

Disaster Recovery Planning Guide

Purpose of the Guide.....	3
DRG's Approach	3
Step 1. Develop the Planning Policy Statement.....	4
Scope.....	4
Resources	5
Roles and Responsibilities	5
Step 2. Conduct a Business Impact Analysis	5
The Outside-In Analysis	5
What are the critical roles in the process?	7
What data is involved and how is it used?.....	7
The Inside-Out Analysis	7
Step 3. Identify Preventative Measures	8
Step 4. Develop Recovery Strategies	9
Step 5. Develop the Plan	10
Plan Section 1: Introduction	11
Plan Section 2: Operational Overview	11
Plan Section 3: Notification/Activation Phase	11
Plan Section 4: Recovery Phase.....	12
Plan Section 5: Reconstitution Phase	14
Plan Section 6: Appendices	14
Step 6. Plan Testing, Training & Exercises	14
Step 7. Plan Maintenance	15
A Final Thought	15

Purpose of the Guide

The primary goal of this guide is not simply to provide a checklist of tasks, but to help you develop an understanding of the disaster recovery (DR) planning process and the principles that underlie it. Before getting into the details, let's consider first just what this guide will do for you and equally importantly, what it will not do.

What this guide *will* do is to lay out a framework for your company's DR planning that keeps things conceptually simple and helps you know what steps must be carried out and why, without a lot of jargon or unnecessary formality. The framework we'll use is equally relevant to disaster recovery planning for a division of a large multi-national corporation and for an operation involving a dozen people in a small office. Of course the scales of the tasks involved in these two cases differ rather drastically.

The focus here will be a practical one. Good disaster recovery planning is about identifying those processes and resources that are truly critical, developing realistic recovery objectives for them and then developing a plan that can achieve those objectives as simply and cost-effectively as possible.

We will also focus on making the planning process *doable*, even if this sacrifices some sophistication. The reality is that a sophisticated DR plan that is too complex or expensive to properly maintain and test is *worse* than a plan that only does the minimum because it gives a false sense of security.

So, this guide is intended to help you negotiate the decisions that you'll need to make in order to develop an effective, executable plan that allows your organization to recover critical processes in order to function after a disaster.

Now to what this guide will *not* do.

First, it will not make you an expert on disaster recovery planning. Nothing but experience and observation can do that. It may well be that you need to hire outside expertise in order to develop a plan of the scale required by your organisation. This guide can still be of value to you since it will help you more effectively participate in the planning process and to understand and evaluate what the outside experts are doing.

Also, it will not teach you about all the different software, hardware and service solutions available as components in a disaster recovery plan. Certainly we will present some general information on the kinds of options available and the trade-offs involved, but the field has many different options and continues to evolve rapidly. There is no substitute for research on the latest offerings that are relevant to your needs and direct discussion with the vendors involved.

Above all, this guide will not make disaster recovery planning "easy." The complexity of the planning at the very least mirrors the complexity of the processes that must be recovered, and the best planning guide in the world cannot change that. Nevertheless, preparing a plan is not easy, the procedures described here will aid significantly in navigating the difficulties and keeping the complexity under control.

DRG's Approach

Whatever the scale and complexity of your organization, it is important to employ an approach to disaster recovery planning that meets the highest standards. Whether the final plan is to be simple or complex, its quality will necessarily reflect the quality of the process used to develop it.

For this reason, this guide adopts the framework that the National Institute of Standards and Technology (NIST) employs in its [Contingency Planning Guide for Information Technology Systems](#). The guide is targeted at government agencies that deal with sensitive information and is fairly long and complex, but the framework is straightforward, consisting of the following seven steps, which we take verbatim from the Executive Summary:

1. **Develop the contingency planning policy statement.** A formal department or agency policy provides the authority and guidance necessary to develop an effective contingency plan.
2. **Conduct the business impact analysis (BIA).** The BIA helps to identify and prioritize critical IT systems and components.
3. **Identify preventive controls.** Measures taken to reduce the effects of system disruptions can increase system availability and reduce contingency life cycle costs.

4. **Develop recovery strategies.** Thorough recovery strategies ensure that the system may be recovered quickly and effectively following a disruption.

5. **Develop an IT contingency plan.** The contingency plan should contain detailed guidance and procedures for restoring a damaged system.

6. **Plan testing, training and exercises.** Testing the plan identifies planning gaps, whereas training prepares recovery personnel for plan activation; both activities improve plan effectiveness and overall agency preparedness.

7. **Plan maintenance.** The plan should be a living document that is updated regularly to remain current with system enhancements.

Over the next seven sections, we will consider each of these steps in turn. We will discuss the purpose of the step, provide some guidelines on what is required to carry it out and discuss how to implement it to match the needs of your business.

