

Quantifying the Business Impact Analysis

by Michael Miora, Copyright (c) 1996, All Rights Reserved.

The Business Impact Analysis (BIA) comprises the heart of the disaster recovery planning process. It is here that the disaster recovery planner determines what is important and in what time frame for inclusion into the Disaster Recovery Plan and what is not relevant to that effort. The BIA determines how far to go in protecting the people, information and equipment that constitute the organization and its functions so that all survive to flourish another day.

The BIA *matrix* is central to the analysis. It is used to establish disaster recovery plan technical objectives, to review the BIA results with management, and to estimate costs of outages. There are two basic methods used to develop this matrix. The first is the traditional Quantitative Risk Model, a version of which is described below. The second is the model developed by the author, Michael Miora, specifically for use in Disaster Recovery Planning.

The author will demonstrate that the traditional Quantitative Risk Model is complex, cumbersome and often results in *analysis paralysis*. The Miora Generalized Cost Consequence (GCC) Model™ (patent pending) is simpler and useful.

Quantitative Risk Model

The Quantitative Risk Model is a formal and rigorous methodology used for analyzing expected losses that will be incurred over a pre-determined time period. This procedure requires a tremendous amount of analysis and research, and is frequently used by management as a delaying or avoidance tactic. Moreover, the risk model assumes a long term view and cannot account for the non-fiscal repercussions of non-preparedness such as management turnover or corporate marauding. It is difficult to build and of marginal use.

The Quantitative Risk Model consists of three main factors: probability of loss, cost of loss, and annual loss expectancy (ALE). The probability of loss is really a sum of the probabilities of different catastrophic events that range from partial outages to severe interruptions. The cost of loss depends upon the level of interruption. For example, a partial building loss affecting computer systems but leaving phone systems in operating condition has a much lower cost than a complete building destruction. Therefore, the cost of loss is dependent upon the type of disaster.

A simplified risk model considers the probability of loss and the cost of the loss. The annual loss expectancy is simply the product of the probability and the cost. For example, let us assume there is a 5% probability (annual) of a major power failure. Stated differently, this means that the facility will experience a major power failure once in twenty years. Let us further assume that the power failure will cause a 72 hour outage, which will cost the company \$1,250,000. The ALE is calculated as $5\% \times \$1,250,000$, or \$62,500. This number is compared to the baseline cost of the recovery plan and the cost of capital. For example, say the cost of the recovery plan is \$2,000 per month and the cost of capital is \$25,000. Then the baseline cost, or comparison figure, is $(12 \times \$2,000) + \$25,000$, or \$49,000. Since the comparison figure is lower than the ALE, recovery planning is justified.

There are some serious shortcomings in this simplified approach. First, the cost of the outage

depends upon the level of loss. A 72 hour power outage is significantly costlier than a 24 hour outage. Therefore, the ALE must reflect the difference in probabilities of different levels of impact. Another problem is defining the probability of occurrence for an aggregation of events.

The more acceptable risk model must consider the different levels of loss and sum the probabilities of all disasters that can cause that level of loss to define the true probability of loss for that loss level. This is accomplished for each loss level. First, a series of disaster events is defined. Each event is then refined into levels. For example, office buildings are susceptible to loss due to fire. Data on numbers of fires and amount of destruction (in predefined ranges) is available from various fire protection services. For a given facility, consider the total number of such buildings in the geographical vicinity. Then calculate the frequency or probability of a fire causing a range of damage. Perform this calculation for each defined range of damage. Assess the level of loss for each range of damage. Then calculate the ALE for each range of damage by multiplying the probability of a fire causing that range of damage with cost of the loss if there is such a fire. Sum all the ALE values to calculate the total fire ALE. This calculation must be performed for all types of disasters that can affect the facility to determine the grand total ALE.

For each level of impact for all disasters, the baseline costs must also be calculated. These figures are also summed to form a total Baseline Cost. The grand total ALE is then compared with this baseline cost figure. This is clearly a complex process that requires tremendous effort to generate and great patience to explain.

There are two great problems that afflict the Quantitative Risk Model. First, calculating all the outage costs is very difficult and subject to great debate among management. This time consuming activity may delay plan development for many months. Moreover, once the cost figures are finalized, they are subject to constant change due to the changing business climate and practices.

