



SANS Institute Information Security Reading Room

The Disaster Recovery Plan

Chad Bahan

Copyright SANS Institute 2021. Author Retains Full Rights.

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

The Disaster Recovery Plan

by

Chad Bahan

June 2003

GSEC Practical Assignment version 1.4b

Option #1

© SANS Institute 2003, Author retains full rights.

Disaster Recovery Plan Table of Contents

Introduction	3
History	3
Relationship to the Business Continuity Plan	4
Disaster Recovery Process	4
Disaster Recovery Plan	6
IT Disaster Recovery Planning Process	6
Identify Applications	7
Identify System Recovery Time (SRT)	8
Identify Data Currency for the Applications	8
Identify Critical Personnel and Recovery Teams	9
Testing the Disaster Recovery Plan	9
Maintaining the Disaster Recovery Plan	10
Summary	10
References	12

© SANS Institute 2003, Author retains full rights.

Introduction

The terrorist attacks on the United States on September 11, 2001 are focusing the attention of organization decision makers on the urgent need to prepare for disaster recovery. The Business Continuity Plan (BCP) describes the steps an organization takes when it cannot operate normally because of a natural or manmade disaster. It may be written for a specific business process or may address all mission-critical business processes. Business continuity and disaster recovery are critical components used to ensure that systems essential to the operation of the organization are available when needed. Before September 11, 2001, most organizations thought of a disaster in terms of a flooding or snowstorm that disrupts operations because essential personnel cannot get to work. Recent events have made it clear that the word "disaster" can mean something far more detrimental. Events may occur, which can take months or even years for an organization to recover from.

History

During the late 1960s, and into the 1970s, consumers of communications, computing, and information technology (IT) began to recognize that their rapidly growing IT operation centers were becoming Single Points of Failure (SPOF). People realized that IT interruptions could potentially have significant impacts on the business continuity of critical operational functions. The continuity of the business itself could even be threatened. Computing hardware, supporting network infrastructures, and software platforms were full of SPOFs during those times.

In the 1960's and 1970's, IT engineers knew how to build resilient computer systems. Actual operational examples of such configurations were developed, and were used in government, military, or research. One of the early IBM mainframes, the System 360 Model 67, had a dual-processor option for improved reliability. The UNIX-based computers that drove Bell System telephone switching throughout that period were also high-availability platforms.

The typical business or government agency of that time could not however, cost-justify the high investment needed to eliminate SPOFs for existing technologies. More economical and practical alternatives were needed as information technology became essential to many public and private organizations. There is a greater risk of computer failure the more an organization depends on computers.

In the late 1970s, through contracts, vendors began offering shared-use access to computing recovery environments. The fees were substantially less than the costs of duplicating critical computing resources by customers. Though SPOFs were not totally eliminated, this cost-effective means to recover from major

outages became the standard for IT recovery during that period and it remains a major sector of the industry today.

Through the 1980s, several companies of various sizes entered the market, offering similar IT operational recovery services. IBM entered the market in a big way in 1990, further adding to the available options. IBM's entry legitimized the industry by effectively endorsing these services as a practical means for customers to recover major and critical portions of their IT infrastructures. The role of business continuity planner or disaster recovery planner evolved from the need to integrate hot site services into a customer's operational environment. This was accomplished by documenting required capabilities and usually included off-site backup data storage services, which were used for testing and in case of a real disaster occurred. Whether documented by the customer's personnel or by consultants, the planner's role was to translate these services into real operational recovery capabilities for a customer's critical IT assets. Disaster recovery is essential in restoring systems and data to a state of normalcy prior to the incident.

Relationship to the Business Continuity Plan

The Business Continuity Plan may be written for a specific business process or may address all mission-critical business processes. The BCP is an umbrella plan whose major sub-components include the Disaster Recovery Plan. Information systems are considered in the BCP only in terms of their support of those business processes. A Business Continuity Plan (BCP) consists of the following component plans:

- Business Resumption Plan
- Occupant Emergency Plan
- Incident Management Plan
- Continuity of Operations Plan
- Disaster Recovery Plan

The Business Resumption Plan, Occupant Emergency Plan, and Continuity of Operations Plan do not deal with the Information Technology (IT) Infrastructure. The Incident Management Plan (IMP), which does deal with the IT infrastructure, establishes structure and procedures to address cyber attacks against an organization's IT systems and generally does not involve activation of the Disaster Recovery Plan.

