look at the recovery of the critical business processes or focus exclusively on the recovery of the systems?
- Do you have a process in place to ensure the synchronization of the recovery objectives for compliance and business requirements?
- How often do you confer with your company's business continuity manager? Do you even know if you have a business continuity manager?

In today's business environment, a business application system is entwined with the business process it supports (especially if dependent on an e-business component). This should also hold true for the recovery plans of those systems and processes. This level of interactivity compounds if there is a level of compliance involved.

## Chapter 3

# Preparing for the worst

When (not “if”) a disaster strikes your company, critical, time-sensitive business processes must continue to ensure your company functions reliably. For example, California companies susceptible to earthquakes must assume worst-case scenarios in which people are injured, freeways are impassable, and phone service is hindered. Organizations in other geographical locations plan for hurricanes, tornadoes, floods, and so on. And while you may have operations remote from a damaged location that will continue to function, what business processes are conducted locally upon which those other sites depend?

Business recovery time objectives are usually determined by a business impact analysis (BIA). On the other hand, system recovery time objectives are defined by the method used and time required to recover the system. Business and system recovery are very different and each deserves special attention. When was the last time you analyzed and compared these two for any gaps? Does your company consider RTOs and RPOs when designing or implementing a new business systems application? Do you investigate the impact to RTOs and RPOs when proposing a major change?

Many companies have system development and/or project management life cycles. Where in that process or methodology do you find your recovery strategy being applied? Best practices require that the BIA and application criticality be defined in the initial assessment, prior to approval. At the next phase, recovery strategy decisions impact the hardware design and dictate data backup solutions. As you know, high-availability solutions are more costly than those addressing a recovery time objective of several weeks. What risks are you willing to incur, and at what cost? Such decisions are based on the criticality of the application and the business process it supports. Your planning needs to be both flexible and testable. For example, if an application already has a disaster recovery plan, changes will probably need to be made based on your analysis. And of course, if the application is new, the disaster recovery plan needs to be created prior to moving to production. In either case, you need to test the DR plan prior to deploying into production.

## Using business objectives to define recovery objectives

A disaster is an interruption of a system for an unacceptable period of time.

A complex business process may utilize multiple computer systems and platforms, and each business process may have a different unacceptable period of downtime. The focus of any recovery plan must be on keeping the business running — not keeping the computers running. Are recovery objectives in sync between the business process and the multiple systems and platforms that support the process?

This reinforces the idea that business continuity requirements should be defined before addressing computer requirements. It's the business owner who needs to determine what is an acceptable risk for the business and what RTO is appropriate, and to agree to the resulting cost of the systems implementation to meet that determination. And if new computer technologies being implemented (such as server, network, and OS upgrades) are not related to a current business application, the business processes that will eventually be dependent on a project's deliverables will determine the recovery requirements. It's helpful to have a decision matrix that describes who has responsibility for functions such as approval, review, and providing critical information upon which recovery decisions will be made. Table 1 shows an example of a decision matrix that captures the role of each person involved in the process. Major players should be identified as information providers, reviewers, or owners.

*Table 1. A sample decision matrix for assigning responsibility during RTO/RPO development.*

Person	Requested RTO/RPO	Cost of Technical Solution	Person
Business Sponsor, responsible for business function	Information provider	Reviewer	Owner, final approval and signature
Technical Architect	Information provider	Owner	(No role associated with final RTO)
Disaster Recovery Manager	Reviewer	Reviewer	Review of plans
Business Continuity Manager	Reviewer	Reviewer	Review of plans
Application Support Rep	Reviewer	Reviewer	Review of plans
Computer Operations Rep	Reviewer	Reviewer	Owner of recovery plans

## Assigning RTO/RPO tiers

RTOs are usually tiered by criticality. You'll need to look at your company's unique requirements as to how many tiers are appropriate for your organization — more than five generally becomes unmanageable. Examples of five RTO tiers might be:

- Tier 1 – Fault tolerant with virtually no impact to the end user if the system goes down. Replication is part of the design of the system/application and usually requires Tier A RPO (see below).
- Tier 2 – RTO of less than 24 hours. Requires hot standby equipment and usually a Tier B RPO.
- Tier 3 – RTO of less than 48 hours. Test and development equipment takes on a production role in the event of a disaster. This usually only applies when a company has a second data center with production running at one site, and test and development running at the other.
- Tier 4 – RTO of two to seven days. Includes lower priority applications than tiers 2 and 3. Supporting hardware can be either remaining capacity at a second data center or hardware available via drop ship arrangements with a third-party vendor.
- Tier 5 – RTO of more than seven days. Requires acquisition of hardware and restoration of systems.

