

The University of Newcastle upon Tyne
Department of Computing Science

**Development of A Methodology and An Expert
System for Disaster Recovery**

PhD Thesis

NEWCASTLE UNIVERSITY LIBRARY

097 50897 8

Thesis L5979

By

Fahhad Al-Harbi

October 1997

Abstract

The number of organisations that rely on computerised systems to perform their day-to-day operations and to help them in making decisions has grown rapidly over the last few years and continues to expand. On the other hand, the destruction or loss of these systems can be a nightmare and, in many cases, may lead to an end of providing services or trading for the organisation. Thus, the growing dependence on computer systems and the fear of being out of business have increased management awareness and understanding of the importance of plans to prevent or recover from a computer failure.

Although senior management and IT directors have begun to appreciate the need for Disaster Recovery Plans (DRPs), they often raise common questions, such as: How long the organisation can tolerate the failure of its computer systems? Are we spending too much or too little on a recovery strategy? What type of recovery strategy is most appropriate for our IT centre?

To look more closely at the effects of disasters on organisations and the importance of adopting DRPs, the researcher carried out a case study involving 111 organisations in Kuwait to examine their DRPs before and after the Iraqi Invasion in 1990 and to identify major problems facing IT managers on disaster recovery issues.

The literature review and the case study show that there is a lack of a comprehensive methodology and of a computerised intelligent system to guide organisations in selecting the most appropriate recovery strategy for their computer centres. Therefore, this research has developed a methodology and delivered an expert system that would assist IT directors to obtain answers to the above-mentioned questions and perform fast recovery from any type of computer disaster. The methodology consists of five phases that provide a step-by-step approach to ensure that the entire recovery strategy selection process is covered. The phases are: Threats Assessment, Business Impact Assessment, Recovery Strategy Analysis, Cost Analysis, and Recommendations.

Acknowledgement

I would like first of all to dedicate this research to my father for his encouragement and support to proceed with this study. Unfortunately, he died before seeing the end of it.

I would also like to express my sincere appreciation to Dr B. N. Rossiter, my supervisor, for his invaluable advice and assistance during the completion of this project. Special thanks is extended to Dr Lindsay Marshall and Paul Ezhilichelvan, members of the thesis committee, for their useful suggestions, specially in the first year.

I Knowledge and am grateful to Dr. Abedlhadi Alotabi, Director General of the Kuwait Institute for Scientific Research, for his support to obtain the scholarship for this degree.

Two great persons need special recognition for their on-going valuable friendship and support. They are Dr Mashan Alotabi and Dr Khalid Hadi. Their companionships in a country that has a totally different culture made the time pass quickly and pleasantly.

Finally, the most sincere appreciation goes to my wife and my children for their patience and continuos support through the long and winding road.

Table of Contents

Chapter 1

Introduction

1.1 Statement of the Problem-----	1
1.1.1 Potential losses-----	3
1.1.2 The Need for Disaster Recovery-----	4
1.2 Purpose of the Study-----	6
1.3 Contribution of the Research-----	8
1.4 Outline of the Research-----	8

Chapter 2

Contemporary Methods in Disaster Recovery

2.1 Introduction-----	11
2.2 Background-----	12
2.2.1 Causes of Disasters-----	12
2.2.2 Impact of Computer Disasters on Organisations-----	13
2.3 Disaster Prevention-----	17
2.4 Insurance-----	20
2.5 Disaster Recovery-----	22
2.5.1 Expected Value Analysis Method-----	23
2.5.2 Security Assessment Questionnaire-----	25
2.5.3 CCTA Risk Analysis and Management Methodology-----	25
2.5.4 CCTA IT Infrastructure Contingency Planning Module-----	26
2.5.5 Subhani Model-----	28
2.5.6 Seminar/Workshop Methodology-----	29
2.5.7 Generic Disaster Recovery Plan Methodology-----	30
2.5.8 Commercial Disaster Recovery Software-----	31
2.6 Expert System Technology-----	32
2.6.1 The Use of ES in Disaster Recovery-----	33
2.6.1.1 AUDIT-----	34
2.7 Summary of Literature Review-----	36

Chapter 3

Disaster Recovery Strategy Options

3.1 Introduction-----	38
3.2 Recovery Strategy Types-----	38
3.3 Null Strategy-----	39

3.4 Internal Strategies -----	40
3.4.1 Manual Procedures -----	40
3.4.2 Withdrawal of Services -----	41
3.4.3 Duplicate Site -----	41
3.5 Mutual Aid Agreements -----	42
3.5.1 Informal Mutual Aid -----	42
3.5.2 Reciprocal Agreement -----	43
3.5.3 Time Brokers -----	43
3.6 Co-operative Recovery Strategies -----	44
3.6.1 Co-operative Cold Site -----	44
3.6.2 Co-operative Hot Site -----	45
3.7 Commercial Recovery Strategies -----	46
3.7.1 Service Bureaux -----	46
3.7.2 Hardware Vendors -----	47
3.7.3 Commercial Cold Site -----	47
3.7.4 Commercial Warm Site -----	48
3.7.5 Commercial Hot Site -----	49
3.7.6 Portable Site -----	50
3.7.7 Mobile Hot Site -----	51
3.7.8 Realtime Recovery -----	51
3.8 Comparison between Strategies -----	55
3.8.1 Availability -----	55
3.8.2 Operational -----	56
3.8.3 Physical Capacity and Cost -----	58
3.9 Prospective Recovery Strategies -----	60
3.10 Concluding Ramarkes -----	62

Chapter 4

Case Study in DRPs

4.1 Introduction -----	63
4.2 Objectives of the Case Study -----	64
4.3 Methodology -----	65
4.4 Content and Distribution -----	66
4.5 DRPs before the Invasion -----	69
4.6 Organisational Losses -----	71
4.7 DRPs Status Now -----	73
4.7.1 Maximum Allowable Downtime -----	74
4.7.2 Recovery Strategy -----	77
4.7.3 Testing -----	79
4.8 Off-Site Backups Storage -----	80
4.9 Summary and Findings -----	83

Chapter 5

The Methodology for Selecting A Recovery Strategy

5.1 Introduction	85
5.2 The Framework of the Methodology	86
5.3 Threats Assessment Phase	88
5.4 Business Impact Assessment Phase	93
5.4.1 Resource Identification	95
5.4.2 Resource Prioritisation	96
5.4.3 Maximum Allowable Downtime	98
5.5 Recovery Strategy Analysis Phase	100
5.5.1 Organisational Characteristics	101
5.5.1.1 Size of the Organisation	101
5.5.1.2 Degree of Dependency	103
5.5.2 Organisation's Recovery Requirements	105
5.6 Cost Analysis Phase	112
5.7 Recommendations Phase	115
5.8 Concluding Remarks	116

Chapter 6

Expert Systems Technology

6.1 Introduction	117
6.2 The Requirements	118
6.3 Feasible Technologies	120
6.4 The Reasons for Selecting Expert System	123
6.5 Expert System Structure	124
6.5.1 Knowledge Base	125
6.5.2 Working Memory	126
6.5.3 Inference Engine	126
6.5.4 User Interface	127
6.6 The Characteristics of Expert Systems	127
6.7 Expert System Development	129
6.7.1 Knowledge Acquisition	130
6.7.2 Knowledge Representation	131
6.7.3 Implementation and Prototyping	131
6.8 Selecting the Tool	132
6.8.1 Languages and Tools for Building Expert Systems	132
6.8.1.1 AI Languages	133
6.8.1.2 SHELLS	134
6.8.1.3 Toolkits or Environments	134
6.8.2 Potential Tools and The Selection Determinants	135
6.8.3 Reasons for Selecting <i>Flex</i>	138
6.9 Concluding Remarks	139

Chapter 7

A Prototype Expert System for Disaster Recovery Strategy

Selection

7.1 Introduction	141
7.2 Overview of the Prototype ESDRSS	141
7.3 Knowledge Acquisition	145
7.4 Computing Maximum Allowable Downtimes	147
7.5 Computation of Investment	149
7.6 Recommending A Recovery Strategy	153
7.6.1 Questions	154
7.6.2 Rules and Inferencing	157
7.6.3 Weighting of Rules	158
7.6.4 Examples of ESDRSS Rules	159
7.7 Testing	166
7.8 Concluding Remarks	167

Chapter 8

Discussion and Conclusion

8.1 Summary of the Research	168
8.1.1 Findings	169
8.1.2 The Methodology	171
8.1.3 The ESDRSS	173
8.2 Who Will Benefit from this Research?	173
8.3 Future Research	174
References	176
Appendix A - System Code	189
Appendix B - Examples of Input and Output Screens	217

List of Figures

Figure 2.1 - Decline in Operational Activities for the Financial Sector - 1978 -----	15
Figure 2.2 - Maximum Allowable Downtime by Organisation Category - 1978 -----	15
Figure 2.3 - Dollar Loss in Manufacturing or Distribution Industry with \$215+ Million Annual Gross Sales in 1978 and 1989 -----	16
Figure 2.4 - Impact of Computer Disasters on Financial Sector -----	17
Figure 4.1- Proportions of Business Activities by Organisation Category -----	68
Figure 4.2 - Percentage Size of Organisations -----	68
Figure 4.3 - Degree of Dependency among Organisations -----	68
Figure 4.4 - Organisations by Processing Type -----	69
Figure 4.5 - DRPs Availability before the Invasion -----	69
Figure 4.6 - Organisational Losses Due to the Invasion -----	72
Figure 4.7 - Availability of DRPs in 1996 -----	73
Figure 4.8 - Excuses for not having Approved DRPs -----	74
Figure 4.9 - Maximum Allowable Downtime by Organisation Category -----	76
Figure 4.10 - Maximum Allowable Downtime by Size for Financial Institutes -----	76
Figure 4.11 - Preferred Recovery Strategies by Organisation Category -----	78
Figure 4.12 - Recovery Strategies by Organisation Size -----	79
Figure 4.13 - Testing Among Available DRPs -----	80
Figure 4.14 - Backups of Critical Systems -----	82
Figure 4.15 - Tests of Backups Per Year to Ensure Readability -----	82
Figure 5.1 - Recovery Strategy Selection Framework -----	87
Figure 5.2 - Threats Magnitude Classification (TMC) Approach -----	90
Figure 5.3 - Criticality in Respect to Tolerance and Cost -----	98
Figure 6.1 - Expert System Structure -----	125
Figure 6.2 - The <i>Flex</i> Environment -----	137
Figure 7.1 - The Design of the Prototype ESDRSS -----	143
Figure 7.2 - The Size Question Example -----	156

List of Tables

Table 3.1 - Types of Recovery Strategies-----	39
Table 3.2 - Major Advantages and Disadvantages of Recovery Strategy Options-----	52
Table 3.3 - Recovery Strategy Options - Availability -----	56
Table 3.4 - Recovery Strategy Options - Operational-----	57
Table 3.5 - Recovery Strategy Options - Physical Capacity and Cost-----	59
Table 5.1 - Types and Examples of Threats in the TMC Approach-----	92
Table 5.2 - Cost of Downtime-----	100
Table 5.3 - Types of Recovery Strategies in Relation to Size-----	103
Table 5.4 - Group Selection Based on Size and Degree of Dependency-----	104
Table 5.5 - Recovery Services Associated with Threat Types -----	106
Table 5.6 - Personnel Support Levels among Recovery Strategies-----	107
Table 5.7 - Work Area Levels among Recovery Strategies -----	108
Table 5.8 - Security Levels among Recovery Strategies-----	109
Table 5.9 - Ability to Modify Platform among Recovery Strategies -----	110
Table 5.10 - Recovery Time for Various Recovery Strategies-----	110
Table 5.11 – Location and Restoration Period -----	111
Table 5.12 - Usage Duration between different Recovery Strategies -----	112
Table 6.1 - Comparison Between Expert Systems and Conventional Systems-----	121
Table 7.1 - Examples of Time Intervals for Some Types of Organisation -----	149

Chapter 1

Introduction

1.1 Statement of the Problem

Recent years have seen a revolution and a tremendous growth in technologies as they have come to play a central role in shaping the evolution of organisations. Computer technology used by organisations nowadays has been one of the most pervasive applications of technology in this revolution. It is rapidly bringing us out of the industrial age into a new epoch, 'the information age'. It has been said that computer technology will prove more important than the steam engine, which in its own time laid the foundations for the industrial revolution (Daler, Gulbrandsen, Melgard, & Sjolstad, 1989). Computers provide a myriad of everyday conveniences and benefits to organisations and customers such as Automated Teller Machines (ATM), controlling telephone networks, diagnosing the body's internal ills, recording the price of groceries at the supermarket check-out, reserving tickets, forms of decision-making, etc. It seems that no other machine in history has so rapidly and so completely changed the world. In short, 'we are living under the very fabric of modern life, making computer avoidance, if not computer ignorance, particularly impossible' (Hassig, 1991).

The introduction of computers, therefore, has created a technology revolution for modern organisations. Since the early birth of computer machines, organisations have been computerising the storage, retrieval, and the processing of huge amounts of data as a technical approach to improving the content and flow of information within and between organisations and society. Year after year, organisations are becoming increasingly aware of the contribution that the computer centre makes to the overall

well being of the company. In fact, in today's environment 'information is the life blood of the business, and the heart that pumps this blood is the computer' (Williams, 1995).

Furthermore, information systems are now considered to be a basic component of nearly all private and government organisations. US companies, for instance, spent close to \$30 billion on their information systems in 1987, and are reaping the fruits of their investments in the form of faster, more refined, more meaningful data - the kind of data that supports decisions and creates wealth (Toigo, 1989). As a result information processing has become the nerve centre of most organisations (Baylus, 1991).

However, there is a side to this symbiosis of organisations and machines that is rarely examined. It is business's dependency on the uninterrupted flow of information from its systems and the consequences for a company if the computer machine was to be suddenly switched off. The occurrence of such unscheduled and inconvenient switching-off is called 'Computer Disaster' (Toigo, 1989). According to Baylus, the term computer disaster means 'any accidental or intentional event that causes disruption to a company's operations' (Baylus, 1991). Toigo, a disaster recovery expert, defines a computer disaster as "the interruption of business due to the loss or denial of the information assets required for normal operations." He adds "it refers to loss or interruption of the company's data processing function, or to a loss of data itself' (Toigo, 1989).

It is human nature to think we will not be hit by a disaster, whether at home or at work, because disasters are perceived as low probability events. But in the last few years a seemingly endless procession of fires, floods, hurricanes, earthquakes and bombings have proved otherwise. Past studies show that organisations are not adequately prepared for computer disasters and they exhibit a reluctance to spend money on acquiring the services needed for a recovery plan. This often occurs when management does not fully understand the risks and exposures that an organisation faces without a recovery capability. Computer managers have tended to ignore the

possibility of disasters occurring. Some managers have the attitude 'it will not happen to us'. Some learn the hard way that this is not so, and few are lucky enough to survive the experience (Smith, 1989).

1.1.1 Potential losses

Although the probability of disaster occurrence may be quite low, the potential business impact is generally large and possibly long lasting. The impact may manifest itself in a number of different ways. With a prolonged and total loss of service, the consequences may be severe - terminal for around 80% of companies (Hiles, 1992). Existing customers may transfer business elsewhere and prospective ones may turn to other competitors. New business is strangled, even loyal customers quickly become disaffected and market share drops. Automated order-taking or telephone-based reservations, supported by on-line transaction systems, often have no alternative manual means of input. For those the equation is quite simple: no computer - no sales. The distribution of products to many large retailers may be severely affected if the computer is down. Many products in the warehouse cannot be despatched; those which are sent off are either late or may arrive in the wrong quantities. This then constitutes a drain on the manufacturer's cash (Heirlein, 1993; Robinson, 1993; Copenhaver, 1997).

Financial losses can also be indirect. Additional staff, hired to take over clerical workloads until the computerised information systems are restored, must be paid. Collection of money owed to the company could slow down significantly if computerised debtors defer settling bills, knowing that credit control systems are not available to pursue them (Hiles, 1992). Important deadlines for payment are missed - payroll, tax, instalments on contracts etc. Understandably, staff loyalty may be severely tested and key staff may fear for their future and join the competition - as, indeed, happened in a well-publicised case involving Hackney Borough Council (Smith, 1990) and in some organisations in Kuwait after the Iraqi invasion (see Chapter 4: a recent survey into the effect of the Iraqi invasion on organisations in Kuwait).

Perhaps an even more important result of losing the computer systems is the loss of control of the organisation by senior management. If management information is corrupted or out of date, poor decisions may be made on vital business issues. The increasing corporate dependence on Decision Support Systems and Executive Information Systems renders this threat even more serious. Thus, the organisation's image and credibility may be damaged beyond recovery (Hiles, 1992; Smith, 1990).

Other indirect losses are related to legal issues. An organisation might have a contract that allows a major customer to impose penalties if goods are not delivered on time. Government regulations might strictly govern a company's business activities, and legally such a company must be available to conduct business. The top management of a company might also have special obligations to its shareholders. Thus, the company officials could be held liable for both criminal and civil damages if they neglect to develop an effective means of protection (Epich & Persson, 1994; Copenhaver, 1997).

1.1.2 The Need for Disaster Recovery

The most widely adopted means of protection, to avoid or minimise the above-mentioned losses, include one or more of the following: 1) transferring the risks via insurance; 2) adopting a disaster avoidance plan to reduce or limit the risks; and 3) adopting a recovery plan to guide the organisation in resuming services and vital business functions.

Being indemnified for a financial loss thorough insurance is not always sufficient to compensate for other indirect losses like loss of market share and goodwill. Moreover, it should be noted that information systems insurance is a complex issue. Comprehensive cover is rarely offered; if it is available, it is very expensive (Haack, 1984; Orr, 1988). However, most disaster recovery experts recommend that the insurance coverage should be part of the disaster recovery plan (but not an alternative to a recovery plan) to fund the recovery efforts following a disaster (Arnell, 1990;

Hearnden, 1993). In fact, organisations that have disaster recovery plans pay lower insurance premiums (Baylus, 1991).

Some organisations apply a disaster avoidance plan by installing some safeguards, particularly for minor threats such as viruses, hackers, and small fires. Despite the high cost of installing safeguards against more serious disasters, even the best avoidance plans cannot prevent every disaster (Redmond, Luongo & Tietz, 1996). There are many disasters, like recent terrorist activities and major earthquakes, which are beyond the control of any type of preventive countermeasures. Thus, according to Ed Devlin, a senior vice-president with Strohl Systems, “we cannot prevent disasters from occurring; therefore we must plan for a fast recovery to minimise their impact” (Devlin, 1996).

Organisations tended in the past to apply the two above-mentioned options: insurance and disaster prevention. In recent years, however, organisations dependent on their IT systems have shifted their focus towards:

- Resuming vital operations within a specified time after the incident occurs;
- Establishing alternative means of operation; and
- Returning to normal operations as soon as practicable (Hyde, 1993; Baylus, 1991).

Therefore, attention began to shift toward disaster recovery. This is not to say that insurance and disaster avoidance measures were completely dismissed; in fact, their relevance in some situations was acknowledged. This turn-around was first observed in the 1990s, when new technological developments changed the way in which organisations worked. Investment in sophisticated corporate networks and business systems has increased the number of points of potential failure within an organisation’s information systems structure. So top management has begun taking more interest in disaster recovery planning. This was demonstrated at many seminars such as the one held in 1993 by CDRS Europe, entitled “Out of the Computer Room and Into the Board Room”, where a healthy proportion of the audience comprised

senior directors and managers at board level (Hyde, 1993). Another indication of senior management's increasing awareness of the importance of disaster recovery was the presence of senior banking and financial executives at a nation-wide teleconference, hosted by SunGard Recovery Services in 1993 to discuss only disaster recovery issues (Datapro, 1993).

Furthermore, the importance of disaster recovery was demonstrated when a major fire destroyed the trading room of the Credit Lyonnais, one of the biggest banks in France, in May 1996. Patrick Hummel, the IT director, proudly announced that although the fire happened on a Sunday morning, their disaster recovery plan along with the alternative site for real-time recovery strategy enabled the bank to conduct 'business as usual' for traders when they arrived at work on the Monday morning (Hars, 1996).

The clear message to emerge from past incidents and recent seminars is that there is a compelling need for adequate disaster recovery plans.

1.2 Purpose of the Study

In today's information intensive economy, survival of a company may depend on management's ability to use its information resources to provide services, to compete effectively, to strategise, to hold market share and to expand, in essence to survive. Thus, organisations, nowadays, cannot tolerate the denial of the computer systems for a long period of time. The consequences of failing to survive after a disaster are so dire that more and more organisations have been forced to recognise the importance of disaster recovery plans.

Although senior management and IT directors have begun to appreciate the need for disaster recovery plans, they often raise common questions, such as: How long the organisation can tolerate the failure of its computer systems? Are we spending too much or too little on a recovery strategy? What type of recovery strategy is most appropriate for our IT centre?

Much of the literature in the field of disaster recovery dealt with the need for disaster recovery planning, how to develop and implement disaster recovery plans, and the consequences of not having one. It shows that the issue of selecting the most suitable recovery strategy, including answers to the above-mentioned questions, has not been fully and adequately addressed. Therefore, a structured methodology is needed to address the issue in a comprehensive way, covering a wider spectrum of possible scenarios in the disaster recovery area.

The aim of this research, therefore, is to address more fully the fast recovery of IT services from unscheduled interruptions caused by computer disasters. To this end the following objectives have been adopted:

1. To develop a comprehensive methodology for IT managers and officers who are responsible for disaster recovery activities that includes:
 - i) A method for classifying threats so that it actually contributes to solving the problem of recovery strategy selection;
 - ii) A full business impact analysis that includes automatic computation of the exact maximum allowable downtime;
 - iii) An approach to estimating how much to spend on disaster recovery; and
 - iv) Methods of constructing organisational requirements and defining recovery strategy characteristics so that they can be utilised by expert system technology to select the most appropriate recovery strategy.

2. Through this developed methodology, to provide a basis for the development and implementation of a structured prototype expert system to assist IT managers in reaching decisions during the disaster recovery selection process.

In addition, the present research complements the work already completed by Tawfig Danish, who obtained his Ph.D. at the University of Newcastle upon Tyne. In his research, he addressed the utilisation of expert systems technology in the field of disaster prevention (Danish, 1994). To provide the complete picture of disaster

preparedness, the present research investigates the later activities - recovering from a disaster. Taken together the two pieces of research will produce a more comprehensive perspective on the disaster preparedness concept using expert systems technology.

1.3 Contribution of the Research

The present research can make several contributions to the information technology and disaster recovery communities. The content of this research, presented in Chapters 2 to 7, contains the following elements:

- The literature review, which highlighting some of the key milestones and methodologies, underpins the proposed solution which is developed.
- A recent case study that shows the importance of disaster recovery planning and identifies some of the key problems facing IT managers regarding disaster recovery issues. The study also makes a significant contribution to the design of the required methodology and the proposed system.
- A methodology for solving the problem of selecting the most appropriate recovery strategy is presented.
- A new classification of threats that can enhance the recovery strategy selection process is developed.
- A review of available technologies to find a suitable tool for implementing the developed methodology concludes that an expert systems approach is feasible to achieve the objectives.
- A prototype expert system to assist IT managers in decision-making on disaster recovery issues is implemented.

1.4 Outline of the Research

The ultimate purpose of this study is to address the concept that IT departments should be secure and reliable and, therefore, those departments should be prepared for

unexpected threats to the continued conduct of their business if they wish to assure survivability. The following chapters set out the various aspects of the study:

Chapter 1 describes the background of computer applications and potential losses expected from a computer disaster. It presents the purpose, contribution, and structure of the research.

Chapter 2 describes the causes of disasters and the impact of computer disasters on organisations. It also identifies and briefly describes the major milestones of development in the field of disaster recovery. The utilisation of expert systems technology in disaster recovery is also reported.

Chapter 3 identifies, briefly describes, and compares major available recovery strategies in the disaster recovery field. Some major prospective recovery strategies are also reported.

A recent study, forming part of the work, that analyses the disaster preparedness of Kuwaiti organisations before and after the Iraqi invasion in August 1990 is reported in Chapter 4. The study highlights the consequences of not having disaster recovery plans and identifies major problems facing IT managers on recovery issues. Some results from the study are used to help in designing the proposed solution.

Chapter 5 presents and explains the proposed methodology for selecting the most appropriate recovery strategy, including threats assessment, business impact assessment, recovery strategy selection and a model for calculating investment.

Potential technologies and tools for implementing and delivering the computerised system are presented in Chapter 6. An analysis comparing and evaluating available technologies and tools is then carried out to enable the most suitable one to be selected.

Chapter 7 describes the proposed prototype Expert System for Disaster Recovery Strategy Selection (ESDRSS). Several examples are introduced to illustrate the operation mechanism of its components.

Finally, Chapter 8 concludes the study by citing possible areas for future research in this field.

Chapter 2

Contemporary Methods in Disaster Recovery

2.1 Introduction

A substantial review of related literature shows that there is a great deal of work that has been carried out in respect of computer disasters and their effects on businesses. The literature also shows that some attempts have been made to prevent some types of disasters and/or mitigate the overall effects of others. This mitigation takes the form of setting a set of safeguards and/or recovery procedures. In this chapter, a full investigation of all areas related to the proposed problem is carried out. The investigation covers the following areas:

- Background of computer disasters. This includes looking at some recent disasters, impact of disasters in businesses and the importance of having Disaster Recovery Plans (DRPs).
- Disaster prevention area. This includes the methods used in risk analysis and installing safeguards.
- The use of insurance as a method of mitigating disasters consequences.
- Disaster Recovery area. This involves looking at methodologies that are used in developing a Disaster Recovery Plan (DRP).
- Expert systems technology. This includes looking at existing expert systems in the disaster recovery area.

2.2 Background

Several catastrophes have dominated the news over the past few years. These disasters range from major devastating threats such as earthquakes, floods, fires and terrorism attacks to relatively minor ones such as power outages, hackers and computer failures. This research, however, will focus only on disasters that affect computer centres and data processing facilities.

2.2.1 Causes of Disasters

A computer disaster, which is the focus of this study, can be defined as “any accidental or intentional event that causes disruption to a company’s operations” (Baylus, 1991). Toigo (1989) however defines computer disaster as “the interruption of business due to the loss or denial of the information assets required for normal operations.” He adds, “it refers to loss or interruption of the company’s data processing function, or to a loss of data itself.”

A considerable amount of attention has been given in the past to natural disasters such as earthquakes, floods and tornadoes. New attention, however, is being directed to the raising incidents of terrorist attacks such as the bombings in Manchester City, Oklahoma City and the World Trade Centre in New York. But there are other types of problem that bring about disturbance to businesses. Fires destroying buildings, power outages, equipment failures, hackers and viruses are some of these problems. Some of these incidents are certainly less dramatic but they are no less disastrous when it comes to grinding the Information Systems to a halt. According to the Computer Disaster Casebook produced by BIS Applied Systems, which covers more than 175 computer disasters in the UK, fires and explosions have the greatest number of disaster occurrences amounting to 36%; software failures 24%; power outages 21%; water damages 9%; other disasters represent 10% (CCTA, 1989).

Another survey which was carried out by the Contingency Planning Research, covering the period between 1982 and 1985, showed that power failures accounted

for almost 28% of US computer outages, followed by storm damage 11.7%, floods 9.6%, hardware error 7.7%, bombings 7.2%, hurricanes 6.3%, fires 5.6%, software error 5.4%, power surge 5.1%, and earthquake 4.9% (DRJ, 1997).

Unfortunately, the number of disasters affecting organisations and businesses is on the increase. The United Nations has designated the 1990s as the International Decade for Natural Disaster Reduction (IDNDR) to reduce the loss of life, property damage and economic disruption caused by natural disasters especially in developing countries (Katayama, 1993). According to an independent survey commissioned by the London-based International Computer Room, specialists in the Hardware Environmental Protection Agency in the UK, an estimated 90% of all computer rooms are at the mercy of an environmental time bomb that has nothing to do with software bugs, network access times or user intolerance (Reed, 1992). Another survey by Arthur Young is quoted to have predicated that 1 in 10 of all companies in the UK will experience a computer disaster (Allen, 1992).

2.2.2 Impact of Computer Disasters on Organisations

Nowadays, organisations in the public and private sectors depend heavily on computer information systems in running their businesses. The number of organisations that rely on computer technology to perform their day-to-day operations and to help them in making decisions has increased rapidly over the last few years and continues to do so. A disruption of computer systems for a few days or even a few hours can, therefore, cause severe financial loss and threaten the survival of the business. When the system is down, the business comes to a halt during the time it takes to recover the system, recreate the lost data and applications, and deal with the backlog of data transactions which occurred during the downtime (Ratliff, 1993).

Recent surveys and studies have shown that the impact of disasters on organisations is enormous. Negative consequences may emerge immediately after the disaster or they may appear gradually in the following years. According to a survey by Price

Waterhouse, 70% of UK companies which have not recovered from a major disaster within 48 hours tend to fail in the following years (Allen, 1992). The Chubb Insurance Group of America commissioned a research project amongst computer-dependent companies which suffered a disaster. The result showed that 9 out of 10 went into liquidation within 18 months (Reed, 1992). This should serve as a serious warning to all computer dependent organisations.

Additional national studies and statistics give further evidence to the increasing impact of catastrophes as organisations continue to evolve into a computerised information dependent economy:

- 50% of all computer dependent businesses that experience a disaster and do not re-establish processing and operations within 10 days never recover (Wesselingh, 1990).
- 60% of companies which are affected by a major disaster in the USA go out of business within two years (DRJ, 1997).
- According to a study by Amedahl Executive Institute, a UK retailer cannot operate its distribution depot for more than 24 hours without computer support (Smith, 1990).
- Each on-line outage, averaging 4 hours, costs companies an average of \$329,000 in lost revenues and productivity (DRJ, 1997).
- Every five minutes, a business catches fire in the US; of these 90% suffer catastrophic losses and 40% never reopen (Wesselingh, 1990).

The first study concerning the impact of computer disasters was carried out in 1978 by the University of Minnesota. According to the study (see Figure 2.1) a data processing failure in a financial institution, one-half day in length, will degrade normal business activity by 13% for the two weeks following the failure. A ten-day outage will result in 96% loss of business activity. The study also examined the relative vulnerability of specific industries and demonstrated the Maximum Allowable Downtime (MAD) allowed by industry before recovery would be nearly impossible. As summarised in Figure 2.2, financial institutions have the lowest tolerance to a

prolonged downtime, while insurance companies have the largest MAD of all. The survey also produced an analysis of dollar loss following a data centre disaster in manufacturing or distribution industries with over \$215 million annual gross sales, Figure 2.3 (Toigo, 1989).

Figure 2.1 - Decline in Operational Activities for the Financial Sector - 1978

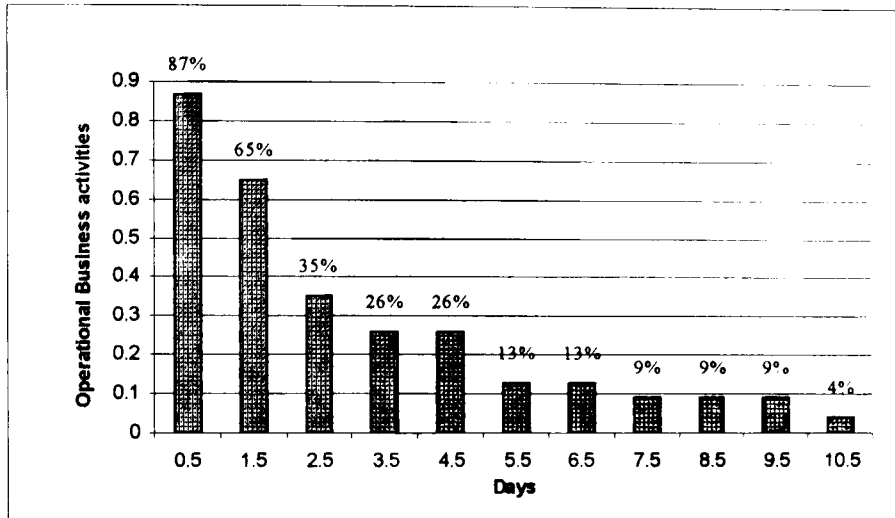


Figure 2.2 - Maximum Allowable Downtime by Organisation Category - 1978






Figure 2.3 - Dollar Loss in Manufacturing or Distribution Industry with \$215+ Million Annual Gross Sales in 1978 and 1989



The University of Minnesota's study is almost two decades old covering a period when only mainframes and midrange computers were available. Since then, substantial changes in the computer technology have occurred, including the proliferation of PCs, networks and telecommunication. Due to these changes and the rising dependency on computers, the results nowadays would not be the same as those in 1978. Percentages in the loss of business activity are enormously greater, organisation's MADs are substantially reduced (this can be seen in the survey presented by the researcher in Chapter 4), and dollar loss following computer disasters is definitely larger. To illustrate this, Alvin Arnell, a disaster recovery expert, compared his observations in year 1989 regarding the cost in dollars following a computer disaster in the distribution industry to those produced by the 1978 study. The cost of denial of computers for 5 days in 1978, according to the University of Minnesota study, is \$94,200, whereas this amount would be lost in less than one day in 1989. Figure 2.3 depicts a comparison between Arnell's observation in 1989 and the University of Minnesota's study in 1978 in terms of the dollar loss. To further illustrate the growing impact, a recent study in 1990 by Price Waterhouse (Figure 2.4) showed the impact of computer disasters on financial industry in the UK (Danish, 1994).

Figure 2.4 - Impact of Computer Disasters on Financial Sector



With this increasing number of threats striking organisations in different parts of the world and the growing impact of computer disasters, awareness of the vulnerabilities of businesses to those unexpected interruptions has dramatically increased. To deal with these risks, there are three types of response available to IT managers:

- Prevention measurements
- Insurance
- Recovery facilities

2.3 Disaster Prevention

Disaster prevention (risk management or risk analysis) seeks to avoid threats in the future by installing protective measures, so that the consequential losses are minimised (Faithfull and Watt, 1991). It is considered to be completely different from the disaster recovery concept (Orr, 1988; Danish, 1994). The process of any disaster prevention policy is summarised below (Orr, 1988; Arnell, 1990; Danish, 1994):

1. Identify assets which need to be protected;

2. Identify threats that may strike followed by risk assessments, and
3. Select appropriate countermeasures to protect the identified assets from the expected threats.

Work in disaster prevention began in the early 1970s in the USA. The first milestone, represented by FIPS PUB 31 in 1974, was enforced by the Public Law B9-306 (Brooks Bill), Part 6 of which is entitled Code of Federal Regulations. Although the Bill provided Federal agencies with a handbook for use when implementing physical security and risk management programs in their IT installation, it was only intended to be a basic reference document and check-list for general use (Danish, 1994).

Following the initiative set in FIPS PUB 31 in 1974, several risk analysis systems and methods were introduced in the disaster prevention area. The systems and methods which are available today fall generally into one of two categories, quantitative and qualitative. Examples of the quantitative methods are the Federal Information Processing Standards publication number 65 (FIPS 65), a method which has been used by IBM draws on FIPS 65 and a software named RISKCALC. Examples of qualitative methods are Los Alamos Vulnerability/Risk Assessment (LAVA), CCTA Risk Analysis and Management Method (CRAMM) and software package called RISKPAC, which was jointly developed in the US by Chemical Bank Information Systems and Profile Analysis Corporation of Ridgefield.

In addition to these attempts, an interesting methodology for disaster prevention in IT centres which uses expert system technology was introduced recently in 1994 by Danish in the Computing Science Department at the University of Newcastle upon Tyne. (Since disaster prevention is not the scope of this research, disaster prevention and risk analysis approaches will not be included and explained here. Further reading of the approaches and methods used in disaster prevention can be obtained from NBS (1985), Jackson and Hruska (1992), and Danish (1994))

As mentioned before, the area of disaster prevention is considered to be completely different from the area of disaster recovery. However, there is one issue that can be of

interest in this research. The issue is the cost-benefit analysis performed in selecting and implementing suitable countermeasures for preventing a threat to occur

The cost-benefit analysis, in disaster prevention, is the process of comparing the estimated expected losses as consequences of an expected threat to the cost of countermeasures to be implemented to prevent such losses. The method for selecting an appropriate countermeasure is described below (Baylus, 1991; Moses, 1992; Danish, 1994):

- Identifying the potential cost of a single occurrence of each identified threat for each resource and asset.
- Estimating the likely frequency of occurrence of identified threat sources.
- Computing the Annual Loss Expectancy (ALE) for each resource.
- Selecting the suitable countermeasures based on cost-benefit analysis.

In disaster prevention, if the cost of installing countermeasures is less than the total ALE, then the proposed countermeasures can be installed. However, if the cost of countermeasures exceeds the total ALE, then other alternatives such as insurance or disaster recovery should be considered (Danish, 1994).

The ALE calculation method was first introduced by the Federal Information Processing Standards publication number 65 (FIPS 65), which was mainly used in the disaster prevention area. It is a traditional and proven way for calculating the investment needed for a typical preventive countermeasure. However, there are several problems associated with this method. These problems are explained in section 2.5.1 of this chapter, when the Expected Value Analysis Method is presented.

In conclusion, there are some disasters which can be avoided by implementing some preventive countermeasures such as power failure and small fires. However, there are other types of disaster which are beyond the control of any type of preventive countermeasures such as terrorism and major earthquakes. To deal with these

uncontrollable risks, IT managers need to look into the other two alternatives insurance and recovery.

2.4 Insurance

Some organisation's management may choose the option of being insured in order to deal with computer disasters. One of the major findings of a recent survey carried out on UK organisations in January 1993 by the University of Loughborough in association with the Computing Services Association and the National Computing Centre, is that many companies lie in the comfort factor of being insured. In fact, over eight out of ten companies in the study claimed to have some form of insurance against computer disasters (Hearnden, 1993).

However, it should be noted that information systems insurance is a complex issue and a comprehensive cover is rarely offered or it is very expensive (Orr, 1988). Although some computer facilities such as hardware, software, network and media can be insured easily, other issues such as the value of data, customer satisfaction, contractual issues are difficult to insure, if not impossible. According to an IBM report in 1993 reviewing the Loughborough survey 'there is some evidence to suggest that the view of what companies believe they are covered for may be more optimistic than the reality.' Although 68% claim to be covered for loss of data, the 1991 Audit Commission Report showed that only 9% of such losses were actually claimed against insurance (IBM Report, 1993b). The Loughborough survey indicates that the level of insurance cover by organisations which claimed to have some form of computer insurance was far from perfect. The survey produced the following warning results regarding information systems insurance coverage:

1. Over 30% of organisations on UK are not covered against loss of software or the cost of reconstituting data.
2. Nearly 40% have no cover against consequential business or financial losses.
3. 60% do not have fidelity bonding (insurance against employee negligence or abuse).

4. Over 75% are not covered against software failure.
5. Nearly half the companies which claimed to have insurance have no emergency recovery facilities.

Furthermore, business interruption insurance resulting from computer disasters is expensive, and could cost an organisation almost four times what it would pay for DP property insurance (Subhani, 1989). To demonstrate how expensive insurance is, Lawrence Cox, president of Cox Insurance Services, gives the following example. Earthquake deductibles can be 10% of the insured value of the damaged building, multiple buildings will mean multiple deductibles. For a large loss of say \$100 million the company may face a deductible expense of more than \$10 million. For a company with a \$100,000 loss, a \$10,000 deductible may shut them down (Cox, 1996).

In addition to the above-mentioned drawbacks, more importantly, insurance *does not put the company back in business* serving the customers, or provide the necessary organisational computer services (Arnell, 1990). 'How ever good the insurance cover, it remains a case of shutting the stable door after the horse has bolted' (IBM, 1993b).

