Chapter Title: Disaster Recovery Planning

Authors: Marco Cremonini and Pierangela Samarati, Università di Milano

Outline

Glossary

Alternate Site. A location, other than the normal facility, used to process data and conduct critical business functions in the event of a disaster.

Business Continuity Planning. A generic term covering both disaster recovery and business resumption planning.

Critical Functions: Business activities that, if unavailable, can result in high cost for the organization.

Disaster. Any event that can cause the inability of an organization to provide critical business functions for a certain period of time.

Disaster Recovery Planning: The preparation of the activities and measures necessary to minimize losses and ensure the continuity of critical business functions in the event of a disaster. SIMILAR TERMS: Contingency planning; business resumption planning; corporate contingency planning; business interruption planning; disaster preparedness.

Disaster Recovery Teams: A structured group of teams responsible of controlling the recovery operations if the event of a disaster.

Electronic Vaulting. Transfer of data to an offsite storage facility via a communication link different a portable media.

Facility. Location containing equipment and infrastructure to perform normal business functions.

Loss. The cost resulting from a disaster.

Network Outage. The interruption of the system availability as a result of a communication failure affecting a network.

Off-site Storage Facility. A secure and remote location where backups of hardware, software, data files, documents, or equipment are stored.

System Outage. An unplanned interruption in system availability resulting from operational problems or hardware/software failures.

ABSTRACT

Historically, enterprises that have survived natural or man-made disasters have had recovery plans in place and were ready to face the risk of business interruptions. Today, the increased dependency on technology and tighter requirements in terms of recovery speed imposed on e-Business and Web-based services have made recovery plans more complex. Disaster recovery planning, which plans for failures of components of the IT infrastructure, system or network outages, mis-configurations, and natural disasters, should be treated as a necessity by conscious system administrators, managers, and CEOs. In this chapter we discuss disaster recovery planning, illustrating the main phases of the planning process and the techniques that can be used to provide for recovery.

INTRODUCTION

For many years, corporations have placed most of their critical information assets into automated systems and, more recently, they have adopted business models largely based on the Internet. While allowing for more efficiency, cost reduction, and competitive advantage, these developments have also increased the dependency on the Information Technology (IT) infrastructure and therefore of the potential economic losses caused by its failure. More than ever today a loss of access to the IT infrastructure will prove to be in many cases fatal for the business of affected companies. It is therefore crucial for companies to consider precise plans for counteracting possible events that may cause such interruptions. Examples of companies that failed to recover their businesses after a prolonged interruption of IT functions, due to floods or fires, for example, are well known and documented (Hiles, 2002) (Kaye, 2001) (Noakes-Fry and Diamond, 2001), (Toigo, 1999), (Toigo, 2001a).

The disaster recovery planning process consists of a set of activities aimed at reducing the *likelihood* and the *impact of disasters* on critical business assets (NIST, 2002), (Toigo, 1999), (Wold, 2002). Note that, while the term disaster recovery has been traditionally associated with protecting against external catastrophic events (e.g., floods, fires, earthquakes), we adopt in this chapter a broader definition of disaster recovery, which includes also the business continuity planning, and therefore the consideration of failures to technological equipments and software. E-business and Internet connections are shortening recovery-time requirements and changing the way we think about Disaster Recovery Planning: full 24x7 operation and “zero downtime” are now commonly seen as a business necessity. While traditional countermeasures to face natural disasters should still be provided, a wide spectrum of other risks (e.g., cyber-attacks) has to be considered. Also, critical business processes must be analyzed and recovered much faster than in the past: downtimes of days (the usual time window for natural disasters), are definitely no longer sustainable for Internet-based businesses.

Especially after the year 2000, companies have made massive investments in re-engineering their business processes and many of them realized that those investments need protection from events other than natural disasters, including outages, power failures, security incidents, and misconfigurations. New methodologies, techniques, and commercial products have been developed to manage backups and to speed up recovery. Thank you to these advancements, it is now possible to plan for a recovery time of just few hours (still far, however, from the “zero downtime” dream, sought by many and sold by others).

Nevertheless, many enterprises still have minimal contingency plans in place to ensure business continuity for their business requirements. Managers fail to plan for disaster recovery in the new project development life cycle, resulting in more risks and exposures than in the past. From a technological standpoint, some managers and planners still do not understand the intrinsic fragility of the new architectures for the e-business to which they are migrating. As a consequence, they do not fully understand the increased requirements for business continuity, do no plan for the inevitable failures, do not mitigate the risks, and are not prepared to sustain the consequences.

In this chapter we address the problem of disaster recovery planning. The remainder of the Chapter is organized as follows. First, we present the risks, that is, the main causes of failures that must be accounted for in disaster recovery planning. Then, we discuss the possible costs that not successfully responding to failures may have on companies. We continue illustrating the main phases of the disaster recovery planning process, including the risk management phase, to which a whole section is then devoted. The chapter continues then presenting backups and alternate sites as main techniques for recovering data and providing for the continuity of the business processes in the occurrence of failures. We also discuss failure issues in the emerging web hosting service scenario. Finally, we provide a reference template for disaster recovery planning.

RISKS, OR THE CAUSES OF FAILURE

Failures in the IT infrastructure may have many causes that, with a broad classification, can be classified as:

- *Natural*, such as , flood, fire, tornado;
- *Human*, such as operator error, sabotage, network intrusion, malicious code, and denial of service;
- *Environmental*, such as equipment failure, software error, telecommunication network outage, and electric power failure.

Different failure causes differ in terms of *likelihood of occurrence* and *potential damage* that they can cause. For instance, events like sabotage, floods, or fire, usually considered in traditional disaster recovery, are rare but, when they happen, they turn out catastrophic. On the other hand, events like equipment failures or software errors are less catastrophic in nature, but may occur frequently and represent almost day-to-day problems for system administrators. Usually, these latter types of failures are not devastating for a company, and in past years (in the pre-

Internet age), they were not included in disaster recovery planning. However, with today's companies relying more and more on networked systems, complex distributed multi-component architectures, and the Internet as a primary business strategy, short but frequent downtime periods can become a major problem in terms of competition, reputation, and revenue losses (Kaye, 2001), (Toigo, 2001b). For each possible cause of failure, likelihood of occurrence and possible damages must both be considered in disaster recovery planning.

Figure 1 shows the results of a study by Gartner Group on the causes of IT failures.

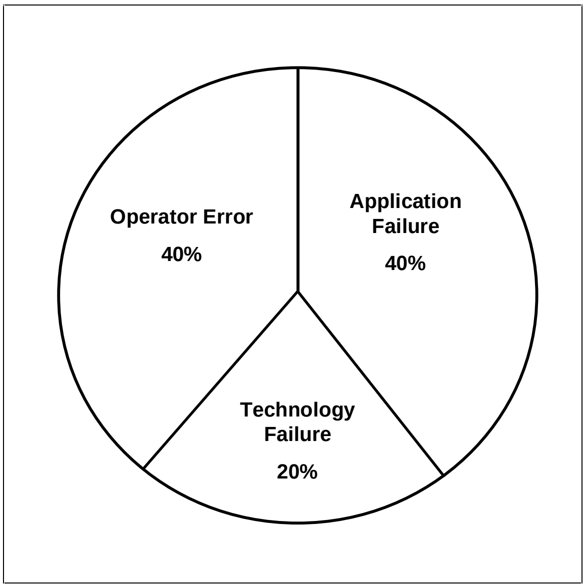

Figure 1. Causes of IT failures. First Study.

