feature

Cyberrisk Assessment Using Bayesian Networks

Organizations are increasingly realizing that the management of cybersecurity risk in complex environments needs to be addressed using suitable decision-making techniques. They are slowly embarking on the journey of quantifying their exposure to cybersecurity threats (operational risk) in much the same way they quantify credit and market risk exposure. While there is a wealth of data and well-established statistical methods for calculating credit and market risk, no such data or methods have been explored for quantifying cybersecurity risk.

Traditional cyberrisk assessment methodologies generally use a likelihood/impact-based risk model to arrive at risk ratings. While useful as a starting point, such models suffer from serious deficiencies including:

- Calculating the probability/impact is often oversimplified and may not put much thought into what lies under the hood.[1]
- Risk is not always independent. For example, speed of delivery and quality of delivery are always linked. Yet poor quality and missed delivery usually appear as separate risk factors in risk registers, giving the illusion that one can be controlled or mitigated independently of the other.[2]
- Visualization tools such as a heat map draw attention to the top right quadrant (high consequence and high likelihood), while items in other quadrants, especially low likelihood and high consequence risk, are generally ignored.[3]
- Risk scoring in the traditional approach represents only one possible outcome. In fact, operational risk can have a wide range of outcomes, i.e., a distribution of outcomes where each potential outcome has a corresponding probability.
- Failure to address key concerns such as:
  - What critical causal factors apply to specific risk factors
  - How to quantify risk reduction by implementing specific controls

This article seeks to address these issues using a causal probabilistic model (called a Bayesian network [BN]) that is based on Bayesian inference. BNs can capture the complex interdependencies among risk factors and can effectively combine data with expert judgment. BNs can provide rigorous risk quantification and genuine decision support for risk management.

Bayesian Networks

BNs, also known as belief networks (or Bayes nets), belong to the family of probabilistic graphical models (PGMs). These graphical structures are used to represent knowledge about an uncertain domain. PGMs with directed edges are generally called a directed acyclic graph (DAG), which is popular in statistics, machine learning and artificial intelligence. "A BN is a visual description of the relationships between cause and effect. It is made up of nodes and arcs, and each node in the network represents a variable, and the arcs represent the causal relationships between the variables."[4] BNs use Bayes' theorem to compute the probabilities in the model. Bayes' theorem is written as:

$$p(A|B) = p(B|A) * p(A)/p(B)$$

p(A|B) is the posterior, i.e., the probability of event A occurring given that event B has occurred.

p(A) is the prior, i.e., the probability of event A occurring.

Venkatasubramanian Ramakrishnan, CISM, CRISC, CHFI
Is a head of information risk management at Cognizant (Chennai, India), focusing on strategy and planning of cybersecurity, information risk management, audit and compliance, and business resiliency programs. Ramakrishnan was instrumental in getting Cognizant shortlisted for the IRM Global Risk Awards 2013 in the category "Risk Management Solution of the Year."

p(B|A)/p(B) is the evidence, i.e., the probability of event B occurring given that event A has occurred, divided by the probability of event B occurring.

The following example uses the Bayes' theorem. In a small town, a particular disease has a 1 in 1,000 (.1 percent) rate of occurrence. A screening test for the disease is done using a diagnostic tool that is 100 percent accurate for those with the disease and 95 percent accurate for those without the disease. The values for the Bayes' theorem are as follows (**figure 1**):

- p(A) = prior probability of the disease occurring (.001)
- p(not A) = prior probability of no disease (.999)
- p(B) = probability of test accuracy
- p(B|A) = probability of test accuracy given the disease (1)
- p(not B|A) = probability of test accuracy given no disease (.95)
- p(A|B) = p(B/A) * p(A)/p(B) = probability of disease given the test is positive
- p(B) = p(B|A) * P(A) + p(not B|A) * P(not A)

The probability of a person having the disease, even if the test is positive, is just around 2 percent, whereas the intuitive answer is 95 percent. This outcome clearly indicates that mere "gut feelings"-based reasoning or back-of-envelope calculation is fundamentally inadequate for risk measurement.

Bayesian Networks in the Cyberdomain

Threats to data security are continuing to evolve. Much of the danger comes through the Internet, which is a vital component of today's business infrastructure. In this risky environment, employees are increasingly using the Internet to communicate, collaborate and access data. Productivity is booming, but this also makes it easier for employees, suppliers or hackers to access, copy or lose intellectual property or customer data that may have a severe impact on organizations.

