

3. Select the relevant process to simulate.
4. Click on *Update Chart* to view the updated computation of a single path or click *OK* to create the process.

Model Results Analysis

For your convenience, the analysis report sheet is included in the model. A stochastic time-series chart and forecast values are provided in the report as well as each step's time period, mean, and standard deviation of the forecast (see Chapter 8 for technical details of using this stochastic process forecast technique and Chapter 6 for the data diagnostic and statistical analysis tools for calibrating the input parameters for the stochastic process models). The mean values can be used as the single point estimate or assumptions can be manually generated for the desired time period. That is, finding the appropriate time period, create an assumption with a normal distribution with the appropriate mean and standard deviation computed. A sample chart with 10 iteration paths is included to graphically illustrate the behavior of the forecasted process.

CASE STUDY: IT INFORMATION SECURITY INTRUSION RISK MANAGEMENT

This case study illustrates an information systems security attack profile and provides decision analysis and support on the optimal investment. The model is contributed by Mark A. Benyovszky, Managing Director of Zero Delta Center for Enterprise Alignment. Zero Delta CfEA is a research and development organization that specializes in helping companies to align their strategic and tactical efforts.

There are several models illustrated in this case study and available in the Modeling Toolkit. First, the IT Bandwidth Requirements for Streaming Media model shows how to create forecasts for media streaming environments for off-peak and peak media consumption cycles. The model explains how to use historical time-series data to establish an understanding of the future demand for media consumption. This model also explores how the Delphi Method can be used when historical data do not exist or when there is significant uncertainty surrounding the demand for new media types that require more bandwidth or have chattier communication channels.

The IT Intrusion Management model shows how to create an information systems security attack profile, determine the probabilities of occurrence of different types of attacks, assess the financial and operational impact that an attack has on an organization, and arrive at the level of investment in

technology security solutions (e.g., intrusion detection systems, intrusion prevention systems, and network behavior analysis solutions) necessary to mitigate the profiled attacks.

Finally, the IT Storage and CPU Demand model illustrates how to create forecasts for future information system storage and CPU demand. The model explains how to use historical time-series data to establish an understanding of the future demand for media storage and computing cycles. This model also explores how the Delphi Method can be used when historical data do not exist or when there is significant uncertainty surrounding the demand for new media types that require more storage space and where CPU cycles are being consumed by expensive encoding and decoding processes.

Organizations of all sizes rely on technology to support a wide range of business processes that span the spectrum from back-office finance and accounting to mid-office manufacturing, distribution, and other operational support functions to front-office sales, marketing, and customer support functions. As a general rule of thumb, larger organizations have more complex system environments and significantly greater volumes of data along with a wide range of different types of information.

If you were to look across industries, you would see different degrees of sensitivity of both the systems and information that are employed. For example, financial and insurance companies store critical and very sensitive information (financial transactions and personal medical histories) about their customers, or that an energy company engaged in gas transmission and distribution relies on critical technology systems that control the flow of gas through complex pipeline networks.

Regardless of the specific industry an organization is involved with or the size of the company, the underlying technology systems and the data and information they consume and produce are significant business assets. Like any asset, they must be protected. In order to protect these assets, we must understand what their individual and collective risk profiles look like.

Protecting these assets is of paramount concern. Technology systems are interconnected across private, semiprivate, and public networks. Every second (perhaps you prefer nanoseconds, picoseconds, or attoseconds, depending on your geekiness factor) of every day, information moves across these networks; most of the time the information moves about intentionally, whereas on other occasions it does not.

We can think of this information and these systems in the context of an information security asset portfolio. It is important for us to quantify the value of each class of system or set of information, which will help us to understand, according to a scale of sensitivity, which assets require greater protection, as higher-value assets are likely to be greater targets for attack (based on the basic risk/reward equation).

We can then apply various methods against the portfolio to determine the composite (high-level view) risk level of the portfolio, risk profiles of categories of assets, and individual asset risk profiles (detailed view). This approach enables us to gain a better grasp on our information and technology asset portfolio, and provides us with the ability to determine how much to invest to protect each class of assets.

While the specific approaches and processes that are required to perform this initial portfolio structuring are beyond the scope of this case study, determining the probabilities of events occurring against these assets and what the resultant outcomes are likely to be is at the center of our discussion. This case study will assume that this structuring process has already been completed. Specifically, there are five steps to undergo:

- Step 1: Create environment details.
- Step 2: Create attack models.
- Step 3: Create attack scenarios.
- Step 4: Determine financial impact.
- Step 5: Arrive at investment decision.