Source: D. Scott, Making Smart Investment to Reduce Unplanned Downtime (Gartner Group, 1999)

According to this study, only 20% of downtime is caused by technology failures including hardware (servers and network devices), environmental factors (e.g., cooling and power outages), and natural disasters. By contrast, 40% of downtime is caused by application failures that include bugs, operating system crashes, performance slowdown, incorrect updates, and changes to software. The remaining 40% of failures is due to human errors. In other words, according to these results, only 20% of downtime is caused by problems in the basic IT infrastructure, while the remaining 80% of failures are due to people and process issues, including software that provides for the business services and management operations (Kaye, 2001), (Noakes-Fry et. al., 2001), (Noakes-Fry and Diamond, 2002).

Another study, by IDC and Robertson Stephens, investigates the technological components that fail. The results are illustrated in Figure 2.

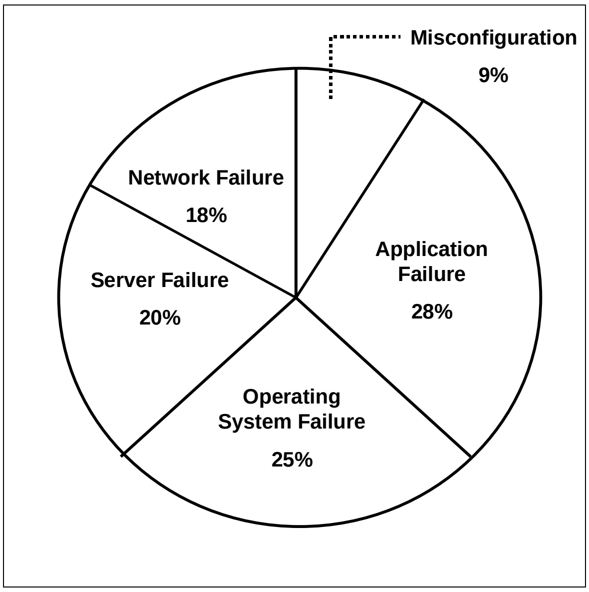

Figure 2. Causes of IT failures. Second Study.

Source: R. A. Juarez et. al., “Virtual Bricks II: Virtual Econ 101 Update, A Comprehensive Guide for Understanding eCommerce Infrastructure Evolution and Convergence”
(Robertson Stephens, May 2000)

According to this study, all the main components of the IT infrastructure, that is, network, servers, operating systems and the applications, fail with comparable frequency, with a predominance of applications and operating systems. Yet, mere human errors or unskilled staff that lead to mis-configurations will cause a significant amount of outages (Kaye, 2001).

While the two studies may appear contradictory, they are not, as they have adopted different classifications: the first study explicitly classifies human errors, while the second study includes them in the involved technological components. Merging the findings of the two studies, and reasoning on the different perspectives underlying them, can result in the statistics illustrated by Kaye (2001) as in Figure 3.

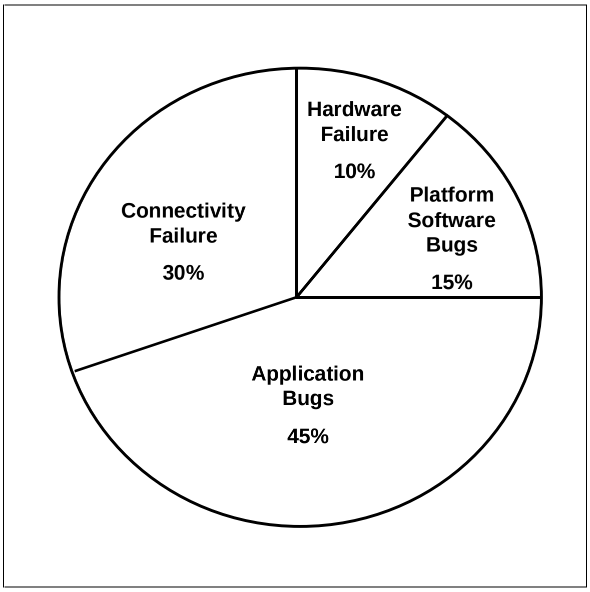

Figure 3. Causes of IT failures. Consolidated results

According to these consolidated results, the most frequent causes of outages are application bugs and connectivity failures. Applications frequently fail because of all sort of problems, including code bugs, bad design, weak requirement analysis, inconsistent tests, misconfigurations, erroneous installation or deployment, incompatibility with the run-time environment, and interference with other applications. As for connectivity failures, this last study estimates that misconfiguration of devices, routers, switches, or DNS, should be expected to be higher than the 9% estimated in the study reported in Figure 2. Other frequent causes of outages are network partitioning and Internet Service Provider (ISP) failures. It is interesting to note how hardware failures represent only a small portion of the risks. As a matter of fact, there are now many reliable hardware vendors in the market and the quality of their products has improved in the years. The suggestion from experienced disaster recovery planners is then to select products and vendors with good reputation and just be skeptical about the latest news from the industry because they often provide exciting new features but at the price of less reliability. Similarly, also the platform software (e.g., operating systems, DBMSs) is not one of the major risks, since products are generally reliable or could be configured and managed in order to maximize their reliability.

LOSSES, OR THE COSTS OF FAILURES

Internet-based companies, which require connectivity and availability of their web sites 24 hours a day, 7 days a week, 365 days a year, have suffered all sorts of system and/or network outages in the last few years.

The cost of a disaster or a service interruption is not only the cost of the loss of business in the downtime frame. There are many associated factors that produce costs and losses, including:

- Brand image recovery;
- Loss of share value;
- Loss of interests on overnight balances;
- Delay in customer accounting;
- Loss of control over debtors;
- Loss of credit control and increased bad debit;
- Delayed benefits of profits from new projects or products;
- Loss of revenue for service contracts;
- Cost of replacement of building, plant, equipment, software;
- Loss of customers;
- Loss of profits;
- Liability claims;
- Additional costs for advertising and marketing.

For instance, the brand value of most companies, while not reflected in their balance sheets, is one of their most valuable assets, reaching dozen of millions of dollars for some companies. Also, statistics have demonstrated that the share price of a corporation suffering a disaster falls by around 5% to 8% within the first few days after a disaster. Recovery of the share price depends on the efficiency in recovering critical business functions. Past cases have shown that efficient recovery has let companies regain the confidence of financial analysts soon, and their share values have rolled back to the prices preceding the disaster, and even increased by 10% to 15% in the following 100 days. Companies that recovered with many difficulties, instead, were penalized with more share values falls, around 15% in average (Hiles, 2002).

A classification of losses should include (Kaye, 2001):

- *Fixed-asset losses*. Fixed-asset for an IT infrastructure generally includes the hardware and software used in the corporate business working. Fixed-asset losses can be caused by computer failures, fire, and thefts. Usually these risks can be mitigated through warranties, service contracts, and insurances.
- *Third-party liabilities*. Depending on the type of business, a company can be exposed to third-party liabilities. These may be caused by an interruption in service provision or by the accidental disclosure of sensitive data after a security incident. Third-party damages are often covered by insurances.
- *Loss of intellectual property*. The loss of data, application code, and sensitive information related to business processes can severely affect the company business. Backup and recovery strategies can be used to mitigate these losses.
- *Losses of revenues and profits*. The loss of revenues and profits is the most evident type of risk in which a company may incur after a disaster or an outage. For companies that rely heavily on the Internet for their business, even short losses of connectivity or unavailability of their corporate portals can result in huge losses of revenue and profits.

As illustrated in Table 1, depending on the type of industrial sector, the cost of downtime can reach noticeable values.