Having said that, there is no way in which absolute protection and recovery from a disaster is possible. Realistically the organisation must minimise possible loss and cover by insurance (Baylus, 1990). In fact, insurance money is needed to fund the recovery efforts after a disaster. An insurance company provides the required money to repair or replace the lost hardware, software and financial expenses which the business would not ordinarily have if there had been no disaster.

Although the computer disaster insurance issue is very complex, most disaster recovery experts recommend that the insurance coverage should be part of the disaster recovery plan, not an alternative to a recovery plan (Arnell, 1990; Hearnden, 1993). In fact, insurance companies recommend disaster recovery to clients and usually consider a reduction in insurance premiums when an adequate recovery plan is in place. Moreover, a company may be prepared to concentrate on certain recovery aspects and simply insure other aspects (Baylus, 1990; Peach, 1991).

As has been explained, insurance is not considered to be a recovery alternative and it does not put the company back in business. Rather it should be an integral part of the recovery plan. Therefore, to insure the continuity of the business during and after an unexpected disaster, a recovery strategy is needed.

2.5 Disaster Recovery

Disaster recovery planning, or what is also called contingency planning, business continuation, or business resumption, is relatively new. It was first introduced in the United States in the late 1970s with the introduction of the mainframe recovery industry (Wrobel, 1990). The defining moment was the establishment of SunGard Recovery Services in Philadelphia, Penn. in the early 1980s. Since then awareness of the importance of disaster recovery has been raised due to the increasing dependency on computers and a number of events affecting the IT environment. This in turn has led to rapid market growth particularly in the US and the UK, but also in other countries (Hyde, 1993). The industry grew to over 100 commercial providers of backup computer centres located throughout the US (Schreider, 1995). A recent report by Datapro (1993) showed that all disaster recovery vendors world-wide who were interviewed by Datapro have enjoyed a good growth rate of around 25%, even during a world economic recession (Hyde, 1993). According to a report by G-2 Research Incorporated, the world business recovery spending in 1995 was approximately \$3.1 billion dollars and is estimated to grow at 20 percent annually through 1999 (DRJ, 1997).

Although the UK is still behind the US, industry analysts agree that the UK is the fastest growing market. Many analysts suggest that the reason for this acute awareness is that, because of recent events, UK organisations are more aware of the threats around them (Hyde, 1993). Hyde adds that other European countries are approximately three years behind the development of the UK. She claims that the problems faced by vendors in these countries amount to general awareness and education in the disaster recovery concept.

Because of the newness of the disaster recovery area, efforts are still made to develop effective methodologies for solving different problems which are related to this area. Most previous attempts in disaster recovery are prescriptive in nature. They detail steps to be taken in order to prevent disasters and steps to develop a workable contingency plan for disasters. Some others are reports of actual distress in businesses describing the nature of the disaster, what preventive measures were in place, damage suffered by the business, and lessons learned from the disaster.

In the sections that follow, brief explanations of existing methodologies in the disaster recovery or related areas are presented. Some of these methodologies are prescriptive and are considered as guidelines and steps. Others are models to be adopted in solving certain problems related to risk analysis, such as the Expected Value Analysis method and the Subhani model.

2.5.1 Expected Value Analysis Method

The traditional approach for risk analysis of computer disasters is based on the expected value analysis. The expected value analysis is also known as the Annual Loss Expectancy (ALE). The ALE approach was first introduced by FIPS Publication 65 in 1979 (FIPS 65, 1979). (The ALE method was explained earlier under the disaster prevention's heading) Although this method has been widely used in the disaster prevention area, some traditional disaster recovery experts apply it in order to decide whether to adopt a particular recovery strategy or not. The principle behind the application of the ALE method in disaster recovery is that the maximum allowable cost of any recovery strategy should not exceed the expected losses (FIPS 87, 1981). A major strength of this method is that it produces a cost analysis which can be easily understood by managers. In other words, it gives dollar estimates of expected losses resulting from a computer disaster.

The ALE approach or one of its variants has been used extensively for some years, particularly in the US, and is very popular with the IT community because of its heavy

emphasis on money figures throughout the method. However, it has been found to be fundamentally flawed. Indeed, the publishers of FIPS 65, National Institute of Standards and Technology (NIST), do not seem to support the ALE approach any longer (Jackson and Hruska, 1992; Moses, 1992).

The main problems and major criticisms associated with the ALE approach are given below:

- Estimating the probabilities of threats is often misleading and inaccurate (Baylus, 1991; CCTA, 1989).
- Threats occurring probabilities are not available in different parts of the world.
- It is not realistic to use cost values for all aspects of losses; for example, a cost value is not appropriate for such issues as embarrassment, loss of goodwill, and legal obligations (Moses, 1992).
- The attribution of cost values to data is very subjective, and thus an ALE based review commences on unsound bases (Moses, 1992; CCTA, 1989).
- A high degree of IT security expertise is required, particularly because no real guidance is offered on security countermeasures (Moses, 1992; Baylus, 1991; CCTA, 1989).
- Full ALE calculations of all application systems and facilities against all possible threats are time consuming and need a lot of manpower effort (Moses, 1992; Baylus, 1991; CCTA, 1989).

Due to these disadvantages, the ALE approach will therefore not be used for cost analysis in the present research. Another approach, which is considered to be more acceptable by disaster recovery experts, is applied (the approach is explained later in section 2.5.8).

2.5.2 Security Assessment Questionnaire

This method was developed by IBM in 1980, and revised in 1985. It consists of fourteen categories, which are divided into three key security areas:

1. Physical Security;
2. Controls and Procedures; and
3. Contingency Planning.

At the end of each of the fourteen categories, a space for rating risk for the entire category is given as extremely low / necessary / acceptable / high. Advantages of this method are that it is brief, and allows the user a quick assessment of an installation's security status. The questionnaire, however, puts more emphasis on security and controls to prevent disasters rather than on recovery issues (NBS, 1985).

In addition to not providing enough emphasis on recovery process, the questionnaire does not give guidance on how to arrive at the risk rating for each category. Also, it does not relate to techniques in calculating maximum allowable downtime and investment required for selecting recovery strategies (NBS, 1985; Danish, 1994). Thus, the questionnaire method does not meet the objectives of the present research.

2.5.3 CCTA Risk Analysis and Management Methodology

The CCTA Risk Analysis and Management Method (CRAMM) was developed by the UK Government's Central Computer and Telecommunications Agency (CCTA) and BIS Information Systems Ltd (BIS), London, UK. The method is embodied in a software support tool which runs on IBM PCs and compatibles. It was produced after examining existing methodologies in order to determine if any, at the time, existed which could be taken for government use. Several methodologies were identified including the ALE approach. However no existing methodology was found to meet their requirements. So the CCTA released its own risk analysis and management approach in 1988 (CCTA, 1989).

The methodology comprises three stages, data for these stages are collected from completed questionnaires. The stages are:

- A. Physical, applications and data asset identification and valuation;
- B. Threat and vulnerability followed by risk assessments; and
- C. Countermeasure identification and selection.

The CRAMM and its software tool are used in many UK government organisations and the CCTA tries to expand its use to private sectors. However, the method and the software have much to be commended, because deficiencies such as difficulties of application to PCs, the need for extensive training, and a brief management summary are addressed (Moses, 1992).

Furthermore, the CRAMM was developed for the concept of risk analysis and management. It does not go into a great level of detail in covering the whole disaster recovery concept. The CRAMM developer, CCTA, acknowledges this drawback by planning to enhance CRAMM in this respect in the future (CCTA, 1989). In addition, CRAMM does not place explicit monetary values on data assets or on the costs of disruption associated with a loss of service. Therefore, the CRAMM approach is not very suitable and will not contribute much to the objectives of the present research.

2.5.4 CCTA IT Infrastructure Contingency Planning Module

This module was developed by the IT Infrastructure Library at the CCTA in 1989. It was developed shortly after the release of CRAMM to be a CCTA guidelines for Contingency Planning (CP). It is aimed at IT Directors, Heads of IT services and senior officers who are responsible for risk management and contingency planning. The module deals with planning to cope with, and recover from, an IT disaster (i.e. loss of service for protracted periods) which requires that work to be moved to an alternative site in a non-routine way. It also provides guidance on safeguarding the existing system (CCTA, 1989).

According to the IT Infrastructure Library (CCTA, 1989), the main application of the module consists of the three following phases: 1) planning, 2) implementation, and 3) post-implementation. The planning phase describes a preliminary task which analyses the risks to a department's IT facilities by using the CCTA Risk Analysis and Management Method (CRAMM), which is explained in the previous section. It also addresses the management approval issue for staffing the contingency planning project and to define terms of reference for it. It also describes the recovery options available and the process of setting up a project team.

The implementation phase covers the development of the contingency plan which includes identifying potential threats and listing critical resources. It also gives some guidance of how and when the plan should be invoked. The final phase, post-implementation, deals with testing and reviewing the plan.

The CCTA IT Infrastructure Contingency Planning Module is currently used by some UK government organisations. Although this module provides acceptable guidelines and steps to follow for contingency planning, it utilises the CRAMM method which suffers from several drawbacks as explained in the previous section. Also, it does not relate to any techniques that help in determining the maximum allowable downtime. This may be because it was developed for the government sector which constitutes non-profit organisations. However, the CCTA announced that, as mentioned earlier, it will enhance the CRAMM approach in terms of the disaster recovery issue in the future. Therefore it would be interesting to see new modifications which may contribute more to the disaster recovery area, especially when some officials of CCTA also indicated that they were thinking of utilising expert system technology in their next version (Danish, 1994). Until these modifications take place, this module does not meet the objectives of the present research

2.5.5 Subhani Model

A recent decision model for calculating the optimal investment required for a disaster recovery plan was developed by Subhani in 1989. The model is made up of two independent steps: 1) determining the Maximum Allowable Downtime (MAD); and 2) applying the developed "Contingency Cost-Response Time Function".

The inputs needed for the first step, MAD determination, are: the loss characteristics, size and risk attitude of the firm, contingency plan cost, threat probability, and contingent loss distribution. The MAD is then substituted in the Contingency Cost-Response Time Function, the second step, to derive an estimate of the optimal investment required for recovery strategy. The two steps are completely independent (Subhani, 1989).

The model is a more accurate approach compared to other disaster recovery methodologies for calculating the required investment. Another major strength of this model is that it does not require a lot of manpower effort for execution. It is also not complex and a high level of recovery skills is not required to implement it.

Having said that, however, the model suffers from the following major drawbacks:

- The probability of threats is one of the inputs used, in the first part of the model, to arrive at the maximum allowable downtime estimate. This is considered to be a drawback because probability of disaster occurrences are not firm and they are not available in many countries (Arnell, 1990; CCTA, 1989).
- It does not include all available recovery strategies and does not present sufficient analysis of the characteristics of those included.
- It does not help the users by giving recommendations in terms of what type of recovery strategy they should select.

The second part of the model, Contingency Cost-Response Time Function, however, is acceptable and it has received no criticisms from disaster recovery experts. In fact, the function was validated, and approved by a panel of 24 disaster recovery experts. Then the function was tested in real-life situations. It was applied to three companies already adopting recovery strategies. The three cases were predicted fairly well (Subhani, 1989).

Since the second part of the model is independent from the first part and is acceptable by disaster recovery experts, the present research will adopt it in calculating the required investment for a recovery strategy (the Contingency Cost-Response Time Function and its application will be explained in more details in both Chapters 5 and 7).

2.5.6 Seminar/Workshop Methodology

The Seminar/Workshop Methodology was created and has been used by DIA*log Management, Inc. It is a comprehensive practical guide for IT managers or those assigned the responsibility of designing, implementing, testing and maintaining Disaster Recovery Plans. It consists of a master plan which is further divided into “miniplans”. The miniplans are listed below (Arnell, 1990):

- Preplanning and Assumption
- Prevention and Security
- Disaster Preparedness
- Disaster Recovery Action Plan
- Training for Disaster Recovery
- Plan Update

Arnell (1990) stated that the primary aim of the Seminar/Workshop Methodology is to enable IT managers and disaster recovery co-ordinators to develop their own in-house disaster recovery plan. Emphasis in the methodology is placed more on plan development. However, the risk analysis issue is by-passed because the methodology

builder believes that 'the need is virtually mandated by an organisation's total dependence on computer' (Arnell, 1990).

Although the Seminar/Workshop Methodology provides skilful guide to disaster recovery co-ordinators, it contains very lengthy procedures and guidelines. The methodology has been developed to deal with all aspects of business interruption preparedness issues. This may be seen as an advantage to some planners, but it is considered to be a time consuming task and it needs much manpower effort. Also, the methodology does not provide management with any indication of how much to spend on disaster recovery planning. Although this is in agreement with some disaster recovery experts who say that the cost analysis of fitting a recovery strategy should be avoided, it is considered to be a potential weakness by others. Therefore, this methodology does not meet the required demands of the present study.

2.5.7 Generic Disaster Recovery Plan Methodology

This methodology was introduced by Carl Jackson, a senior manager with Ernst & Young in Houston, USA. The author cited that today's organisations are making ever-increasing use of information systems technologies in order to provide the most cost effective and efficient services. He added, "while increases in productivity and efficiency are the desired result, we often overlook the pitfalls associated with this dependence on sometimes fragile computer systems to support time-critical business functions." Jackson's methodology is composed of five fundamental steps which are given below (Jackson, 1994):

- 1) Project Initiation;
- 2) Vulnerability Assessment;
- 3) Recovery Alternatives;
- 4) Recovery Plan Development; and
- 5) Recovery Plan Testing and Maintenance.

The Generic Disaster Recovery Plan methodology provides only guidelines just like the previous method (Seminar/Workshop Methodology). However, an obvious advantage of this methodology is that it concentrates only on disaster recovery, not on other aspects such as disaster prevention or security.

The above methodology is prescriptive in nature since it only contains steps and guidelines to follow. It does not provide any recommendations at the end. Furthermore, this methodology lacks any techniques for calculating the required investment on disaster recovery strategies. Another disadvantage of this approach is that it does not show the organisation how to conceive and calculate its maximum allowable downtime. Thus, this method does not fulfil the purpose of the present research.

2.5.8 Commercial Disaster Recovery Software

There are a few vendors who sell disaster recovery planning software that are based on PCs and larger CPUs. These software packages can be used for in-house development. They help to reduce the learning curve and the costs of the disaster recovery planning project. Also, some of them are easy to use (i.e. word processor driven). However commercial software has several drawbacks as shown below (Toigo, 1989; Robinson, 1993):

1. Most software requires the organisation to adopt the methodology of the software author. This is a benefit for organisations whose requirements dovetail with the software features, otherwise it will be inefficient for those who do not.
2. Many commercial software systems set forth a system recovery strategy that presumes the use of hot site services (recovery strategies are explained in chapter 3). Other recovery strategies such as service bureaux or reciprocal agreements are not considered.
3. These software systems are mostly just guidelines and plans. They do not estimate the investment required for adopting a recovery strategy.

4. Most of them are based on what the vendor offers. For example, if a company needs only a mobile site and the vendor does not provide this, the option is excluded from the software system.

Although these software systems are not available for investigation due to their high cost and their location in the US, the above-mentioned criticisms made by disaster recovery experts indicate that they do not meet the objectives of the present research.

2.6 Expert System Technology

Since the arrival of the computer, solutions to problems are usually implemented using conventional systems technology. Since then, individuals have been developing programs to perform rapid calculations, to access data, or to perform modelling of complex process. In the last decade or so, the use of Artificial Intelligence (AI) in many fields has increased. Durkin (1994) defines Artificial Intelligence *as a field of study in computer science that pursues the goal of making a computer reason in a manner similar to a human*. Another special purpose computer programs, subset of AI, called Expert Systems were also applied in many areas. These are programs or systems that employ human knowledge which is captured in a computer to help solve problems that usually require human expertise (Turban, 1992).

Expert systems, or knowledge-based systems, are used to give advice and to make decisions in the light of evidence given to them in much the same way as human experts would be consulted. They can preserve knowledge, increase productivity, capture scarce expertise and make it widely available, improve and speed up decision making, enhance problem solving and provide training with the explanation facility.

With the proliferation of PCs and the introduction of easy-to-use expert systems in recent years, the growth rate of the use of this technology in many areas has been tremendous. According to Durkin (1994), an expert and author of many books and articles in expert systems technology, the opportunity to expand the expert system technology in many areas will be enormous in the near future. The application, benefit

and growth of expert systems have been extensively addressed in the literature and need not be expanded on the present research. (Additional readings of the subject can be obtained from materials located in the Reference section at the end of this document.) However, what most concerns this research is the utilisation of expert systems technology in disaster recovery.

2.6.1 The Use of ES in Disaster Recovery

Both disaster recovery and expert systems are considered to be relatively new fields. In addition to their novelty appearances, their applications have grown tremendously over the last few years. Much of expert systems work in the past was related to medical consultations, computer configurations and engineering. They have been applied with great success in disease diagnosis, fault finding and design. In recent years, expert systems technology was introduced to areas related to the disaster recovery such as risk analysis and disaster prevention.

An example of utilising expert systems in risk analysis is the introduction of an approach called LAVA, which is an acronym for 'Los Alamos Vulnerability and Risk Assessment'. The Los Alamos National Laboratory under the auspices of the US Department of Energy introduced LAVA in 1988 (Moses, 1992). Then several prototype systems were introduced to utilise the expert systems technology in the disaster prevention area. Such attempts are the Prototype Expert System for Disaster Mitigation in the Caribbean (Chin, 1992) and the Knowledge-Based System for Computer Disaster Prevention in IT Centres (Danish, 1994). Looking at these attempts and talking to the author of the latter prototype gave a strong indication to the researcher that utilising expert system technology in disaster recovery is feasible. In fact, the work produced by Danish plays a major role in the decision to proceed with expert system technology in this research.

The literature also shows that there are some new research projects going on, particularly in the US and the UK, to utilise expert system technology in the disaster recovery area. As mentioned earlier, the UK Government's Central Computer and

Telecommunications Agency (CCTA) will utilise expert systems in the next version of its Contingency Planning Module. The recent research which fully applies the expert system technology in disaster recovery has just been introduced during the writing up of this project. The following section explains this research.

2.6.1.1 AUDIT

A prototype expert system for disaster recovery planning auditing called AUDIT was presented in late 1996 in the US. According to Marcella and Rauff (1996), the prototype system was constructed by using the product Level 5, that is an expert system shell. The knowledge base consists entirely of IF-THEN rules. The rules encode the heuristic knowledge of one expert concerning the protection of off-site backup and retrieval of critical data, applications of software and documentation, and system support software. Most of the rules in the system were deterministic, and thus ask for a yes-no or multiple choice response.

The two professional developers of the AUDIT prototype system came up with interesting results to this research. They say that:

'The results of our exploration into the possibility of using expert systems for auditing DRPs are encouraging. We believe that automated auditing DRPs utilising expert systems, is feasible and cost effective in at least three contexts. First, expert systems along the lines of AUDIT could be quite effective as an aid in developing a DRP. Second, an auditing expert system could serve as an inexpensive pre-auditing tool for an organisation. It would provide a means by which a firm could get the "bugs" out of its DRP before incurring the expense of a professional audit. Finally, an expert system for DRP auditing can provide almost continuous investigation of DRPs' (Marcella and Rauff, 1996).

Although the results of AUDIT's development were introduced towards the end of the present research, it presents the following interesting conclusions which would assist in meeting some of the objectives of this study:

- A) Utilising expert systems in disaster recovery area is feasible and cost effective.
- B) Expert systems could serve as inexpensive pre-auditing tools for an organisation.
- C) Expert systems can be used to give some indications of what to expect before seeking, contracting and incurring the expense of a professional audit.

At a first glance, the author of the present research thought that the AUDIT prototype system is similar to the end product of this research. However, after further investigation the AUDIT system is found to be different in the following aspects:

- 1) The AUDIT prototype system does not use or follow any DRP methodology, whereas one of the objectives of the present research is to produce a comprehensive methodology and then to deliver the end-product system based on the developed methodology.
- 2) The system does not address the threats assessment and the business impact assessment issues.
- 3) The AUDIT system does not provide IT managers with a model of estimating the required investment on disaster recovery.
- 4) It also does not provide IT managers with any techniques of estimating the maximum allowable downtime.
- 5) More importantly, it does not include the characteristics of available recovery strategies and, hence, does not assist in recommending a suitable strategy.

Therefore, unless future developments take place, the AUDIT prototype system, which was introduced in the US last year, does not meet most of the objectives of this research. However, it provides the researcher with more confidence that he is on the right path.

2.7 Summary of Literature Review

Organisations are becoming increasingly aware of the value of the computer centre to the overall success and continued operation of a company. They are also becoming more aware of the tremendous liability it imposes when a disaster occurs. Because most businesses nowadays depend heavily on technology and computer systems, disruptions for even a few hours can cause a severe financial loss and can threaten the survival of the company. The growing dependence on computers and the increasing awareness of threats have therefore created the need for disaster recovery.

It was not until as late as 1979 that the concept of disaster recovery was introduced. It first began to appear in the United States with the introduction of the mainframe recovery industry. Since then, several attempts to avoid or minimise the impact of disasters have been made. Disaster prevention and insurance have been adopted as approaches to protect the survival of organisations. To that end several methods have been introduced, particularly in the prevention area. These methods may protect the business from some minor threats, or provide a means of compensation by insuring the assets, but they do not contribute much to keeping the service or business running during, or shortly after, the disaster.

As a result, more attention was devoted to recovery plans rather than the two previous approaches. At the end of the 1980s, some disaster recovery plans and guidelines were introduced, followed by the development of automated approaches utilising the computer systems. Research into applying expert systems technology in the field of disaster recovery was not, however, introduced until the mid-1990s. Indeed, some officials of the main UK government computer agency, CCTA, have indicated that they will use expert system technology in the next version of CRAMM.

After an extensive analysis of the literature given above and after examining the requirements and objectives of this research, it is clear that the previous work does not adequately and fully address the recovery strategy selection problem. The

published work has therefore a number of deficiencies in terms of meeting the objectives of this research as follows:

1. There is no comprehensive methodology for IT managers and officers who are responsible for contingency planning that includes:
 - i) a method for classifying threats in which a greater contribution to the recovery strategy selection problem would be obtained;
 - ii) a full business impact analysis including the calculation of the exact maximum allowable downtime;
 - iii) an approach to estimating how much to spend on disaster recovery; and
 - iv) methods of constructing organisational requirements and recovery strategy characteristics in such a way that they can be utilised by expert systems technology to select the most appropriate recovery strategy.

2. There is a still a lack of a full use of expert systems technology in the disaster recovery area, to assist IT managers in decision-making.

Therefore, the present research is intended to present a solution to the recovery strategy selection problem. The description of available recovery strategies, the proposed methodology, the expert systems technology and the proposed prototype system are all parts of that solution and they are presented in the following chapters.

Chapter 3

Disaster Recovery Strategy Options

3.1 Introduction

The aim of this research is to help organisations select a suitable recovery strategy to keep their businesses going during an unexpected outage. But before developing the required methodology to help IT managers and others more towards the above aim, it is important to explain the recovery strategy options currently available in the field of disaster recovery and those which are expected to become available in the near future.

Any successful disaster recovery plan should include adequate backup procedures for recovering the lost applications and systems. Selecting an alternative recovery site to recover the computer systems and keep the business running during the outage is part of these backup procedures. The alternative recovery site is called (in disaster recovery term) *recovery strategy*. In this chapter, major recovery strategy options available in the disaster recovery area are briefly described, with their respective advantages and disadvantages outlined. Then a comparison between all the various options is carried out. In addition, since the computer-related recovery industry has only been around for a relatively short period of time, there are several recovery strategy options the feasibility of which are still under investigation. These strategies are also briefly described in this chapter.

3.2 Recovery Strategy Types

There is a wide range of recovery strategy options available for consideration. They may be grouped into the following five types (see Table 3.1): null strategy, internal,

mutual aid, co-operative, and commercial. The *null strategy*, as the name may indicate, is having no backup procedures at all. *Internal* recovery strategies are those which can be performed within the organisation. The third type, *mutual aid*, is signing an agreement between two companies or through a third party to use each other's computer facilities in the event of a disaster. The *co-operative* type is when two or more companies share the cost of managing and maintaining an alternative site. However, *commercial* recovery strategies are the most commonly adopted type. It is a recovery strategy that is provided by a commercial company which spreads the cost across a number of subscribers (Baylus, 1991).

Table 3.1 - Types of Recovery Strategy

Recovery Strategy Type	Examples
Null	Doing nothing
Internal	Manual procedures, withdrawal of service, duplicate site.
Mutual Aid	Informal mutual aid, reciprocal agreement, time broker
Co-operative	Co-operative hot site, co-operative cold site.
Commercial	Service bureau, commercial hot site, commercial warm site, commercial cold site, hardware vendor, realtime recovery, mobile hot site, portable site.

A full account of all the strategies would take a book by itself. Thus, in the following sections the researcher identifies and briefly explains major recovery strategies to give the reader some knowledge of potential recovery options in the disaster recovery field.

3.3 Null Strategy

As would be the case with any alternatives, there is always the null strategy option. The null strategy option in the disaster recovery area means simply that there is no backup at all. In this case, management has made a decision to ignore all potential

hazards, and will not fund any strategy to backup their data processing facilities. Those who chose this option usually believe in the saying 'it will not happen to me'. There are few, if any, organisations that can justifiably adopt this option. They justify the null strategy because they rely on insurance coverage. As explained in Chapter 2, this is not considered by many disaster recovery experts to be a genuine recovery alternative. It should be noted, as described in section 2.4, that information systems insurance is a complex issue; comprehensive cover is rarely offered or may be prohibitively expensive (Orr, 1988).

Other organisations may rely on preventive measurements to minimise a disaster's impact on any computer facility without considering the recovery procedures. Although the disaster prevention measurements may seem sound and robust for some organisations, even the best avoidance plans cannot prevent every disaster (Redmond, Luongo & Tietz, 1996).

In summary, this strategy may appear financially attractive, but any organisation that is able to function without its computer services for a long time after a disaster must ask itself whether it needs them at all.

3.4 Internal Strategies

Internal recovery strategies are those which can be performed within an organisation. The possible internal recovery strategies are manual procedures, withdrawal from computer services, and establishing a duplicate site.

3.4.1 Manual Procedures

The manual procedure strategy cannot be directly dismissed. It holds promise, particularly for those organisations that have little dependency on computers. Manual procedures may be in place for the short term disaster. For example, banks always have manual procedures for making deposits and handling withdraws when their computer services are down for any reason (Arnell, 1990).

If manual operation is the strategy to be employed, it must be thought out completely. Manual forms must be ready. Provision for temporary staffing may also be necessary. However, if functions have been supported by automation for a long period of time, manual procedures may have been forgotten. Even if the data and paper forms are available, the workload may be too large, the time too short, and the staff's memory of the old manual procedures too dim.

3.4.2 Withdrawal of Services

The withdrawal of services strategy is applied when there is no urgent need to recover and run the computer systems during or immediately after a catastrophe. This strategy is normally considered for long-range work which can be transported to another location and run whenever feasible. There are some application systems which may fit here, such as long-range analytical and planning work, small business programs, and some types of research and development work, where the obvious strategy is simply not to perform the job until the computer facilities have recovered from the disaster. However, this strategy is considered to be unacceptable by customers. Also, since the job will not be performed until the computer is up again, staff productivity will decrease dramatically (Baylus, 1991).

3.4.3 Duplicate Site

The duplicate (redundant) site strategy is setting up an entire alternative data centre in another location. Many organisations consider their security and dependence on computer systems to be so great that they cannot afford the time it takes to re-establish service at a commercial backup site. Such organisations are those involved in operations related to national security, defence production, and critical financial activities. For these organisations, building another duplicate site is the only acceptable alternative for disaster recovery backup, despite its high start-up cost (Hyde, 93).

In the event of a disaster, redundant systems at a separate facility are brought on-line. Users are either transported to an operations centre that is co-located to the backup site or are provided with remote access terminals and printers and connected to the backup CPU via communications. For example, Sears (the giant retailer in the USA) has a data centre in Chicago suburb (north of USA) dedicated to running the Sears nation-wide computer systems. A mirror site in Dallas (south of USA) is also in place, ready to take over command of the computer systems in the event of disaster hitting the Chicago centre (Burch & Grudnitski, 1989).

Organisations which have built their own recovery site have justified the often considerable expense of the second computer centre by using it for research and development, training, and overflow work such as large batch processes. Besides being the most reliable method of systems backup, the duplicate site strategy is also the most expensive (Hyde, 1993; Toigo, 1989).

3.5 Mutual Aid Agreements

The mutual aid strategies entail having a formal or informal agreement directly between two companies or through a third party, to use each other's computer facilities in the event of a disaster. The possible recovery strategies within this category are informal mutual aid, reciprocal agreement, and time broker.

3.5.1 Informal Mutual Aid

It is possible to have agreements on an informal basis. This is usually done when there is no complexity involved in recovering systems. There are some organisations which provide assistance when a neighbouring company is hit by a disaster. Another example is when one company is a major customer of another. The vendor party might not wish to sign a contract, but would certainly like to help during an emergency. This strategy is becoming more feasible as vendors increasingly provide computer services to their clients (Baylus, 1991).

3.5.2 Reciprocal Agreement

The shared contingency, reciprocal or mutual aid, agreement is when two or more organisations, usually in the same industry, having identical or similar computer environments, formally agree to use each other's computer resources if either of them suffers a disaster. One of the major factors which has to be considered when adopting this strategy is location. Organisations using this strategy cannot be within the same locality. They have to be geographically far apart. For example, they may be in the same city but not in the same street, if they are to have confidence that they are both protected from a regional disaster (Hyde, 1993).

The issue of security also has to be considered when running applications in the "backup" organisation computer centre (either during testing or after a real disaster). An obvious disadvantage of a reciprocal agreement is testing, as this usually generates unwelcome disruptions in the "backup" organisation. However, security and testing issues are usually covered in the contract and with the disaster recovery co-ordinators (Baylus, 1991). The cost of this strategy is relatively low because each organisation uses the other firm's resources.

Despite its complications and disadvantages, this strategy is still the only realistic and affordable option for some organisations. It is certainly feasible for medium and small organisations and for larger firms with numerous subsidiaries (Toigo, 1989; Hyde, 1993).

3.5.3 Time Brokers

The time broker strategy involves a third party keeping a list of firms, with available computer resources and time, who would be willing to provide such resources and time to other organisations on a temporary basis. The broker enters into contracts with all the parties involved and can guarantee the availability of predetermined computer facilities. All decisions in the use of the facility and the contractual terms are

made with the broker. There is no direct negotiation between the other two parties (Baylus, 1991).

This is a type of mutual aid agreement in which the computer users do not have to take the time to search for compatible computers. The broker usually charges a monthly fee. This strategy is an inexpensive approach and has very favourable logistics if several companies in an area are brought into the agreement. It can work if good relationships develop between the companies that are involved. However, there are serious problems with maintaining system compatibility over time, with being assured of availability when testing, and if all parties are subject to the same disaster (Hyde, 1993; Baylus, 1991).

3.6 Co-operative Recovery Strategies

The co-operative recovery strategy options category covers the situation when two or more companies share the cost of owning, managing and maintaining an alternative site. This normally involves one of two options: co-operative cold site and co-operative hot site.

3.6.1 Co-operative Cold Site

This strategy becomes an option when several organisations, usually within the same industry, form a group and agree to build an alternative “empty shell” site for long term occupation. The co-operative empty shell, another term, is similar to the commercial cold site provided by disaster recovery vendors. It is a building or a computer room on a fixed site. The room is equipped with the necessary power, environmental controls and telecommunications connections. The site is managed and maintained by a dedicated team assembled from the group (Arnell, 1990).

When a disaster strikes one member of the group, the affected organisation starts the recovery procedures by ordering and installing the required hardware and software to resume its business. The affected organisation can stay as long as needed, normally

until it rebuilds its original computer centre site. Although the size of the co-operative cold site is designed so that it can accommodate two or three organisations at the same time, the members of the group are carefully selected to be geographically remote from one another in order to avoid the chances of two organisations being affected by the same disaster.

The adoption of the co-operative recovery strategy option is usually feasible for large organisations. Small organisations cannot afford the cost of managing and maintaining the alternative site. Although this option is considered to be suitable as a long term strategy and is relatively inexpensive for large organisations, the process of finding suitable members to form the required group is not easy. It may take a considerable period of time to reach a mutually acceptable agreement with all the members.

3.6.2 Co-operative Hot Site

This strategy employs the same concept as the co-operative cold site in that several organisations form a group and agree to build an alternative site. However, the co-operative hot site is a building or a computer room which is fully equipped with necessary hardware and software for fast recovery. In addition, the site is only utilised for a short period not exceeding six, or at the most, eight weeks.

The co-operative hot site's membership comprises organisations that use identical computer platform environment. For some co-operative hot sites, the recovery services, however, are not limited only to those organisations which originally formed the group, but are also available to secondary level members. This is executed with a contractual stipulation that if an original member experiences a disaster, the secondary member is asked to leave the facility at short notice. Members of the co-operative hot site carry out this technique, on allowing secondary memberships, to cover the expenses of managing and maintaining the hot site. Therefore, subscribing to a co-operative hot site as a secondary member is much cheaper than subscribing to a commercial hot site (Arnell, 1990).

The virtue of the co-operative hot site strategy is that cost-sharing and the further acceptance of secondary level membership reduce the expenditure of the primary members. Also, the secondary members benefit from the relatively low cost because the co-operative hot site is established as a non-profit site. However, having to be ejected from the site at short notice is a drawback for the secondary members. Finally, as with the co-operative cold site strategy, finding members to form the group and reach a co-operative agreement may be difficult and take a long time.

3.7 Commercial Recovery Strategies

This category contains the majority of the recovery strategy options. Commercial recovery strategies are those provided by commercial companies which spread the cost across a number of subscribers. The various possible commercial recovery options in the disaster recovery market are described in the following sections.

3.7.1 Service Bureaux

Service bureaux are commercial data centres which offer shared use of computer systems. They offer both batch services and on-line or time-sharing services. They vary from two-person operations, with little equipment, to large corporations with multiple equipment. Some of the problems of using service bureaux are similar to those found with mutual aid agreements (e.g. security and location). Some companies may not even have a service bureau located in the same city to provide needed recovery services. Another problem is that there may be configuration changes from time to time which cause extra effort and cost (Blair, 1987; Baylus, 1991; Hyde, 1993).

However, service bureaux have a unique capability to satisfy many processing requirements and should be given careful consideration as one of the strategies to be used. Risk, cost, and effectiveness are moderate for this option. The service bureau strategy would probably be an acceptable arrangement for a company that does not require its systems to be recovered within a very short time scale (Hyde, 1993).

3.7.2 Hardware Vendors

The hardware vendor strategy involves an agreement with a vendor to utilise its facilities for regular backup procedures in the event of a disaster. Vendors always have equipment available to be used at the time of a disaster. The equipment may be at the demonstration area centre, test centre, or at internal or sales sites.

Historically, computer vendors' efforts have been excellent in providing the necessary equipment when it is needed. Also, hardware vendor strategies are often considered to be attractive because of their pre-tested compatibility and support. Although the systems may possibly be compatible at the time of the agreement, vendor facilities are subject to continuous modifications because of the rapid changes in technology (Arnell, 1990; Baylus, 1991).

In addition, there are some problems associated with this strategy. Vendors' facilities are usually showcases or sales sites, which make them subject to serious security problems. Another problem is that the hardware vendor does not have the relevant experience in disasters and recovery procedures. In fact, some of them indicate in advance that they will not agree to participate in the disaster recovery procedures.

3.7.3 Commercial Cold Site

The commercial cold site is a commercial computer-ready room held in reserve for the subscriber's system. It usually contains a power supply, a raised floor, and air-conditioning units. It does not contain computer processors or peripherals, although it may be equipped with dial-up lines for a communication network. Cold sites can also be private. Any company, if it can afford to, would build its own empty room at a site remote from all other data centres. The empty site needs not to be unproductive space because it can be used as a warehouse for supplies and equipment (Baylus, 1991; Hyde, 1993; , 1995).

In this strategy, separate arrangements need to be made to acquire the necessary hardware and software to run at the cold site. Ordering, shipping and installing the required equipment may take several days. One of the advantages of this option is that access should be available to the site almost immediately. The cost is relatively low compared to other options such as service bureaux and hot sites. In the commercial cold site strategy, there are many companies who may subscribe to the same cold site. This will reduce the cost to any individual member. However if one disaster affects several subscribers, the services may not be as promised.

A cold site is best utilised in conjunction with a hot site. The hot site is used while the new equipment is shipped and installed at the cold site. When the new system is brought up at the cold site, processing can be transferred from the hot site to the cold site location. In fact most hot site providers also offer the cold site option. The cold site is usually available for a period from twelve to sixty months (Schreider, 1995).

However, this strategy has several problems. One of the more serious problems is that there is no way the organisation can test its disaster recovery plan. When the disaster actually happens, they can only follow the steps laid out in the plan. Since it will be the first time they have operated the plan, they will be confronted with many unexpected issues. This will affect the performance of the plan, which will, in turn, affect the overall business. Another problem is the inevitable delay in obtaining replacement equipment. When ordering equipment after the disaster, it must be realised that equipping a cold site will take at least a week.

3.7.4 Commercial Warm Site

The warm site strategy is relatively new to the market. It has been established to take its place between the cold site and hot site strategies. It is similar to the service bureau strategy and can be relatively expensive. Contracts are usually for a one year period. Warm sites are better equipped than cold sites. In addition to the usual cold site capabilities, they have telecommunications capabilities. They may also have lower specification hard disks and peripherals. IBM has taken warm sites further than most

and has provided services very close to hot site specification. IBM can occupy their warm sites for up to twelve weeks (Hiles, 1992; IBM Report, 1993a).

3.7.5 Commercial Hot Site

A commercial hot site is a complete data centre, from the commercial point of view, fully equipped with different sizes of processors, peripherals, communications, networks, and any necessary equipment. Hot sites are usually equipped to run any application that is compatible with its hardware and operating system (Robinson, 1993). Professional personnel are usually available to assist the organisation's operation team in their efforts to get the system up and running again. In addition, specialised equipment can be provided to satisfy the customer's backup requirements.

Once a disaster is declared, the affected organisation sends its backup media to the hot site. Applications are mounted and tested, users are provided with terminal and modems at the location, and a data processing service is restored. Location of hot sites is not an issue. Network capabilities allow a subscriber to communicate with the hot site remotely, thereby eliminating the need to move employees away from their families (Hyde, 1993).

Maintaining a commercial hot site is an expensive disaster recovery strategy, but a necessary one for many organisations with critical applications. Financial-oriented organisations tend to utilise hot sites more frequently due to the critical nature of their operations. In fact, over 52% of all hot site recoveries involve this type of organisation (Schreider, 1995).

The hot site strategy has proven to be an effective strategy for recovering computer centres. Since 1982, 582 successful recoveries were completed at over 25 hot site locations in the USA. Until now, the industry has comprised 31 companies, representing the majority of the hot site providers and generating subscription fees of \$620+ million annually in the USA alone. The majority of recoveries have occurred at the major hot site vendors: IBM, SunGard, and Comdisco. These three vendors alone

have supported over 67% of all disaster recoveries in the USA. (Robinson, 1993; Schreider, 1995)

Hot sites are increasingly becoming the recovery option preferred by medium to large organisations. Also, hot site vendors are beginning to provide end-user and PC/LAN-based recovery services for smaller organisations which can afford the cost of having a hot site. Hot sites are recommended for any organisation which cannot manage without its computer system for more than one day (Toigo, 1989; Hyde, 1993; Schreider, 1995).

3.7.6 Portable Site

A portable site is transported to the affected organisation, in the event of a disaster, and built on a site prearranged with the subscriber, normally the car park. The amount of accommodation provided is tailored to the size of the configuration required. Electricity supply and telecommunication links may be required from the original site to the portable site (CCTA, 1989; Hiles, 1992).

