◆ A Model for Quantifying Business Continuity Preparedness Risks for Telecommunications Networks

Ahmad Jrad, Thomas Morawski, and Louise Spergel

In the wake of recent events, a sense of awareness has emerged among businesses and government agencies of the need to continue their operations and provide critical services after a disastrous strike. The need for disaster recovery has always existed. However, the emphasis is shifting from reactive (recovery) to proactive (preparedness) to minimize damage from disasters and limit disaster impact through proper planning. We present a new model for business continuity preparedness (BCP) planning for telecommunications networks and a taxonomy for the quantification of the BCP readiness compared to similar businesses and industry practices. We categorize disasters as natural events, technical failures, and human threats. We discuss how each type of disaster is modeled and show how these models are used to quantify risks and potential impacts. We also show how the model is used to make specific recommendations to minimize exposure while maximizing the return on the investment in the BCP planning.

Introduction

In the wake of recent events, a sense of awareness has emerged among businesses and government agencies of the need to continue their operations and provide critical services during and after a disaster strike or unplanned event. The need for disaster recovery has always existed, yet there is much greater awareness of it now with the paramount importance that has been placed on homeland security. However, the emphasis is shifting from a reactive (recovery) approach to a proactive (preparedness) approach in order to minimize damage from disasters and limit their impact on the business.

Naturally, any serious discussion of disasters and unplanned events must consider in its scope the high-profile events reported in the media such as the September 11, 2001, attack on the World Trade Center, the August 2003 large power failure in the Northeastern United States, and major natural disasters. However, the scope of this paper goes well beyond those disasters and also includes smaller disasters, which must be included in any business continuity preparedness (BCP) planning effort, as they, too, may pose considerable risks to businesses. These smaller disasters and events include local power failures, lightning strikes, cable cuts, and vandalism. In fact, these smaller, more local disasters can have a more profound effect on a business than larger ones, as they may not affect the competition in the same way as

Bell Labs Technical Journal 9(2), 107–123 (2004) © 2004 Lucent Technologies Inc. Published by Wiley Periodicals, Inc. Published online in Wiley InterScience (www.interscience.wiley.com). • DOI: 10.1002/bltj.20029

wide-ranging disasters do. This is particularly true for telecommunications service providers because their business depends on their networks being operational 24 hours a day, 7 days a week, independent of any external occurrences. Their customers have very little patience with network downtime, particularly if the competition is still operational.

In this paper, we present a new model for BCP planning for telecommunications networks and a taxonomy for the quantification of the network BCP readiness. We show how to combine a mathematical model for the network with a model for the disaster events and use that model in assessing the risks of the various potential disasters on the network and the services that it provides. We explain how disasters are categorized as natural events, technical failures, and human threats and how disasters in each of those categories are modeled. We then show how those models are used to quantify the risk and the potential financial impact for each disaster at each location within the network. Finally, we illustrate how the model is used to justify specific recommendations to minimize exposure while maximizing the return on the investment in BCP planning.

Panel 1. Abbreviations, Acronyms, and Terms

AEC—availability environment classification
BCP—business continuity preparedness
CO—central office
FIT—failure in time (1 FIT = 1 failure in 10^9 hours)
MSC—mobile switching center
MTBF—mean time between failures
MTTF—mean time to fail
MTTR—mean time to repair/restore
PSTN—public switched telephone network
ROI—return on investment
UMTS—Universal Mobile Telecommunications System

Background

Business continuity preparedness comprises the set of processes by which a business prepares itself for disasters and unplanned events. These disasters, which are described in more detail below, include not only natural disasters such as earthquakes and hurricanes and high-profile disasters such as terrorism, but also everyday disruptions such as local power outages and service disruptions caused by events like plumbing failures. Since it is not possible to predict precisely the nature, timing, and severity of these disasters, it is necessary to be prepared for a variety of disasters that may occur.

BCP is different from disaster recovery in that BCP occurs before a disaster event happens, and disaster recovery occurs immediately after the event takes place. However, planning for the actual recovery is an integral part of the BCP process.

Most previous work on BCP takes a qualitative rather than quantitative approach. While this approach is valuable, it is limited, especially when it comes to making decisions on how to spend a limited BCP budget. One such qualitative approach is described in this section at a high level. We believe that this type of qualitative approach is both important and necessary. However, it is not comprehensive, especially when the business has assets such as a telecommunications network that need to continue to operate during the disaster. The following sections will describe a more quantitative approach, including a model of both the network and the disasters, which we believe is an important complement to the traditional qualitative BCP approach. Only when both of these methods are applied, can a business achieve the highest level of disaster preparedness for future events.

The traditional qualitative approach to BCP planning generally includes a life cycle with several stages. Although the number and names of these stages may differ, the approach is similar. An example of this life cycle contains the following six stages:

- Plan validation,
- Risk assessment,
- Business impact analysis,
- Plan design and development,
- Plan testing, and
- Plan maintenance.

It is only natural that the plan validation stage comes first, as one has to know the baseline state

before starting to analyze and proposing changes. In the risk assessment stage, the purpose is to learn as much as possible about the potential risks and what their impact is on the business. The model and methodology that we present in this paper generally focus on two stages: risk assessment and business impact analysis. Using the information from these two stages, the next stage involves creating the BCP plan. The final stages are the testing and maintenance of the BCP plan.

The BCP plan covers all aspects of the business. For example, a BCP plan includes human issues, such as how to maintain safety and evacuate employees during a disaster event, or how to reach employees in the event of a facility closure. Another area of high focus is the preservation of data. This includes ensuring that data is backed up in diverse locations in a timely manner suited for the particular business and that software and system configurations are kept in diverse locations and not necessarily in the immediate area where they are used. Other examples of areas that would receive a similar qualitative assessment are facilities, operating systems, application systems, and policies and procedures. What is not included here is an assessment of whether a network will remain operational and a quantitative way of determining how much benefit to the network can be achieved through mitigating actions relative to their cost. We believe that this additional quantitative approach is essential to complement the qualitative approach for businesses that operate a network (e.g., a telecommunications service provider) or that depend heavily on having a reliable network at all times.

Several other methods with some amount of BCP quantification have been previously used in the industry. One example is the work from the Harvard Research Group [4], which defines six availability environment classification (AEC) levels. The definitions of the AEC levels include the amount of downtime that can be tolerated and whether or not data is lost. The method by which the classification level of a particular server or function is determined is not specified, but such a determination may be made by qualitative inspection, questionnaires, observation, or another method specific to the system in question. Another method that has been used involves surveys of personnel within the business and benchmarks of their scores against industry-specific practices. Although both of these methods assign scores, they are indeed largely subjective, and neither provides any way to model complicated systems or networks such as telecommunications or data networks. In addition, neither includes a method of quantifying the financial implication of disaster impact in terms of equipment damage and service downtime.

The Need to Plan for Disasters

Why plan for disasters? The simple answer is that sooner or later all businesses are eventually bound to experience some kind of a disastrous event. They can either try to be prepared to deal with the event or be caught off guard and suffer the consequences. In critical industries, such as telecommunications, the importance of being prepared takes on added significance because people and businesses rely on the services provided by these industries to meet their everyday needs, especially during disaster strikes.

Disasters happen, usually with disastrous consequences. Below we list a few examples of disaster events experienced by telecommunications companies [1]:

- May 1988—A fire destroyed a central office (CO) in Hinsdale, Illinois, causing complete and massive loss of service to customers and affecting thousands of businesses and their employees. In addition, both Chicago airports lost crucial communication lines to the Federal Aviation Administration for one day, resulting in flight delays [5]. It is fortunate that the affected customers themselves were not exposed to the disastrous event (i.e., the fire), but we can easily imagine a scenario where loss of service could mean the difference between life and death.
- September 1989—Hurricane Hugo ravaged the city of Charleston, South Carolina, destroying everything in its path. This led to massive loss of telecommunications services, ironically due to the loss of electrical power that resulted from the hurricane strike.

- October 1989—A major earthquake shook the city of San Francisco, California, and caused massive devastation to the infrastructure. Fortunately, loss of telecommunication was limited and isolated to areas that suffered from cable cuts.
- October 1990—Illinois Bell suffered a major cable cut. This led to massive disruptions to financial institutions; radar services were also disrupted.