Step 1. Develop the Planning Policy Statement

This is an easy step to skip or skim over. Don't — it is probably the most important of the seven steps.

The NIST description is a bit dry: the “policy provides the authority and guidance necessary to develop an effective contingency plan.” Dry or not, a clear statement of the “authority and guidance necessary” is vital to the success of the planning venture.

The policy statement is really about communication between management and those responsible for developing the plan.

By making clear the driving goals of the project, the level of financial and other resources the effort commands and the particular people who are to be responsible, the policy statement gives planners everything they need to work out options that can achieve the organization's goals. It also provides a basis for planners to communicate back to management either their success or the need to reassess the goals or the resources, should that be necessary.

The importance of this step extends well beyond the stage of DR plan development and implementation. Why? Because much, sometimes most of the cost is incurred *after*

the initialization phase, during testing and maintenance and of course, in the worst case, during and after a disaster that proves the inadequacy of the plan.

This is probably a good time to point out that you may need a couple of cycles through the steps. The first version of the policy may set goals that turn out to be impossible under the resource constraints specified. You will need to re-evaluate the policy and scale down goals, scale up resources, or attempt some radical rethinking. The important point to remember always in disaster recovery planning is that reality is your partner and like it or not, you must cooperate with it, not fight it.

With that introduction, here are the key points that the policy statement should address:

Objectives

First, what kinds of disasters do you intend to cover?

Large-scale disasters like earthquakes, power-grid failures, hurricanes, etc., account for only about 5% of IT system downtime incidents overall. Other causes include user errors, software errors, virus and other attacks, maintenance that leads to unplanned downtime, etc.

Different threats require different types of solutions. Equally important, what do you want to accomplish? Is the goal simple survival, or to resume functioning within a day or a few hours or a few minutes? Possibly, even to take advantage of the crisis to improve your organization's image with stakeholders, a good way to turn DR planning from cost to strategic investment? Try to express the purpose of the plan in terms of *external* requirements, e.g., impact on customers or other stakeholders outside the IT department, or the need to satisfy statutory or regulatory requirements, since this will maximize your flexibility in meeting the goals. Make sure also those goals are quantified — it is difficult to hit a goal you can't measure. Stay as high-level as possible. It may be necessary to mandate low level requirements such as test frequency, but there should always be good reason to do so, such as industry standards or legal requirements.

Scope

Where does the responsibility of the plan and the planners end? For example, are only IT

hardware and software systems to be covered? What about layers of staff and non-IT systems between the IT infrastructure and the external functions that depend on them all?

In general, it is helpful to break up an organization's overall DR plan into a number of more limited and simpler plans. An outside-in approach works best, with the internal resources covered by one plan becoming the source of external requirements for the next level in.

Resources

What is the maximum level of resources that the plan can command during preparation, implementation, testing *and* maintenance? Make sure you include constraints on all relevant resources: financial, staffing, equipment, space, etc.

Roles and Responsibilities

Who is responsible for the various components of the plan and what is their authority? Who has the right to make final decisions? Who is responsible for ongoing testing and maintenance activities?

Obviously, there is a lot to consider in this step. Two final comments before we move on. First, it is important that the policy statement and indeed the entire plan, be written down. Above all else, the plan is a set of instructions and instructions are best kept in writing — if they are not to be forgotten. Remember, this policy will not only guide the initial development of the plan, it should also guide maintenance and review in the future, since it states the purpose and the expectations for resources on an ongoing basis.

At the same time, the statement need not be long or complicated. A half page of bullets covering the key points is sufficient for many organizations. Best is to start small and add as questions arise that are not addressed by the version you have.

Step 2. Conduct a Business Impact Analysis

The purpose of this step is to ensure that you are protecting everything that you need to protect, without wasting resources on systems that are of secondary importance. The

goal is to determine what must be recovered and how fast — this information will be used to develop recovery strategies.

The output of this step is a prioritized list of critical data, roles and IT resources that support your organization's business processes, together with maximum outage times for each of the critical systems.

To begin, you must identify the key business processes that underlie the organization's ability to carry out its business and the requirements that drive these processes. It is very important that this be done from the outside-in, starting from the standpoint of external stakeholders, whether they be customers of the company, outside suppliers, or internal departments within the company that depend on the IT services you provide. It is also important those actually involved in the business processes be engaged in the planning process, including external stakeholders, the internal staff who deal with them and those charged with operational support of the process.

The remainder of the analysis is carried out for each of the processes identified, with two distinct phases, one that works from the outside-in, the other from the inside-out. At the end of this guide is a template System Information Form (SIF) that you may find useful for collecting and organizing the information gathered during the analysis. Whether you use this, simple text documents, spreadsheets or other forms, one key purpose served by recording the information in a standard format is that it is available for easy reference both during the remainder of the initial plan development and during later plan maintenance and revision.

