COMPASS CONSULTING INTERNATIONAL
233 Heaths Bridge Road, Suite 102, Concord, MA 01742
978-243-2413
info@compassconsulting.com

Business Continuity Planning – An Overview

By Geoffrey Tritsch and Dr. Robert Kuhn

Introduction

No campus is immune. In the wake of every tragic event (and sadly there are too many of them lately), there is sure to be a resurgence of disaster preparedness planning. Towards that end, the Association of Telecommunications Professionals in Higher Education (ACUTA) asked Compass Consulting International, Inc. to put together a series of articles on this significant and timely issue. The purpose of this compilation of those articles is to give you a framework to think about disasters, preparedness, avoidance, mitigation, recovery, and planning.

Here's an outline of what we will cover:

- Where Do You Start?
- Did the Lights Go Out? (Events, Results, and Impacts)
- What Are You Afraid Of? (Risk Analysis)
- What You've Got and What You Need. (Information Gathering)
- Writing and Selling the Plan
- A Living Disaster Plan (not Living with a Disastrous Plan)

Part 1 - Where Do You Start?

What you need for a plan depends a great deal on what you're trying to accomplish. To start, we suggest you ask yourself a few questions:

- **Why are you developing this plan?**

  (We're really concerned. We need it for a risk audit. Somebody said so. Where have you been since September 11? ...)

- **Who is driving the planning process?**

  (IT, Risk Management, External auditors, Senior Executives ... Every successful project has a champion, someone high up to shepherd it along. Who will be your champion?)

- **What are the business issues related to this planning process?**

  (Is the scope of the disaster preparation a single service (like telecommunications or the financial system), all campus technology, or the institution as a whole? Is there a commitment to provide the resources (people and dollars) that are needed for the disaster planning process, as well as a commitment to address the shortcomings that you find?)

- **What are the desired end results?**

  (Real protection, a workable plan, getting them off my back, a "cookbook" with call-out lists and step-by step procedures, a methodology...)

How you answer these questions can make a significant difference in what you do, how you do it, how long it will take, what resources you need, and, ultimately, how successful you might be.

General Dwight D. Eisenhower said, “In preparing for battle, I have always found that plans are useless, but planning is indispensable.” This is also true for disaster planning. It is important to understand that the real value in disaster planning lies not in the report that is produced (although call-out lists and procedures are definitely of value), but in the following three areas:

1. **The decision-making/assessment process.** (What could happen? What if it did? How do you prevent, mitigate, or resolve the impact? Is it worth it?)
2. **The data gathering process**. (What do you have? Where is it? Who uses it? For what? How is it at risk?)
3. **The increased awareness** that results from such a project.

Next, we will begin to look at the assessment process.

Part 2 - Did the Lights Go Out? Events, Results, and Impacts

One of the hardest parts of disaster planning is trying to establish some conceptual framework for the plan. There is so much that could happen. You clearly can’t plan for every eventuality so how do you decide where to put your resources?

Talking loosely, we often speak of hurricanes, earthquakes, mudslides and the like as natural disasters; and fires, explosions, terrorism, vandalism, etc. as man-made disasters. But these events are not in and of themselves disastrous, no matter how large.

Think of a monster hurricane in the middle of the ocean, or the thousands of acres of bush fires that renew remote Australian forests every summer. Unless there are actually negative effects from the event, there is no disaster. In fact, we don’t even measure disasters by the result or chain of results so much as by the final impact. Who or what was hurt? What was lost? How long will it take before full operation is restored? How much will it cost to fix?

To make this abstract discussion more concrete, let’s take a (somewhat improbable) example. Suppose a tornado tore through and destroyed a whole neighborhood which had already been evacuated for demolition in order to build a giant new ballpark. The results in terms of physical effect would normally qualify as a major disaster, but the lack of impact – few injured, none killed, little expense – means that there was no disaster.

There are essentially three components to any disaster: An *Event* has *Results* which in turn have *Impact* on the user community. When deciding on a business continuity strategy in the face of possible disasters, one can target any or all of these components for prevention, mitigation, amelioration, or restoration.