The Need to Quantify Disaster Events

The approach we describe in this paper is indeed applicable to many industries; yet it is best suited, and was initially developed, to focus on the unique needs of the telecommunications industry. The unique aspect of the telecommunications industry to which we refer is the fact that a telecommunications company gets its revenue from a network that is expected to operate 24 hours a day, and generally does so automatically. Furthermore, depending on the type of disaster, parts of the network may be utilized at a higher rate during the disaster. This was true during several recent disasters in the United States including the San Francisco Earthquake, the September 11 attack on the World Trade Center, and the more recent power failure in the Northeast. Therefore, it is most critical for telecommunications companies to ensure that their networks continue to provide services during a disaster strike. However, making a network fully redundant, both logically and geographically, may be too costly and too technically difficult to achieve. Therefore, it behooves the telecommunications companies to know which improvements would best protect their networks and provide the highest possible level of availability during likely disaster strikes, while balancing these improvements with their respective costs.

The equipment used to provide the telecommunications services can be very costly to repair or replace. Therefore, telecommunications service providers generally want to make improvements that can minimize damage to this equipment. Often, these improvements may be disaster specific. For example, the actions needed to protect against a future flood would not necessarily protect against an earthquake and vice versa. Since service providers cannot afford to protect against every type of natural disaster, it is most useful to know which disasters pose the highest risks in their respective regions of operation and the expected magnitude of the impact of these disasters on their networks.

The model described below quantifies the probability of occurrence of disasters and the expected impact that they would have in terms of both lost revenue due to network downtime and the expected cost to repair and replace damaged equipment. This is then used to calculate the loss exposure for each disaster at different sites in the network. Mitigation approaches for disasters and regions with high loss exposure are then examined in light of the cost to implement and the expected return, in financial as well as overall network availability measures. The disaster impact information produced by this method can be effectively utilized to make intelligent and cost-effective decisions on how to allocate limited BCP planning resources.

Description of Model

The model begins by quantifying the reliability of components in networks based on parameters such as mean time between failures (MTBF) and mean time to repair (MTTR). These parameters are then aggregated to assess the reliability of networks comprising multiple components.

We then add probabilistic models of the disasters that may affect the components in these networks based on their function and their vulnerability to the various threats.

Based on the combined analysis of normal component downtime and vulnerability to disaster threats, we quantify the level of preparedness for the overall network by defining and measuring the following parameters:

- Aggregate risk of disaster impact based on the individual risk to individual network components.
- The likely impact of a disaster event to the existing network. This is based on the individual impact on each network element in terms of physical damage to the equipment as well as service interruption downtime.
- A better overall network availability number that takes into account the network availability under

normal operating conditions, as well as the expected downtime due to the expected disaster occurrences that will likely impact the network operations.

- A cost-based analysis of alternative mitigation strategies that takes into account their expected costs and benefits on the overall network availability.

Network Models

In order to assess the availability of a network and the impact various events have on it, we first develop an encompassing view of the network and its individual network elements. We accomplish this by examining the network though three distinct lenses. Each of these views presents some aspects of the network that, when combined, lead to an overall understanding of the network in its totality.

To illustrate how this is accomplished, we take a simple wireless network and show how it is modeled. This network is a simplified version of a real, existing wireless network. The choice of a wireless network is inconsequential; we could easily show the same analysis given any communication network such as a public switched telephone network (PSTN), a data network, or an optical network.

The first view of the network is called the *building blocks view*. As indicated in **Figure 1**, this view shows how the various network components are interconnected and which components interact in establishing and carrying out network services.

The second view of the network shows the actual geographic layout of the network elements and is referred to as the *map view*. As we see in **Figure 2**, the map view gives a tremendous amount of information

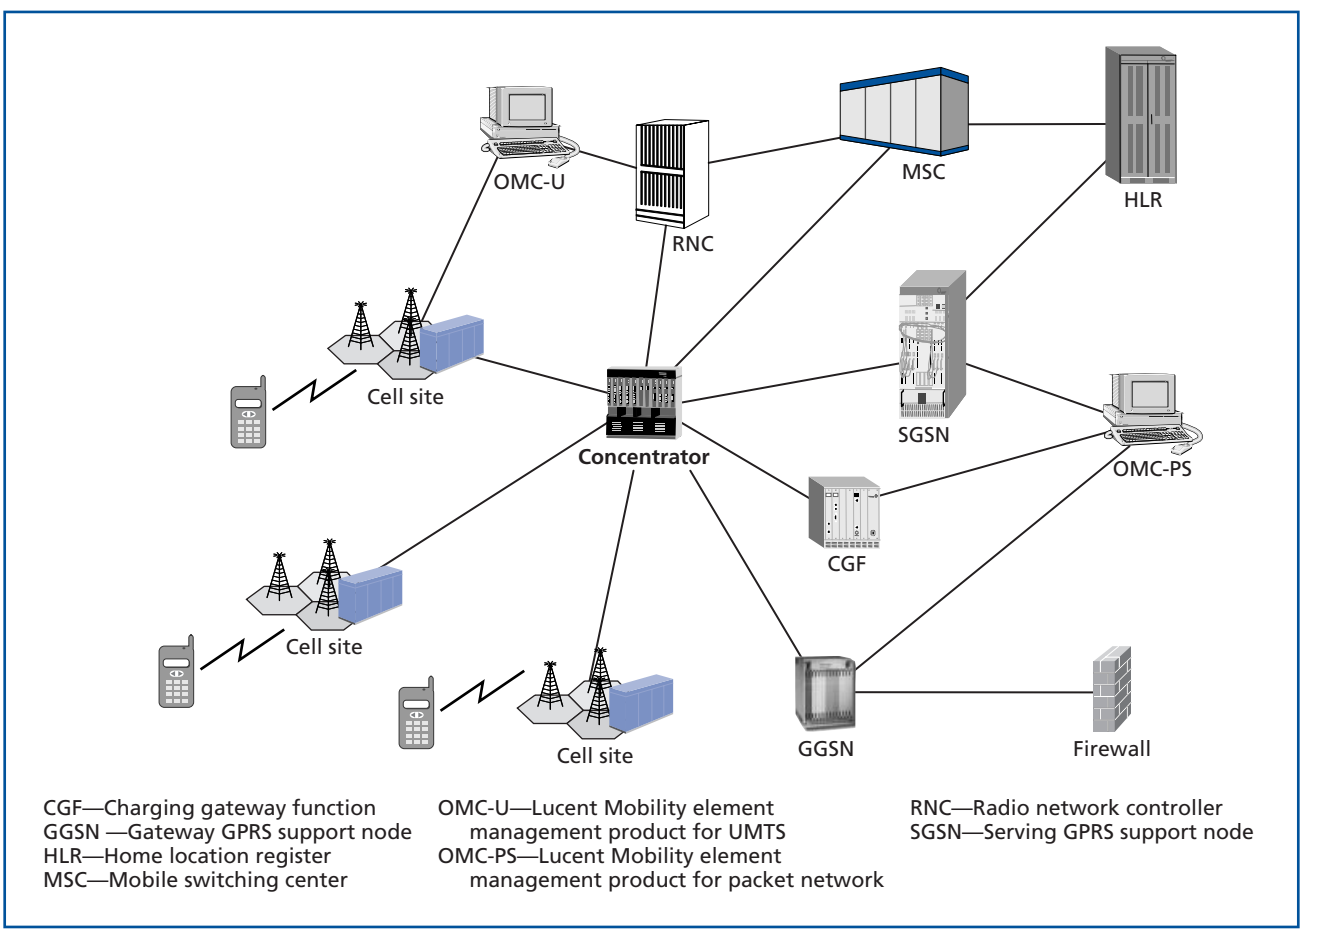

Figure 1.
Network view: building blocks view.

