5. Risk Analysis and Assessments

Learning objectives of this chapter: Risk management framework, methods of risk analysis and assessments — quantitative and qualitative risk analysis — fault tree analysis, event tree analysis, failure mode effects analysis (FMEA), continuous risk assessment, vulnerability assessment, penetration/security testing.

Fear, uncertainty and doubt (FUD) have over time proven to be improper drivers for an organization to determine a proper security posture. A salubrious alternative is to carry out an assiduous assessment of risks to determine mechanisms to manage the said risks, be it to invest in controls to mitigate some or them, or to make a conscious decision to accept some others.

To secure, or not to secure

There are numerous dilemmas on whether and what to invest for security, arising from a plethora of reasons. Donald Rumsfeld's quote on 'unknown unknowns' aptly encapsulates one of the dilemmas. We may not even know what are all the security threats, and what are all the consequences — so, how do we determine the amount of resource to dedicate for security, and how do we ascertain if it is adequate? One may argue that security is priceless, and we must commit as much resource as possible. A contrary opinion may be forwarded by arguing that, given that perfect security is impossible to achieve in any case, so why bother at all? Instead of such a pessimistic line of reasoning, one may also get complacent based on past success. If the existing security measures have so far prevented major security incidents, then one may assume from the lack of security incidents that the corresponding expenses to deploy and operate security controls are in fact needless, and thus may consider them for avenues for budget cuts.

There are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns ? the ones we don't know we don't know. - Donald Rumsfeld

Deficient security may also result from a decision process which sees security as a obligation than a need aligned with one's business objectives. This naturally leads an organization to aim for a lowest common denominator — for instance, by being standard compliant but not doing anything beyond, or by doing something just about better than a competitor, and thus have a false sense of security by falling in the trap of the fallacy of relative deprivation. Such self-delusionary claims of best practices may work as a public relations gimmick, but do not provide robust security, nor facilitate a systematic way to address security concerns.

Beyond the question of how much resources ought to be assigned for security purposes, one also needs to determine how much of the said resource is to be allocated across controls that are preventive, detective or responsive in nature.

Revisiting the maxim that there is no perfect security, one may then conclude that a pragmatic way to approach security is to prioritize — to decide resource allocations within a constrained security budget. This, assessing and managing risks, by prioritizing the allocation of resources to optimize the cost-benefits, is the crux of information security. Blakely et al [1] rationalize that since information security concerns the protection of business-critical or sensitive information and related IT systems and infrastructure, failures of information security will trigger adverse events, resulting in losses or damages that will exert negative impacts on a business. Information security must be a risk management discipline that manages risks by considering their costs and/or impacts on a business. In other words, 'information security is information risk management'. It is worth adding a caveat that viewing information security from such a pragmatic cost-benefit trade-off is relatively easier for commercial enterprises than say, an issue of national security, though ultimately, the principles of-course apply — and the challenge in practice emanates from the difficulty of putting value to things.

[1] Bob Bakely, Ellen McDermott, and Dan Geer. Information security is information risk management. In *Workshop on New security paradigms*, 2001

At this juncture it is also worth emphasizing that, though we have recurrently stated that standards are regulations provide a minimal baseline, they also provide guidance to cary out risk assessment in a structured manner, by identifying critical assets that need to be protected, or the typical threats to be guarded against, and so on. Furthermore, the risk analysis also needs to take into account the cost of being in breach of a regulatory requirement. For instance, if electronic protected health information (ePHI) under Health Insurance Portability and Accountability Act of 1996 (HIPAA) is breached, an organization may have to endure fines - and the quantum can be directly attributed as part of the consequent loss.

Risk analysis

There are various ways to carry out risk analysis. But in all cases, the basic issues to consider include identifying what (asset) needs to be protected and the nature of associated threats and vulnerabilities. Accordingly, one needs to determine the consequences of a security breach, typically by assigning a value to the asset in question, and accordingly planning the necessary controls to minimize the loss or damage. The desired outcome from a risk analysis is a set of recommendations to maximize the protection of security objectives while still providing functionality and usability. Such an analysis also needs to take into account the budget constraints — to determine how to allocate limited resources to maximize gains, but also to justify the necessary budget to achieve specific security goals.