This section examines some hypothetical cases related to the risk of data loss from the perspectives of insiders and external cyberattackers. One of the values of doing this is that it enables common causal drivers of each to be identified and prepared for adequately. Once the structure is fleshed out, it can readily be used to actively and dynamically assess risk. The practical application of such a model can be used for a cyberrisk such as data leakage.

Figure 1—Probability of Disease Using the Bayes' Theorem

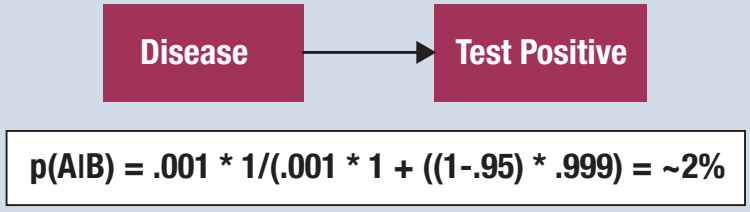

Source: V. Ramakrishnan. Reprinted with permission.

The following models (**figures 2**, **3** and **4**) highlight the different risk levels for data leakage, assuming different levels of control effectiveness and the corresponding interactions.

Figure 2—Best Case (Controls Functioning at High Levels of Effectiveness)

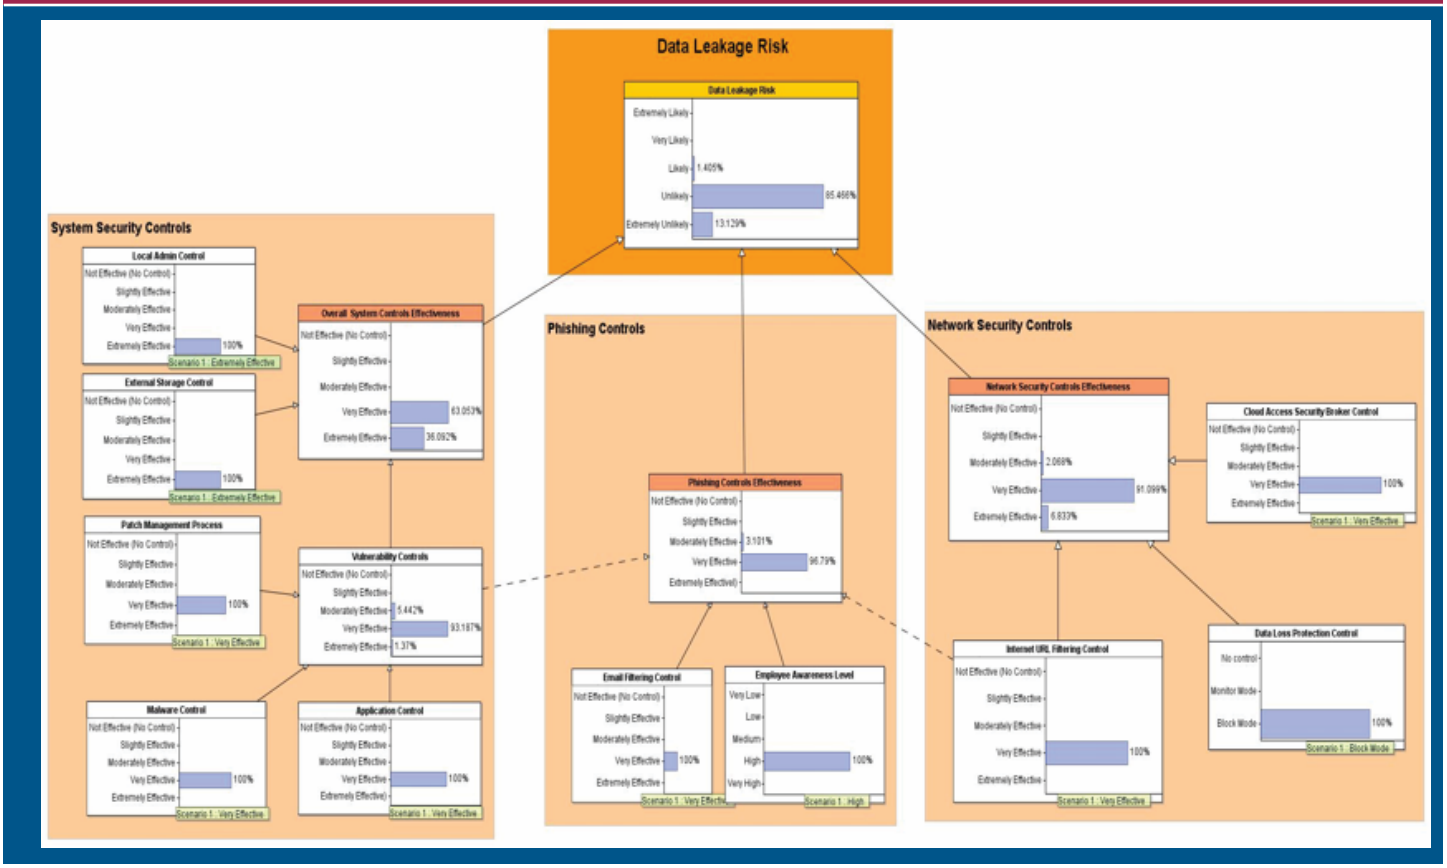

Source: V. Ramakrishnan. Reprinted with permission.

Figure 3—Moderate Case (Controls Functioning at Moderate Levels of Effectiveness)

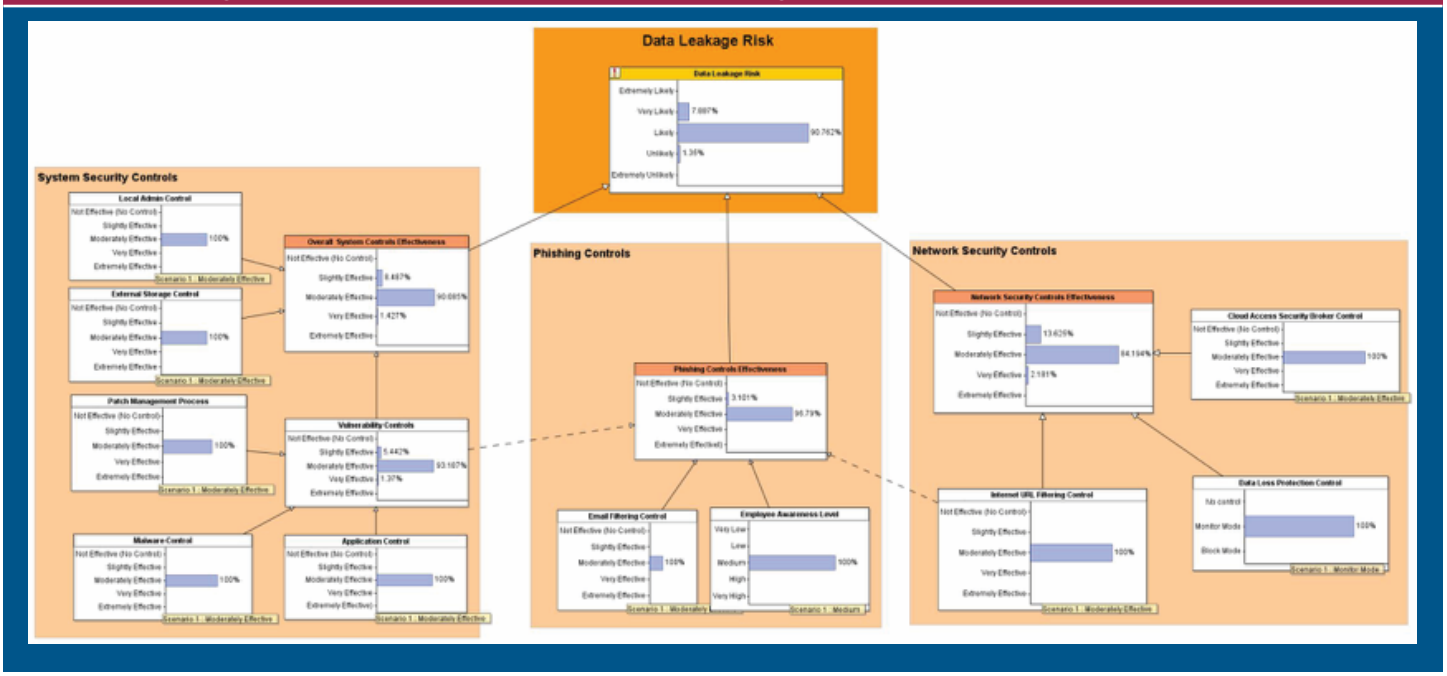

Source: V. Ramakrishnan. Reprinted with permission.

