Communicating Technology Risk to Nontechnical People

Helping Enterprises Understand Bad Outcomes

Também disponível em português
www.isaca.org/currentissue

Too often, well-meaning technology professionals attempt to explain risk to their enterprises and fail to achieve their objective. These professionals fully understand the state of the computing environment and the importance of securing it. They may even have a relevant third-party affirmation of their beliefs through the US National Institute of Standards and Technology (NIST), the International Organization for Standardization (ISO), COBIT® or other standards and frameworks. However, they are unable to motivate their nontechnical colleagues to recognize the importance of what they are trying to communicate.

Failures in communication have been studied extensively and are included in any introductory course on the subject. To identify the reasons for failure to communicate technology risk in particular, the body of knowledge related to cybersecurity must be temporarily abandoned to explore what makes communication in general work well and what causes it to fail.

Jack Freund, Ph.D., CISA, CRISC, CISM, CISSP
Is a leading voice in cyberrisk measurement and management. He is the coauthor of an award-winning book on cyberrisk quantification and holds a doctorate in information systems. Freund is an International Association of Privacy Professionals (IAPP) Fellow of Information Privacy and Fellow of the FAIR Institute. In 2018, he was the recipient of ISACA's John W. Lainhart IV Common Body of Knowledge Award and the FAIR Institute's FAIR Champion Award.

Models of Communication

No survey of communication would be complete without reviewing the seminal model proposed by Aristotle.[1] This model has three parts: the sender, the message and the receiver. Most important is the receiver, who ultimately determines whether communication has taken place. This simple model identifies at least one part of the failure to communicate technology risk. Absent the executive team's reception of the message, communication cannot happen, regardless of intentions.

In the context of information technology, it is appropriate to consider Claude Shannon's classic information theory model, published in 1948.[2] Shannon applied mathematical theory to communications, leading to concepts such as signal-to-noise ratio. Indeed, the noise part of the model can help explain problems in technology risk communication. Business executives have so many competing priorities that it is often difficult for technology professionals to rise above the noise and get their point across. Business management is largely risk management, so for an individual security-related message to resonate with decision makers, it has to compete with market risk, credit risk, competitive risk, regulatory risk, conduct risk, reputational risk and all other forms of operational risk. Numerous reports indicate that executives consider cybersecurity a top priority, so clearly they are not ignoring it.[3,4,5] However, the voluminous number of bad things routinely brought to their attention by IT professionals belies their experiences. In fact, when executives weigh these potential calamities against actual incidents, many of them conclude that IT

professionals are prone to "Chicken Little-ism," or prognosticating doom that never happens (fear, uncertainty and doubt [FUD] in other words).[6,7,8]

> "THE MESSAGE BECOMES WATERED DOWN WHEN NUMEROUS SO-CALLED CRITICAL MATTERS ARE COMMUNICATED BUT RARELY RESULT IN ACTUAL PROBLEMS OR INCIDENTS."

IT professionals are rarely told to their faces that they are Chicken Littles. Instead, they have to interpret the feedback they receive. Thus, the modern communications model differs from the Shannon model in that it contains an explicit feedback loop (**figure 1**).[9]

Business executives employ technology professionals to identify problems and raise important issues. If all technology problems are treated as critical, the result may be apathy. The message becomes watered down when numerous so-called critical matters are communicated but rarely result in actual problems or incidents. Executives are surely aware that bad things can happen, and they may even have peers in other organizations or industries who have experienced bad outcomes, but they probably have little personal experience. Executives desire better information about cyberrisk, but they often assume that the issue is so complex that even the people they hire to deal with it are incapable of doing better. Executives tend to understand the systems that need to be online to serve their customers and the systems that cause regulators to get upset. However, experience tells them that even if they ignore the critical broken things, nothing bad is going to happen.

Too often, executives' subtle and not-so-subtle messages are poorly received by technology professionals. Instead of changing the message to ensure that the receiver better understands it (by casting the message in terms the receiver cares about), technology professionals may become petulant and secretly wish for a security breach to prove them right. Such sullenness may compel IT professionals to send decision makers articles about bad things that have happened elsewhere.

Rectifying these communication failures involves looking at how risk is communicated and how cybersecurity can be made relevant to business executives, starting with how potential risk is communicated to them.