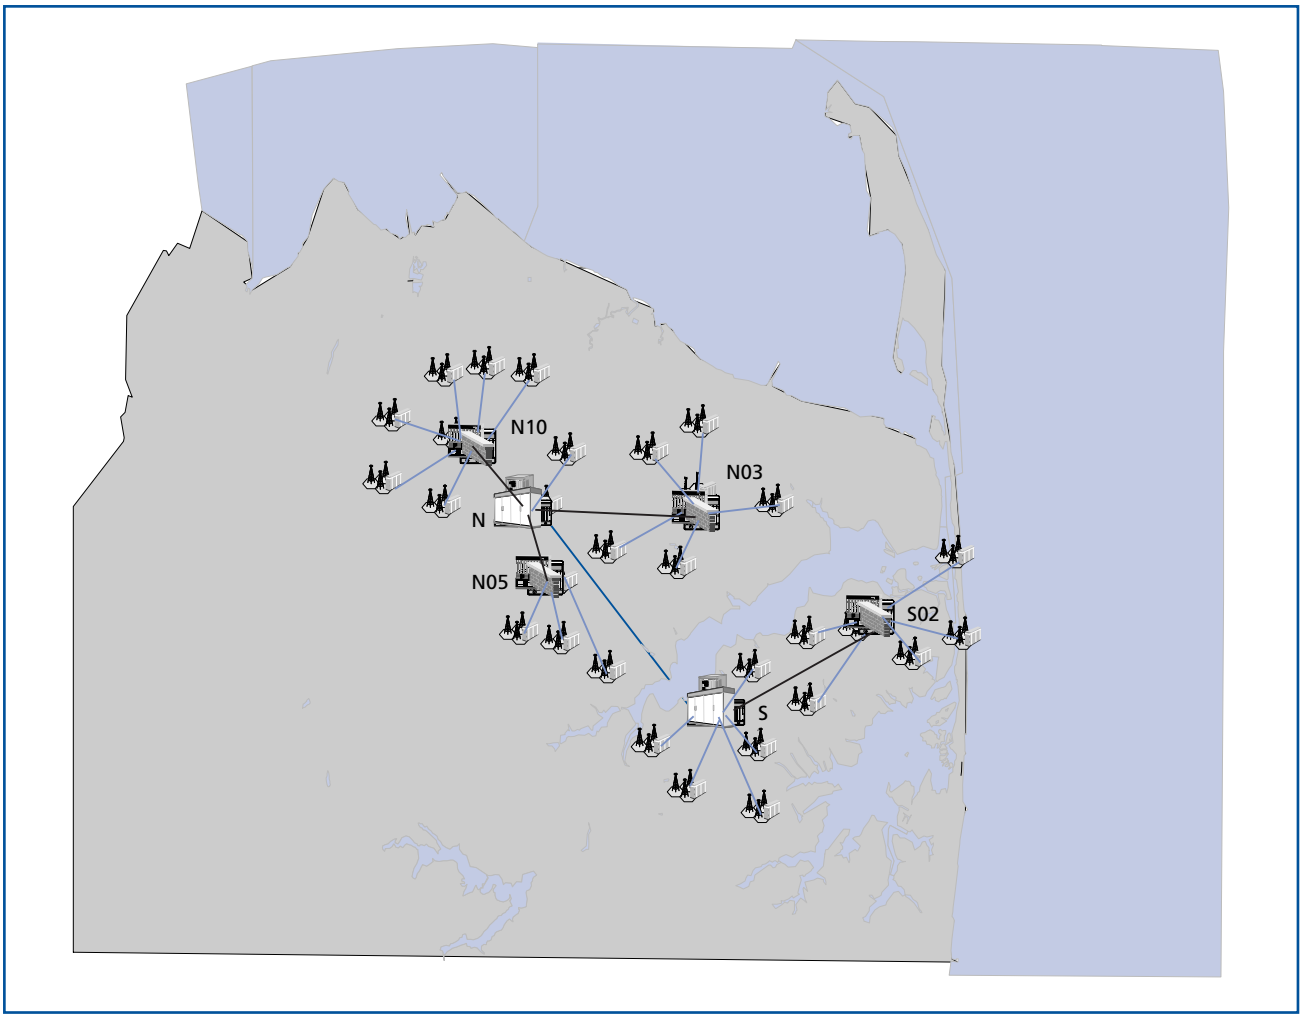

Figure 2.
Network view: map view showing site names.

that plays an important role in determining the types of disasters that may affect the network as well as how these disasters are best modeled for the various network elements. For instance, by looking at the map view, we are able to quickly determine which parts of the network are in the proximity of bodies of water. This knowledge is useful in modeling flooding as well as hurricanes and other disasters specific to coastal areas.

In Figure 2, we show the map background as a plain background for easy viewing; however, by superimposing layers that add features to the map view, we are able to determine the elevation information as well as the proximity to certain targets that may be at a higher risk with regard to human threats. These targets may include airports, nuclear power plants, and other high-risk infrastructure.

Figure 2 also shows the site names for the various sites in this network. The two major sites are named N (for North) and S (for South). Sibling sites are named by taking the host site name and following it with a number that corresponds to its position (with respect to the host) on a clock dial. Thus, N03 is the site at roughly 3:00 o'clock from the host site N.

The third view of the network serves as a basis for determining the overall network availability. We refer to it as the *call path view*. In order to evaluate the network availability, we examine the network in a service-centric fashion by following the path of a particular service

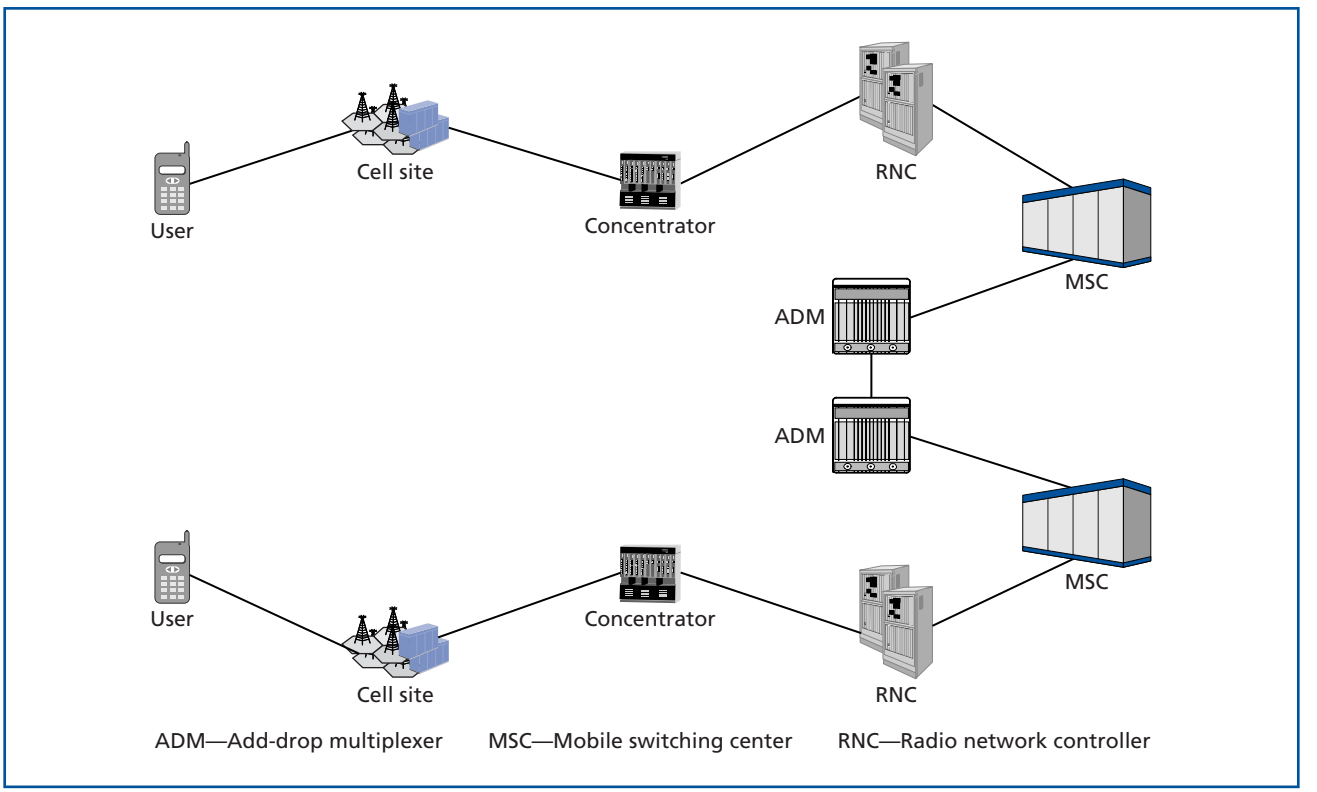

Figure 3.
Network view: call path view.

within the network. For instance, given our wireless network, if we consider a call from one wireless subscriber in the network to another wireless subscriber in the network where the two subscribers are served by distinct mobile switching centers (MSCs), then the end-to-end network view is shown in **Figure 3**. In most cases, the network will have more than one call path view. Typically, there is one call path view per type of service offered. As a result, the network availability will vary subject to the type of service that is being considered.

