

Mukul Pareek, CISA, ACA, AICWA, PRM, is a risk professional based in New York, USA. He has more than 20 years of audit and risk experience in industry and financial services. He is copublisher of the Index of Cyber Security, www.CyberSecurityIndex.org. He can be reached at mp@pareek.org.

Using Scenario Analysis for Managing Technology Risk

In the world of market and credit risk, scenario analysis is used as a part of stress testing. Stress testing is mandated by national regulators and central banks, and takes the form of asking financial institutions to consider the effect of adverse scenarios on their capital and solvency. Scenarios include historical events, such as market crashes and debt defaults, and hypothetical scenarios, such as larger-than-expected moves in interest rates, housing prices or foreign exchange rates.

In the world of operational risk, scenario analysis is used in combination with the loss distribution approach to estimate operational risk capital under the Basel framework.¹

For technology risk managers, scenario analysis can be a useful tool to identify, understand and articulate the technology risks faced by their organizations. Taken a step further, it can also be used as a tool to quantify and express a technology-value-at-risk number by expressing future losses in the form of a loss distribution.

In essence, scenario analysis consists of identifying future “what-can-go-wrong” situations that can cause a loss to an enterprise. This is something most technology risk managers already do as part of their daily task of explaining controls to business managers (e.g., when explaining risks or audit issues or when requesting new investments in security). What scenario analysis allows us to do is to consciously understand what adverse events can occur, explain how controls prevent (or, in some cases, are unlikely to prevent) unfavorable outcomes and explain how bad circumstances can get within a reasonable range of probability.

SCENARIO ANALYSIS AND THE TECHNOLOGY RISK MANAGER

There are a number of reasons why technology risk managers and analysts need to consider scenario analysis as part of their risk management tool kit:

- **Risk and control comprehensibility**—Scenarios put controls in the context of real-life situations that profit-and-loss (P&L) managers can comprehend. Scenario analysis helps create conversations that are in plain business language, as opposed to discussions about arcane control frameworks. Scenarios transform the discussion from, for example, talking about the control benefits of an identity management system to a discussion about business data that could be stolen by a competitor.
 - **Completeness of scope**—If scenarios are comprehensive and cover the risk universe against which the technology risk function provides protection, they become a useful tool for a coherent explanation of the value of the technology risk function to the enterprise. They also help set boundaries for what the function does and protects against, and set expectations for senior management.
 - **Response preparedness**—Scenarios can help enterprises plan for how to react in the event that the scenario transpires.
 - **Identification of risk drivers**—When constructed methodically, scenarios can help isolate the drivers of risk, allowing for focused action.
 - **Control effectiveness**—The identification of scenarios necessarily involves a consideration of the controls in place to prevent them from occurring, which allows a qualitative assessment of the effectiveness of controls themselves. Scenarios help us understand how controls interact with and reinforce each other.
- At its essence, scenario analysis is not all that different from risk analysis, with the notable difference being that multiple individual risks are required to combine together to create a comprehensive and plausible scenario. A scenario follows in the tradition of storytelling, whereas enumerating risk at a granular level is more of an exercise for the risk- and control-literate risk manager.



Do you have something to say about this article?

Visit the *Journal* pages of the ISACA web site (www.isaca.org/journal), find the article, and choose the Comments tab to share your thoughts.

Go directly to the article:



Enjoying this article?

- Learn more about, discuss and collaborate on risk management and risk assessment in the Knowledge Center.

[www.isaca.org/
topic-risk-management](http://www.isaca.org/topic-risk-management)

HISTORICAL VS. HYPOTHETICAL SCENARIOS

Broadly, there are two ways to identify scenarios: first, through an analysis of historical events, and second, through the construction of hypothetical yet plausible adverse events that may reasonably occur. Taken together, scenarios should

be comprehensive and updated from time to time.

Scenarios based on historical events may include real events that happened to the organization or its peers (e.g., a large compromise of its billing systems revealing sensitive

personal information). Generating hypothetical scenarios requires judgment, skill and a good understanding of the business. Hypothetical scenarios are important because they allow for the completion of the gaps left by historical events.

Scenarios based on historical experience need no explanation as to their plausibility. Hypothetical scenarios can be made plausible by seeking input from business managers who should play an active role in identifying them.