The process to determine the risks should be inclusive in nature, and derive its inputs from the various organizational levels. This is important because different stakeholders may be aware of or perceive different threats. For instance, even if a hospital's website server is in a demilitarised zone serving simple web pages, and does not expose any confidential business or patient information, a CEO may consider possible defacement of the webpage bad for public relations and brand image. The CIO in the meanwhile may be worried about ensuring HIPAA compliance, and thus emphasise on encrypting backup tapes before storing hem offsite. In the meanwhile, a medical equipment operator may be worried that a drug pump in the hospital can be hijacked to tweak the amount of insulin to cause patient fatality.[2] Each of these issues are of genuine concern, though with varied severity, but may not be the concern that is naturally foremost in mind for every member of the organization. Once the inputs from different stakeholders is aggregated, one needs to determine whether and what remedial or mitigating measures to pursue for each concern. The risk analysis process can then be seen as a composition of (i) scoping, (ii) data collection, (iii) vulnerability analysis, (iv) threat analysis and ultimately (v) identification and analysis of acceptable risks.

[2] In June 2015, security researcher Billy Rios reported that several models of drug infusion pumps made by Hospira can be hacked to surreptitiously and remotely change the amount of drugs administered to a patient. http://www.wired.com/2015/06/hackers-can-send-fatal-doses-hospital-drug-pumps/

Scope of risk analysis: Oftentimes, there is a specific purpose for crying out a risk analysis, which then determines its **scope**. Depending in what the analysis is for, it will also determine which all aspects and degree of detail in which the analysis needs to be carried out. Which systems and applications to consider. The scope also depends on who the analysis is for. If the analysis is for the CEO and CFO to approve IT security budget, the necessary details could be pretty high level. For instance, they may need to know that the organization is

not using any firewall, and accordingly a budget to procure the same is required. In contrast, an IT system administrator will require low level actionable information, for instance, even if the firewall already exists, may be it is not properly configured, and recommendations to set it up correctly will be interesting for that purpose. The scope of the analysis also needs to determine whether to consider threats from the outside, by attackers with no a priori special privileges, or from the inside, with various possible extent of access rights.

DATA COLLECTION: Access to the right kind of information is essential for a proper risk analysis. Accordingly, a major part of the risk analysis process is about data collection. This includes determining the existing procedures and policies in place, interviewing the personnel to obtain their respective perception of threats and vulnerabilities, but also to capture known and typical issues, as well as identify if they meet the existing policy requirements. Manual as well as automated tools driven tests to determine if all systems are updated with latest software patches, or whether there are any known and common vulnerabilities (say, determined using a repository of known vulnerabilities) also need to be carried out. Penetration testing, to determine if a specific resource can be compromised, with or without prior knowledge of and access to the system internals also need to be conducted.

VULNERABILITY ANALYSIS: Ultimately, the purpose of the vulnerability analysis is to determine the likely impact it is likely to have, which can then be used to rank the vulnerabilities, and prioritize remedial actions. One way to determine the likely impact of a vulnerability is to determine how easy or difficult (in terms of the needed resources) it is for an adversary to exploit the said vulnerability (**severity**), and if the vulnerability is in fact exploited, then how bad are the consequences (**exposure**). These two aspects can then be combined to determine an overall vulnerability rating. FMEA under ISO 27000 adapts and extends this idea, which we shall discuss subsequently.

THREAT ANALYSIS: It is also important t understand what the threats might be. These may be human or non-human threats, targeted specifically with purpose, accidental or just random. For instance, a hacker trying to steal credit card data from a particular retailer is a human threat targeted specifically at a victim, while, a misconfigured software losing data is an accidental event. Data server destroyed in a flooding accident is an example of non-human threat, while the data server destroyed by a random virus infection is an example,

where the virus was man-made, but the attack was not specifically targeted at the particular business or server. All these kinds of threats need to be considered and categorised. This helps in determining the motivation and resources that adversaries may have, to cross-reference these against the vulnerabilities, in order to establish the likelihood of actual security incidents.