Industry Sector	**Revenue/Hour**	**Revenue/Employee-Hour**
Energy	$2,817,846	$569.20
Telecommunications	$2,066,245	$186.98
Manufacturing	$1,610,654	$134.24
Financial Institution	$1,495,134	$1,079.89
Information Technology	$1,344,461	$184,03
Insurance	$1,202,444	$370.92
Retail	$1,107,274	$244,37
Pharmaceuticals	$1,082,252	$167.53
Banking	$996,802	$130.52
Food/Beverage processing	$804,192	$152.10
Consumer products	$785,719	$127.98
Chemicals	$704,101	$194.53
Transportation	$668,586	$107.78
Utilities	$643,250	$380.94
Health Care	$636,030	$142.58
Metals/Natural resources	$580,588	$153.11
Professional services	$532,510	$99.59
Electronics	$477,366	$74.48
Construction and engineering	$389,601	$216.18
Media	$340,432	$119.74
Hospitality and Travel	$330,654	$38.62
Average	$1,010,536	$205.55

Table 1. The cost of downtime.

(*Source*: IT Performance Engineering & Measurement Strategies: Quantifying Performance Loss, Meta Group, October 2000.)

These figures show how important it is for companies to have a good disaster recovery plan in place. Statistics have shown that a typical organization may spend three or more times its normal annual marketing budget in the aftermath of a disaster to retain customer confidence and to retain or regain market share.

THE DISASTER RECOVERY PLANNING PROCESS

Disaster recovery planning consists in a set of activities aimed at reducing the *likelihood* and the *impact of disaster events* on critical business assets (NIST, 2002), (Toigo, 1999), (Wold, 2002). Disaster recovery planning is as complicated as any major system development project, with many stages of analysis and an overall activity encompassing all critical assets, business functions, and technological solutions.

Phases of disaster recovery planning

A disaster recovery planning project begins by considering the *risks* and how to provide for their management so to make their occurrence tolerable to the organization (Noakes-Fry et al., 2001), (Noakes-Fry et al., 2002), (Toigo, 1999), (Wold, 2002). The disaster recovery planning, which needs senior management approval to start, consists in *i)* planning the actions to contain and recover from an emergency (e.g. procedures, teams, objectives) and *ii)* preparing the agreements and the procedures for the collaboration among the teams from different branches of the organization and with the external partners. The collaboration procedures can be defined by a cross-functional team that includes top management, a purchasing representative, IT team members, maintenance personnel, and software user departments.

Another cross-functional issue that the Disaster Recovery planner should define together with top management and the public relation staff concerns possible interviews and news coverage. The importance of this issue depends on the exposure of the organization and can then vary depending on the business sector. In certain cases (e.g. financial institutions, government agencies or Internet-based companies) misleading press information can result in extremely high losses, affecting brand image and share values, and imposing costs for subsequent advertising/marketing campaign. Consequently, the Disaster Recovery Plan should make sure that everyone in the organization understands who the media contact person is so that uninformed members of the organization do not make comments that are not in line with the public relations goals of the organization.

Once the project's development phase is completed, all procedures should be tested and personnel be should adequately trained. Testing helps identifying possible weaknesses in the plan and evaluating the ability of recovery teams to implement the plan efficiently. Plan testing for the IT infrastructure should include (NIST, 2002):

- System recovery on alternate platform from backup media or mirrored sites;
- Internal and external connectivity;
- System performance using alternate equipment;
- Restoration of normal operation;
- Notification procedures.

In addition to these IT specific actions, all procedures for physical risk containment and recovery should be tested, including:

- Personnel evacuation plan;
- Alternate facility activation and recovery.

Training should make personnel aware of the Disaster Recovery procedure and provide those teams more involved in handling an emergency with sufficient experience to be able to execute their tasks with minimal support of documentation (as documentation might be unavailable in the first moment of an emergency). Plan elements that could be subject of training include (NIST, 2002):

- Purpose of the plan;
- Cross-team coordination and communication;
- Reporting procedures;
- Security requirements;
- Team-specific and individual-specific processes.

When the plan is effective, it must be maintained, with periodic review, revision and maintenance. Figure 4 summarizes the phases of a typical disaster recovery planning project (Kaye, 2001).

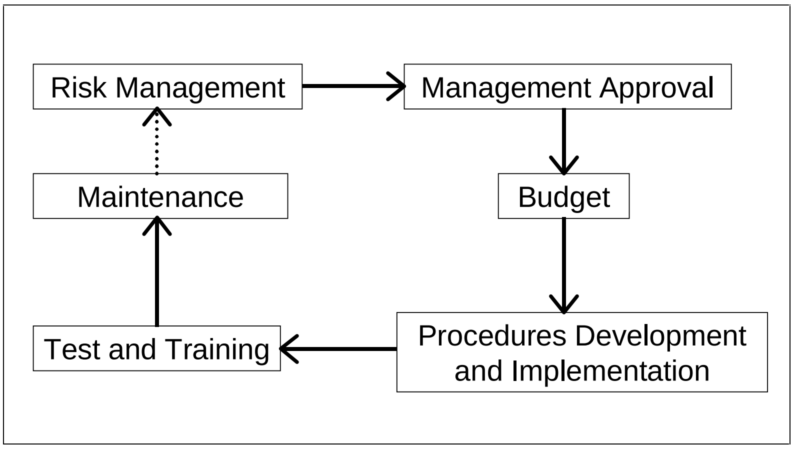

Figure **4.** Disaster Recovery Planning Project

A Disaster Recovery Planning should encompass all the IT infrastructure's components. In particular, all data repositories and all critical assets should be considered. More and more data are stored in a variety of supports and media, not just on mainframes and database servers; desktops and portable systems often hold important information (NIST, 2001), (Takemura and Taylor, 1996). IT platforms to be considered in the planning should then include:

- Desktop computers and portable systems;
- Servers;
- Web sites;
- Local area networks;
- Wide area networks;
- Distributed systems;
- Mainframe systems.

For each IT platform type, the *requirements* to recover each platform should be formally documented and considered in the analysis stage, and the *solutions* for addressing the event of a disaster should be detailed. The technical solutions should encompass *preventive* countermeasures (to avoid the failures) and *recovery* actions (to minimize the damage and reestablish the normal system operation in case of failures). Related documentation should describe:

- the frequency of backup and offsite storage of data;
- the production and safe storage of an up-to-date copy of all applications and operating systems, together with their configurations, patches and installation;
- the redundancy of critical components;
- the documentation of system installation, configuration, and requirements;
- the interoperability between system components and between local and remote site equipment;
- the requirements for power management systems and environmental controls.

Business Impact analysis

Convincing corporate management to invest in disaster recovery planning is often a greater challenge than dealing with the technical problems of backing up critical systems, minimizing the downtime, and maintaining the network connectivity. As it is often noticed, unfortunately, security investments usually receive low priority by managers, who prefer to invest in facilities and features that improve the business productivity, and realize the importance of protecting against disasters and system's disruptions only after they happen. To complicate matters, the management decision of investing in disaster recovery capabilities, is not a one time event, but must be sustained over the time, since every disaster recovery plan needs to be regularly tested, reviewed, and updated to be effective. Also, disaster recovery planning needs to be managed, improved, and updated with respect to the modification of the IT infrastructure and critical assets. For these reasons, the management consensus is vital for every disaster recovery plan that, in the best of circumstances, will provide for disaster avoidance, detecting and reacting to potential problems before they become disasters. Experienced disaster recovery planners usually observe that there are typical issues that should be justified to senior management and, among them, the most relevant is to carefully justify the plan's expenditures. Disaster recovery plans should be adequately documented with in-depth analysis of the risks addressed in the plan, the benefits for the corporation, the possible impact on business of a weak or insufficient disaster recovery plan, the estimates of the cost of downtime, and the possible reductions of insurance costs (Kaye, 2001), (Toigo, 1999).

The impact of a disaster from a financial and an economic standpoint (Hiles, 2002) can be determined by means of a *Business Impact Analysis* (BIA), whose objectives are to correlate each system component with the critical services it provides so to characterize the consequences of the disruption of the component. A BIA therefore should:

- *Identify critical IT resources.* Different components, applications, services, and processes can have a different importance on the system services and functionalities. It is therefore important to determine priorities on the system functions and consequently on the resources for their support. Identification of critical resources includes PCs and portable systems (especially those of top management and system administrators).