The cost of this option is usually the same as the cold site, which is inexpensive. The advantage of this strategy, apart from the relatively low cost, is that the portable site can be constructed adjacent to the home site. The portable site strategy can support organisations with smaller hardware configurations which can manage without its computer systems for as long as one week or more (Schreider, 1995). This strategy is commonly used as a long term strategy alongside another short-term strategy such as commercial or mobile hot sites.

A potential disadvantage of the portable site lies in the need to provide a suitable secure location near the home office. This strategy also has similar disadvantages to the cold site strategy, such as the time needed to construct and commission the site, which can vary from seven to ten days. In addition, there is no way an organisation can practice and rehearse its disaster recovery plan in advance (CCTA, 1989).

3.7.7 Mobile Hot Site

A mobile hot site strategy is a stand-alone unit on a mobile trailer. A recovery vendor contracts to deliver an agreed system to the customer site, within a certain time limit. The trailer is fitted out as a computer room with the necessary environment services. The organisation needs to provide a secure area on which the trailer can be parked. Like the portable site, electricity supply and telecommunication links are probably required from the original site to the mobile site (CCTA, 1989).

In addition to the advantage of having the mobile site adjacent to the home site, it can be brought into use as a very quick response to a call for help, probably by the next day. A major disadvantage of this option is that it can only accommodate a limited amount of hardware platforms. Also, a location has to be provided near the original site. This option is usually recommended to small businesses as well as bank branches.

3.7.8 Realtime Recovery

Traditional forms of disaster recovery strategies such as hot and cold sites and mobile options remain the correct solutions for many companies. There are, however, a growing number of organisations with business continuity requirements which demand more immediate systems recovery.

Currently, systems recovery is based on creating backup tapes (vital records) at some particular point in time. This affords the organisation recovery to the last backup. Traditionally, these backups have been sent off-site by road. At the time of a disaster, these tapes must be retrieved and then dispatched to the recovery site. This can be financially hazardous, time-consuming, devastating to the reputation of the organisation, and vulnerable to human error.

To avoid all these problems a new concept of realtime recovery has been introduced to the field of disaster recovery. It is a strategy which has yet to be employed on a major scale in the US and Europe. The primary reason is cost. Although many clients

would like to take advantage of this type of service, the high cost has kept all but a few from implementing it. The realtime recovery can be done by either remote journaling or electronic vaulting which provides the ability to capture the intra-day transactions and transmit them to a commercial hot site. By using the realtime recovery strategy, the customer has the ability to recover to the actual point of failure, minimising the manual effort needed to recreate information (Ratliff, 1994). Realtime recovery is emerging as a popular service. However, it will take a few more years for it to become cost-effective and then more widely accepted (Schreider, 1995).

To have an easy overview of the various options, list of the major advantages and disadvantages of most recovery strategies are shown in Table 3.2.

Table 3.2 - Major Advantages and Disadvantages of Recovery Strategy Options

Rec. Strategy	Advantages	Disadvantages
Null Strategy	<ul style="list-style-type: none"> •No preliminary cost •Good for long range analytical and planning work 	<ul style="list-style-type: none"> •Very high risk that may leads to loss of business •Unacceptable by customers
Manual Procedure	<ul style="list-style-type: none"> •Keeps some customer service •Preliminary cost is low 	<ul style="list-style-type: none"> •Impossible for on-line operations •Procedures could be forgotten •Labour intensive
Withdrawal of services	<ul style="list-style-type: none"> •Relieves work load •Fits long range analytical work 	<ul style="list-style-type: none"> •Bad productivity •Unacceptable for customer
Duplicate site	<ul style="list-style-type: none"> •Company security standards •Under management control •Familiarity with work loads •Test runs any time 	<ul style="list-style-type: none"> •Very expensive •More sites to manage and maintain
Informal mutual aid	<ul style="list-style-type: none"> •Little cost •Good relationships between firms may be established 	<ul style="list-style-type: none"> •No contractual arrangement •Not reliable

Table 3.2 - Continued

Reciprocal agreement	<ul style="list-style-type: none"> • Little or no cost • Good for small to medium firms • Good relationships between firms may be established 	<ul style="list-style-type: none"> • Limited short-term occupancy • Unwelcome testing from other firm • Security and location should be carefully considered
Time broker	<ul style="list-style-type: none"> • Good contractual arrangement • Inexpensive 	<ul style="list-style-type: none"> • May be configuration changes • Unwelcome testing from other firm • Security is not under control
Co-operative cold site	<ul style="list-style-type: none"> • Good for large organisations • Under management control • Long-term strategy • Can be leased to secondary members 	<ul style="list-style-type: none"> • Difficult to form group of participant companies • No pre-testing • Not for critical applications
Co-operative hot site	<ul style="list-style-type: none"> • Fast recovery • Allow pre-testing • Under management control • Can be leased to secondary members 	<ul style="list-style-type: none"> • Expensive • High cost of management and maintenance • Difficult to form group participant companies
Service bureau	<ul style="list-style-type: none"> • Low cost until disaster • Service bureau staff support • Telecommunications capabilities 	<ul style="list-style-type: none"> • May be busy when needed especially in major disasters • Operations not under control • May be configuration changes • Security is not under control
Hardware Vendors	<ul style="list-style-type: none"> • Available pre-testing • Good relation with the vendor 	<ul style="list-style-type: none"> • Compatibility may change • Little experience in disaster recovery procedures • Security problems

Table 3.2 - Continued

Commercial cold site	<ul style="list-style-type: none"> • Relatively inexpensive • Under management control • Long-time occupancy 	<ul style="list-style-type: none"> • Takes long time to get computer environment ready • No pre-testing • No control of hardware shipment • Not for critical applications
Commercial warm site	<ul style="list-style-type: none"> • Better equipped than cold sites • Telecommunication capabilities 	<ul style="list-style-type: none"> • Relatively expensive • Not for critical applications • Not for large equipment
Commercial hot site	<ul style="list-style-type: none"> • Immediate access • Short time to restart computer environment • Pre-setting can be done 	<ul style="list-style-type: none"> • Expensive • Limited short-time occupancy • Not under management control • Security is not under control
Portable site	<ul style="list-style-type: none"> • Relatively inexpensive • No need to relocate staff • Long-time occupancy 	<ul style="list-style-type: none"> • Space may not be available • Takes long time to get computer environment ready • No pre-testing • No control of hardware shipment • Not for critical applications
Mobile hot site	<ul style="list-style-type: none"> • No need to relocate staff • Can be brought in a short notice • Cheaper than fixed hot site • Good solutions for small and medium companies 	<ul style="list-style-type: none"> • Adjacent space may not be available • Limited number of hardware • Limited number of staff at site • Usually only for small to medium companies
Realtime recovery	<ul style="list-style-type: none"> • Recovery at the actual point of failure • Good solution for companies with very vital applications 	<ul style="list-style-type: none"> • Very expensive • Not available in every country

3.8 Comparison between Strategies

In the previous sections, all recovery strategies were described, with the advantages and disadvantages in each case to provide a full picture of the range of possible options. As an additional step towards the aim of this research, a full comparison among the strategies has also been carried out across the full spectrum of the disaster recovery strategy selection process. The comparison includes several issues which must be considered when a company comes to choose between the available recovery strategy options. These issues are divided into three categories: 1) availability, 2) operation; and 3) capacity and cost.

3.8.1 Availability

Table 3.3, Recovery Strategy Options - Availability, shows that only sites internally managed or controlled can be counted on for immediate availability. There is no time limit for utilising these strategies. They can be used as long as it takes until the original site is ready. Also, internal strategies are available for testing at any time and as many times as desired. Several commercial sites and co-operative sites can be made available within 24 hours. Any of the strategies that require adjusting or movement of equipment will naturally take several days. Some of the commercial sites and mutual aid sites usually have a time limit on the duration of using the site.

The testing availability varies from one strategy to another. As stated before, internal strategies have the best testing availability. However, there are also other strategies which provide good testing availability. In fact, many commercial sites will not continue contractual arrangements unless testing has been performed (Arnell, 1990). There are some options which have poor testing availability, such as hardware vendor strategy, or others where tests can not be performed, such as portable and cold sites.

Table 3.3 - Recovery Strategy Options - Availability

Recovery Strategy	Availability for use	Usage duration	Testing availability
Withdrawal of service	6 hrs or less	Long	Good
Duplicate site	6 hrs or less	Long	Good
Reciprocal agreement	24 hrs	Short	Poor
Time broker	2 days	Short	Poor
Co-operative cold site	More than 7 days	Long	No test
Co-operative hot site	24 hours	Short	Good
Service bureau	24 hrs	Short	Good
Hardware vendor	3 days	Short	Poor
Commercial cold site	7 days or more	Long	No test
Warm site	2 days or more	Long	Fair
Commercial hot site	24 hrs or more	Short	Good
Mobile hot site	24 hrs or more	Short	Fair
Portable site	7 days or more	Long	No test
Realtime recovery	6 hrs or less	Short	Good

3.8.2 Operational

Table 3.4, Recovery Strategies Options - Operational, shows that those strategies that are managed internally have the full measure of management control and security guaranteed. This does not mean that other strategies, such as the commercial type, may not have good security and control. In fact, because of competition, their actual controls might be superior to internal controls. However, they are not under the subscriber management's control. On the contrary, mutual aid strategies usually have the least security and management control since the alternative computer centre is controlled and managed by non-specialised and non-profit making organisation (Baylus, 1991; Arnell, 1990).

The scope of modifying hardware and operating systems is usually limited in commercial and mutual aid strategies. This is because there are other subscribers who are using the same equipment and software. Any type of modification on equipment

or software, however, can be allowed in co-operative strategies. In strategies which are selected for long term recovery, such as cold site, the subscriber has the choice to install whatever equipment he needs.

Full service support is available in many commercial strategies. There are other commercial strategies which limit their service supports, such as service bureaux and hardware vendors. This is because they are not fully dedicated to the disaster recovery business (Arnell 1990). Because of the limited recovery manpower that is associated with mutual aid agreements, it is not possible to provide service support for the recovery activities. Similarly, in the co-operative strategies group, external recovery personnel support is not available unless the personnel assigned to the co-operative site are well-trained to overcome any recovery problems. Finally, the service support issue is not a problem with internal strategy options.

Table 3.4 - Recovery Strategy Options - Operational

Recovery Strategy	Security	Control by own staff	Ability to modify SW/HW	Service support
Withdrawal of service	Very good	Strong	No	Yes
Duplicate site	Very good	Strong	Yes	Yes
Reciprocal agreement	Poor	Fair	No	No
Time broker	Poor	Weak	No	No
Co-operative cold site	Good	Good	Yes	No
Co-operative hot site	Good	Good	Yes	No
Service bureau	Poor	Weak	No	Limited
Hardware vendor	Poor	Weak	No	Limited
Cold site	Good	Good	Yes	Yes
Warm site	Good	Good	Little	Yes
Hot site	Good	Good	Little	Yes
Mobile hot site	Good	Good	Little	Yes
Portable site	Good	Good	Yes	Yes
Realtime recovery	Good	Good	Little	Yes

3.8.3 Physical Capacity and Cost

They are other criteria that should be considered when selecting a recovery strategy, including physical capacity and associated costs. Large organisations usually need larger recovery work areas for their computer centre staff. Co-operative strategies and some commercial strategies usually have sufficient work space to accommodate staff from the affected site. On the other hand, mutual aid agreements and small recovery strategies, such as the mobile hot site, offer very limited working space. Table 3.5 shows the physical capacity offered by each strategy.

Another important factor is cost, which not only varies between strategies but also between vendors offering the same recovery strategy. The present research, however, only analyses and compares costs between strategies. It is not one of the objectives to analyse price variation among vendors, although this could be recommended as a possibility for future research.

There are three basic costs associated with adopting a recovery strategy. The first two costs are pre-disaster expenses; the third is a post-disaster expense (Toigo, 1989; Arnell, 1990).

- *Initial cost.* The cost of initial set-up, which includes membership, construction, additional equipment and additional software.
- *Ongoing cost.* The cost of maintaining and operating the facility, including rent, ongoing backup operations, and additional testing.
- *Activation cost* (also called usage cost). This involves the actual use of the facilities, including disaster notification, service support, and overtime.

Table 3.5 gives a general indication of the levels of cost that might be expected for all recovery strategy options. It is clear that the initial cost of building an additional duplicate site for a company is very high, whereas ongoing and activation costs are low because the site is under the company's management. Similarly, the initial cost of having a co-operative site is high and the ongoing and activation costs are relatively

low because the expenses are distributed among several members. In commercial hot sites, all types of cost are high. On the other hand, mutual agreement options have low costs because they are based on using each other's facilities.

It is important; however, to say that costs alone should not determine the choice of a recovery method. As has been stated before, the crucial consideration is to ensure the continuation of critical processing and to provide the time necessary to recover from an adverse incident (Arnell, 1990; Baylus, 1991).

Table 3.5 - Recovery Strategy Options - Physical Capacity and Cost

Recovery Strategy	Physical capacity	Initial cost	Ongoing Cost	Activation cost
Withdrawal of service	-	No cost	No cost	Low
Duplicate site	Good	Very high	Low	Low
Reciprocal agreement	Limited	No cost	Low	Low
Time broker	Limited	Medium	Low	High
Co-operative cold site	Good	Medium	Low	Low
Co-operative hot site	Good	High	Medium	Low
Service bureau	Limited	Medium	Low	High
Hardware vendor	Limited	Low	Low	Medium
Cold site	Good	Low	Low	Medium
Warm site	Good	Medium	Medium	Medium
Hot site	Good	High	High	High
Mobile hot site	Limited	Medium	Low	High
Portable site	Limited	Low	Low	Medium
Realtime recovery	Good	Very high	Very high	High

3.9 Prospective Recovery Strategies

Since the computer-related recovery industry has only been around for a relatively short period of time, there are several recovery strategy options which are still at an early stage in the investigation of their feasibility. These strategies have only just been

introduced during the preparation of this dissertation. They are not yet fully recognised by disaster recovery experts, but may become potential recovery strategies in the near future. Therefore, they are not included in the selection process described in Chapter 5. These new strategies are explained in the following paragraphs.

OmniCentric Hot Site

The OmniCentric strategy is a new concept, and term, in the disaster recovery industry. It is a term created from two words with opposite meanings. According to Powel (1997), Omni means all, every thing, present in all places, having no limits. Centric means centre, having a centre, focused. Powel presents the OmniCentric architecture as a jigsaw puzzle. Omni represents the overall image of the jigsaw puzzle picture, and Centric represents the individual puzzle pieces. An example of an OmniCentric hot site might have the central data processing servers at a recovery site in city A, the system's operations and applications recovered at a site in city B, a large number of user departments recovered at sites in city C, and network servers and technical support at a recovery site in city D. The required recovery elements at the various recovery sites would be interconnected using a backbone network, as well as dial up access. If one of the locations (pieces) is affected, for example user departments in city C, then they can be reconnected back to the organisation's unaffected pieces to complete the whole corporate jigsaw puzzle (Powel, 1997).

The OmniCentric strategy deals with the organisation's locations as one entity. Therefore, integrity is assured when recovering between locations. Recovery can be achieved without revealing that the company has experienced a disaster. However, the strategy only fits organisations with multiple locations. Although the strategy looks promising for such organisations, issues such as security, reliability and full integrity need to be clarified in the near future.

Internet

Because of the rapidly growing use of the World Wide Web, some companies are already offering services for backing up and restoring data and applications over the Internet. The idea is that an Internet Service Provider, either in the recovery business or in related areas, would accommodate some backup services through the World Wide Web to small businesses or PC type of application holders. Subscribers to these services would download the backup programs from a service provider's WWW site, and then register on-line. Once registered, users would specify a daily backup schedule, after which the service would begin performing the on-line backups automatically (Schreider, 1996).

In the event of a disaster, users could restore their backup data over the Internet and download it at their recovery location. The cost of this type of strategy usually involves a reasonable monthly service fee based on the quantity of compressed bytes of data backed up.

Although this type of recovery strategy appears financially attractive for some users, the service is not yet adequately recognised by disaster recovery experts. Moreover, there are several problems associated with the Internet. The sanctity of the data once it has been sent across the Internet, the threat of hackers accessing the company through the Internet, virus-attacks, legal issues, and communication issues are all potential problems that need to be further clarified and resolved (Schreider, 1996).

Employees' Homes

For organisations that have thousands of employees located in one area, finding a work area large enough to accommodate them in the event of a disaster is a major problem. This issue was discussed at a nation-wide teleconference hosted by SunGard Recovery Services in 27 of May 1993, which brought together financial executives and disaster recovery experts to identify and discuss recovery issues. The participants expressed the hope that in the very near future a new recovery strategy would be

devised to overcome the problem of relocating employees during a disaster. They believe the future solution to the above-mentioned problem may be recovery from the homes of key personnel. Business can be resumed from the homes of staff using telecommuting or telecommunication facilities already built into certain employees' homes. Employees do not need to go their offices but they can be connected to the recovery site or to a large server in one of the employees' homes (Datapro report, 1993).

3.10 Concluding Remarks

In this chapter, major recovery strategy options in the disaster recovery area were described, with their respective advantages and disadvantages outlined. Then a comparison between the various options was carried out. In the following chapter, a case study is carried out to show the importance of adopting recovery strategies and to identify major problems facing IT managers regarding disaster recovery issues.

Chapter 4

Case Study in DRPs

4.1 Introduction

In the previous chapters, several surveys were reported as showing the importance and the strong need for Disaster Recovery Plans (DRPs). Some of the data in those surveys, particularly that related to maximum allowable downtimes, which was carried out in 1979 by the University of Minnesota, are important to the design of the methodology and the prototype expert system produced by this research. However, most of these studies were done a long time ago. Since then substantial changes in computer technology have occurred and corporate dependency on computers has grown. Therefore, a new survey is required to produce up-to-date data on which to develop a methodology that would be realistic and acceptable to organisations today.

This chapter explains the results attained from a study recently conducted in Kuwait. It looks at the disaster preparedness of Kuwaiti organisations before and after the Iraqi invasion disaster in August 1990. It explores many issues such as maximum allowable downtimes, the losses that organisations may face, recovery strategies, the necessity of testing, off-site backup strategy, and other related issues. Some of the material reported in this study, such as testing and off-site storage, may seem peripheral to the present research but they are included to present a global picture of disaster recovery plans.

4.2 Objectives of the Case Study

On the dawn of the 2nd of August 1990 an international crisis started in the Middle East. More than 150,000 troops of the Iraqi regime crossed the border of the state of Kuwait toward its capital, Kuwait City. Never since the Second World War has a country invaded another independent state, a member of the United Nations, seeking to eliminate its very name and identity.

The Iraqi occupation lasted more than seven months, applying day by day a firmer and more aggressive grip on the people, property and natural environment of Kuwait. During the occupation large scale destruction of the information system infrastructure occurred. Computer machines were taken to Iraq. Fires were set in many computer sites.

The situation in Kuwait, before and after the invasion, provides a good opportunity to carry out an analysis of the organisational effects arising from the large scale destruction of the information system infrastructure. This study, therefore, looks at the disaster preparedness of Kuwaiti organisations before and after the invasion. It examines the effectiveness of recovery planning when subjected to events far exceeding the normal range of anticipated norm scenarios. In summary, the main objectives of the study are to:

- 1) reveal the Iraqi invasion's effects on organisations and the consequences of their not having disaster recovery plans;
- 2) clarify whether organisations learned from the invasion disaster and whether they now realise the importance of disaster recovery plans;
- 3) identify major problems facing IT managers regarding disaster recovery issues;
- 4) establish maximum allowable downtimes for different categories of organisations by size;
- 5) determine the most appropriate recovery strategies for organisations with different sizes and different degrees of dependency on computers; and

- 6) determine the optimum time-scale for providing each recovery strategy in hours, days or weeks.

4.3 Methodology

The method used to collect the necessary information is through distributing questionnaires and conducting follow-up interviews. The questionnaire is one of several data collection tools that can be used for research. It is perhaps the most popular of all such tools employed in statistical work (Wilson and McClean, 1994). As with any other tool, the questionnaire technique has some advantages and disadvantages. According to Wilson and McClean (1994), the advantages are that it

- Provides a useful method of obtaining information in a structured format;
- Can be administered without the direct support of an interviewer; and
- The responses to questions may be sought in a particular format to facilitate pre-determined analysis techniques.

The disadvantages are that it:

- Requires a lot of time to design and develop;
- Suffers from the “form filling” syndrome, especially if administered by post; and
- Has limited flexibility in terms of response format.

In selecting the country in which to undertake the study, the only available choices to the researcher were Kuwait and UK. Kuwait is thought to be more useful than the UK for the following reasons:

- Kuwait has just come out of the Iraqi invasion disaster, and therefore it is certain that all organisations in Kuwait will have taken stock of that experience; in the UK there is no such assurance.
- Follow-up interviews can be done more easily in Kuwait because of the small size of the country and close geographical proximity between the relevant

organisations, whereas in the UK a satisfactory sample of follow-up interviews would be virtually impossible to conduct in the available time.

- Questionnaires can be distributed and collected by hand in Kuwait, allowing face-to-face explanations in any cases of misunderstanding. This would increase the return rate of respondents. To apply this method in the UK would require a great deal of time, effort and expense.
- The researcher's native language is that spoken in Kuwait, thus facilitating better communication than would be achieved in the UK.
- The researcher's well-established links with IT managers in Kuwait can lead to a fuller co-operation with the study than could be expected from their counterparts in the UK.

4.4 Content and Distribution

A literature review shows that there has not been adequate research into the disaster recovery implications for Kuwait of the Second Gulf War. Because of the magnitude of the destruction and the large number of organisations which experienced the catastrophe in one form or another, there is a great opportunity to capture and analyse these experiences. The selection of relevant organisations was accomplished by reference to a variety of sources. One major source was a list provided by the Ministry of Commerce and Industry in Kuwait. The Business Phone Directory and the 1996 INFO Exhibition, which was held in April 1996 in Kuwait, served as secondary sources for locating additional organisations. Kuwait is a very small country and the number of organisations with an organised computer environment is not large. However, the researcher aimed for a target sample of 100 respondents. Accordingly, 140 questionnaires were distributed to different types and sizes of organisation. Fortunately, the responses exceeded the target number by 11. The residue of non-respondents are unlikely to have any effect on the final conclusion of this study because most of them are considered to have a low dependence on computers.

The questionnaires were distributed by hand to many organisations in Kuwait. They were passed to carefully selected respondents, either computer managers or disaster

recovery co-ordinators - if available. Follow-up interviews were conducted with most respondents to explain the objectives and to provide further clarification, if needed. Then the completed questionnaires were either handed to the researcher during the meetings or returned by post, using self-addressed and postage-paid envelopes which were provided by the researcher. The latter method was used in case the respondent did not desire his or her organisation to be recognised by the researcher. Anonymity was guaranteed for responding organisations.

The questionnaire contained 29 questions. It was divided into three parts to determine: 1) the general characteristics of the organisation; 2) the recovery strategies employed; and 3) the off-site storage backup strategies installed.

From the 140 questionnaires distributed in February 1996, a total of 111 usable responses were received after 12 weeks, representing a return rate of 79.3%. This good rate is at least partly due to the methodology of distributing the questionnaires and to the researcher's personal contacts with some of the respondents. Most respondents were keen to be helpful because, given their own experience of the Second Gulf War, they thought the subject was very important.

The organisations in the sample vary in size, degree of dependency on computers, the processing type they adopt, and the type of businesses they conduct. The study sought to cover a wide range and achieve a fair balance among different types of organisation. The four key characteristics among the obtained sample can be seen in Figures 4.1 through 4.4.

Figure 4.1- Proportions of Business Activities by Organisation Category

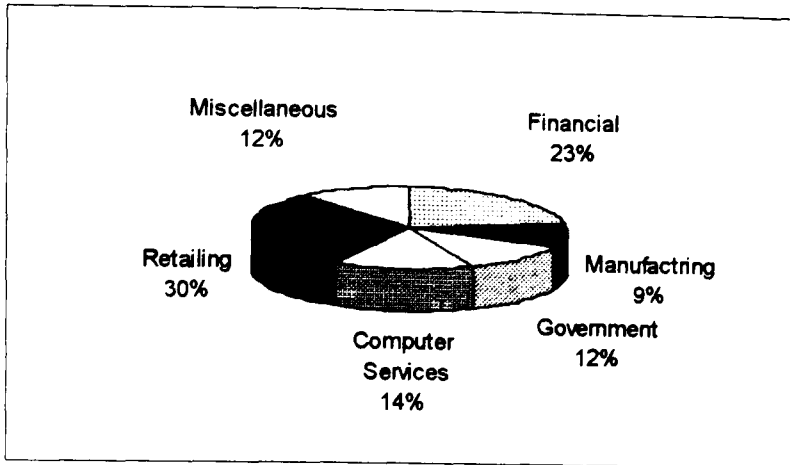


Figure 4.2 - Percentage Size of Organisations




Figure 4.3 - Degree of Dependency among Organisations






Figure 4.4 - Organisations by Processing Type



4.5 DRPs before the Invasion

Figure 4.5 shows that only 25% of organisations had Disaster Recovery Plans (DRPs) before the invasion. Nevertheless, only 68% of those which had DRPs had actually activated them either partially or fully. The rest did not operate their plans because roads were blocked and they could not reach their organisations due to the fast pace of the military occupation of the country.

Figure 4.5 - DRPs Availability before the Invasion



The high percentage, 66%, of organisations which did not have DRPs reflects the fact that many of them had not appreciated the importance of recovery plans. Most of them admitted that they lacked the knowledge of the importance of DRPs and never thought a disaster would strike them. Others understood this importance but did not make the time and effort to establish a DRP.

Even then, most of the organisations which had activated their plans had only saved some of their vital data, critical applications and documents. One of the respondents said that even though his organisation had a DRP, they had not tested it for more than a year. So by the time they had grasped the gravity of the situation, it was too late to enter the computer floor. However, they managed to locate one of the employees who had worked there for several days during the invasion who smuggled out some of the tapes, hidden under his clothing.

Another disaster recovery co-ordinator said that because of the lack of testing, it took him several hours to call the right people and meet at the site to collect the backup media. Luckily, the building was not yet occupied and the weekly routine backup was performed just the night before the day of the invasion. (*The invasion happened at the weekend*).

However, there are a few organisations such as the Public Authority for Civil Information and the Public Institution for Social Security, which saved all their data, applications and documents. In fact, the latter organisation recovered all its resources, except fixtures like large hardware, because their recovery plans were well tested. The testing procedures were done regularly in these organisations because the nature of the data is very important to the government and to the individual citizens whose personal details are held. Some of these organisations ran their applications and continued to carry out part of their activities in neighbouring countries or in the UK.

4.6 Organisational Losses

According to Dr. Adel Assem, the Director General for the Public Authority for Compensations Resulting from the Iraqi Invasion, which was established by the United Nations, the losses for private organisations in Kuwait totalled approximately \$35.6 billion while the total losses for Kuwaiti Government institutions was around \$69 billion. These figures do not cover personal losses for individuals, such as death, injuries, damage to homes, automobiles, furniture, etc. (Al-Watan, 1996). All the losses resulted either from stolen and damaged property or from disruption to business. All those who lost their data, information and applications, had to restart from scratch. Acquiring the hardware and software was not too much trouble, but redeveloping applications and capturing data again proved an unwelcome and demanding experience. Usually, employees are not motivated to do the same work again.


Losses resulting from the Iraqi invasion with respect to computer centres (Figure 4.6) vary from one sector to another. The computer centres in government organisations, for example, were completely damaged. Computers (large, mini and PCs) were disassembled under the supervision of specialists and carefully shipped to Baghdad. Connection wires, air-conditioning and communication equipment were dismantled and also transferred. The damage sustained by the government institutions often entailed complete loss, resulting from the transfer of the contents to Iraq or from the destruction of the buildings and the remainder of their contents in order to render them useless in future, as happened to the Kuwait Institute for Scientific Research (Centre for Research and Studies on Kuwait, 1994).

The education sector also suffered from the aggression. Schools and other educational institutions were used for accommodating the troops. Losses here extended beyond the computer centres. The laboratories, research equipment and the furniture of lecture halls were also dismantled and transferred to Baghdad. UNESCO Mission Report to the UN indicated that Kuwait University lost no less than 95% of its

computer facilities and databases (UNESCO, 1991; Centre for Research and Studies on Kuwait, 1994).

However, the private sector, especially banking organisations, did not suffer as badly as government organisations regarding equipment destruction. There was some damage to equipment and the loss of documents in some branches, but most of their main computer centres were not damaged. Nevertheless, they experienced major losses such as discontinuity of revenues and the failure of a significant proportion of their skilled personnel to return after the Liberation.

Figure 4.6 - Organisational Losses Due to the Invasion




The results show that many organisations suffered mostly from the loss of revenue, hardware, skilled personnel and software. The loss of revenue was most serious because of the long period during which the businesses were unable to function. The occupation lasted for about 7 months. However, clearing the remaining troops from the country, restoring electricity and communication lines, and allowing citizens who had fled the invasion to return to the country took a further three months.

The loss of skilled personnel came third in importance to organisations. This is not surprising because many Palestinians, Jordanians, and some other nationalities, in addition to Iraqis, were either deported or not allowed to enter the country after the Liberation because their countries had supported the Iraqi aggression. In addition, many skilled personnel from other nationalities did not return because they had found employment with other organisations world-wide during the occupation. Kuwait's population in 1990, before the invasion, was a little more than two million. In the 1995 census, it was just over 1.5 million. The loss of approximately one quarter of the population is devastating for a country as small as Kuwait.

4.7 DRPs Status Now


Did the organisations learn from the war disaster? Figure 4.7 shows that some Kuwaiti organisations recognised the importance of disaster recovery plans. This was clear when comparing the number of organisations which had DRPs before the invasion (25%) with those which now have approved DRPs (49%) or waiting for approval by top management (29%). But why still waiting for approval? One of the IT managers in a medium-sized financial institution replied: 'it is hard to convince top management, especially when our organisation came out from the invasion disaster without any equipment damages.'

Figure 4.7 - Availability of DRPs in 1996



In spite of the magnitude of the disaster, the results show that there are some organisations in Kuwait (22%) still neglecting to employ DRPs. The reasons are quite similar to the types of excuses used in every co. Lack of budget, as expected, was the most frequent excuse advanced because the whole country is operating under extreme financial stringency due to the huge losses resulting from the invasion. However, lack of knowledge of the importance of recovery plans came in third place - surprisingly, in a country just emerging from such a major disaster. This was perhaps because the sample included several organisations which were only established after liberation and several others are private organisations which were not affected much by the invasion. Figure 4.8 produces a breakdown of the excuses, with percentages, for not having an approved DRP.

Figure 4.8 - Excuses for not having Approved DRPs



4.7.1 Maximum Allowable Downtime

Maximum Allowable Downtime (MAD) is the period for which an organisation can be maintained without computer services, and when computer backups must be provided for them. Other names are Response Time or 'Drop Dead Time'. It is

important that organisations identify their various resources and define a MAD for each one. Then these resources are prioritised in terms of criticality to indicate which should be recovered first, second, third, etc. Depending on the maximum allowable downtimes, it is possible to choose an appropriate recovery strategy (Arnell, 1990).

The utilisation of computers is not new in Kuwait. It started slowly in the 1970s, but expanded fast in the late 1980s. Because of the invasion many organisations lost their information systems infrastructure, so they had to rebuild their computer environment again. This rebuilding gave them the chance to get rid of the old systems and adopt the most up-to-date technology in the world. This was especially true when many western organisations came to the country after the liberation to offer technology services and products. Therefore most organisations in Kuwait now depend heavily on computer services to conduct their operations and make decisions.

Figure 4.9 gives an indication of the extent to which organisations are computer dependent by illustrating the MADs for several organisational categories in Kuwait. It was found that financial organisations have the lowest MAD (1.82 days) and manufacturing organisations have the highest MAD (2.94 days). The average among all organisations, regardless of size, is 2.43 days.

However, when organisations are categorised by size, it is found that as the size of the organisation gets larger, the MAD gets lower. This can be seen in Figure 4.10 for the financial sector category.

Figure 4.9 - Maximum Allowable Downtime by Organisation Category





Figure 4.10 - Maximum Allowable Downtime by Size for Financial Institutes



Comparing these findings with the University of Minnesota survey, mentioned in Chapter 2 in Figure 2.2, the variation in the results is understandable. For example, it is reasonable that organisations now should have lower MADs than in 1978 because firms nowadays depend more heavily on computers in every aspect of their business.

In addition, back in 1978, only medium to large organisations could afford to have large computer processors. The use of personal computers started to spread widely in the 1980s. Since then many smaller organisations have come to rely on computerised systems to conduct their business. For example, most large financial organisations which used computer processors in 1978, had a MAD of 2 days. Since then, as the finance sector's dependency on computers has grown, this study shows that the MAD is now less than 24 hours for giant financial organisations.


4.7.2 Recovery Strategy

A recovery strategy defines the interim ability to process data while a full recovery of the primary computer site is underway (Arnell, 1990). That is, it is the selection of an alternative recovery site for running the business until the original site is ready once more. Since some organisations cannot afford to be without computer services for as long as one or two days, they must have an alternative site to run their critical operations if the original computer site is subject to a disaster.

There are not too many disaster recovery vendors in Kuwait because the country is small and the disaster recovery market is still in its infancy. Some alternate sites are in neighbouring Gulf countries (Saudi Arabia, Bahrain, UAE). Some organisations have branches, or are themselves branches of organisations, in these countries. When organisations are asked what type of alternative site they have adopted or will adopt, they responded as shown in Figure 4.11.

The results show that the hot site strategy is the preferred strategy for financial organisations. This is understandable in view of the critical nature of their businesses. The preference for hot sites by financial organisations is virtually true in every country. In fact, over 65 percent of all hot site providers in the USA involved financial organisations (Schreider, 1995). The preferences of other types of organisations are more variable but the hot site is the commonly preferred strategy for all categories.

Figure 4.11 - Preferred Recovery Strategies by Organisation Category




Categorising organisations by size, Figure 4.12, shows that hot site and duplicate site strategies are preferred by large and giant organisations. This is understandable because these organisations can afford the high cost of these types of strategy. However, small organisations prefer either mutual aid strategies or other less costly strategies such as mobile and portable sites, hardware vendor, or just manual procedures.

These results conform to the findings of a survey undertaken by CHI-COR Information Management, Inc in 1986 called 'Computers in Banking Survey' (Toigo, 1989). Both conformed that the size of the organisation is an important factor in selecting a recovery strategy. Both studies showed that larger organisations prefer the hot site strategy for the recovery of their businesses. Small to medium size businesses prefer the mutual assistance strategy, along with other strategies such as mobile and portable sites. The interesting difference, however, between the 1986 survey and the Kuwaiti survey is that in 1986 many large banks were adopting the mutual assistance strategy while in this survey it is clear that large organisations no longer favour this

recovery option. This is understandable for many reasons, such as the declining cost of hot sites, the testing problems accompanying mutual assistance agreements and the fact that many disaster recovery experts do not recommend this option for large organisations, unless other options are unavailable (Hyde, 1993).

Figure 4.12 - Recovery Strategies by Organisation Size




4.7.3 Testing

Even the best laid plans never work out during an actual disaster as one thinks they will. It is impossible to determine if a plan is really capable of recovering the business until it is tested. It is therefore important to have continual testing and evaluation of plans. Experts strongly recommend that DRPs must be tested at least every six months (Baylus, 1991).

The Kuwaiti study also showed that even if organisations recognise the importance of having disaster recovery plans, they still lack the commitment to test these plans more often. As shown in Figure 4.13, it was found that from the 49% of organisations which have DRPs, 47% did not test their disaster recovery plans because of budget and time constraints and 35% tested their DRPs only once a year. Only 18% of organisations tested their plans more than once yearly. This is a very low percentage,

bearing in mind that disaster recovery experts recommend that an effective plan should be tested at least twice a year. However a survey, undertaken on UK organisations in January 1993 by the University of Loughborough in association with the Computing Services Association and the National Computing Centre, showed that only 22% of UK organisations have a viable recovery plan, indicating that a complacent attitude towards testing is almost a universal phenomenon (Hearnden, 1993).

Figure 4.13 - Testing Among Available DRPs




4.8 Off-Site Backups Storage

Many DRP experts believe that effective off-site storage of critical resources is certainly one of the most important components of any effective and successful disaster recovery plan. Also, a plan may be in perfect condition, having been securely stored in a fireproof safe, but will still be useless if the fireproof safe is buried beneath the rubble. To guarantee that off-site backups and source documents are not consumed in the same disaster, thus rendering production systems unusable, these items should be stored at a safe location.

The study showed that 16% of the organisations do not backup their applications and data. This is because some of them did not depend heavily on computers but in other cases it was just bad management. It also showed that many organisations in Kuwait (68%) were storing their backups in the same computer building. Some of them say that they do this because they have a fireproof safe in the basement of the building. This is not a safe decision because in case of a disaster, an organisation cannot reach its backups because the area is usually sealed off and no one is allowed inside until the building is investigated and declared safe. This might take several days. In this situation an organisation cannot operate its alternative recovery site if the critical applications, data, and documents are not available. However, in addition to storing in the same building, some of these organisations were storing other copies of their backups in separate buildings (37%) and outside the country (14%).


Organisations which store their backups outside Kuwait started doing that just after the invasion. They learnt from the invasion's experience and still fear the threat of the near neighbour, especially after the October 1994 Iraqi army build-up on the Kuwaiti-Iraqi border. Organisations adopting this strategy are those which can afford the cost of so doing. They are usually government and multinational organisations. Many, the so called local organisations, store their backups inside the country and do not have any copies outside. They believe that if the country is occupied, the outside backups will not be of any practical use because they could not run their businesses from abroad.

The study also shows that many organisations (65%) are considering the importance of their critical applications and data by backing up their work daily. However, as Figure 4.14 demonstrates, 13% and 22% only do their backups every other day and weekly, respectively.

Figure 4.14 - Backups of Critical Systems

Having outdated or unreadable backups is as bad as not having backups at all. No matter how much detail is provided on the activities, unless tests are performed to determine their usability, the backups will often not work, contain rubbish, or just not be up-to-date.

The study shows, Figure 4.15, that only 27% of organisations are testing the readability of their backups three times or more a year. Disappointingly, 28% of organisations are not checking backup readability at all. This is a high percentage and these organisations could face an unpleasant situation, if and when a disaster strikes and the backup media is needed for recovery.

Figure 4.15 - Tests of Backups Per Year to Ensure Readability

4.9 Summary and Findings

The study obtained data for 111 organisations in Kuwait. It covered a wide spectrum of different sizes and types of organisation. It covered organisations which heavily depend on computers as well as those which do not. It highlights, for the first time, the scale of the destruction inflicted by the Iraqi invasion disaster with respect to computer centre losses. It shows the consequences and losses of computer centres that come from large scale human-made disasters such as wars. It also found that government organisations, rather than private organisations, tended to be the main targets for Iraqi destruction.

The study shows that the awareness of the importance of DRPs is rising in Kuwait, particularly after the distress that most organisations experienced from the invasion. This was clear when comparing the number of organisations which had DRPs before the invasion (28%) with those which now have approved DRPs (49%) or are in the process of finalising one (29%). However, even though organisations recognise the importance of disaster recovery plans, they still do not recognise the importance of testing these plans more often.

The study also shows that although some IT managers recognise the need for adopting DRPs, it appears that they do not have a methodology to follow for the recovery strategy selection process. IT managers have selected their alternative recovery strategies based on outside recommendations or simply on similar projects performed for comparable businesses. This approach is not efficient because requirements differ from one organisation to another. The methodology presented late in this research can, in the researcher's opinion, provide a solution to this problem.

Another finding is that most of the respondents are grappling with a common question: how much should the company be spending on the DRP? This appears to cause a good deal of concern and needs further investigation. Some IT managers stated that they would greatly appreciate efforts to provide a solution to the above question.