Identification of (acceptable) risks

Having identified the potential impact emanating from breaches of various vulnerabilities, it is then pertinent to categorise these in two (or three) groups — vulnerabilities with unacceptable consequences (very high severity and exposure) which need to be eliminated, risks with high impact, which call for risk reduction, and finally accepting some risks (low severity and exposure) with only baseline protection. It is to be noted that the assignment of quantum for severity and exposure itself is a subjective issue, and the organization carrying out the risk analysis may need to make a call on that. When prioritising mitigating controls, one may also need to take into account other aspects, such as risk urgency and any interdependencies and collaterals.

Quantitative risk analysis

The discussion above provides a qualitative treatment for risk analysis. However, that may not be adequate for decision support, and one may need a more objective and quantified way to approach the analysis. Ideally, it should be derived from some irrefutable facts, offering direct projection of costs and benefits of any mitigating actions. An objective and quantified risk analysis is also a powerful selling tool to the management, since it is less prone to disagreements as is bound to happen with qualitative measures. In practice, the challenge may lie in finding all the necessary facts, and make the correct estimates to put a monetary value to everything in order to get a precise cost-benefit analysis.

As an example, consider the following hypothetical setup. BuyAnyTime Inc. is an online retailer, gearing up for Christmas sale, and for the increased traffic at its site for a period of one month (30 days), it expects the following: An average of 100 transactions per minute and an average profit of $10 per transaction. Based on the current IT infrastructure, and the expected high load during the festive season, an average downtime of 30 minutes per day is expected, when the online site will not be usable. It was also determined than an upgrade of the IT infrastructure will however cost $800,000 and it would reduce the system downtime to 5 minutes per day. So the management

of BuyAnyTime Inc. has a very simple issue to resolve - should they, or shouldn't they invest in the system upgrade?

Clearly, the security attribute under compromise here is availability. But the question at hand is, is it worth fixing the problem, even though it is acknowledged that there is a problem. A simple calculation in this hypothetical setup can help the management make up their minds. 30 minutes of system downtime suggests 30×100 unsuccessful transactions per day, so $30 \times 30 \times 100$ unsuccessful transactions over the month, which in turn means $\$30 \times 30 \times 100 \times 10$ of missed profit, which essentially means an effective cost to business of $900,000. At close to a million dollars, that's doubtlessly some serious amount of money. So one then needs to determine the situation of intervening action is taken. With only 5 minutes downtime per day, lost profits will be of the tune of $\$5 \times 30 \times 100 \times 10$, which is $150,000, however, we also need to account for the upfront investment of $800,000, and thus the effective cost to business works out to be $950,000. So finally, in this hypothetical scenario, the decision is easy to make — based on the information at hand, it is not worth fixing the problem.

But risk analysis is not a straightforward issue in real life, so why shall we settle for something too simple even for a hypothetical example. So let's instead add an additional element to the problem. Say, BuyAnyTime Inc. does a consumer survey to determine, that in fact, if they have more than 10 minutes downtime per day, their reputation will suffer, leading to 2% customer attrition, which then is expected to lead to a proportional decrease in the volume of transactions throughout the day. In light of this new information, how does the cost analysis work out? The effective cost to business with the IT infrastructure stays unchanged, since it does lead to a downtime of less than 10 minutes. However, a 2% attrition implies an additional cost of $\$1410 \times 30 \times 2 \times 10$ (1410 remaining minutes per day when the system is up, 30 days, 2 lost transactions per minute). This cost, in addition to the $900,000 lost during the downtime, adds up to $1,746,000. Equipped with this additional information, the management can revisit and alter their original decision.

This extremely simple hypothetical example already illustrates several interesting challenges of quantitative risk analysis. Firstly, incomplete information can still lead to misleading decisions. Secondly, there could be interdependencies, that may need to be identified, and it may not in fact always be easy to do so. Third, it may be difficult to quantify everything. For instance, in the example above, the dependency was with respect to the company's image - however, we just implicitly assumed that this image will decay by 2% because of downtime excessive of 10 minutes per day, a figure that may not be

so easy to estimate precisely in practice.

In general, quantitative risk analysis also needs to take into account different countermeasure strategies which can lead to different payback and cashflow scenarios, considering also any resulting secondary effects. It may need to discern the impacts of long term benefits versus short term benefits, as well as discern one time investments (e.g., infrastructure upgrades) versus recurrent costs (operational costs, regular penetration tests by security consultants, etc.).