The Outside-In Analysis

The outside-in phase of the analysis focuses on whole systems and is like peeling layers of an onion. At each layer, we consider the current process or system as distinct both from the users or other systems that depend on it and from other systems on which it depends.


Depending on the overall complexity of your business and how it makes the best sense to divide things up in your context, you may end up with just a single layer or with many of them.

For each layer, we must answer the following key questions:

What are the availability and performance requirements for the current process, system or resource? What is the cost if these requirements are not met?"

These requirements are of course driven by the business process or other system that depends on this one. For example, one transaction processing system may be required to handle 25,000 transactions an hour, essentially 24 X 7, while a CRM system supporting a field sales force with strong customer relationships may only see a few hundred interactions a day and could be down for as much as a week before the consequences would become dire. It may be a good idea to determine an absolute minimum performance and availability level that enables the business simply to survive versus one that allows it to function adequately through the course of the crisis.

With a range of availability options, greater flexibility is available for spending decisions. The estimate of costs is important both for determining how much to invest in your organization's ability to recover from a disaster and in determining the maximum allowable outage that can be tolerated. The NIST guide suggests a useful framework for thinking about this, as illustrated in the following figure from the NIST guide:



In general, the cost of downtime is an increasing function of time, as shown in the blue curve. For example, in a typical organization the cost of email server downtime ranges from annoyance during the first minutes to a small dip in productivity as

workers are forced to shift to lower priority alternative activities to actual lost revenue and loss of customer confidence as downtime extends to the point that the organization is unable to fulfil its obligations.

The green curve represents the cost of recovery solutions as a function of the time to recovery that they support. Not surprisingly, this curve exhibits the opposite behaviour from the first. Solutions that enable recovery over days or weeks are typically much less expensive than solutions that enable a business to recover in hours, minutes or seconds.

The intersection of the two curves yields a rough estimate of the maximum allowable outage when projected onto the time axis. Projected onto the cost axis, it offers a target amount that it makes sense to spend on the solution.

Obviously, at this stage we lack the information in the green curve, which represents the general tendency for the cost of recovery solutions to rise as they are able to more rapidly recover and so we estimate allowable outage times on the basis of best estimates. When we look at specific recovery strategies and their costs in step 4, we will begin to complete the picture. As we have noted before, the entire set of steps that we are covering may need to be reiterated in order to determine the best solution.

What are the points of contact between the current process or system and any other systems on which it depends?

The answer to this question gives us the next layer of the onion and how it interfaces with the current one. For example, the sales process may involve personnel working with a CRM system — the point of contact there may be via a web portal into the CRM application. The CRM application is itself a system that may consist of webservers, CRM application servers and database servers, all with the portal as a point of contact with the sales process. Later you may consider the CRM layer and its points of contact with, for example, the storage systems, if they are administered separately.

At the end of the chain, of course, are basic infrastructure resources and systems such as electric power, telecommunications connections and environmental control systems — these must be considered as well since they may well be impacted by some of the disruptions for which contingency planning is being undertaken.

What are the critical roles in the process?

The point of this information is two-fold. First, in many cases, these roles must be taken into account for recovery. If a process requires the intervention of a person for monitoring, management, analysis, or maintenance on a regular basis, how long can the process run without that role in a crisis? For IT systems, in particular, there may be IT administrator roles that will play a critical part both in the recovery and in running systems for the duration of the crisis.

The second reason these roles are important is that the people in them represent critical sources of information for determining system dependencies and requirements. These are the people who must be closely involved both in this phase of the planning and in subsequent testing.

What data is involved and how is it used?

Data that is used or generated in the process must be accounted for separately from equipment and roles for the simple reason that, unlike a system or role, data cannot be replaced and usually cannot be repaired — it must somehow be copied and the copy later used to restore the data. You should consider as well the role each process plays in the life cycle of the data: which process creates the data, which processes use it and where and through which process is it stored and maintained?

Note that the points of division between layers or processes are different for different organizations. Examples of criteria that may be used to determine how best to subdivide the overall process into system and subsystems include departmental lines; oversight, maintenance and management roles; substitutability by manual or other systems; complexity and size; and of course, common sense.

One helpful way to work through this phase of the task is to perform a mental or actual walkthrough of each process, with the participation of all those who are normally involved. This can help ensure that you do not miss critical dependencies that may not be immediately obvious. The result of this outside-in phase of the analysis is, for each business process, a complete list of systems involved, points of contact between systems and performance and availability requirements at each point.

The Inside-Out Analysis

The inside-out phase focuses on resources that are required in each layer in order to provide the services that have been identified in the previous phase. Beginning from the deepest system or layer, list all IT and infrastructure resources that are required for it to function. Next, for each of these resources, determine the impact of a disruption in the availability of the resource on the functioning of the system and its ability to deliver the services on which outer layers depend. In particular, determine the maximum allowable outage time for each resource before it causes unacceptable disruption in essential functions — essentially the point at which the availability of the system falls below the most stringent requirement of all the systems which depend on it. Be sure to include in the analysis any indirect impact that may occur through related or dependent systems.