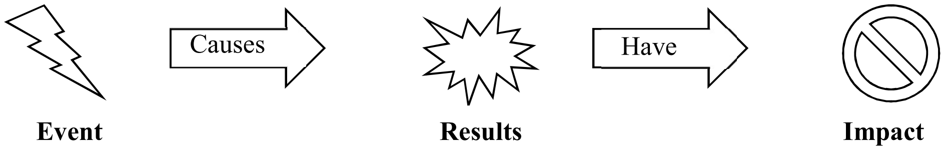

Let's look at a simple example. "Backhoe fade" is the term given to the effect on the network that one experiences when an excavation contractor severs a cable. Good processes, accompanied by accurate maps, reduced the risk that anyone will unwittingly start digging over one of your conduits. Prevention is the best approach where it is possible.

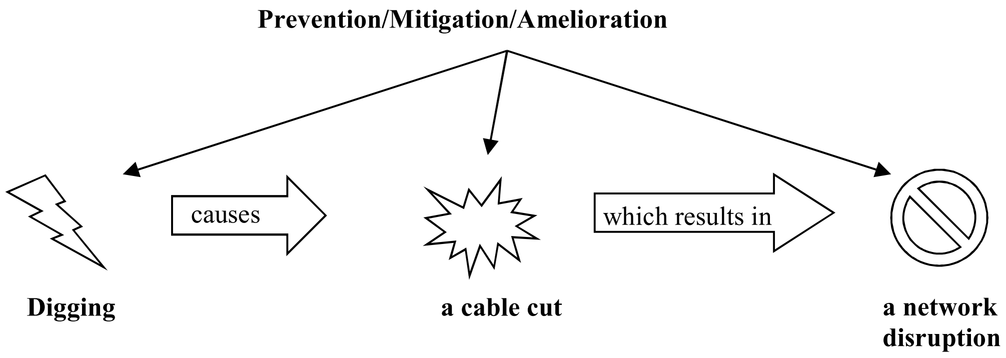

However, through ignoring procedures, misreading maps, or risk-taking ("I won't be digging that deep") digging will still take place. Burying warning tape above the route of network conduits is a measure to prevent the backhoe getting to the conduit – in other words, preventing the results. Concrete-encasing the conduit is designed to mitigate the result if the backhoe does dig that far. If the conduit (and its contents) is severed, then diverse routing can mitigate the impact. If your network can adapt to the loss of the connection, then the impact is ameliorated. Otherwise, your choice is to fix it (recover) or abandon it.

All disaster planning or business continuity strategies can be reduced to these three categories:

- prevention or avoidance strategies (stop it happening)
- mitigation or amelioration strategies (it may happen, but it won't matter so much)
- and recovery or restoration strategies (fix it afterwards).

For each disaster (or category of disaster as you probably can't plan for every eventuality), ask yourself the following:

(Pick one from each column):

Column A	**Column B**	**Column C**
Can I	Prevent/Avoid	the Event
Should I	Mitigate/Ameliorate	the Results
Must I	Restore/Recover from	the Impact

Your answers need to be prioritized using a healthy dose of common sense: first, protect lives; second, limit physical property damage; and third, mitigate/limit interruption of business operations.

The common wisdom is that prevention is much, much cheaper than recovery. On the other hand, major events tend to be very, very unlikely. The next section "What Are You Afraid Of? (Risk Analysis)," will look at the balancing of expensive measures against unlikely events.

Part 3 - What Are You Afraid Of? (Risk Analysis)

When is prevention is worth the cost?

There are several attitudes to valuing disaster prevention strategies. Generally, one tends to be either a **Grasshopper** or an **Ant.** Grasshoppers play the odds. A disaster is unlikely to occur. Why worry, be happy. Ants take precautionary measures, no matter what the "odds," figuring that if they don't, they are going to have the disaster. While few institutions are wholly Ants, too many institutions are Grasshoppers by default – they have no structured approach to disaster planning.