The computations used in the above example are in fact construed (though relevant for the specific setup), and we will next discuss a well recognized (though with its own limitations) metric for quantitative risk analysis, namely annualized loss expectancy (ALE).

ANNUALIZED LOSS EXPECTANCY: The expected monetary loss in one year due to a risk is defined as the Annualized Loss Expectancy (ALE), which in turn is determined by how often a loss event occurs, i.e., the Annual Rate of Occurrence (ARO) and the monetary loss expected from the occurrence of the risk once, i.e., Single Loss Expectance (SLE), specifically $ALE = ARO \times SLE$. SLE itself is in turn determined by the Asset Value (AV) being affected by the risk, and degree of affectation (e.g., percentage of asset lost) known as Exposure Factor (EF), so $SLE = AV \times EF$.

By assigning a monetary value to a risk, and particularly by doing so for a period of an year, ALE acts as a simple by effective way to aid decision support. Particularly, it provides a good guideline as to the overall annual budget (typically capped at the ALE figure) for security controls which will be worth putting in place, to eliminate a given risk.

There are however several shortcomings with this approach - the ARO may be hard to predict and may have high variance (the later making average a less meaningful metric), and AV and EF may be hard to quantify in many cases, and are ultimately determined in a subjective manner. In addition to these aforementioned shortcomings rooted in uncertain or imprecise information, even if/when one is equipped with precise ARO, AV and EF information, ALE blurs the distinction between high-frequency, low-impact events from low frequency, high-impact events. In practice, such distinction may however have profound impact, and thus may need to be accounted for in the decision making process.

Ultimately, any one of these means is inadequate in making robust risk analysis, and hence a hybrid approach deploying multiple qualitative as well as quantitative techniques for risk analysis in conjunction is typical.

Vulnerability assessment

As discussed previously, rigorous vulnerability and threat analysis are key to a proper risk analysis. We will next discuss some standard mechanisms which are often used for the purpose of identifying vulnerabilities, but also to prioritise corrective measures, based on factors such as the likelihood of occurrence and the degree of adverse impact from a risk.

When one wishes to ascertain if a specific (set of) resource(s) is/are secure, one may commission a penetration testing, whereby, someone (typically a security consultant) tries to find ways to compromise the target(s). Since it is a goal oriented exercise, once (if) a compromise of the target is achieved, the test can be terminated. At this stage, a report detailing what all was attempted, which security controls were found to be adequate, and what vulnerabilities were discovered and exploited is presented. Additionally, recommendations are suggested to prevent further attacks using same/similar modus operandi. However, neither the test, nor the recommendations systematically identify or eliminate all vulnerabilities, and it is possible that some other means to compromise the target remain undetected. Even if at the end of a penetration testing, the target has not been compromised, the same issue of not having explored systematically all possible ways to attack the target persists. So to say, penetration testing does not lend itself any method for exhaustive exploration of vulnerabilities in a self-contained manner. We next elaborate a few generic techniques, which provides the tools to do so. The caveat being, how well the tool is/can then be used in turn still remains a somewhat open-ended issue.

Fault tree analysis:

Fault tree analysis (FTA) is a top down, deductive failure analysis in which an undesired state of a system is analyzed using Boolean logic to combine a series of lower-level events.

For instance, consider the undesired system state of unauthorized access to an individual's emails. The question at hand then is to deduce the different manners in which this may occur. In this case, we consider the proximate cause being a compromise of the user's logic/password credential (password guessed), and two factor authentication being rendered ineffective (ineffective 2FA). Since both of these events must occur before unauthorized access of emails become feasible, it is demonstrated using an AND gate. We will like to point out that the example fault tree is in fact not exhaustive in nature. For instance, instead of compromise of the individual user's login/password and 2FA, if the mail server itself is compromised

in other ways, then too, an individual's emails will be exposed. We do not show such other situations, which a more rigorous fault-tree analysis needs to consider. We then continue the deduction by determining in turn the proximate cause for each of these factors. In the chosen example, we claim that password is guessed if the attacker succeeds in achieving any one of the following three — use a key-logger against the victim, crack the password (e.g., using a dictionary attack) or obtain the password from the victim using a social engineering technique such as phishing. Since any one of these mechanisms suffice, we join them using an OR gate, and so on and so forth.