Disaster Types

We classify different disasters according to the way in which they occur and their overall impact. Accordingly, we group the disasters into three distinct categories:

- Natural disasters,
- Technical failures, and
- Human threats.

Natural disasters. These disasters occur naturally throughout the world, and they are largely driven by geographical location and the natural environment. Some examples of natural disasters are hurricanes, tornados, floods, earthquakes, and tsunamis. In any given region, it is well known which disasters occur, and the history of these is often well documented.

Technical failures. Disasters may result when there is a failure associated with some technology we depend upon (e.g., a failure associated with power equipment leading to a power outage) or when there is a failure of some technology that causes an environmental threat (e.g., a failure of equipment leading to chemical leaks or to nuclear radiation from a power plant, as in the Chernobyl disaster of 1986). The probability of occurrence of these events is determined based on a combination of historical events in confluence with the current conditions of the technology in question.

Human threats. These disasters are typically characterized by a purposeful act to disrupt or cause damage. A prominent example of this type of disaster is the September 11 attack on the World Trade Center. Although the network under study may not be the intended target, it may well be impacted by a nearby

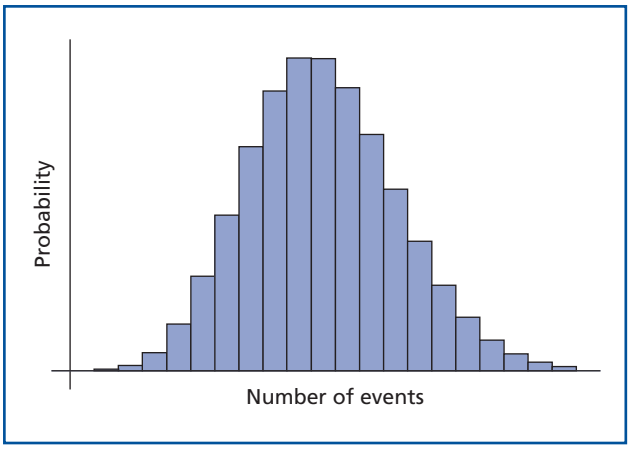

Figure 4.
The Poisson distribution.

occurrence. This category also includes such acts as sabotage, vandalism, and anti-corporate attacks.

Disaster Models

For simplicity, we initially assume a normal Poisson distribution [3] for both natural disasters and technical failures as shown in **Figure 4**. We also assume that various disasters occur independently of one another and that the rate of past occurrences is a good indicator of the expected future behavior.

As mentioned above, this assumption is made initially; however, the disaster models, especially the ones for the technical failures, are then individually adjusted based on the knowledge of any known relevant parameters. For example, in the case of power failures, the length and duration of power failures over the past few years is initially used as a predictor for the rate and duration of future power failures. However, the future rate of occurrence is then adjusted to account for known changes, such as upgrades installed by the power company or additional load on the local power grid due to new construction. This is then further adjusted to account for power failures that may result from other natural disasters such as adverse weather conditions, as appropriate.

In Figure 4, we assume that the probability of a disaster occurrence, Pd, which defines the probability of occurrence of a given disaster (Dj) a given number of times (a) during a time period (T) at a given network element site (Ei) is based on two essential parameters. The first, Rd, is defined as the historical rate of occurrence of the given disaster (Dj) in the given area of the network where the element (Ei) is located. The second parameter, Ph, is defined as the rate of a direct hit of the disaster (Dj) on the precise location of (Ei). For instance, the rate of occurrence of a tornado for a given city may be established based on historical data; this is what we call Rd. However, the likelihood of a given tornado striking a particular block or building within the city is then defined as Ph. Assuming that all neighborhoods within the city have equal likelihood of being struck by the tornado, Ph may simply be computed as the area of the neighborhood divided by the area of the city.

Now given Rd and Ph, the probability of occurrence Pd is computed as:

$$Pd_{(Ei,Dj)}(a,T) = e^{-(Rd_{(Ei,Dj)}*Ph_{(Ei,Dj)}*T)}\left[\frac{(Rd_{(Ei,Dj)}*Ph_{(Ei,Dj)}*T)^a}{a!}\right].$$

We also define the probability of having at least one occurrence (but potentially many more) of a disaster during the time interval T. This probability can be computed as:

$$Pd_{(Ei,Dj)}(1+, T) = \sum_{k=1}^{\infty} Pd_{(Ei,Dj)}(k,T),$$

or simply as:

$$Pd_{(Ei,Dj)}(1+, T) = 1 - Pd_{(Ei,Dj)}(0,T)$$
$$= 1 - e^{-(Rd_{(Ei,Dj)}*Ph_{(Ei,Dj)}*T)}.$$

Disasters in the category of human threats are treated in a different manner. For these types of disasters we assume a small but finite probability of occurrence and, instead of computing that probability to the last digit, we focus our interest primarily on the impact of such occurrences on the network and facilities.

We make no assumption on the distribution and no attempt to quantify in any precise manner the likelihood of occurrence. Such events are usually driven by extremely complicated factors that are difficult to quantify in a precise way. We also believe that the likelihood of occurrence changes from time to time, not unlike the threat assessment level that is defined by the U.S. Department of Homeland Security. As we

have seen, that threat assessment level can change on a daily basis. However, we do use an assumed rate of occurrence largely based on input from local authorities and a subjective assessment of the network area and its environs.

Yet, given that the rate of occurrence is not zero, we make assumptions about the nature of the potential threat and attempt to quantify the possible and likely levels of impact. We believe that such an assessment is helpful in making the network better prepared, irrespective of the precise likelihood of occurrence.

Quantifying the BCP Readiness

Once the model has been developed for both the network and the disasters of interest, it can be used to generate a number of BCP indicators, which can be used in guiding the overall BCP effort and focusing the BCP planning in the most productive directions.

We take the sample network described earlier, and we model various disaster occurrences on this network to illustrate how the concept is applied to the real world. We consider four natural disaster types, one technical failure, and one human threat. We depict the various disasters using representative icons for easier representation (see **Figure 5**).

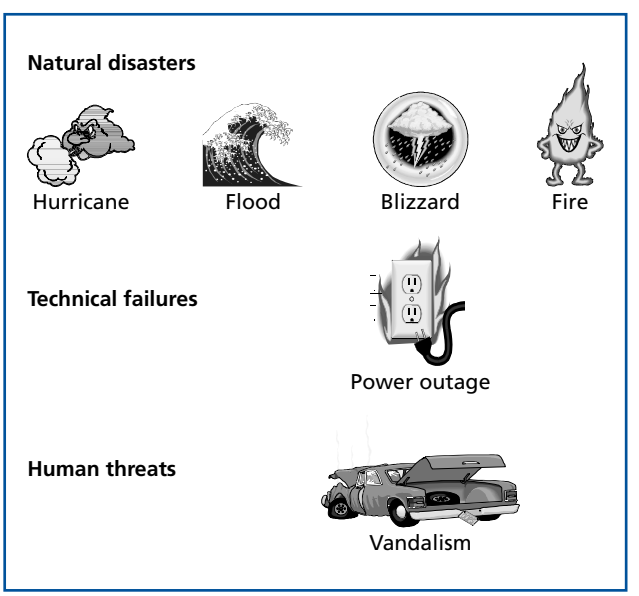

Figure 5.
Disaster icons.

