

# Risk Management and Business Continuity

Overview and Perspectives

White Paper

*By Dennis Wenk*

February 2005

## *Executive Summary*

For organizations struggling with the issues of disaster recovery and business continuity, the underlying uncertainties and strong emotions can be overwhelming. Risk management concepts can be applied to help organizations understand the salient issues, identify the key characteristics, and come to grips with the fundamental problem.

However, while expert opinions, risk-audits, best practices, or performance standards may identify significant risk points, they only represent an initial step. It is difficult to know whether a risk is worth reducing without having sense of its size and the cost to reduce or mitigate it.

This white paper provides some background on risk management and traditional vulnerabilities, and offers an overview regarding pitfalls to avoid. More importantly, it reveals the need for a quantitative risk management approach that will reduce uncertainty. The key is to objectively evaluate alternatives and/or competing solutions while providing the business justification, through cost-benefit analysis, for selecting the optimal answer for each particular environment.

# Contents

<b>Introduction .....</b>	<b>1</b>
<b>Background .....</b>	<b>2</b>
Uncertainty Creates Complexity .....	2
Emotions Cloud the Facts.....	2
<b>Business Continuity, Not Recovery .....</b>	<b>3</b>
<b>Demystifying Risk Management: Pitfalls .....</b>	<b>4</b>
Perfect Security .....	4
“Iron Doors in Paper Walls” .....	4
The “Vulnerability” Misconception .....	4
<b>Going Beyond Best Practices.....</b>	<b>5</b>
Quantitative Risk Management .....	5
<b>Case Study: Company A.....</b>	<b>6</b>
<b>Summary .....</b>	<b>7</b>
<b>Appendix A: About the Author .....</b>	<b>8</b>

# Risk Management and Business Continuity

## Overview and Perspectives

White Paper

*By Dennis Wenk*

### *Introduction*

Business continuity requirements are changing, forcing organizations to take a hard look at their round-the-clock operations and growing service-level expectations. Add to this equation the emergence of closer regulatory scrutiny and stringent out-of-region data protection requirements, and it becomes clear that the increased sensitivity to loss of information assets is not going to subside any time soon. Is this necessarily a bad thing? After all, the challenge is to understand and reduce risk...and increase business resilience.

At the heart of this issue is data. Data is an irreplaceable corporate asset. Replacement copies of vital corporate records are not available on the open market. Without a copy of the data, loss will be irreversible and permanent. Data stored on many systems is viewed as a strategic asset, and as the value of the data increases, the reliability and quality of storage solutions must be considered. There is no question that storage is the primary technical component of all business continuity solutions.

Many organizations—along with stakeholders and regulatory authorities—are focusing on operational risk assessment, and one key area of focus is information systems and data protection. There is, however, a general uncertainty that many stakeholders feel about disaster recovery and business continuity because it is not necessarily obvious what the risks are, what the exposure to data loss is, and what the best option is in order to mitigate risks and improve data protection.

For organizations struggling with the issues of disaster recovery and business continuity, the underlying uncertainties and strong emotions can be overwhelming. Risk management concepts can be applied to help organizations understand the salient issues, identify the key characteristics, and come to grips with the fundamental problem.

This white paper provides some background on risk management and traditional vulnerabilities, and offers an overview regarding pitfalls to avoid. More importantly, it reveals the need for a quantitative risk management approach that will reduce uncertainty. The key is to objectively evaluate alternatives and/or competing solutions while providing the business justification, through cost-benefit analysis, for selecting the optimal answer for each particular environment.

## Background

Business continuity planning involves consideration of risk potential, trade-offs, choices, and resource allocation. Typical binary thinking that appraises what did or did not happen must be replaced with methods of measuring what might happen and how likely it is to happen, and, based on this assessment, discerning how to best choose the IT infrastructure and technologies that will protect the data.

Decisions concerning business continuity involve interruption risk. There is no “on-off” switch that can tell whether an organization is in the domain of safety or danger. Interruption risk is a matter of *degree*, which can range from zero availability (the catastrophic event, total long-term outage) to one hundred percent (24/7 continuous availability).