- *Identify disruption impacts and sustainable time window* in which recovery should take place. As we have already noticed, timeliness is a key factor in recovery. The BIA must then identify, for each resource, the maximum time window for which the resource can be denied and the possible cascading effects that the outage can have on other processes. This includes identifying financial and non-financial costs associated with a disruption.
- *Develop recovery priority.* Based on the criticality of resources and on the associated sustainable time window and cost consideration, recovery strategies can be prioritized so to optimize allocations and expenditures for the recovery process.

The BIA can provide input to the Risk Management process (see next section) on business risk that may not otherwise be identified. It can then be used to raise awareness of disaster recovery planning on senior management and focus individuals on their potential responsibilities, possible solutions, and costs.

THE RISK MANAGEMENT PROCESS

Risk is defined as the possibility of suffering harm or loss. Risk management is the process of planning for such eventuality. The goal of risk management is not to eliminate all risks (which would be impossible in practice), but rather to find an *acceptable balance between reducing risk and preparing the organization for losses*. The risk management process can be summarized in three steps:

- *Risk analysis*: *i)* identify the types of losses that may occur; *ii)* determine the causes and the likelihood that a loss will occur; and *iii)* estimate the resulting cost for the organization.
- *Risk mitigation*: *i)* identify actions to reduce or eliminate the likelihood of occurrence of each risk; *ii)* determine the cost of those actions; and *iii)* produce a cost/benefit analysis of the actions with respect to each risk and loss.

- *Risk transfer: i)* Address the losses that may occur in spite of preventive actions and *ii)* consider transferring them to insurers or third parties like service providers, vendors, or outsourcers in general.

We next examine each of these three steps in more details.

Risk Analysis

Risk analysis has two objectives:

1. *Identify business processes and their associated IT infrastructure resources*, such as the data, applications, systems, and networks that are used in supporting the corporate business. Business processes should be classified with respect to criticality and downtime;
2. *Identify existing risks* to the business processes and IT infrastructure.

Basically, risk analysis is concerned with identifying the potential losses and the risks that might cause them. As already discusses, risks include failures of connectivity, hardware and software, as well as security incidents and natural disasters (Kaye, 2001), (Toigo, 1999), (Wold and Shriver, 2002). All the types of risks that can cause system disruption, that is, natural, human, and environmental causes should be analyzed and prioritized in disaster recovery planning. For each failure cause, a disaster recovery planner must define approaches to manage the risk and estimate the trade-off between sustaining the possible losses and implement a disaster recovery strategy.

Risk Mitigation

Risk mitigation is a process aimed at limiting the likelihood of risks and the potential losses that risks can cause (Kaye, 2001). As sketched in Figure 5, the process is composed of the following phases.

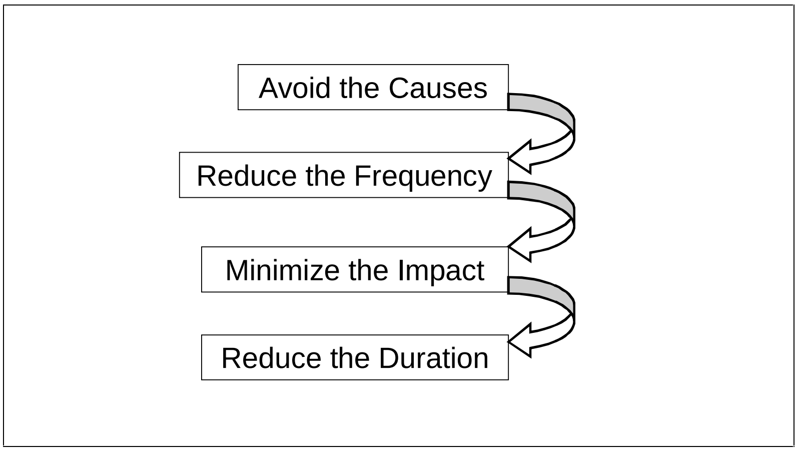

Figure 5. Risk mitigation process.

- *Avoid the causes*. Make conservative and prudential technology choices for critical assets: avoid technology risks by developing services in a robust manner and with the support of reliable platforms instead of adopting the latest technology (usually riskier), unless it has proven to give a sensible competitive advantage.
- *Reduce the frequency*. All IT infrastructure components should be chosen among those ensuring a low rate of failure. Operating systems and software/hardware platforms of different vendors (or different versions produced by the same vendor) often have different failure frequency. Reliability, and not just performance and features, should be carefully evaluated when technology is purchased to enforce the risk management strategy.
- *Minimize the impact*. Since the frequency of failures and outages can never be reduced to zero, a risk mitigation strategy should analyze all possible single points of failure, whose failure may lead to major losses in terms of downtime and extent of service interruption. Redundancy is the typical strategy that can be applied to servers, network devices, and internal components.
- *Reduce the duration*. Minimize the duration of downtime is the ultimate goal of every disaster recovery plan: the more the downtime, the more the revenue loss. Many strategies can be adopted to speed up the recovery of data and functions. These may involve on-site

hardware spares, on-site staff, efficient backup solutions, and alternative operative locations.

Risk Transfer

Risks that cannot be sustained by a company can be transferred to a third-party which assumes the risks in exchange of a service fee (Kaye, 2001). Hiles (2002) estimated that, in average, up to 40% of actual losses could be transferred, but no more than that. Risk transfer can assume two forms: *insurance* or *outsourcing*.

Insurance is the traditional way of transferring risks that cannot be mitigated with in-house solutions or transferred by purchasing the service outside. Most insurers offer coverage even for risks related with Internet-based threats like intrusions or denial of service attacks. With *outsourcing*, a third-party provides a specific service in place of an in-house solution. This third-party can then assume the risk of losses that the company may incur in, refunding the fee paid by the company (as in case of web-hosting services) or reimbursing the losses (as in case of vendors). Risk transfer is established via a contract that usually includes a Service Level Agreement (SLA) that should guarantee the company that purchases the service. This is a suitable solution in many situations, but the trade-off between the cost of possible downtime, if the service is realized in-house, and the cost of the purchased service with a certain quality level should be carefully evaluated. A careful evaluation of the return on the outsourcing investment should therefore be done in order to figure out whether the strategy is convenient (downtime costs vs. service fees). Consider, for example, the case of Web Hosting Services, where outsourcing is the common choice today. When a critical service is purchased from a vendor, instead of managing it in-house, it is typical to sign an agreement concerning the level of service guaranteed. In case of networked services, web hosting, mainframe hosting and so forth, the parties agree on a certain uptime guarantee, expressed as a percentage of time of availability (usually a value between 99% and 100%) of the purchased service. Clearly, the more the uptime guarantee, the more the cost of the service. But, how should a company decide which is the proper uptime guarantee value?

Kaye (2002) shows, as illustrated in Figure 6, the typical relationship between the amount paid for a certain service level and the value (i.e., the benefit that the company receives) of the increasing percentage of uptime guarantee (or, conversely, the decreasing downtime).