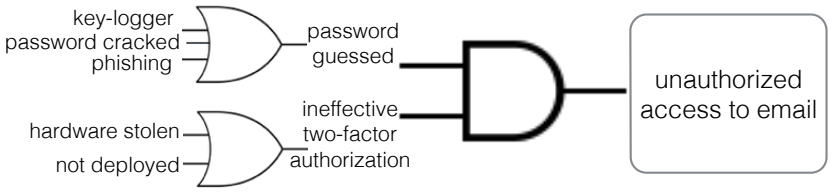

Figure 1: A non-exhaustive fault tree analysis example, using a scenario for unauthorized access to an individual's emails.

ATTACK TREE ANALYSIS:

Attack trees are similar to fault trees, but there are several nuanced differences. It is a conceptual diagram showing how an asset, or target, might be attacked, possibly qualifying an attack in multiple dimensions, making in information rich. Numerous kinds of information, for instance, whether an attack is motivated or opportunistic in nature, what kind of access (internal or external) is needed to carry out an attack step, resources or skills needed on an attacker's part, and so on, may be captured in an attack tree representation. Such information can then be used to better plan controls. For instance, it is useful to determine if a specific attack requires a phishing email, in which case the attacker is more likely to do it, than if it requires the attacker to carry out a physical burglary. The controls will thus be determined based on the expected degree of risk aversion on the attacker's part.

The structure of the tree is determined in a fashion similar to a fault tree. Starting with the undesired (top) event, the possible proximate causes of that event are identified at the next lower level. If each of those contributors could produce the top event alone an OR gate is used (not shown explicitly); while if all the contributors must act to result in the top event an AND gate is used (shown explicitly). Then continue to the next level. We show several variations of attack trees in Figures 2-3.

In fact, attack trees afford various benefits. It is flexible, such that it is possible to combine multiple qualifiers, e.g., time it may take an attacker to successfully carry out a sub-event, or the resources

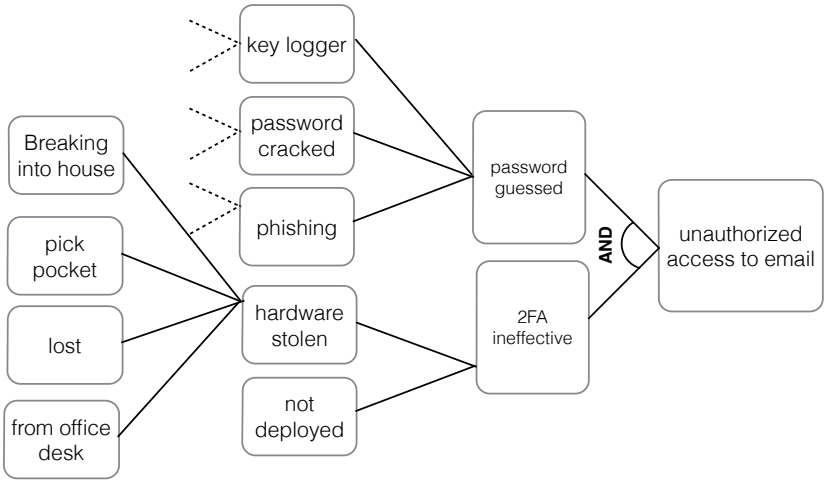

Figure 2: A non-exhaustive attack tree analysis example, showing a scenario (similar to the example for fault tree shown in Figure 1) for unauthorized access to an individual's emails.