DISTINGUISHING BETWEEN EXPECTED AND UNEXPECTED LOSSES

Expected losses are losses that are considered part of the cost of doing business, and arise year after year. They are characterized by a high frequency of occurrence and a low impact. An example would be average annual credit card fraud events experienced by a bank. Up to a point, these are just ordinary losses that are absorbed as part of the cost of doing business. The product is priced to include the occurrence of expected losses. These are governed by the law of large numbers. Unexpected losses include events such as a large-scale data breach.

Scenario analysis should not cover expected losses. Scenarios should be directed toward high-severity, exceptional and infrequent events. The nature of the technology risk universe means it rarely has to deal with expected losses; nonetheless, this is an important point to make so that technology risk managers do not become too engrossed with the details of ongoing transactional events.

“Scenarios should be comprehensive and updated from time to time.”

WHAT SHOULD A SCENARIO INCLUDE?

At the very minimum, a scenario should include:

- **The situations**—An explanation of the sequence of events that leads to an adverse outcome. These may be industry- and organization-specific, but must include things such as:
 - BCP events
 - External attacks by hackers, competitors or nation states
 - Malicious insiders stealing information
 - Accidental release of confidential information
 - Vendors and third parties mishandling data
- **The outcomes**—Clearly identified outcomes that are unfavorable to the organization and are a result of the event. An event may have multiple outcomes. For example, the same scenario may result in the loss of revenue, legal costs and regulatory fines. Each outcome should be explicitly laid out, describing its impact.
- **Controls in place**—Controls work as separate lines of defense—at times in a sequential way, and at other times interacting with each other—and help prevent the occurrence of the adverse event. Often, the correct operation of just one control may provide adequate protection or mitigation. If the controls operate independently of each other, as they often do, the combined probability of all of them failing simultaneously tends to be significantly lower than the probability of failure of any one of them. An attacker, for example, who is trying to get into a network may first have the intrusion detection/prevention system (IDS/IPS) to deal with, which may have a failure rate of 10 percent. But even after the attacker gets in and tries to install a rogue program, there may be protection provided by the antimalware protection, operating at a failure rate of 10 percent, for example. The probability of both controls failing together will be only 1 percent, showing how multiple controls acting together may create a 10-fold improvement in the security, even though on their own each

may be a pretty coarse control. A third control, e.g., restricting users from administrative privileges, may further reduce the effectiveness of the attack to extremely unlikely.

- **Frequency of occurrence**—The frequency, or likelihood, of the scenario actually being realized should be a part of the scenario analysis, and is best estimated during a discussion with the business managers. What the technology risk manager is truly interested in is a probability, but the question is better framed in terms of how likely the scenario is over a long enough period (e.g., 10 years). This question is best answered by business managers, in partnership with the technology risk manager. If the answer is that the scenario might materialize once over 10 years, the probability of its occurrence each year is 10 percent.
- **Severity of the outcomes**—Much in the same way as frequency, the severity of each of the adverse outcomes should be estimated separately. Now what does severity mean? Is it the worst-case loss, or the most likely or median loss? In some cases, the absolute worst case may not be knowable, or may mean something as catastrophic as the end-of-game for the organization. Such scenarios should be modeled separately from scenarios that are expected to occur over a long-term period.

In some cases, estimating the worst case may be a meaningless exercise, as the technology risk team may not have the mandate to manage for the truly catastrophic. While this may sound surprising, it is generally not expected that technology risk management provides for events such as nuclear attacks or meteor strikes.

The technology risk analyst must strive to get at least two data points for severity—one at the 50th percentile and another at a higher percentile, such as the 90th. The question for the 50th percentile is easily posed as: What is the median expected loss level if the scenario in question were to materialize, i.e., the loss right in the middle? For the 90th percentile, it may be better to pose the question: Of all losses possible, what would be the loss if the enterprise were in the top 10 percent of the category of such losses? Precision is not desired nor should it be pursued, as it is neither achievable nor meaningful.

ADJUSTING FOR BIAS

People have a generally optimistic bias toward their perception of their own competence and good fortune. This bias is likely to be reflected in any scenario-analysis session

that a technology risk manager organizes—in the form of lower expected frequency of occurrence or severity.