Assessing Risks and Impacts

The first type of output we will discuss is the risk-impact chart. This scatter plot shows the probability of a disaster striking the network (risk) on the horizontal axis and the impact of this disaster on the vertical axis. The risk is defined as the probability of the disaster occurring at least once over a given period of time. This period of time can be any number of years, although five years is the most common as it represents a typical planning cycle for a telecommunications service provider. The impact is defined as the expected value of the sum of lost revenue due to service downtime and cost to repair or replace damaged equipment as a result of a single disaster occurrence. The impact is measured in monetary units such as dollars.

Initially, we compute the risks versus impact numbers for the various network elements on an individual basis. To do so, we assume that individual network elements are struck by a disastrous event on their own. Once those numbers are computed, then we can discuss the parameters of a real-world disaster strike.

We define two regions for each hypothetical disaster strike: a physical region and a logical region. The physical region of impact is defined as the area where the immediate impact of the disaster takes place. We assume that all equipment in that area may be damaged by the disaster, although not necessarily destroyed. Some equipment may be partially damaged, and other equipment may escape harm completely, depending on the parameters that are defined to govern the disaster occurrence. For instance, in a flood event, any equipment located on higher floors within a building will be assumed to have less risk of direct water damage from the flood than other equipment on the ground floor.

The logical region of impact defines the scope of the network that may suffer loss of service. So while some equipment outside the immediate region of the disaster strike will escape physical damage, that equipment may be subject to downtime. This may result when such equipment is dependent on other equipment in the area of the strike that may have sustained damages.

For our case, let us consider the equipment in one region and for one disaster, flooding in this case.

Table I. Probability of flood damages to individual network elements.

Element	Risk	Damage cost	Service cost
PSA×4	1.4598%	$ 14,400	$560,000
B5-PSA×4	1.4598%	$ 34,000	$900,000
B4-PSA×4	1.4598%	$ 34,000	$900,000
B3-PSA×4	1.4598%	$ 34,000	$900,000
B2-PSA×4	1.4598%	$ 34,000	$900,000
B1-PSA×4	1.4598%	$ 34,000	$900,000
Site N03	1.4598%	$184,400	$900,000

We show in **Table I** the actual probabilities and associated costs of flooding on the equipment within that region. The table shows that the overall risk to region N03 is the average of the risks to the network elements and the total damage costs are the sum of the costs to the individual elements. The downtime cost is the maximum of the costs that are caused by each of the element's failures.

Similarly, these numbers are computed for all the regions within the network, and then the data is aggregated to show a complete flooding risk and cost assessment, as in **Table II**. For simplicity, in this example we consider site N05 as part of site N with respect to flooding. This would imply that the two sites are close enough and geographically linked such that a flood event in the one site would simultaneously affect the other site as well.

Table II shows how the overall disaster risk is the weighted average of the individual site risks, and the weight we used in this example is the number of major network elements at each site; however, other weighting factors can also be used. Therefore, the overall risk of flooding is roughly 11%. This can be seen also by looking at **Figure 6**. The overall potential cost of a disaster strike is computed as the maximum sum of the cost of damage to equipment plus the cost due to service interruption over the various regions that are subject to the disaster. In the case of flooding, the maximum cost is computed as:

Maximum flood strike cost is for region N,
$Cost_N$ = $675,640 + $900,000 = $1,575,640.

Figure 6 shows a sample risk-impact plot for all disaster types within our study and for the entire network under study. This shows the worst case (or it can also show the weighted average) risk and impact for all the sites in the network. This begins to give a good indication of where the significant risks are, both in terms of likelihood of occurrence and in terms of expected potential damage caused by a disastrous event.

In addition to the plot of Figure 6, it is also possible to zoom in on individual possible disasters for additional analysis. By computing the risks for each site separately, we can examine the risk-impact by site and by disaster type. We are also able to show the individual disaster impact on a network element basis. This helps us to determine where disaster mitigation

Table II. Computing the overall disaster risk and cost.

Site	Weight	Probability of occurrence	Cost of equipment damage	Cost of loss of service	Expected loss value
N	18	0.014598	$675,640	$900,000	$ 17,518
S	12	0.229958	$582,114	$900,000	$275,949
N10	8	0.014598	$252,400	$900,000	$ 13,635
N03	6	0.014598	$184,400	$900,000	$ 13,635
S02	6	0.406744	$184,400	$900,000	$379,899
Overall		0.11334192	$675,640	$900,000	$700,636

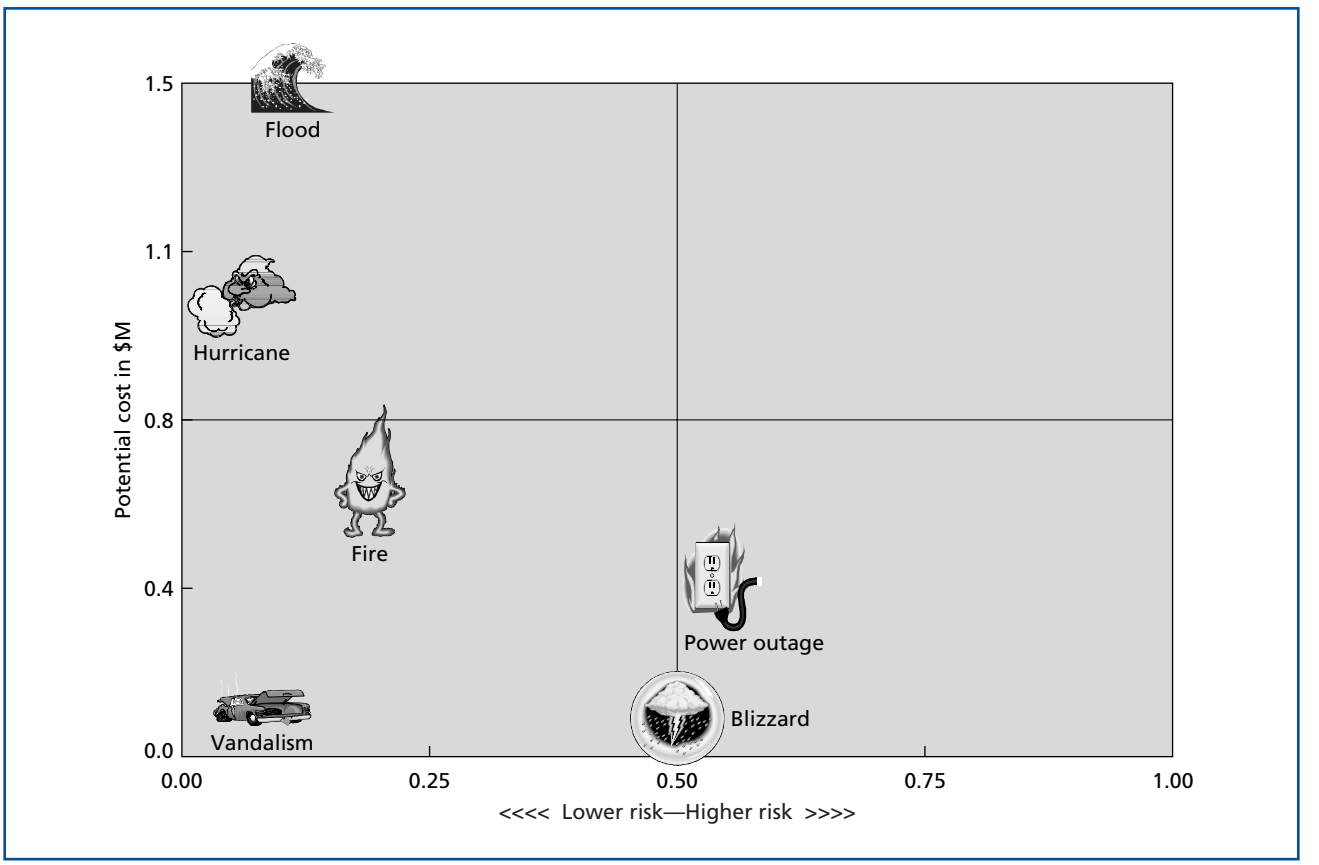

Figure 6.
Risk versus impact by disaster type.

is most needed. **Figure 7** shows a risk-impact plot for a single disaster type (in this case, flood) for each site in the network.