Figure 4—Worst Case (Controls Functioning at Poor Levels of Effectiveness)

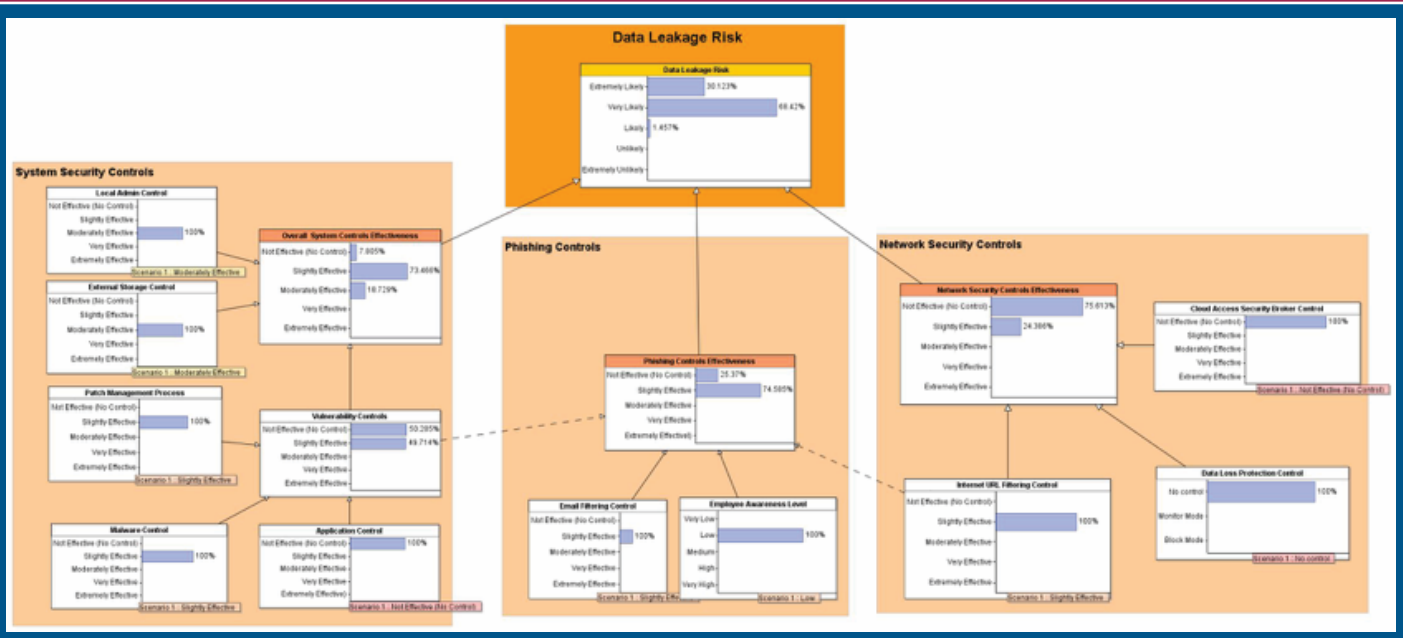

Source: V. Ramakrishnan. Reprinted with permission.

Using different levels of control effectiveness (and adjusting corresponding probabilities), the probabilities illustrated in **figure 5** can be derived.

The models shown in **figures 6**, **7** and **8** highlight the indicative loss due to data exfiltration by an attacker due to vulnerability and control weakness on the web server. Essentially, the model separates three important types of uncertainty: the uncertainty of attack success, the uncertainty of attacker choice, and the uncertainty from a security information and event management (SIEM) system.[5] Though the probability of compromise is the same for all scenarios, the loss amount reduces significantly from US $33 million to US $7 million when the effectiveness of the security operation center (SOC) team is high.

Using different levels of SOC team response effectiveness, the loss amounts shown in **figure 9** can be derived.

Figure 5—Scenarios and Data Leakage Probability

Scenarios	Control Effectiveness	Probability of Data Leakage (approximate)	Risk Rating
Best	• Very effective controls with a few extremely effective controls • Data loss prevention (DLP) in block mode • Employee awareness high	85 percent unlikely; 13 percent extremely unlikely; Less than 1.5 percent likely	Low
Moderate	• Moderately effective controls with a few not effective controls • DLP in monitor mode • Employee awareness medium	91 percent likely; 7.9 percent very likely; Less than 1.4 percent unlikely	Medium
Worst	• Not effective controls with a few controls slightly effective • No DLP • Employee awareness low	68 percent very likely; 30 percent extremely likely; Less than 1.5 percent likely	High

Source: V. Ramakrishnan. Reprinted with permission.