Now, let us get on with the heart of our discussion. Monte Carlo simulation provides us with an effective way to estimate losses associated with a given attack. Monte Carlo simulation addresses the “flaw of averages” problem that plagues many single-point estimates or estimates based on standard averages.

For the sake of this discussion, we will explore how we applied this approach to a large gas transmission and distribution company. The company (which we refer to as Acme T&D) is one of the largest natural gas transmission and distribution companies in North America. Acme T&D has an extensive gas pipeline network that supplies natural gas to wholesalers and retailers in some of the largest markets throughout North America.

Energy companies fit in a unique category of organizations that use technology at the core of their business operations. Acme T&D relies on an extensive industrial control system known in the industry as Supervisory Control and Data (SCADA) and Process Control Monitor (PCM) systems. These systems are composed of a number of devices that are distributed throughout the gas pipeline network; these components are used to control the flow of gas through the network. They supply critical information, such as gas flow rate, temperature of gas, and pressure at various points through the network, to a system operator who then makes certain decisions about what to do to keep the pipeline at an operationally efficient level—always

supplying gas where and when it is needed in a dynamic environment that changes continually.

These systems are critical not only to the operations of Acme T&D but also to the greater infrastructure of the United States. If the transmission and distribution of natural gas are interrupted for a significant period of time, the interruption can have downstream effects that could be devastating economically (the suspended operations of manufacturing companies that rely on natural gas) or personally (lack of gas to run a furnace in the cold of winter).

Clearly, these SCADA system(s) would be categorized as business-critical assets with the highest priority placed on their protection.

Step 1: Create Environment Details

When we consider the extent to which an attack will cause damage, we must identify the factors that drive the top end of our model. These factors will be different for each company (with similarities for companies within the same industry).

For Acme T&D our greatest concerns, from an operational perspective, are the count and types of networks in the environment, and employee productivity (we will take into account separately how operations are impacted when a threat impacts a SCADA network). The reason for using employee productivity as a factor is that when networks are down or systems are unreachable (for whatever reason), employees are directly impacted (we use this in this example because of its universal relevance across industry domains).

Acme T&D Network Counts

Enterprise Network Count	16
SCADA Network Count	4
PCM Network Count	1
Total Networks	21

As an aside, and as previously alluded to, the factors that drive the model will change based on industry characteristics. For example, a financial institution may wish to use the economic losses associated with stolen credit card data as a primary factor to drive the model, in addition to employee productivity losses and so forth.

Acme T&D has approximately 10,000 employees. We must determine the payroll expenses (fully burdened) per hour. We are simplifying this model

intentionally—it is not likely that 10,000 employees are working all at once (e.g., some percentage of employees may be on a shift rotation schedule). A sample computation is shown next:

$$\text{Total Employee Cost per Hour} = \text{Employee Count} \times \frac{\text{Salary}}{2,000}$$

where 2,000 is the number of hours worked per employee each calendar year (2,080 less 80 hours for holidays); and the salary input is the fully burdened amount at the average of all employees.

The model is based on various types of attack. We determine the probability that each attack will occur and to what extent it will cause damage (economic and operational) to the organization. We then create a separate model (our attack portfolio), which will allow us to simulate multiple attacks occurring against different networks in the environment and the resultant impacts in aggregate. We classify attacks based on two variables—the frequency and impact of the attack.

An attack as profiled in Class I is considered an average attack (see Table 14.8). An average attack could be considered a low-impact worm, a Trojan horse, or a virus that may affect various network systems and employee computers. Acme T&D has a variety of tools deployed in the network to mitigate these types of attacks; however, no tool is 100% effective. This is where the value of Monte Carlo simulation is realized.

Minimum 0.7

Most Likely 0.8

Maximum 1.0

Now we construct the remaining elements of the model. We will use standard (and very conservative) estimates for the probability of occurrence of an attack.

Table 14.8 illustrates how the top end of the model comes together. We place the attack types across the columns of the model and we create the network structure and impact structure components.

Step 2: Create Attack Models

We must first create a common attack model and categorize the different types of attacks that exist. The classes of attacks are based on the severity level of the attack (from average to extreme). We also indicate the extent