The factors that benefit business operations—speed of processing and access to information—also increase risks of computer intrusion, fraud, and disruption. Technology has long been at some risk from malicious actions or inadvertent user errors and from natural and man-made disasters. In recent years, business systems have become more susceptible to these threats because computers have become more interconnected and, thus, more interdependent and more accessible to a larger number of individuals. In addition, the number of individuals with computer skills is increasing, and intrusion or “hacking” techniques are becoming more widely known via the Internet and other media.

### Uncertainty Creates Complexity

Disaster recovery is a perplexing subject filled with many uncertainties. The underlying uncertainty is whether organizations are adequately protected against an ever-increasing number of threats. There is uncertainty regarding the types of threats, their potential harm, and the likelihood of their occurrence. Which threats are the really important threats and which are innocuous? There is uncertainty regarding vulnerabilities, the magnitude of loss, and which solutions will provide the greater benefit.

Uncertainty is what makes recovery a complex business problem. Under conditions of uncertainty, both rationality and measurement are essential to decision making. Given the number of variables, it is very difficult to know intuitively whether organizations are *spending the right amount on the right things*. There are no *best practices* that eliminate risk entirely and no *expert* that can predict the future with absolute certainty. Organizations can, however, use science to provide a set of guiding principles to make good decisions about what is acceptable and to make sound investments. By using statistical data, organizations will reach much more accurate results.

### Emotions Cloud the Facts

In addition to the complexity caused by all the uncertainty, disaster recovery is also an emotional problem because catastrophic events evoke strong emotions. Disasters like the events of September 11, 2001, create fear. It is seldom a surprise that in many domains, emotions have a large effect on judgment and decision making. Relying on emotion, however, can lead to serious errors of judgment, both in the form of excessive fear of small risks and neglect of large ones. Strong emotions cloud otherwise quantitative judgments, including judgments about probability and making the best case or the worst case seem highly salient.

The lack of information has the same effect as chance. The aim should be to reduce the scope of chance. Organizations will make better choices by using statistical and probability data. A quantitative approach to business continuity will help organizations understand risk, measure it, and weigh its consequences. Science shows more promise of accuracy than intuition and hunch, where inconsistency and myopia so often prevail. A quantitative approach will identify threats that adversely affect critical operations and assets, and estimate the likelihood that such threats will materialize.

## *Business Continuity, Not Recovery*

Over the past 10 to 15 years, a change in the character of IT has increased the need for continuous access to business data. Organizations have integrated IT into the fabric of their operations. This integration has contributed to the explosive growth of Internet and intranet computing, which increases an organization's reach globally and extends the hours of operations to 24/7. Worldwide operations have no time boundaries and virtually eliminate inherent time zone differences. Interruptions to key business applications in this environment can have significant negative impact on revenue, productivity, and expenses.

The reconstruction and restoration of data either on the premises or at another location (for example, a hot site) from tape backup simply doesn't meet the performance requirements of most businesses. CEOs recognize that what is needed is the ability to resume IT operations following any interruption quickly, if not instantly. The use of the term "business continuity," with its suggestion of no interruption, in lieu of "disaster recovery," which implies a gradual resumption, emphasizes the importance of this change in IT.

Traditional disaster recovery addresses long-term, catastrophic events and may seem to be a costly and disruptive solution to an unlikely occurrence. This view is changing rapidly as it becomes apparent that in reality, the same protections can eliminate the short-term disruptions that have significant and continual negative economic impact on business operations. The accumulation and frequency of these limited disruptions is such that their aggregate cost outweighs the extremely rare event (and cost) of the destruction of IT facilities. Most organizations are reluctant to make large investments for uncertain and rare events. To many it appears as if there are no tangible benefits for disaster recovery. They think it is expense that brings only peace-of-mind, and investments for peace-of-mind lose out to those investments that have a positive business benefit.