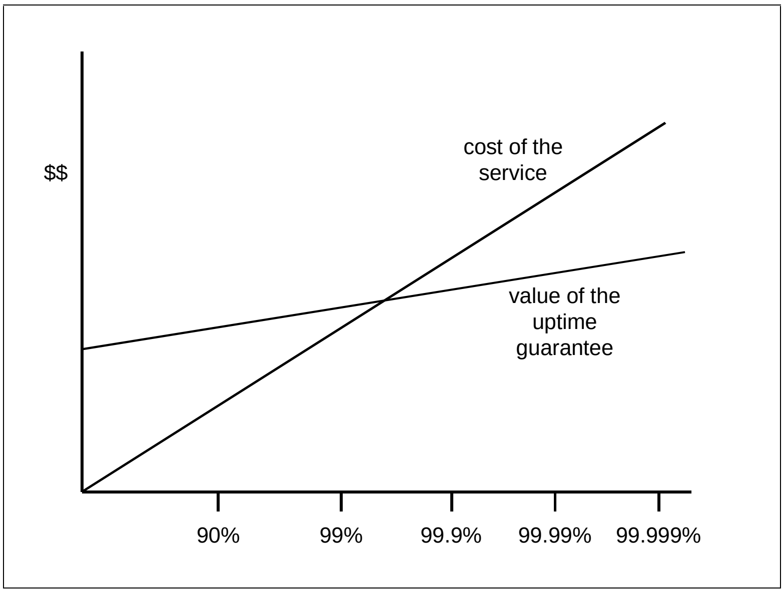

Figure 6. Uptime value vs. Cost

This reference relationship is meant to say that generally, after a certain percentage of uptime guarantee, the cost that a company would pay for a given service level is not worth the benefit gained. In fact, the revenue losses would be less than the cost paid and the investment would not yield a positive return.

Let us discuss this more with an example taken from Kaye (2001). Consider a company with $50 million in revenues and 15% pretax net profit. Increasing from 99% to 99.9% uptime will lead to increase revenues by $183,960 per month and profits by $27,594. If the cost of the service is less than $27,594, then the investment is worthwhile. If the company wants to upgrade from 99.9% to 99.99%, than the increase in monthly revenue is just $18,396 and $2,759 for profits. Hence, in this case the raise of the service cost should be much smaller, less than $2,759, to be convenient.

The example, clearly, is over-simplified since in practice a company should consider its business strategy and all the costs involved (e.g., loss of competitive advantage, market positioning, reputation, third-party liabilities) in order to evaluate the investment. However, the decreasing value of upgrading the service level is usually confirmed in most situations and should be carefully evaluated, possibly with formal models, to establish the return of the investment (Toigo, 2001b).

BACKUP AND RECOVERY

The goal of disaster recovery is to provide the ability of recovering the critical business functions of an organization in the occurrence of a disaster. One of the most important requirement for an effective recovery is the *availability* of *all the company's data* as well as of *the configuration of all critical components of the IT infrastructure* (e.g. backup of servers and network devices configuration, patches, passwords, routing tables and firewall's rules). While the need of maintaining backup copies of the data is well recognized, the need of a copy of the system configuration, in order to effectively restore the functionality of the IT infrastructure, is often overlooked. The reason why this requirement is so frequently overlooked is due to its intrinsic organizational nature: system management and day-to-day administration of the IT infrastructure (like change in routing tables or firewall rules, installation of patches, etc.) must be fully documented. Also, procedures should be put in place to make sure that every modification to system components (the critical ones, at least) is reported in a backup copy. Effective recovery requires also that *media for custom and purchased software* together with their *license information* should also be kept offsite. Last, but not least, the Disaster Recovery Plan itself should be backed up. It may seem obvious that the Disaster Recovery Plan will be available when an emergency calls for its activation, but this means that a complete copy of the disaster recovery plan should be kept in a safe place (e.g. off line at an alternate site). The need to back up all the information and elements needed to fully restore a system applies even when a comparable system is kept at a hot site: the hot site should maintain the exact level of operating system patches and operating system "state".

Some considerations should be devoted to backup requirements for desktop computers and portable systems, which have some peculiarities with respect to servers. In particular, when the PC data backup process is not automated (e.g. saving data on a fileservers regularly backed up), which is frequently the case, users duty of saving critical data on a backed up server should be ruled by an appropriate Usage Policy (i.e. a documentation possibly subscribed by each user stating, among the others, that each user should care of saving his/her critical data on a backed up server and not locally). The policy should be enforced by intermediate managers and users should be trained about the importance of this recovery requirement. Providing an automatic backup process and restoration for PCs requires *interoperability and platform standardization*. It is then usually recommended to have homogeneous platform operating systems, configurations and applications among the organization's PCs and portable device (NIST, 2002).

Backups should be kept in a safe place, where they cannot be affected by the consequences of a disaster, so to enable their *recovery* if the original is compromised or unavailable. This has been for decades the golden rule of disaster recovery, especially for mainframe-centric IT infrastructures. Traditionally, backup has been seen as a batch process running during night hours and storing on tapes all relevant information produced that workday. While this solution is still used in a large number of organizations and is well suited in many circumstances, in many others cases this scenario is now anachronistic. With the increase of Web-based services and extranets, applications and business process tend to be active on a 24x7 basis. As a consequence, there is no longer the notion of "night hours", which was meant to be that window of inactivity of business transactions perfectly suited to take a snapshot of data. In many real world scenarios, transactions are always running, information is always accessed and possibly updated. In addition, data are growing every year at a terrific rate. Many commentators are commonly referring to the ever growing amount of data to manage, backup and restore as "the explosion of data", to stress both the impressive growth and the chaotic nature (Campbell, 2000), (Kaye, 2001), (Nicolett and Berg, 1999), (Toigo, 2001a).

Backup and recovery planning should achieve two fundamental goals:

- *Backup* should be made with a minimal impact on the production environment and, ideally, should be performed without blocking transaction execution.
- *Recovery* should be efficient so to minimize the time frame necessary to restore the data and bring the system back to its operational state. The efficiency of the recovery process is a key factor for the company effective restoration of its business functions as a delay in recovery could cause a severe loss of market shares.

Backup and Recovery Techniques

The definition of data recovery strategies requires effective backup strategies to be in place. These in turn, require the analysis of data to defining priorities and criticality, and to provide means to perform the backup process while maintaining the system operating. The classification of business-critical data is needed to better organize their backup, optimize recovery and therefore minimize downtime that should follow the evolution of the company's business processes and the corresponding flows of information. Table 2 illustrates a simple typical data classification schema (Kaye, 2001).

Classification	Definition
Critical	Data that must be preserved for legal reasons, for use in business critical processes. Must be restored as soon as possible in the event of a disaster.
Vital	Data needed for important, but not critical, business processes. A loss of this data would represent a severe damage and economic loss for the company. This data may have privacy requirements.
Sensitive	Data used in normal business operations, data that could be reconstructed from alternative sources in the event of loss but at some cost.
Noncritical	Data that could be reconstructed with minimal costs in the event of loss. No privacy or security requirements.

Table 2. Example of Data Classification for Backup and Recovery.

To better maintain a clear identification and classification of data, a formal policy should be adopted. The goal of the policy is to:

- Establish who is in charge of formally classifying data for each business unit, department, or workgroup.
- Define a uniform classification methodology for the whole organization corresponding to the recovery priorities and requirements in the event of a disaster.

Some of the main strategies and technologies for backing up data are summarized by Toigo (1999) as follows:

- *Server Image Backup*. Image backup creates a physical image of an entire disk. It operates at physical disk level by transferring sectors of physical data storage blocks from hard disk to tape. It has fast transfer rates and a provide complete server backup by transferring the entire content of the hard disk as a single unit. This technology requires the system to be inactive during the backup.
- *Snapshot or Versioning Backup*. This technique is fast because after an initial complete image backup, it proceeds incrementally by backing up the modifications only. It works with

block-level copies on a partition basis. It is widely used in traditional tape-based backup systems.