TABLE 14.8 Qualitative Assessments of Attack Classes

Attack Class	Severity Level of Attack	Type of Attack	Extent of Damage	Recovery Approach
Class I	Average	Benign worm, Trojan horse, virus, or equivalent	Limited. Most damage occurs at host level.	Mostly automated, but may require some human intervention.
Class II	Slightly above average	Worm, Trojan horse, virus, or equivalent designed to create some damage or consume resources	Limited. Damage can occur at the host and network level.	Human intervention is required. Humans use tools that require interaction and expertise.
Class III	Moderately above average	Worm, Trojan horse, or equivalent designed to create significant damage and consume resources	Noticeable damage at host and network levels. Automated tools have limited effect to combat attacker.	Significant human intervention is required. Personnel require physical access to host machines and network environments.
Class IV	Significantly above average	Concentrated attack by hacker using a variety of tools and techniques to compromise systems	Significant damage to important/sensitive data. May also include damage to host machines as Trojans and other tools are used to circumvent detection and mitigation techniques.	Extensive human intervention is required. Data and systems recovery is necessary. Multiple techniques and methods are necessary to fully recover.
Class V	Extreme case	Concentrated attack by hacker or groups of hackers who are trying to compromise information/systems and have malicious intent	Critical damage to important/sensitive information. Irreversible damage to systems/hardware.	Extensive human intervention is required. External experts are required to assess and recover environment.

of damage that an attack produces and the recovery details associated with each class of attack. This classification structure provides us with a basic framework we can leverage throughout the analysis exercise. We have five classes of attacks structured in our model. The descriptors are qualitative in nature (see Table 14.8).

We create current state and future state models for the classes of attacks. The purpose for creating current and future state models is so that we can compare the models to each other. The current state model is based on the technology and approaches that are currently in use (our preexisting investments) to detect, mitigate, and recover from each respective type of attack. The future state model is based on a set of new technologies (our future investments) that can be deployed in the environment to enhance the security of the environment, mitigate a wider range of attacks, and more rapidly recover from various types of attacks.

These types of attacks will be consistent across our current and future state models. There are a number of variables that are a part of our attack models. They include:

- % of network impacted.
- % of employees impacted.
- Productivity loss (hours/employee).
- Costs to recover employees.
- Hours to recover employees.

Note that the models are populated with static values that are single-point estimates and averages. For example, a Class I attack in the current state attack model has a 10% Network Impacted value and a 5-hour Productivity Loss value.

How can we be absolutely certain that a Class I attack will always impact 10% of the networks and result in a productivity loss of 5 hours per employee (along with the other variables included in the model)? We cannot be sure with a reliable degree of confidence. Therefore, any analysis based on single-point estimates or averages is flawed.

Monte Carlo simulation allows us to refine our assumptions and provides us with a mechanism to perturb these variables in a dynamic fashion. While we have solved the problem of dealing with averages, we are faced with a new challenge: What are the appropriate ranges to use to perturb the values, and how should these perturbations behave throughout the simulation?

To gather these values, we leveraged the Delphi Method. Following the Delphi Method approach, we interviewed a number of technical experts

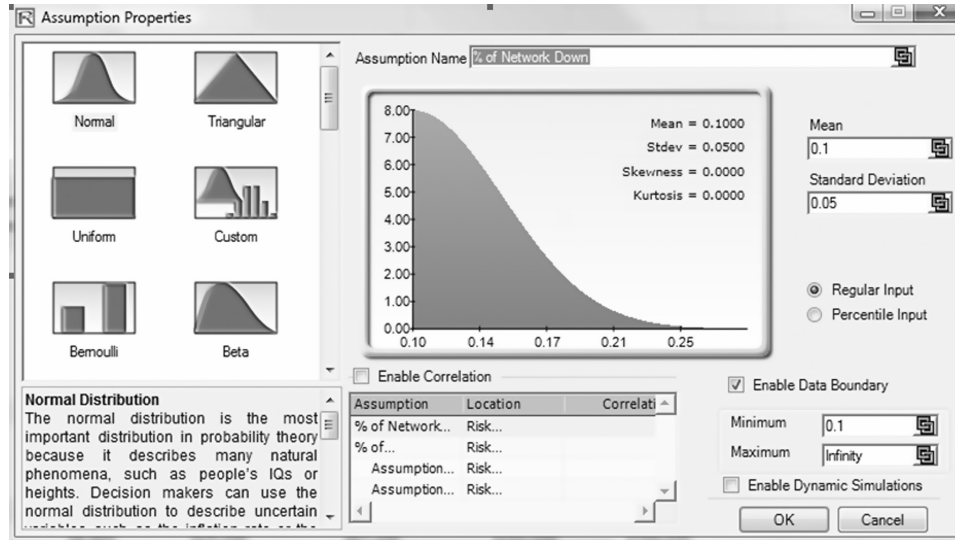


FIGURE 14.114 Truncated percent of Network Impacted simulation assumption.

in the environment who had knowledge of prior attacks and the extent to which tools were used to mitigate them. The expert panel provided the details necessary to determine how the model variables might behave and what their respective upper and lower boundary values may be.