In addition to the specific findings mentioned above, the results of this study can contribute significantly to the development of a methodology for recovery strategy selection, introduced in the following chapter, and of a prototype expert system, introduced in Chapter 7, by constructing the following tables:

- Table 5.4 - Recovery strategies selection by size and degree of dependency.
- Table 5.10 - Recovery time among recovery strategies.
- Table 7.1 - Examples of time intervals for different organisation categories.

Finally, it is hoped that the study increases awareness among organisations, and not just computer centre managers in respect to disaster recovery plans. Also, its very recent findings should strengthen the claims made late that the methodology and prototype expert system are realistic and acceptable approaches worthy of serious consideration by organisations in this field.

Chapter 5

The Methodology for Selecting A Recovery Strategy

5.1 Introduction

Much of the work published in the disaster recovery area has dealt with the need for disaster recovery planning, how to develop and implement a recovery plan, effects of actual disasters, and the consequences of not having a plan. Moreover, most of the recovery plans which have already been developed deal with issues such as management support, choosing the recovery team, risk analysis, emergency procedures, testing and maintenance (Jackson, 1994; Robinson, 1993; Brown, 1993; Orr, 1988). A literature review also shows that these analyses only point to specific features of the various recovery strategies without recommending the most suitable one. The full implications of selecting the most suitable recovery strategy have not been addressed adequately in the literature.

The aim of this research therefore is to develop a methodology and a computerised system that fully addresses the issue of selecting an appropriate recovery strategy. The study hopefully will help IT managers and disaster recovery co-ordinators to estimate the optimal investment required and to recommend one or more recovery strategies for a particular organisation. The solution consists of two stages: 1) development of a structured methodology for the recovery strategy selection process; and 2) development and implementation of a computerised system to make use of the structured methodology from (1). The methodology is explained in this chapter whereas the computerised system is discussed in Chapter 7.

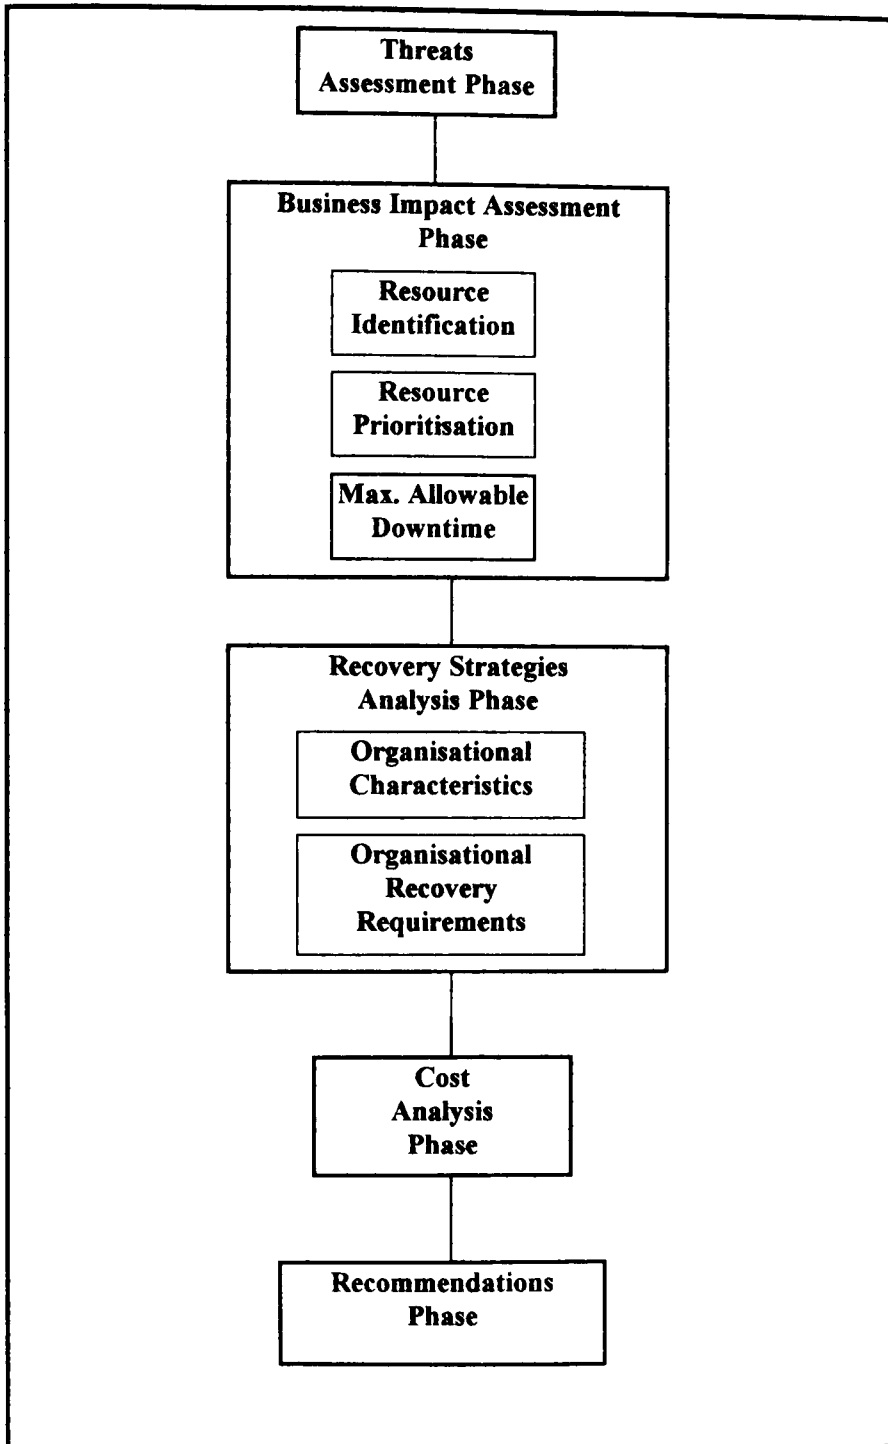
5.2 The Framework of the Methodology

The framework of the methodology is displayed in Figure 5.1, Recovery Strategy Selection Framework. It explains how the process of selecting the appropriate recovery strategy is undertaken. The methodology consists of five phases that provide a step-by-step approach to ensure that the entire recovery strategy selection process is covered. The phases are: Threats Assessment, Business Impact Assessment, Recovery Strategy Analysis, Cost Analysis, and Recommendations. The following paragraphs briefly describe these phases; then each phase is explained in more detail in the following sections.

The methodology will aid IT managers to identify potential disasters that are threatening their companies. A new approach for classifying threats is presented. The threats are classified in such a way that they can contribute more usefully in the recovery strategy selection process. This analysis is done in the first phase and is called Threats Assessment Phase.

Once the threats assessment has been finalised, the second phase, Business Impact Assessment, is presented. In this phase, computerised systems and applications are identified and then prioritised in terms of criticality to the organisation. More attention is given to those that are deemed critical in terms of their importance to the survival of the organisation after a disaster. Then, the overall maximum allowable downtime for which an organisation can tolerate the failure of its computer systems is calculated. The key organisational requirements are also identified in this phase.

The IT manager, the disaster recovery co-ordinator/team or whoever is in charge has the responsibility for analysing and selecting the most suitable and efficient recovery strategy. This strategy must fit the true recovery requirements of the organisation. A number of factors affect the choice of recovery strategy. Some of these factors are related to the organisation itself. Others relate to the characteristics of different recovery strategies. The third phase, Recovery Strategy Selection, explains these factors and shows how the recovery strategy selection process is undertaken.

Figure 5.1 - Recovery Strategy Selection Framework

Disaster recovery experts state that the cost analysis of accommodating a recovery strategy must only be used for budget purposes, not to make a decision on whether to adopt a recovery strategy or not. Disaster recovery plans should not be evaluated on the basis of cost-effectiveness (Robinson, 1993; Baylus, 1991). The present research, however, does not ignore the fact that management needs some indication of how much they need to spend on a disaster recovery. It provides IT directors with a method for calculating the investment required for a disaster recovery strategy. This is carried out in the fourth phase.

In the final phase, a computerised system is developed to provide recommendations based on several inputs from the user. The recommendations cover the following three aspects:

- 1) determining maximum allowable downtimes for organisations;
- 2) calculating the investment required to adopt a recovery strategy; and
- 3) providing recommendations in terms of selecting the most suitable recovery strategy.

5.3 Threats Assessment Phase

The first step of the methodology is to identify what threats exist to normal information processing activities. This is an extensive and difficult phase, fraught with uncertainty and the need to apply judgement (Orr, 1988). However, no reasonable recovery planning can be done without reaching agreement within the organisation as to what types of threat could realistically affect the operation, and what are the most probable disaster occurrences to expect. This step is fundamental in deciding upon the types of preventive measurement that should be installed or recovery strategies that should be selected (Arnell, 1990; Baylus, 1991; Toigo, 1989).

There are many threats that can have serious consequences on computer centre operations. There are various schemes in the literature for classifying threats. One

scheme divides threats by causal origin, either natural or man-made. Another division is by looking at the phenomena such as water, fire, power failure, mechanical breakdown, etc. (Toigo, 1989). According to Parker (1981) and Carroll (1984), threats may also be divided into intentional, accidental, and natural. A fourth division is to distinguish between threats according to their effects on computers: those affecting software and data are called logical, whilst threats affecting hardware are called physical (Elbra, 1992; Danish, 1994).

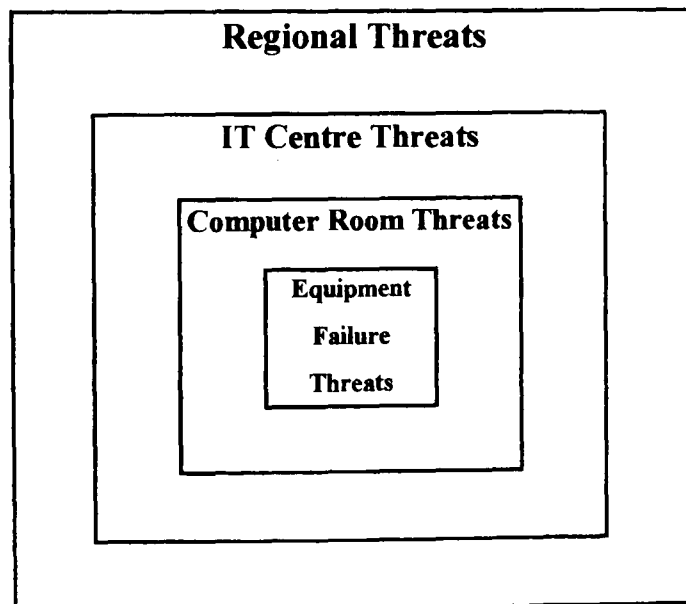
The above-mentioned classifications, however, do not contribute adequately to the decision-making process of selecting the most suitable recovery strategy. They do not assist in deciding what types of recovery services should be adopted in order to recover and save the business after a disaster. Rather, they help in installing safeguards to protect existing assets against probable threats (Arnell, 1990; Danish, 1994). These safeguards are installed either inside or outside the computer centre to reduce the risk of a threat occurring (Orr, 1988). For example, in the causal origin classification mentioned above, if a natural disaster such as a flood is anticipated, a countermeasure may be taken such as erecting barriers. Another example is that if terrorist activity is expected in a certain area, then a suitable countermeasure would be to deploy additional security officers. Classifying threats by phenomena can also contribute to installing proper safeguards. For example, a sensitive fire detector may be installed in every room in the computer centre to detect fires as early as possible.

However, if a company wants to invest not only in prevention but also in recovery, these classifications are of no assistance in deciding the type, recovery time and location of a recovery strategy required for business continuation. The bottom line is that the above-mentioned types of classification are helpful in making decisions about suitable preventive measurements but not in actually selecting the recovery strategies. Therefore, it is important that threats should be classified in such a way that the system of classification itself contributes significantly to the recovery strategy selection process.

In this research, a new approach named “Threats Magnitude Classification” (TMC). Figure 5.2, is presented to assist in the selection process. The severity and longevity of the disaster will determine, in addition to other factors, the proper recovery strategy, or combination of two strategies, to be adopted (Arnell, 1990). For example, for massive disasters, a short-term strategy could be selected that would provide the temporary use of a service bureau, while a long-term strategy, cold site, is outfitted with hardware.

Threats are categorised in the new approach according to the size of an area that they are expected to damage. For example, some organisations are more concerned with regional threats such as earthquakes, hurricanes, and major floods. Others are concerned with threats that affect only the computer centre premises, such as building fires and terrorist activities. Others may feel that they have preventive measurement for the more likely disasters but still feel that they must prepare for recovery because so many disasters are unpredictable and uncontrollable.

Figure 5.2 - Threats Magnitude Classification (TMC) Approach



In the Threats Magnitude Classification (TMC) approach, threats are classified into four types according to the size of damaged area that the threat may cover. Descriptions and examples of the four types are illustrated in Table 5.1. The objective of this new classification is to assist in determining the most suitable recovery strategy by looking at the following issues:

- Whether a short or long term recovery strategy is required. For example, equipment breakdowns that need to be repaired or replaced might only require short-term strategy, whereas building fires require both short and long-term recovery strategies.
- The physical capacity of the alternative site. For example, in threats damaging only the computer room, a large work area is not needed for personnel because their offices are not affected and they can be readily connected, if necessary, through communications lines to the alternative site. On the other hand, in regional threats, sufficient space in the alternative site is essential so that key staff can process and recover critical functions.
- The availability of additional personnel provided by the alternative site's vendor. In major disasters where a wide area is affected, employees are pre-occupied for days or even weeks with the safety of their homes and families rather than recovering the business of their employers. For organisations expecting these threats, additional assistance from external staff is necessary to run the business during that time.
- The location of the recovery site. If a regional threat is expected, there is no point in having a recovery site near the original site because both sites may be affected by the same disaster. In the case of computer room or equipment failure threats, a mobile site parked in an open area adjacent to the firm might be suitable.

Table 5.1 - Types and Examples of Threats in the TMC Approach

Threat Type	Description	Examples
Type I Regional	Threats that affect countries, regions, cities.	Hurricanes, major floods, wars, tornadoes, earthquakes.
Type II IT Centre building	Threats that affect smaller area such as buildings in a street, computer centre. Includes IS offices, computer room, equipment.	Building fires, major terrorist activity, falling aircraft, minor floods.
Type III Computer room	Threats that affect only the computer room or floor. Includes floor, equipment, communications, processors, air conditions, etc.	Minor fires, burst pipes, heat, humidity, smoke.
Type IV Equipment failure	Threats related to equipment failures. Hardware, software, interface, etc.	Mainframe failure, sabotage, software failure, hacker, theft, communication failure.

The distinction between types of threats may be easy and straightforward for some organisations, but it might be difficult and unclear for others. For instance, for organisations that are located in an area that is exposed to well-known threats the identification of threat type may be very clear. Examples of this include earthquakes in the West Coast of USA; hurricanes on the East Coast of USA; or IRA bombing in major cities in the UK. However, for other organisations the identification may not be easy. To ensure that the type of threat is correctly identified, the information collected should include, but not be limited to answers to the following questions:

- Is the organisation located in an area that is exposed to a natural disaster? *major river, earthquake fault.*
- Is the organisation in or near a building that has the potential for being attacked by a terrorist? *federal building, embassy.*
- Is the organisation near a takeoff or landing runway of an airport?

- Is the organisation located in or near a country that is considered to be politically and militarily unstable?
- In general, what threats are expected? *e.g. fire, flood, hurricane, intrusion, earthquake.*
- What is the area that would be affected? *e.g. country, city, street, building, floor, offices.*
- What resources are affected? *e.g. hardware, software, communication lines, WAN, LAN, data, documents.*
- Will all organisation sites be affected? *e.g. main office, branches, computer centre.*
- How much will the computer centre be affected? *e.g. fully, partially, or not affected.*
- Are there any preventative measurements installed to eliminate or reduce the expected threat(s), and, if so, which one from the four threat types they can prevent? *e.g. flood barriers, additional security around the building, water detectors, Halon flooding, mantrap.*

The Threats Assessment phase is extensive and not an easy task to perform. However, it is important that it should be done before proceeding to the following phase, Business Impact Assessment.

5.4 Business Impact Assessment Phase

Once the Threats Assessment phase has been finalised and the type of threat expected has been determined, the next step is to perform the Business Impact Assessment (BIA). The ultimate purpose of the Business Impact Assessment phase is to calculate the maximum allowable downtime for the organisation. Since all systems are not equally important, the BIA should thoroughly and objectively examine all of the organisation's resources, identifying and prioritising critical ones. The BIA is the foundation on which the overall recovery strategy selection rests (Wold, 1996). 'It is considered to be the cornerstone of the recovery plan' (Fisher, 1996). According to Robinson (1993) and Fisher (1996), the BIA serves several purposes:

- It identifies all resources in the organisation;
- It helps to distinguish the resources which are critical to the survival of the business;
- It determines the cost of downtime and the maximum allowable downtime for each resource;
- It recognises possible exposures and liabilities to internal and external entities, such as unions and regulatory agencies; and
- It determines whether there are intangible issues, such as public image and political embarrassment.

The Business Impact Assessment process can be done using one of the well-known data collection methods: conducting interviews, distributing questionnaires or a combination of the two. An interview may be conducted with each department. A form sheet containing several questions about each system that the department has will be reviewed with the department manager or with a senior member of staff who has been in the department long enough to be able to answer all the questions. Alternatively, a questionnaire may be sent to every department of a specific system or application to identify the extent of its usage in the performance of normal work (Toigo, 1989; Fisher, 1996).

The reasons for collecting the data and information are to:

1. identify each resource that needs to be recovered;
2. prioritise the identified resources in terms of criticality; and
3. determine the downtime cost for the denial of these critical resources.

If the organisation is too large with multiple locations domestically and internationally, the BIA process will certainly be a huge task. To solve this dilemma, several teams are assigned to do the job. Each team may take on particular functions or locations. This automatically narrows the process of interviewing critical departments and individuals (Fisher, 1996; Wold, 1996).

5.4.1 Resource Identification

The collection of data should include a comprehensive list of computer and telecommunications hardware a complete inventory of applications and systems software. All resources must be identified individually and in detail before any priorities are set and an assessment made for the recovery strategy selection. This identification helps the company to look for any resource that is critical and might be forgotten because it is not operated regularly. Such a resource is customised operating system software.

As mentioned earlier, the ultimate goal of the Business Impact Assessment phase is to determine the actual maximum allowable downtime for the critical resources. Therefore, questions should be constructed carefully to achieve this goal. Direct questions such as “Is your application critical?” should not be asked. All departments like to believe the work they do or the systems they use are critical. It is human nature to want to be needed. However, if the questions ask not how critical a systems is, but rather what steps a department would take to perform the same function if the system was unavailable, then subjective views about criticality becomes less problematic for the assessment. Many disaster recovery co-ordinators who have used these types of question have discovered that departments will provide a surprisingly fair assessment of their system’s criticality (Toigo, 1989).

There are several questionnaires and form sheets presented in related literature which can assist in capturing the required data. Some of the questions that can be asked should include, but not be limited to:

- What does the interviewee’s department do?
- What software and hardware do you need to run your department?
- What would happen if these software and hardware were not available to you?
- Can the department perform the job manually if the computer is not in service?

- Does the department use any critical specially-tailored equipment or software?
- To what degree can the department tolerate the interruption of the application?
- What is the financial loss for several time intervals? (For example: 6 hours, 12 hours, a day, two days, one week, etc.)
- What is the minimum staff and floor space needed to continue operations at another facility?
- What communication devices would be necessary to continue operations? (i.e. telephone, facsimile, switchboards)
- What are the revenue producing functions of your organisation?

The scope of the recovery should be limited to those systems which are deemed critical in terms of importance to the commercial survival of the business. One of the significant steps of the recovery process will be identifying and prioritising these critical data and applications.

5.4.2 Resource Prioritisation

After identifying all the resources in the organisation, the resources are then prioritised in terms of criticality and which should be recovered first, second, third, etc (Jackson, 1994). The continuation of a large percentage of the information systems operations at the alternative recovery site immediately after a disruption is rarely logistically, technically, or economically feasible. The resources that an organisation has are not all of equal importance. Attention should be focused on time critical resources requiring recovery as soon as possible while placing non-time critical resources at a lower priority for recovery (Toigo, 1989; Baylus, 1991; Jackson, 1994; Fisher, 1996; Wold, 1996).

Systems criticality can be measured in several ways. It may be measured by degree of tolerance. Tolerance is defined as the ability to cope with system interruption. Tolerance may be expressed in many ways. It may be commonly expressed as a monetary value: the loss of revenue to the company from system outages of specific duration. If there is a very low tolerance within the company to the loss of a system or

to the interruption of an application, this low tolerance is expressed as a high money value or cost. If, on the other hand, the company can tolerate to a significant extent the loss or interruption of an application, this high tolerance is expressed as a low money value or cost. Applications whose losses would entail substantial costs for the organisation are termed critical. Conversely, high tolerance applications are referred to as noncritical (Toigo, 1989).

Tolerance may also be based upon the length of time that the system or application is unavailable for use, or upon the time of the day or month an outage occurs (Toigo, 1989). For example, the general ledger is not considered critical until the end of an accounting period, specifically a quarter. However, the Business Impact Assessment process assumes that an outage will always occur at the worst possible time.

After gathering information about each system, then they are divided into three categories: critical, semicritical, and noncritical (see Figure 5.3). The definitions of the three categories are stated below:

- **Critical:** A disruption of service of these systems would seriously jeopardise the operation of the organisation (usually one day). Their tolerance to interruption is very low and the cost of interruption is very high.
- **Semicritical:** Systems and applications for which suspension can be tolerated for a short period of time (usually a week). They have higher tolerance and lower cost than critical systems.
- **Noncritical:** Systems that may be interrupted for an extended period of time, at little or no cost to the organisation.

Figure 5.3 - Criticality in Respect to Tolerance and Cost

	High	Noncritical		
Tolerance	Medium		Semicritical	
	Low			Critical
		Low	Medium	High

Cost of Interruption

Since the aim of this project is to assist IT managers, or others having the same responsibility, in selecting a recovery strategy to save the business from great loss or even the cessation of trading, it is important to concentrate only on systems and applications that are crucial for organisational survival. Management must carefully review the list of critical systems to ensure that only the really critical ones are included. This is because, based on these critical systems, the maximum allowable downtime can then be determined.

5.4.3 Maximum Allowable Downtime

The objective of this section is to determine how long an organisation can tolerate the interruption of its systems at a time of adverse incidents (maximum allowable downtime). The maximum allowable downtime contributes significantly in the choice of an appropriate recovery strategy (Arnell, 1990; Jackson, 1994). For example, organisations with lower maximum allowable downtimes, *i.e. one day*, should adopt the hot site recovery strategy option. On the other hand, organisations with higher maximum allowable downtimes, *i.e. one week*, should adopt the cold site recovery strategy option. In addition, the maximum allowable downtime is also needed as an

input for the model applied in this research for calculating the investment required on a recovery strategy, as will be explained in the Cost Analysis phase later in this chapter.

The work done in the previous two sections (collecting data and information for identifying and prioritising resources) is essential to arrive at the maximum allowable downtime estimate. Calculating the maximum allowable downtime is based on the cost of the computer downtime with respect to the organisational revenue. For selected time intervals (e.g. 6 hours, 12 hours, one day, two days, etc.) the consequences of the denial of computer systems for each critical resource is estimated. Each loss or potential exposure is quantified and the cost effects are aggregated for each time interval. Then the aggregate cost is compared to the income for the first selected time interval. The maximum allowable downtime (one of the selected time intervals) is reached when the cost of downtime exceeds the revenue. This process will be explained more fully by an example when the prototype system is described in Chapter 7.

Table 5.2, which shows the cost of downtime, is an example of how to collect the cost of downtime for each resource for a specific time interval. Assuming that the entire IT centre has been damaged (worst-case scenario), users are restricted from entering the damaged facilities, and the data and applications stored off-site can only recover the system to midnight of the previous day. The time interval in the table can be adjusted depending on the type and size of the company. For example, a giant-size financial organisation might begin with a small time interval e.g. 3 or 6 hours, whereas a small research centre might start with 2 or 3 days.

Table 5.2 - Cost of Downtime

Time after disaster	Downtime Cost of Resources					Total Cost	Income
	res1	res2	res3	res4	res5		
3 hours							
6 hours							
12 hours							
One day							
Two days							
One week							

At the end of this phase, Business Impact Assessment, the proposed system automatically computes the maximum allowable downtime for a particular organisation. However, there are some conservative top management who may want to change the maximum allowable downtime for one reason or another. Top management looks at the organisation from a different perspective. Intangible issues such as political embarrassment, public image and media criticism are more important to them and may reduce the overall maximum allowable downtime for an organisation. Therefore, the developed computerised system will need to be flexible to allow for this type of change.

5.5 Recovery Strategy Analysis Phase

The process of selecting a recovery strategy is carried out after the Threats Assessment and the Business Impact Assessment phases have been accomplished, the critical systems have been identified and the maximum allowable downtime for the organisation has been approved.

The aim of this research is to assist in selecting an appropriate recovery strategy for organisations. The Information Technology (IT) manager, the disaster recovery coordinator/team or whoever is in charge has the responsibility for analysing and selecting the most suitable and efficient recovery strategy. The recovery strategy must fit the true recovery requirements of the organisation and the maximum allowable

downtime necessary to recover from an adverse incident. Available recovery strategies, with their respective advantages, disadvantages and a comparison between them were listed and explained in Chapter 3. The approach for selecting the most appropriate recovery strategy is described in this phase.

There are a number of factors that influence the decision to select a particular recovery strategy. Some of them are related to the organisation itself. Others are related to the characteristics associated with the recovery strategy. The decision is based on the following:

- 1) organisational characteristics (size of organisation, degree of dependency on computer and maximum allowable downtime);
- 2) organisational requirements and recovery services required (outside personnel support, work area, special-tailored platform, security, usage duration); and
- 3) recovery strategy's characteristics (discussed in Chapter 3).

5.5.1 Organisational Characteristics

There are three major organisational characteristics that are highly significant in the selection process: 1) size of the organisation; 2) level of degree of dependency on computers; and 3) maximum allowable downtime. The latter was explained earlier in section 5.4.3. The first two are explained in the following two sections.

5.5.1.1 Size of the Organisation

The larger the organisation, the larger and more complex its computer centre and the more complex and advanced the strategy it needs. Recovery requirements for large organisations are not the same as those for smaller organisations. For instance, for giant organisations, the alternative site should be capable of providing space not only to perform computer operations, but also for computer centre personnel who need to be there to run the complex data processing (Arnell, 1990; Toigo, 1989). On the other hand, small organisations usually need to run few systems and only one or two of their

key personnel need to be at the alternative site. Moreover, Epich and Persson (1994) state that the selection of a recovery strategy depends largely on the size of the company (number of hardware, software, number of original sites in the company, staff).

Classifying organisations by size is not an obvious approach from any theoretical literature. Sizes conceivably vary across industries and even across countries. Organisations may be divided either according to their turnover, value of assets, or number of employees. However, because of inflation and the fast pace of corporate development, classifications according to turnover or value of assets may change every ten years or less. A company with a turnover of £1.5 million was considered to be medium size ten years ago, now it is classified as a small firm (Shafto, 1991). Classification according to the number of employees also varies from one industry to another. For example, in the manufacturing sector, a company with 200 employees is classified as small whereas in other sectors it is regarded as a medium size company (Storey, 1988; Acs & Audretsch, 1993; Harrison, 1994).

Nevertheless, in the real world organisations are normally divided into four sizes: giant, large, medium, and small (Shafto, 1991; Storey, 1988). Most organisations usually define themselves according to this classification or at least they know to which category they belong. The objective of this research is not to guide organisations on how to classify themselves. Therefore, the above classification is applied in the present research and it is assumed that organisations know into which category they would fall.

Size can also determine the type of recovery strategy (internal, co-operative, mutual agreement and commercial) the organisation should adopt. For example, giant organisations can afford the cost of handling their own duplicate site or sharing a co-operative site with other companies where the cost of operating the site is distributed between them. On the other hand, small firms adopt mutual agreement or manual procedures and few would choose commercial recovery strategies as an alternative.

Table 5.3 shows the types of recovery strategy in relation to the size of organisation. (The types were explained in Chapter 3, Table 3.1).

Table 5.3 - Types of Recovery Strategy in Relation to Size

Size of organisation	Type of recovery strategy
Giant	Internal, co-operative, commercial.
Large	Co-operative, commercial.
Medium	Commercial, mutual agreement.
Small	Commercial, mutual agreement, manual.

5.5.1.2 Degree of Dependency

The second organisational characteristic that influences the selection decision is the degree of dependency on computers. Organisations vary in their level of dependency on computers. Organisations that are highly dependent on computers in their daily operations, such as financial institutions cannot sustain the denial of computers for a long period of time. On the other hand, some work can be postponed, such as certain types of research and development tasks, until the computer system is up again (Baylus, 1991). An organisation can be categorised to be highly dependent on computer by looking at its maximum allowable downtime. Arnell (1990) stated that organisations with maximum allowable downtimes ranging from 1 to 5 days are considered to be *highly* dependent. Organisations with a *medium* level of dependency usually have maximum allowable downtimes ranging from 6 to 30 days. Organisations which have maximum allowable downtimes greater than 30 days are considered to be a *low* dependent (Arnell, 1990).

The levels of dependency in this research follow Arnell's classification. Organisations are divided into three levels: high, medium, and low. This division helps to determine which group of recovery strategies is appropriate for each level of dependency within different categories of size.

Table 5.4, showing the Group Selection Based on Size and Degree of Dependency, illustrates how groups of recovery strategy options are in respect to size and degree of dependency. For each size category and degree of dependency, a group of recovery strategies may be recommended. The recommendations made here are based on information gathered from the literature, related surveys and the recent case study of organisations in Kuwait (see Chapter 4). Organisations are classified by 4 sizes (*giant, large, medium, low*). Then within each size, organisations are further divided into three categories according to their degree of dependency on computers (*high, medium, low*).

Table 5.4 - Group Selection Based on Size and Degree of Dependency

Size	Degree of Dependency	Recovery Strategy Selection
Giant	High	Duplicate site, realtime recovery, commercial hot site.
	Medium	Warm site, co-operative cold site, commercial cold site, portable site.
	Low	Withdrawal of service, manual procedure.
Large	High	Realtime recovery, commercial hot site, service bureau, mobile hot site, co-operative hot site.
	Medium	Warm site, hardware vendor, co-operative cold site, commercial cold site, portable site.
	Low	Withdrawal of service, manual procedure.
Medium	High	Realtime recovery, commercial hot site, mobile hot site, time broker, service bureau, reciprocal agreement.
	Medium	Hardware vendor, portable site, withdrawal of service.
	Low	Withdrawal of service, manual procedure, null strategy.

Table 5.4 - Continued

Small	High	Mobile hot site, warm site, reciprocal agreement, time broker, hardware vendor.
	Medium	Withdrawal of service, manual procedure.
	Low	Manual procedure, null strategy.

However, in order to select the most suitable recovery strategy within a suggested group, other important information must be obtained. This information is related to the types and levels of recovery services required by organisations and the recovery strategy's characteristics.

5.5.2 Organisation's Recovery Requirements

Organisational requirements usually influence the recovery strategy selection process. The decision cannot be taken until all the organisation's requirements are fully identified. There are some requirements which can be instantly identified such as security, recovery time, and specially-tailored hardware or software employed by the organisation. There are others which cannot be fully appraised and understood until the Threats Assessment phase is completed and the threats type is identified. Such services are allowable length of time to utilise the alternative site, external personnel support, work area, and the location of the alternative site.

The Threats Magnitude Classification (TMC) approach which was developed in this research and explained earlier in the Threat Assessment phase contributes a great deal to decisions about which recovery services are required. Each threat type is associated with specific recovery services that are needed only when this particular type is anticipated. Table 5.5 illustrates the recovery services needed for each threat type.

Table 5.5 - Recovery Services Associated with Threat Types

Threat Type	Recovery Services
Regional	Personnel support, work area, replacement of HW&SW, short-term recovery, long-term recovery, remote location.
Building	Work area, replacement of HW&SW, short-term recovery, long-term recovery, within the city area.
Floor	Replacement of HW&SW, short-term recovery, adjacent location.
Equipment failure	Short-term recovery, adjacent location.

As explained in Chapter 3, Recovery strategies vary in the level of recovery services they provide. Identifying the required level of a particular recovery service assists in determining the exact recovery strategy that suits an organisation. The tables that are explained in the following paragraphs are a rearrangement of the tables presented in Chapter 3. The reason for this rearrangement is to facilitate the process of identifying the required level of a particular recovery service. The tables classify recovery strategies depending on the recovery service they provide. The proposed expert system, which is developed by the present research, employs a series of linguistic values (for example high, medium, low) to capture the exact level of recovery service required by organisations. Based on the user responses, the system recommends the most suitable recovery strategy for a particular organisation.

It may be worth mentioning here that it is not an objective of this research to help IT managers in deciding which vendor to select for recovery services. Vendors differ, within each recovery strategy, depending on many parameters such as: hardware and software compatibility; reputation; communications facilities; reliability and supplies. While a comparison between vendors within each recovery strategy is not a focus for this research, it can be recommended as a possible area for future study.

External Personnel Support

External personnel assistance is required by some organisations, especially those which are exposed to major disasters. Many organisations tend to believe that their own staff will always be available after a disaster. This may be true in minor disasters but it is not the case in regional disasters. In major disasters, such as earthquakes or wars, employees are more concerned with their own and their family's safety rather than their company's welfare (Baylus, 1991). Companies which are exposed to major disasters should carefully consider the availability of outside help from a recovery vendor. These services may include the provision of experienced telecommunications specialists, system programmers, customer support representatives, and recovery operations specialists. The extent of such assistance varies significantly among recovery strategies. There are strategies which provide full personnel support. Others provide partial support or do not provide any support at all. Table 5.6 shows the levels of personnel support that are provided by recovery strategies (Hyde, 1993).

Table 5.6 - Personnel Support Levels among Recovery Strategies

Personnel Support	Recovery Strategies
Full	Duplicate site, commercial hot site, warm site, mobile hot site, realtime recovery.
Partial	Service bureau, hardware vendor, commercial cold site, portable site.
Not available	Reciprocal agreement, time broker, co-operative hot site, co-operative cold site.

Work Area

It is important for some organisations that there is a sufficient workspace at the alternative site to accommodate staff from the affected site. This requirement depends on the size and the complexity of the computer centre. The larger and more complex

the computer centre, the more space is needed for working staff. Some organisations may require space not only to perform computer operations but for office and administrative functions (Arnell, 1991). However, there are other organisations which only need a few key personnel to be at the alternative site. Table 5.7 shows the levels of work area available among recovery strategies.

Table 5.7 - Work Area Levels among Recovery Strategies

Work Area	Recovery Strategies
Sufficient	Duplicate site, commercial hot site, warm site, commercial cold site, realtime recovery, co-operative hot site, co-operative cold site.
Limited	Service bureau, hardware vendor, portable site, reciprocal agreement, time broker, mobile hot site.

Security

Recovery sites have different levels of security, both logical and physical. There are some organisations that have to satisfy a very high security level. These include: sensitive military installations, air-transportation command and control centres, air traffic control centres, government electronic mail centres, etc. For these organisations, it may be economically feasible to set up an entire alternative site in a geographical location far enough away not to be subject to the same disaster (Arnell, 1990). Some organisations perform activities requiring a lower security level. If the security issue is a very important requirement, a company should not consider the mutual or co-operative strategies. It is virtually impossible to protect the integrity of data in a mixed processing environment (Arnell, 1990; Hyde, 1983). Table 5.8 shows security levels among the recovery strategy options.

Table 5.8 - Security Levels among Recovery Strategies

Security	Recovery Strategies
Very high	Duplicate site, co-operative hot site, co-operative cold site.
High	Realtime recovery, commercial hot site, warm site, commercial cold site, mobile hot site, portable site.
Medium	Service bureau, hardware vendor, reciprocal agreement, time broker.

Special-Tailored Platforms

Often, IT managers involved in designing and implementing an open system type data centre and in-house developed software require an alternative site that can help them manage their special-tailored platforms. This is not a problem if a company subscribes to an open shell site, such as a cold or portable site, where equipment can be shipped after the disaster episode has settled down. However, the only sure means of ensuring an equipped site that handles this type of request is to have a duplicate centre.

Currently, there are no commercial vendors that are capable of handling full multivendor recoveries. There are some commercial strategies which can partially handle this request but, according to Robinson (1993), these are not without some disadvantages:

- They may increase the price of a commercial site considerably.
- Vendors may not have the knowledgeable personnel skilled in the various systems.
- They may lack the capability to handle various types of telecommunications needs.

The remaining strategies, co-operative, mutual agreements and some commercial ones, have neither the capacity nor the ability to meet this type of request. Table 5.9 shows the levels of the ability to change the platforms among recovery strategies.

Table 5.9 - Ability to Modify Platform among Recovery Strategies

Modify hardware/software	Recovery Strategies
Good	Duplicate site, co-operative hot site, co-operative cold site, commercial cold site, portable site.
Limited	Realtime recovery, commercial hot site, warm, mobile.
None	Service bureau, hardware vendor, reciprocal agreement, time broker.

Recovery Time

The ability to provide fast recovery services varies among recovery strategies. The decision as to which of the strategies to select depends on the organisation's maximum allowable downtime, as calculated in the Business Impact Assessment phase. Table 5.10, showing the recovery time for various recovery strategies, shows that the duplicate site option, which is internally managed and controlled, can be available immediately. The other option is realtime recovery which is a new concept and just introduced recently. Although the latter option is very expensive, its cost is expected to decrease as the technology progresses (Hyde, 1993). A number of other strategies can offer recovery within one to two days. Any other options that require the moving of equipment will naturally take several days.

Table 5.10 - Recovery Time for Various Recovery Strategies

Recovery Time (MAD)	Recovery Strategies
Immediate	Duplicate site, realtime recovery.
One to two days	Commercial hot site, co-operative hot site, time broker, service bureau, reciprocal agreement, mobile hot site.
3 days or more	Warm site, hardware vendor, co-operative cold site, commercial cold site, portable site.

Location

To avoid a situation where the same disaster strikes both the company and the alternative site (Threat Type I), many companies look for alternative sites at a distant location. Some companies have learned the need for this separation the hard way. One Mexico City company maintained its hot site in the same city prior to the earthquake there, and both the company's and the hot site's computer systems were destroyed (Hyde, 1993). Similarly, many disaster recovery vendors have established sites in and around London, UK. In the wake of the IRA bombing campaign, organisations need to ensure that their alternative sites are not in areas which are also vulnerable either from the effects of bomb damage or in accessibility by recovery personnel (Hyde, 1993). If a threat of Type III or IV is anticipated, then the alternative site may be within the same street. Therefore, determining the type of threat is a very significant factor in deciding upon the location of the recovery site. Table 5.11 suggests, according to the threat type, where to locate the alternative site.

Table 5.11 – Location and Restoration Period

Threat Type	Location	Restoration period after the disaster	Type of time-strategy required
Regional	Remote	Long time (months, years)	Combination of strategies
Building	Within the city	Long time (months)	Combination of strategies
Floor	Adjacent	Medium (weeks)	Short-term strategy
Equipment failure	Adjacent	Short (days)	Short-term strategy

Usage Duration

Recognising the type of threat to which a company may be exposed helps in deciding how long the alternative site is needed and if a combination of both short and long-term strategies is necessary. Regional threats usually have a long aftermath (months or

years) and may therefore require a combination of strategies to be developed to provide the temporary use of service bureau or hot site (usually up to six weeks), while a shell site is outfitted with software and hardware equipment until the company rebuilds its own original site (Arnell, 1990). Table 5.11 also shows appropriate time-strategy required for each type of threat.