Essentially the same analysis should be performed for data and roles as well. Again, what is the impact that results from unavailability or loss of data or from the inability for someone to fulfil a specific role? In the latter case, this may occur because the person is injured or otherwise prevented from performing their duties, but it may also be a result of the lack of access. If an epidemic were to occur, for example, so that people were required to work from home, as happened during the SARS epidemic in 2003 in Asia, or perhaps North Shore Harbour bridge collapse – were people might be quite capable of working, but unable to do so because they lack access to the resources they need.

Carrying out this sequence of analyses yields, in effect, a full chart of dependencies that runs from the outermost layer of business processes to the innermost layer of core

infrastructure on resources, people and data. This is a very valuable tool for later test development and maintenance and should be included in the disaster recovery plan in the System Description and Architecture section described later.

Once this analysis is complete, it is time to develop recovery priorities for IT systems and individual components, beginning with the latter. This task is straightforward if the work described in this section has been done thoroughly since the priorities follow naturally from the outage impact and allowable outage times recorded for each component.

There are many possible scales that may be used for labelling priorities, from a simple high-medium-low qualitative scale to a numerical scale to a scale more focused on business impact, such as “customer-facing high” versus “management and control” versus “low priority maintenance”. Whatever scale you use, it is important that the scale be uniform across all systems based on business impact (in some cases it may be all right to use a different scale internally within a process or system, as long as system-level values remain mutually consistent).

Recovery priorities must be developed at the system level as well. Consistency is obviously vital — it won't work for one system to have a higher priority than another system on which it critically depends, unless it can continue to function without the dependency at an acceptable level. It is convenient to transfer system level priorities to a *Master SIF* which lists each system together with a very brief description of its purpose, the recovery priority, maximum outage time and business impact, major dependencies on other systems and a brief description of the recovery strategy after it has been developed in Step 4.

Now, before turning to recovery strategies, we will briefly consider prevention, since prevention can often be significantly cheaper than recovery after something has gone wrong.

Step 3. Identify Preventative Measures

A simple formula for estimating the financial risk associated with a given type of disaster (and thus how much is worth investing in a plan to mitigate that risk) is $R\$ = P \times C \times T$

where P is the probability that the disaster will occur, C is the hourly or daily cost of downtime in lost productivity, lost revenue, etc. and T is the time that systems are expected to be down. For example, if the probability of a major hurricane hitting your place of business in the next 3 years is 20% and it will cost you roughly \$100,000 per day that you are down and you expect that you are likely to be down for a week, then your financial risk is $0.2 \times \$100,000 \times 7 = \$140,000$.

One way to minimize this risk is to reduce T , the time you are down — that is basically the purpose of the disaster recovery planning exercise. However, it is not the only way. The risk can be reduced as well by reducing the probability that the disaster will occur or by reducing the cost that will be incurred if it does. Both of these are types of preventative measures. It is very often the case that the cost of preventing a problem is far lower than the cost of fixing it after it occurs.

Measures that reduce the probability of a disaster occurring range from fairly drastic, like physically moving the organization out of reach of threats such as hurricanes or floods, to the fairly mundane, such as: ensuring that regular maintenance is performed on critical systems; that redundant components are built in; that sensors are installed to monitor environmental factors; that performance monitors are installed to give early warning of server malfunction; even something as simple as keeping plastic tarps available to throw over computer equipment to protect it from water damage.

It is sometimes even possible to reduce the cost of downtime by reducing your organization's dependence on the system. The basic idea is to examine the potential win of removing or replacing a system entirely. What is sometimes forgotten when new equipment and systems are implemented is that the total cost of any system includes not just the upfront cost and the ongoing maintenance, but also the *risk* associated with it. There are times that it is better to replace a system with one that, while lower in performance, exposes the organization to significantly lower risk.

While we don't have any particular procedure to offer, it is potentially very useful to spend some time in this step both for all types of disasters that you wish to protect against and

for all the systems being protected, at both the full-system and component levels.

Step 4. Develop Recovery Strategies

The primary task of this step is to determine how you will achieve your disaster recovery goals for each of the systems and system components that were identified in

Step 2. It is here that you do the core work of balancing costs and benefits of the available approaches, before diving into the complexities of the full plan.

This step is not about selecting specific vendors, determining exact costs, or developing detailed procedures. Rather, the purpose in this stage is to select the *types* of solution that you will use and to determine the *scales* of the costs involved.

Thus, for example, you may determine that a small, critical subset of systems require a fully-mirrored and staffed alternate site ready to take over in minutes, while other systems can utilize a more traditional backup strategy which trades longer recovery time for much reduced expense.