- *Full Volume Backup on a File-by-File Basis*. Complete datasets are copied over the production network to backup devices. A process controlled by the server initiates the transfer to backup devices.
- *Incremental or Differential File Backup*. Files are compared with their version in the preceding backup and only changed files are backed up again. *Differential* backup is meant to keep separate versions of same backed up files, while *incremental* backup just keeps the newer version.
- *Backup using Object Replication*. Logical objects are defined by software, including disk partition tables, boot volume, security information, system volume, and user data volumes. Logical objects are regarded as single logical unit and all data can be copied at once.
- *Electronic Tape Vaulting*. It replaces the manual procedures for handling tapes stored in offsite storage facilities and for moving tapes between locations. It exploits a Wide Area Network connection between the company facility and the off-site storage facility where tapes are stored to provide for immediate access to data stored on tapes. Technologies like shadowing and multiplexing controllers allow duplicating backup data streams in two or more tape devices simultaneously. This way a company could make both a local backup (for complete recovery in the event of a disaster) and an electronic vaulting on a remote storage facility (to preserve data from disaster, like flood or fire, affecting the company site).
- *Remote Disk Mirroring***.** It is another alternative to physically handle tapes. With Remote Mirroring, the content of disks is duplicated in near real time to a remote disk volume. The same concept is applied to RAID (Redundant Array of Inexpensive Disks) technology as a means of fault tolerance instead of disaster recovery.

Selecting an *offsite storage* facility is the safest choice for preserving backups. The remote storage facility should be chosen to be unaffected by the same disasters that might affect the main company site. That site could be a branch of the company or a commercial system recovery facility offering this service. While companies with decentralized sites could choose to keep backups among their different sites, several commercial storage facilities offer backup storage as a service, including media transportation, safe storage, and recovery (NIST, 2002), (Nicolett at al., 1999), (Toigo, 1999). When selecting an offsite storage facility and vendor, there are some key factors to keep in mind:

- *Geographic area.* Distance from the organization and the possibility of the storage site being affected by the same disaster as the organization (locations geographically close to the organization are likely to be affected by the same natural disasters).
- *Accessibility.* Length of time to retrieve the data from storage and the storage facility's operating hours.
- *Security.* Security capabilities of the storage facility and employee confidentiality, accordingly to the data's sensitivity and security requirements.
- *Environment.* Structural and environmental conditions of the storage facility, like temperature, humidity, fire prevention and power management control.
- *Cost.* Cost of shipping, operational fees, and recovery services.

ENSURING CONTINUITY OF OPERATIONS: ALTERNATE SITES STRATEGY

Disaster Recovery plans have traditionally accounted for natural disasters as the first threat to an organization. Hence, beside the safe storage of data backups, it has been necessary to provide an organization with a plan for the continuation of its business functions while severe damages occurred to its site and hardware assets. A disaster recovery plan should then include a strategy to recover and perform system operation at an alternate facility. Generally, three types of alternate site strategies are possible (NIST, 2002):

- Dedicated site owned or operated by the organization;
- Reciprocal agreement with an external entity (e.g., a business partner);
- Commercial leased facility.

As prospective alternate sites are evaluated, the Disaster Recovery planner should ensure that the system's security, management, operational, and technical controls are compatible with the prospective site. Controls may include measures for network security, such as firewalls and intrusion detection systems, and physical access controls. If the service is purchased from a vendor, a Service Level Agreement (SLA) must be negotiated carefully. The SLA must clearly state all conditions like testing time, workspace, security requirements, hardware requirements, telecommunication requirements, support services and recovery time frame. The Disaster Recovery planners should be aware that a commercial facility usually hosts many customers, which may all be simultaneously affected by the same natural disaster. Hence, the commercial facility could be asked to recover different customers simultaneously and it may not have enough resources to do it properly. In this case, the commercial facility would tend to set priorities among its customers and not treat all of them at the same level of service. Unfortunately, this scenario is realistic and for this reason it is extremely important to negotiate a clear and extremely detailed SLA with the commercial vendor, in order to set once and for all how possible disasters should be addressed and how priority status is determined. Regardless of the type of alternate site strategy, the alternate facility could provide different services and have a different level of readiness with respect to the goal of immediate resumption of all business functions. There is a typical classification that also reflects the different commercial services offered on the market (NIST, 2002), (Nicolett et al., 1999), (Toigo, 1999):

- *Cold Site*. It typically consists of a facility with adequate space and infrastructure (electric power, telecommunication connections, and environmental control) to support the IT infrastructure. The cold site contains neither IT equipment nor office automation equipment, like telephones, fax machines, or copiers. The organization leasing a cold site should provide for all the necessary equipment.
- *Warm Site*. It is a partially equipped office space containing some hardware, software, telecommunication, and power sources. The warm site is kept in an operational status to be promptly made fully operative in the event of a disaster at the company's main site.
- *Hot Site*. A hot site is a fully functional and immediately ready facility equipped with all the needed hardware, software, telecommunication, and power sources. Hot sites are usually staffed 24 hours a day, 7 days a week. Recovery of business functions is guaranteed to be extremely fast.
- *Mobile Site*. It is a self-contained, transportable shell equipped to restore immediately the subset of the most critical business functions, including telecommunication facilities. Major commercial vendors lease these mobile sites and keep them located in strategic places to let them reach most of the industrial plants in a relatively short time. However, to guarantee an on-time recovery by means of a mobile site, a company should sign agreements with vendors to have the mobile site ready in a time frame compatible with the business requirements.
- *Mirrored Site*. It maintains an almost exact replica of the IT infrastructure of the company's main site and is always ready to recover all the business functions. Recovery is almost immediate through a mirrored site. It is usually owned and operated by the company, given the high degree of customization and specificity.

Table 3 summarizes the characteristics of the different alternate sites.

Site	Cost	Hardware Equipment	Telecom.	Setup Time	Location
Cold	Low	None	None	Long	Fixed
Warm	Medium	Partial	Partial/Full	Medium	Fixed
Hot	Medium/High	Full	Full	Short	Fixed
Mobile	High	Dependent	Dependent	Dependent	Not Fixed
Mirrored	High	Full	Full	None	Fixed

Table 3. Alternate site characteristics

(*Source*: Contingency Planning Guide for Information Technology Systems, NIST Special Publication 800-34, June 2002)

The above solutions differ in terms of cost and timeliness. Cold sites are less expensive to maintain but require considerable time to acquire and install the necessary equipment. Going down in the list, we find solutions that result more expensive but that provide greater availability. Mirrored sites are the most expensive solution and provide almost 100% availability (i.e., minimal downtime). Remote mirroring has gained large popularity as a reliable solution for disaster recovery, we therefore discuss it in more detail next.

Remote Mirroring

Advancements in networking technology have put the basis for many commercial systems based on remote mirroring for safeguarding corporate data and providing fast recovery. However, sometimes the expectations that planners put in remote mirroring solutions are not realistic. Also, solutions based on remote mirroring tend to be complex, both architecturally and in data management, to satisfy to the high requirements from companies. This may lead to a lack of robustness of the recovery system, misconfiguration, and improper array maintenance. The consequence could be the loss of synchrony between arrays or the storage of corrupted data. If this happens, in the event of a recovery, the result could be a multiplication of the disaster, since the risk is to recover corrupted or incoherent data (Kaye, 2001), (Toigo, 2001a), (Toigo, 2001b). Also, remote mirroring does not provide for zero downtime, which is considered a requirement by some senior managers. Remote mirroring, if configured properly and implemented with high-level technology, could be very fast, much faster than many other recovery alternatives, but it cannot reduce the downtime to zero. Hence, remote mirroring is not the panacea that makes the nightmare of downtime disappear: risk management and disaster recovery planning are still mandatory. Moreover, remote mirroring is subject to data gaps, since some sort of asynchrony is usually implemented, (e.g., the mirroring between the secondary and the tertiary arrays of Figure 7). Possible consequences of incoherent recovered data should be carefully analyzed in the planning stage. These may provoke the loss of commercial orders, payments, or shipments, if data gaps are not identified and handled properly. Solutions designed to be completely synchronous exist but come at high prices in terms of performances and system latency. In fact, the problem in remote mirroring has been traditionally the latency caused by the distance, possibly many hundreds of miles, between the sites, in particular when the mirroring is synchronous (i.e., writes to production disks waits for the mirrored writes to be completed and acknowledged). Starting from mid-90s, the cost of high-speed bandwidth WAN decreased and vendors realized architectures that made remote mirroring one of the best solution at affordable costs for many companies. The typical improvement was to provide for three arrays, adding an intermediate site between the company and the final remote storage facility (Toigo, 1999). Figure 7 shows the traditional and the enhanced schema for remote mirroring. The benefit of the enhanced schema is twofold. First, the link between the company site and the secondary array could be a high-speed serial one, thus reducing the latency in synchronous mirroring. Second, the mirroring between the secondary and the final arrays, connected with a WAN, could be realized asynchronously, hence optimizing the mirror operation, while not affecting the robustness of the disaster recovery solution.