Lists of Risk vs. Risk Scenarios

Too often, technology professionals use confusing terminology to discuss risk. As a result, a risk assessment often looks like a collection of things that are broken; groups of people who could do harm; and abstract, esoteric or even existential notions of consequences.[10] Such a list of risk factors might look like this:

- Privileged insiders
- Reputation
- Untested system recovery process

Figure 1—Modern Communications Model

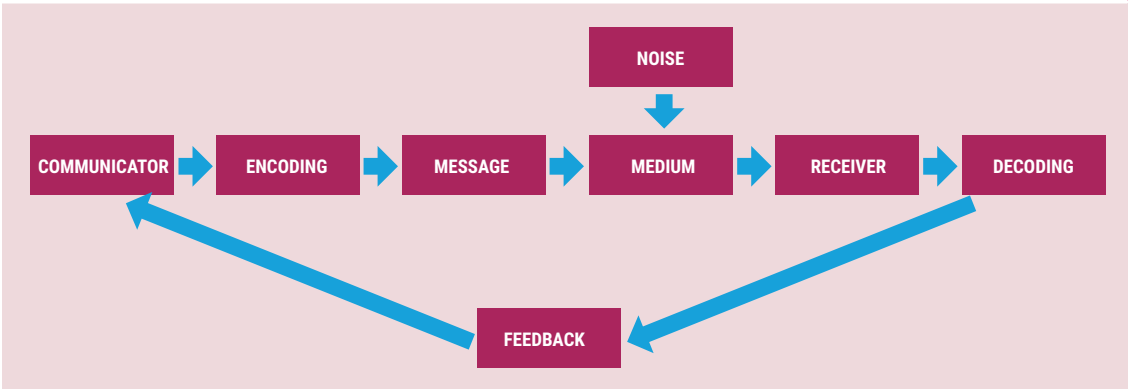

> WHEN COMMUNICATING RISK, IT IS IMPORTANT TO REMEMBER THAT MOST PEOPLE HAVE AN INCOMPLETE UNDERSTANDING OF STATISTICS, SO STATISTICAL LITERACY CANNOT BE ASSUMED.

- Cloud data shares with sensitive data
- Short passwords
- Cybercriminals

It is easy to see that each item on the list is something that might cause concern. However, technology professionals use a kind of shorthand when communicating with other similarly trained and liked-minded professionals, whereas business executives are forced to fill in the blanks with their imaginations (guided by experience). In a fully qualified risk statement, these missing parts are articulated so that they are easily understood by individuals who are unfamiliar with the shortcuts of the profession.

It is important to clearly communicate to the target audience which items on the list are threats, assets and controls (however weak they may be). Executives must understand how the combination of these categories of things can be manipulated to cause harm to the enterprise.

The first step in improving risk communication is to ensure that there is a fully defined risk scenario to which a risk formula can be applied.[11] Each risk scenario statement should tell a story that is instantly accessible to nontechnical people. For example, such a scenario might be: "Privileged insiders leverage legitimately granted credentials to steal data from critical applications." It specifies who is doing something bad, what methods are being employed to do it and how the organization will be impacted once it is done. A proper risk scenario needs to be forward looking. It should describe a series of bad things that might come to pass, not necessarily something that is happening currently. A good risk statement is also relatively perpetual; if an item can be removed from the risk register after something has been fixed, it is a control deficiency, not an actual risk.

The Classic Risk Formula

The classic risk formula (probability multiplied by impact) can be confusing to those receiving technology risk communications. Consider the compounded problem associated with determining both the probability and the impact of privileged insiders (from the earlier sample scenario). Asking executives to interpret the probability of insiders as 0.45 does nothing to improve communication. The probability of what, exactly? This statement does not help the receiver understand the problem.

When communicating risk, it is important to remember that most people have an incomplete understanding of statistics, so statistical literacy cannot be assumed. As a result, the use of concepts such as the basic risk formula can lead to incorrect calculations along with imperfect communication. The first problem is that the terms "likelihood" and "probability" are used interchangeably when speaking and writing. This does nothing to further mutual understanding.

Next, probability is not temporally bound.[12] It is entirely unhelpful to tell executives that the probability (or likelihood) is 40 percent. Alongside the "probability of what?" question mentioned earlier is the obvious question of when. Is it 40 percent probable that this event will happen today? This week? This year? This decade? Time matters, and taken by itself, this value does not effectively communicate what executives need to know about the probability or likelihood of risk realization.

To overcome this problem, many people apply fixed timelines to their estimates. They describe these values as representing annualized probabilities. Unfortunately, there is a fundamental mathematical problem with these kinds of assessments: What if the event could happen more than once per period? It is mathematically unsound to assert that something has a 200 percent likelihood of happening in the next year, as probability is a value between 0 and 1. And that value is non-inclusive: Probability can never be 0 or 1, because a future event can never be ruled in or ruled out with 100 percent certainty.