Figure 6—Best Case (SOC Team Skills and Response Effectiveness High)

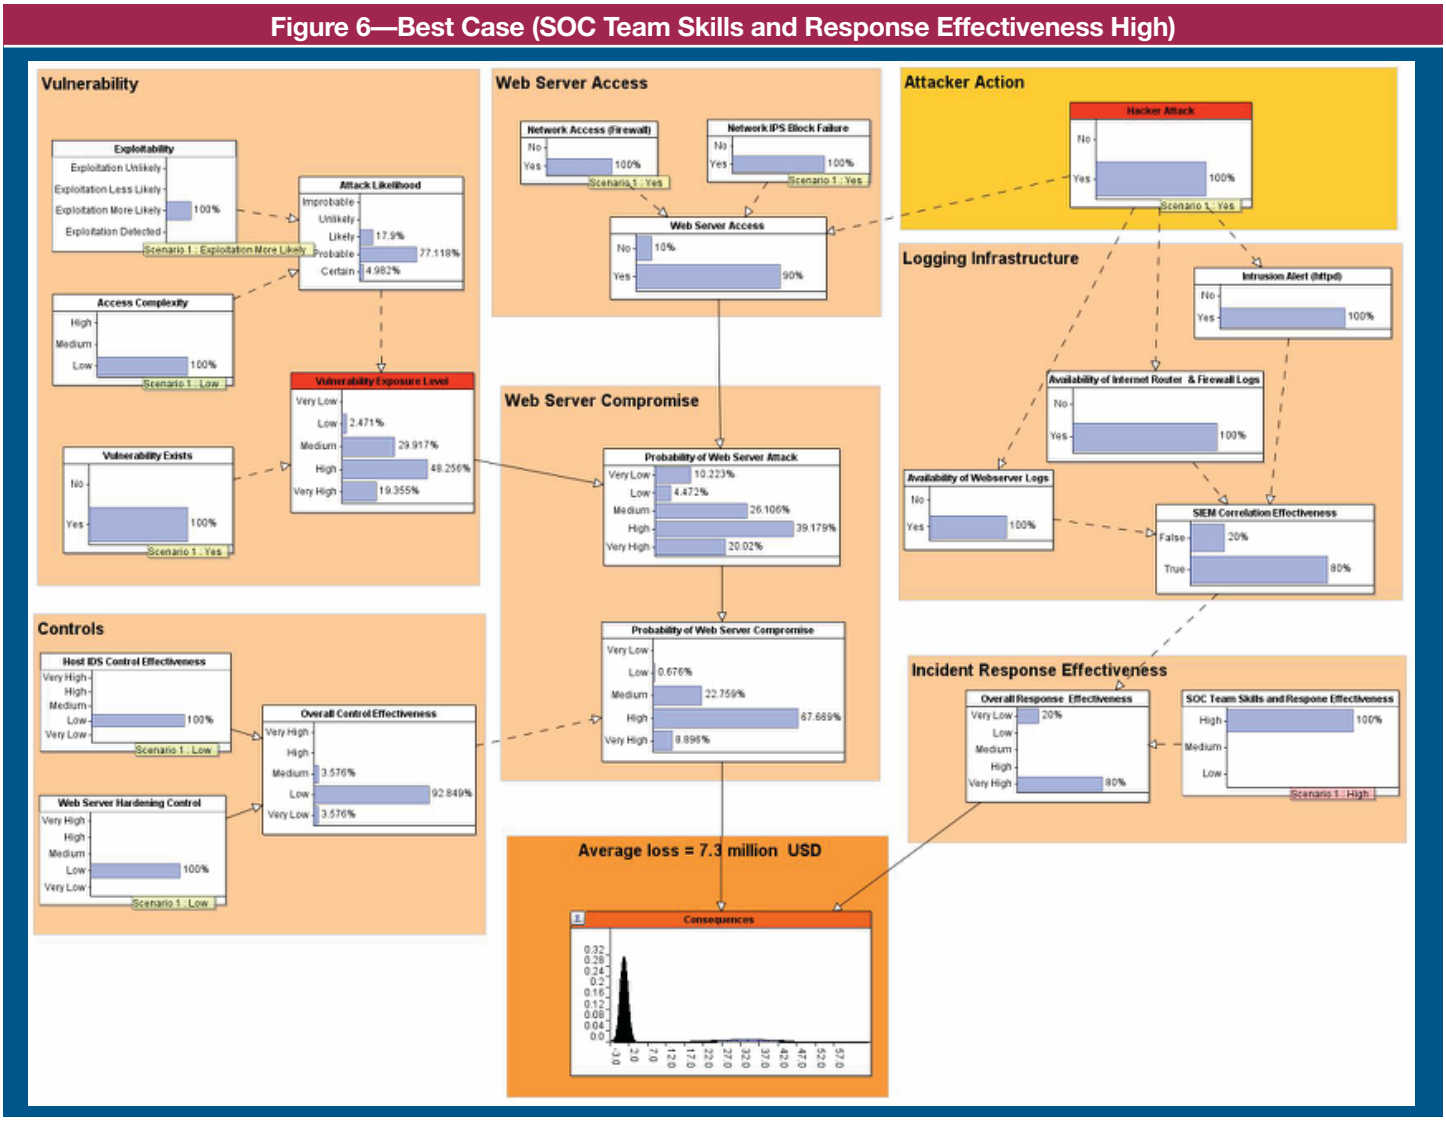

Source: V. Ramakrishnan. Reprinted with permission.

Figure 7—Moderate Case (SOC Team Skills and Response Effectiveness Medium)

