WHITE PAPER:
DISASTER RECOVERY PLANNING GUIDE

The Business persons Handbook for ensuring Business Continuity

Disaster Recovery Planning Guide

Who should read this paper

CIOs, VP of IT, and anyone involved in Business Continuity Planning

Content

Introduction

"What matters most is not the size of the "storm," but the scale of the impact"

For many organizations, a "disaster" usually means something that impacts the data center from the outside, such as the wrath of a storm or of a violent terrorist act. While newsworthy events should inspire reflection on the state of our preparations, disaster recovery assessment shouldn't be limited to the consequences of a hurricane, earthquake or similar catastrophe. Lower profile but nevertheless important events – from software bugs to hardware failures – that may be every bit as consequential as fire or flood, need to be considered as well.

From a business perspective, a disaster isn't just what makes the news, but anything that makes the ordinary conduct of business difficult or even impossible. If an event, at any scale, can interrupt our operations, it poses a threat we cannot ignore. Whatever is at stake, be it the loss of revenues, reputation and customers – or even, for the security forces and medical professionals who serve and protect, the potential loss of lives – any unexpected IT interruption represents a potential disaster which we must either be prepared to avoid or from which we must be prepared to recover.

This *Disaster Recovery Planning Guide* offers a business perspective on what is often mistakenly considered a technological issue. As you'll see in subsequent pages, the most crucial considerations are determined more by business needs than IT requirements. In fact, the most important disaster recovery decisions are not about technology, per se, but are about the business demands that *drive* technology choices.

While the technologies for data recovery and application availability evolve over time, the underlying business reasoning, the core of any effective disaster recovery plan, remains consistent year after year:

- Assessing your business exposure to disaster
- Reviewing your options for cost-effective preparation and recovery
- Setting the expectations for performance that direct technology decisions
- Testing your plan for vulnerabilities

The first step toward recovery planning is disaster awareness: understanding what a disruption would mean to your business and what you can do to prevent or mitigate disastrous consequences.

Assess Your Business Needs

From a business perspective, a disaster isn't just what makes the news, but anything that makes the ordinary conduct of business difficult or even impossible.

If you were to ask your IT resources to assess your vulnerabilities, chances are, you would get a reasonably accurate accounting of your data and applications with pages of documentation about redundant drives, backups and, possibly, remote data centers. While such a report might expose the technological consequences of a disaster, it would not reveal the *business consequences* of lost hours, lost data and lost applications, leading to lost revenues, profits, customer confidence or worse outcomes.

To expose those consequences, you might conduct a "business impact analysis" (BIA) that calculates the dollars and cents costs of a single event on your business – such as a hurricane, hardware failure, sudden flood or software bug – taking into account the extent of the damage (how much data lost, how many interactions broken) and the duration of the disruption (how long it takes to restore data, applications and

operations.) Using this kind of math, you arrive at a number that represents your "potential loss" – the quantifiable sum of everything that may be at risk in the event of a sudden disaster.

Are You Missing the True Impact of 'Silent' Disasters?

In 2008/2009, the United States suffered a major financial meltdown, one with an impact that many economists have estimated at $1.8 trillion[1]. While we intuitively understand the consequences of a loss at that scale, most of us fail to recognize the extent of a "silent" IT disaster unfolding under our virtual noses. According to IT complexity expert and ObjectWatch founder, Roger Sessions, organizations in the United States lose $1.2 trillion from IT failures *every year*. Worldwide, the total comes to $6.2 trillion. Although Sessions' numbers have been challenged by other economists, their calculations remain sobering, concluding that threat worldwide is "only" $3.8 trillion!

The most notable aspect of Session's math is this: the overwhelming majority of the annual $1.2 trillion loss is *not* caused by the low-probability/high-consequence catastrophes that capture attention, but by *high-probability/low-consequence* events that occur frequently, such as software bugs, hardware failures and security breaches. Worse, as applications become more complex, involving an ever-larger tangle of codes, data nodes and systems networks, the exposure to these "smaller" events becomes more frequent and their impact more costly.