One possible way to correct for this may be for the risk analyst moderating the scenario analysis not to focus on the enterprise, but to talk about similar organizations or competitors (i.e., how likely is such a scenario at the top four or five competitors, and if they were to suffer a loss, how much is it likely to be?). Any internal loss data or anecdotes of actual occurrences may help further align perceptions to reality.

COMPLETENESS OF SCENARIOS—MAPPING THE ENTIRE RISK UNIVERSE

Scenarios should cover the range of known technology risks that the business is likely to face. Documented controls should address one or more of these scenarios, and if the technology risk managers find controls that do not address a scenario, then either the universe of scenarios is incomplete or the control is redundant. Under each of the broad categories, such as process and workflow errors, information leakage events, business continuity events and external attacks (these may differ across organizations), there would be a number of scenarios.

Compiling a list of acceptable risk scenarios including all the attributes described previously is not a trivial task and requires sponsorship, cooperation from P&L managers and an understanding of the business by the technology risk manager. Scenario building may be carried out in a conference room setting with the technology risk manager or analyst leading the agenda.

In many cases, the scenario-analysis exercise is a valuable end itself. In some cases, the risk manager may choose to perform additional quantitative analysis by calculating a technology-value-at-risk number, as detailed in the next section.

COMPUTING TECHNOLOGY VALUE AT RISK

Once scenarios have been identified, together with their expected frequency and severity, as explained in the previous section, these estimates can be converted to estimates of losses at different confidence intervals, similar to value at risk. We will call it the technology value at risk to distinguish it from the more common measure of financial risk.

The steps to determine a technology-value-at-risk number are:

1. **Assume a distribution for the frequency and severity estimates.** Generally, the Poisson distribution for frequency and the lognormal distribution for severity are reasonable

choices. Using a distributional approach recognizes that both frequency and severity are not single-point estimates, but can cover a wide range of possible values. This makes the calculation process more acceptable in a management discussion, as the technology risk analyst is not claiming certainty in any calculations.

The remainder of this article will proceed with these distributional choices (i.e., Poisson and lognormal), though the overall process would be quite similar even if other distributions were selected.

For the frequency distribution modeled by the Poisson distribution, there is only a single parameter, the mean, that is required to build the distribution. The mean was estimated as part of the scenario analysis exercise.

For the severity distribution modeled by a lognormal distribution, two parameters are needed to describe the complete distribution: the mean and the standard deviation. Estimating these will require the availability of

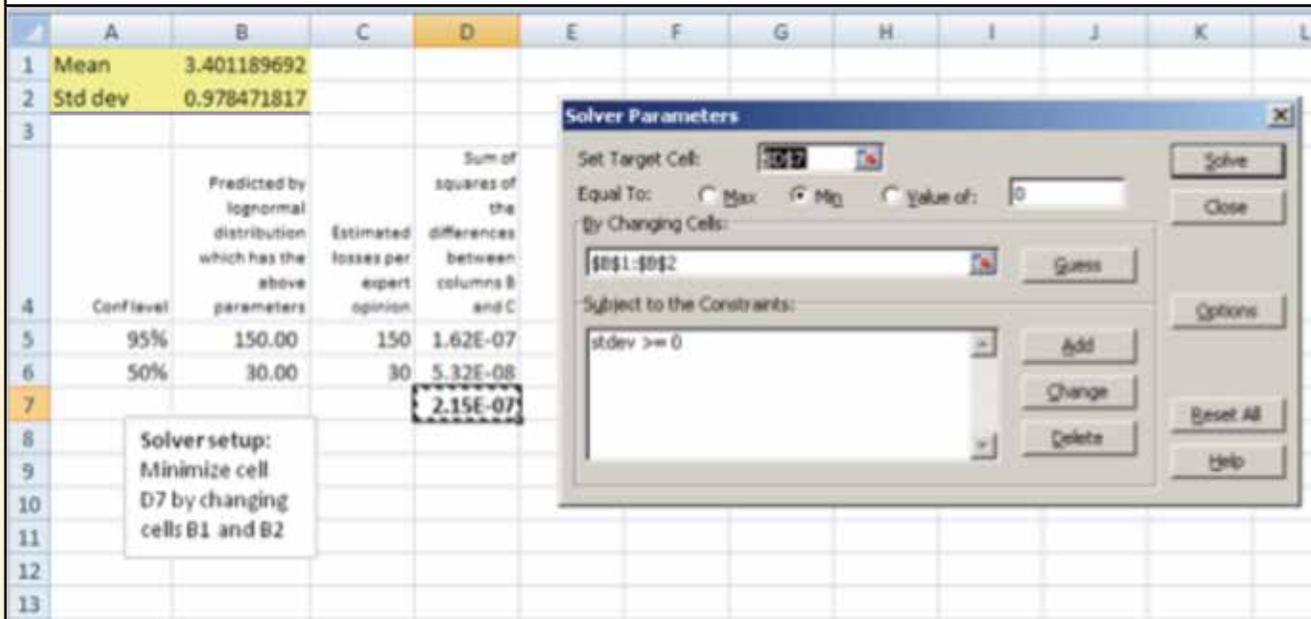
two (or more) data points for the losses that were estimated as part of the scenario analysis exercise—the most likely loss and the loss at the 90th (or another) percentile. Using the method of least squares, the two-point estimates can be used to estimate the best fitting mean and standard deviation for the severity distribution. This can be done in Excel, using Solver² (see **figure 1**), or using a mathematical package such as R³ (see **figure 2**).