The heart of risk analysis is the concept of "expected cost." Expected cost is the probability of any event multiplied by its associated costs. (The Ants' expected costs are pretty much fixed by the annual and annualized costs of their preventive measures. On the other hand, Grasshoppers either pay nothing if there is no disaster or the full cost of recovery if there is one. Averaged over enough time and enough Grasshoppers, this all-or-nothing approach would yield the expected cost for the Grasshoppers.) The statistically rational person (or insect) acts to minimize expected cost. From a purely statistical viewpoint, if expected cost of prevention is less than that of loss, you prevent; if greater, you don't.

One can justify being more "Antish" than statistically rational by thinking of prevention as insurance. When you buy insurance, your costs are likely to be higher, but you won't be wiped out by a disaster. Applying this thinking to disaster preparation, prudence encourages us to go with prevention even if the purely statistical rationale points the other way. Redundancy, diverse routing, fire suppression, alternative sites, and other preventative measures all cost money and offer little return on the investment if they are never needed. But they provide insurance against costly, high profile, and embarrassing results.

On the other hand, you can't prevent all possible disasters, and as you get closer to doing so the costs rise asymptotically. So when the cost of prevention gets too high (when compared to the expected loss), one option is actually taking out insurance against the disaster instead of attempting to prevent it. Insurance companies write millions of policies, and, so long as the premiums are lower than the cost of prevention, taking out an insurance policy could be better protection against disaster. However, keep in mind that insurance policies don't prevent disasters. They only help you address the financial aspect of recovery.

Of course, the discussion above assumes you can measure the cost of a disaster. For profit-making businesses this is mostly true (although how do you value human life?) It is more problematic in a non-profit arena, where costs and benefits don't come neatly in dollars. Public safety, public image, competition, inconvenience and even politics also come into play. But the financial impact is always a good place to start.

That, in a nutshell, is the nature of risk analysis. Is it better to invest in preventing the event or mitigating its impact or is better to let the event happen and deal with recovery if and when necessary? The answer varies by institution, by disaster, by risk profile, and even by manager.

Part 4 - What You've Got and What You Need. (Information Gathering)

Now that we've discussed general business continuity planning issues and analysis tools, now let's move on to some of the more specific aspects of protection of your technology assets and the use of those assets in the case of an emergency, specifically we focus on information and communication. Of all the aspects of planning, the where and how of gathering information for your planning process is the one that is most often underrated in terms of value and underestimated in terms of time.

Good records are vital to rapid response to a disaster. "Good" records are complete, accurate, up-to-date, and most importantly available. If you have your records on the database server and in hard copy in the machine room in which the server resided, and that room just went up in smoke, you are now in trouble.

What records do you need? You need information on:

- Spaces and pathways, rooms, conduits
- Equipment/Physical components
- Logical and physical network connectivity (including cabling: where it goes and what is on it)
- People: Normal roles, Disaster Preparedness roles, who backs up whom, who uses what to do what, contact information
- Procedures
- Priorities

Besides the basic information on things, and people, such as locations, descriptions, contact information, you need a lot of relational information.

- Who knows about what?
- What services run on which servers?
- Who uses what services for what purposes?
- Which vendors provide which services and equipment?
- What have vendors agreed to do in an emergency?
- If you need to restore services, what should take priority over what?
- How should decisions be made?

These last two questions of priorities and procedures need to be settled to protect both the institution and its officers. Public safety has to be the foremost concern, followed closely by legal requirements, then business priorities. Remember that business priorities may change by season (think of the Admissions Office in the spring), and by time of day (communication with the outside happens 9 am to 5 pm, but purely internal functions, however important, can be offset in time during a disaster to alleviate congestion on scarce resources). The most important procedures to put in place are those that allow your people to respond to events without delay. In part, that means having responsibilities shared among teams of people.