Going Beyond Business Impact Analysis

If your only assessment of loss is a business impact analysis, you may be missing the real cost of disasters, and failing to adequately plan for recovery.

While it is still important to conduct the business impact analysis, decision-makers must not allow the results to blind them to the consequences of multiple high-probability events that, year after year, impose losses on their enterprises. Your recognition of "potential loss" from a catastrophic event must be complemented with a deep understanding of "expected loss" – a more realistic figure that factors in two critical elements:

- **Probability:** "Expected loss" includes calculations for probability, the likelihood of a loss event that "potential loss" neglects. When probability is accounted for, the significance of multiple, small events becomes visible, allowing you to direct budgets and resources to the events that really matter: the ones that regularly impact your bottom-line.
- **Current investments:** In its doomsday calculations, the business impact analysis often fails to account for current investments in recovery – such as backups and automated failovers – that would temper overall losses.

Although the process may seem intimidating, there's a silver lining to the cloud of more sophisticated loss assessment. No matter how much you invest in preventative or corrective action, "potential loss," because it measures the overall value of business at risk, never goes down. But when you shift your focus to expected losses, which include accommodations for probability and corrective action, you can actually see a reduction in your loss exposure – and measure the value of your disaster recovery investments.

[1] 1. These and subsequent figures are from The IT Complexity Crisis: Danger and Opportunity, Roger Sessions, November 8, 2009

Action steps:

1. Include probability into your risk calculations to arrive at realistic "expected loss" figures.
2. Shift your focus so that high-probability/low-consequence events figure as or more prominently in your disaster recovery planning than low-probability/high-consequence catastrophes.
3. Identify and protect the "hidden" dependencies (such as supplier networks, access to physical buildings or even availability of personnel during a disaster) that must be taken into account to recover critical data and applications.
4. **Establish priorities:** Not all data and applications are equal; the bulk of your disaster recovery planning should be directed toward the top 20% of expected losses.

Review Your Options

All applications and data are not equal: in view of their business impact, some merit much greater investment in disaster recovery while for others, you may tolerate lower standards for recovery. Technology choices should mirror business objectives: the priorities you established in the previous section should dictate the level of investment you make in disaster recovery.

Two key decisions: Recovery Point Objective and Recovery Time Objective

At the heart of your disaster recovery plan are two critical decisions that reflect your tolerance for loss:

- The **recovery point objective (RPO)** that determines the *moment* in time, before the disrupting incident, that you restore to. The closer the RPO to the incident, the lower the data loss.
- The **recovery time objective (RTO)** that establishes the *amount* of time it takes to restore operations. The lower the RTO, the less time it takes to recover.

Match Your Service Level Agreements to Your Priority Tiers

1. RPO/RTO of Seconds to Minutes: This category includes data and applications so important – by measures of public safety (health, military, police) or financial impact (banking, insurance, trading) – that they demand a zero RPO and a zero or near zero RTO. Meeting your obligations will require investments in automated solutions that can respond instantly to disaster.
2. RPO/RTO of Minutes to Hours: Here, the data and applications are important, but not mission-critical: think Enterprise Resource Planning, Customer Relationship Management and email for example. Automation still plays a role, but you can accept some minor data loss from your RPO, and can endure a few hours of recovery time delay.

US Bank Saves Millions by Addressing Expected Losses

It was a CEO's ugliest nightmare: after his staff ran a business impact analysis, the leader of a major US bank was told that the business' potential loss, reflecting 24 hours of downtime to its trust system, was a whopping $80 billion a day; worse, the bank anticipated a recovery time of 48 hours, doubling anticipated loss to $160 billion. Seeking relief, the CEO turned to an IT vendor who promised to shrink the recovery time objective (RTO) from 48 hours to six for a cost of $3.5 million. Weighed against the potential loss of $160 billion, a $3.5 million investment seemed reasonable.