- 2. Build the loss distribution.** The loss distribution is a product of the frequency and severity distributions, much in the same way as loss equals frequency multiplied by severity. While there is no way to formulaically multiply the Poisson (for the frequency) and the lognormal (for the severity) distribution, one can use a Monte Carlo simulation⁴ to obtain the loss distribution. This requires picking a random number from each of the frequency and severity distributions and multiplying them to get a single data point representing a loss. This process is then repeated thousands of times to get enough data points to produce a loss distribution.

Figure 1—Practical Modeling Using Excel

Step 1: Determine distributional parameters for the lognormal distribution using ordinary least squares.

In Excel, the mean and standard deviation for the lognormal distribution can be obtained using the Solver add-in. An illustrative example appears here.



Step 2: Use Monte Carlo simulations after distribution parameters have been estimated.

Random numbers from both the Poisson and lognormal distributions to simulate frequency and severity may be generated in Excel using the Analysis Toolpak, a standard Excel add-in. The data points for the loss distribution are obtained by multiplying severity with frequency. The technology value at risk can then be calculated at the desired confidence level, e.g., at the 99th percentile the loss will be =PERCENTILE(data_range_loss_column,0.99).

3. **Calculate the technology-value-at-risk number for the scenario.** Once the loss distribution has been obtained as a large set of data points, the technology value at risk for the particular scenario can be determined by calculating the quantiles in which one is interested. For example, if one wishes to calculate the loss at the 99th percentile, one would look at the loss level below which 99 percent of all losses lie.
4. **Aggregate the technology value at risk.** As an additional step, an aggregate technology-value-at-risk number that includes all the scenarios may also be calculated by doing a Monte Carlo simulation for all the scenarios simultaneously. This is assuming that the loss events are independent and not correlated.

These steps can be performed in Excel, or in a mathematical package such as R. While Excel is a great environment for prototyping and solving less-complex problems, R is more suitable to heavy-duty work. The decision of which to use would depend upon how widely and repeatedly the technology risk manager needs to use the risk model, and available skill sets.

To summarize, the technology-value-at-risk calculation includes the following steps, as visualized in **figure 3**:

1. Identify scenarios.
2. For each scenario:
 - Determine frequency as a single-point estimate. Use this estimate as the mean for a Poisson distribution that models the likelihood of the scenario occurring.
 - Determine severity as a point estimate at two quantiles or more. Using these data points, calculate the mean and standard deviation of the closest lognormal distribution. This lognormal distribution now defines our severity distribution.

- Simulate the loss distribution, picking one point each simultaneously from both the frequency and severity distributions
3. For each scenario, calculate the appropriate percentile (usually the 95th or 99th) as the technology value at risk.