The allowable duration of utilising an alternative site varies between different recovery strategies. There are some recovery strategies which can be used for up to six or seven weeks and they are called short-term strategies. Others can be used as long as needed and they are called long-term strategies. Table 5.12 shows examples of different strategies in each of these categories.

Table 5.12 - Usage Duration between different Recovery Strategies

Duration	Recovery Strategies
Short-term	Realtime recovery, commercial hot site, co-operative hot site, service bureau, reciprocal agreement, mobile hot site, time broker, hardware vendor.
Long-term	Duplicate site, warm site, commercial cold site, portable site, co-operative cold site.

It is obvious from the previous criteria and tables that selecting the most appropriate recovery strategy is not a straightforward decision, but rather an extremely complex exercise. Several factors and inputs have to be weighed and carefully considered. Therefore, an expert system which can handle this methodology is recommended to interact with the user to simplify the recovery strategy selection process. The system and the mechanism, which is used to decide between different recovery strategies, are introduced in Chapter 7.

5.6 Cost Analysis Phase

According to Arnell (1990), Baylus (1991) and most disaster recovery experts, the cost justification for adopting a recovery strategy can be a long and misleading

process. It thus should be avoided. IT managers can skip this process and convince top management that a disaster recovery strategy is required for statutory reasons. The pressure of government regulations has made disaster recovery arrangements virtually a mandatory requirement. This indeed, enables the cost justification step to be skipped (Baylus, 1991). Various Acts of Parliament have been introduced in the UK which have ensured that advances in computer technology do not jeopardise the requirement for confidential data to be adequately protected (Kerby, 1990). In the USA the Foreign Corrupt Practices Act made disaster recovery a requirement. The scope of this Act extends to UK and European subsidiaries of US companies who operate outside the USA as well as UK and European companies operating in the USA. In many countries, the banking industry is currently the most rigorously controlled sector in terms of disaster recovery arrangements. An organisation's failure to comply with government regulations could expose it to negligence claims against the company, its directors, and its officers (Hyde, 1993; Baylus, 1991; Kerby, 1990; Arnell, 1990).

Disaster recovery experts state that the cost analysis, if it can not be avoided, of fitting a recovery strategy must only be used for budget purposes, not to reach a decision on whether to adopt a recovery strategy or not. Therefore, disaster recovery plans should not be evaluated simply on the basis of cost-effectiveness (Robinson, 1993; Baylus, 1991; Arnell, 1990).

This research, however, will not ignore the fact that management needs to have some indication of how much they need to spend on a disaster recovery strategy. This research will try to provide an answer to a very common question usually posed by top management: Are we spending too much or too little on a recovery strategy? The literature shows that many IT directors usually justify the investment for adopting a recovery strategy by showing top management how much the company will lose on an hourly, or daily, basis if the computer systems go down (Mercorella, 1995; Baylus, 1991). This is an excellent approach if one wants to convince the executives of the need for a recovery strategy. However, it does not provide an estimate of how much money is required to accommodate one. For example, one day's loss of their systems

for large financial companies may cost several million pounds, thereby justifying even the most expensive recovery strategy.

To present IT management with an indication of how much they need to spend on a disaster recovery strategy, this research uses the second part of the model developed by Subhani. The Subhani model and reasons for preferring it over the ALE approach were explained in Chapter 2.

The Subhani model introduces a function called 'Contingency Cost-Response Time' which will be used in this research. The contingency cost means the cost of a recovery strategy, and the response time means the Maximum Allowable Downtime. For consistency, we will stay with our naming convention for recovery strategy and maximum allowable downtime. The function states that the cost of a recovery strategy is inversely proportional to the MAD. According to Subhani (1989), the relationship between the cost of a recovery strategy and the MAD can be expressed by:

$$R = R^0 e^{-nt}$$

Where:

R = Cost of a recovery strategy with a maximum allowable downtime of t days;

R^0 = Cost of recovery strategy with instantaneous or zero MAD. In practice, such a strategy would be a duplicate site;

n = Parameter measuring the intensity with which the recovery strategy cost declines with the maximum allowable downtime;

t = Maximum allowable downtime in days; and

e = Exponential constant.

Some additional data are needed from the disaster recovery market for the above-mentioned function to calculate the investment required for a recovery strategy. The company will be asked by the proposed system to provide two maximum allowable downtime estimates, and two annual cost estimates for two types of disaster recovery strategy. The cost estimates are quoted from the recovery strategy industry. Since the

commercial hot site is the most expensive strategy, it is selected to be the upper bound. Similarly, the cold site strategy is selected to be the lower bound because it is the least expensive. At the final stage, the maximum allowable downtime (which is an output of the Business Impact Assessment phase) is substituted in the model to calculate the required investment for adopting a recovery strategy. An example is presented in Chapter 7 to illustrate the mechanism of the model and how the final result is reached.

5.7 Recommendations Phase

The development of a computerised system to assist IT managers in selecting the most appropriate recovery strategy is the end-product of this research. The computerised system performs some computations and rule-based decisions to provide some recommendations regarding the continuation of business activities after a disaster. The system, which is presented and explained in more detail in Chapter 7, will cover the following aspects:

- computation of Maximum Allowable Downtimes (MAD);
- computation of the required investment for fitting a recovery strategy; and
- recommending a disaster recovery strategy.

The above computations and recommendations are based on the following criteria:

- 1) the type of threat that an organisation is exposed to;
- 2) organisational characteristics (size, degree of dependency, type of business, revenue, cost of computer downtime);
- 3) organisational requirements, where some of these requirements are deduced from the expected threat type (security, work area, platform modification, service support, usage duration, location); and
- 4) data and information collected from the disaster recovery market.

To develop the required system that fits the requirements of the developed methodology, a full investigation of technologies that can support this effort is required. The investigation should include a comparison between available technologies and tools that may play a major role in developing the required solution. The technology and tools which are selected must meet the methodology requirements introduced in this chapter. The process of selecting the most suitable technology and tool is described in the following chapter.

5.8 Concluding Remarks

As mentioned earlier in this chapter, the literature on disaster recovery and related areas shows that the complete process (investment, recovery time and alternative recovery site) of selecting the suitable recovery strategy has not been addressed adequately. Therefore, a structured methodology to address solutions to the selection recovery strategy problem was developed in this chapter. The developed methodology also provides a basis for the development and implementation of a structured prototype computerised system. The prototype system and the methodology serve the IT managers in the following ways:

- Identify anticipated threats that might jeopardise the organisation.
- Classify threats in a way that can contribute more effectively to the recovery strategy selection process by introducing the Threats Magnitude Classification approach.
- Identify critical systems and applications that are crucial for a company to recover and resume its business.
- Calculate the maximum allowable downtime.
- Calculate the investment needed to adopt a recovery strategy.
- Finally, recommend a recovery strategy, or combination of strategies, for organisations.

Chapter 6

Expert Systems Technology

6.1 Introduction

The end-product of this research is to develop a computerised system that will assist IT managers in selecting an appropriate recovery strategy. The proposed system should do the following: (1) determine the Maximum Allowable Downtime; (2) compute the amount of investment required; and (3) recommend a recovery strategy. Having developed the methodology (described in the previous chapter), the next step is to search for an implementation tool to deliver the required system. Therefore, a full investigation of the available technologies which can support this enterprise is necessary. Such an investigation should include a comparison between available technologies and then between the tools that are capable of playing a major role in developing the required solution. Both the technology and the tool must fit the methodology requirements introduced in the previous chapter. To achieve these objectives, the following issues need to be examined:

- The functionality needed for the technology
- Feasible technologies
- Selecting the appropriate technology
- Expert Systems structure and development life cycle
- Languages and tools analysis
- Selecting the suitable tool

6.2 The Requirements

After identifying the problem and developing a methodology to provide a meaningful solution, the next step is to select the most suitable technology. The technology has to satisfy some functionality requirements. These requirements are explained in the following paragraphs.

Store Knowledge

A large amount of data and information related to the disaster recovery area needs to be stored in the system. This data and information are called facts. These facts relate to threat types, characteristics of recovery strategies, and services accompanying the recovery strategies. For example, there are several facts which are associated with each threat type, as indicated in Table 5.5. An important requirement of the technology to be selected is its capability of storing a large number of facts.

Flexible for Modification

The disaster recovery field is relatively new in comparison with other fields. It has only come into prominence in the late 1980s and the early 1990s. In fact, some big computer companies started to invest in IT disaster recovery to provide recovery solutions only a few years ago. For example, IBM entered the disaster recovery market in 1991 (IBM report, 1993). Others are thinking of providing Internet recovery services within the year 1997 (Shreider, 1996). Therefore, a requirement of the selected technology is the flexibility for modification to accommodate additions or changes to the knowledge base.

Reasoning Rule Decision

As was seen in the previous chapter, the problem-solving approach employed to select a recovery strategy was presented in the form of condition-action pairs: IF this *condition* occurs THEN an *action* is recommended. For example, one rule says: for a

company, that requires *very high* security requirement such as sensitive military installations, a *duplicate site* strategy is recommended. This problem-solving method is called Rule-Based Decision. Therefore, the technology to be selected should utilise this type of approach.

User-friendliness

Since the system is going to be used mostly by IT managers who have other responsibilities and little or no experience in the disaster recovery area, it should, if it is to be efficient, be developed to be user-friendliness, incorporating explanation facilities. The technology to be selected, therefore, should run under Windows and provide an easy-to-use pull-down menu environment. It also should use the point-and-click technique and a multiple choice method in order to reduce the effort of entering text. The technology should have the facility to provide explanations to steps or questions asked by the system. These and other similar facilities can save the user time and effort, thereby making the system understandable, effective and efficient in use.

Prototype

The design, development and coding rules and facts of a full system require a great deal of time from a project team. Its membership must include a disaster recovery expert who is knowledgeable about all the relevant rules, regulations, guidelines and methods of solving problems. It must also include a designer and a programmer to work closely with the expert to design and code the system. It is impossible for one person to produce a full system within the time scale assigned to this dissertation. A small-scale system (prototype) is, however, practicable in order to demonstrate how the final system will work. The technology selected to deliver the solution therefore should be capable of producing a prototype system during the development phase.

6.3 Feasible Technologies

After specifying the functional requirements, an extensive investigation was carried out to find a suitable technology to implement a computerised system for the disaster recovery selection process. Since the arrival of the computer, solutions to problems are usually implemented using conventional systems technology. Individual programs are developed to perform rapid calculations, access data, or perform modelling of complex processes. However, in the last decade or so, a new technology has been introduced to the field of computing science, called Artificial Intelligence (AI). Durkin (1994) defines artificial intelligence as *a field of study in computer science that pursues the goal of making a computer reason in a manner similar to humans*. Then, a new development of special purpose computer programs, a subset of AI called Expert Systems (ES), was introduced. These are programs or systems that employ human knowledge captured in a computer to solve problems that ordinarily require human expertise (Turban, 1992).

The above-mentioned two types of technology can be used for implementing the proposed system. To select the most appropriate one, they both need to be described and compared. The following paragraphs briefly describe the two candidate technologies: conventional systems and expert systems. Table 6.1 also summaries the main similarities and differences between the two possibilities.

Conventional Systems

In conventional systems technology, the computer is told how to solve the problem. It is given data and a step-by-step program that specifies how the data should be used to reach an answer. The conventional systems are based on an algorithm, which is a clearly defined sequential procedure, that produces a unique solution. They address problems where the information is complete and exact, such as database management systems or accounting programs. If data is faulty or missing, a conventional system cannot provide any results. The output that conventional systems produces must be correct or it has no meaning (Waterman, 1985).

During the development of a conventional system, the programmer receives the tasks from the designer and works largely alone, interacting with others only when difficulties arise or new directions are needed. The specifications defined during the design phase are assumed to be fixed and no changes are expected. If changes do need to be made, the task is sent back to the designer for the required modification. The user interaction is only with the analyst and perhaps the designer most of the time, there is no interaction between the programmer and the user (Durkin, 1994).

Conventional systems follow a three-step development process of design, code, and debug. The system is not deliverable until the programmer has completed all three phases.

Table 6.1 - Comparison between Expert Systems and Conventional Systems

Dimensions	Expert Systems	Conventional
Processing	Mainly symbolic	Primarily computing
Nature of input	Can be incomplete	Must be complete
Search	Heuristic (mostly)	Algorithms
Explanation	Provided	Usually not provided
Major interest	Knowledge	Data and information
Nature of output	Can be incomplete	Must be correct
Solution	Unique solution	May produce several solutions
Maintenance and update	Relatively easy	Usually difficult
Reasoning capability	Yes	No

In summary, the conventional programmers' sphere of interest is limited to a set of data. Their focus is on the problem's data from which they try to find ways to process it to reach a unique solution (Durkin, 1994).

Expert Systems

Expert systems applications are developed in many fields to assist or replace an expert to solve a particular problem. Therefore, they are required to capture a great amount of knowledge about the area related to that particular problem. In expert systems, the computer is given knowledge about the subject area plus some inferencing (reasoning) capability. The expert system program determines the specific procedure for arriving at a solution. Expert systems are based on symbolic representation and manipulation. A symbol is a letter, word, or number that is used to represent objects, processes, and their relationships. By using symbols, it is possible to create a knowledge base that states facts, concepts, and the relationships among them. This knowledge is captured by an expert in the subject area. Then various processes are used to manipulate the knowledge to generate advice or a recommendation for solving problems (Jackson, 1990; Turban, 1992; Durkin, 1994).

Expert systems address types of problem that are less structured than conventional systems. The information available may not be sufficient to arrive at an exact solution. However, an expert system may still arrive at some inexact reasonable solution.

During the development of an expert system, the designer works closely with the expert throughout the project, endeavouring together to uncover the key points of knowledge. A little amount of knowledge is added to the system and tested to evaluate the solution. Therefore, a small prototype system can be built and presented to the expert at any stage to validate the problem-solving approach. Expert systems technology is explain more detail later in this chapter.

No one can say that one particular technology is better than the other. The choice of the right technology depends upon the requirements of the problem and how it needs to be solved. However, based on the problem requirements, which have been addressed by this research, expert systems is thought to be the most suitable

technology. The reasons for selecting this technology are discussed in the next section.

6.4 The Reasons for Selecting Expert System

The previous section explained the two types of technology that may be used to provide a solution to the problem introduced by this research. It also highlighted the major distinctions between them. Looking at the characteristics of both types and the functional requirements of the problem, it was found that expert systems technology is more suitable for implementation for the following reasons.

A Rule Based Technology is Needed

Expert systems is a technology that commonly represents knowledge in the form of condition-action rules. Many sets of rules are stored in the knowledge base that describes how to solve a problem. The knowledge base can store as many rules as are required for a particular subject area. As was seen in Chapter 5, the developed methodology has many condition-action pairs. Therefore, expert systems technology is thought to be capable of producing a solution by drawing very extensively on the rule-based method.

Revision is Required in the Future

Since the field of disaster recovery strategy is a relatively new one, some new strategies have only just been introduced, such as realtime recovery; and others may be introduced in the near future (see Chapter 3). In fact, some Internet Service Providers are investigating the use of a new recovery strategy called Internet Recovery Strategy. They found that this type of recovery is, indeed, feasible and very soon the Internet will provide recovery activities (Shreider, 1995). This means that the technology to be selected must be flexible and have the ability to add more recovery strategies or modify existing ones in the future. Expert systems technology has the ability to ensure this type of update.

Fast Decision Making is Required

Expert systems technology deals with jobs that involve the processing of a large amount of complex rules and facts. This can be done much more quickly than would be possible by a human expert. Indeed, an expert system can recommend a solution within a few minutes. It is therefore very suitable for the recovery strategy selection process.

A Technology with an Explanation Feature

The tangible product of this research is an expert system developed so as to assist IT managers in selecting an appropriate recovery strategy. IT managers have other responsibilities and cannot be expected to be experts in the field of disaster recovery strategy. Therefore, the system should act as an expert to provide explanations at each step or for each question it asks. At the end, the system will provide recommendations based on some facts about the organisation's requirements and recovery strategy characteristics and explain these facts.

A Technology with a Prototype Feature

Building a complete system takes a long time and requires more time than that available for this research. Therefore, due to the time constraint, it is essential to adopt a technology that has the ability to provide a prototype system. A significant advantage of expert systems technology is that it supports the development of a small-scale system for presentation, when it is needed. This feature is important in this research because only a prototype system will be delivered at the end of this project.

6.5 Expert System Structure

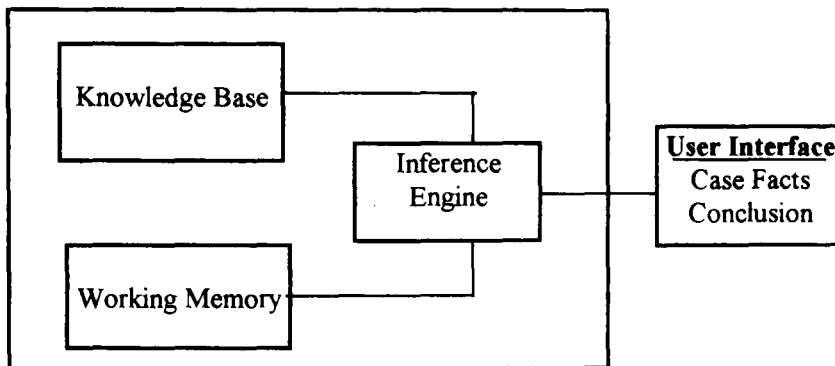
The 1990s were declared the decade of the brain by the US Government. Along with other scientific issues such as biological and biochemical ones, artificial intelligence and its derivatives are to be a primary focus for research (Turban, 1992; Durkin,

1994). The *expert system*, a derivative from the artificial intelligence field, is another name for the term *knowledge-based expert system*. Efraim Turban, a specialist in expert systems technology from the California State University at Long Beach and one of the leading researchers in the technology, defines an expert system as follow:

'It is a system that employs human knowledge captured in a computer to solve problems that ordinarily require human expertise. Well-designed systems imitate the reasoning processes experts use to solve specific problems. Such systems can be used by non-experts to improve their problem-solving capabilities' (Turban, 1992).

A typical expert system is composed of four basic elements. These are shown in Figure 6.1. They are: knowledge base; inference engine; working memory, and user interface. The following paragraphs explain briefly the function of each element.

Figure 6.1 - Expert System Structure



6.5.1 Knowledge Base

The knowledge base is the part that holds the domain knowledge necessary for understanding and solving problems. It holds guidelines, regulations and rules that link a solution to a specific problem in a specific area. The knowledge and human skills within a narrow area are obtained from an expert and organised and coded in the

knowledge base, using one or more knowledge representation techniques. These techniques are different from the procedures used in a conventional program (Turban, 1992; Durkin, 1994).

A number of effective techniques for representing knowledge in a knowledge base have been developed over the years. The most common techniques used in the development of an expert system are: logic; rule base; frames; and semantic networks.

6.5.2 Working Memory

The working memory is another part of an expert system that contains the information and facts about a problem that is either supplied by the user or inferred by the system. The system matches the facts entered in the working memory with knowledge contained in the knowledge base to infer new facts. The new facts are entered into the working memory and the matching process continues until a conclusion is reached, which is also entered into the working memory for future use, if necessary.

Expert systems can also utilise information contained in external storage such as databases, spreadsheets and sensors. The system may load this information into the working memory at the beginning of the session or access it when it is needed during the consultation phase (Turban, 1992; Durkin, 1994).

6.5.3 Inference Engine

The inference engine is the control of the expert system. This element is a computer program that provides a methodology for reasoning about knowledge in the knowledge base and in the working memory in order to formulate a conclusion. It works with the facts contained in the working memory and rules and knowledge in the knowledge base to derive new information. It searches the rules for a match between their conditions and information in the working memory and applies the rule with the highest priority.

For a rule-based system, two general methods of inferencing are used: forward chaining and backward chaining. In forward chaining the inference engine analyses the problem by looking for the facts that match the IF condition of its IF/THEN rules. When the IF part matches the information in the knowledge base, the rule *fires* and the THEN part of the rule is added to the knowledge base. The process continues until no matches exist between the IF part and the facts in the knowledge base. This inference mechanism is also referred to as data-driven or data-directed.

The other type, backward chaining, also called goal-driven or goal-directed, is used to prove a particular goal or hypothesis for specified data. The process starts with an initial goal and searches backwards through rules in the knowledge base from their THEN parts to their IF parts. The process ends when the inference mechanism reaches a conclusion which may or may not exist in the knowledge base.

6.5.4 User Interface

User interface is another component of expert systems that manages the interaction and communication between the system and the user. This interaction is conducted and carried out in a natural language style. A basic design requirement of the interaction is to ask questions. To obtain reliable information from the user, the designer needs to pay special attention to the question's design. The interface may be supplemented by menus, graphics and a special tailored screen.

6.6 The Characteristics of Expert Systems

Experts systems have many characteristics that distinguish their technology from conventional systems technology and human experts. In addition to the expert systems features which were explained in preceding paragraphs, the following are some of the major characteristics.

Knowledge and Control are Separates

In expert systems, the knowledge base and inference engine (control) are separate modules. This separation eases the tasks of modifying and maintaining the system. The user can easily locate and change some particular piece of knowledge, or add new knowledge at any location within the knowledge base without reviewing the control. This is a feature that differs from conventional systems where knowledge and control are intermixed. In cases of a change in knowledge being required, the code (control) has to be reviewed and understood first (Waterman, 1986).

Possesses Expert Knowledge

Another characteristic of expert systems is that they capture and code the expertise of a human expert. This includes both knowledge and problem-solving skills. The facts, concepts and rules about a subject area are gathered from experts, regulations and literature and stored in the knowledge base (Durkin, 1994).

Depth Expertise

Expert systems, like human experts, are designed to solve problems within their narrow area of expertise. They have limited ability to solve problems beyond the subject area. By concentrating on one area the expert systems technology can achieve depth and capture a great amount of knowledge to be processed for solving problems.

Permits Inexact Reasoning

Expert systems technology supports applications that involve uncertain facts, rules or both. This occurs when the user cannot provide a definite answer or offers only incomplete information when prompted for a response. Expert systems technology treats incomplete answers with techniques that deal with uncertainty such as the certainty factors, the Bayesian method and the Dempster-Shafer theory (Turban, 1995).

Flexibility

Expert systems have an efficient and modular storage capability for handling rules. They provide users with an efficient mechanism to add, change, and delete knowledge. This is a very important characteristic because expert systems contain very large amounts of knowledge.

Provides Explanations

An expert system can give explicit and detailed reasons for the questions it asks (why) and reasons that lead to its recommendation (how). This increases the confidence that the right decision has been reached. It also makes each step of the system understandable. This is difficult to achieve if reliance is placed on human experts. A human expert may be too tired, unwilling for one reason or another, or simply unable to provide an explanation in every case.

6.7 Expert System Development

The development of an expert system is similar to the development of any other system. It passes through the system development life cycle phases: project initialisation (which includes problem identification, assessment, alternatives, and managerial support), systems analysis and design, implementation, testing, and maintenance. However, the nature of the specific system determines which phases or tasks are to be performed, in which order, and to what depth. For example, a large scale expert system is developed according to a complex life cycle process, whereas a small scale system for end-users includes only a few tasks.

According to Harmon and King (1985) and Turban (1992), most expert systems specialists observe the following steps when developing an expert system application:

- Identify the problem.

- Identify the knowledge required to be included in the system. The process is called knowledge acquisition.
- Organise the knowledge by specifying what type of knowledge representation technique is suitable to solve a particular problem.
- Select a tool and implicitly commit yourself to a particular consultation paradigm.
- Implement by developing a prototype of the system using a tool. This includes creating a knowledge base and testing it by running a number of tests
- Expand, test, and revise the system until it does what the user wants it to do.
- Maintain, train and update as needed.

The following sections describe briefly the major steps for developing an expert system.

6.7.1 Knowledge Acquisition

As mentioned earlier, expert systems are also known as knowledge based systems, and it is clear that they are only useful in so far as they contain knowledge. This knowledge needs to be extracted and obtained from several sources and transferred to the knowledge base, and sometimes to the inference engine. This process is called knowledge acquisition. It is usually done throughout the entire development process, and even afterwards when a new knowledge is recognised.

Acquiring knowledge from the expert is a complex task for complex applications and it is well-known as the bottleneck for expert systems construction. For such complex applications, it is usual to have a knowledge engineer to assist the expert in making his knowledge explicit. The knowledge engineer interacts with the expert and helps him structure the problem area by interpreting and integrating human answers to questions (Turban, 1992; Durkin, 1994). However, in many smaller applications, experts can learn to use expert systems building tools and can engineer their own knowledge, as is the case in this project. Most experts are motivated to do so as the system may ultimately help them in their work. More importantly, knowledge will

change, new knowledge will need to be added to the knowledge base, and old rules may be amended as ideas change. All this points to the good sense of the expert continuing to be responsible for the updating of the expert system (Jackson, 1992).

Human experts are not the only source of knowledge. Other potential sources include textbooks, articles in journals, databases and special research reports.

6.7.2 Knowledge Representation

After acquiring knowledge from one or more of the above sources, the next step is to encode and organise it into the knowledge base. This is done by applying one of several knowledge representation techniques. The major techniques for representing knowledge, as mentioned before, are rule base, semantic networks, frames and logic.

The prototype system developed by this research uses the rule base technique. A rule is an IF THEN structure that logically relates the condition contained in the IF part to an action contained in the THEN part. The rules are matched to the facts about a problem contained in the working memory by the inference engine. This technique is probably the closest to the way a human expert would solve a problem (Waterman, 1986; Turban, 1992; Durkin, 1994).

6.7.3 Implementation and Prototyping

After selecting the type of knowledge representation, the next step is to implement the expert system. Expert systems technology encourages the use of an incremental prototype approach in system implementation. Turban (1992) states that prototyping is crucial to the development of many expert systems. Therefore, most expert system projects begin the implementation effort by building a small prototype system to determine the structure of the knowledge base before devoting the substantial amount of time necessary to build more rules. Turban (1992), Waterman (1986), and Durkin (1994) recommend that, for example, in a rule-based system the prototype may include only fifty rules. This small number of rules is sufficient to produce

consultations of a limited nature. The prototype approach has the following advantages:

1. It allows project developers to determine whether it is feasible to proceed with the full application using expert systems technology.
2. It provides a means through which to examine the effectiveness of the knowledge representation and the development tool as a whole.
3. It gives an idea of what the final application will do and what it will look like to the users.
4. It gives an opportunity to impress management or system founders and gain their commitment and increase their support for the project.
5. It allows the possibility of an early correction to the project direction, based on the feedback from management or potential users.

6.8 Selecting the Tool

Once the technology has been chosen, the next step is to select the proper programming language or tool to implement an expert system application. Although an expert system application can be built in any programming languages such as COBOL or FORTRAN, expert system developers prefer to use specific artificial intelligence programming languages and comprehensive integrated development packages. This preference is due to the fact that AI languages and tools are designed for symbolic processing i.e. for programming logical problems which involve knowledge.

6.8.1 Languages and Tools for Building Expert Systems

The following paragraphs illustrate briefly some AI languages and expert systems tools that are used for building expert system applications.

6.8.1.1 AI Languages

There are two major languages that are used for programming and debugging procedures in expert systems. The two languages are LISP and PROLOG. Their major features are briefly described below.

LISP

LISP stands for List Processing language. It was developed by John McCarthy in 1958 at MIT. Although it is an old programming language, many expert systems developers are still using it, particularly in the US. The basic data structure of LISP is the list; for example, an object can be presented as a list of words. LISP is oriented toward symbolic computation; the programmer can assign codes to terms like 'disaster' and 'flood'. Although such terms have no direct meaning in LISP, the LISP program can conveniently manipulate such symbols and relationships (Harmon and King, 1985; Turban, 1992; Jackson, 1992).

PROLOG

PROLOG (an acronym for PROgramming in LOGic) was initially developed by Colmerauer and Roussel at the University of Marseilles in 1975. The first efficient PROLOG compiler was developed at the University of Edinburgh. It is now the most popular expert system language in Europe and Japan (Harmon and King, 1985).

PROLOG is structured in terms of objects and relationships between objects (predicates). Knowledge is expressed in the form of facts about the objects and rules, showing how new facts are inferred from other facts. Usually goal statements declare what the PROLOG program has to prove. In order to prove a goal statement, PROLOG applies a pattern matching search to the database. The search is guided by forward or backward chaining theorem-proving.

6.8.1.2 SHELLS

Just as any word processor is a tool for producing a document, an expert system shell is a tool used for developing an expert system. A typical shell consists of some form of knowledge representation technique and a ready-made inference mechanism. Jackson (1992) classifies building-shells according to the knowledge representation techniques they use. The main categories include:

- Inductive shells, which uses a number of established facts to draw some general conclusion.
- Rule-based shells, where rules are entered in the form of IF THEN.
- Hybrid shells, which combine rules and induction.

6.8.1.3 Toolkits or Environments

Toolkits are integrated expert system packages that are developed to support several different ways of knowledge representation and handling inferences. Unlike the shells, which contains only one knowledge representation technique, toolkits may use several techniques such as frames, rules, semantic networks, object-oriented programming, and different types of chaining (forward, backward, bi-directional).

Toolkits permit a programming environment that allows complex specific systems to be built. They are more specialised than languages. Therefore, they can increase the productivity of expert system builders. Although toolkits require more programming skills than shells, they are more flexible. Because they were expensive in the past, they were mostly used for research rather than applications. However, as familiarity with expert systems technology has grown and the prices of hardware and software have fallen, they have become ideal vehicles for building expert system applications (Jackson, 1992).

From the above descriptions of languages and tools, it becomes apparent that expert system toolkits are potentially the most suitable tool for building an expert system

application because of the functionality they offer, which matches the requirements of this research.

6.8.2 Potential Tools and The Selection Determinants

Although toolkits are clearly more appropriate for this project, the decision as to precisely which one should be selected as the most suitable one is not easy. After extensive investigation, the following two toolkits were available to the researcher: CLIPS and *flex*. Although they are both expert system toolkits, they have different features and characteristics. The following paragraphs briefly explain each toolkit, and elaborate upon some of the issues involved in selecting an expert system toolkit in general.

CLIPS

CLIPS is an expert system tool which stands for C Language Integrated Production System. It was designed at NASA Johnson Space Centre, using the C programming language. CLIPS provides support for rule-based, object-oriented and procedural programming. The procedural programming capabilities provided by CLIPS are similar to the capabilities found in languages such as C, Pascal and Ada. CLIPS is syntactically very similar to LISP, which has been used mostly in USA. The inferencing and knowledge representation capabilities provided by CLIPS's rule-based programming language are similar to those in other expert system tools. However, CLIPS only supports forward-chaining. Backward-chaining is not supported by CLIPS (Giarratano & Riley, 1994).

CLIPS has been installed in a wide variety of computers, ranging from PCs to supercomputers. It is available on Windows 3.1, Macintosh and MS-DOS environments. The tool is available through the Computer Software Management and Information Centre (COSMIC) in Georgia, USA, which is the distribution point for NASA software. A major advantage of CLIPS is that it can be downloaded free of charge from several sites on the Internet. However, technical support is not provided.

To give an idea of how rules are written in CLIPS, the following example is presented. (The same rule is presented later using *flex*)

The pseudocode of an example rule is: If the emergency is a fire then the response is to activate the sprinkler system. Converting this pseudocode to a rule in CLIPS gives:

```
(deftemplate emergency (slot type))
```

```
(deftemplate response (slot action))
```

```
(defrule fire-emergency
```

```
  (emergency (type fire))
```

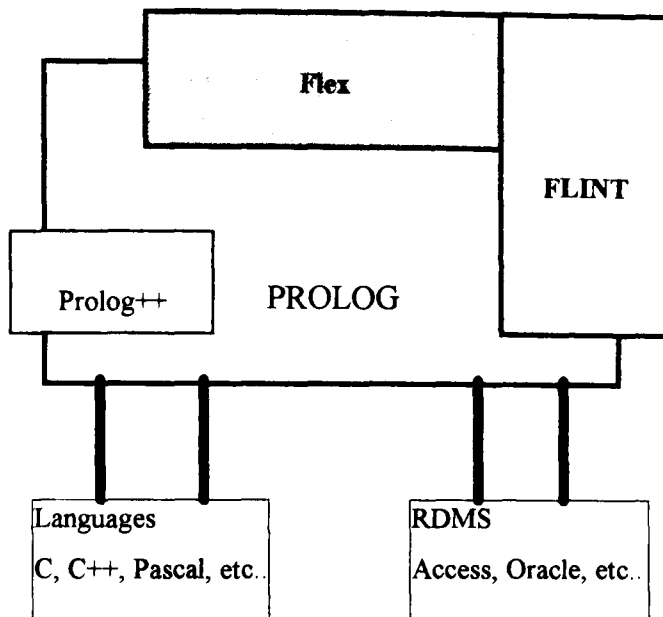
```
=>
```

```
  (assert (response (action activate sprinkler system))))
```

Flex

Flex (Forward Logical Expert system) is an expert systems toolkit which was developed by a UK company called Logic Programming Associates in 1988. It is a powerful toolkit which supports frame-based reasoning with inheritance, rule-based and data-driven procedures fully integrated with a logic programming environment. An important feature of *flex* is that it contains its own dedicated English-like Knowledge Specification Language (KSL). The KSL enables developers to write simple, concise and English-like statements about the expert's world and produce virtually self-documented knowledge-bases which can be easily understood and maintained by non-programmers.

Flex has direct access to Prolog and has support for procedures written in C, C++ and Pascal languages. The *flex* toolkit can be used on its own or in conjunction with FLINT, which provides support for fuzzy logic inferencing. Prolog also provides the means to access compliant databases and facilities for communicating with other programming languages. Figure 6.2 presents the *flex* environment and its interfaces.

Figure 6.2 - The *Flex* Environment

The *flex* toolkit is available as a portable solution across a wide range of different hardware and operating platforms. It is available on Windows 3.1, Macintosh and MS-DOS machines, and has been licensed to other Prolog providers on UNIX.

To give an idea of how rules are written in *flex*, the previous example using CLIPS is presented again, this time using the *flex* language.

```
rule fire_emergency
if emergency is fire
then activate the sprinkler system.
```

Comparing the two languages, it is clear that the language used in *flex* is easy to construct and more understandable than CLIPS by computer and non-computer users. However, a disadvantage of the *flex* toolkit is that, unlike CLIPS, it does not come free of charge.

In proceeding to identify which one of these two toolkits would be more appropriate, it was necessary to look at some check lists for selection. The following representative issues in software selection for expert system development have been identified by Turban (1992):

- Can the tool be easily obtained and installed? (*cost, compatibility*)
- How well is the tool supported by the vendor? (*stability, reputation, technical staff, availability, accessibility*)
- How difficult will it be to expand?
- What kind of knowledge representation schemes does the tool provide? (*rules, frames, etc.*)
- How well do the knowledge representation schemes match the intended application?
- Do the inference mechanisms provided match the problem? (*forward-chaining and backward-chaining*)
- What is the track record of success of the package?
- Is the tool capable of interfacing with other software and languages?
- Can the language of the tool be easily understood and maintained after implementation by non-programmer users?

6.8.3 Reasons for selecting *Flex*

After a full examination of: 1) the features of the two possible toolkits, *flex* and CLIPS; and 2) the representative issues in software selection for expert system development identified by Turban, it was clear to the researcher that the *flex* toolkit is more suitable than CLIPS for the following reasons:

1. Since there is always a possible need for modification and the addition for new knowledge, the language used in *flex*, English-like Knowledge Specification Language (KSL), is clear and understandable for revision by non-programmers.

2. CLIPS provides one inference mechanism only: forward-chaining. In contrast, *flex* provides both forward- and backward- chaining.
3. The *flex*'s vendor (LPA) is a British company which can be easily contacted by the researcher. Its main office is in London. In contrast, most CLIPS providers are to be found in the USA.
4. Technical support for the *flex* toolkit is available, when it is needed, but CLIPS support is not always available.
5. *flex* has an efficient and effective user interface (explanation facility, graphical display, on-line help).
6. Training sessions in *flex* are easily obtained from LPA and many other providers.
7. The *flex* toolkit was highly recommended by Dr. Tawfig Danish, a previous Ph.D student at the University, who developed a similar system using *flex*.
8. The proposed system, Expert System for Disaster Recovery Strategy Selection, can be used as a continuation to the Knowledge-Based Decision Support System for Computer Disaster Prevention, delivered by Dr. Danish a few years ago using the *flex* toolkit. This is beneficial in terms of issues such as compatibility, training, technical support and maintenance.

Additional *flex* toolkit features are presented and explained in more detail when the prototype system is described in the next chapter.

6.9 Concluding Remarks

After a full examination and analysis of the functionality requirements to support the methodology explained in Chapter 5 and the technologies, languages and tools that would be feasible to implement the proposed system, it became clear which technology and tool would be most suitable to deliver the solution. It was apparent that conventional languages and procedures had limitations which made them unsuitable for implementing the required system. For example, inference mechanisms and knowledge representation schemes require to be programmed and developed in conventional languages, whereas these mechanisms and schemes are already available

in expert systems technology. Therefore, expert systems technology was thought to be the most appropriate for meeting this project's requirements

After comparing and analysing languages, shells, and toolkits, the *flex* toolkit was found to be more suitable than CLIPS for the proposed system. The utilisation of *flex* and some of its additional features are explained when describing the prototype Expert System for Disaster Recovery Strategy Selection in the next chapter.

Chapter 7

A Prototype Expert System for Disaster Recovery Strategy Selection

7.1 Introduction

The previous chapter explained expert systems technology, its features and how to develop an expert system. It was also clear that expert systems technology would meet the functionality requirements of the methodology developed in Chapter 5. Moreover, the *flex* toolkit was found to be the most suitable tool for developing the proposed prototype system.

This chapter describes the proposed prototype system which is called: Expert System for Disaster Recovery Strategy Selection (ESDRSS). It consists of an overview of the proposed prototype system, its transactions, how the expert systems technology is applied, which knowledge representation technique is used and why. Then, the methods employed to acquire knowledge are stated and explained. Finally, the transactions and the mechanism of how they would work are explained in more detail. The contents of this chapter can be used as documentation for the proposed system.