Assessing the Loss Exposure

From the risk-impact data, we calculate another parameter, called the loss exposure, as the product of the risk and the impact. This shows the actual expected disaster cost over the time interval under study. By adding the loss exposure for all the sites in the network, we are able to deduce the overall expected cost of the various disasters to the network under study. **Figure 8** shows the loss exposure for the data in the risk-impact plot shown earlier. The loss exposure is broken down to separately indicate its two components: lost revenue due to network downtime and cost to repair or replace damaged equipment. Figure 8 also shows that the total expected loss due to all considered disaster occurrences is statistically expected to reach $3.174M over the next five years.

As with the risk-impact plot, the loss exposure calculations can be broken down by site or further broken down by component.

Incidentally, this analysis takes into account the possibility that multiple occurrences of the same disaster may strike the same site. This is especially true when the time interval chosen is quite long. To accomplish this, a recursive algorithm is employed to compute the sum of the costs of disaster occurrences over the site. The algorithm is repeated until it converges.

Assessing the End-to-End Network Availability

We can determine the end-to-end network availability by first computing the expected downtime of the various components in a given call path view under normal operating conditions. We do this by taking the reliability data for the network components.

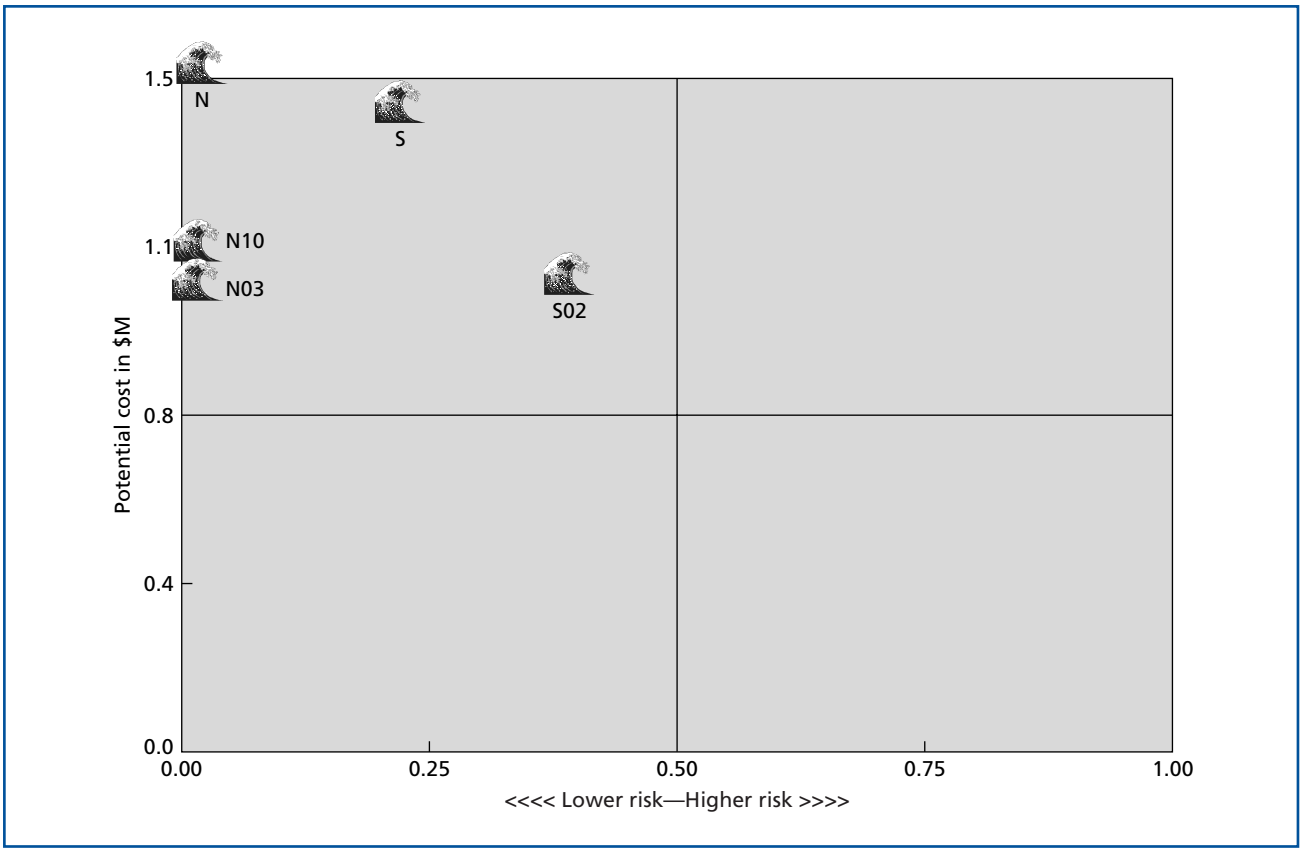

Figure 7.
Risk versus impact of flooding by network site.

This data includes the mean time to fail (MTTF) and MTTR in case of failure. Given MTTF and MTTR, we compute the MTBF as:

$$MTBF = MTTF + MTTR.$$

We also compute the failures in time (FIT) rate as:

$$FIT = \frac{10^9}{MTTF},$$

where 1 FIT is defined as 1 failure in 1,000,000,000 hours of operation.

In addition, we consider the cases where there are active standby components available as well as the cases where parallel paths are available in case of some network element failure. For all cases of single and parallel subsystems, we use an extension of the Markov model presented in Telcordia SR-1171 [6].

In case of a redundant system, we assume an N + 1 system where N is the number of active units (see **Figure 9**). We compute the downtime of each subsystem on its own using the continuous-time Markov Chain [2]. Then we combine the individual downtimes to derive the end-to-end downtime of the network as a whole.

After determining the expected downtime under normal conditions, we adjust the computation to include the expected downtime due to disaster occurrences. This is computed by examining the downtime per occurrence per disaster strike on each site in the network. This downtime is multiplied by the probability of occurrence and the algorithm is again recursed over possible multiple occurrences.

By adding the downtimes above we are able to determine a more realistic number for the actual network availability.