The second great problem is calculating the probabilities. This is also very difficult and often requires many subjective conclusions. For example, what is the effect of modernizing the sprinkler system on the level of damage experienced by a particular type of fire? Each countermeasure can significantly alter both the cost and the probability. Moreover, the probability of any particular event tends to be quite small, often less than one percent. Such low probability figures tend to cause management to decide against disaster recovery planning by fostering the idea that no disaster will befall them -- many believe that disasters only happen to others.

The Quantitative Risk Model is an interesting actuarial exercise. It also can be improved by using greater mathematical complexity to refine the estimates. However, for the disaster recovery planner, this risk model requires tremendous effort and tends to delay the process.

Generalized Cost Consequence Model

The Generalized Cost Consequence (GCC) Model™; (Patent Pending) does not consider probabilities of specific disaster events. Instead, it estimates the total cost of outages as a function of time after an event. This model is significantly simpler than the Quantitative Risk Model: it is easier to build and simpler to explain. The GCC estimates the cost of an outage for each function and applies that cost to the total disaster cost after the maximum allowable down

time has been exceeded.

For instance, assume the cost of delaying the Treasury Department's bank management function is \$25,000 per day after the first day. Let us also assume that the cost of delaying the Law Department's general contract review is \$5,000 per day after seven days. For the bank management function, we calculate the cost to the company as \$25,000 per day beginning on the first day. For the contract review function, we calculate the cost as \$5,000 per day beginning on the eighth day. Therefore, the contract review function does not contribute to loss during the first seven days. We perform this calculation for each function and collect the costs by category. This category cost summary is used to develop and present a graph that shows the total cost losses for each category level once they are activated and the total for all categories over time.

A sample graph of the contribution of functions aggregated by Category level is shown in Figure 1. In this example, Category I functions cause slightly more than \$120,000 of loss on a daily basis once the maximum allowable downtime has been exceeded. Category II functions contribute slightly under \$60,000 in this example. It is likely that functions will commence their loss contribution at various times after the disaster event. Therefore, the true Category I loss contribution may begin at a lower level and increase to its full level. That distribution of effect will occur beginning on the first day any function exceeds its allowable downtime and continue growing until the last day any function exceeds its downtime, at which point the effect will have achieved its full loss contribution. Since the categorization clusters functions with similar downtimes, we can present the loss as a single, or point, value rather than a value that varies over time.

Notice the shape of the set of bars: Category I is a higher value than Categories II and III. Category IV is highest. This is a characteristic shape for this graph. No matter the number of categories chosen, the Category I functions tend to be quite high in their contributions. The next one or two categories are lower, but increasing in values until the lowest priority category, which tends to be quite high. This shape results from the manner of the categorizations. The Category I functions are highly critical with great effects on corporate operations. Category I functions tend to number in the few, with each function contributing significantly. The middle categories tend to be larger in number than Category I, with the number increasing with each succeeding category, though still small compared to the overall set of organizational functions. The final category tends to contain the largest number of functions. Though many contribute less on a daily basis than the Category I functions, the sheer magnitude of this set often causes it to be the largest overall contributor to loss. Recall, however, Category I functions begin their effects quickly, while loss of Category IV functions may be benign for weeks.

A sample graph showing the accumulation of losses from all categories following a disaster is shown in Figure 2. Note that this figure uses the data from the previous figure but in a different format. Neither this nor the previous figure measures physical losses such as real property, capital equipment, and the like. The horizontal axis (x-axis) represents the number of days since the disaster event caused damage and cessation of corporate functions. The vertical axis (y-axis) represents the total losses the company will have sustained on the corresponding day represented on the x-axis.

At the point of the disaster occurrence, no losses will have accumulated as the result of the non performance of any function. Hence, the accumulated loss begins at zero on the first day.