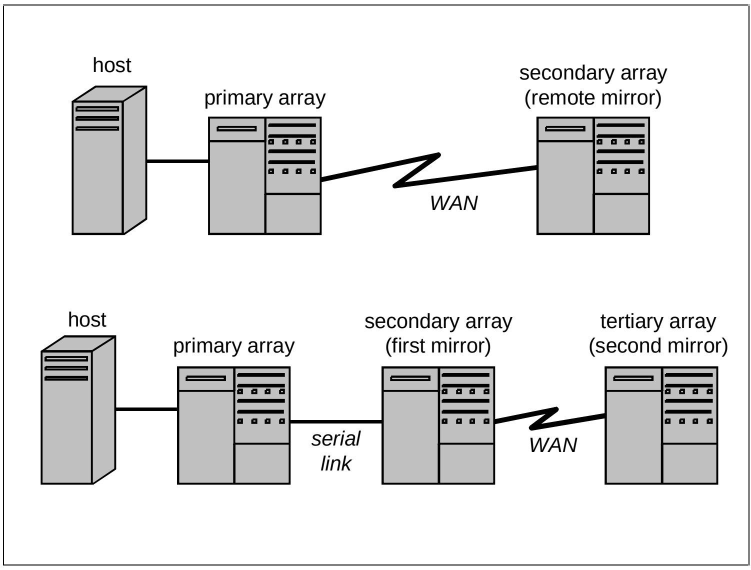

Figure 7. Remote mirroring configurations.

BACKUP AND RECOVERY FOR WEB-BASED HOSTING SERVICES

Web Hosting Service (WHS) providers are commercial vendor that may store hundreds of web sites in their infrastructure. Moreover, the service effectively hosted may be complex, not just a web site. WHS can be composed of n-tiers architectures with web portals, application servers, middleware layers, and community services. Due to this increasing complexity, nowadays it has become common for an organization to purchase the web management and hosting service instead of implementing it in-house. However, like other outsourced services, even for web hosting, customers should evaluate carefully how the service is provided and consider which guarantees the vendor provides in the event of an outage.

The primary deployment options for Web-based applications are *Internet data centers*. A growing number of enterprises are considering dual-site application architectures that split application traffic between two active sites. The active sites could be corporate data centers, commercial data centers, or both. One aspect to consider carefully when services from commercial data centers are purchased is the large number of simultaneous recoveries they should provide for all their customers in the event of an outage. Thus, similarly to what we observed for hot site vendors, the customers of a web hosting service should make sure that the outsourcer has a robust disaster recovery plan and, through clear Service Level Agreements, that she has a contractual responsibility to effectively execute it (Kaye, 2001), (Toigo, 2001b).

Different technological options can be used for storage management of web sites and Internet applications (Kaye, 2001):

- *Direct Attached Storage (DAS)*. The disk drives or arrays are directly connected to an individual server.
- *Network Attached Storage (NAS)*. High-performance general-purpose file servers connected to web-sites via LANs.
- *Storage Area Networks (SAN)*. High-performance special-purpose storage systems connected to web servers via dedicated fiber optical links.
- *Global Storage Systems (GSS)*. A technology for geographically replicated file systems.

Table 4, from Kaye (2001), summarizes the principal characteristics of each storage solution.

Capacity	For Static Content	For Shared Reads/Writes	For Databases
DAS <= 200 GB	Yes (very few web servers)	No	Yes (good reliability using RAID)
NAS <= 10 TB	Yes (more web servers)	Yes	Yes (very good reliability if clustered)
SAN 100 GB-100 TB	No (too expensive)	No	Yes (high reliability)
GSS 300 GB-100 TB	Yes (more locations)	Yes	No

Table 4. Web-hosting storage technologies.

Networked storage solutions, however, can pose special difficulties that significantly degrade the speed of most tape-based solutions. For instance, in a SAN, domain servers, routers and software delivering virtual volumes to SAN servers often manage physical disk devices. This networked infrastructure poses an increased management burden on the IT staff and is more prone to failures than simpler DAS.

Most e-business applications hosted in Internet data centers are based on n-tiers architectures. These multi-tiers platforms typically include one or more Web servers, applications servers used to integrate Web technology with legacy systems, connectors and adapters, database servers storing raw data, and authentication and authorization servers filtering the user accesses and incoming connections. This complex architectural design and the interconnection of several components is clearly more prone to failures than centralized monolithic systems. These failures, often implying short but frequent outages, must be handled carefully in the planning stage or in negotiating SLAs with vendors. Selection of components and replication strategies should be required with the goal of minimizing downtime in mind.

In addition to this class of system failures, e-business needs to face the growing rate of network attacks and denial of service. Many commentators are now commonly talking about “the convergence of Disaster Recovery and Network Security” for e-business applications, meaning that the two areas have many in common and more and more interdependencies. From a Disaster Recovery planning viewpoint, intrusions and security incidents are another cause of downtime to be considered in the risk management stage, in the BIA, and in the selection of vendors. Internet data centers should provide high-level staff skilled in network security and provide for robust security architectures and policies (Kaye, 2001), (Toigo, 2001a), (Toigo, 2001b).

DISASTER RECOVERY PLANNING TEMPLATE

This section provides a reference template for Disaster Recovery Planning derived from Brooks et. al. (2002). It is only intended as one possible example of Disaster Recovery Planning template, since an actual plan should meet the peculiar requirements of every organization. Further examples and interesting readings for the definition of a Disaster Recovery Plan could be found at Toigo (1999) and DRJ (2002).

A plan could have the following schema:

Introduction

Purpose: The purpose of the plan should be shared and fully understood by the personnel involved. For example, the plan recall that in the event of a disaster, the plan provide for precise procedures to activate the plan, to recover from the disaster and to restore normal operations. Other key actions to be introduced are: the identification and classification of assets and services, the assignments of responsibilities and roles to personnel to be carried out during the interruption of normal operations and the coordination required with other staff involved in the recovery phase, e.g. vendors, local or government agencies.

Scope: The scope of the plan describes which measures the organization has put in place in the event of a disaster. For example, which alternate site strategy has been set, whether a hot site is ready to recover critical operations, how long the alternate facility could be effectively used to backup the IT infrastructure and how long should be the estimated unavailability of the main IT functions to activate the alternate facility (this is strictly dependent on the business sector so the range could stem from 48 hours or more of inoperability to immediate switching to a mirrored site in case of a failure).

Assumptions: The introduction could even highlight what is assumed to be fully operational and be organized for the success of the plan. For instance, it should be assumed that preventive safety measures have been followed, in particular in case of a natural disaster. Thus, the sprinkler system, fire extinguishers, generators etc. must be fully operational and properly maintained. Assumptions related to the IT infrastructure usually refer to the availability of UPS (uninterruptible power supply) to backup a temporary power failure and to the safety of data backups in a site not affected by the disaster. Other assumptions could be related to the effectiveness of SLAs with service providers, e.g. which actions are supposed to be carried out by the providers, which resources should be available, how communication should be done, etc. Of overall importance is the availability of a consistent copy of the configuration of all critical components of the IT infrastructure, which must be backed up in a safe place.