The foregoing issues can be overcome by utilizing frequency in place of probability in the equation.[13] This accomplishes several things. First, frequency is a much more accessible concept to represent

future events. For those who are uncomfortable with statistics, it is better to ask them how often something might happen rather than the probability of its happening. Second, this variable is better able to capture events that occur more than once per year (or period). A frequency of two per year is easy to comprehend, whereas a 200 percent probability is not only mathematically incorrect but also difficult to understand practically. Additionally, probability values of less than 1 (e.g., 0.5) are more easily recognizable as frequency values and can be communicated in plain language (e.g., once every two years).

Finally, and most important, the simple risk formula does not contain guidance on exactly of what one should assess the probability and impact. Knowing what to measure is just as important as knowing how to measure it. Risk is about loss, so whatever is being measured must be a complete statement of loss relevant to the enterprise. The list of technology-related risk presented earlier is a classic example of things that are not business risk factors because they do not express a complete loss scenario.

Business Process Mapping

Some enterprises may be unfamiliar with business process mapping. However, business continuity teams may already have some version of it. If so, their mappings and process inventories are a good place to start, requiring fewer resources and supporting a single source of information on processes in the enterprise. To initiate business process mapping from scratch, the sampling approach should be followed, and key products and services and the critical processes for each should be the focus. The first year, a sample size that is doable should be chosen, and a plan for increasing the number of samples each year and determining the resources required should be created.

Business process mapping is the first step in creating a fully qualified risk scenario. This requires understanding how enterprises operate and connecting technology to business offerings. It also requires a list of the products and services the enterprise offers (or reasonable groupings of them). This list can often be compiled by considering what is offered in each line of business or some other category in large enterprises (such as geographic location). Then the parts of the enterprise that help deliver each product or service are linked. Considering the business processes that enable each part of the enterprise is helpful. Finally, a connection is made between those business processes and the technology that enables them. The result looks something like **figure 2**.

> "BUSINESS PROCESS MAPPING IS THE FIRST STEP IN CREATING A FULLY QUALIFIED RISK SCENARIO."

Once there is a connection between technology (the tech stack) and products and services (the business stack), it is time to develop the risk scenarios that affect each. This helps decompose the process map into more detailed scenarios. In general, applications are the primary interface between enterprises and their technology and, as such, they serve as the nexus that connects the tech stack to the business stack. Some business processes are enabled by simple applications, such as email. In this case, the supporting infrastructure that enables email is also aligned with the business process and, ultimately, with the products and services that process enables. This provides a sense of what kind of technology-related problems can arise and how they can affect the enterprise and its offerings. Incidentally, this model works for both for-profit and nonprofit, and public- and private-sector enterprises. In all cases, an enterprise exists to offer something, and technology is aligned with those offerings to enable them. In some cases (such as the email example), technology is aligned with multiple business processes and corresponding products and services. Once this mapping of offerings and technology is complete, risk scenarios can be created.

Developing Fully Qualified Risk Scenarios

There are different levels of scenarios, depending on where in the business process map the scenario exists. For instance, at the very top (e.g., board reporting), there are likely to be only a handful of aggregate scenarios. Scenarios in the middle parts of the enterprise (e.g., senior management, heads

Figure 2—Business Process Mapping That Connects Products and Services to Technologies