There are several considerations to keep in mind as you work through this step.

First, this is the point at which it becomes important to consider exactly what types of disasters you need to prepare for and to classify them by the extent and type of impact they have. The reason is straightforward — the recovery strategies available to you necessarily depend on what you must recover from.

A hardware component failure or a water leak over a couple of servers is very different matters than a site-wide disaster like a flood, fire, or regional blackout. In one, it may be possible to depend on a vendor to deliver a replacement. In the other, an alternate facility may be required.

The second consideration that arises from this is the need to consider solutions of differing breadth of coverage.

Obviously, a solution that can address a site failure will serve as well to recover from failures of individual systems or components.

There are a number of reasons not to depend on a single system-wide solution to address all issues, however. The most obvious is that such a solution is certainly costly to implement and probably costly to activate — the disruption from failing over all of your systems to a secondary site is unlikely to be commensurate with a problem arising from failure of a single component.

There are two other reasons not just to consider but actually to implement several alternative approaches that address different levels of problems. First, no matter how good your planning and testing are, depending on a single solution, especially a complex one, means that your recovery is all or nothing. Anyone with much experience in complex systems knows that this is not a good idea. It is much better to have a series of backup solutions so that, if one fails, another is in place to recover at greater, but still not completely devastating cost. Second, alternative solutions that are actually implemented can give you significant flexibility in responding to an actual disaster.

For example, if a problem occurs in the middle of a business day, it may be important to go immediately for an expensive solution in order to recover quickly, but if it happens at night or over a weekend, you may be able to institute slower, but less disruptive or less costly recovery procedures.

A final consideration is that you need to take into account the particular characteristics of the infrastructure, human and data aspects of recovery. Each of these must be considered differently, with different fundamental drivers of the decision of what type of solution to invest in and how much to spend.

Infrastructure is the simplest. While there may be more or less significant costs involved, the salient characteristic of infrastructure is that it can be replaced. A server can be replaced with another server, an alternate provider can be found for network connectivity and so on. In many cases, the replacement system need not even be an exact duplicate, as long as it interfaces with other components and systems in a compatible way: manual systems may replace IT-based automated systems, or you may be able to temporarily outsource an internal system. The fundamental driver when

considering recovery strategies for infrastructure is typically that of cost versus performance.

People (roles) represent a more difficult factor. Particular roles usually require special skills and knowledge. If a recovery strategy requires, for example, that a given role be duplicated at an alternate site, there may be additional costs associated with hiring or training personnel to fill that duplicate role. Similarly, the recovery solution itself may require special training or skills. The key driver for the consideration of recovery strategies from the standpoint of roles, then, is the degree to which they require you to duplicate or acquire special skills, which then impacts the long-term cost of hiring and training.

Finally, data is potentially the most difficult issue because data is usually unique — it is not generally possible to replace data with other data that has the same or similar properties. Either you have it or you don't. The driving question from the standpoint of data then is: how much data are you willing to lose? It is also important here to recall that data may not only be destroyed, for example through the loss of the system it is stored on, but may also be corrupted through user or administrator errors or through deliberate sabotage, for example, through a virus attack.

Keeping all these considerations in mind, this step consists in reviewing each of the systems characterized in Step 2 and determining the system and component-level strategies to apply that can achieve recovery within the maximum outage time, while remaining roughly in the bounds of budget and other resource constraints that have been established. In considering cost, it is always important to characterize the *total* cost of a solution, which means not only the cost of any hardware, software or services purchased, but also the level of disruption caused by installation and the on-going costs in both money and personnel to maintain and test the solution.

Keep in mind too that the different systems are not independent of one another and in many cases require that the recovery strategies be compatible. This may be a minor consideration for purely local solutions, but if you wish to set up an alternate site, you need to ensure that all the systems on which a given system depends are also duplicated.

For example, if you wish to have a failover system for your email servers at an alternate site on the other side of the country, in addition to the email server you will need facilities to house the server, power, network connectivity, sufficient bandwidth to operate throughout a crisis, access or personnel to maintain the facilities and the servers, DNS servers and other auxiliary systems that are critical to the operation, etc.

When you are done, add a brief summary of the recovery strategies for each of the systems on the Master System Information Form.

Step 5. Develop the Plan

This step is the culmination of all your work. It is not, unfortunately, an easy step, but neither is it too complicated, as long as you have been thorough in the previous steps and you approach it systematically.

The outcome of the step is both a documented plan and the completed implementation of all the infrastructure required to enable the plan. The documentation includes background information on the assumptions and constraints that went into making the plan, as well as written documentation on specific procedures. The implementation side includes purchasing and installing hardware and software, setting up alternative locations, contracting for alternative sources of network or other communication services and so on.