7.2 Overview of the Prototype ESDRSS

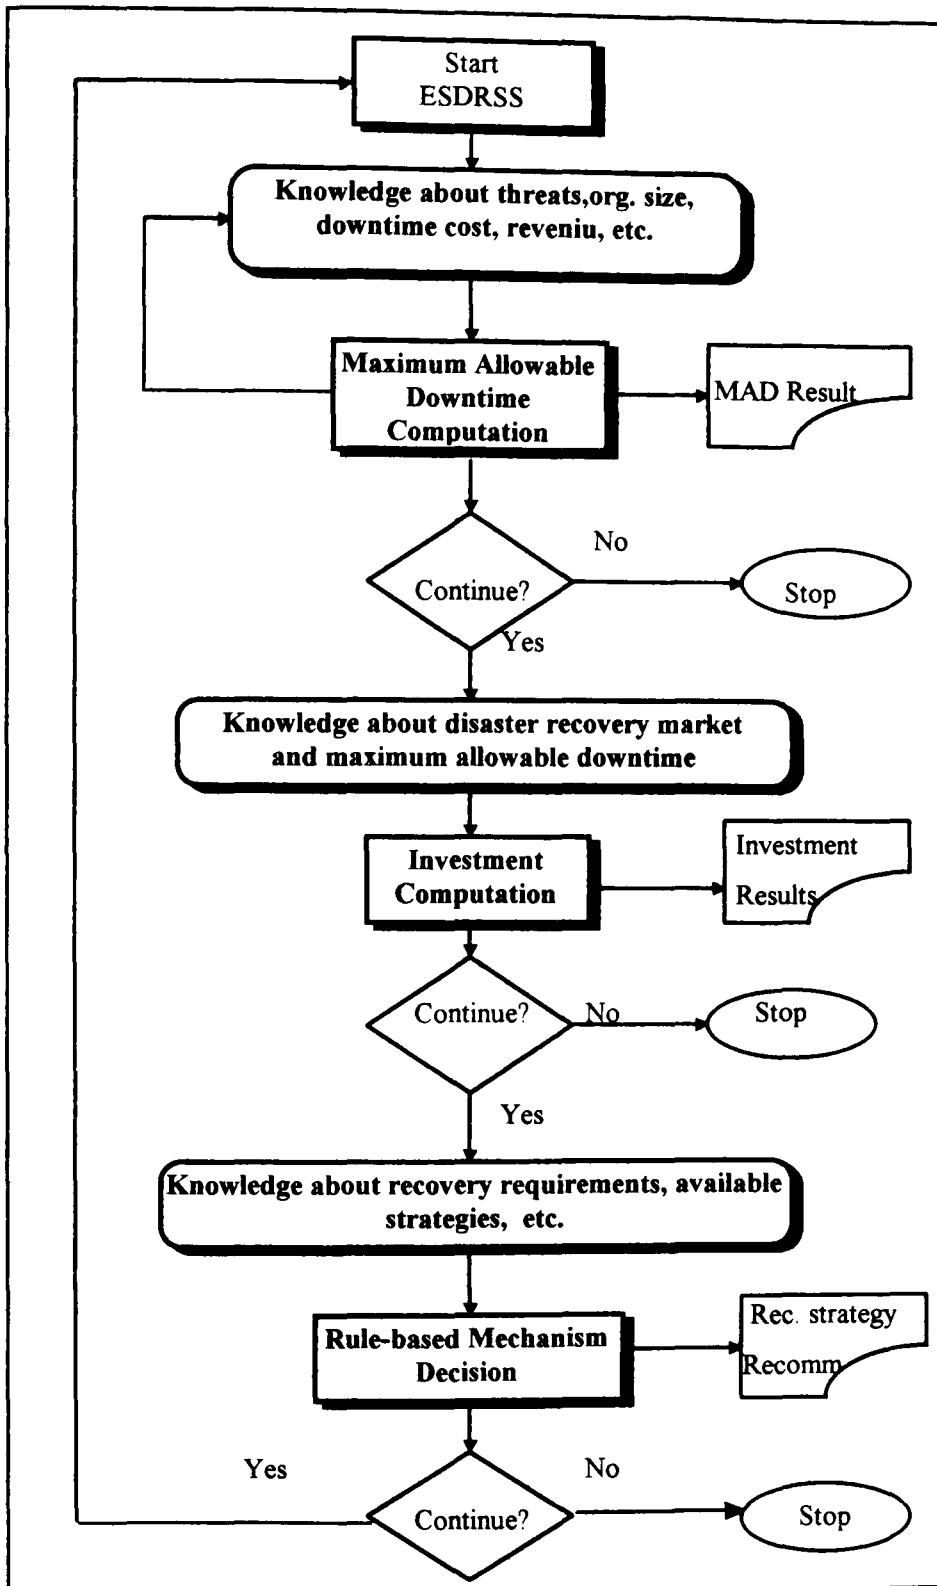
The prototype ESDRSS is developed to assist IT managers, disaster recovery coordinators, disaster recovery consultants and others to perform some computations and reach some rule-based decisions regarding the continuation of business activities after a disaster. The development of the prototype ESDRSS was accomplished via two stages: 1) the development of a structured methodology for selecting a recovery

strategy (Chapter 5); and 2) the development and implementation of the prototype ESDRSS, based on the methodology from (1) above, using expert systems technology. This system, after expansion and several tests, can be used to be part of the master disaster recovery plan for any organisation. The prototype ESDRSS (see Figure 7.1) consists of the following three major components (or transactions):

- computation of Maximum Allowable Downtimes (MAD);
- computation of the required investment for fitting a recovery strategy; and
- recommending a disaster recovery strategy based on the organisation's requirements and the recovery strategy's characteristics, using a rule-based knowledge representation mechanism.

The maximum allowable downtime (the output of the first transaction) is very significant because it is used as an input to the second transaction. It also assists in choosing the correct recovery time category in responding to one of the questions in the third transaction. The recommendation of a recovery strategy in the third transaction is, however, independent of the investment calculated in the second transaction. This is because the cost of a particular recovery strategy varies as between the many vendors, depending on the size and reputation of the provider, calibre of the consultant and the extent of the services provided. For example, the price of a hot site strategy ranges from \$10,000 to \$120,000 a month (Datashield Report, 1993; IBM Report, 1993a; Schreider, 1995).

Figure 7.1 - The Decision Structure in the Prototype ESDRSS



When designing the proposed system, it is important that it should be as explicit and understandable as possible without impairing the quality and efficiency of the contents and the results. This is an important goal because:

1. the system is delivered to IT directors have other responsibilities and have little or no experience in the disaster recovery area;
2. it is also delivered to disaster recovery co-ordinators who are considered to have little or no programming experience;
3. more information may need to be added in the future to maintain and enhance the knowledge base, as the field of disaster recovery expands; and
4. in order to be user-friendly, the system employs the point-and-click and multiple choice techniques to reduce the effort of entering text.

To satisfy the above demands, the knowledge in the prototype ESDRSS is presented as production rules in the form of condition-action pairs. The rule representation technique is especially applicable when there is a need to recommend a course of action, as is the case for the present research's objectives, based on observable events. According to Turban (1992) and Durkin (1994), the rule technique has the following major advantages:

- rules are easy to conceive, they are understandable because they are a natural form of knowledge;
- inference and explanations are easily derived;
- future modifications and maintenance are relatively easy for non-programmers;
- uncertainty is easily combined with rules;
- each rule is usually independent of others; and
- it is possible with rules to get a prototype system running quickly for budget approval or because of time constraints.

7.3 Knowledge Acquisition

Acquiring knowledge for any expert system application is not an easy task. Indeed, it is known as the bottleneck for expert system construction and continues to present many difficulties in developing an expert system. It takes time and requires several sources to feed a knowledge base. Sources for knowledge can be experts, books, reports, articles, regulations, guidelines, etc. Expert knowledge is, however, considered to be the primary source for most expert system projects.

Expert system developers have recommended that a prototype system should be created first, with relatively little knowledge inserted into the knowledge base, to test the feasibility of the project. Then, future additions to enhance the knowledge base can be incorporated at a later stage. Accordingly, due to the limited time scale for this project and to the proposal to develop only a prototype system, the researcher has limited the knowledge sources for the prototype ESDRSS to the following sources.

Literature

Some of the knowledge was extracted from specialised literature sources such as books, articles and reports in the disaster recovery area. Guidelines, case studies, actual experiences and research reports also fall into this category. The literature which was used is listed in full at the end of this dissertation. However, the following sources are thought by the researcher to have been the most useful:

- Disaster Recovery Handbook by Chantico Publishing Company, 1991;
- Handbook of Effective Disaster Recovery Planning by Alvin Arnell, 1990;
- IT Infrastructure Library Contingency Planning Module by IT Infrastructure Management Services, CCTA, 1989;
- Writing Disaster Recovery Plans for Communications Networks and LANs by Leo Wrobel, 1993;
- Contingency Planning by Information Systems Guide, 1989; and

- The Disaster Recovery Journal (DRJ) which specialises in the disaster recovery field. The DRJ is published quarterly and contains many articles, post-disaster outcomes, surveys and research reports.

Seminars and Workshops

There are many seminars and workshops world-wide in the field of disaster recovery, business continuity, and contingency planning, such as the International Disaster Recovery Symposium and Exhibition which has taken place annually in the US since 1989. The researcher had the opportunity to attend similar seminars and workshops on themes which are closely related to the content of this research. Some of the knowledge collected from those seminars contributed in the knowledge acquisition process. Some of the seminars and workshops attended by the researcher are listed below:

- Developing Disaster recovery strategy for Banking & Financial Institutions in Dubai, 12 - 14 December 1995;
- Crisis Management and Disaster Prevention, Bahrain, 19 - 21 April 1995; and
- Backups and Recovery by IBM in London, January 1992.

Researcher's Experience

After the Iraqi invasion in 1990, the researcher was appointed to take responsibility for collecting information concerning the destructive effects of the invasion on the information systems environment in the organisation for which he works (the Kuwait Institute for Scientific Research). The data and information collected were duly presented to the Public Authority for Compensation, established by the United Nations. Moreover, he was a member of the team which performed fast-recovery activities for critical applications and long-term recovery for other departments of the Institute. He was also assigned to take full responsibility of backups and recovery activities in his department (System Development) for more than three years. Those

activities and responsibilities gave him experience and insights into the recovery area, on which he was able to draw in establishing the ESDRSS knowledge.

Surveys and Case Studies

The results of several surveys and case studies in the disaster recovery field were collected and used to form part of the ESDRSS knowledge base. These include surveys carried out by the Disaster Recovery Journal and the Amedahl Executive Institute. In addition, the results of the fieldwork undertaken for this research, which covered large number of organisations in Kuwait, also proved very valuable for knowledge acquisition. Questionnaires were distributed to many disaster recovery coordinators and IT managers, supplemented by one or more follow-up interviews. The objectives, analysis and results of the fieldwork are presented and explained in Chapter 4.

7.4 Computing Maximum Allowable Downtimes

Since the aim of this project is to help in selecting a recovery strategy to save the business from great loss - or even being forced out of business - the main concern has been to concentrate only on these systems and applications that are crucial for organisational survival. As mentioned in Chapter 5, top management must carefully review the list of critical resources to ensure that only the truly critical ones are included.

The goal of this transaction is to determine how long the organisation can tolerate the interruption of its critical resources at the time of an adverse incident (Maximum Allowable Downtime or MAD). The MAD contributes significantly to the decision-making process for selecting the most appropriate recovery strategy and the amount of investment required to accommodate it (Arnell, 1990; Jackson, 1994). For example, organisations with a lower MAD (*e.g. one day*) would adopt the hot site strategy option, whereas organisations with a higher MAD (*e.g. two weeks*) would adopt the cold site strategy option.

The work done in the Business Impact Assessment phase of the methodology presented in Chapter 5 (collecting data and information for identifying and prioritising resources) is essential to arrive at the MAD estimation. Calculating the MAD is based on the cost of the denial of the computer centre and its resources with respect to the organisation's income. For several selected time intervals (e.g. 3 hours, 6 hours, 12 hours, one day, two days, etc.), the consequences of the denial of computer systems for each critical resource is estimated (see Table 5.2, Cost of Downtime). Each loss or potential exposure is quantified and the cost effects are aggregated for each time interval. Then the aggregate cost is compared to the revenue from the first selected time interval. The MAD (one of the selected time intervals) is reached when the cost of downtime exceeds the revenue.

The prototype ESDRSS starts by asking the user questions about his or her organisation such as: type of business; size of organisation; and average daily income. Then, the system proposes a particular time interval in response to the user's answers. Table 7.1 shows some examples of time intervals for different sizes and types of organisation. These time intervals have been derived from the fieldwork described in Chapter 4, surveys mentioned in Chapter 2 and other related materials. Then, the system asks the user to estimate the total consequences and downtime costs for all critical resources for the proposed time interval. The MAD is reached when the total downtime cost of all critical resources is equal or greater than the income for the specified time interval. If the total downtime cost is less than the income for that specified time interval, the system will double the time interval and ask the user to provide a new downtime cost for the doubled time interval. The process is performed again and again with a new downtime cost for each new time interval until it reaches the MAD (total time intervals). As already explained, the MAD is reached when the cost of downtime exceeds the income for the selected time interval. A special algorithm has been developed to calculate the MAD. This code can be found in Appendix A. An example of how the MAD is calculated is provided at Appendix B.

Table 7.1 - Examples of Time Intervals for Some Types of Organisation

Organisation type	Giant	Large	Medium	Small
Financial	3 hrs	6 hrs	12 hrs	One day
Manufacturing	One day	2 days	5 days	7 days
Government	One day	2 days	5 days	7 days
Computer services	One day	2 days	3 days	4 days
Retailing	6 hrs	12 hrs	One day	2 days
Telecommunication	3 hrs	6 hrs	12 hrs	One day
Education	2 days	3 days	5 days	7 days
Health Care	One day	3 days	5 days	7 days
Insurance	One day	2 days	3 days	5 days

In summary, the objective of the first transaction of ESDRSS is to automatically compute the MAD for a given organisation. This calculated MAD is applied later in the other two transactions to calculate the required investment and to select the most appropriate recovery strategy.

7.5 Computation of Investment

As explained in Chapter 5, disaster recovery experts highly recommend that the cost analysis of introducing a recovery strategy must only be used for budget purposes, not to make a decision on whether to adopt a recovery strategy or not (Robinson, 1993; Baylus, 1991; Arnell, 1990). A recent survey by Price Waterhouse showed that 70% of all UK organisations which did not recover from a major disaster within 48 hours failed to continue their businesses (Allen, 1992). Therefore, disaster recover plans should not be evaluated on the basis of cost-effectiveness.

The ESDRSS uses the second part of the mathematical model, Contingency Cost-Response Time Function, to calculate the required investment (see Chapter 2, Section 2.5.5). The Contingency Cost in the function means the cost of adopting a disaster recovery strategy, and response time means Maximum Allowable Downtime. The

function states that the cost of a recovery strategy is inversely proportional to the MAD. According to Subhani (1989), the relationship between the cost of a recovery strategy and the MAD can be expressed by the following equation; introduced earlier in section 5.6:

$$R = R^0 e^{-nt}$$

Where:

R = Cost of a recovery strategy with a maximum allowable downtime of t days;

R^0 = Cost of recovery strategy with instantaneous or zero MAD. In practice, such a strategy would be a duplicate site;

n = Parameter measuring the intensity with which the recovery strategy cost declines with the maximum allowable downtime;

t = Maximum allowable downtime in days; and

e = Exponential constant.

To illustrate the working mechanism of the above-mentioned function and how to apply it to compute the required investment, the following example is provided.

The user is asked to provide the system with two MAD estimates, and two annual cost estimates for two types of disaster recovery strategy that exist in its region or country. For example, let us assume it is a giant company with a multi-million pounds revenue and it gives the following answers: the annual cost for recovering within a half-day by subscribing to a hot site strategy is £200,000; the annual cost for recovering within three days by subscribing to the same strategy is £80,000; the annual cost for recovering within a week by subscribing to a cold site strategy is £10,000; and the annual cost for recovering within two weeks by subscribing to the same strategy is £2,000. Assume that the actual MAD for this organisation, as calculated in the first transaction, is 1.5 days.

From the above estimates the average recovery time, t , and the average annual cost, R , for the hot site strategy are: 1.75 days and £140000 respectively. The average

recovery time, t , and the average annual cost, R , for the cold site strategy are: 10.5 days and £6000 respectively.

Hot site: $t = 1.75$ days, $R = £140,000$

Cold site: $t = 10.5$ days, $R = £6,000$

Then we need to calculate, the cost of a recovery strategy with instantaneous or zero MAD, R^0 , and the parameter of measuring the intensity, n .

By taking the natural log of both sides of the “Contingency Cost-Response Time Function”, $R = R^0 e^{-nt}$, we get

$$\ln R = \ln R^0 - nt \dots \dots \dots \text{Relation 1}$$

By substituting the values of t and R for both strategies (hot and cold) in relation 1, we obtain the following two equations:

$$\text{For hot site:} \quad \ln(140,000) = \ln R^0 - 1.75n$$

$$\text{For cold site:} \quad \ln(6,000) = \ln R^0 - 10.5n$$

Solving the above two equations for two unknown variable R^0 and n , we find :

$$R^0 = £262,859 \text{ and} \quad n = 0.35998$$

Then, we take the MAD which was calculated for the company from the first transaction, 1.5 days for the purpose of this example, along with R^0 and n and substitute them in the Contingency Cost-Response Time Function to generate the investment required for a recovery strategy.

The optimal investment for having a recovery strategy is calculated by substituting 1.5 for t , 0.35998 for n , and 262,859 for R^0 in the Contingency Cost-Response Time Function $R = R^0 e^{-nt}$

$$R = 262,859 e^{-0.35998 * 1.5}$$

After a simple mathematical calculation, we come to a final figure of $R = \text{£}153,183$.

In conclusion, a company with a MAD of 1.5 days should invest $\text{£}153,183$ to accommodate a recovery strategy. However, as explained in Chapter 5 and mentioned again earlier in this chapter, this cost should not contribute to the decision of whether to adopt a recovery strategy or not. It should only be used for budget purposes.

According to Subhani (1989), the above model was validated and approved by a panel of 24 disaster recovery experts. Then, the model was tested in a real-world situation. It was applied to three companies which have already invested in recovery strategies. The model predicted the three cases fairly well (Subhani, 1989). As with any expert system application, information from the disaster recovery market needs to be collected before running the above function in ESDRSS. Such information, as shown in the above example, is the cost of two types of recovery strategy: hot site and cold site. These two strategies are chosen as reference points because:

- it is easy to obtain annual cost answers for these two strategies;
- the hot site strategy is known to be almost the most expensive whereas the cold site strategy is known to be the least expensive in commercial disaster recovery; and
- these two strategies are available in almost every country or region.

Although it is possible that the MAD, calculated in the first transaction, can be transferred directly to the Contingency Cost-Response Time Function in the second transaction without the intervention of the user, the researcher intentionally avoided

this approach. Rather, the system allows the user to modify the value of MAD and to input whatever value he wants. This is the case because there are some conservative top executives who may want to make changes to their maximum allowable downtimes for one reason or another. Top management views the organisation from a different perspective. Intangible issues such as political embarrassment, public image and media criticism may be important to top management and may reduce the overall MAD which would be tolerable for the organisation. The proposed expert system was therefore designed and implemented to have the flexibility needed to incorporate changes of this kind.

7.6 Recommending A Recovery Strategy

The third transaction of the ESDRSS involves the recovery strategy's recommendations. After computing the MAD and the investment needed, the next step is to recommend one or more recovery strategies. The output given by the first transaction, MAD, is extremely important in the third transaction so that an organisation may know its recovery time (see Table 5.10). The ability to provide recovery services with respect to response times varies among recovery strategies. The organisation's MAD is a major consideration in choosing an appropriate strategy. For example, organisations requiring immediate recovery (e.g. less than 1 hour) should own a duplicate site or subscribe to a disaster recovery vendor who provides a realtime recovery. On the other hand, an organisation with a high MAD (e.g. 10 days) value should (depending on other elements) subscribe to a commercial or co-operative cold site, or adopt another low time-response strategy. Some examples of how the decision is made are presented later in this chapter.

The elements that guide the decision-making process for selecting the most suitable recovery strategy were explained in more detail in Chapter 5. It may be helpful if they are mentioned here again:

- organisational characteristics (size of organisation, degree of dependency on computer technology and MAD);

- organisational requirements and recovery services required (external personnel support, work area, special-tailored hardware or software (hw/sw) platform, security, usage duration, location);
- characteristics of recovery strategies that meet the organisational requirements; and
- type of disaster that jeopardises the organisation.

The approach used to select the most suitable recovery strategy was also explained in Chapter 5. As mentioned before, the problem-solving mechanism for selecting a recovery strategy which was developed by the methodology was presented in the form of condition-action pairs: IF this *condition* occurs, THEN an *action* is recommended. Clearly, a technology that handles this type of problem-solving approach is required. In Chapter 6, two types of technology, conventional and expert systems, were evaluated to determine which is more suitable for delivering the required solution. The expert systems approach was selected as the most suitable technology for implementation. After investigation, the LPA *flex* toolkit was found to be the most suitable tool for implementation. The reasons for selecting expert systems and the *flex* toolkit were explained in the previous chapter.

The *flex* toolkit employs the rule-based knowledge representation technique, which suits the problem-solving approach presented by the developed methodology. The following sections illustrate how this technique is utilised by the ESDRSS. In addition, the capability of the *flex* toolkit's mechanism for questions is explained and some examples of questions which are used in selecting a disaster recovery strategy are presented. Then, *flex* rules, forward-chaining inference engine and some examples of constructed rules used by the ESDRSS are described.

7.6.1 Questions

Most expert system applications involve some type of communication with the user. This communication may be illustrated in many ways, such as graphs, pictures or questions. However, asking questions and providing answers is the most convenient

method for interacting with the user in expert systems technology. In the *flex* toolkit, this is achieved by invoking pre-defined questions. These questions may involve making single and multiple choice menus and typing information at a keyboard.

The prototype ESDRSS starts by asking several questions related to the organisation's characteristics and requirements and the type of threat to which the organisation may be vulnerable. The questions are associated with size, degree of dependency on computers, work area required, external personnel support, available recovery strategies, threats, etc. A list of all the questions used by the ESDRSS can be found in Appendix B. To give a flavour of how the question mechanism is used in *flex*, some examples are listed below:

Beginning with the *size* question, the question menu seen by the user is depicted in Figure 7.2. However, the actual structure of the question's coding is illustrated in Example (1).

Example (1)

question size

'What is the size of the organisation?';

choose of sizes ;

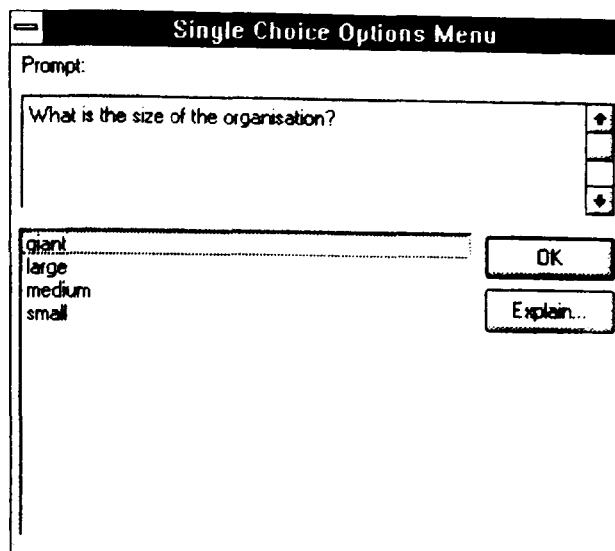
because size will help in deciding what MAD and recovery strategy are appropriate to the organisation .

group sizes

giant, large, medium, small .

The first line in example 1 indicates the name of the question: *size*. Next, the actual question is posed, *'What is the size of the organisation?'* The question menu consists of two parts (see Figure 7.2). The top part is the question sentence. The lower part contains a list of options provided by the system to choose from (*choose one of sizes*). The options (*giant, large, medium, small*) are grouped together in a function called *group sizes*.

Figure 7.2 - The Size Question Example



Moreover, *flex* provides a facility for attaching an explanation to questions, using the *because* clause. The explanation can either be some typed text to explain why a question has been asked and how to collect the required data, or it can be a name of a file to be browsed over. The explanation is presented whenever the user requests it (see the Explain button in Figure 7.2).

Another question mechanism used by *flex* is through single field keyboard *input*. The data entered can be either a text item, a floating-point number, an integer, or a set of such items. As in the Example (2), the *total_cost* question, the user is requested to enter a *number*.

Example (2)

question total_cost

'What is the total downtime cost for the critical resources for the proposed time interval?';

input number ;

because The value of this input should be collected in the Business Impact Assessment Phase.

Constraints can be added to keyboard input questions by using keywords **such that**. For instance, in Example (3) only yes or no responses are allowed. Otherwise a default, or customised, message will appear requesting the user to try again.

Example (3)

question hot_co-operative

'Is there a possibility that your organisation can establish a co-operative HOT site with other nearby organisations?';

input k such that yes_or_no_answer (k) .

relation yes_or_no_answer (yes) .

relation yes_or_no_answer (no) .

Questions in *flex* are invoked by typing **ask** whenever there is a request to ask a question. This can be part of the main program. Defining questions in *flex* is not a difficult task because *flex* uses its own language. The language is called Knowledge Specification Language (KSL). This language clearly distinguishes *flex* from other expert systems tools, especially when structuring rules. This is demonstrated in the next section.

7.6.2 Rules and Inferencing

Rules are considered to be the life-blood of expert systems technology. Most applications implemented by expert systems technology use the rule-based technique very extensively (Turban, 1992). There are two approaches for controlling inferencing in the rule-based technique: forward-chaining and backward-chaining. Forward-chaining is a data-driven approach. It starts from the available information as it comes and then tries to draw conclusions from it. Backward-chaining is a goal-driven approach. It starts from an expectation of what will happen (hypothesis), and then seeks evidence that supports the expectation. *Flex* supports both approaches by using the IF THEN format. The forward chaining rules are indicated in *flex* by the keyword

rule, and the backward-chaining rules are indicated by the keyword **relation** (Durkin, 1994).

It has been established from many expert systems applications that forward-chaining has proved itself very suitable to configuration problems. This occurs when it is not known what the final configuration will be, but it is known how to combine certain facts together according to some combining rules. Then, if the rules continue to be applied, everything is combined accordingly (Vasey, 1996).

It has become apparent to the researcher that the forward-chaining approach is the more suitable inferencing mechanism for solving the present problem because, basically, there is no goal or hypothesis to be drawn nor evidence to be derived to support the goal or hypothesis (backward-chaining). Rather, it is a problem where many facts are known about an organisation, the threats to it, and the recovery strategy options, and a conclusion needs to be drawn, based on these available facts.

Facts and knowledge in *flex* are organised in the knowledge base, separate from the control (inference engine), in the form of IF THEN rules. A rule is triggered when all of the antecedents of the implication are satisfied (i.e. when these antecedents are present in the memory.) An important feature of *flex* is that it does not have a limit to the number of rules. It can handle as many as several thousands rules. However, the ESDRSS contains only around sixty rules because it is only a prototype. The benefits of building a prototype system are explained in Chapter 6 under the heading Implementation and Prototyping.

7.6.3 Weighting of Rules

In rule-based systems there are always choice points where one rule is preferred to another. Attaching weights to rules is an option in *flex* which can assist in making these preferences. The weight of a rule reflects its relative importance with respect to the other rules in the system. Whenever two or more rules are simultaneously applicable, their relative weights can be compared to decide which one to use. Most

weighting systems in *flex* are static, with each rule being assigned a specific score. The more important the rule, the higher its score should be. Furthermore, *flex* allows for dynamic weighting systems, whereby the score attached to a rule is not fixed when the rule is defined, but is dependent upon some changing information. This can be done by asking the user to input a value for each item, which means adding more questions. In the prototype ESDRSS, the static weighting system is used for simplicity and to avoid adding more questions. The technique of weighting rules is explained in the following section, where examples of some rules are illustrated.

Flex also allows the attachment of an optional explanation to rules. This is used to explain why a rule was triggered. The explanation can either be some text displayed on the screen or information obtained by the user being allowed to browse through a file.

7.6.4 Examples of ESDRSS Rules

In the following paragraphs, several rules are presented to illustrate how recommendations are executed. First, however, it is important to note that the researcher has introduced a specific method for naming rules. Some of the names, such as *security*, can be read and understood easily whereas others need to be explained further. For instance, *giant_high_build_1*, means:

- the size of an organisation is *giant*,
- the degree of computer dependency is *high*;
- the anticipated threat type is *building*; and
- number *1* means that there are more rules to be applied with the same three features listed above.

Likewise, a rule with the name *small_high_floor_7* means that:

- the size of the organisation is *small*;

- the degree of computer dependency is *high*;
- the anticipated threat type is *floor*, and
- number 7 means that there are more rules to be applied with the same three features.

After explaining the naming method applied in the proposed system, the following are some examples of rules used in the ESDRSS and their descriptions.

A)

rule security_1

if security_requirement is 'very high'

and degree is high

and threat is regional

then short_strategy becomes 'Duplicate Site'

and long_strategy becomes 'The Duplicate Site should be used as a long term strategy'

and location becomes remote

score 100 .

The foregoing rule is called *security_1* which deals with the security issue. There are some organisations that have to satisfy a very high security level such as sensitive military installations, air-transportation command and control, air traffic control, government electronic mail centres, etc. For these organisations, it may be economically feasible to set up an entire alternative site in a geographically remote location in order to escape the same disaster.

B)

rule security_2

if security_requirement is 'very high'

and degree is high

and threat is building

then short_strategy becomes 'Duplicate Site'

and long_strategy becomes 'The Duplicate Site should be used as a long term strategy'
and location becomes 'within the city area'
score 100 .

The second rule, *security_2*, also deals with the same issue of security. However, this rule recommends a solution to an organisation which would be exposed to threats covering only a smaller area; for example, a fire in the computer building or a bomb in a busy street. In this scenario, it is a Type II threat where the buildings or even the streets around the affected area may be evacuated (see Table 5.1). Therefore, the alternative site should be located several miles away from the original site but within the city area to facilitate employees' transportation and customer satisfaction.

In the previous two rules; *security_1* and *security_2*, a high weighting score is given to both rules because the security issue is a **very high** factor. Therefore it is given a score of 100 which means that these rules have priority over other rules.

C)

rule giant_high_build_1
if size is giant
and degree is high
and recovery_time is immediate
and modification is yes
and threat is building
and personnel_support is no
and hot_co-operative is yes
and cold_co-operative is yes
then short_strategy becomes 'Duplicate site or co-operative hot site with realtime recovery'
and long_strategy becomes 'co-operative cold site'
and location becomes 'within the city area' .

The rule, *giant_high_build_1*, deals with organisations that have the following characteristics and requirements (or antecedents):

- ⇒ Giant size
- ⇒ High degree of dependency on computer
- ⇒ Requires immediate recovery
- ⇒ Requires special hw/sw installation arrangements
- ⇒ Organisation is exposed to Type II threat; building
- ⇒ Outside additional personnel support is not required
- ⇒ Possibility that organisation can establish a co-operative cold site with other nearby organisations.

In the above rule, the recommendations are given only if all of the above antecedents of the implication are satisfied. The short-term strategy is recommended to be either a duplicate site or a co-operative hot site with a real time recovery because immediate recovery and special hw/sw arrangements are required. The long-term strategy is recommended to be a co-operative cold site for two reasons:

(See co-operative cold site strategy in Chapter 3)

- external personnel support is not required; and
- a cold co-operative alternative site can be established and managed with other organisations with the same business and hardware/software platform.

At the end of the rule, the location of the alternative sites (short and long-term) should be several miles away from the original site but within the city area because the organisation is exposed to a Type II threats: building.

The following two rules, *small_high_floor_1* and *small_high_floor_2*, deal with organisations with similar characteristics but different requirements. The characteristics of both are:

- ⇒ Small size;
- ⇒ High degree of dependency on computer;
- ⇒ Organisation exposed to Type III threat: floor; and
- ⇒ Fast recovery, less than 2 days.

However, the requirements in the first rule (see example D), *small_high_floor_1*, are: hardware and/or software modifications and external personnel support. The mobile hot site strategy, which serves these requirements and fits the above-mentioned characteristics, is therefore recommended as a short-term strategy. A portable site is recommended as a long-term strategy which also meets the necessary requirements. The location is recommended to be adjacent to the organisation because the threat is a Type III threat: floor.

D)

rule small_high_floor_1
if size is small
and degree is high
and [recovery_time immediate or recovery_time is '1 to 2 days']
and threat is floor
and [modification is yes or personnel_support is yes]
then short_strategy becomes 'mobile hot site'
and long_strategy becomes 'portable site'
and location becomes 'adjacent to the organisation'
score 30 .

While the characteristics of the second rule (see example E), *small_high_floor_2*, are similar to those in the first rule, the requirements are different. The organisation does not need special-tailored hardware or software modifications. Therefore, a service bureau is recommended which can provide a relatively fast recovery (less than two days). The external personnel support factor is not included here because a service bureau can meet this requirement. A portable site is also recommended as a long term recovery strategy because it is only a small organisation. Again, the location is

recommended to be adjacent to the organisation because the anticipated threat is a Type III threat: floor.

E)

rule small_high_floor_2
if size is small
and degree is high
and [recovery_time is immediate or recovery_time is '1 to 2 days']
and threat is floor
and modification is no
and service_bureau is yes
then short_strategy becomes 'service bureau.'
and long_strategy becomes 'portable site'
and location becomes 'adjacent to the organisation' .

It can be seen that there are some questions which have not been asked because they do not apply to a particular situation. For example, the question related to work area has not been asked in the previous two rules because the scenario is of a Type III threat: floor. Since the anticipated damage area is only expected to spread over one floor or less, and employees will not be located in another location, there is no need to ask a question regarding the need for an additional working area. This has been explained in the methodology developed in Chapter 5.

F)

rule small_high_build_4
if size is small
and degree is high
and [recovery_time is immediate or recovery_time is '1 to 2 days']
and threat is building
and personnel_support is no
and modification is no
and work_area is no

*and service_bureau is no
 and reciprocal is yes
 then short_strategy becomes 'reciprocal agreement.'
 and long_strategy becomes 'commercial cold site'
 and location becomes 'within the city area'.*

Example F applies to small organisations that need to achieve recovery within 2 days. It fits organisations that do not have any requirements except for an alternative site which provides similar hardware and software platforms. This solution is only recommended if there is the possibility of making a reciprocal agreement with another nearby organisation.

G)

*rule medium_low
 if size is medium
 and degree is low
 then short_strategy becomes 'withdraw of service or manual procedure.'
 and long_strategy becomes "
 and location becomes ".*

Example G illustrates the situation when there is a medium-size organisation that does not depend on computers in running its business. The feasible and economically sound recommendation for this type of organisation is to withdraw from the computerised service until everything is back to normal, or to provide those services manually until the computer is up again.

The above examples are given only to illustrate the user interface method, the rules structure and the adaptability of the *flex* expert system toolkit. The examples also show and explain how recommendations are given in the prototype ESDRSS. All the rules and the program code for the whole system can be found in Appendix A.

7.7 Testing

Disaster simulation exercises are often used to test the activities and decision-making of disaster recovery plans (Rosenthal & Sheiniuk, 1993; Rosenthal & Himel, 1991). However, such simulation is a costly exercise, and there is a general reluctance among senior managers within organisations to undertake this exercise because of the expense and the disruption of normal business activities (Toigo, 1989; Doughty, 1993; Jackson, 1994). The cost of testing any methodology or plan requires several man-months (CCTA, 1989). According to Barbara DePompa, a certified disaster recovery planner, 'the effort of the first integrated test will cost as much as \$10,000' (DePompa, 1995). In order to carry out a test for a disaster recovery methodology to check its validity on a real-world situation, the following criteria must first be satisfied:

- many interviews have to be carried out with key personnel or department managers to collect the required data (See examples of questions in the Threat Assessment and Business Impact Assessment Phases in Chapter 5);
- the cost of downtime for all applications and systems must be recorded. This may take considerable time and may be regarded as sensitive data by some organisations;
- employees should be knowledgeable about the importance of DRPs so they can co-operate with the test process; and
- individuals who are assigned to do the simulation test should be trained in how to collect the information needed and how to perform the required activities.

Due to the above-mentioned obstacles, it is difficult to locate organisations which are willing to conduct a simulation test of the present research on their applications and systems. Furthermore, like any other disaster recovery plan, the true soundness of the recommendations which are given by the methodology and expert system cannot be

really tested until an actual disaster occurs. However, the normal type of test for any system was carried out to check the logic of the proposed system and to check against any programming errors.

7.8 Concluding Remarks

The prototype ESDRSS presented above shows that expert systems technology is capable of accommodating the heuristics required to produce solutions to the problems and issues involved in disaster recovery strategy selection. Furthermore, it is clear from this project that expert systems technology and the rule-based mechanism can be used to test and validate the developed methodology and therefore can provide IT managers or disaster recovery co-ordinators with a workable system in a usable form.

Chapter 8

Discussion and Conclusion

8.1 Summary of the Research

The number of organisations who rely on computerised systems to perform their day-to-day operations and to help them in making decisions has grown rapidly over the last few years and continues to expand. Thus, information systems are now considered to be a basic component of nearly all private and public organisations. They are here to stay and their uses will continue to grow in all sectors in the years ahead. On the other hand, the essence of good management is the rational use and protection of resources. Next to personnel, an organisation's most important resource nowadays is information. Therefore, effective management of the computer centre and information resources will be an essential determinant of business success. The destruction or loss of these resources can be a nightmare and in many cases, unless restored promptly, may lead to an end of trading for the business.

In recent years, several significant disasters have occurred, which received extensive news coverage. This has increased management awareness and understanding of the need for a means of protection to survive such disasters. This, in turn, brings into sharp focus the necessity for a carefully constructed disaster recovery plan to assure the continuation of service and business. Therefore, the time and attention devoted to disaster recovery planning (DRP) have increased dramatically over the last few years.

To look more closely at the effects of disasters on organisations and the importance of adopting DRPs, the researcher carried out a case study on organisations in Kuwait to identify major problems facing IT managers on disaster recovery issues. The study

identified a number of problems and also identified factors which have contributed a great deal to the development of the proposed solution. The case study and its findings were explained in Chapter 4.

The literature in this field and the above mentioned study show that senior management and IT directors have begun to realise the need for disaster recovery plans. However, several questions are often raised by them in this context, such as: How long can the organisation tolerate the failure of its computer systems? Are we spending too much or too little on a recovery strategy? What type of recovery strategy is best appropriate for our IT centre?

Seeking answers to these questions was the main target of part of this research. Several findings relevant to an understanding of the disaster recovery field were analysed before finalising the proposed solutions. The findings and the solutions are briefly explained in the following sections.

8.1.1 Findings

The major findings of the present research can be summarised as follows:

- 1) insurance is not considered to be an acceptable alternative to recovery because it does not put the company back in business. Instead, it should be part of the disaster recovery plan. Indeed, insurance money is needed to fund the recovery efforts following a disaster;
- 2) disaster avoidance countermeasures are valid to prevent some types of threats. But because there are other types of disaster which are beyond the control of any preventive countermeasures, another method of protecting the business and assuring the survival of the organisation is essential;
- 3) much of the literature in the disaster recovery area has dealt with the need for disaster recovery planning, how to develop and implement disaster recovery

plans, and the consequences of not having one. The wider issues of selecting the most suitable recovery strategy and answering the questions posed by IT managers have not been fully and adequately addressed; and

- 4) the utilisation of expert systems technology in the field of disaster recovery has still to be tested and further investigations to check feasibility of employing this new technology is needed.

The literature review presented in Chapter 2 shows that few computerised systems in utilising expert systems technology in the field of disaster recovery were introduced recently and others are under development. For example, the AUDIT system, introduced in late 1996, and the CCTA IT Infrastructure Contingency Planning Module developed by the UK Government's Central Computer and Telecommunications Agency (CCTA), which is under development. The AUDIT system which is similar to the end-product of the present research, as explained in section 2.6.1.1, is an outstanding attempt in this area. However, the present work provides better solutions to the disaster recovery problems since it uses a constructed methodology and applies a very powerful toolkit, *Flex*.

Based on the increasing demands of disaster recovery planning, including answers to important questions raised by IT managers and the above findings, a comprehensive approach has been developed in this research to assist in the process of selecting the most suitable disaster recovery strategy. The approach provides a range of methods enabling the continuation of services and the provision of guidance for fast recovery in the event of a disaster that affects the IT installation. In addition, a prototype expert system is developed as a practical end-product to assist IT managers in the strategy selection process.

The decision to undertake this particular area of research was influenced by the following facts:

- a) the potential business impact on organisations of disasters that cause services interruptions is considerable and the possibility of business failure, if not bankruptcy, cannot be discounted (see Chapters 1, 2 and 4). Thus, more attention should be given to this area;
- b) The researchers' past experience in the disaster recovery area notably after the Iraqi invasion of Kuwait (see Chapter 7, Knowledge Acquisition) enabled him to cite some of the weak points associated with existing disaster recovery methodologies. These weaknesses presented an obvious focus for new research. An extensive investigation of the disaster recovery literature and lengthy interviews with well-informed individuals in the field conformed his initial view that there is a strong need to develop a new methodology in order to find solutions to these weak points; and
- c) checking the feasibility of using expert systems technology, as successfully applied in many other fields, to make a useful contribution in the disaster recovery area.