Disaster recovery planning is focused around the "smoking crater." Invoking a disaster recovery plan can take longer than the vast majority of real disruptions an organization experiences, such as system failures, software defects, power failures, equipment breakdowns, hacker attacks, and operator errors. The disaster recovery plan can be ineffective for these more familiar problems. Even planned downtime for maintenance takes precious time away from valuable business functions. These short-term disruptions are far more common and, given the 24/7, global operating mode, can be very costly even though they cause relatively short-term outages. As a matter of fact, persistent short-term failures can be extremely detrimental to a business.

By implementing a robust IT infrastructure that supports business continuity, the CFO will see a payoff as often as several times a year through more rapid recovery than, for example, tape backup. Every organization needs to be able to recover from a disaster if it occurs, but the only irreplaceable element is the data. Consequently, it may make sense to copy data to tape daily as the ultimate offsite backup, but the important requirement is to keep the business running with data replication in the face of the far more frequent but less disastrous events.

## Demystifying Risk Management: Pitfalls

Before spending more time on identifying some techniques that support improved data protection through better management of risks, let's take a quick look at some pitfalls to avoid.

### Perfect Security

A few IT managers believe that the objective is perfect security. That is to say, any and all security failures are unacceptable. This is not a rational risk management goal for a simple reason. *Perfect security is infinitely expensive* because there is no obvious point at which one can stop spending. In any situation, there is always something more one can do to increase security. However, no organization has unlimited resources.

### “Iron Doors in Paper Walls”

The risk management technique that bears the above colorful title addresses perceived security exposures one by one. For example, the risk manager notes that there is a wooden door that could be forced open relatively easily. The obvious solution is to replace the door with a sturdy iron door with a combination lock, a door-open alarm, a magnetic card reader, and a CCTV camera. However, the risk manager overlooked the fact that the door is installed in a paper wall. Overall security would be much stronger if the same resources were spent replacing the paper walls with two layers of 5/8-inch drywall. This fanciful example illustrates an important point. Rational risk management *requires* that the risk manager consider the totality of the risks before developing a risk management strategy for allocating the limited resources available to implement the selected security measures.

### The “Vulnerability” Misconception

Some “experts” espouse the use of “vulnerability analyses.” The concept is that an examination of an IT facility will reveal “vulnerabilities” which “attackers” will “exploit.” This concept probably was spawned by the concern over hacker attacks. Indeed, hacker tactics that take advantage of discovered software vulnerabilities are referred to as exploits. To understand why a vulnerability analysis is not a sufficient basis for risk management, consider the following example:

The risk manager observes that the IT facility, located in a small town in Iowa, is not enclosed in an elephant-proof fence. This creates a vulnerability to being attacked by enraged bull elephants. The vulnerability serves as the justification for building the elephant-proof fence.

Now ask yourself why you know this is absurd. The obvious answer is because there are no wild elephants in Iowa. Exactly, but notice that the “vulnerability analysis” alone did not satisfy you. You immediately did an informal risk analysis, and noted that the occurrence rate of the threat—enraged bull elephants in Iowa—was zero, and therefore *the vulnerability was irrelevant in Iowa, but might be significant in parts of Africa.*

## Going Beyond Best Practices

The concept behind best practices is that there is a set of silver bullets that will shoot down all the IT risks: "Adopt my set of best practices and you will be safe." In September of 2002 the United States Homeland Security Office published a "strategy" for Cyberspace Security that presented a list of about 86 recommendations to enhance IT security. The problem was that the document was not a strategy. It was a list of 86 "best practices" in disguise, and it was silent on the question of *exactly which* measures one should implement in a specific situation. Early in 2003 rumor had it that a new draft was coming out with only 60 recommendations.

There are two basic flaws in the best practices concept:

1. Who says a given set of best practices is, in fact, best? How can any list be expected to be best for all possible circumstances? Surely some desired measures will be omitted, and others will waste resources.
2. Since resources are always limited, it is essential to prioritize the recommendations and defer the items that don't make the cut; but, interestingly, no one has ever published a best practices list that explains how to prioritize. The reason for the omission is simple. It is *impossible* to devise a set of prioritizing rules for the simple reason that each IT facility has its own unique set of risks, loss exposures, and mitigation measures.