Figure 2—Practical Modeling Using R

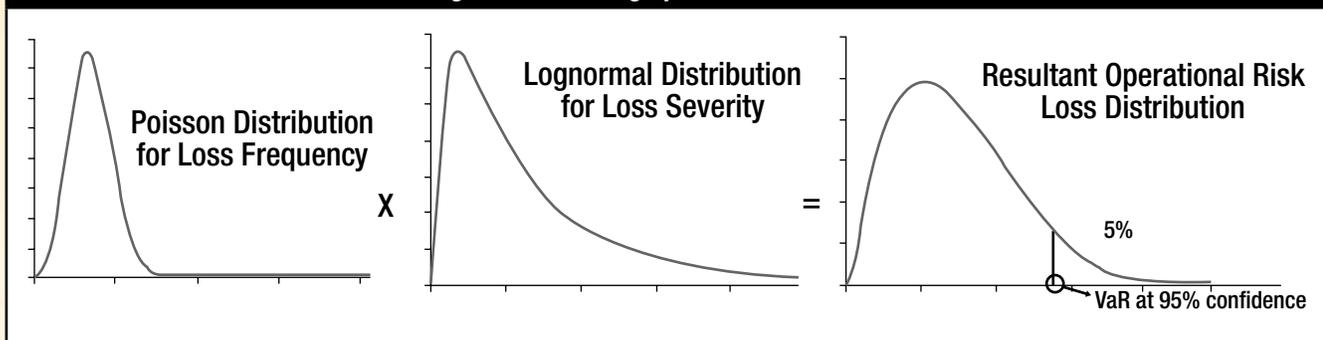
In R, the steps can be performed using the following commands. The initial variables will need to be set by the risk analyst before running these commands:

```
#Initial variables
simulations <- 1000000
lambda <- 0.1
exp_est1 <- 10
exp_est2 <- 80
conf1 <- 0.5
conf2 <- 0.95

#Two estimates for Scenario 1 placed in dataframe s1
s1 <- data.frame(conf=c(conf1,conf2), expert_est=c(exp_est1,exp_est2))
#Setting a function to calculate the sum of squares
ss <- function(x) {
  x1 <- x[1]
  x2 <- x[2]
  sum((qlnorm(s1$conf,x1,x2) - s1$expert_est)^2)
}

#Minimizing the sum of squares function a
pp <- optim(c(0,1),ss)
s1mean <- pp[[1]][1]
s1stdev <- pp[[1]][2]
#Calculating the loss distribution function
ld <- rlnorm(simulations,s1mean,s1stdev)*rpois(simulations,lambda)
qt <- quantile(ld, probs=c(0.95, 0.99, 0.995, 0.999))
#Publish everything we calculated thus far
s1
s1mean
s1stdev
ss(c(s1mean,s1stdev))
qt
```

Figure 3—Modeling Operational Risk Losses



4. To calculate an aggregate technology value at risk that includes all scenarios, simulate a loss from all scenarios simultaneously.

CONCLUSION

Scenario analysis, even if carried out without any additional quantification, can be a useful exercise to bring together technology risk practitioners and the business that they serve. It can generate the right conversations and engagement and focus management on issues that truly matter to the organization. It can also help evaluate controls in the context of real business situations, and help identify controls that can be safely dropped without an inordinate increase in the risk. If scenarios are converted to a technology-value-at-risk number, the enterprise gets the additional benefits of being able to evaluate the monetary impact of adding or removing controls.

Yet the approach is not without limitations. Real life is complex, and adverse outcomes inevitably compound. Additionally, the impact from scenarios often extends beyond technology. It is difficult to successfully model strategic, legal

and reputational risk areas that often accompany technology risk events. A modeler would need to bear these limitations in mind as part of any scenario analysis.

ENDNOTES

- ¹ Basel Committee on Banking Supervision, *Basel II: International Convergence of Capital Measurement and Capital Standards: A Revised Framework—Comprehensive Version*, www.bcbs.org
- ² Solver is a native Microsoft Excel add-in that allows complex problems to be solved using optimization routines. It may be enabled under the Add-Ins menu in Excel.
- ³ R is a popular open-source software used for mathematical and statistical analysis. It can be downloaded from cran.r-project.org.
- ⁴ Monte Carlo simulations are a statistical method where data points are obtained by repeated random sampling. This allows for simulating complex systems and interactions that may be difficult to express analytically (e.g., as a clean formula).