Suspecting that the real risk was much lower, the CEO brought in a team of Symantec risk analysts to assess the situation. Their investigation found that the bank already had manual mitigations in place that brought the trust system's true risk, when accounting for probabilities, to a mere $150,000. Suddenly, the $3.5 million "fix" no longer seemed like an efficient investment.

Further, the Symantec team discovered feeder systems previously ignored, representing an expected loss of $10 billion if they remained unprotected. For only $750,000, they calculated, the bank could reduce its risk by $9 billion.

3. RPO/RTO of Hours to Days: Consider this the place for your less critical, but nice to have applications, such as your internal website or human resource functions. Time is not of the essence, and much of your disaster recovery can be managed through inexpensive manual efforts.

Action steps:

1. Discriminate: rank and categorize your data and applications by their business or safety significance.
2. Assign different RPO and RTO performance requirements to your tiers.
3. Budget unequally, anticipating a higher spend on your most critical tiers.

Set Your Expectations

By definition, your IT team has technology expertise, but as the business decision-maker, you must set the *objectives* the technology must achieve. Chief among these is simplicity: when disaster strikes, recovery must be simple and easy if it is to be rapid and effective.

Fast, easy recovery requires:

Automation: In the event of a disruption or emergency, you will not have time to assemble teams, coordinate meetings and distribute responsibilities. To meet your previously determined RPOs and RTOs, you need events-driven application management, an automated process that eliminates or minimizes manual intervention.

Comprehensive fit: Your IT infrastructure wasn't built in a day, but took shape over time, incorporating a mix of environments (physical and/or virtual), platforms and operating systems. Regardless, your disaster recovery technology must work across *all* your components, capable of communicating and coordinating events among disparate pieces.

Availability and reliability: Data storage resilience – the ability to recover quickly from failure – must be accompanied with data ubiquity that ensures that all systems that require the recovered data can find and access it.

Consider Recovery *Capacity* Objective

Your recovery point and time objectives are the backbone of your disaster recovery plan. But if you're obligated to fulfill service level agreements (SLAs) for your customers, you should consider a third metric: the recovery capacity objective (RCO), the acceptable amount of functionality you need, not only to recover, but to return to the *contracted standard* of service you are obligated to fulfill. Your RCO represents a level of performance that can vary from a compromised "brown-out" to a complete return to full service.

Simple restoration of complex applications: That one purchase on an ecommerce site or that one withdrawal from an ATM? Behind the scenes, these single activities represent a complex, multi-tiered *stack* of technology that often includes application code, stored data access, middleware connectivity, and other functional layers. Effective recovery requires technology that can not only restore each layer, but restore them in the right order and re-integrate their activities to recover the entire application. If your current recovery technologies cannot restore the entire stacks of multi-tiered applications, they cannot recover your most business-critical technology functions.

Action steps:

1. Assess your current disaster recovery components. Are they integrated and automated for rapid action, or will your recovery be delayed by the need for coordinated manual interventions?

2. Dig into your application layers to be sure that *every* tier can and will be restored, in the right order, in the event of disaster.
3. Conduct an IT inventory to expose the system elements and dependencies that must be restored together to effect a rapid recovery.

A Dive into Data Replication

"Data replication" refers to the process by which data in one site is mirrored in another, typically the backup location designated for disaster recovery. There are different types of data replication, each with its strengths and weaknesses:

- **Synchronous:** With synchronous data replication, data tasks at the primary site are not acknowledged as completed until they have been replicated at the secondary site. While synchronous replication closes the RPO gap, it comes with some expense in systems performance, and its application is limited to sites within a wide area network typically no more than 60 kilometers apart.
- **Asynchronous:** Asynchronous replication accommodates mirroring data across any distance and allows the primary center to write to disc without waiting for acknowledgment from the secondary. Although asynchronous replication allows for greater speed and distance, it opens up a gap in the data record between the two sites, potentially compromising RPO.
- **Hybrid:** The hybrid approach applies both methods, using synchronous replication for almost instantaneous availability in the event of localized failures, and asynchronous replication to a distant data center to provide restoration in the event of a catastrophe. The client determines the failover threshold from one site to the other, and the application of the hybrid solution requires sophisticated planning.