Each inventoried item should list the alternatives. Server flooded – load backup on alternative; system manager in traction – call her backup. In fact, disaster preparation necessitates thinking in terms of clusters of equipment and teams of people. The research that goes into creating this database serves a double purpose. It is critical to timely response when disaster does occur, and it also tells you what equipment and people are critical to your operation and need to be protected by redundancy.

You can think of preparing the information as

Identify...		Teams
Understand...	} … {	Clusters, their uses, (inter)dependencies, and users
Validate...		Potential loss, its probability & impact
Quantify...		Cost/impact of prevention/mitigation vs resolution/recovery

The more decisions that can be made in the course of data gathering, the easier the planning will be.

In the next section -- *Writing and Selling the Plan* -- we'll discuss two things, one whose importance is vastly overrated (the physical document) and one whose importance is vastly underrated (selling the plan to the constituencies).

Part 5 - Writing and Selling the Plan

By now, you have of course gathered all the information on current spaces, systems, people, priorities, and procedures and are all ready to write your plan. (… What, you haven't?) A disaster preparedness plan consists largely of organizing that information and making it readily accessible to you and your various teams.

Before you get to the "meat," there are some organizational elements that it is best to state explicitly up-front. You need to define the place of the plan in the institution, i.e. set the scope by saying what is and is not covered and define the relationship to any more general business continuity planning for other institutional functions. For technology, you need to specify support for emergency services, public safety, public information, and so on.

Once your scope is defined, you can get down to business. One organizational principal comes from the NIST Contingency Planning Guide for Information Technology Systems: think in terms of the phases you must go through after an incident.

- First, you need to let the appropriate people know that an event has occurred (Notification/Activation Phase).
- Next you need to get the priority systems running with what is available, which may mean temporary locations or facilities (Recovery Phase).
- Finally, you need to restore to (a possibly new level of) normal operations (Reconstitution Phase).

The information you gathered is critical to each of these phases.

Another way of organizing the material is functionally: policies, procedures, people (internal and external), and systems (internal and external). This approach will only make sense for a highly integrated or converged operation. If your Telecom/IT operations are not unified, you may want to structure your plan primarily with sections for each unit: telecommunications, networking, administrative systems, academic systems… One disadvantage of this stovepipe structure applied to your disaster preparedness plan is that it can conceal opportunities for cooperation. If your academic systems are not physically co-located with the administrative ones, then you have the opportunity for redundant sites by sharing both of the spaces for both functions. If spare systems (or systems capacity) are pooled between administrative and academic systems, you will be more resilient to incidents and have a more flexible set of response strategies than if the resources are segregated.

The data-gathering process discussed earlier amounts to turning over all your rocks and writing down what you found there. Inevitably there are going to be consequences:

- You are going to identify weak points where a simple loss could have disastrous consequences, e.g., data networks in a star topology, servers concentrated in a single location, all trunks in one cable. Either protecting those single points of failure or redesigning to create redundancy costs money without providing a concrete return on investment. So if you do nothing else, you must document those points so that you can address them later.
- The reasons you prioritized services are not going to be comforting for those whose functions get assigned a lower priority. When the music stops, they may be without a chair. Better to fight those battles now than during a disaster.

A disaster plan is a major project involving both one-time and ongoing commitment of resources. Data gathering, consensus building, and writing the plan are going to take resources and require an executive champion. So you will have to sell your plan to executive management and get buy-in for political and financial support.

Part 6 - A Living Disaster Plan (not Living with a Disastrous Plan)

By now you have probably realized that preparing a disaster plan isn't going to be easy. While the process will not be effortless, it is also not impossibly complicated either – just time-consuming. But only under very limited circumstances does having written the plan mean you are done. Think of the disaster plan as the documentation for an *ongoing* set of processes – the journey, not the destination.

The success of disaster response depends on good communication with internal and external constituencies, especially (but not only) ensuring that the emergency personnel (on- and off-campus such as telephone operators, public safety, fire/rescue, police, facilities and technical staff) can all reach and be reached by anyone involved. Don't forget to involve your human resources, legal, and public relations departments in both the planning and the response. The information gathered in the fourth article in this series about priorities, procedures, places, people, "parts", and their relationships all needs to be kept current. Whatever you thought worth including in your plan, it must be kept up-to-date to be of any use.