This step is a major project all by itself, even if the previous steps have been carried out perfectly. It will require a significant amount of time on the part of the person or team responsible for leading development of the plan, but it will also require time and effort by everyone whose systems are involved since their expertise will be required both to develop recovery procedures and of course, to test them.

We will cover all aspects of the plan, roughly following the organization suggested in the NIST guide. The organization itself is not important — that should be adapted to best serve your needs — but all the types of information we will discuss should be present in the plan. The sections we will use are as follows:

1. Introduction
2. Operational Overview

3. Notification/Activation Phase
4. Recovery Phase
5. Reconstitution Phase
6. Appendices

For each section, we will briefly discuss its contents and the purpose for their inclusion, then review some of the major issues or potential pitfalls to keep in mind and finally offer some suggestions on how to approach the actual development of that section.

Keep in mind that this part of the work in particular is likely to be iterative. As you select specific solutions and work out step-by-step procedures, you may discover new dependencies, errors in your initial assumptions, or simply that your planned approach exceeds the resources available, so that you need to revisit and re-evaluate your recovery strategies and perhaps even your recovery targets. In fact, you may wish to do this deliberately in preparing the entire plan — to make several passes through all of the steps in this guide, but starting with a much shallower effort and deepening gradually. This approach can significantly reduce the likelihood of surprises that entail rethinking major parts of your plan.

Plan Section 1: Introduction

The main purpose of this section is to document the goals and scope of the plan, along with any requirements that must be taken into account whenever the plan is updated.

Preparing this section is simple since it is essentially the Planning Policy Statement that you prepared in Step 1 — you may as well just include its contents verbatim here.

This is also an appropriate place to include the record of changes made to the plan whenever it is updated.

Plan Section 2: Operational Overview

The purpose of this section is to provide a concise picture of the plan's overall approach.² It contains essentially two types of information: (1) a high-level overview of the systems being protected and the recovery strategies employed and (2) a description of the recovery teams and their roles.

The section provides a context for understanding the plan as a whole. It should be possible to read just the introduction

and this section and have a good idea of the overall approach to disaster recovery for all the systems involved.

The first type of information in this section, once again, has already been prepared by you during the previous steps: the business impact analysis and the development of recovery strategies. An easy way to do this part of this section is simply to include a copy of the Master System Information Form (SIF), which lists all the major systems with individual SIFs and annotates them with recovery priority, maximum outage time and business impact, major dependencies on other systems and a brief description of the recovery strategies employed. It is recommended to include all the individual SIFs in an appendix of the plan as well.

These serve as a useful reference during reviews of the plan and ensure that all information is together during future updates.

The operational overview should also contain a description of the teams who will be responsible for activating and carrying out the plan. Detailed information on the teams and the succession plan will be prepared in the next section and so the information here should be brief, probably no more than a list of key teams and roles together with a few words describing their role in the overall recovery plan. It is probably best to develop this part after completing all of this phase of the work to ensure that it reflects the actual structure of the recovery teams after all decisions have been made.

Plan Section 3: Notification/Activation Phase

This is the first of three plan sections that document actual recovery operations.

According to the NIST guide, this section “defines the initial actions taken once a system disruption or emergency has been detected or appears to be imminent. This phase includes activities to notify recovery personnel, assess system damage and implement the plan. At the completion of ... [this] phase, recovery staff will be prepared to perform contingency measures to restore system functions on a temporary basis.” It is common for organizations to plan this phase inadequately.

Most of your work on the plan is focused on the actions that you must take to recover. It is very easy to forget that *declaring* the

emergency and deciding that it is time to initiate operations under the disaster recovery plan can be difficult and requires advance planning as well. In fact, there is a very natural tendency, absent clear guidelines, for people to delay action until they are "certain" that the disaster is imminent or is bad enough to react. Unfortunately, by the time certainty is achieved, it is generally too late.

This section must answer the following questions:

- On what basis will the decision to activate the plan be made?
- What information is required?
- What additional guidelines are there for making the decision?
- Who is responsible for performing the damage or threat assessment that will provide the information above?
- What restrictions or guidelines are there for how long the assessment may take or how certain the information must be?
- Who is responsible for the go/no-go decision?
- What are the rules of succession, if that person is not available?
- How will teams be notified?
- How will recovery teams communicate?

There are several things that should be kept in mind when preparing this section of the plan. When preparing guidelines for the activation decision, account for the fact that there is a natural tendency to be conservative in assessing risks to avoid being seen as an alarmist if the feared damage doesn't actually happen.

Wherever possible, provide clear, positive and specific guidance on when the person in charge *should* or *should not* activate the plan. Make sure that the person responsible for decisions is in a position to know what they need to know. For example, if the decision requires knowledge of conditions at a site, it is best to designate someone located at that site, in case communications are impossible.

