University of BRISTOL

CyBOK

Cyber Security Body of Knowledge: Risk Management & Governance

bristol.ac.uk

University of BRISTOL

CyBOK

The CyBOK project would like to understand how the CyBOK is being used and its uptake. The project would like organisations using, or intending to use, CyBOK for the purposes of education, training, course development, professional development etc. to contact it at contact@cybok.org to let the project know how they are using CyBOK.

Introduction

- Fundamental principles of cyber risk assessment and management, and their role in risk governance
- Explain why, as humans, we need effective risk assessment and management principles to capture and communicate factors that may impact our values
- Describe different perspectives on cyber risk management – from individual assets to whole-system goals and objectives
- Study some of he major risk assessment methods and highlight uses and limitations

Introduction

- Discuss security metrics – which features to measure for risk, how to measure risk, and why measure risk at all?
- Explain why effective governance is necessary to uphold cybersecurity, including some social and cultural factors that are essential to consider when developing governance frameworks
- Finally, we discuss incident response and its link to the risk governance process

What is risk?

- Renn’s working definition of risk is *the possibility that human actions or events lead to consequences that have an impact on what humans value*
- Grounded in human value, which applies to many different scenarios
- How to define value and capture indicators to measure and manage risk?
  - Outcomes that have impact on what humans value
  - Possibility of occurrence (uncertainty)
  - Formula to combine both elements
- These elements are at the core of most risk assessment methods

What is risk?

- Key challenge is making assumptions explicit and finding the balance between subjective risk perceptions, and objective evidence
- Risk assessment is therefore the process of collating observations and perceptions of the world that can be justified by logical reasoning or comparisons with actual outcomes
- Risk management is the process of developing and evaluating options to address the risks in a manner agreeable to people whose values may be impacted – with consideration for a spectrum of rejection to acceptance
- Risk governance is the overarching set of ongoing processes and principles that aim to ensure awareness, education, responsibility and accountability to all involved in managing it

CyBOK

Why is risk assessment and management important?

- Risk assessment involves three core components:
  - Identification and, if possible, estimation of hazard (events and strength of outcome)
  - Assessment of exposure (aspects open to threat e.g. people, devices, databases) and/or vulnerability (attributes of aspects that could be targeted e.g. susceptibility to deception, hardware flaws, software exploits)
  - Estimation of risk, combining likelihood and severity (impact of outcomes e.g. quantitative or qualitative)

- Without any of this information we have no basis from which to understand our exposure to threats nor devise a plan to manage them

Why is risk assessment and management important?

- The risk management process involves reviewing the information collected as part of the risk assessment, and leads to one of three possible decisions:
  - *Intolerable:* the aspect of the system at risk needs to be abandoned or replaced, or if not possible, vulnerabilities need to be reduced and exposure limited.
  - *Tolerable:* risks have been reduced with reasonable and appropriate methods to a level as low as reasonably possible (ALARP) or as low as reasonably allowable (ALARA). A range of choices may include mitigating, sharing, or transferring risk, selection of which will depend on the risk managers' (and more general company) appetite for taking risks.
  - *Acceptable:* risk reduction is not necessary and can proceed without intervention. Furthermore, risk can also be used to pursue opportunities (also known as 'upside risk'), thus the outcome may be to accept and embrace the risk rather than reduce it. Hillson discusses this perspective in further detail [4].

Why is risk assessment and management important?