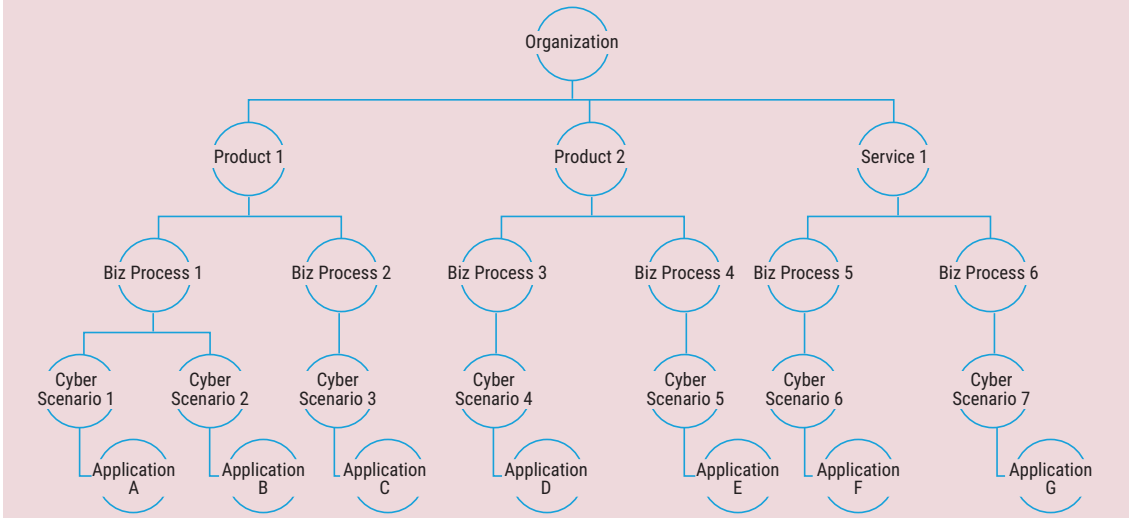

of various lines of business) will include additional decompositions of those aggregate scenarios that are linked to specific products and services. At the very bottom, there will be many versions of cyberscenarios that trigger upper-level scenarios.[14] An example of this kind of decomposition is presented in **figure 3**.

When designing top-tier risk categories, it is important to consider the specific business in which the enterprise is engaged. However, one can start with the following Basel II event categories, even for enterprises that are not involved in financial services:[15]

1. Internal fraud
2. External fraud
3. Employment practices and workplace safety
4. Clients, products and business practice
5. Damage to physical assets
6. Business disruption and system failures
7. Execution, delivery and process management

Most enterprises will have some version of categories 1, 2, 5, 6 and 7 that covers their technology risk. An example of applicable risk categories (based on **figure 3**) would be the following:

1. Data loss and theft
2. Data reliability
3. System availability
4. Fraud

"WHEN DESIGNING TOP-TIER RISK CATEGORIES, IT IS IMPORTANT TO CONSIDER THE SPECIFIC BUSINESS IN WHICH THE ENTERPRISE IS ENGAGED."

These four risk categories are suitable for board-level reporting. They can then be decomposed into product- and service-specific versions that reflect scenarios in a particular line of business as shown in **figure 3**.

Although such labels are helpful for grouping risk, these categories need to be decomposed one more level to get a fully qualified risk scenario that provides a greater degree of precision in the risk assessment. For instance, "Theft of data from critical applications" is a useful category, but it does not provide enough detail about what is happening, who is doing it, and how to assess risk factors and the efficacy of controls. A fully qualified risk scenario might be: "Privileged insiders leverage

Figure 3—Decomposition of Risk Categories to Scenarios

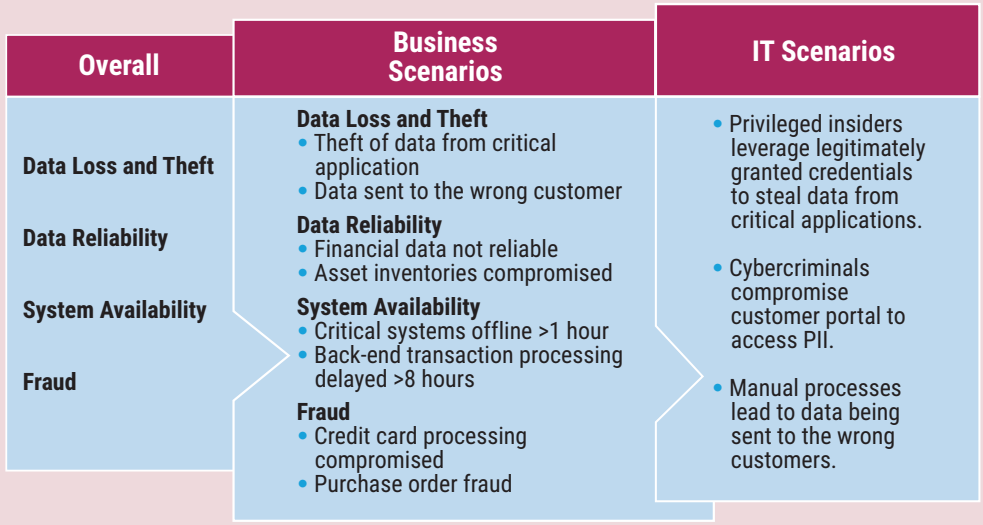

legitimately granted credentials to steal data from critical applications."