Disaster Recovery Process

A disaster is defined as a sudden, unplanned catastrophic event that renders the organizations ability to perform mission-critical and critical processes, including the ability to do normal production processing of systems that support critical business processes. A disaster could be the result of significant damage to a

portion of the operations, a total loss of a facility, or the inability of the employees to access that facility.

The disaster recovery process consists of defining rules, processes, and disciplines to ensure that the critical business processes will continue to function if there is a failure of one or more of the information processing or telecommunications resources upon which their operations depends. The following are key elements to a disaster recovery plan:

- Establish a planning group
- Perform risk assessment and audits
- Establish priorities for applications and networks
- Develop recovery strategies
- Prepare inventory and documentation of the plan
- Develop verification criteria and procedures
- Implement the plan

Key people from each business unit should be members of the team and included in all disaster recovery planning activities. The disaster recovery-planning group needs to understand the business processes, technology, networks, and systems in order to create a DRP. A risk and business impact analysis should be prepared by the disaster recovery planning group that includes at least the top ten potential disasters. After analyzing the potential risks, priority levels should be assigned to each business process and application/system. It is important to keep inventory up-to-date and have a complete list of equipment, locations, vendors, and points of contact.

The goal is to provide viable, effective, and economical recovery across all technology domains. The following chart can be used to classify organization applications and/or systems:

Classification of Application/System

Classification		Description
1	Mission Critical	Mission Critical to accomplishing the mission of the organization Can be performed only by computers No alternative manual processing capability exists Must be restored within 36 hours
2	Critical	Critical in accomplishing the work of the organization Primarily performed by computers Can be performed manually for a limited time period Must be restored starting at 36 hours and within 5 days
3	Essential	Essential in completing the work of the organization Performed by computers Can be performed manually for an extended time period Can be restored as early as 5 days, however it can take longer

Classification		Description
4	Non-Critical	Non-Critical to accomplishing the mission of the organization Can be delayed until damaged site is restored and/or a new computer system is purchased Can be performed manually

The disaster recovery process will identify the risks and exposures to mitigate their consequences to a level acceptable to senior management. These risks and exposures will assist in identifying the level of recovery required. Requirements will determine which recovery strategy option is needed to support those requirements.

Disaster Recovery Plan

In its full context, the focus of a Disaster Recovery Plan (DRP) is to restore the operability of systems that support mission-critical and critical business processes. The objective is for the organization to return to normal operations as soon as possible. Since many mission-critical and critical business processes depend on a technology infrastructure consisting of applications, data, and IT hardware, the DRP should be an IT focused plan. Every organization should develop a Disaster Recovery Plan for all applications. Restoration of systems does not necessarily imply technology redundancy. The DRP may call for some procedures to be completed manually. The decision to revert to manual procedures, rather than to build and maintain an IT infrastructure is a cost-driven decision made by the organization. Having a DRP in place reduces the risk that the length of time that a disruption in a business process does not go beyond what has been determined to be acceptable by management in the organization. During the recovery phase, the focus is on establishing controls over occurring events to limit the risk of any additional losses.

IT Disaster Recovery Planning Process

Developing a technical disaster recovery strategy is just one step in the overall IT Disaster Recovery Planning process. This process is common to all IT systems and utilizes the following six steps:

1. Develop the Business Contingency Planning Policy and Business Process Priorities
2. Conduct a Risk Assessment
3. Conduct the Business Impact Analysis (BIA)
4. Develop Business Continuity and Recovery Strategies
5. Develop Business Continuity Plans
6. Conduct awareness, testing, and training of the DRP
7. Conduct Disaster Recovery Plan maintenance and exercise

The objective is to design a technical recovery strategy in step 4. Since this step is being accomplished before a Business Impact Analysis (BIA) can be performed in step 3, the recovery strategy is developed into a standard suite of service offerings that can be activated after the BIA has been completed. A BIA can take months to complete and some organizations do not have the budget for this. However, management should understand the potential return on investment for conducting a BIA.

The goal of the BIA is to define objectives for the recovery of host computing systems that run the applications that support the business processes. These objectives are stated as the Recovery Time Objective (RTO) and Recovery Point Objective (RPO). RTO is the number of hours or days management has put on resuming a business process or a system. RPO describes the age of the data you want the ability to restore to in event of a disaster. For example, if the RPO is 8 hours, systems should be restored in the state they were in no longer than 8 hours ago. The technical disaster recovery strategy depends upon meeting RTO and RPO specifications. The RTO and RPO requirements determine which option of disaster recovery plan to implement.