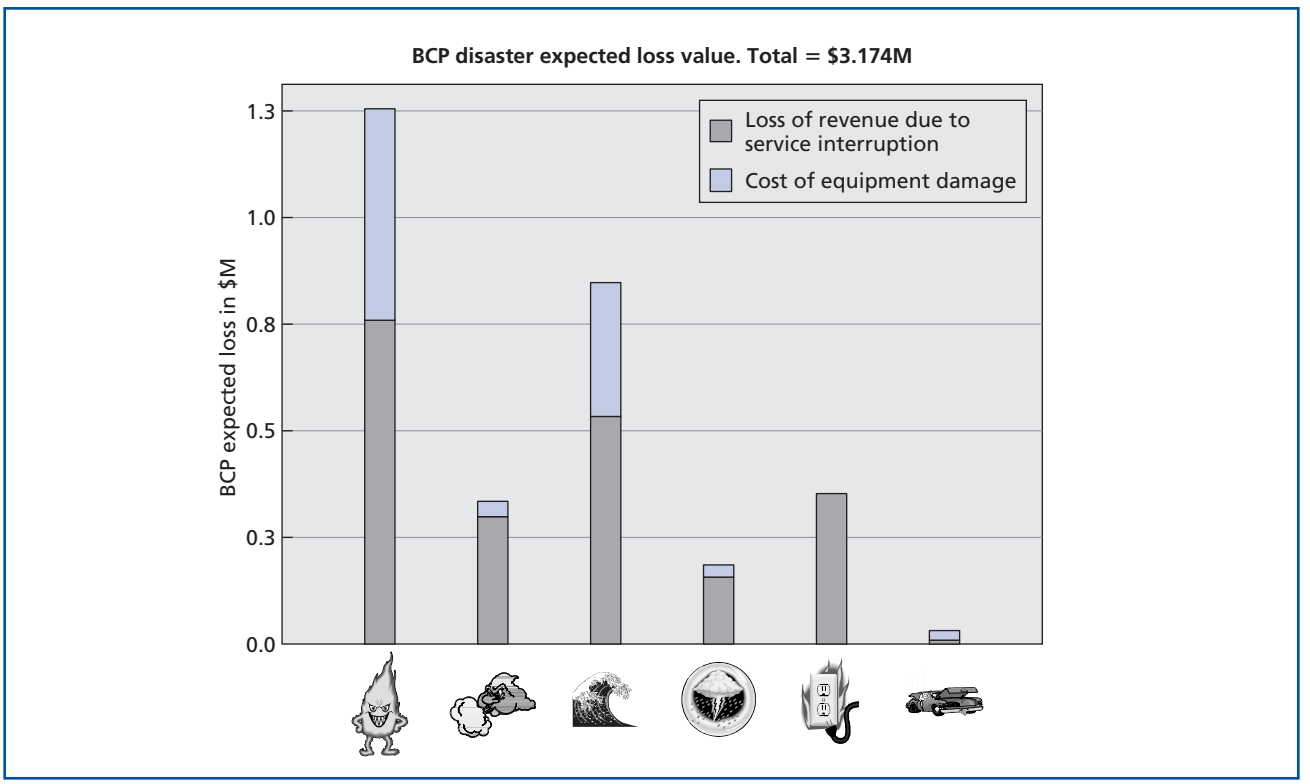

Figure 8.
Expected loss value by disaster type.

Mitigation Approaches

From the data in the risk-impact plot and the loss exposure calculations, we can determine which disaster-site or disaster-component combinations would be good candidates for mitigation. Mitigation proposals generally fall into one of two categories:

- *Network redundancy approaches*. With these approaches, the network is made more redundant in certain high-risk or single point-of-failure areas. The advantage of these approaches is that they will allow the network to continue providing service during a variety of types of disasters. The disadvantage is that they do not reduce the probability of damaged equipment that will need to be repaired or replaced.
- *Disaster-specific approaches*. These approaches generally focus on keeping equipment operational and preventing damage to equipment from a specific type of disaster. These approaches are most useful for natural disasters, which are well understood and occur in a somewhat predictable manner. For example, a disaster-specific approach to earthquakes is to make sure that earthquake bracing according to the recommended guidelines is implemented in all offices in an earthquake zone. An approach to floods would be to raise all equipment off the floor or to move it to a higher floor within the building.

In most cases, the mitigation approach will reduce the impact of a disaster while the probability of the disaster will remain the same. However, in some cases, it is possible to reduce the probability of the occurrence of the disaster. For example, if the disaster of concern is a power failure, the probability of occurrence may be reduced by switching to a more reliable power company or by improving the wiring inspection practices so that the wiring is inspected and repaired more often.

The mitigation approaches are often technical in nature and should be proposed by someone familiar with the network technology and/or the given disaster type. In some cases, it is advisable to propose more

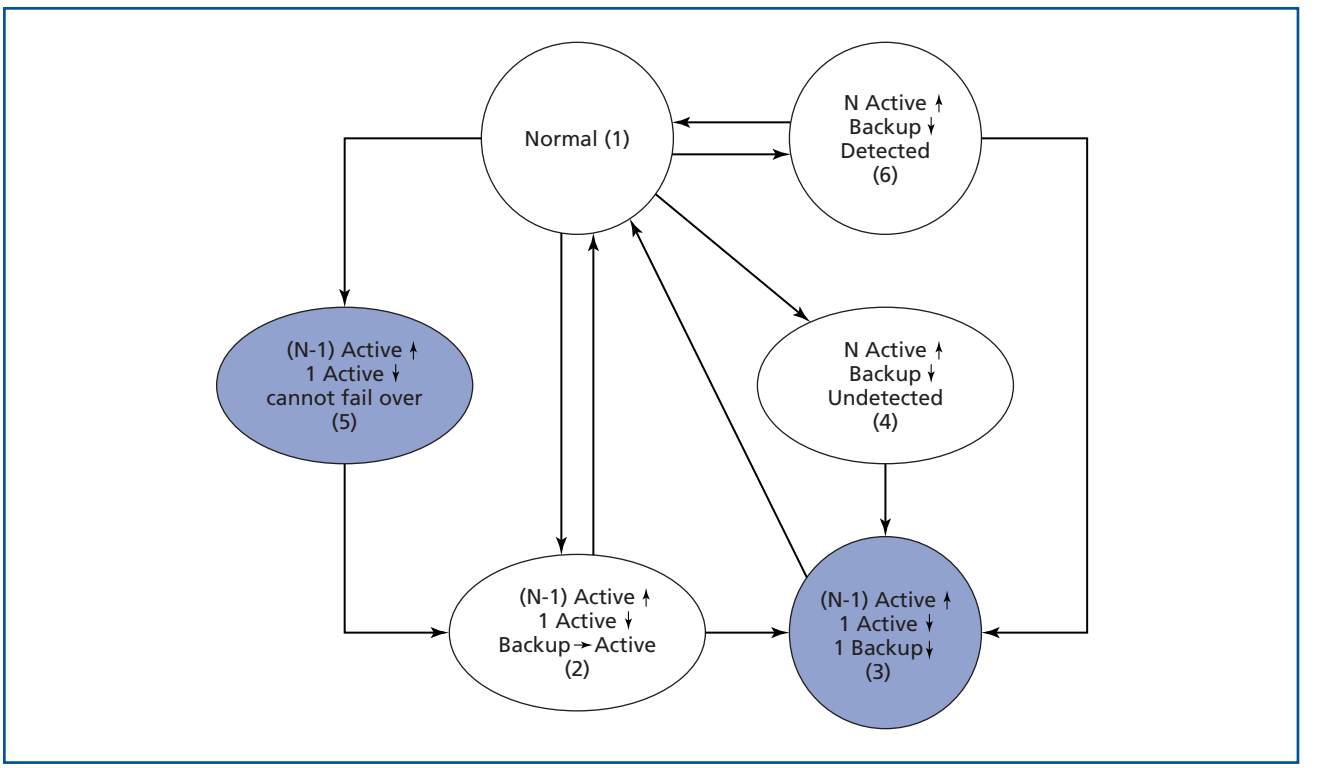

Figure 9.
Markov model for a N+1 protected system.

than one mitigation strategy for a given disaster-site combination with a high loss exposure. When proposing mitigation strategies, it is also necessary to investigate the cost of implementing each strategy.

For each mitigation strategy that is proposed, the new parameters are entered into the model and the new risk, impact, and loss exposure are calculated. These can be compared with the cost to implement, and a simple cost-benefit analysis can be performed.

For our example, **Table III** shows the loss exposure over a five-year study window for the highest six disasters within the network and the cost of several mitigation proposals. By mitigating for these particular disasters, we are able to achieve a better return on investment (ROI) on the BCP dollars that are spent.

A sample set of mitigations along with their expected costs is shown in **Table IV**. Implementing all the mitigations shown would require an estimated

Table III. The six costliest disaster strikes.

Event name	BCP site	Probability of occurrence	Cost of equipment damage	Cost of loss of service	Expected loss value
Flood	S02	0.406744	$184,400	$749,600	$379,899
Flood	S	0.229958	$582,114	$617,884	$275,949
Power outage	N	0.527633	$0	$400,000	$211,053
Hurricane	S02	0.228697	$ 90,400	$826,601	$209,715
Fire	S	0.178698	$408,514	$101,487	$ 91,136
Fire	N	0.170971	$457,240	$ 52,760	$ 87,195

Table IV. Possible disaster mitigation actions.

Disaster	Recommended action	Estimated cost
Fire	Add fire resistant walls	$ 40K
Flood	Move to higher floor	$160K
Hurricane	Fortify structures and add wind breakers	$ 40K
Power outage	Add battery backup to deficient network element	$ 10K
Total		$250K

total investment sum of $250K. However, we can compute the expected loss value for the network after the mitigation is implemented by adjusting the disaster models to reflect the mitigations. The resulting loss exposure of the network is shown in **Figure 10**. As we can see, the overall expected cost of the network disasters is reduced to $1.9M. This is a saving of $1.2M in return for the BCP investment.