8.1.2 The Methodology

The methodology developed in Chapter 5 contained five phases that provide a step-by-step approach to ensure that the entire recovery strategy selection process is covered. The phases are: Threats Assessment; Business Impact Assessment; Recovery Strategy Analysis; Cost Analysis, and System Recommendations. The following paragraphs briefly explain these phases.

The methodology is intended to assist IT managers in identifying the potential disasters that threaten their companies. A new approach for classifying such threats was developed so that the nature of the threats can contribute to a more informed decision-making process for selecting a recovery strategy. This analysis was carried out in the first phase of the methodology and is called the Threats Assessment.

Once the threats assessment has been finalised, the second phase, Business Impact Assessment, was presented. In this phase, computerised systems and applications are identified and then prioritised in terms of criticality to the organisation. The focus is, however, shifted to systems that are deemed critical in terms of importance to the survival of the organisation following a disaster. Then, the overall maximum allowable downtime for which an organisation can tolerate the failure of its computer systems is determined. The main organisational requirements were also identified in this phase.

The IT manager, the disaster recovery co-ordinator/team, or whoever is in charge has the responsibility for analysing and selecting the most suitable and efficient recovery strategy. This strategy must meet the true recovery requirements of the organisation. There are a number of factors that influence the right decision in terms of selecting the recovery strategy. Some of these factors are related to the organisation itself. Others relate to the characteristics of different recovery strategies and how will they fit the true recovery requirements of the organisation. The third phase, Recovery Strategy Analysis, explains these factors and shows how the recovery strategy selection process is undertaken.

Disaster recovery experts state that the cost analysis of accommodating a recovery strategy must only be used for budget purposes, not to make a decision on whether to adopt one or not. Disaster recovery plans should not therefore be evaluated on the basis of cost-effectiveness (Robinson, 1993; Baylus, 1991). This research, however, does not ignore the fact that management needs to have some sort of indication of how much they need to spend on a disaster recovery strategy. This research provides IT directors with a method for calculating the amount of investment required to spend on a disaster recovery strategy. This was carried out in the fourth phase of Cost Analysis

In the final phase, a prototype computerised system was developed to provide IT managers with some recommendations regarding strategy selection based on several inputs from the user. This system is called the Expert System for Disaster Recovery Strategy Selection (ESDRSS).

8.1.3 The ESDRSS

A major objective of this research was to develop and deliver a computerised system as an end-product, which could test the validity of the previous developed methodology and be provided to IT managers, disaster recovery co-ordinators, disaster recovery consultants and others in a usable form. Therefore, the prototype Expert System for Disaster Recovery Strategy Selection (ESDRSS) was developed to perform some computations and rule-based decisions regarding the continuation of business and services following a disaster. The development of the prototype ESDRSS was accomplished via two stages: 1) the development of a structured methodology for the selection process of a recovery strategy in Chapter 5; and 2) the development and implementation of the proposed prototype system, based on the previous methodology using the expert systems technology. This system, after expansion and several tests, can be used as part of the master disaster recovery plan for any organisation. The prototype ESDRSS (see Figure 7.1) consists of the following three major components (or transactions):

- computation of the Maximum Allowable Downtime (MAD);
- computation of the required investment for fitting a recovery strategy; and
- recommending a disaster recovery strategy, based on the organisation's requirements and each recovery strategy's characteristics using a rule-based knowledge representation mechanism.

8.2 Who Will Benefit from this Research?

Of course, everyone hopes that they will never have to exercise such a disaster recovery methodology, but it is important that they should feel secure if eventually does arise that there is a means of protection to assure the continuity of work. Therefore, this research has developed, implemented and delivered a methodology and a prototype expert system which reveal what information is required and how that information is managed as well as guiding decision-makers regarding investments and

the selection of the most suitable recovery strategy. The following individuals are expected to benefit from both the methodology and the expert system:

- 1) IT directors in organisations that depend heavily on computers and where the availability of computers is critical for their revenue. The computer information systems are used to perform daily work and/or to make critical decisions, such as in financial institutions;
- 2) IT managers in organisations which perform activities that utilise computer facilities but can tolerate interruption of systems for a period of time, such as research institutes;
- 3) Disaster recovery co-ordinators who need to submit reports to IT managers or top management regarding investments and recommendations on disaster recovery strategies;
- 4) Disaster recovery consultants who are hired by organisations to carry out an analysis about the organisation's recovery requirements eventually to provide the necessary recommendations for future action;
- 5) Disaster recovery vendors who provide different strategies for different organisations. This research helps them to decide which of the recovery strategies they hold would be most suitable for existing or prospective subscribers; and
- 6) Insurance underwriters and other risk assessors.

8.3 Future Research

The time has finally arrived for IT managers to give adequate attention to the concept of disaster recovery, or what others call business continuity (Iyer and Diez, 1997). Therefore, disaster preparedness is considered to be a subject which deserves more

research studies in the future, as the increasing awareness of the importance of this field continues to grow. Drawing on the experience gained through this research, the following observations are suggested for future work in this field:

- 1) The development of a methodology and expert system to provide informed recommendations as between various vendors who offer different strategies. Vendors differ, within each recovery strategy, depending on many parameters such as: hardware and software compatibility, reputation, communications facilities, reliability and supplies;
- 2) The development of a methodology and expert system to compare the various disaster recovery planning software system. The main comparison issues are likely to be on product design, flexibility, user friendliness, price, product scope, product support, the complexity of DRPs used, etc.;
- 3) Implementing the developed methodology in Chapter 5 using alternative mechanisms, in place of the rule-based mechanism, such as frame, object oriented programming or neural networks;
- 4) The development of a relational database that contains all the relevant information about disaster recovery vendors such as: strengths, limitations, market position, location and other characteristics. Such a database assists organisations to select among these vendors; and
- 5) Investigating the utilisation of the available WWW-site interfaces for recovery to support the many emerging Internet and Intranet applications.

References

Aasgard, D. (1979) 'An Evaluation of Data Processing Machine Room Loss and Selected Recovery Strategies', University of Minnesota Management Information Systems Research Centre, working papers.

AL-Watan Newspaper, Friday 22 March 1996, no 7222/1668 - Year 35, p. 3. Kuwait.

Acs, Z. and Audretsch, D. (1993) 'Small Firms and Entrepreneurship: an East-West Perspective', Cambridge, MA: Cambridge University Press., USA.

Allen, C. (1992) 'Protecting the Business', Share Europe Anniversary Meeting, Switzerland, pp. 1-11.

Anderson, R. (1992) 'Information & Knowledge Based Systems: An Introduction', Prentice Hall.

Arner, P. (1970) 'Computer Application in Government', The Computer Impact, Englewood Cliffs, Prentice-Hall, pp. 12-20.

Arnell, A. (1990) 'Handbook of Effective Disaster Recovery Planning: A Seminar/Workshop Approach', McGraw-Hill Publishing Comp., New York.

Bannister, D. (1982) 'How to Cope When Disaster Strikes', Electron Weekly, 7th July, p. 7.

Bramer, M. (1990) 'Practical Experience in Building Expert Systems', John Wiley & Sons Ltd, England.

Barrett, M. and Beerel, A. (1988) 'Expert Systems in Business: A Practical Approach', Ellis Horwood Limited, Chichester, England.

Bates, R. (1992) 'Disaster Recovery Planning: Networks, Telecommunications, and data communications', McGraw-Hill.

Baur, G. and Pigford, D. (1990) 'Expert Systems for Business: Concepts and Applications', Boyd & Fraser Publishing Company, Boston, USA.

Baylus, E. (1991) 'Disaster Recovery Handbook', Blue Ridge Summit, Chantico Publishing Company, PA.

Beerel, A. (1993) 'Expert Systems in Business: Real World Applications', Ellis Horwood Limited, Chichester, England.

Biasiotti, A. (1988) 'Making Allowances for Nature', Planning for Disaster Recovery, November, pp. 58-59.

Blair, G. (1987) 'Confronting The Contingency Planing Dilemma', Australian Computer Conference, Melbourne, Australia, pp. 579-88.

Blinkhorn, S. (1993) 'Its Just One Damn Crisis After Another', 1993 International Emergency Management & Engineering Conference - 10th Anniversary: Research & applications, 29 March - 1st April, pp. 253-8.

Brown, R. (1993) 'What You Need to Know to Plan for Disaster', Networking Management, Vol. 11, Pt. 4, April, pp. 25 - 27.

Burch, J. & Grudnitski, G. (1989) 'Information Systems: Theory and Practice', New York, John Wiley & Sons.

Carroll, J. (1984) 'Managing Risk: A computer-Aided Strategy', Stoneham, M.A. . Butterworth Publisher.

Carpentar, L. (1993) 'Learning from Hurricane Andrew: Arby's Disaster Plan and Recovery', Journal of Systems Management, June, pp. 8 - 11.

Carter, Roy (1988) 'Dependence and Disaster Recovery from EDP Systems Failure', Management Services, December, pp. 20-22.

CCTA (1989) 'Contingency Planning', Central Computer Telecommunication Agency, UK, Government Publication.

The Centre for Research and Studies on Kuwait (1994) 'The Iraqi Aggression on Kuwait: The Truth and the Tragedy', 1994.

Chin, M. (1992) 'Disaster Mitigation in the Caribbean Using Expert System', Expert Systems Application, Vol. 5, Pt. 3, pp. 437-40.

Cline, C. (1988) 'Contingency Disaster Planning at Fidelity', Planning for Disaster Recovery, November, pp. 55-57.

Copenhaver, John (1997) 'Contingency Planning As Asset Management: Legal Issues', Disaster Recovery Journal, Vol. 10, Issue 2, P. 26, Spring 1997.

Cox, Lawrence (1996) 'How Do You Pay For It?', Disaster Recovery Journal, Vol. 9, Issue 2, Spring.

Crawford, W. (1993) 'Not A Total Disaster', LAN Magazine, August, pp. 79-87.

Daler, T.; Gulbrandsen, R.; Melgard, B. and Sjolstad, T. (1989) 'Security of Information and Data', New York, Ellis Horwood, Ltd.

Danish, T. (1994) 'A Knowledge-Based Decision Support System for Computer Disaster Prevention in IT Centres', Ph.D. Thesis, University of Newcastle Upon Tyne.

Datapro Report (1993) 'Business Recovery Issues: A Time of Challenge', A nationwide teleconference on disaster recovery hosted by SunGard Recovery Services, May 27, 1993.

Datashield Report (1993) 'Datashield Disaster Recovery Services', October, Datapro.

Davies, D. (1993) 'Up the Creek? A Major New Survey of the Preparedness of UK Companies for Computer Related Disaster', Computer Law Security Report, Vol. 9 Pt. 4 pp. 195 - 6.

Department of Trade and Industry (1990) 'Expert Systems Opportunities - Guidelines for the Introduction of Expert Systems Technology', London: HMSO

DePompa, B. (1995) 'Disaster Strikes! Are You Ready?' Disaster Recovery, May 15, Internet.

Devlin, E. (1996) 'The Perspective of Ed Devlin', Disaster Recovery Journal, Vol. 10, Issue 1, P. 9, Winter 1997.

Doughty, K. (1993) 'Auditing The Disaster Recovery Plan', EDPACS, Vol. 21, Pt. 3, pp. 1-12, September.

DRJ (1997) 'Disaster facts', Disaster Recovery Journal, Vol. 10, Issue 1, Winter, p. 43.

Durkin, J. (1994) 'Expert Systems: Design and Development', Macmillan Publishing Company, New York.

Elbra, R. (1992) 'Computer Security Handbook', NCC Blackwell Limited, Oxford, England.

Epich, R. & Persson, J. (1994) 'A fire Drill for Business', Information Strategy Executive, Vol. 10, Pt. 2, Winter, pp. 44-7.

Everett, D. (1988) 'How to Asses Risks to your Technical Resources', Planning for Disaster Recovery, November, pp. 62-78.

Faithfull, M. and Watt, S. (1991) 'The survivor's Guide to IT Centre Design', Elsevier Science Publishers Ltd.

FIPS Publication 65 (1979) 'Guidelines for Automated Data Processing Risk Analysis', US Department of Commerce, National Bureau of Standards.

FIPS Publication 87 (1981) 'Guidelines for Automated Data Processing Contingency Planning', US Department of Commerce, National Bureau of Standards.

Fisher, Patricia (1996) 'How to Conduct A Business Impact Analysis', Disaster Recovery Journal, Vol. 9, Issue 3, Summer 1996.

Frenchman, K. (1988) 'UBF Protects itself against Disaster', Planning for Disaster Recovery, November, pp. 38-42.

Goldblum, E. (1982) 'Computing: Planning For Disaster', Architects Journal, Vol. 176, 27th October, pp. 73-76, 79.

Giarratano, J. and Riley, G. (1994) 'Expert Systems: Principles and Programming', PWS Publishing Company, Boston, USA.

Haack, M. (1984) 'Insuring the Data Processing Risk', Bests Review, January, pp. 44-50.

Harmon, P. and King, D. (1985) 'Artificial Intelligence in Business Expert Systems', John Wiley & Sons Ltd, New York.

Harmon, P; Maus, R. and Morrissey, W. (1988) 'Expert Systems Tools and Applications', John Wiley & Sons Ltd, England.

Hars, Adele (1996) 'Trial By Fire: Credit Lyonnais' Battle to Save the Data - the Day', Disaster recovery Journal, Vol. 9. Issue 4, Fall, pp. 10-11.

Harrison, B. (1994) 'Lean and Mean: The Changing Landscape Corporate Power in the Age of Flexibility', HarperCollins Publishers, Inc. New York

Hassig, L. (1991). 'Computer Basics', Alexandria, Time-Life Books.

Hayes-Roth, F; Waterman, D. and Lenat, D. (1983) 'Building Expert Systems', Addison-Wesley, London.

Hearnden, K. (1993) 'Corporate Computing 1993: Key Business Issues in Contingency Planning for Business Recovery, A Research Study of 421 UK Organisations', Security Journal, Vol. 4, No. 4, October, pp. 205-220.

Heirlein, E. (1993) 'Recovery Management', Computer Security, Vol. 12, Pt. 4, June, pp. 334-7.

Hiles, Andrew (1992) 'Surviving A Computer Disaster', Computing & Control Engineering Journal, May, pp. 133-6.

Hirshleifer, Jack (1988) 'Price Theory and Application', Prentice-Hall, New York.

Hyde, J. (1993) 'Disaster Recovery sites: an Overview', Datapro, July 1993.

IBM Report (1993a) 'IBM Business Recovery Services', Datapro, May 1993.

IBM Report (1993b) 'Up the Creek: The Business Perils of Computer Failure', In association with Loughborough and the Computing Services Association, UK.

IBM Report (1995) 'IBM Business Recovery Services: The Ultimate Hot Site', Summit' 95, May 7-10.

Iyer, R. and Diez, R. (1997) 'Enhancing Senior Management Awareness and Gaining its Commitment to Business Continuity Planning', Disaster Recovery Journal, Vol. 10, Issue 2, Spring, pp. 11-12.

Jablonowski, Mark (1995) 'Scenario-Based Risk Analysis', Systems Support Inc., Internet, October 6.

Jackson, C. (1994) 'Business Continuity Planning: The Need and tApproach', Datapro, February.

Jackson J. (1988) 'How to Select a Hot Site Vendor', Planning for Disaster Recovery, November, pp. 28-33.

Jackson, K. and Hruska, J. (1992) 'Computer Security Reference Book', Butterworth-Heinemann, Oxford, UK.

Jackson, M. (1992) 'Understanding Expert Systems Using Crystal', John Wiley & Sons Ltd, England

Jackson, P. (1990) 'Introduction to Expert Systems', Addison-Wesley Publishing Company.

Katayama, T. (1993) 'International Decade for Natural Disaster Reduction: Are We Chasing a Dream', Proceedings 19th Annual International Conference on Industrial Electronics Control and Instrument, pp. 1-6.

Kerby, J. (1990) 'Disaster Recovery: The State of The Market', Computer Fraud Security Bulletin, pp. 17 - 18.

King, J. (1993) 'Contingency Plans & Business Recovery', Information System Management, Vol. 10(4), pp. 56-9, Fall 1993.

Lauletta, Patricia (1993) 'Disaster Recovery in a Client Server Environment: Who's Responsible?', Datapro, Security, September, pp. 1-3.

Marcella, A. and Rauff, J. (1996) 'Automated Disaster Recovery Plan Auditing: Prospects for Utilising Expert Systems to Evaluate Disaster Recovery Plans', Disaster Recovery Journal, Vol. 9, Issue 4, Fall, pp. 70-75.

Mauch, J. and Birch, J. (1993) 'Guide to the Successful Thesis and Dissertation', Marcel Dekker Inc., New York.

Mercorella, R. (1995) 'The Importance of a Disaster Recovery Contingency Plan', Disaster Recovery Journal, October.

Morris, D. (1988) 'Contingency Planning at Colonial Mutual', Planning for Disaster Recovery, November, pp. 50-54.

Moses, R. (1992) 'Risk Analysis and Management', Computer Security Reference Book, pp. 243-262, Butterworth-Heinemann, Oxford, UK.

Mower, Mark (1991) 'Planning for a Disaster', Management Services, pp. 21-23.

Mueller, D. (1990) 'The Dynamics of Company Profits. An International Comparison', Cambridge University Press.

Musgrave, P. (1988) 'Assessing and Protecting against Risk', Planning for Disaster Recovery, November, pp. 22-27.

Naylor, C. (1988) 'Building Your Own Expert System', John Wiley & Sons Ltd, England.

NBS Special Publication, 500-133 (1985) 'Technology Assessment: Methods for Measuring the Level of Computing Security', Us Department of Commerce, National Bureau of Standards.

Orr, J. (1988) 'Guide to Contingency Planning', Planning for Disaster Recovery, p 11 - 21.

Oscar, N. and Chien, Y. (1991) 'Knowledge-Based Systems: Fundamentals and Tools', IEEE Computer Society Press, California.

Own, Jeffrey (1993) 'Network Disaster Recovery', Datapro, Network Planning, April, pp. 1-9.

Parker, D. (1981) 'Managers Guide to Computer security', Reston, VA, Reston Publishing.

Pashigian, B. (1995) 'Price Theory and Application', McGraw-Hill, New York.

Patterson, S. (1994) 'Good Planing Saves Information and Ensures Speedier Recovery', IS Audit Control, Vol. 1, pp. 14-16.

Peach, S. (1991) 'Disaster Recovery: An Unnecessary Cost Burden or an Essential Feature of any DP Installation', Computer & Security, Vol. 10, pp. 565-8.

Perryman, M. (1988) 'Business Recovery Planning at Manufacturers Hanover', *Planning for Disaster Recovery*, November, pp. 43 - 49.

Powell, Jeanne (1997) 'OmniCentric Hot sites: Local Access to Global Continuity', *Disaster Recovery Journal*, Vol. 10, Issue 2, Winter, pp. 26-28.

Ramsey, S. (1988) 'Taking the 'Hot' out of Hot Site Contracts, *Planning for Disaster Recovery*, November, pp. 34 - 37.

Ratliff, J. (1994) 'Realtime Recovery From Concept to Reality', *Datapro*, January.

Redmond, M.; Luongo, J. and Tietz, J. (1996) 'Recovery Planning Can Speed Your Return to Business', *Disaster Recovery Journal*, Vol. 9, Issue 4, Fall, pp. 38-39.

Reed A. (1992) 'Computer Disaster: The Impact On Business in the 1990's', *Computer Sci. Technology*, Vol. A-15, pp. 13-21.

Robinson L. (1993) 'Contingency Planning and Disaster Recovery', *Faulkner Technical Reports*, June, pp. 1-9.

Rosenthal, Paul and Himel, Barry (1991) 'Business Resumption Planning: Exercising Your Emergency Response Teams', *Computer & Security*, Vol. 10, pp. 497-514.

Rosenthal, Paul and Sheiniuk, Gene (1993) 'Business Resumption Planning: Exercising The Disaster Management Team', *Journal of Systems Management*, June, pp. 12-16, 38-42.

Schreider T. (1995) 'US Hot Site Industry Market Analysis & Forecast', *Disaster Recovery Journal*, October.

Schreider T. (1996) 'White Paper: The Internet', *Disaster Recovery Journal*, Vol. 9,

Issue 3, Summer.

Shafto, Tony (1991) 'The Foundations of Business Organisation', Stanley Thornes Ltd, England.

Smith, M. (1989) 'Commonness Computer Security', McGraw-Hill Book Company.

Smith, D. (1990) 'Computer Systems Disaster Report is Unequivocal: Plan Now or Risk Corporate Roulette', Information Technology & Public Policy, Vol. 9, No. 1, pp. 21-5.

Storey, D. (1988) 'Entrepreneurship and the New Firm', Antony Rowe Ltd, Chippenham, Wiltshire, UK.

Subhani, S. (1989) 'Decision Model for Optimal Selection of Recovery Plans for Computer Outages', Ph.D. Thesis, University of Texas at Arlington.

Thompson, S. and Wright, M. (1988) 'Internal Organisation: Efficiency and Profit', Philip Allan Publishers, Oxford, UK.

Tilley K. (1993) 'Business as usual', Computer Bulletin, Vol. 5, Pt 6, December, pp. 7-9.

Turban, E. (1992) 'Expert Systems and Applied Artificial Intelligence', Macmillan Publishing Company, New York.

Turban, E. (1995) 'Decision Support and Expert Systems', Prentice-Hall, Inc. Englewood Cliffs, NJ.

Toigo J. (1989) 'Disaster Recovery Planning: Managing Risk and Catastrophe in Information Systems', Prentice-Hall, New Jersey.

UNESCO (1991) Report Mission to investigate damages inflicted on Education, Culture, and Scientific Research in Kuwait during the Iraqi occupation at the request of the Secretary General of the United Nations; 30 March 1991.

Valley, Ian (1992) 'Think About the Unthinkable You Could be Next', Works Management, December, pp. 28-31.

Vasey, Phil (1996) 'Flex Expert System Toolkit: Technical Reference', Logic Programming Associates Ltd, London, UK.

Vijayaraman, B. and Ramakrishna, H. (1993) 'Disaster Preparedness of Small Businesses with Micro-Computer Based Information Systems', Journal of Systems Management, June, pp. 28-32.

Waterman, D. (1985) 'A guide to Expert Systems', Addison-Wesley Publishing Company.

Watt, S. (1988) 'The Implications of no Formalised Contingency Plan', Planning for Disaster Recovery, November, 60 - 61.

Wesselingh, E. (1990) 'Disaster Recovery Contingency Planning', 13th National Computer Security Conference - Information System Security, 1 - 4 October, Vol. 2, pp. 385-391.

Williams, J. (1995) 'Minimising Damage with a Disaster Recovery Plan' Cyber Business Journal, Internet, October 1995.

Wilson, N. and McClean, S. (1994) 'Questionnaire Design: A Practical Introduction', Audio Visual Services, University of Loughborough, UK.

Wold, G. (1996) 'Some Techniques for Business Impact Analysis', Disaster Recovery Journal, Vol. 9, Issue 4, pp 27 - 33, Fall 1996.

Wrobel L. (1990) 'Disaster Recovery Planning for Telecommunications', Artech House.

Wrobel L. (1993) 'Writing Disaster Recovery Plans for Telecommunications Networks and LAN', Artech House.

Zeidman, B. (1996) 'An Introduction to Remote Backup', Disaster Recovery Journal. Vol. 9, Issue 3, Summer.

Appendix A

System Code

```
%%%%%%%%%%
% this part is the main body of the program
%%%%%%%%%%
```

Actions

```
action flex_starter
do ti
and ask_continue1
and inv
and ask_continue2
and rule_q
and ask_continue
and ri0 .
```

```
%%%%%%%%%%
%% allocating time intervals
%%%%%%%%%%
```

```
ti('Financial', giant, 0.125).
ti('Financial', large, 0.25).
ti('Financial', medium, 0.5).
ti('Financial', small, 1).
```

```
ti('Manufacturing', giant, 1).
ti('Manufacturing', large, 2).
ti('Manufacturing', medium, 5).
ti('Manufacturing', small, 7).
```

```
ti('Government', giant, 1).
ti('Government', large, 2).
ti('Government', medium, 5).
ti('Government', small, 7).
```

```
ti('Computer Services', giant, 1).
ti('Computer Services', large, 2).
ti('Computer Services', medium, 3).
ti('Computer Services', small, 4).
```

ti('Retailing', giant, 0.25).

ti('Retailing', large, 0.5).

ti('Retailing', medium, 1).

ti('Retailing', small, 2).

ti('Telecommunication', giant, 0.125).

ti('Telecommunication', large, 0.25).

ti('Telecommunication', medium, 0.5).

ti('Telecommunication', small, 1).

ti('Education', giant, 2).

ti('Education', large, 3).

ti('Education', medium, 5).

ti('Education', small, 7).

ti('Health Care', giant, 1).

ti('Health Care', large, 3).

ti('Health Care', medium, 5).

ti('Health Care', small, 7).

ti('Insurance', giant, 1).

ti('Insurance', large, 2).

ti('Insurance', medium, 3).

ti('Insurance', small, 5).

%%%

% This procedure is to calculate investment required for adapting a recovery strategy.

%%%

calc_cost :-

lookup(hot_time_1, global, HT1),

lookup(hot_time_2, global, HT2),

Hot_time is (HT1 + HT2) / 2 ,

lookup(cold_time_1, global, CT1),

lookup(cold_time_2, global, CT2),

Cold_time is (CT1 + CT2) / 2 ,

lookup(hot_cost_1, global, HC1),

lookup(hot_cost_2, global, HC2),

is(X1, ln((HC1 + HC2) / 2)),

lookup(cold_cost_1, global, CC1),

lookup(cold_cost_2, global, CC2),

is(X2, ln((CC1 + CC2) / 2)),

M3 is $X1 - X2$,
 Time1 is (Cold_time - Hot_time),

N is (M3 / Time1),
 Y is ($X1 + (\text{Hot_time} * N)$),

is(AX, aln(Y)),

lookup(max_all_time, global, MAXAT),
 is(AX1, (AX * aln(- N * MAXAT))),
 new_slot(i, global, AX1).

start :-

repeat,
 getb(INPUT),
 (INPUT = 13 ; INPUT = 27).

%%%

user_message(TITLE, TEXT) :-

wdcreate(ud1,TITLE,30,50,500,200,[ws_caption,dlg_modalframe]),
 wccreate((ud1,1000),static,TEXT,5,5,490,180,[ws_child,ws_visible,ss_left]),
 wccreate((ud1,100),button,'OK',10,140,100,30,[ws_child,ws_visible,ws_tabstop,bs_ pushbutton]),

call_dialog(ud1, ok).

%%%

%%%

% This part is for recommending a recovery strategy

%%%

action rule_q

do ask size
 and ask degree
 and long_strategy becomes "
 and short_strategy becomes "
 and if degree is low
 then invoke ruleset low_degree
 and display_results
 else
 ask security_requirement
 and ask threat
 and ask recovery_time
 and ask modification


```

and if [ size is giant or size is large ]
  then do size1_questions
  else do size2_questions
  end if
end if .

```

```
%% This action for giant and large size organisations
```

```

action size1_questions
do if [ threat is regional or threat is building ]
  then ask hot_cooperative
  and ask cold_cooperative
  and ask personnel_support
  and ask work_area
  and invoke ruleset rec_strat
  and display_results
else
  ask hot_cooperative
  and invoke ruleset rec_strat
  and display_results
end if.

```

```
%% This action for medium and small size organisations
```

```

action size2_questions
do if [ threat is regional or threat is building ]
  then
  ask personnel_support
  and ask work_area
  and do med_small_rec
else
  do med_small_rec
end if .

```

```

action med_small_rec
do  ask service_bureau
  and ask reciprocal
  and ask time_broker
  and ask hardware_vendor
  and invoke ruleset rec_strat
  and display_results .

```

```
%%%%%%%%%
```

```
% Printing to the screen section
```

```
%%%%%%%%%
```

```
action display_results
```

```
do open_windows    % defn written in data.pl file to create window
and writer(result, 'For an organisation of size ') and writer(result, size)
and writer(result, ' and degree of dependency on computer ')
and writer(result, degree) and nlr(result)
and writer(result, ' ') and nlr(result)
and writer(result, 'The recommended short term strategy is ')
and writer(result, short_strategy) and nlr(result)
and writer(result, 'and the recommended long term strategy is ')
and writer(result, long_strategy) and nlr(result)
and writer(result, 'and the recommended location is ')
and writer(result, location )
and writer(result, ' ') and nlr(result) .
```

```
%%%%%%%%%
```

```
% This part is to calculate the Maximum Allowable Downtime
```

```
%%%%%%%%%
```

```
action ti
```

```
do ask organisation_type
and ask size
and ti( organisation_type, size, TIVALUE )
and ti becomes TIVALUE
and open_windows
and writer(info, 'For this type and size of organisation,
the proposed time interval is ') and writer(info, ti )
and writer(info, ' day(s)') and nlr(info)
and writer(info, ' ') and nlr(info)
and msgbox('Time Interval Recommended','Please take a note of the proposed time
interval',48,Code)
and ask daily_income
and ask total_cost
and total_ti becomes ti
and clear( result )
and calc_MAD
and writer(result, 'The Maximum Allowable downtime for your organisation is: ')
and writer(result, total_ti ) and writer(result, ' days(s)') and nlr(result) .
```

```
action get_next_cost
```

```
do repeat
```

```
writer(info, 'We have been unable to establish a MAD in the previous time interval')
and nlr(info)
```

```

and writer(info, 'Please provide the total downtime cost for the NEXT ')
and writer(info, ti ) and writer(info, ' days') and nlr(info)
and writer(info, ' ') and nlr(info)
and msgbox('MAD is not reached','Further calculation is required',48,Code)
and ask next_total_cost
and total_cost becomes (total_cost + next_total_cost)
and total_ti becomes total_ti + ti
until the total_cost >= daily_income * total_ti
end repeat
and mad becomes total_ti .

```

```

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% This part is to calculate the invesment required.
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

```

```

action inv
do ask hot_time_1
and ask hot_cost_1
and ask hot_time_2
and ask hot_cost_2
and ask cold_time_1
and ask cold_cost_1
and ask cold_time_2
and ask cold_cost_2
and ask max_all_time
and open_windows
and clear( result )
and do calc_cost % this procedure is placed in data.pl
and writer(result, 'The investment required for your organisation is: ')
and fnwriter(result, i ) and nlr(result) .

```

Relations

```

relation calc_MAD
if total_cost >= daily_income * ti
and mad becomes ti .

```

```

relation calc_MAD
if get_next_cost .

```

```

relation threat_check1
if threat is regional

```

```

relation threat_check
if threat is building

```

Questions

question size

'What is the size of the organisation?';

choose one of sizes ;

because Size will help in deciding what MAD and recovery strategy are appropriate to the organisation .

question degree

'What is the degree of computer dependancy?';

choose one of degrees ;

because Degree will help in deciding what MAD and recovery strategy are appropriate to the organisation .

question recovery_time

'How fast does your organisation need to be recovered?';

choose one of times ;

because The answer can be obtained from the first transaction output (MAD) .

question security_requirement

'What level of security does your organisation required?';

choose one of levels ;

because If you have very high security requirements eg national defence, a specific recovery strategy has to be adapted .

question modification

'Does your organisation need special-tailored hardware or software arrangement?';

choose one of mods ;

because Some organisations install unique hardware or software such as air control system .

question threat

'What type of threat are you exposed to?';

choose one of threats ;

because This important to decide on the location of the alternative site .

question personnel_support

'Does the organisation need outside personnel support?';

choose one of yes_no ;

because Some organisations need outside help for example: staff in organisations exposed to regional threats might be busy with their personnel affairs and the company needs outside help to run the IT centre .

question hot_cooperative

'Is there a possibility that your organisation can establish a cooperative HOT site with other near by organisations?';
choose one of yes_no .

question cold_cooperative

'Is there a possibility that your organisation can establish a cooperative COLD site with other near by organisations?';
choose one of yes_no .

question reciprocal

'Is there a possibility that your organisation make a mutual agreement with another organisation that have similar platform?';
choose one of yes_no .

question service_bureau

'Is there a service bureau facility available that provides recovery arrangements?';
choose one of yes_no .

question time_broker

'Is there a time broker available who can provide recovery arrangements?';
choose one of yes_no .

question hardware_vendor

'Does your hardware vendor provide recovery arrangements?';
choose one of yes_no .

question work_area

'Does your organisation need large working area (more than 10 employees) in the alternative site?';
choose one of yes_no .

%%
% questions for time intervals calc.
%%

question organisation_type

'What type of business does your organisation perform?';
choose one of organisations .

question total_cost

'What is the total downtime cost of the denial of computer facilities for all critical resources for the proposed time interval?';
input number ;
because The value of this input should be collected in the Business Impact Asseement Phase.

question daily_income

'What is the average daily income of the organisation?';
input number .

question next_total_cost

'Please provide the total downtime cost of the denial of computer facilities for all critical systems for the next time interval?';
input number .

%%%%%%%%%%
% questions for investment calculation
%%%%%%%%%%

question hot_time_1

'Provide an estimate of minimum response time for a commercial HOT site';
input number ;
because Hot site vendors usually have minimum response time that they can not provide services before it.

question hot_cost_1

'Provide an estimate of the annual cost for the previous minimum response time when subscribing to a commercial HOT site';
input number ;
because This can be collected from the disaster recovery market .

question hot_time_2

'Provide an estimate of maximum response time for a commercial HOT site';
input number ;
because Just select any time that you think is appropriate to the hot site provider that you know .

question hot_cost_2

'Provide an estimate of the annual cost for the previous maximum response time when subscribing to a commercial HOT site';
input number .

question cold_time_1

'Provide an estimate of minimum response time for a commercial COLD site';
input number .

question cold_cost_1

'Provide an estimate of the annual cost for the previous minimum response time when subscribing to a commercial COLD site';
input number .

question cold_time_2

'Provide an estimate of maximum response time for a commercial COLD site';

input number ;

because Cold site vendors usually have minimum response time that they can not provide services before it .

question cold_cost_2

'Provide an estimate of the annual cost for the pervious maximum response time when subscribing to a commercial COLD site';

input number .

question max_all_time

'What is the maximum allowable downtime for your organisation which was calculated from the previous transaction';

input number ;

because You can see it in the Results window above this question .

Groups

group sizes

giant, large, medium, small .

group degrees

high, medium, low, 'Do not know !!' .

group times

immediate, '1 to 2 days', 'more than 3 days', 'Do not know !!' .

group levels

'very high', high, medium, 'Do not know !!' .

group mods

yes, no, 'Do not know !!' .

group threats

regional, building, floor, 'disk failure only', 'Do not know !!' .

group organisations

'Financial', 'Manufacturing', 'Government', 'Computer Services',
'Retailing', 'Telecommunication', 'Education', 'Health Care', 'Insurance', 'Others' .

group yes_no

yes, no, 'Do not know !!' .

Rules

ruleset rec_strat

ruleset rec_strat

contains security_1, security_2, security_3, giant_high_reg_1, giant_high_reg_2,
giant_high_build_1, giant_high_build_1, giant_high_floor_disk_1, giant_high_reg_3,
giant_high_reg_4, large_high_build_1, md_med_reg_1, md_med_reg_2,
d_med_build_1, md_med_build_2, md_med_room, md_med_floor, small_high_reg_1,
small_high_reg_2, small_high_reg_3, small_high_reg_4, small_high_reg_5,
small_high_reg_6, small_high_reg_7, small_high_reg_8, small_high_build_1,
small_high_build_2, small_high_build_3, small_high_build_4, small_high_build_5,
small_high_build_6, small_high_build_7, small_high_build_8, small_high_floor_1,
small_high_floor_2, small_high_floor_3, small_high_floor_4, small_high_floor_5,
small_high_floor_6, small_high_floor_7, small_med ;
initiate by doing location becomes " ;
terminate when location is not " .

/*****

In rules 1 - 3: duplicate site is recommended because the company has very high security requirements.

*****/

rule security_1

if security_requirement is 'very high'
and degree is high
and threat is regional
then short_strategy becomes 'Duplicate Site'
and long_strategy becomes 'The Duplicate Site should be used as a long term strategy'
and location becomes remote
score 100 .

/*****

If the threat covers the whole building, the site should be located several miles away because the buildings or even the streets around the affected area may be evacuated

*****/

rule security_2

if security_requirement is 'very high'
and degree is high
and threat is building
then short_strategy becomes 'Duplicate Site'
and long_strategy becomes 'The Duplicate Site should be used as a long term strategy'
and location becomes 'within the city area'

score 100 .

rule security_3

if security_requirement is 'very high'

and degree is high

and [threat is 'disk failure only'

or threat is floor]

then short_strategy becomes 'Duplicate Site'

and long_strategy becomes 'The Duplicate Site should be used as

a long term strategy'

and location becomes 'adjacent to the company or within the city area'

score 100 .

/*****

The following rules deal with organisations that have the following characteristics:

- giant size,
- high degree of dependency on computer, and
- need immediate recovery
- need special HW/SW installation arrangements

*****/

rule giant_high_reg_1

if size is giant

and degree is high

and recovery_time is immediate

and modification is yes

and threat is regional

and [personnel_support is yes

or cold_cooperative is no]

and hot_cooperative is yes

then short_strategy becomes 'Duplicate site or cooperative hot site

with realtime recovery'

and long_strategy becomes 'commercial cold site'

and location becomes remote .

rule giant_high_reg_2

if size is giant

and degree is high

and recovery_time is immediate

and modification is yes

and threat is regional

and personnel_support is no

and hot_cooperative is yes

and cold_cooperative is yes

then short_strategy becomes 'Duplicate site or cooperative hot site

with realtime recovery'
 and long_strategy becomes 'cooperative cold site'
 and location becomes remote .

rule giant_high_build_1
 if size is giant
 and degree is high
 and recovery_time is immediate
 and modification is yes
 and threat is building
 and [personnel_support is yes
 or cold_cooperative is no]
 and hot_cooperative is yes
 then short_strategy becomes 'Duplicate site or cooperative hot site
 with realtime recovery'
 and long_strategy becomes 'commercial cold site'
 and location becomes 'within the city area' .

rule giant_high_build_2
 if size is giant
 and degree is high
 and recovery_time is immediate
 and modification is yes
 and threat is building
 and personnel_support is no
 and hot_cooperative is yes
 and cold_cooperative is yes
 then short_strategy becomes 'Duplicate site or cooperative hot site
 with realtime recovery'
 and long_strategy becomes 'cooperative cold site'
 and location becomes 'within the city area' .

rule giant_high_floor_disk_1
 if size is giant
 and degree is high
 and recovery_time is immediate
 and modification is yes
 and [threat is 'disk failure only'
 or threat is floor]
 and hot_cooperative is yes
 then short_strategy becomes 'Duplicate site or cooperative hot site
 with realtime recovery'
 and long_strategy becomes 'No need for long term strategy because the
 original site should be rebuild within short period of time'
 and location becomes 'adjacent to the organisation' .

/******

The following rules deal with organisations that have the following characteristics:

- giant size,
- high degree of dependency on computer, and
- need recovery time from 1 to 2 days
- need special HW/SW installation arrangements

*****/

rule giant_high_reg_3

if size is giant

and degree is high

and recovery_time is '1 to 2 days'

and threat is regional

and [personnel_support is yes

 or hot_cooperative is no]

then short_strategy becomes 'commercial hot site'

and long_strategy becomes 'commercial cold site'

and location becomes remote .

rule giant_high_reg_4

if size is giant

and degree is high

and recovery_time is '1 to 2 days'

and threat is regional

and hot_cooperative is yes

and cold_cooperative is yes

then short_strategy becomes 'cooperative hot site'

and long_strategy becomes 'cooperative cold site'

and location becomes remote .

rule large_high_build_1

if size is large

and degree is high

and recovery_time is '1 to 2 days'

and threat is building

and hot_cooperative is no

and cold_cooperative is yes

then short_strategy becomes 'commercial hot site'

and long_strategy becomes 'cooperative cold site'

and location becomes 'within the city area' .