This statement reveals several critical things. First, it states who is taking the action. Next, it states how they are accomplishing it. In this case, the enterprise has already granted these individuals the tools they need to perpetrate bad acts, which are also clearly identified as stealing data from critical applications. Most important is that the statement tells a story, and this type of narrative ensures that communication is clear and complete. As a category of loss scenario, "data loss" is useful, but the phrase may conjure different images to different people. A fully developed risk scenario articulates the specific way in which data loss occurs.

The next step is to connect the loss scenarios to the relevant technology assets. To accomplish this, it is necessary to identify the inherent attributes of those assets that connect them to the scenario. For instance, the preceding sample scenario would require only a single attribute: users permitted to see sensitive data. This is similar to the way insurance underwriters use demographic information to determine insurance premiums. Here, these inherent attributes link the right risk scenarios to the assets that could bring them about. Also, because the risk scenarios are worded in such a way that the risk formula can be applied accurately, technology assets can be linked, via their demographics, to risk ratings that represent how loss could occur in that system. The scenario tells a narrative that is specific to the tech stack and that can be aligned with the risk categories reported up through the enterprise.

Top-Down vs. Bottom-Up Risk Assessments

Bottom-up risk assessments are typically acknowledged to be far more complete than top-down assessments. However, because bottom-up assessments require the collection of large amounts of information from various technologies and individuals, most enterprises consider them overly time consuming, possibly resulting in an incomplete assessment before the due date for reporting.

Top-down risk assessments, in contrast, have the reputation of being fast and easy. They require fewer resources to accomplish and can provide meaningful results. They are, however, subject to the bias of the people conducting them at the top, who are usually disconnected from the day-to-day problems and risk scenarios that are well known to those at the bottom.

In practice, those performing audit functions typically do not suffer from these either/or scenarios. They acknowledge that they cannot possibly assess everything, and they select samples at the bottom for the categories at the top on which they want to report. Such a sampling approach can be very helpful for enterprises trying to bridge the gap between top-down and bottom-up risk assessments. Sampling, in

Enjoying this article?

- Read *Bridging the Digital Risk Gap. www.isaca.org/Bridging-the-Digital-Risk-Gap*
- Learn more about, discuss and collaborate on risk management in ISACA's Online Forums. *https://engage.isaca.org/onlineforums*

conjunction with the risk scenario decomposition outlined earlier, provides the tools needed to confidently report on the state of risk in an enterprise.

The Sampling Approach to Risk Assessment

To use the sampling method:

1. Select a handful of samples from the lowest level to inform each top- and intermediate-level risk category. For example, in the data loss and theft category, correlate several intermediate-level risk statements from each business unit (e.g., theft of data from critical applications).
2. Select the cyberscenarios and tech stacks linked to them. This results in several specific, low-level resource stacks to assess. These can become the risk ratings used to justify the ratings applied at other levels.

This approach allows a top-down-style risk assessment with the benefit of assessing risk at the bottom to validate those ratings.

Initially, these samples should cover critical and key applications and infrastructure, but over time, an enterprise can sample most of its technology environment. For example, it can sample top applications in the first year, followed by second- and third-tier applications in the following years. The resulting risk assessments should include a scoping statement indicating that the rating is based on a sample (e.g., 15 percent) of critical applications and infrastructure. Such scoping can also be included in annual strategic plans, and any additional sampling requested by an enterprise can help security and risk leaders prepare better budgets for resources to support these requests.

Addressing "Broken Things"

Too often, lists of "broken things" find their way onto organizational risk registers. A good rule of thumb is that if an item in the register can be checked off, removed or completed with the right configuration change, technology or process implementation, it is probably not a risk scenario and does not belong on the risk register. However, these lists of broken things are very important to the overall risk management capability of an enterprise. Alongside each level in the risk scenario hierarchy, there should be a corresponding list of broken things, at increasing levels of detail as one goes down the list.

For example, a list of missing patches or misconfigured servers should not be on the risk register. Instead, they should reside on their own list of problems requiring attention, such as an issue management register or a break/fix register. These individual items can be categorized at an aggregate level in a way that allows them to be linked to cyberscenarios. For instance, several users may have been overprivileged with access to critical applications. This can be categorized as unnecessary permissions or privilege creep. That category can be aligned with the cyberscenario of "privileged insiders misusing legitimately granted permissions," for example. At a higher level, such broken things can be grouped in a category called "identity and access management."