- Deciding which to select will be dependent on a number of factors, for example (as suggested in ISO 31000:2018 [8]), tangible and intangible uncertainty, consequences of risk realisation (good or bad), appetite for risk, organisational capacity to handle risk etc.
- Renn also defines four types of risk that require different risk management plans [3]. These include:
  - *Routine risks:* these follow a fairly normal decision-making process for management. Statistics and relevant data are provided, desirable outcomes and limits of acceptability are defined, and risk reduction measures are implemented and enforced. Renn gives examples of car accidents and safety devices.
  - *Complex risks:* where risks are less clear cut, there may be a need to include a broader set of evidence and consider a comparative approach such as cost-benefit analysis or cost-effectiveness. Scientific dissent such as drug treatment effects or climate change are examples of this.
  - *Uncertain risks:* where a lack of predictability exists, factors such as reversibility, persistence and ubiquity become useful considerations. A precautionary approach should be taken with a continual and managed approach to system development whereby negative side effects can be contained and rolled-back. Resilience to uncertain outcomes is key here.
  - *Ambiguous risks:* where broader stakeholders, such as operational staff or civil society, interpret risk differently (e.g., different viewpoints exist or lack of agreement on management controls), risk management needs to address the causes for the differing views. Renn uses the example of genetically modified foods where well-being concerns conflict with sustainability options. In this instance, risk management must enable participatory decision-making, with discursive measures aiming to reduce the ambiguity to a number of manageable options that can be further assessed and evaluated.

Why is risk assessment and management important?

- Management options, therefore, include:
  - a risk-based management approach (risk-benefit analysis or comparative options)
  - a resilience-based approach (where it is accepted that risk will likely remain but needs to be contained, e.g. using ALARA/ALARP principles)
  - a discourse-based approach (including risk communication and conflict resolution to deal with ambiguities).
- Without effective consideration of the acceptability of risk and an appropriate risk reduction plan, it is likely that the response to adverse outcomes will be disorganised, ineffective, and likely lead to further spreading of undesirable outcomes.

Why is risk assessment and management important?

- Effective risk management through structured assessment methods is particularly important because, although our working definition of risk is grounded in consequences of interest to people, we (as a society) are not very good at assessing this risk.
- Slovic's article on risk perception highlights that perceptions related to *dread risk* (e.g., nuclear accidents) are ranked highest risk by lay people, but much lower by domain experts who understand the evidence relating to safety limitations and controls for such systems.
- Expert risk ranking tends to follow expected or recorded undesirable outcomes such as deaths, while lay people are influenced more by their intuitive judgment (a nuclear accident could impact my whole family).
- There is, therefore, a mismatch between perceived vs. actual risk. As people we tend to exaggerate *dread-related* but rare risks (e.g., nuclear incidents and terrorist attacks) but downplay common ones (e.g., street crime and accidents in the home) – even though the latter kill far more people.

Why is risk assessment and management important?

- This is also why concern assessment is important in the risk management process alongside risk assessment. Schneier's book *Beyond Fear* [5] notes that we have a natural sense of safety in our own environment and a heightened sense of risk outside of this. For instance, we feel safe walking down a street next to our house but on edge when arriving in a new city.
- As a society, we rarely study statistics when making decisions; they are based on perceptions of exposure to threat, our perceived control over threats, and their possible impact.
- Risk assessment helps us capture quantitative and qualitative aspects of the world that enable us to put a realistic estimate of how certain we can be that adverse events will come to pass, and how they will impact on what we value most.
- This applies to us personally as individuals, and as groups of people with a common aim – saving the planet, running a business, or educating children. We need to capture our goals, understand what could lead to the failure to achieve them, and put processes in place to align realistic measures to reduce harms inflicted upon our objectives.

Why is risk assessment and management important?

- When done well, risk assessment and management enables decision makers, who are responsible, to ensure that the system operates to achieve the desired goals as defined by its stakeholders.
- It can also ensure the system is not manipulated (intentionally or otherwise) to produce undesired outcomes, as well as having processes in place that minimise the impact should undesirable outcomes occur.
- Risk assessment and management is also about presenting information in a transparent, understandable and easily interpreted way to different audiences, so that accountable stakeholders are aware of the risks, how they are being managed, who is responsible for managing them, and are in agreement on what is the acceptable limit of risk exposure.
- This is absolutely crucial to successfully managing risk because, if the risks are not presented clearly to decision makers (be they technical, social, economic or otherwise), the impact of not managing them will be overlooked, and the system will remain exposed.