## Quantitative Risk Management

The objective of rational risk management is to optimize the allocation of limited resources to IT security. This can only be done by balancing the *cost of mitigation* against the *reduction in future losses* that the mitigation is expected to generate. In other words, we want to find \$10 solutions to \$1 million risk exposures, not the other way around. This means that the analysis *must* be quantitative since the resources are allocated quantitatively. It is absurd for an IT risk manager to say to a CFO: "I have identified a *big* risk exposure, so I need a *big* resource allocation!" In short, a qualitative risk analysis is insufficient as a risk management tool.

The *only* technique that can meet rational business management objectives is a quantitative risk management procedure. The major advantage of a quantitative risk approach is that it provides a measurement of the magnitude and probability of the disruption, which can be used in a cost-benefit analysis of recommended alternatives. No organization is able to afford the same level of protection for all functions and assets. Since there are a vast number of threat events that could cause damage, a ranking of events by their relative economic impact permits a better allocation of finite resources to areas of highest priority.

This is a three-step process as follows:

1. Construct a quantitative model of risks. The model will be the loss-potential experience relating to the current environment, as well as a ranking of events that are most likely to generate material losses.
2. Determine the options available to reduce a risk or mitigate the consequences. Each option will be characterized by its cost to implement and maintain as well as its expected impact on the material threat events.

3. Use the quantitative model to evaluate each option and identify the option measures with a favorable return on investment (ROI), and provide a basis for prioritizing their implementation.

## *Case Study: Company A*

A financial services company was operating two IT facilities several hundred miles apart. Incremental backups were taken nightly and shipped to the alternate site daily. The stated goals of this program were a recovery time objective (RTO) of 48 hours and a recovery point objective (RPO) of 24 hours. They had a serious exposure in their business continuity program using this traditional backup/restore strategy and wanted to improve but were having difficulty justifying additional investment in technology.

After an initial review of the current situation, we found the nightly backups often continued well into the business day. For this reason, the backups did not leave the building until the end of the business day. The travel time to the alternate site took 17 hours. It took six hours to recover the operating environment and another 36 hours to restore the database. This produced a worst-case RTO of 59 hours, which caused an opportunity loss of two business days.

During staff interviews we discovered that most of the business functions had been automated to increase productivity, speed the processing and access to information, and facilitate interconnectivity with suppliers, partners, and customers. In this effort to automate, most of the company's hard-copy audit trails were replaced by electronic data. This increased the severity of any data loss because it would be very difficult for the company to reconcile its business activity. The fact that backups did not leave the premises until the end of the business day meant that the backup tapes could be exposed to the same catastrophic event as the facility, such as fire. The conclusion was that the worst-case RPO, a critical business continuity component, was actually 72 hours of lost data.

Any protracted interruptions could make it difficult for the company to meet its obligations for the days in question or reconcile the days' records. This would lead to significant liquidity bottlenecks by impacting the ability of the company to communicate its own financial position as well as that of its customers. Such interruptions would also lead to an increased exposure to fraud.

Market-based and geographic concentration of a customer organization intensifies the impact of operational interruption. Since our customer played a significant role in the critical financial markets of a number of countries, its sudden absence would present a systematic risk to the financial markets of the affected countries and threaten the stability of those markets. Additionally, considering the large volume and value of transactions/payments that are cleared and settled on a daily basis, failure to complete the clearing and settlement of pending transactions within a business day would create systemic liquidity dislocations, as well as exacerbate credit and market risk for critical markets.

This information was used to quantify both the loss potential and expected loss for our customer. It was determined that the loss potential for a 48-hour (the stated RTO) interruption was approximately US\$27 million. This information was important, but we still needed to determine the annualized loss expectation (ALE) based on the customer's unique risk profile. The ALE was over US\$12 million.