Recovery time, and how current data is are key components in determining the level of service a business process requires in the event of a major disruption. To properly implement a disaster recovery plan, one must know the RTO and RPO that the organization is willing to accept in case of a disaster. The technical disaster recovery strategy of different options of recovery is based upon a combination of these requirements.

Often times, in regards to business continuity, business and IT units are not on the same page. "As companies become more dependent on information, the business-continuity tolerance for information loss becomes less and less, particularly in e-business," says Don DeMarco, Director, IBM Business Continuity and Recovery Services. Although recovery management (maintaining an IT-based contingency plan and IT recovery plan) is an element of the systems management discipline, DeMarco says, "The decision as to the acceptable amount of risk for information loss must come from upper management."

For example, IBM uses RTO and RPO to classify the two objectives management must consider in business continuity. RTO is used by management to determine the amount of time needed to set up IT capabilities in order to resume critical business processes. RPO is something that management tends to forget. During an outage when business processes cannot be performed, how much data can the organization afford to lose and how current must data being recovered be? A manager of a bank cannot afford to lose six hours worth of data. Management must decide what are the acceptable levels of risk.

Identify Applications

Business processes that have been classified with RTOs and RPOs must be mapped to supporting applications systems along with their data. As stated before, applications are classified as mission-critical, critical, essential, or non-critical. As is often the case, a business process will depend on multiple applications that support the business process. These applications systems need to be identified with the same RTOs and RPOs as their supported business process.

Identify System Recovery Time (SRT)

System Recovery Time (SRT) activity takes place after a disaster is confirmed. The organization must plan the order of priority that it will use to recover hardware systems and components, in order to meet business process RTO. Within the short timeframes of RTO, the DRP often entails the setting up of recovery host systems and related components. The host systems and related components execute the applications that perform the business processes. Hardware infrastructure components needed by the application systems and the data required to support the business process must be identified. All application interdependencies, network infrastructure components, and support staff need to be identified as well. The RTO is used to determine the host component SRT and is carried through from the business process to the application.

There are several ways to bring the SRT within specification if a host component's system recovery time can not meet expectations set by the business process. A different disaster recovery suite can be selected that will accommodate the business process RTO. The time for alert notification, assessment, and disaster declaration can be reduced. Business process RTOs can be reassessed based on a cost/benefit analysis. Since a hardware component may host several applications that support several business processes, the SRT is determined from the shortest RTO for an application that is dependent upon that host component. The host component consists of all items of hardware. This includes host processing system, applications, data storage systems, local and wide area network, and security infrastructure to include firewalls. The components to identify and successfully determine the SRT are as follow:

- Identify business processes
- Determine RTO requirements
- Determine RPO requirements
- Identify application systems that support business processes
- Identify host systems that support applications
- Determine recovery time for host systems and applications

Identify Data Currency for the Applications

Operational activity that supports the application RPO must take place before the business disruption or a disaster event. After the disaster occurs, applications and data at the affected site are unavailable to execute the DRP. Therefore, the normal state of computing center operations must ensure that applications and data are copied offsite from the production-computing center to an offsite storage facility. The offsite storage facility can be located either at the recovery-computing center, or within distance and transport parameters that permit recovery of the data from the storage media to the host infrastructure system within SRT. Based upon the recovery criticality, applications and data will be copied from the production-computing center to the recovery-computing center using one of two methods:

- Electronic replication of application and data from the production-computing center over the Wide Area Network (WAN) to the recovery computing center or electronic vault facility. This replication will initially consist of entire replication of applications and data, followed by time-scheduled replication of files.
- Copy data to physically removable media such as tape, which can be physically shipped away from the production-computing center to an offsite storage facility.

Copies of production applications and data on the production system must leave the production computing center site within the specified RPO timeframe in order to fulfill business RPO requirements. This is especially true with physically removable media, such as tape. The applications and data should be backed up from the production systems to the physical media and removed from the production site within the RPO.

Identify Critical Personnel and Recovery Teams

The applications and host systems that support the critical business processes are dependent upon personnel with a unique knowledge, skills, and abilities. Identifying the staff that has the knowledge to recover the infrastructure that supports the business processes is key to a DRP. Essential support personnel should be identified along with their skill sets. Support knowledge should be dispersed geographically throughout the organization. This will ensure that the organization has sufficient trained staff available to execute the DRP in event of a declared disaster. Employee skill set information should also be documented and included in the DRP, and updated with the same frequency as the DRP.