Why is risk assessment and management important?

- Likewise, if the purpose of risk management is not made clear to the people at the operational level, alongside their own responsibilities and accountability for risk impacts, they will not buy in to the risk management plan and the system will remain exposed.
- More broadly, if wider stakeholder concerns (e.g., civil society) are not heard or there is lack of confidence in the risk management plan, there could be widespread rejection of the planned system being proposed.
- As important as it is to convey risks clearly to stakeholders, it is equally as important to stress that risks cannot always be removed. There is likely to be some residual risk to the things we value, so discussions must be held between decision makers and those who are involved with the operations of a system.

Why is risk assessment and management important?

- Ultimately, decision makers, who will be held to account for failure to manage risk, will determine the level of risk tolerance – whether risk is accepted, avoided, mitigated, shared, or transferred.
- However, it is possible that wider stakeholders, such as those involved with system operations, may have differing views on how to manage risk, given they are likely to have different values they are trying to protect.
- For some, saving money will be key. For others, reputation is the main focus. For people working within the system it may be speed of process or ease of carrying out daily tasks.
- The purpose of risk assessment and management is to communicate these values and ensure decisions are taken to minimise the risks to an agreed set of values by managing them appropriately, while maximising ‘buy in’ to the risk management process.

Why is risk assessment and management important?

- One of the major drivers for risk assessment and management is to demonstrate compliance. This can be a result of the need:
  - to have audited compliance approval from international standards bodies in order to gain commercial contracts;
  - to comply with legal or regulatory demands (e.g., in Europe the Network and Information Systems (NIS) directive [9] mandates that operators of essential services (such as critical national infrastructure) follow a set of 14 goal-oriented principles [10]);
  - or to improve the marketability of a company through perceived improvements in public trust if certification is obtained.
- This can sometimes lead to 'tick-box' risk assessment whereby the outcome is less focused on managing the risk, and more about achieving compliance.
- This can result in a false sense of security and leave the organisation exposed to risks. This bring us back to Renn's working definition of risk.
- These examples focus on managing risk of failing compliance with various policy positions, and as a result, they may neglect the broader focus on impact on values held by wider organisational, societal or economic stakeholders.
- The context and scope of risk management must take this broader outcomes-view in order to be a useful and valuable exercise that improves preparedness and resilience to adverse outcomes.

Why is risk assessment and management important?

- Based on these factors, risk assessment and management is most certainly a process not a product.
- It is something that, when done well, has the potential to significantly improve the resilience of a system. When done badly (or not at all) it can lead to confusion, reputational damage, and serious impact on system functionality.
- It is a process that is sometimes perceived to be unimportant before one needs it, but critical for business continuity in a time of crisis.
- Throughout the process of risk assessment we must remain aware that risk perception varies significantly based on a variety of factors, and that despite objective evidence, it will not change.

What is cyber risk assessment and management?

- The introductory sections have made the case for risk assessment and management more generally, but the main focus of this document is to frame risk assessment and management in a cyber security context.
- Digital technology is becoming evermore pervasive and underpins almost every facet of our daily lives. With the growth of the Internet of Things, connected devices are expected to reach levels of more than 50 billion by 2022 [15].
- Further, human decision-based tasks such as driving and decision-making are being replaced by automated technologies, and the digital infrastructures that we are increasingly reliant upon can be disrupted indiscriminately as a result of, for example, ransomware [16].
- Cyber security risk assessment and management is, therefore, a fundamental special case that everyone living and working within the digital domain should understand and be a participant in it.

Risk Governance

- Risk assessment and developing mitigation principles to manage risk is only likely to be effective where a coordinated and well communicated governance policy is put in place within the system being managed. Millstone et al. [19] proposed three governance models:
  - *Technocratic:* where policy is directly informed by science and evidence from domain expertise.
  - *Decisionistic:* where risk evaluation and policy are developed using inputs beyond science alone. For instance, incorporating social and economic drivers.
  - *Transparent (inclusive):* where context for risk assessment is considered from the outset with input from science, politics, economics and civil society. This develops a model of 'pre-assessment' – that includes the views of wider stakeholders – that shapes risk assessment and subsequent management policy.