In fact, communication is a key aspect that must be carefully thought through. Depending on the type and scale of the disaster, common means of communication, such as cell phones, email and land-line phones may not be operational.

Either alternative methods must be developed, like satellite phones or radios, or local teams

need to be empowered to make decisions autonomously and given the tools they need to implement those decisions.

This section may be an appropriate place to keep contact information on each team member; alternatively, it may be included as an appendix. Keep in mind that you can never have enough contact information. Especially for key recovery personnel, you may wish to include work and home phones and email addresses, faxes, cell phones, neighbours, relatives, etc.

This is the first point in the development of the plan where you must design actual procedures to be followed during a crisis. It is impossible to develop good procedures of any complexity without playing them out. For this reason it is critical that the procedures developed for this section be tested in actual use, just like the recovery procedures themselves. It is also useful, during the development of the plan, to role-play a crisis, preferably with the people who will actually be involved, in order to think through the steps necessary.

Plan Section 4: Recovery Phase

This is the second of the three major sections documenting actual recovery operations, but it is the one that most of us have in mind when we talk about a DR plan. This is the section of the plan that documents in detail the solutions to be used to recover each system and the procedures required to carry out the recovery and restore operational capabilities.

The organization of this section is simple. For the organization as a whole and for each system individually, the plan identifies a sequence of recovery goals (for example, to restore internet connectivity, or to switch email services to a backup system at a secondary site) and provides documentation on the procedures required to accomplish each of them.

Procedures may be as simple as a couple of bullets or may be many pages of instructions and checklists, depending on the complexity of the recovery solution and of course, of the system. There are two aspects to your work in developing this part of the plan: actually implementing solutions that align with the recovery strategies you identified in step 4 and then documenting the procedures required during recovery.

Implementing your recovery strategy is, of course, the major work here. This includes everything from evaluating and purchasing hardware and software solution components, designing and implementing the solution around these components, equipping alternate sites for your systems, negotiating with managed service facilities, vendors and IT consulting companies and so on. We could not possibly lay out everything you need to know here in a few paragraphs.

That would require a substantial book, one that would be out of date as soon as it was published. This leads us to the most important point to be made about doing this section of the plan. More than anything else you need to make sure that you are finding and using any information, resources and help that you can. You may wish to consider outsourcing some or all of this work to a professional services organization that specializes in disaster recovery system development — although it is important to understand that it will still take a lot of effort on your part to ensure that they have the information they need about your organization.

Whether you are doing all of it yourself, outsourcing parts or even everything, it is a good idea to utilize other resources to help ensure that you understand the options available and the tradeoffs involved. Here are a few suggestions:

- A minimal effort searching on the web will turn up lots of articles, guides, checklists, etc. Specialized sites like the Disaster Recovery Journal organize some of it for you, but it is worth an independent search as well. Make use of all this information.
- Talk to other companies in your area or similar to you who have implemented disaster recovery plans and learn from their experience.
- Talk to your vendors and insist on their help. Obviously they have particular interests and points of view, but they are also specialists in the technologies that they offer. Make sure that you are able to speak with knowledgeable technical people, not just sales and marketing.
- Hire professional help.
- Try to consider a diversity of options at every level. When thinking about the overall strategy, ask yourself if a radically different strategy might work as well. When you are considering specific technologies or types of solutions,

make sure that you look at a variety of vendors that differ not just in feature sets, but also in company type. Ask them how they are differentiated from competitors, how they are similar and to keep them honest, under what circumstances they would recommend a different type of solution.

- Don't assume a solution will work as advertised — seek proof. Evaluate technologies in a lab whenever possible before committing. Ask for references with installations that are similar to yours.

When you have selected and implemented the recovery solutions that you will use, the remaining task is to develop and document the actual procedures to be followed for the recovery of each system and for coordination between them, since the recovery of one system often depends on what is happening with another.

There is no substitute for action in this phase — *do* the recovery and write down what you did. Then do it again, following the instructions and see if they are correct and complete. It is important that the level of detail in the instructions correspond to the level of knowledge of the *least* knowledgeable personnel who might need to carry them out.

To minimize disruptions, much of the *doing* can be simulated or role-played and procedures for individual subsystems can be worked out independently. Nevertheless, it is extremely important that both final documentation and final testing include full scale tests as they are expected to occur during a real crisis. This is the only way to discover subtle inter-system dependencies that might otherwise be missed.

It is very important to highlight any points in the procedures that require coordination with other teams or other systems

— this is information that should be easy to see by skimming quickly over a given procedure. If the procedure is long and complex, it may be helpful to include an overview of the major steps in an introductory section, or to break out some of the details into separate checklists or sub-procedures. Of course, while completeness and correctness are key, the shorter and simpler the procedure, the better. Remember that these procedures will be carried out in a high-stress environment, which makes mistakes and confusion much more likely than during a practice run.