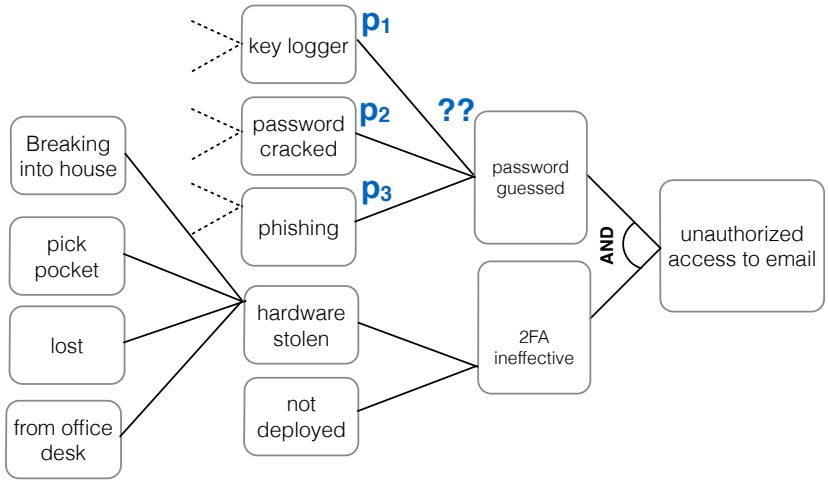

Figure 3: The example from Figure 2 repeated, qualified with additional attributes, in this case probability of the individual events happening (shown partially). Since the operation involved is disjunctive in nature (OR gate), the probability of the resulting event, marked as ?? in the figure is in fact $1 - (1 - p_1)(1 - p_2)(1 - p_3)$. This demonstrates the principle of how the probabilities of individual events can be combined to determine the probabilities of the higher levels, to ultimately determine the probability of the undesired event happening.

required by the attacker to do so, and so on. These can then be used to identify the preferable and likely attack paths (recall the principle of easiest penetration) that the attacker might choose, and, depending on the severity and likelihood of the individual compromises leading to the ultimate compromise of the target, it also allows the analyst to prioritize and deploy layered defense as counter measure. Another benefit of the analysis approach is that it yields lego-block like reusable modules. Thus sub-trees of analysis can be reused wherever applicable, and these sub-trees can be refined with as many details as one wishes or has information about.

Event tree analysis Event trees are somewhat a dual to attack trees, applying a forward, bottom up, logical modeling technique for both success and failure that explores responses through a single initiating event and lays a path for assessing probabilities of the outcomes and for overall system analysis.

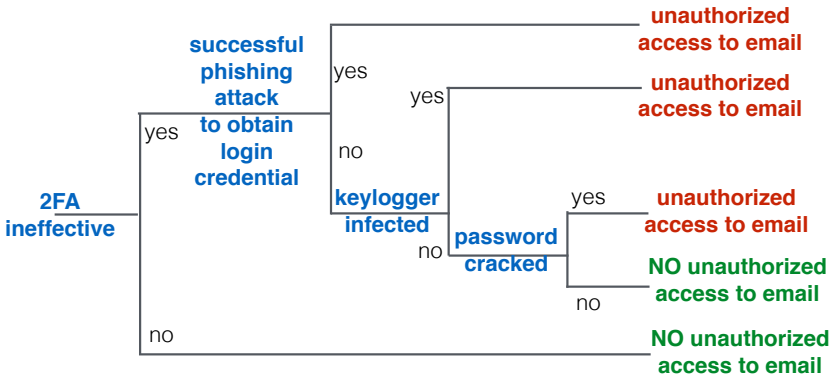

Figure 4: An event tree example for the same running example of unauthorized email access, starting from an initial event which renders two factor authentication ineffective. This tree structure, equipped with the necessary information, can again be leveraged to compute the probabilities of each branch, and thus the chances of arriving at undesirable states.

Failure mode and effects analysis (FMEA)

FMEA originated from the literature of reliability analysis, but it is also used in the context of information security as a risk analysis framework under the ISO 27k family. FMEA is an inductive reasoning (forward logic) single point of failure[3] analysis to review as many components, assemblies, and subsystems as possible to identify failure modes, and their causes and effects.

[3] A single point of failure is a part of a system, which, if it fails, will lead to the failure of the whole system.

The crux of FMEA is an extension of the ideas of severity and exposure[4] we had previously discussed for vulnerability analysis. Specifically, in FMEA each failure mode is attributed a numeric score determined by three factors: (i) likelihood (**prob**ability) that the failure will occur, (ii) likelihood that the failure will not be **det**ected and (iii) the amount of harm or damage the failure mode may cause to a person or to equipment (**sev**erity). The additional aspect we notice here is the likelihood that a failure goes unnoticed, and indeed, lack of detection would perpetuate the problem without triggering any reaction, which is undesirable. A higher score for any of the three component expresses the fact that the problem is worse. [5] The assignment of the numeric scores for these three individual components are somewhat qualitative and subjective in nature. Nevertheless, FMEA provides a reasonably objective manner to determine the magnitude of problem, and rank them, by computing the product of these factors to determine what is called Risk Priority Number ($RPN = prob \times det \times sev$).