'Automatic' vs 'Automated'-What's the difference?

- "Automatic" requires no manual intervention whatsoever; the triggering event initiates a sequence of activities almost instantaneously.
- "Automated" refers to processes that, once initiated by a manual action or decision, run without further need for intervention.

For localized recovery, such as disc to disc or even a failover to a nearby data center within a wide area network, automatic solutions are preferred. But for failovers to distant data centers that might impose disruptions to data streams, automated processes give businesses the power to make informed choices.

Test Your Plan

You've assessed your needs, established priorities, matched service levels to those priorities, and set the expectations for your recovery solutions. Once the technology has been identified and purchased, you're prepared for disaster recovery. Right?

Wrong. Your final step, the one you take to make sure your plan truly meets your needs, is one you'll repeat time and again: *testing*. You do not test to prove that your plan works; you test to expose your vulnerabilities, to make the unknown known BEFORE disaster strikes.

The truth is, if you have never failed a disaster recovery test, you do not have a comprehensive disaster recovery plan. By actively searching for and finding the holes in your plan, you can make informed business decisions:

- If the probability of a particular failure is low, or the consequences of that failure minor, you might decide that additional protection is not worth the added expense.
- If, however, you find vulnerabilities that are probable, or could have significant consequences, or both, you now know precisely *where* to direct your investments.

Because it's not a matter of "if," but "when": Put your disaster recovery plan to the test

If you have never failed a disaster recovery test, you do not have a comprehensive disaster recovery plan.

Use the following checklist to ensure you have determined the recovery needs and technology objectives your IT team must execute effectively:

- Have you linked IT functions to business consequences – and assigned a dollar value to their significance?
- Does your definition of disaster include the high-probability/low-consequence events that cause the majority of catastrophic business disruptions?
- Can you calculate, not just potential loss, but expected losses? Do your calculations reflect both current mitigations and event probabilities?
- Have you used your expected loss figures to focus your disaster recovery priorities?
- Do your RPO and RTO service levels reflect your priorities?
- Have you created a hierarchy of tiers that allow you to make recovery investments matched to the business significance of your applications and data?
- In addition to your RPO and RTO, have you set a recovery capacity objective (RCO) that acknowledges gradations in recovery status?
- Do you have the appropriate data replication model for your recovery needs?
- Are your recovery solutions automated to facilitate rapid, coordinated recovery in the event of disaster?
- Can your current recovery solution embrace your entire technology environment, regardless of platform, operating system, and other variables?
- Will your stored data be ubiquitous upon restoration, available to every application and system that needs it?
- **Crucial:** Can your recovery solutions restore *every* layer in your complex, multi-tiered applications, automatically and in the correct order?
- Do you regularly test your disaster recovery plan, not to prove efficacy, but to expose vulnerabilities?

If you cannot answer "yes" to every checklist question, you have areas in your disaster recovery plan that may need more attention.

More Information

For expert help in recovery planning and execution:

Visit our website

http://go.symantec.com/business-continuity

To speak with a Product Specialist in the U.S.

Call toll-free 1 (800) 745 6054

To speak with a Product Specialist outside the U.S.

For specific country offices and contact numbers, please visit our website.

About Symantec

Symantec protects the world's information, and is a global leader in security, backup, and availability solutions. Our innovative products and services protect people and information in any environment – from the smallest mobile device, to the enterprise data center, to cloud-based systems. Our world-renowned expertise in protecting data, identities, and interactions gives our customers confidence in a connected world. More information is available at www.symantec.com or by connecting with Symantec at go.symantec.com/socialmedia.

For specific country offices and contact numbers, please visit our website.

Symantec World Headquarters
350 Ellis St.
Mountain View, CA 94043 USA
+1 (650) 527 8000
1 (800) 721 3934
www.symantec.com