Testing the Disaster Recovery Plan

To obtain the most value from a disaster recovery test, explicit test objectives and success criteria are required. The use of test objectives and success criteria enable the effectiveness of each DRP element and the overall Business Continuity Plan to be assessed. The two major test criteria are the recovery of

the business process within its RTO with data currency within the RPO. The criteria are as follows:

- Recovery Time Objective
 - Alert Notification/Assessment/Disaster Declaration timeframe is tested and confirmed outside of the context of a system disaster recovery test. Standard operating procedures and call trees are tested to verify that the procedures and personal contact information is current and accurate.
 - System Recovery Time is contained within the application RTO. SRT is tested through a system recovery exercise to determine if recovery operations can be completed within the stated objective.
- Recovery Point Objective confirms that the RPO can be met at any time outside the context of a system disaster recovery test. Standard operating procedures and operational logs are inspected and verified to determine if the data has been transported off site from the production-computing center so that RPOs can be met within the stated objective.

Testing a DRP can be a very complex engagement. The overall objective of the Business Continuity Plan is to continue business processes while the overall objective of a Disaster Recovery Plan is to replicate parts of or the entire existing IT production environment at an alternate site until normal operations have been resumed.

Maintaining the Disaster Recovery Plan

Maintenance of the DRP and the other plans within BCP is critical to the success of an actual recovery. The plans must reflect changes to the environments that are supported by the plans and be up to date. It is critical that existing change management processes are revised to consider recovery plan maintenance. In areas where change management does not exist, change management procedures should be recommended and implemented. Many recovery software products consider this a requirement.

Summary

The world is changing and organizations need to prepare for natural or man-made disasters that could disrupt business processes. Customers and millions of dollars could potentially be lost and never recovered if business processes are disrupted, IT systems exceed their Recovery Time Objective (RTO), or data exceeds the Recovery Point Objective (RPO). The Business Continuity Plan resumes business processes and the Disaster Recovery Plan resumes the IT systems. The objective of a DRP is to restore the operability of systems that

support mission-critical and critical business processes to normal operation as quickly as possible. Business continuity planning integrates the business resumption plan, occupant emergency plan, incident management plan, continuity of operations plan, and disaster recovery plan.

Personnel from each major business unit should be included as members of the team and part of all disaster recovery planning activities. These people need to understand the business processes, technology behind those processes, networks, and systems in order to create the disaster recovery plan. Applications and systems are identified by the team that is mission-critical and critical to the organization. The systems Recovery Time Objective (RTO) and the data's Recovery Point Objective (RPO) are identified by management during the planning phase and are an integral part of developing business resumption strategies. The recovery plan and strategies are designed to meet RTO and RPO. The disaster recovery team will be responsible for training, implementing, and maintaining the plan. They will possess unique skills, knowledge, and abilities that should be updated in the plan. A Disaster Recovery Plan that is well developed, trained on, and maintained, will minimize loss and ensure continuity of critical business processes in the event of disaster.

© SANS Institute 2003, Author

References

- 1) Tipton, H. and Krause, M. Information Security Management Handbook 4th Edition. NY: Auerbach Publications, 2000. 581 – 596.
- 2) Bell, Judy. "Why Some Recovery Plans Won't Work." Disaster Recovery Journal. Spring 2003: 30 - 32.
- 3) Fitz-Gibbon, Thomas P. "Disaster Recovery Planning for Call Centers." Contingency Planning and Management. March 2003: 26 – 28.
- 4) Kunene, Glen. "Create a Disaster Recovery Plan." URL: http://archive.devx.com/enterprise/articles/drecovery/IBM_BCRS/Demarco-1.asp (07 June 2003).
- 5) Herriott, Larry. "Business Contingency Planning Is..." PHH Corporation. URL: http://www.drj.com/new2dr/w3_006.htm (05 June 2003).
- 6) "Disaster Recovery Planning: Project Plan Outline." Computing & Networking Services, University of Toronto. URL: <http://www.utoronto.ca/security/drp.htm> (05 June 2003).
- 7) Kirvan, Paul. "What's Wrong with BCP?" January 2003. URL: <http://www.contingencyplanning.com/PastIssues/janfeb2003/1.cfm> (06 June 2003)
- 8) Bounds, Gene. "Preparing for the Worst: A Best Practices Guide to Disaster Recovery Planning." April 2003. URL: <http://www.contingencyplanning.com/PastIssues/apr2003/5.cfm> (07 June 2003).