Thereafter, the accumulated loss is increased each day by the cost contribution of the category whose earliest start time has already been surpassed. Let us assume that Category I functions begin to contribute to corporate losses of \$120,000 on the first day, and Category II functions begin to contribute losses of \$60,000 on the seventh day. Further, assume that Category III functions contribute \$80,000 daily beginning on the fourteenth day. In this case, the cumulative loss begins at zero and grows by \$120,000 per day for the first six days. On the seventh day, the Category II functions begin to contribute \$60,000 per day along with the ongoing contribution of the Category I functions. Therefore, beginning on the seventh day, the cumulative loss grows by the sum of \$120,000 and \$60,000, which is \$180,000 daily. On the fourteenth day, the daily loss increases by another \$80,000, representing the Category III contribution. This brings the total daily loss to \$240,000, which is the sum of the contributions of Categories I, II and III.

The cumulative loss summary shows at a glance the loss the company will experience over time following the disaster if no disaster recovery planning is performed. The active simplifications of this model are the grouping of functions by category and the subsequent representation of each category as a single value beginning at a fixed point in time. This may render the estimate slightly inaccurate at some local points, but the overall values are as accurate as the underlying estimates.

The Generalized Cost Consequence (GCC) Model summarizes at a glance the effects of loss of functions. The Law of Large Numbers helps assure that the overall estimate is reliable. These loss figures can be moderated by insurance reimbursements, legal liabilities, and overall management objectives.

A second model can be developed assuming the disaster recovery planning is in place, thereby showing the residual loss proposed. That model would be developed in a similar manner, but with the assumption that certain functions are restored within established parameters. There would be residual loss only if the restoration occurs later than the allowable downtime. This will almost certainly be the case for some functions. Figure 3 illustrates a graph that shows cumulative losses with and without a disaster recovery plan in place. The *Without DR Plan* curve is similar to the previous figure. The *With DR Plan* curve reflects the residual loss that will occur even if a plan is in place. This loss is normally dramatically lower, but is seldom zero. Typically, a reduction to zero residual loss would require extraordinary and expensive measures such as parallel processing with real time backup to off site locations. Most organizations can benefit more from a significant lowering of residual losses than from a full reduction to zero losses.

The estimation process itself is much simpler than the Quantitative Risk Model. For each function, estimate the loss based on three criteria or types of losses. These three losses are: tangible and direct losses, tangible and indirect losses, and intangible losses.

The tangible and direct losses are the easiest to calculate. These losses can be traced to specific revenue producing functions. The results are direct because the loss occurs as a first order effect, meaning that revenue stops because the function cannot be performed. The results are tangible because they can be easily measured. An example of such a function is automated production control of an assembly line. If the systems exceed their allowable downtime, then production will cease. The cost is the resultant loss of sales after inventory is depleted. Another example is loss of order entry functions. In this case, the result is similarly calculable: lost sales after in stock items are depleted and existing orders are produced. Tangible, direct losses include lost sales,

lost manufacturing, lost deliveries, and other lost opportunities.

The tangible and indirect losses are the most common and slightly more difficult to estimate. Support functions generally produce tangible results whose privation would cause a fiscal loss indirectly. For example, a public corporation issues quarterly earnings reports that, if late, could have significant consequences on the company's stock value. Though this is not a direct loss resulting from cessation of sales or other production, the loss can be calculated using accounting standard practices. The internal accounting personnel are in the best position to provide this estimate to the disaster recovery planner. Tangible, indirect losses include penalties, fees, fines, market share, and other issues that can be directly calculated.

Intangible losses are the most difficult to calculate. These intangible effects include reduced public confidence, compromised customer satisfaction, promises not kept, damaged reputation, and other losses that are general in nature and not easily calculable. Sometimes, these losses are not translated into specific fiscal losses and are, therefore, not represented in the cost graphs. In such cases, prominent notations should be made explaining the additional, but not quantified losses.

The Generalized Cost Consequence Model can solve the problem of cost justification. This model shows the potential, catastrophic losses without engaging in the analysis paralysis that can stem from a detailed Quantitative Risk Model development effort.

Michael Miora is ICSA's Director of Consulting Services. If you have any questions about ICSA's consulting services, you may direct them to Michael Miora via e-mail at mmiora@miora.com or by phone at 310/306-1365. You may also get information by visiting our Internet site at <http://www.icsa.net> or our CompuServe forums (GO ICSA).