Figure 14.114 illustrates how we have adapted the % of Network Impacted value for a Class I attack. The original value was based on an average of 10%. Upon closer inspection and after some discussion, our panel of experts determined that such an attack is unlikely to impact less than 10% of the network and may in fact impact a greater percentage of the network before it is identified and stopped, preventing further damage. Using Monte Carlo simulation, we create an assumption for this value and select a normal distribution. We truncate the left side (or minimum value) of the distribution to take into account the 10% floor and provide some movement toward the right side (or maximum value) of the distribution. We set the mean to 10% and standard deviation to 5%. The resultant distribution indicates a minimum value of 10%, a mean of 10% (our average), and a maximum value of approximately 25%.

We have introduced into our model a very powerful feature. Our model better reflects reality by taking into account the uncertainty associated with this value. We use this same approach for the other values and select and adjust the distributions accordingly. To further illustrate this point,

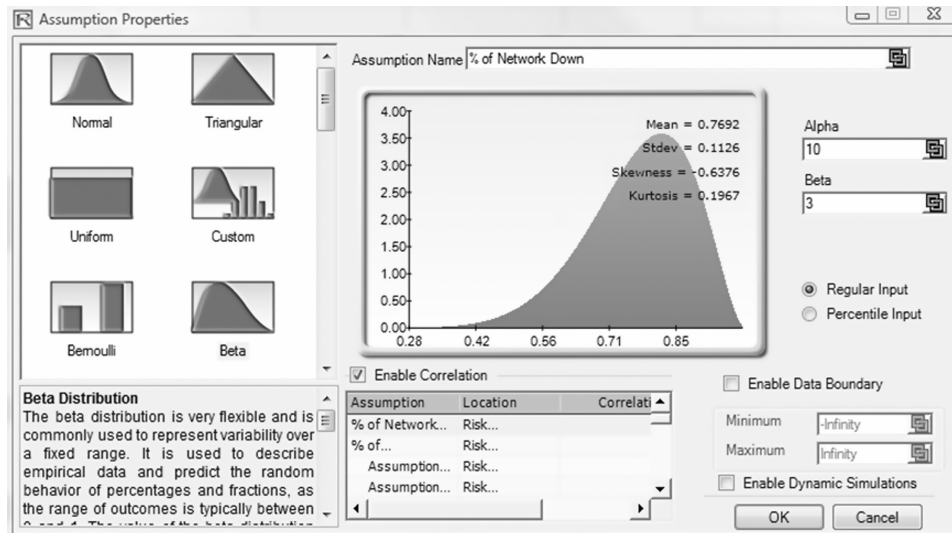


FIGURE 14.115 Percent of Network Impacted simulation assumption of a Class V attack.

Figure 14.115 is taken from the Class V attack column. A Class V attack is considered an extreme event. The probability of occurrence is very low, and the damage caused is expected to be extreme or catastrophic in nature. An analogous event would be a volcano eruption or an earthquake (they may evoke a tsunami wave, for example, especially if located in the South Pacific) that occurs once every 100 years.

The Gumbel Maximum Distribution is ideally suited for this type of catastrophic event. This distribution model is positively skewed and is designed to produce a higher probability of lower numbers and a lower probability of extreme values. We set the alpha value to 70 and the beta to 10. This results in a mean of 75.7722 and a standard deviation of 12.8255. It is important to note the third and fourth moments of this distribution. Skewness coefficient is 1.1395 (indicating the positively skewed nature of the distributions) and kurtosis coefficient is 2.400 (indicating the extent to which extreme events should occur in the distribution).

This distribution model better reflects reality vis-à-vis extreme attacks. We can see in Figure 14.116 that there are higher probabilities to the left of the mean than to the right. However, the model has taken into account the extreme distributions to the far right of the median.

The original analysis based on averages indicated that for this scenario, the total financial losses are \$21,741,176. If we follow our “1 in 3”

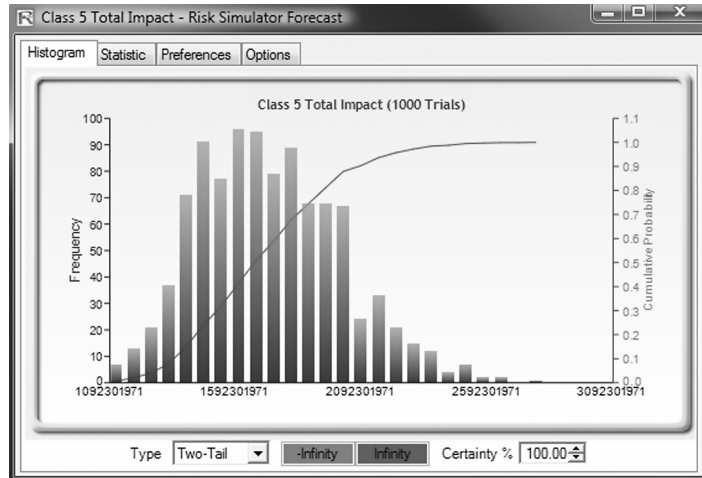


FIGURE 14.116 Forecast distribution of a Class V impact.

approach, we find that the number is adjusted downward to \$18,0174,729 or by a little over 12%. As you explore the model in more detail, you will note the use of various distributions for each class of attack.