/******

The following rules deal with organisations that have the following characteristics:

- medium size,

- medium degree of dependency on computer

*****/

```
rule md_med_reg_1
if size is medium
and degree is medium
and threat is regional
and hardware_vendor is yes
then short_strategy becomes 'hardware vendor'
and long_strategy becomes 'portable site'
and location becomes remote .
```

```
rule md_med_reg_2
if size is medium
and degree is medium
and threat is regional
and hardware_vendor is no
then short_strategy becomes 'portable site'
and long_strategy becomes 'portable site'
and location becomes remote .
```

```
rule md_med_build_1
if size is medium
and degree is medium
and threat is regional
and hardware_vendor is yes
then short_strategy becomes 'hardware vendor or reduction of service'
and long_strategy becomes 'portable site'
and location becomes 'within the city area' .
```

```
rule md_med_build_2
if size is medium
and degree is medium
and threat is regional
and hardware_vendor is no
then short_strategy becomes 'portable site or reduction of service'
and long_strategy becomes 'portable site'
and location becomes 'within the city area' .
```

```
rule md_med_floor
if size is medium
and degree is medium
and threat is floor
then short_strategy becomes 'portable site or reduction of service'
and long_strategy becomes 'portable site'
```

and location becomes 'adjacent to the organisation' .

```
rule md_med_room
if size is medium
and degree is medium
and threat is 'disk failure only'
then short_strategy becomes 'reduction of service'
and long_strategy becomes 'reduction of service until the disk failure is fixed'
and location becomes 'no need' .
```

/******

The following rules deal with organisations that have the following characteristics:

- small size,
- high degree of dependency on computer, and
- need immediate or recovery time from 1 to 2 days

*****/

```
rule small_high_reg_1
if size is small
and degree is high
and [ recovery_time is immediate or recovery_time is '1 to 2 days' ]
and threat is regional
and [ work_area is yes or modification is yes ]
then short_strategy becomes 'commercial hot site.'
and long_strategy becomes 'commercial cold site'
and location becomes remote
score 30 .
```

```
rule small_high_reg_2
if size is small
and degree is high
and [ recovery_time is immediate or recovery_time is '1 to 2 days' ]
and threat is regional
and modification is no
and work_area is no
and service_bureau is yes
then short_strategy becomes 'service bureau.'
and long_strategy becomes 'commercial cold site'
and location becomes remote .
```

```
rule small_high_reg_3
if size is small
and degree is high
and [ recovery_time is immediate or recovery_time is '1 to 2 days' ]
and threat is regional
```

and personnel_support is yes
and work_area is no
and service_bureau is no
then short_strategy becomes 'mobil hot site.'
and long_strategy becomes 'commercial cold site'
and location becomes remote
score 20 .

rule small_high_reg_4
if size is small
and degree is high
and [recovery_time is immediate or recovery_time is '1 to 2 days']
and threat is regional
and personnel_support is no
and work_area is no
and modification is no
and service_bureau is no
and reciprocal is yes
then short_strategy becomes 'reciprocal agreement.'
and long_strategy becomes 'commercial cold site'
and location becomes remote .

rule small_high_reg_5
if size is small
and degree is high
and [recovery_time is immediate or recovery_time is '1 to 2 days']
and threat is regional
and personnel_support is no
and work_area is no
and modification is no
and service_bureau is no
and reciprocal is no
and time_broker is yes
then short_strategy becomes 'time broker.'
and long_strategy becomes 'commercial cold site'
and location becomes remote .

rule small_high_reg_6
if size is small
and degree is high
and [recovery_time is immediate or recovery_time is '1 to 2 days']
and threat is regional
and reciprocal is no
and time_broker is no
and service_bureau is no

then short_strategy becomes 'commercial hot site.'
and long_strategy becomes 'commercial cold site'
and location becomes remote
score 5 .

rule small_high_reg_7
if size is small
and degree is high
and recovery_time is 'more than 3 days'
and threat is regional
and [work_area is yes or modification is yes
or hardware_vendor is no]
then short_strategy becomes 'warm site.'
and long_strategy becomes 'warm site or commercial cold site'
and location becomes remote .

rule small_high_reg_8
if size is small
and degree is high
and recovery_time is 'more than 3 days'
and threat is regional
and modification is no
and work_area is no
and hardware_vendor is yes
then short_strategy becomes 'hardware vendor'
and long_strategy becomes 'commercial cold site'
and location becomes remote .

rule small_high_build_1
if size is small
and degree is high
and [recovery_time is immediate or recovery_time is '1 to 2 days']
and threat is building
and [work_area is yes or modification is yes]
then short_strategy becomes 'commercial hot site.'
and long_strategy becomes 'commercial cold site'
and location becomes 'within the city area'
score 30 .

rule small_high_build_2
if size is small
and degree is high
and [recovery_time is immediate or recovery_time is '1 to 2 days']
and threat is building
and modification is no

and work_area is no
and service_bureau is yes
then short_strategy becomes 'service bureau.'
and long_strategy becomes 'commercial cold site'
and location becomes 'within the city area' .

rule small_high_build_3
if size is small
and degree is high
and [recovery_time is immediate or recovery_time is '1 to 2 days']
and threat is building
and personnel_support is yes
and work_area is no
and service_bureau is no
then short_strategy becomes 'mobil hot site.'
and long_strategy becomes 'commercial cold site'
and location becomes 'within the city area'
score 20 .

rule small_high_build_4
if size is small
and degree is high
and [recovery_time is immediate or recovery_time is '1 to 2 days']
and threat is building
and personnel_support is no
and modification is no
and work_area is no
and service_bureau is no
and reciprocal is yes
then short_strategy becomes 'reciprocal agreement.'
and long_strategy becomes 'commercial cold site'
and location becomes 'within the city area' .

rule small_high_build_5
if size is small
and degree is high
and [recovery_time is immediate or recovery_time is '1 to 2 days']
and threat is building
and personnel_support is no
and modification is no
and work_area is no
and service_bureau is no
and reciprocal is no
and time_broker is yes
then short_strategy becomes 'time broker.'

and long_strategy becomes 'commercial cold site'
and location becomes 'within the city area' .

rule small_high_build_6
if size is small
and degree is high
and [recovery_time is immediate or recovery_time is '1 to 2 days']
and threat is building
and service_bureau is no
and reciprocal is no
and time_broker is no
then short_strategy becomes 'commercial hot site.'
and long_strategy becomes 'commercial cold site'
and location becomes 'within the city area'
score 5 .

rule small_high_build_7
if size is small
and degree is high
and recovery_time is 'more than 3 days'
and threat is building
and [work_area is yes or modification is yes
or hardware_vendor is no]
then short_strategy becomes 'warm site.'
and long_strategy becomes 'warm site or commercial cold site'
and location becomes 'within the city area' .

rule small_high_build_8
if size is small
and degree is high
and recovery_time is 'more than 3 days'
and threat is building
and modification is no
and work_area is no
and hardware_vendor is yes
then short_strategy becomes 'hardware vendor'
and long_strategy becomes 'commercial cold site'
and location becomes 'within the city area' .

rule small_high_floor_1
if size is small
and degree is high
and [recovery_time is immediate or recovery_time is '1 to 2 days']
and threat is floor
and [modification is yes or personnel_support is yes]

then short_strategy becomes 'mobil hot site'
and long_strategy becomes 'portable site'
and location becomes 'adjacent to the organisation'
score 30 .

rule small_high_floor_2
if size is small
and degree is high
and [recovery_time is immediate or recovery_time is '1 to 2 days']
and threat is floor
and modification is no
and service_bureau is yes
then short_strategy becomes 'service bureau.'
and long_strategy becomes 'portable site'
and location becomes 'adjacent to the organisation' .

rule small_high_floor_3
if size is small
and degree is high
and [recovery_time is immediate or recovery_time is '1 to 2 days']
and threat is floor
and personnel_support is no
and modification is no
and service_bureau is no
and reciprocal is yes
then short_strategy becomes 'reciprocal agreement.'
and long_strategy becomes 'commercial cold site'
and location becomes 'adjacent to the organisation' .

rule small_high_floor_4
if size is small
and degree is high
and [recovery_time is immediate or recovery_time is '1 to 2 days']
and threat is floor
and personnel_support is no
and modification is no
and service_bureau is no
and reciprocal is no
and time_broker is yes
then short_strategy becomes 'time broker.'
and long_strategy becomes 'commercial cold site'
and location becomes 'adjacent to the organisation' .

rule small_high_floor_5
if size is small

```

and degree is high
and [ recovery_time is immediate or recovery_time is '1 to 2 days' ]
and threat is floor
and service_bureau is no
and reciprocal is no
and time_broker is no
then short_strategy becomes 'mobil hot site.'
and long_strategy becomes 'portable site'
and location becomes 'adjacent to the organisation'
score 5 .

```

```

rule small_high_floor_6
if size is small
and degree is high
and recovery_time is 'more than 3 days'
and threat is floor
and [ modification is yes
or hardware_vendor is no ]
then short_strategy becomes 'warm site.'
and long_strategy becomes 'portable site'
and location becomes 'adjacent to the organisation' .

```

```

rule small_high_floor_7
if size is small
and degree is high
and recovery_time is 'more than 3 days'
and threat is floor
and modification is no
and hardware_vendor is yes
then short_strategy becomes 'hardware vendor'
and long_strategy becomes 'portable site'
and location becomes 'adjacent to the organisation' .

```

```

rule small_med
if size is small
and degree is medium
then short_strategy becomes 'withdraw of services or manual procedure.'
and long_strategy becomes '---'
and location becomes '----' .

```

```

ruleset low_degree
contains giant_low, medium_low, small_low ;
initiate by doing location becomes " ;
terminate when location is not " .

```


ntr(WIN) :-

```
wedttxt(( WIN,1), `~M~J` ).
```

clear(WINDOW) :-

```
wtext((WINDOW,1), ``).
```

```
%%%%%%%%%%
%% This is the start of the system
%%%%%%%%%%
```

ri0 :-

```
wbload(ult, 'ulti2.bmp'),
wcreate(ri0,`Knowledge-based Disaster Recovery Strategy
System`,10,10,600,460,[dlg_ownedbyprolog,ws_sysmenu,ws_caption]),

wcreate((ri0,900),grafix,``,90,10,400,270,[ws_child,ws_visible,ws_border]),
```

```
wcreate((ri0,1000),static,`Please wait while the files are loading.`
,90,295,400,50,[ws_child,ws_visible,ss_left]),
```

```
wcreate((ri0,102),button,``,
30,340,490,55,[ws_child,ws_visible,ws_tabstop,bs_groupbox]),
```

```
window_handler( ri0, ri0_handler ),
show_dialog( ri0 ),
wflag(1), wait(0),
wgfx( (ri0,900), [bits(0,0,400,450,55,8,ult)] , 0,0,600,600),
```

```
%reconsult_rules( control ),
%reconsult_rules( quest ),
%reconsult_rules( rules ),
wtext( (ri0, 1000),`WELECOME TO THE KNOWLEDGE-BASED
DISASTER RECOVERY SYSTEM
```

```
Please press NEXT to continue or EXIT to leave the system.` ),
```

```
wcreate((ri0,101),button,`EXIT`, 40,355,
90,36,[ws_child,ws_visible,ws_tabstop,bs_pushbutton]),
wcreate((ri0,100),button,`NEXT`,415,355,
90,36,[ws_child,ws_visible,ws_tabstop,bs_pushbutton]), !.
```

```
ri0_handler( (ri0, 100), msg_button, _, _ ) :-
wclose( ri0 ),
r1 .
```

```

ri0_handler( ri0, 101), msg_button, _ _ ) :-
    wclose( ri0 ), abort . %halt .

ri1 :-
    wdcreate(ri1,`Knowledge-based Disaster Recovery Strategy
System`,10,10,600,460,[dlg_ownedbyprolog,ws_sysmenu,ws_caption]),

wcreate((ri1,1000),static,`
,90,40,400,300,[ws_child,ws_visible,ss_left]),

wcreate((ri1,102),button,`
30,340,490,55,[ws_child,ws_visible,ws_tabstop,bs_groupbox]),

    window_handler( ri1, ri1_handler ),
    show_dialog( ri1 ),
    wflag(1), wait(0),
    wtext( (ri1, 1000),`Welecom to the knowledge based system for disaster
recovery strategy selection.
The system contains the following three transactions :

1. Calculation of Maximum Allowable Downtime (MAD)
2. Calculation of the required investment for fitting a recovery strategy, and
3. Recommendation of disaster recovery strategy(s).

The system will ask you several questions about your organisation. Some of the
questions are self explained.
For unclear questions, an explanation of the question or why it is been asked can be
found by clicking in the Explain button next to the question.

At the end of the first transaction, please take a note of the calculated Maximum
Allowable Downtime because you need it as an input in one of the questions in the
second transaction.

Please press NEXT to start first transaction or EXIT to leave the system.` ),

wcreate((ri1,101),button,`EXIT`, 40,355,
90,36,[ws_child,ws_visible,ws_tabstop,bs_pushbutton]),
wcreate((ri1,100),button,`NEXT`,415,355,
90,36,[ws_child,ws_visible,ws_tabstop,bs_pushbutton]), ! .

ri1_handler( (ri1, 100), msg_button, _ _ ) :-
    wclose( ri1 ),
    flex_starter .

```

```

ril_handler( ril, 101), msg_button, _ _ ) :-
    wclose( ril ), abort . % halt .

```

```

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%% continuation screen procedures
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

```

```

ask_continue :-

```

```

    WIN = cont,

```

```

    wdcreate(WIN, `Continue`, 100, 230, 400, 200, [dlg_ownedbyprolog, ws_sysmenu, ws_
caption]),

```

```

    wcreate((WIN, 1000), static,

```

```

    `Do you wish to continue to the next transaction ?

```

```

Please press Yes to continue or No to leave the
system.` , 90, 30, 350, 60, [ws_child, ws_visible, ss_left]),

```

```

    wcreate((WIN, 102), button, ``,
    5, 125, 390, 50, [ws_child, ws_visible, ws_tabstop, bs_groupbox]),

```

```

    wcreate((WIN, 100), button, `Yes`, 300, 135,
    90, 36, [ws_child, ws_visible, ws_tabstop, bs_pushbutton]),

```

```

    wcreate((WIN, 101), button, `No`, 10, 135,
    90, 36, [ws_child, ws_visible, ws_tabstop, bs_pushbutton]),

```

```

    window_handler( WIN, cont_handler ),

```

```

    show_dialog( WIN ),

```

```

    repeat,

```

```

        wflag(1), wait(0),

```

```

        retract( continue ),

```

```

    wclose( cont ).

```

```

cont_handler( (cont, 100), msg_button, _ _ ) :-
    assert( continue ).

```

```

cont_handler( (cont, 101), msg_button, _ _ ) :-
    wclose( cont ), abort . % halt .

```

```

ask_continue1 :-

```

```

    WIN = cont1,

```

```

    wdcreate(WIN, `Second Transaction: Investment
Calculation`, 70, 120, 500, 350, [dlg_ownedbyprolog, ws_sysmenu, ws_
caption]),

```

```
wcreate((WIN,1000),static,
```

```
`The cost of alternative back up sites vary inversely with respect to the response time
```

The user should investigate the disaster recovery market and provide :

- 1 Two response time estimates (minimum and maximum)
2. Annual cost estimates for each of the previous response times for two types of commercial strategies: hot and cold sites
3. The Maximum Allowable Downtime (generated from the last transaction)

```
Please press Yes to continue or No to leave the
system.` ,70,40,400,220,[ws_child,ws_visible,ss_left]],
```

```
wcreate((WIN,102),button,``,
15,265,435,50,[ws_child,ws_visible,ws_tabstop,bs_groupbox]],
```

```
wcreate((WIN,100),button,`Yes`, 350,275,
90,36,[ws_child,ws_visible,ws_tabstop,bs_pushbutton]),
wcreate((WIN,101),button,`No`, 25,275,
90,36,[ws_child,ws_visible,ws_tabstop,bs_pushbutton]),
```

```
    window_handler( WIN, cont_handler ),
    show_dialog( WIN ),
    repeat,
        wflag(1), wait(0),
        retract( continue ),
    wclose( cont1).
```

```
cont_handler( (cont1,100), msg_button, _ _ ) :-
    assert( continue ).
```

```
cont_handler( (cont1,101), msg_button, _ _ ) :-
    wclose( cont1), abort . % halt .
```

```
ask_continue2 :-
```

```
    WIN = cont2,
    wcreate(WIN,`Third Transaction: Recovery Strategy
Recommendation`,70,120,500,350,[dlg_ownedbyprolog,ws_sysmenu,ws_caption]),
```

```
wcreate((WIN,1000),static,
```

```
`The recovery strategy is the ability to process data while a full recovery of the original
site is underway.
```

The recommendation is based on :

1. Characteristics of the organisation
2. Threat Type
3. Organisation's requirements
4. The availability of some recovery strategies.

Please press Yes to continue or No to leave the system.` ,70,30,400,200,[ws_child,ws_visible,ss_left]),

```
wcreate((WIN,102),button,`,
15,265,435,50,[ws_child,ws_visible,ws_tabstop,bs_groupbox]),
```

```
wcreate((WIN,100),button,`Yes`, 350,275,
90,36,[ws_child,ws_visible,ws_tabstop,bs_pushbutton]),
wcreate((WIN,101),button,`No`, 25,275,
90,36,[ws_child,ws_visible,ws_tabstop,bs_pushbutton]),
```

```
    window_handler( WIN, cont_handler ),
    show_dialog( WIN ),
    repeat,
        wflag(1), wait(0),
        retract( continue ),
    wclose( cont2).
```

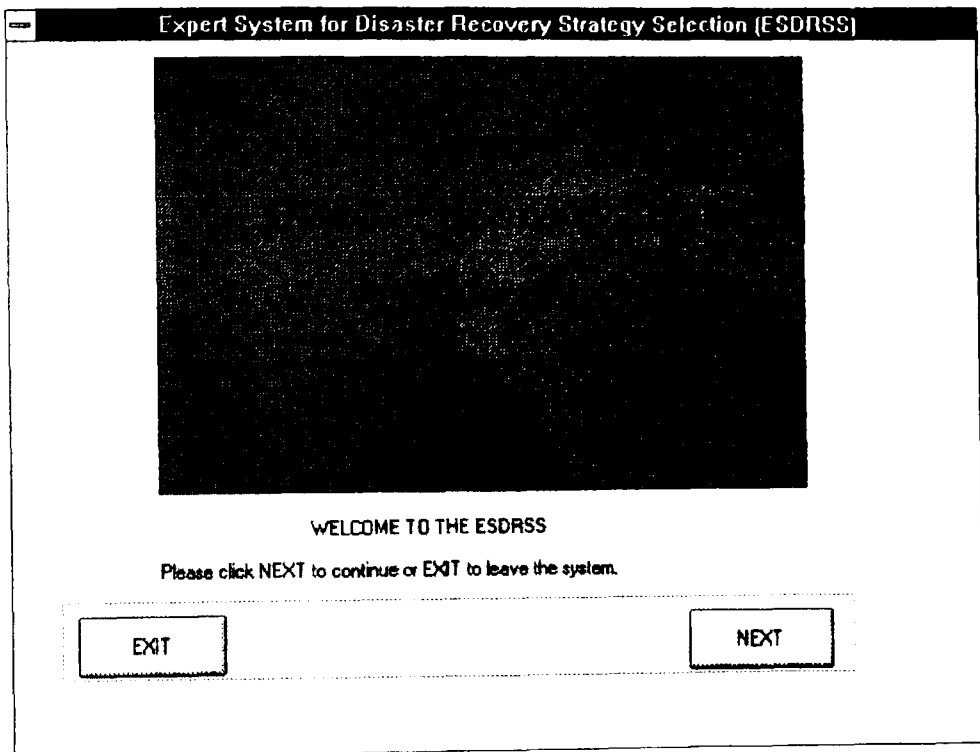
```
cont_handler( (cont2,100), msg_button, _, _ ) :-
    assert( continue ).
```

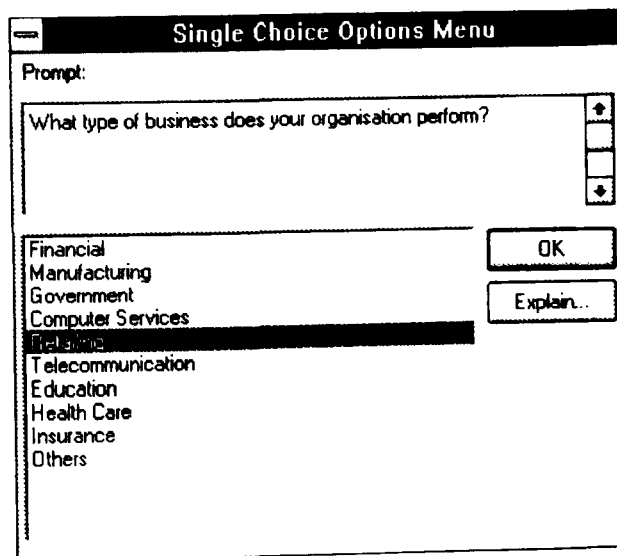
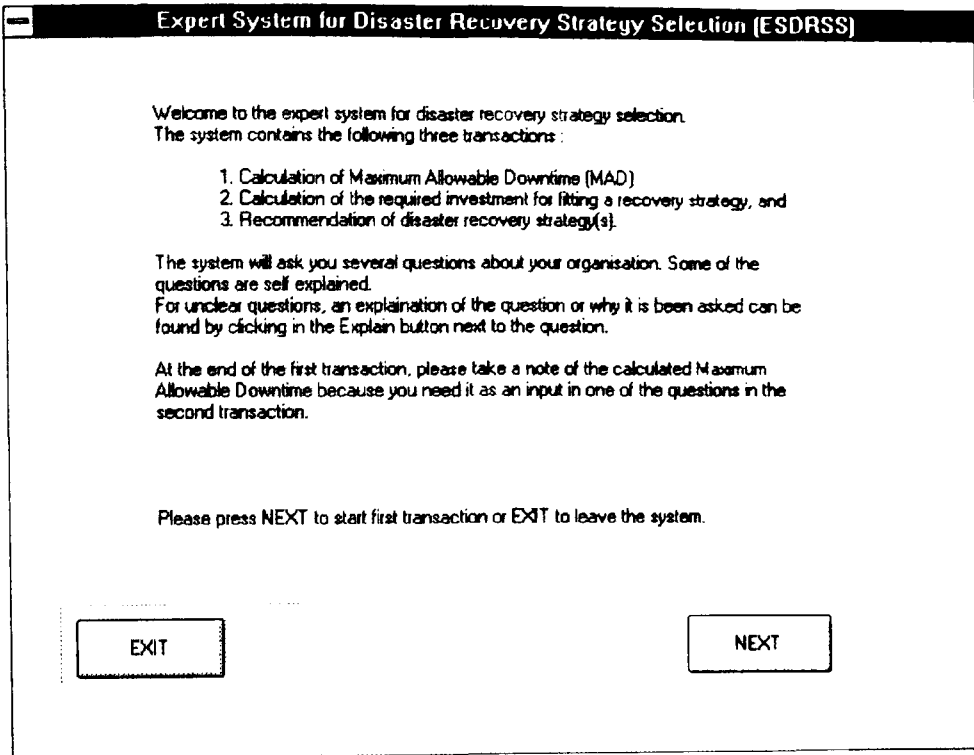
```
cont_handler( (cont2,101), msg_button, _, _ ) :-
    wclose( cont2), abort . % halt .
```

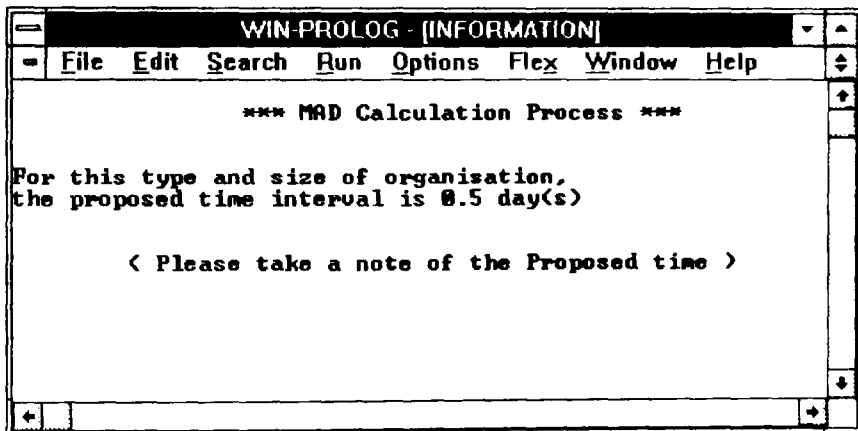
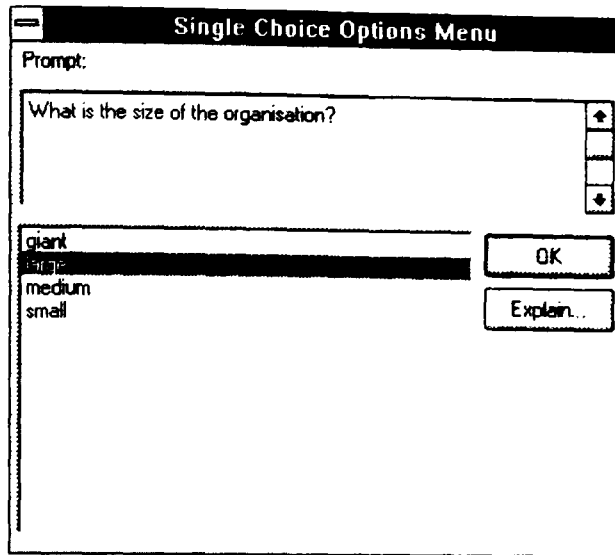
```
%:- ri0 .
```

Appendix B

Examples of Inputs and Outputs Screens







Typed Input

Prompt:

What is the average daily income of the organisation?

250000

OK

Explain...

Typed Input

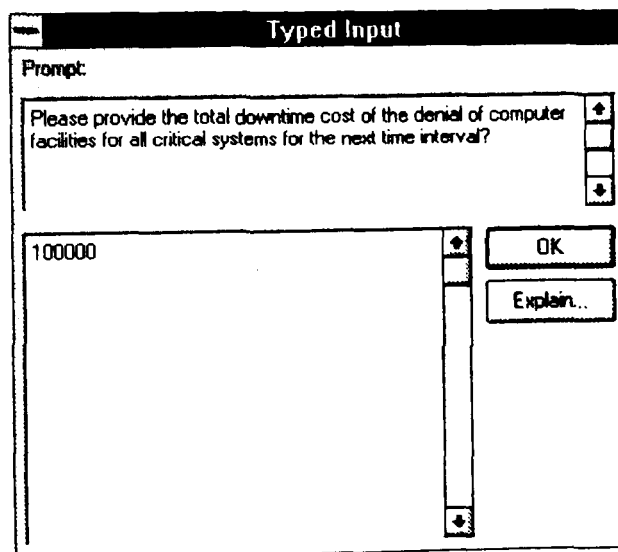
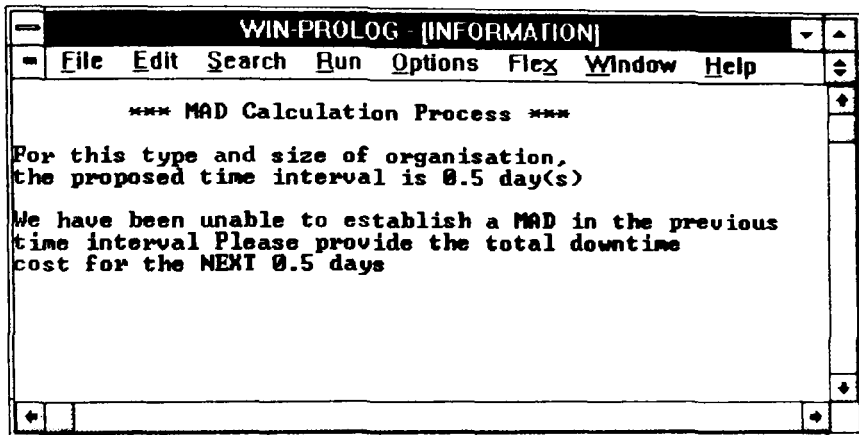
Prompt:

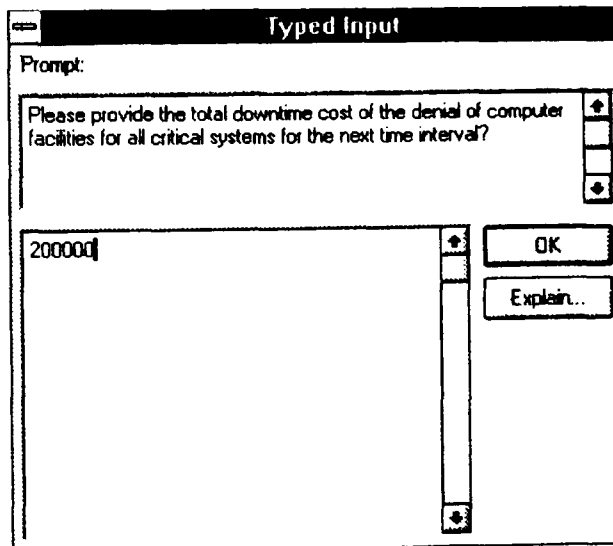
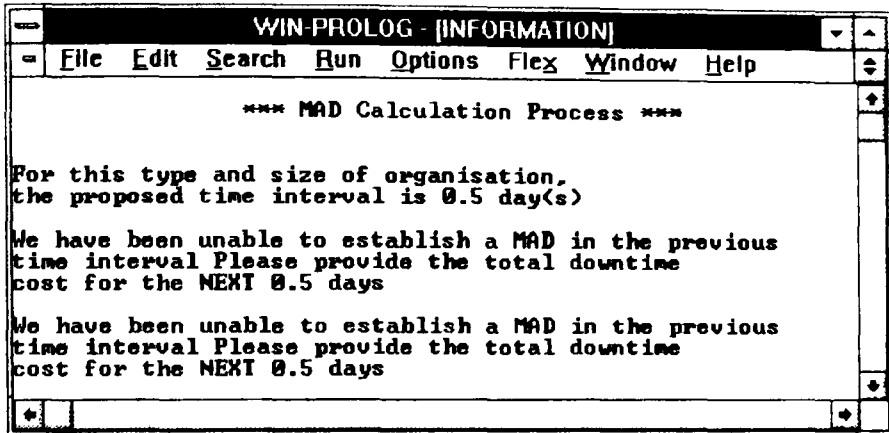
What is the total downtime cost of the denial of computer facilities for all critical resources for the proposed time interval?

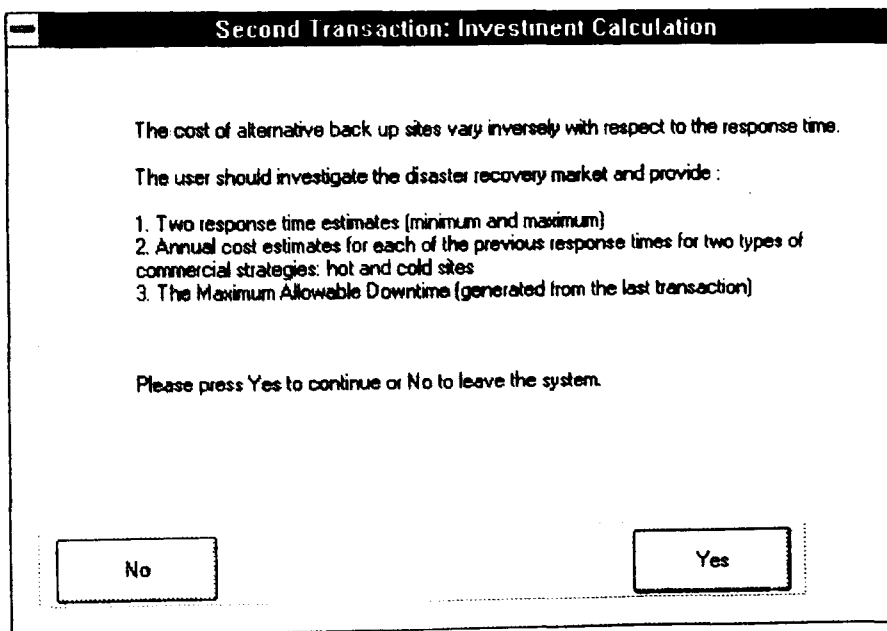
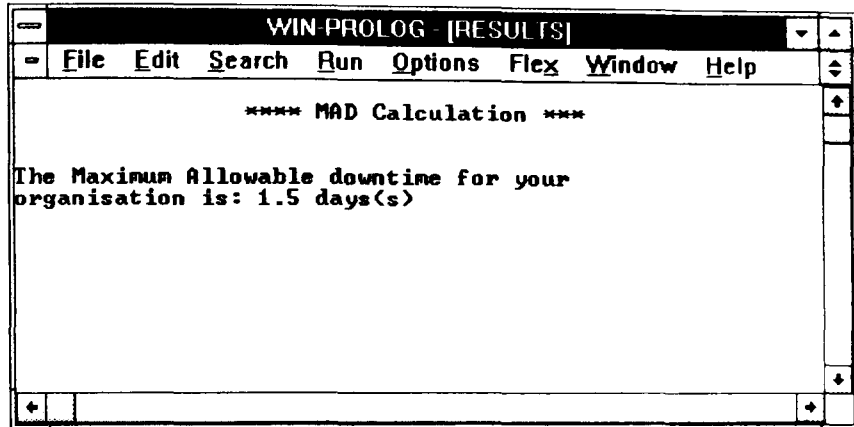
80000

OK

Explain...







Typed Input

Prompt:

Provide an estimate of minimum response time for a commercial HOT site

0.5

OK

Explain...

Typed Input

Prompt:

Provide an estimate of the annual cost for the pervious minimum response time when subscribing to a commercial HOT site

200000

OK

Explain...

Typed Input

Prompt:

Provide an estimate of maximum response time for a commercial HOT site

3

OK

Explain...

Typed Input

Prompt:

Provide an estimate of the annual cost for the pervious maximum response time when subscribing to a commercial HOT site

90000

OK

Explain...

Typed Input

Prompt:

Provide an estimate of minimum response time for a commercial COLD site

7

OK

Explain...

Typed Input

Prompt:

Provide an estimate of the annual cost for the pervious minimum response time when subscribing to a commercial COLD site

10000

OK

Explain...

Typed Input

Prompt:

Provide an estimate of maximum response time for a commercial COLD site

14

OK

Explain...

Typed Input

Prompt:

Provide an estimate of the annual cost for the pervious maximum response time when subscribing to a commercial COLD site

2000

OK

Explain...

Typed Input

Prompt:

What is the maximum allowable downtime for your organisation which was calculated from the previous transaction

1.5

OK

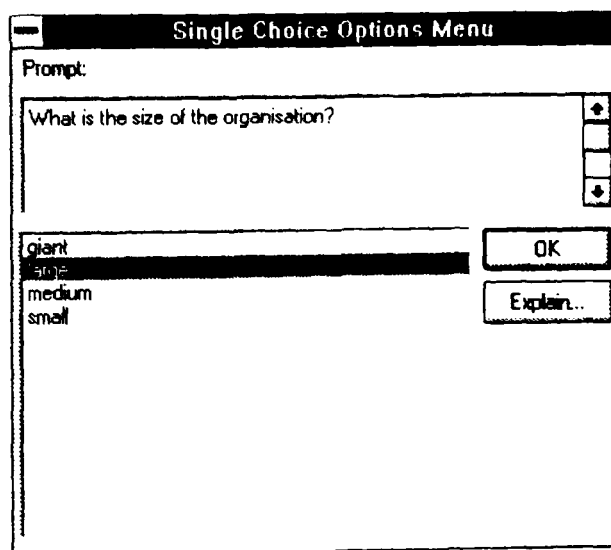
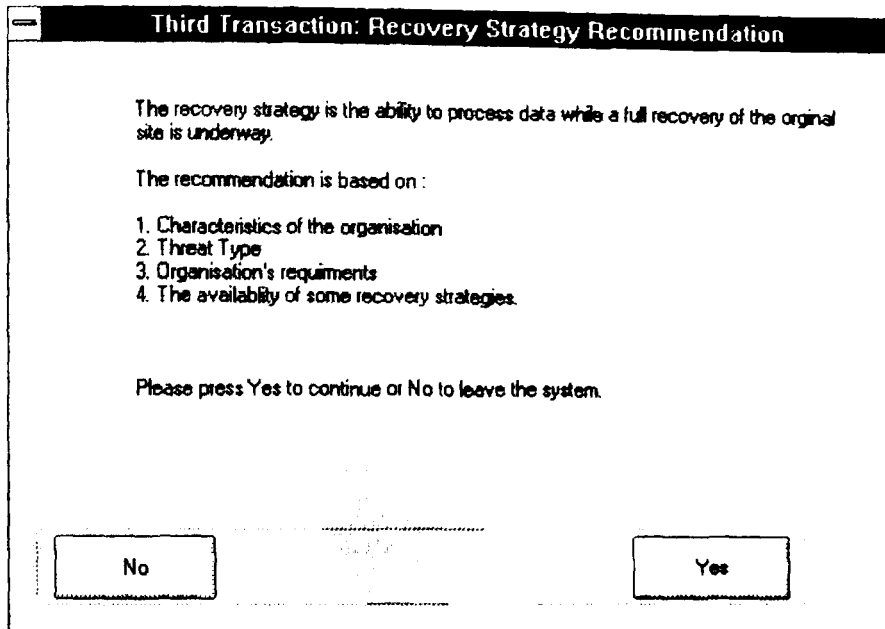
Explain...

WIN-PROLOG [RESULTS]

File Edit Search Run Options Flex Window Help

*** Investment Computation to adopt a DRP ***

The investment required for your organisation is: 153183.887



Single Choice Options Menu

Prompt:

What is the degree of computer dependency?

low
medium
low
Do not know !!

OK

Explain...

Single Choice Options Menu

Prompt:

What level of security does your organisation required?

very high
medium
Do not know !!

OK

Explain...

Single Choice Options Menu

Prompt:

What type of threat are you exposed to?

regional
building
floor
disk failure only
Do not know !!

OK

Explain...

Single Choice Options Menu

Prompt:

How fast does your organisation need to be recovered?

immediate
1 to 3 days
more than 3 days
Do not know !!

OK

Explain...

Single Choice Options Menu

Prompt:

Does your organisation need special-tailored hardware or software arrangement?

yes

no

Do not know !!

OK

Explain...

Single Choice Options Menu

Prompt:

Is there a possibility that your organisation can establish a cooperative HOT site with other near by organisations?

yes

no

Do not know !!

OK

Explain...

Single Choice Options Menu

Prompt:

Is there a possibility that your organisation can establish a cooperative COLD site with other near by organisations?

yes

no

Do not know !!

OK

Explain...

Single Choice Options Menu

Prompt:

Does the organisation need outside personnel support?

yes

Do not know !!

OK

Explain...

Single Choice Options Menu

Prompt:

Does your organisation need large working area (more than 10 employees) in the alternative site?

no
Do not know !!

OK
Explain...

WIN PROLOG - [RESULTS]

File Edit Search Run Options Flex Window Help

***** Recovery Strategy Recommendations *****

For an organisation of size large and degree of dependency on computer high

The recommended short term strategy is commercial hot site and the recommended long term strategy is cooperative cold site and the recommended location is within the city area.