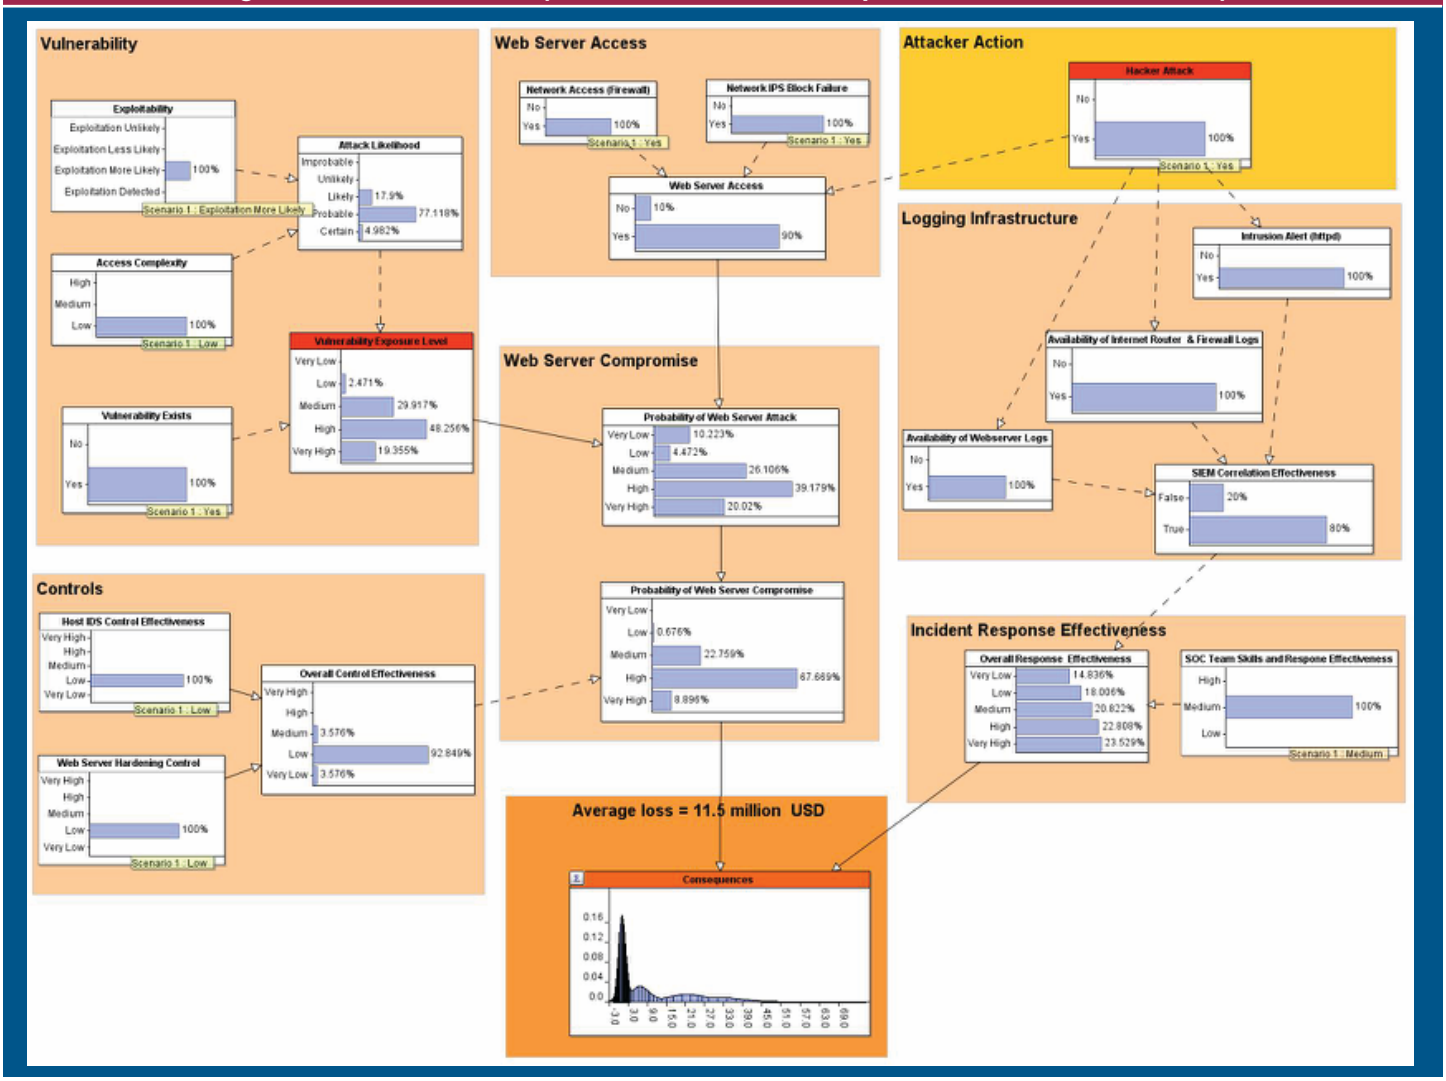

Source: V. Ramakrishnan. Reprinted with permission.

Figure 8—Worst Case (SOC Team Skills and Response Effectiveness Low)