[4] See the note on inconsistent terminology at the end of this topic.

[5] A detailed example of how FMEA can be used may be found at `http://www.iso27001security.com/ISO27k_FMEA_spreadsheet.xlsx` (accessed on 21st September 2015).

Note on inconsistent terminology: In previous discussion, we referred as 'exposure' to what we will refer as 'severity' in the context of FMEA, while, we use 'likelihood' now for what we roughly referred as 'severity' in previous discussion. This anomaly of terminologies is arising because the former terminology is adapted from [6], while the later is as used in FMEA. The later is thus a more mainstream usage.

[6] Bayne, 2002

Information security continuous monitoring (ISCM)

Information security continuous monitoring (ISCM) is defined as maintaining ongoing awareness of information security, vulnerabilities, and threats to support organizational risk management decisions, and has been proposed in [7] for US federal organizations and their information systems (statutory under FISMA). The emphasis is on continuous risk assessment, where continuous means frequent for all practical purposes, and the frequency itself determined by the criticality of issues at hand. ISCM creates three tiers for risk assessment at different levels.

[7] Dempsey et al., 2011

Tier 1 risk management activities address high-level information security governance policy as it relates to risk to the organization as a whole, to its core missions, and to its business functions. For instance, assessment at this tier delves into issues such as whether the business processes are aligned to the necessary regulations.

Tier 2 criteria for continuous monitoring of information security are defined by how core mission/business processes are prioritized with respect to the overall goals and objectives of the organization, the types of information needed to successfully execute the stated mission/business processes, and the organization-wide information security program strategy. So to say, at this tier one determines things like whether there are well defined processes on how to do things - to meet the compliance requirements, either for government standards, industry standards, or the business? self-defined needs. e.g., What are the SLAs (service level agreements) for outsourcing?, etc.

ISCM activities at Tier 3 address risk management from an information system perspective. These activities include ensuring that all system-level security controls (technical, operational, and management controls) are implemented correctly, operate as intended, produce the desired outcome with respect to meeting the security requirements for the system, and continue to be effective over time. So it is at this tier that one will ponder if the actual deployed information systems are robust: whether the firewall is rightly configured, if the right security patch updates are installed, and so on.

Concluding remarks

In the absence of perfect security, and under the constraints of limited resources, the problem of security in general, and that of information security, comes down to a matter of identifying risks, and prioritizing which of these risks to eliminate, mitigate or accept. The risk analysis process is encumbered by many factors - lack of complete information, diverse perspective among different stake holders, difficulties in attributing specific value to assets, or determining the quantum of losses, identifying secondary (knock on) effects, etc.

Nevertheless, the exercise of risk analysis is one of the more realistic ways to approach the issue of information security, and in this module we explored a set of general tools that are typically used. How and which specific of these tools to apply under a specific situation is where hands on experience, both with using the tools, and the business (sector) for which the analysis is being carried out, comes handy.

Acknowledgements

Some of the definitions are derived from other sources, most prominently from Wikipedia, and at times quoted rather verbatim given that the definitions were particularly succinct. A lot of the rest of the material of this chapter are derived from multiple sources including the following enumerated articles[8].

[8] Bayne, 2002; Tan, 2002; and Dempsey et al., 2011

Bibliography

Bob Bakely, Ellen McDermott, and Dan Geer. Information security is information risk management. In *Workshop on New security paradigms*, 2001.

James Bayne. An overview of threat and risk assessment. Technical report, SANS Institute, 2002.

Kelley Dempsey, Nirali Shah Chawla, Arnold Johnson, Ronald Johnston, Alicia Clay Jones, Angela Orebaugh, Matthew Scholl, and Kevin Stine. Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations. Technical Report Special Publication 800-137, National Institute of Standards and Technology, 2011.

Ding Tan. Quantitative risk analysis step-by-step. Technical report, SANS Institute, 2002.