We adjust these figures for each scenario to take into account the greater variability of more advanced and staged attacks. We know that as attacks gain more sophistication there are more unknowns about how far-reaching or to what extent damage will occur. Hence, the mean and standard deviation parameters can be adjusted to take into account this variability.

Model Results

Impact to Operational Productivity We have determined that the average fully burdened salary per employee is \$80,000. For Scenario I, we estimate that an attack that affects each employee results in 5 hours of lost productivity. It costs Acme T&D \$39.22 per employee per hour of lost productivity. For an attack profile we modeled in Scenario I where 10% of the networks and 10% of employees are impacted results in a total productivity loss of \$196,078.43 (see Table 14.9).

Recovery Costs Attacks generally result in some form of damage (more often than not the damage is nonphysical in nature). It is often necessary

TABLE 14.9 Modeling Results from Scenario I

Lost Revenues	
Impact to Operational Productivity	\$196,078.43
Assumption—Average Salary/Employee (Fully Burdened)	\$ 80,000
Assumption—Total Time to Fully Recover/Employee (Hours)	5
Productivity Cost/Hour	\$ 39.22
Costs to Recover/Employee	\$50
Assumption (Hours to Recover)	1
Costs to Recover Networks	\$ 4,800
Assumption—Hours to Recover	12
Resources per Network	5
Cost per Hour	\$ 50
Total Costs to Recover Employees	\$ 50,000
Total Costs to Recover Networks	\$ 4,800
Total Impact	\$246,078.43

to deploy technical personnel to remediate the impacted environments and systems. There are two dimensions to this remediation. There is network remediation (resetting/reconfiguring network routers, switches, firewalls, etc.) and client remediation (ghosting client machines, patching software, reinstalling/reconfiguring software, etc.).

Our model takes into account the number of resources and the time necessary to recover the networks and the same for recovering employees. For Scenario I the costs are \$50,000 and \$4,800, respectively.

Total Impact We now sum up all of the separate loss components of the model:

$$Loss(Productivity) + Loss(Network Recovery) + Loss(Employee Recovery)$$

For Scenario I, we have total losses of \$147,647.

In the model, there are four additional scenarios. For each scenario we tweak the assumptions to better fit the attack profiles. The percentages of networks down and employees impacted increase for each scenario.

Exposing the Flaw of Averages

	Class I Attack	Class II Attack	Class III Attack	Class IV Attack	Class V Attack
Total Impact (Original)	\$147,647	\$616,471	\$1,933,235	\$5,223,529	\$21,741,176
Total Impact (Revised)	\$310,616	\$714,145	\$1,679,616	\$7,507,908	\$23,817,256
Variance (%)	210.38%	115.84%	86.88%	143.73%	109.55%

The next step of our modeling efforts involves creating a portfolio of attacks. This step will provide us with the answer to the question: How much should Acme T&D invest in security solutions to mitigate the risks associated with the attacks profiled?

Step 3: Create Attack Scenarios

Now that we have determined the estimated costs associated with different types of attacks, we are ready to move on to creating the attack scenarios. The attack scenarios will provide us with the total losses realized during a specified period of time.

We have created six attack scenarios. The attack scenarios consider the occurrence of different types of attacks over a 5-year period. By creating different scenarios, we can consider different foreseeable futures. This approach allows an organization to determine how it wishes to view the world from a risk planning and risk mitigation standpoint.

The degree to which an organization will tolerate risk varies greatly. Some organizations are more tolerant of risk and will invest less in mitigating technologies and approaches, whereas other organizations that are more risk averse will invest substantially more in order to reduce their risk profiles.