> "BOTTOM-UP RISK ASSESSMENTS ARE TYPICALLY ACKNOWLEDGED TO BE FAR MORE COMPLETE THAN TOP-DOWN ASSESSMENTS."

Risk Ownership

It is a popular notion that a business entity "owns" risk. What this means in practice is that every item in the risk register must be aligned with an owner who is not in IT, risk management, cybersecurity and so forth. The risk should be aligned with someone responsible for the products and services articulated in the business process map. This represents a significant culture shift for most enterprises. For many, it is anathema to think that an IT professional does not "own" a data loss and theft risk. More to the point, IT may "own" a series of what operational risk professionals call risk triggers or a causal taxonomy, such as those

"unnecessary permissions" mentioned earlier. Ultimately, the loss is owned by those responsible for the products and services affected. An important side effect of allocating risk ownership this way is that the assessment of a risk scenario varies from one business unit to the next. The amount of loss associated with customers is likely to vary significantly from one product to another. Thus, a risk statement can appear on multiple internal risk registers, likely with different risk ratings. For all such risk factors aligned with business units, it is important to assign a liaison person to act as a bridge between IT and the business unit to help with communication and translation of IT terminology and to assist with risk treatment decisions, including following up on fixing the broken things aligned with these risk statements.

Conclusion

Ultimately, technology programs exist so that enterprises can deliver the products and services for which they are chartered. With rare exceptions, enterprise leaders are not experts in delivering technology solutions. Every profession has its own language, acronyms and shorthand that enable professionals to communicate with one another expediently. However, IT is a profession that exists to serve an organizational objective and, as such, it needs to adjust its communications to help organizational leadership achieve their goals.

Being better aligned with the enterprise allows for better value creation, facilitates the perception of competence and alleviates internal feuds that distract from delivering on customers' expectations. Rearranging IT risk reporting to better align with the enterprise's understanding of its purpose and priorities improves communication and provides decision makers with the information they need to be better managers.

Endnotes

1 Griffin, E.; *A First Look at Communication Theory, 6th Edition*, McGraw-Hill, USA, 2006
2 Shannon, C. E.; "A Mathematical Theory of Communication," *Bell System Technical Journal*, vol. 27, iss. 3, 1948, p. 379–423, 623–656

> "A RISK STATEMENT CAN APPEAR ON MULTIPLE INTERNAL RISK REGISTERS, LIKELY WITH DIFFERENT RISK RATINGS."

3 ISACA®, *State of Cybersecurity 2019*, *www.isaca.org/cyber/Documents/state-of-cybersecurity_res_eng_0316.pdf*
4 Risk.Net, "Top 10 Operational Risks for 2019," 14 March 2019, *https://www.risk.net/risk-management/6470126/top-10-op-risks-2019*
5 Marsh & McLennan, Microsoft, "2019 Global Cyber Risk Perception Survey," September 2019, *https://www.microsoft.com/security/blog/wp-content/uploads/2019/09/Marsh-Microsoft-2019-Global-Cyber-Risk-Perception-Survey.pdf*
6 *Ibid*.
7 Copeland, J.; "No Time to Talk Cyber Risk, Senior Executives Say," Fair Institute Blog, 19 September 2019, *https://www.fairinstitute.org/blog/no-time-to-talk-cyber-risk-senior-executives-say*
8 Jones, J.; "Jack Jones: Quit Blaming Executives for Cybersecurity Problems," Fair Institute Blog, 19 August 2019, *https://www.fairinstitute.org/blog/quit-blaming-executives-for-cybersecurity-problems*
9 Laws, S. M.; "Corporate Communication: Identity, Image and Reputation," *International Journal of Business Competition and Growth*, vol. 3, iss. 4, 2014, p. 344–349
10 Freund, J.; J. Jones; *Measuring and Managing Information Risk: A FAIR Approach*, Butterworth-Heinemann, USA, 2014
11 Maurice, D. R.; J. Rathod (ed.); "Cybersecurity and Technology Risk," *Operational Risk Perspectives: Cyber, Big Data, and Emerging Risks*, Risk Books, UK, 2016
12 *Op cit* Freund, Jones
13 *Ibid*.
14 Freund, J.; "Keep It Simple: How to Avoid Drowning in Cyber Risk Information," Risk.net, 2017, *www.risk.net/risk-management/3938516/keep-it-simple-how-to-avoid-drowning-in-cyber-risk-information*
15 Bank for International Settlements; " QIS 2–Operational Risk Loss Data," 2001, *https://www.bis.org/bcbs/qisoprisknote.pdf*