IT Infrastructure and Responsibilities

System description and architecture: A clear and comprehensive documentation about all the critical components and functions of the IT infrastructure is necessary to speed up the recovery after a disaster. The IT architecture should be described with details about: physical location of critical servers and network devices, functionality of each component, LAN/WAN/SAN topologies and telecommunication connections. Also, for each critical component all the information regarding, their operating systems and applications should be immediately available. The plan should indicate the location where the backup copies, with software and actual configurations, are safely stored.

Responsibilities: Key roles in the organization's chart in the event of a disaster should be indicated, as long as the line of succession (e.g. if the Chief Information Officer is unable to function, his/her functions must be delegated to someone else). Aside this, the plan describes roles and responsibilities of all the teams specifically dedicated to contain the effect of a disaster, organize the recovery of critical functions, and restore the normal operations. For each team, its members, leadership, roles, responsibilities and coordination with other teams should be indicated.

Activation

Damage assessment: The first duty of a Disaster Recovery plan is to contain the damages in the initial phase of an emergency. The top priority, especially in the event of a natural disaster, is the safety of the personnel; hence proper evacuation plans and directives must be enforced. Then, the plan should state detailed procedures to determine the cause of the disaster, the extent of damages, affected physical area and the status of physical infrastructure and of the IT equipment. After a first inspection, a Disaster Recovery team should proceed to a detailed inventory, following the priorities set for the organization's assets.

Criteria for the activation: Based on the results of the damage assessment phase, the Disaster Recovery coordinator could decide immediate actions like notify civil emergency personnel, the provider of the alternate facility or the vendors for replacement items. The plan should define criteria for the activation. For example, how long the facility is estimated to be unavailable and how long critical components of the IT infrastructure cannot be replaced. If the plan is to be activated, the coordinator should inform all team leaders, which in turn should activate all team members.

Recovery

Priorities: The key factor in the recovery phase is to follow priorities. Unless the disaster was confined to a single or few components, the recovery should be focused on the most critical assets and business functions. The goal of restoring as fast as possible those functions necessary for the organization business should be clearly stated and supported. Priorities has been already set during the Risk Analyses and the Business Impact Analyses. The results of these two key phases of the Disaster Recovery Planning should be reflected in primary, secondary and tertiary recovery objectives that the disaster recovery teams should pursue.

<u>Return to Normal Operation</u>

Original or Alternate Site Restoration: This phase of the plan presents the activities required for restoring operations at the original or at the alternate site. Technical details may include: which teams operate first and on which components, coordination activities among teams (for example between network and system administrators, between system administrators and application administrators, and or between personnel of organization and personnel of the provider). Testing procedures could be also described.

Deactivation: Procedures should outline how to return to the normal activity, like: how gracefully deactivate the alternate site or temporary measures put in place, and restore a proper backup system and the safe storage of backup data.

<u>Appendices</u>

Documentation: The appendices of the plan should include all the detailed documentation mentioned in the plan itself and needed during an emergency. For example: Personnel contact list, vendor contact list, equipment and specifications, SLAs, IT standard operating procedures, BIA, evacuation plan, equipment inventory and configuration data.

CONCLUSIONS

The scope of Disaster Recovery has become wider in recent years than in the past due to a change in business functions of many enterprises. Internet connections have driven the development of new business models which require immediate recovery and continuous operations and have the “zero-downtime” as their ultimate goal. While the “zero-downtime” goal is still impossible to achieve, and this should be stated clearly in each Disaster Recovery project, the Disaster Recovery area has developed new planning methodologies and techniques to match these requirements. As a consequence, the process of putting in place an effective Disaster Recovery plan has now sophisticated technologies to exploit but it is also more complex for its interconnections with both the IT infrastructure and the business functions of a company. Disaster Recovery plan could result in severe losses when one of the many risks becomes reality. Such risks today are not only the traditional disasters like fire, floods or sabotage, but also the “normal” failures of the IT infrastructure, such as misconfigurations, failures of applications or security incidents, that may lead to frequent downtimes. E-business models can no longer sustain those frequent downtimes or network outages like in the past when companies could accept recovery-time of some days in the event of a IT failure. Processes that allow achieving a better efficiency of the recovery phase in the event of a disaster, such as the Risk Analysis, the Business Impact Analysis and the innovative Backup and Recovery techniques assume therefore great importance.

BIBLIOGRAPHY

Campbell, R. (2000). Continuity Planning In The New Millenium - The Convergence of Disciplines. Disaster-Resource.com. Available at:
http://www.disaster-resource.com/cgi-bin/article_search.cgi?id='16'
(Date of access: October 10, 2002).

Hiles, A. (2002). Business Impact Analysis: What's Your Downside?. Rothstein Associates Inc.;. Available at:
http://www.rothstein.com/articles/busimpact.html
(Date of access: October 10, 2002).

Brooks, C., Bedernjak, M., Juran, I. and Merryman, J. (2002). Appendix A. DR and Business Impact Analysis Planning Templates. In Disaster Recovery Strategies with Tivoli Storage Management. IBM Redbook SG24-6844-00. Available at:
http://www.redbooks.ibm.com/redbooks/SG246844.html
(Date of access: October 10, 2002).

DRJ (2002). DRJ's Sample DR Plans and Outlined. Disaster Recovery Journal. Available at:
http://www.drj.com/new2dr/samples.htm
(Date of access: October 10, 2002).

Kaye, D. (2001). Strategies for Web Hosting and Managed Services. John Wiley & Sons.

NIST (National Institute for Standard and Technology) (2002). Contingency Planning Guide for Information Technology Systems. NIST Special Publication 800-34. Available at:
http://csrc.nist.gov/publications/nistpubs/800-34/sp800-34.pdf
(Date of access: October 10, 2002).

Nicolett, M. and Berg, T. (1999). Storage Backup and Recovery for Distributed Systems. Strategic Analyses Report R-09-3148; Gartner Group.

Noakes-Fry, K. and Diamond, T. (2001). Business Continuity Recovery Planning and Management: Perspective. Technology Overview DPRO-100862. Gartner Group.

Noakes-Fry, K. and Diamond, T. (2002). Business Continuity Planning Software: Perspective. Technology Overview DPRO-100469. Gartner Group.

Scott, D. and Natis, Y. (2000). Building Continuous Availability into E-Applications. Research Note COM-12-1325. Gartner Group.

Takemura, R. and Taylor, R.M. (1996). The Increasing Need For Client/Server Contingency Planning. Disaster-Resource.com. Available at:
http://www.disaster-resource.com/cgi-bin/article_search.cgi?id='36'
(Date of access: October 10, 2002).

Toigo, J.W. (1999). Disaster Recovery Planning: Strategies for Protecting Critical Information Assets (2nd Edition). Prentice Hall PTR.

Toigo, J.W. (2001a). Fighting Fires on the Web. Enterprise Systems. Available at:
http://esj.com/features/article.asp?EditorialsID=60
(Date of access: October 10, 2002).

Toigo, J.W. (2001b). Storage Disaster: Will You Recover?. Network Computing. Available at:
http://www.networkcomputing.com/1205/1205f1.html
(Date of access: October 10, 2002).

Wold. G.H. (2002). Disaster Recovery Planning Process - Part I, II, and III. Disaster Recovery Journal. Available at:
http://www.drj.com/new2dr/w2_002.htm, http://www.drj.com/new2dr/w2_003.htm,
http://www.drj.com/new2dr/w2_004.htm
(Date of access: October 10, 2002).

Wold, G.H. and Shriver, R.F. (2002). Risk Analysis Techniques; Disaster Recovery Journal. Available at:

http://www.drj.com/new2dr/w3_030.htm.

(Date of access: October 10, 2002).