One can think of this type of investment as an insurance policy—juggling premium with payout—or from a strategic real options perspective of risk mitigation. The scenarios provide us with a landscape view from lowest to highest possible losses. We explore two different approaches to determining the probability of attacks occurring across a specified time line. The first approach involves the use of the Delphi Method. We interview a number of subject matter and technical experts who are asked to produce five different likely scenarios for various attack profiles. We provide some guidance and suggest to each expert that the scenarios should range from a most likely

	Scenario I					
	Year 1	Year 2	Year 3	Year 4	Year 5	Totals
Class I Attacks	1	0	1	1	0	3
Class II Attacks	0	0	0	0	1	1
Class III Attacks	0	0	0	0	0	0
Class IV Attacks	0	0	0	0	0	0
Class V Attacks	0	0	0	0	0	0
Class I Attack Impact CS	309,579	0	309,579	309,579	0	928,737
Class I Attack Impact FS	44,632	0	44,632	44,632	0	133,896
Class II Attack Impact CS	0	0	0	0	713,288	713,288
Class II Attack Impact FD	0	0	0	0	303,871	303,871
Class III Attack Impact CS	0	0	0	0	0	0
Class III Attack Impact FS	0	0	0	0	0	0
Class IV Attack Impact CS	0	0	0	0	0	0
Class IV Attack Impact FS	0	0	0	0	0	0
Class V Attack Impact CS	0	0	0	0	0	0
Class V Attack Impact FS	0	0	0	0	0	0
Impact Based on Current State	309,579	0	309,579	309,579	713,288	1,642,025
Impact Based on Future State	44,632	0	44,632	44,632	303,871	437,767
Variance	85.58%	#DIV/0!	85.58%	85.58%	57.40%	73.34%
Risk Adjustment	264,947	0	264,947	264,947	409,417	1,204,258

FIGURE 14.117 Scenario I attack profile of a future state.

scenario to a least likely scenario. This team of experts is then brought together to discuss the ranges across the spectrum. After various conversations and some debate, the team collectively determines to reduce the total scenarios ($5 \text{ experts} \times 5 \text{ scenarios} = 25$) to the final 5. These scenarios are reflected as Scenarios I through V on the Attack Scenarios spreadsheet.

Figure 14.117 illustrates a Scenario I attack profile. On our defined scale of least likely to most likely, this scenario is most likely to be realized. The experts provided the count of each type of attack that occurs within our 5-year period and further determined the years in which the attacks will occur.

We have carried over our financial impact information from our previous exercise. For each class of attack we have current state and future state impact costs.

The first section of the model includes the classes of attacks. For Scenario I, we have determined that three Class I attacks will occur in years 1, 3, and 4. In addition we have determined that one Class II attack will occur in year 5.

The second section of the model includes the impact values from the attack models. We include in this model both the current state and future state impact values. These values are computed for each year and are summed in the totals column. The variance value indicates the percentage reduction from current to future state loss values. By investing in the proposed technologies, we can reduce by 73.34% the total losses for this scenario. The risk adjustment value is the difference between the current state impact and future state impact values. This value is carried over to the next step of our analysis.

We use this same model to create the other attack scenarios. Figure 14.118 illustrates the Scenario IV attack profile. This scenario represents the

Scenario IV						
	Year 1	Year 2	Year 3	Year 4	Year 5	Totals
Class I Attacks	0	0	0	0	0	0
Class II Attacks	0	0	0	0	0	0
Class III Attacks	0	0	0	0	0	0
Class IV Attacks	0	0	0	0	0	0
Class V Attacks	0	0	0	0	1	1
Class I Attack Impact CS	0	0	0	0	0	0
Class I Attack Impact FS	0	0	0	0	0	0
Class II Attack Impact CS	0	0	0	0	0	0
Class II Attack Impact FD	0	0	0	0	0	0
Class III Attack Impact CS	0	0	0	0	0	0
Class III Attack Impact FS	0	0	0	0	0	0
Class IV Attack Impact CS	0	0	0	0	0	0
Class IV Attack Impact FS	0	0	0	0	0	0
Class V Attack Impact CS	0	0	0	0	23,791,472	23,791,472
Class V Attack Impact FS	0	0	0	0	16,095,255	16,095,255
Impact Based on Current State	0	0	0	0	23,791,472	23,791,472
Impact Based on Future State	0	0	0	0	16,095,255	16,095,255
Variance					32.35%	32.35%
Risk Adjustment	0	0	0	0	7,696,217	7,696,217

FIGURE 14.118 Scenario IV attack profile of a future state.

opposite end of the spectrum. In this scenario the company is successful in preventing all classes of attacks until year 5, where a Class V attack occurs. This is the scenario of the infamous “hacker with malicious intent” who makes a concentrated effort to circumvent intrusion management technologies with the specific desire to cause significant harm to the organization. For Acme T&D this scenario could perhaps reflect the sentiments of a terrorist who has a desire to gain access to the critical gas pipeline systems in order to cause a catastrophic failure to the pipeline network.