Using the identified threats and the ALE information, we developed four alternative strategies to mitigate the most harmful risk to the customer. We evaluated the alternatives based on their ability to reduce the risk. While all the alternatives reduced the customer's ALE, the Hitachi TrueCopy™ Remote Replication software solution reduced the customer's ALE well over US\$8 million a year, taking into account TrueCopy software's implementation and operating costs of US\$2.6 million. The annual net benefit to the customer was over US\$5 million.

## Summary

The purpose of quantification is to give decision makers a clearer picture about the exposures likely to be experienced so that they can make prudent allocation of limited resources. While expert opinions, risk-audits, best practices, or performance standards may identify significant risk points, they only represent an initial step. It is difficult to know whether a risk is worth reducing unless we have both a sense of its size and the cost to reduce or mitigate it. Risk cannot be eliminated entirely; it is a matter of degree. Just recognizing that a risk exists does not explain how to minimize the risk cost-effectively. There is neither an obvious stopping point to spending nor a specific indicator of more detrimental under-spending.

While it is impossible to predict the future with absolute certainty, science provides a set of guiding principals to rank the serious issues and facilitate sound management decisions. Quantifying the high risks and critical exposures will enable organizations to fully understand how function-rich technologies from Hitachi Data Systems (such as TrueCopy Remote Replication software, Hitachi ShadowImage™ In-System Replication software, and Hitachi Universal Replicator software) can improve their operation and which of these solutions will offer the greatest benefits for their particular environment. A quantitative risk approach demonstrates how best to take advantage of Hitachi technology to reduce or mitigate the consequence of disruptions.

## *Appendix A: About the Author*

### **Dennis Wenk, Senior Global Architect, Business Continuity and Resiliency, Hitachi Data Systems**

Over the last 33 years, Dennis Wenk has served in plethora of information systems and industry positions, ranging from application programmer to systems development director, from computer operator to data center manager, and from senior IT auditor to vice president of information processing. His extensive areas of experience include quantitative risk assessments, continuous availability reviews, emerging technology solutions, complex technology consolidations and integration, facility construction and relocation, centralized versus distributed alternatives, platform selection and implementation, outsourcing evaluation, business impact assessments, capacity planning and performance analysis, business contingency planning, open systems architecture, and multiprotocol networking solutions. He gained much of this experience at positions with IBM Global Network, Comdisco, KPMG, Consumer Systems Consulting, Zurich Insurance, Heller International, and Disaster Avoidance & Recovery Services.

 **Hitachi Data Systems Corporation****Corporate Headquarters**

750 Central Expressway  
Santa Clara, California 95050-2627  
U.S.A.  
Phone: 1 408 970 1000  
**www.hds.com**  
**info@hds.com**

**Asia Pacific and Americas**

750 Central Expressway  
Santa Clara, California 95050-2627  
U.S.A.  
Phone: 1 408 970 1000  
**info@hds.com**

**Europe Headquarters**

Sefton Park  
Stoke Poges  
Buckinghamshire SL2 4HD  
United Kingdom  
Phone: + 44 (0)1753 618000  
**info.eu@hds.com**

Hitachi Data Systems is registered with the U.S. Patent and Trademark Office as a trademark and service mark of Hitachi, Ltd. The Hitachi Data Systems logotype is a trademark and service mark of Hitachi, Ltd.

TrueCopy and ShadowImage are trademarks of Hitachi Data Systems Corporation.

All other company names are, or may be, trademarks or service marks of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, express or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems. This document describes some capabilities that are conditioned on a maintenance contract with Hitachi Data Systems being in effect, and that may be configuration-dependent, and features that may not be currently available. Contact your local Hitachi Data Systems sales office for information on feature and product availability.

Hitachi Data Systems sells and licenses its products subject to certain terms and conditions, including limited warranties. To see a copy of these terms and conditions prior to purchase or license, please go to [http://www.hds.com/products\\_services/support/license.html](http://www.hds.com/products_services/support/license.html) or call your local sales representative to obtain a printed copy. If you purchase or license the product, you are deemed to have accepted these terms and conditions.

©2005, Hitachi Data Systems Corporation. All Rights Reserved.

WHP-183-00 February 2005