You wouldn't expect actors to perform in a play without rehearsal, so you also need to train all those involved to fulfill their roles in response to a disaster. You will need to train internally (team leaders, technology and support staff) and to some extent externally (public safety personnel, facilities personnel.) The training should be used to update the planning, and should itself be documented. (Train on the document and document the training.)

Like any muscle, a plan that isn't exercised, atrophies. So to avoid the worst embarrassments, your plan will need to be tested regularly. Testing should run the gamut from walking through a hypothetical scenario, to simulations, to (announced or unannounced) live exercises. Testing will reveal flaws in planning and areas requiring additional training.

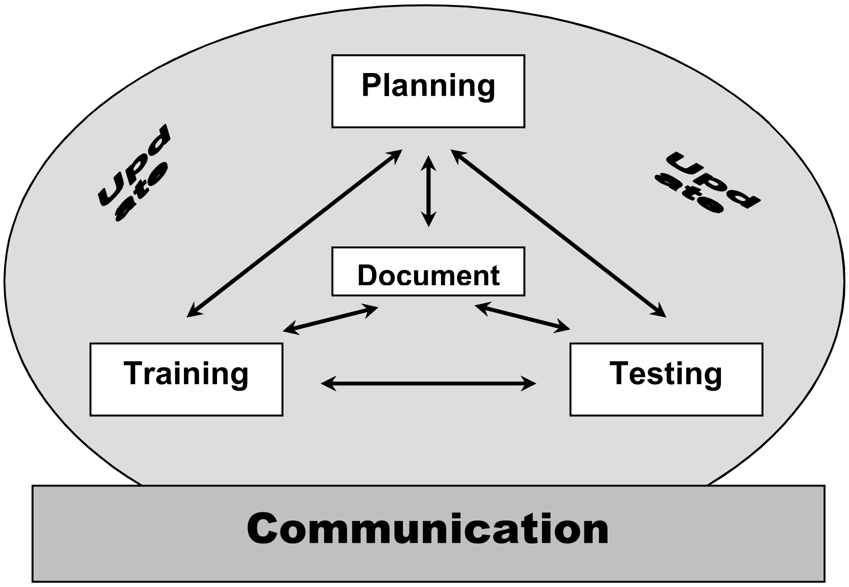

If your sole reason for developing a disaster plan was to satisfy your external auditors, then your plan can sit on a shelf to be dusted off annually for them. If you want the plan to be useful **in an emergency** (What a concept!), then it needs to be routinely tested, all participants trained, and everyone needs to be involved in the process. According to John Toigo in his excellent book, *Disaster Recovery Planning*

> "Even those who have successfully managed recovery efforts are quick to point out that their recovery strategies failed in as many parts as they succeeded. Almost invariably, they attribute successful recovery to luck, hard work, on-the-spot ingenuity and innovation, and God. While all believe that outcomes would have been very different had no contingency planning been undertaken, to a person they concede that their plans were – and could only be – imperfect."

As we said in part one of this series, there is more value in the **process** (planning – training – testing) than there is in the document. If you try to write a plan for every possible scenario and variation, you will fail! No matter what you plan for, the disaster will happen differently. The real value is have the information (people, places, stuff, …) readily available so make the correct decisions on-the-fly as the situation requires.

So, after reading this article you are scratching your head and saying, "I agree, but *NOW* what?", here are some next steps:

- Develop your champion – the person high up in the organization who will support the project and smooth the political rough spots for you
- Establish the role of communications and technology at your institution
- Determine the relevant business issues
- Ascertain the most critical factors for the institution
- Determine your institution's exposures should a disaster occur
- Put together the planning team
- Develop the desired end results
- Gather, sort, and organize data
- Test and train

However, it is not so important how you start but that you start. As Confucius said, "If one does not have long-range considerations, one will surely incur imminent afflictions."