Plan Section 5: Reconstitution Phase

This is the last of the three sections of the plan that documents actual operations and it is again one that does not immediately leap to most peoples' minds when thinking about disaster recovery planning.

Disasters eventually end and there is a need to return operations to normal. Per the NIST guide, in this phase "recovery activities are terminated and normal operations are transferred back to the organization's facility. If the original facility is unrecoverable, the activities in this phase can also be applied to preparing a new facility to support system processing requirements."

There are two reasons that are very important to think through and document reconstitution procedures just as carefully as recovery procedures. The first is that it will help you ensure that the solutions you select are *supportive* of a relatively painless return to normal. A solution that gives you easy recovery, but makes it hard to get back is not a particularly good solution.

The second reason is even more important. Even though stress levels may be lower during the return to normal, doing it without well documented and well tested procedures, risks mistakes that can transform the return to normal into *another* disaster.

The type of content and suggestions for preparing it are the same as for the recovery section, so we won't repeat them here.

Plan Section 6: Appendices

The appendices should contain any information that (a) is necessary as reference material during recovery, (b) may be necessary during any revision of the plan, or (c) documents legal agreements. Examples include:

- Team contacts
- Vendor contact (including offsite storage)
- SOPs and checklists for system recovery or processes
- Lists of equipment, system requirements for hardware, software, firmware, etc.
- Vendor SLAs, reciprocal agreements, etc.
- Description of/directions to alternate site
- Worksheets used to develop the plan

Step 6. Plan Testing, Training & Exercises

Over time, things change. Hardware components are replaced, software is upgraded, networks are reconfigured, data sizes grow, people come and go. All this is a normal part of the life of an IT environment. And all of it can impact the performance of your disaster recovery systems. Although these systems were fully tested when first installed, the dynamic nature of the environment makes it critical that testing continues to take place regularly and that personnel training be up to date.

There are many different types and levels of testing. Generally speaking, they span two key dimensions: scope and realism. Scope refers to the degree to which you are testing a full system or just individual components. Realism refers to the degree to which you are performing exactly the procedures that you would during a disaster — a class room role-playing test in which you talk through steps without actually doing them is one extreme and a full execution the other. In both cases, one side of the spectrum tends to be less expensive and less disruptive to day-to-day operations but also less reliable in its results.

In general, it is a good idea to do a mix. Less disruptive tests can be carried out more often. Problems found and fixed that way avoid the typically higher impact and cost of finding the issues during a live test. The NIST guide states that "it is important that a test never disrupt normal operations." We would modify that. If never disrupting normal operations means never performing a full disaster recovery exercise, then it is necessary to occasionally disrupt normal operations. Such disruptive tests should be kept to a minimum, perhaps once or a few times a year, as long as component and subsystem testing are carried out more regularly. It is important to remember that it is always less expensive to expose yourself to the cost of a full test in a planned way than to discover during a disaster that a subtle missed dependency leads you to be unable to recover at all.

We should note finally that it is important to be thinking about the testing side of your plan throughout the previous steps since testing

represents a significant part of the total cost of your disaster recovery plan. In particular, when considering particular solutions and vendors, make testability part of the evaluation process. Proper training is equally vital. Training in disaster recovery procedures should be considered part of the regular orientation of new hires if they have any role at all in implementing the plan. Key disaster recovery personnel should undergo frequent enough training that they are intimately familiar with the procedures that they will have to carry out under the plan. As noted in the NIST guide, ideally they should be well enough trained that they can execute their responsibilities without the aid of the actual disaster recovery plan document.

you have any questions and would like to discuss your DRP requirements, please don't hesitate to ring us on: 0800 99 88 20.

Step 7. Plan Maintenance

If it is worth the money and effort to develop a disaster recovery plan, it is also worth the effort to ensure that the plan accurately reflects current requirements and systems.

Otherwise, it is only a matter of time before the two diverge sufficiently to put your capacity to recover from a disaster in real jeopardy.

For the most part, this step is beyond the scope of this document, but we offer one thought. There are three natural points at which the plan can be reviewed: during testing, in a regular annual or semi-annual review devoted specifically to the task of review and when changes are made in either the IT systems being protected or in the business processes they support.

The first two fall directly under the purview of those responsible for disaster recovery planning and so can be planned for directly. The last requires that consideration of the impact of changes on the disaster recovery plan be introduced as a standard consideration in procedures that are *outside* the scope of direct concern of those responsible for the DR plan. As a result, it requires that, one way or another, those responsible for changes in the systems take on a certain level of responsibility for DR plan impact. While this may be difficult, it makes maintaining a correct DR plan significantly less costly than discovering changes later, through testing or an annual review.

A Final Thought

We hope that this guide is useful to you in your own disaster recovery planning efforts. Should