In the example above, we have chosen specific recommendations for each of the disaster areas on its own merit. However, oftentimes certain mitigations, such as adding redundancy to some network elements, would improve more than a single disaster event. This must be taken into account when computing the post mitigation expected loss on the network. It is also possible that some mitigation actions that have a positive impact on certain disasters may have a counter and negative impact on other disaster events. As a simple example, adding a sprinkler system to a facility would have a positive impact on mitigating fire occurrences; however, sprinkler systems are known to malfunction, and therefore that may raise the risk of water damage and flooding. In that case, other flood mitigating actions may have to be taken to counter that possible increase in risk. Alternately, a dry, gas-based fire extinguishing system, which could be much more appropriate based on the equipment involved, would solve the possible negative complications entirely.

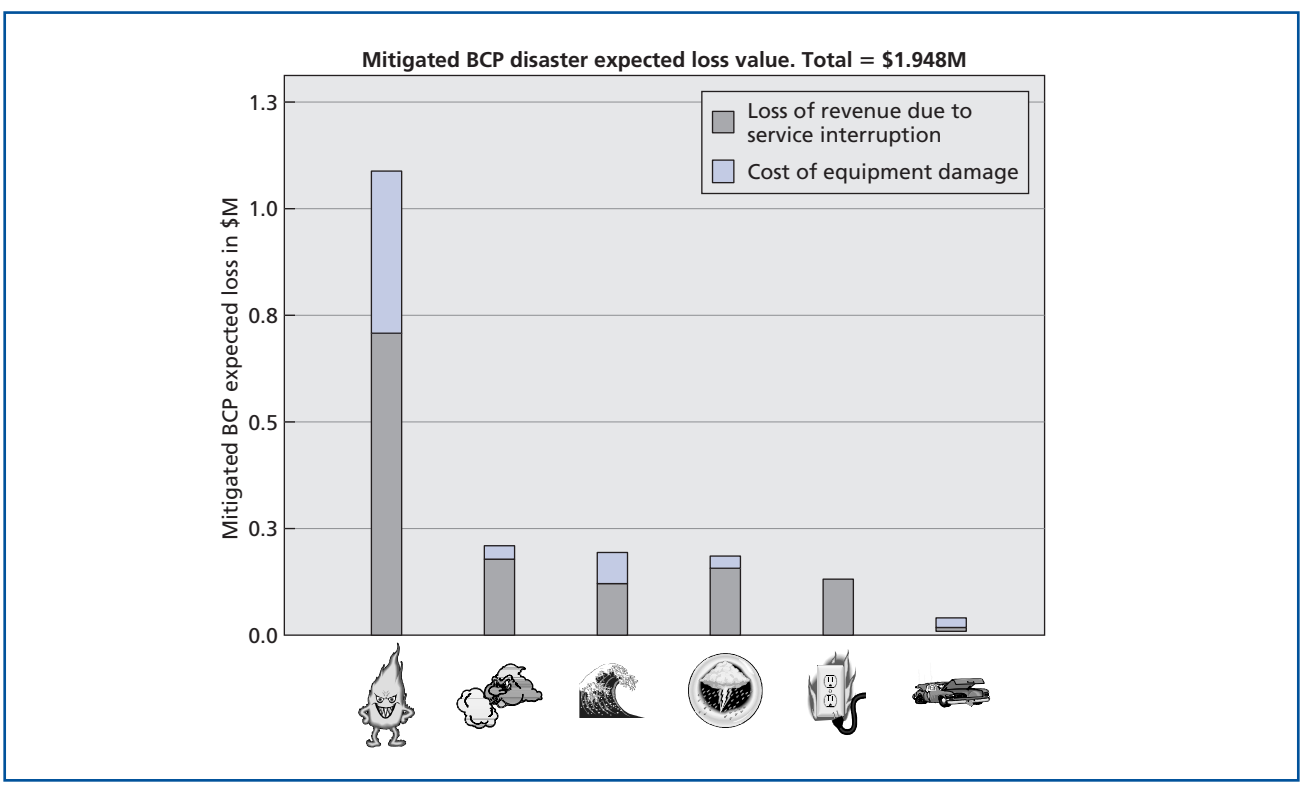

Figure 10.
Post-mitigation expected loss value.

Conclusion

Disasters can occur anytime, anyplace. Their nature and timing cannot be precisely predicted, yet they can be planned for. Indeed, and in the spirit of homeland security, they must be planned for. Businesses that are prepared for disasters and unexpected events generally fare better than those without a plan. In particular, businesses that operate networks, such as the telecommunications service provider business, need to plan for their networks to continue operation during the disaster in addition to the traditional business continuity planning that all businesses need to do.

In this paper, we have described a model for both the networks and the potential disasters as well as a methodology that is used with the models to help telecommunications service providers more effectively plan for continuous operation of their networks during a disaster. We have shown how the model can be used to quantify the probability of a disaster strike to any given location within the network and the financial impact that the disaster would be expected to cause. The model can be used, in addition, to quantify the improvement to these parameters that each proposed mitigation strategy would allow. We have shown how this information, combined with the expected cost to implement each mitigation, will give service providers the knowledge that is needed to most intelligently decide how to spend their BCP budgets.

It is possible to compute the various aspects of information about disasters and their potential impact manually; however, that would be quite laborious and expensive to do. Therefore, we have developed a specialized tool that automates much of the processing involved.

Possible future extensions of this work may focus on extending the model to cover other aspects outside the network infrastructure and services. This would likely include the impact of disasters on business processes and operations and how to quantify the cost of that impact and the benefits that may be achieved through alternate mitigation strategies.

Acknowledgments

We would like to acknowledge Jay Etris, Blesson Mathews, and Gerry O'Reilly for their help in the work that led to this paper. We would also like to acknowledge Chun K. Chan for his work on the reliability model.

References

[1] R. J. Bates, Jr., Disaster Recovery Planning: Networks, Telecommunications and Data Communications, McGraw Hill, New York, NY, 1992.

[2] R. Billinton and R. Allan, Reliability Evaluation of Engineering Systems: Concepts and Techniques, Pitman Publishing Inc., Marshfield, MA, 1983.

[3] G. Grimmett, and D. Stirzaker, Probability and Random Processes, 2nd ed., Oxford Univ. Press, Oxford, UK, 1992.

[4] Harvard Research Group, "HAS Tracking/High Availability Challenge," <http://www.hrgresearch.com/ha/>.

[5] Insure Egypt, Horus Newsletter, 2:13 (Mar. 20, 2001), <http://www.insureegypt.com/Download.html>.

[6] Telcordia (Bellcore), "Methods and Procedures for System Reliability Analysis," Iss. 1, SR-TSY-001171, Jan. 1989.

(Manuscript approved February 2004)

AHMAD JRAD is a member of technical staff in the Network Planning and Business Modeling Department of Bell Labs Advanced Technologies in Holmdel, New Jersey. He is responsible for developing new models and techniques for network planning and design professional services and associated tools. He holds B.S., M.S., and Ph.D. degrees in electrical engineering from the University of Pittsburgh in Pennsylvania. Dr. Jrad's current areas of focus are in business continuity planning and network reliability.

THOMAS MORAWSKI is a technical manager in the Network Planning and Business Modeling Department of Bell Labs Advanced Technologies in Holmdel, New Jersey. He holds a B.Sc. degree in experimental psychology from Towson University in Maryland and an M.S. degree in operations research from the State University of New York at Buffalo. His current responsibilities include developing and delivering network planning and design services. Mr. Morawski and his group also develop network planning and design automation tools.

LOUISE SPERGEL is a senior manager in the Network Planning and Business Modeling Department of Bell Labs Advanced Technologies in Holmdel, New Jersey. Her current activities focus on the development of Bell Labs powered professional services including the Business Continuity Preparedness service. She holds B.S. and M.S. degrees in electrical engineering from Cornell University in Ithaca, New York, and the Massachusetts Institute of Technology in Cambridge, Massachusetts, respectively. ◆