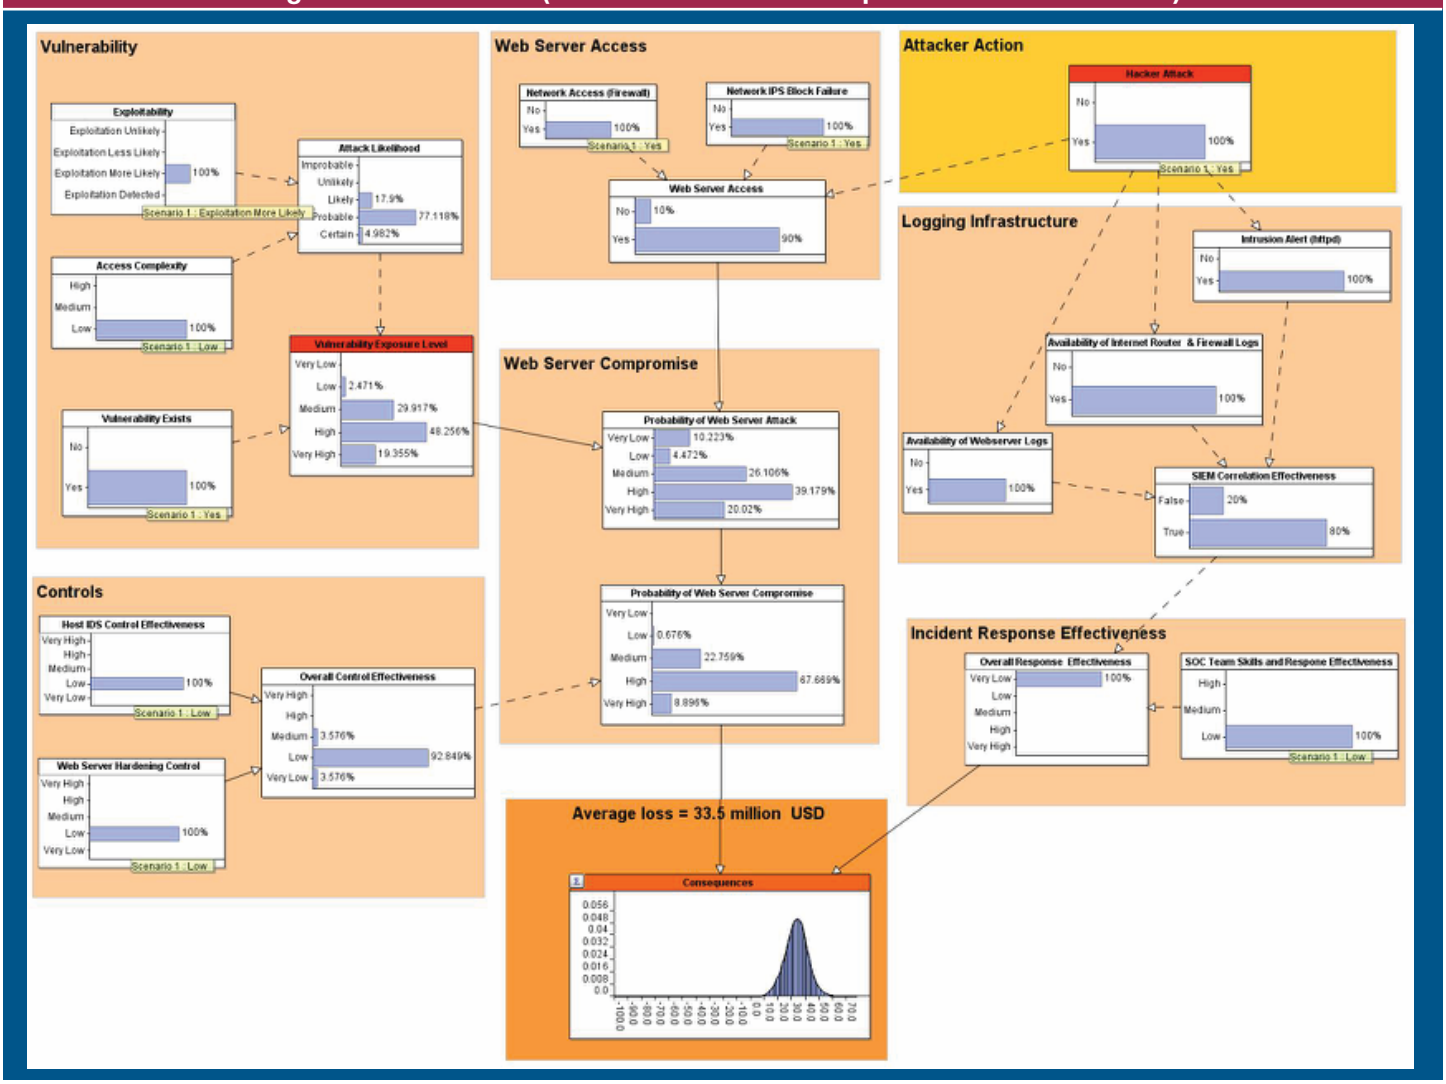

Source: V. Ramakrishnan. Reprinted with permission.

Figure 9—Loss Amounts Based on SOC Team Response Effectiveness

Scenarios	SOC Team Response Effectiveness	Loss Amount (in Millions of US Dollars)
Best	High	7.3
Moderate	Medium	11.5
Worst	Low	33.5

Source: V. Ramakrishnan. Reprinted with permission.

Conclusion

The BN approach helps to identify, understand and quantify complex interrelationships and can help make sense of how risk factors emerge and are connected, and how to represent control and mitigate them. By thinking about the hypothetical causal relations among events, alternative explanations can be investigated, and it is possible to evaluate the consequences of actions and identify unintended or undesirable side effects. Having said that, it is important to make judgments about how deeply some risk factors are modeled and how quickly this analysis informs actions.[6]

> **"The BN approach helps to identify, understand and quantify complex interrelationships and can help make sense of how risk factors emerge and are connected."**

Until now, usage of such an approach has been explored in the areas of operations risk. Extending such an approach to the cyberrisk domain can result in the following key benefits:[7]