One could argue that such an approach to determining these probabilities lacks scientific rigor or can be significantly biased—either intentionally or unintentionally. Consider technical experts who firmly believe that their skills are second to none with respect to effectively deploying and managing an armory of intrusion management technology. They may be biased to create scenarios that are on the conservative end of the spectrum, significantly coloring the reality of the environment or threat landscape. If we pin our decision on this approach, a crafty hacker who has superior skills to these individuals may easily circumvent these technologies and successfully realize his or her attack objectives and goals.

Conversely, consider the doomsday character who is constantly pondering the worst-case scenario and has a strong voice in the room. He or she may be overly aggressive with the attack scenarios, creating unrealistic models that result in doom and gloom.

How can one test for these biases? Is there a way to independently determine the probabilities and likelihoods of events? Indeed there is a way, and it is again found in Monte Carlo simulation.

Scenario VI represents our independent attack scenario. You may consider this the control model. This is our expert who is neutral to all biases. The probabilities of occurrence are factually driven and leverage a distribution model that is focused on the discrete nature of these events—an event either happens or it doesn't.

The Poisson distribution provides us with the ability to address the unique aspects of occurrence probabilities. Figure 14.119 illustrates how we can leverage the Poisson distribution for modeling an attack. These events are discrete in nature—they either occur or don't occur. For a Class I attack we set the lambda value to 1.5984. This creates a distribution model that ranges from 0 to 6. Note on the left side of the model the probability scale. We can see that this lambda value results in a 20% nonoccurrence outcome. Or, in other words, 80% of the time a class I attack will occur at least one time (at a rate of approximately 33%) and may occur up to 6 times within our time interval at a rate of, say, 0.01%.

Compare this to a Poisson distribution model for a Class V extreme attack. We set the lambda value to 0.0012 to reflect this. It results in a

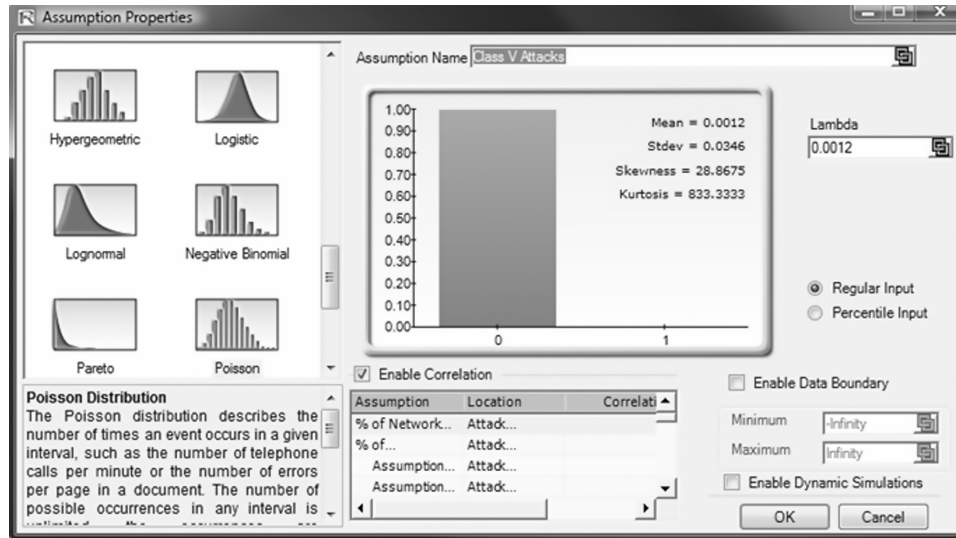


FIGURE 14.119 Poisson distribution assumption.

distribution model where this event will not occur 99.9988% of the time. There is only a 0.0012% chance that the event will occur in any given trial.

You may wonder why, if Monte Carlo simulation can be used reliably to arrive at probabilities of occurrence, we choose to use two different methods for determining probabilities. There are three primary reasons:

1. **To Reduce the “Fear of the Black Box” Phenomenon.** People who are not familiar with analytical techniques or the details associated with statistical methods have a tendency to treat analysis and the results outputs as black box-generated values. There is a natural fear to mistrust the unknown. By leveraging both statistical methods and expert opinion interviews, laypersons observing the analysis and output can rationalize in their minds how the results were generated and how they were validated or refuted. It also provides an avenue for the layperson to provide input (vis-à-vis his or her own opinions) into the equation.
2. **To Spur Additional Dialogue and Debate.** The interview process inherently spurs additional dialogue among the expert panel. My experience has been that the more divergence of opinions, the more debate occurs, which results in more robust and more refined models. The process may require more work, but, more often than not, the value of the outcome is greater than the additional effort.

3. *As a Litmus Test of Expert Opinions.* Conversely, if we rely solely on the input of expert opinions without thinking through and modeling out the statistical side of the equation, we may fall victim to tunnel vision.