Organizations determine RPOs based on the amount of data or transactions that they can afford to lose. Possible RPO tiers include:

- Tier A – No data loss
- Tier B – RPO of less than 24 hours
- Tier C – RPO of last backup (24–36 hours in most cases)

You must define RTOs and RPOs whether you are recovering at your own alternate data center, or at a cold or hot site operated by a third party. Third-party providers now have advanced recovery services that can meet high-availability requirements for RTOs and RPOs.

Tiers allow you to identify the relative criticality of systems to business processes. Tiering is the first step to identifying actual RTOs and RPOs. Setting actual RTOs and RPOs requires a clear understanding of business processes and the ability to balance costs.

## Chapter 4

# Finding RTOs and RPOs

In most cases, defining actual RTOs and RPOs is an iterative process. With no absolute formula available, there is a negotiation process with the business owner to balance the risk with the cost. That is, while there may initially be a requirement for short RTOs and RPOs, after weighing the costs of the solution, business owners may have to accept longer, less-costly RTOs and RPOs. It's a question of how much risk you are willing to take and for what cost.

So how do you go about determining RTOs and RPOs? It requires analysis. There is a wide range of financial and non-financial considerations to keep in mind. For example, it's the business owner's responsibility to answer (at least) the following questions:

- What does this business process use to do its work?
- What resources (people, skill sets, other tools) are needed for this process to continue to function in a disaster mode?
- What vital information flows through this business process, either from another process and/or to another process? What other business processes are dependent on this process?
- What activities of the process can be done manually, if needed? What manual workaround procedures could be put in place to minimize either the financial or non-financial impacts?
- What would be the direct financial loss to your company if this business process were not available for hours, days, or weeks? How is this loss calculated? What components contribute to this loss?
- Does this business process have business cycles? Would a significant loss to your company be different at different times of the year? What months are critical? Are there times of the month that are more critical than others?
- What is the business recovery plan? Are there subject-matter experts outside of the affected area that could process the work if critical employees are not available?
- What are the negative impacts of the following non-financial concerns if this process does not function for hours, days, or weeks? Examples include:
  - Cash flow (generation of revenue)
  - Public image
  - Shareholder confidence
  - Financial reporting
  - Managerial control (for example, approval levels)
  - Productivity
  - Competitive advantages
  - Industry image
  - Customer service

Do you really know the full impact of a certain business process going down? Figuring that out is the only way to ensure you apply the appropriate recovery methods.

- Vendor relations
- Legal/contractual violations
- Regulatory requirements
- Employee morale
- Consumer confidence
- For each day of outage, how long will it take to handle the critical backlogged work—in addition to other daily work—when this process is back in operation?
- What expenses would be incurred if this process were disrupted?  
Examples include:
  - Temporary employees
  - Emergency purchases (supplies, office machines, etc.)
  - Rental/lease of equipment
  - Wages paid to idle staff
  - Overtime
  - Temporary relocation of employees to alternate business recovery location
- What other vulnerabilities and exposures exist with this business process?

Of course, this is merely a sampling of possible general issues. Business impact analysis (BIAs) are usually customized to specific companies and industries. Best practices require that each organization develop its own set of questions. There are numerous resources available on the Internet relating to BIA development.

As with other business continuity plan components, conduct an annual review of your RTO and RPO requirements to capture changes to both the business environment and the systems environment.

## Conclusion

Linking RTOs and RPOs to business processes and overall organizational compliance is in the best interest of the entire organization. You must ask hard questions about the nature of these processes and the risks associated with them. The next step is to ensure the coordination of IT and business owner personnel. Enterprise planning is the only way to assure your plan will be resilient enough to meet or exceed your recovery and restoration of business operations.

## Contact Information

Frank Leonetti, CBCP, is the Sun Microsystems Professional Services Manager responsible for the Business Continuity and Disaster Recovery Practice of the Sun Storage Group. Frank is a Certified Business Continuity Professional with over 22 years in the IT industry specializing in the BC/DR area. He is also the President of the Association of Contingency Planners Liberty Valley Chapter and is a frequent keynote speaker at seminars and industry conferences.

You can contact Frank at [frank.leonetti@sun.com](mailto:frank.leonetti@sun.com) OR 609.267.9245