- A meaningful explanation of how outcomes are directly related to the risk drivers
- Identifying and capturing where and how risk mitigation actions can reduce likelihood and impact
- Helping perform a what-if analysis to test alternative strategies to reduce overall risk and ways to measure impact of the strategy adopted
- Helping quantify the loss amount
- To visualize the dependency factors
- To build a storyboard for communication

Author's Note

All views expressed in this article are those of the author and do not necessarily represent the views of his employer. The author does not necessarily endorse or recommend the use of any particular software through this discussion, but merely proposes a method for risk assessment using BN.

Endnotes

1 Neil, M.; N. Fenton; *Risk Assessment and Decision Analysis With Bayesian Networks*, CRC Press, United Kingdom, 2012
2 *Ibid*.
3 Samad-Khan, A.; "Modern Operational Risk Management," *Emphasis*, 2008
4 *Ibid*.
5 Xie, P., *et al.*; "Using Bayesian Networks for Cyber Security Analysis," IEEE/IFIP International Conference on Dependable Systems & Networks (DSN), 2010, *http://people.cis.ksu.edu/~xou/publications/dsn10.pdf*
6 *Op cit,* Neil
7 *Ibid*.

Enjoying this article?

- **Learn more about, discuss and collaborate on risk assessment in the Knowledge Center.** *www.isaca.org/topic-risk-assessment*