While it is beyond the scope of this case study, these models could be further enhanced by creating forecasts for different types of attacks and determining the probabilities of becoming a victim for each attack. These enhancements could be realized by using historical data (what is published in the public domain along with an organization's own data).

For the purposes of simplicity, we leveraged the Delphi Method to create the various attack scenarios. The attack scenario total impact values range from \$1,547,895 to \$23,791,472—quite a significant range. How do we determine how much to invest to mitigate the risks associated with attacks?

Step 4: Determine Financial Impact

We are now ready to explore different investment scenarios to offset the risks of attacks. We now have more reliable estimates for the various classes of attacks and can take this financial information and turn it through a classical net present value (NPV) and discounted cash flow (DCF) analysis.

Our NPV/DCF analysis also has six different scenarios that will follow the same scenario structure as those previously defined. We follow this same approach through the entire analysis. It allows us to see multiple sides of the problem and arrive at a more reliable outcome.

We return to our original investment estimate (as provided by the client) of \$2,000,000, which was previously arrived at through a variety of network and systems analyses. This amount reflects the investment necessary to upgrade and enhance the intrusion management systems currently distributed throughout the environment.

At a high level, this investment will result in:

- The replacement of intrusion detection systems (IDSs) with intrusion prevention systems (IPSs).
- An increased deployment of IPS devices at additional points throughout the network—from network perimeter to network core.
- The deployment of Network Behavior Analysis (NBA) solutions at various points throughout the network along with data collection and analysis engines necessary to detect anomalies and suspect network and traffic behavior.

The logical question is: Does a \$2,000,000 investment adequately address the risks associated with the attacks and their likelihood of occurrence in this environment? Add to this: Is it too much or too little?

as our capital investment in intrusion management technologies, we can see that this scenario results in a positive NPV of \$2,228,925.15, which results in a 35.32% IRR on our investment. Clearly, this model supports a \$2,000,000 investment. This model in isolation would suggest that we could nearly double the initial investment and still have a positive NPV and IRR (the threshold to negative NPV is \$3,666,000 year 1 expense following standard computations for all other variables in the model).

Step 5: Arrive at Investment Decision

We are now near the end of our analysis. We have a solid understanding of what our current and future risks are vis-à-vis the losses we are likely to incur across a variety of attack scenarios. We know that a \$2,000,000 investment is within the range of reason. However, we also know that we could invest more and as a result further reduce our risk of losses. Alternatively, we could invest less and rely on the relatively low probability of being a target of a severe or catastrophic event.

We are at a crossroad. There is no absolute right or wrong decision for any organization. The decision makers in your organization must choose the best decision based on all of the available facts, on expert opinion, and in light of the organization’s culture.

Consider that the analysis is relatively conservative in nature. Consider that the most conservative and least biased model (the model generated by our independent expert) suggests that 80% of the time losses will be greater than \$1,857,474 (current state), and if we implement our proposed future state technology plan these losses will reduce to \$267,792, resulting in a total loss reduction of \$1,589,742. Follow this same mode of thinking and be on




FIGURE 14.122 Risk tolerance levels.

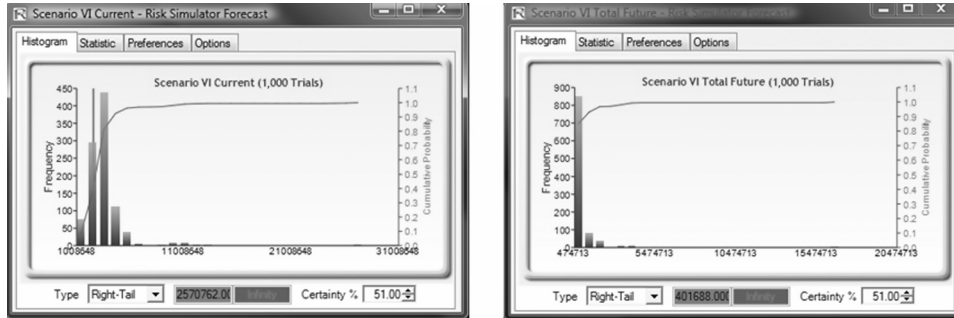


FIGURE 14.123 Simulation forecast risk tolerance levels.

the greater side of a betting man—51% of the time losses will be greater than \$2,570,762 (current state) and \$401,688 (future state), yielding \$2,169,074 in loss reductions. Figure 14.122 illustrates an example set of risk tolerance and required investment levels, and the resulting simulation forecast distributions shown in Figure 14.123 further illustrate the probabilistic levels of these risk tolerances.