Risk Governance

- None are correct or incorrect. There is a fine balance between the knowledge and findings of scientific experts, and perceptions of the lay public.
- While the technocratic approach may seem logical to some risk owners who work on the basis of reasoning using evidence, it is absolutely crucial for effective risk governance to include the wider stakeholder view.
- Rohrmann and Renn’s work on risk perception highlights some key reasons for this [20]. They identify four elements that influence the perception of risk:
  - intuitive judgment associated with probabilities and damages;
  - Contextual factors surrounding the perceived characteristics of the risk (e.g.,familiarity) and the risk situation (e.g., personal control);
  - semantic associations linked to the risk source, people associated with the risk, and circumstances of the risk-taking situation;
  - trust and credibility of the actors involved in the risk debate.

Risk Governance

- These factors are not particularly scientific, structured or evidence-based but, as noted by Fischoff et al. [21], such forms of defining probabilities are countered by the strength of belief people have about the likelihood of an undesirable event impacting their own values.
- Ultimately, from a governance perspective, the more inclusive and transparent the policy development, the more likely the support and buy-in from the wider stakeholder group – including lay people as well as operational staff – for the risk management policies and principles.
- A major principle is ensuring that the governance activity is tightly coupled with everyday activity and decision-making. Cyber risk is as important as health and safety, financial processes, and human resources.
- Cyber security should be thought of as a clear set of processes that reduce the risk of harm to individuals and the business. Everyone involved in the daily running of the system in question must understand that, for security to be effective, it must be part of everyday operational culture. The cyber risk governance approach is key to this cultural adoption.

The Human Factor and Risk Communication

- Sasse and Flechais [22] identified human factors that can impact security governance, including people:
  - having problems using security tools correctly;
  - not understanding the importance of data, software, and systems for their organisation;
  - not believing that the assets are at risk (i.e., that they would be attacked);
  - or not understanding that their behaviour puts the system at risk.
- This highlights that *risk cannot be mitigated with technology alone*, and that *concern assessment* is important. If risk perception is such that there is a widely held view that people do not believe their assets will be attacked (as noted by [22]), despite statistics showing cyber security breaches are on the rise year-on-year, then there is likely to be a problem with the cyber security culture in the organisation.

The Human Factor and Risk Communication

- Educating people within an organisation is vital to ensuring cultural adoption of the principles defined in the risk management plan and associated security governance policy.
- People will generally follow the path of least resistance to get a job done, or seek the path of highest reward.
- As Sasse and Flechais note, people fail to follow the required security behaviour for one of two reasons:
  - they are unable to behave as required (one example being that it is not technically possible to do so; another being that the security procedures and policies available to them are large, difficult to digest, or unclear)
  - they do not want to behave in the way required (an example of this may be that they find it easier to work around the proposed low-risk but time consuming policy; another being that they disagree with the proposed policy).

The Human Factor and Risk Communication

- Weirich and Sasse studied compliance with password rules as an example of compliance with security policy [23] and found that a lack of compliance was associated with people not believing that they were personally at risk and or that they would be held accountable for failure to follow security rules.
- There is thus a need to ensure a sense of responsibility and process for accountability, should there be a breach of policy. This must, of course, be mindful of legal and ethical implications, as well as the cultural issues around breaching rules, which is a balancing act.

The Human Factor and Risk Communication

- Risk communication, therefore, plays an important role in governance [24] [1] including aspects, such as:
  - *Education:* particularly around risk awareness and day-to-day handling of risks, including risk and concern assessment and management;
  - *Training and inducement of behaviour change:* taking the awareness provided by education and changing internal practices and processes to adhere to security policy;
  - *Creation of confidence:* both around organisational risk management and key individuals – develop trust over time, and maintain this through strong performance and handling of risks.
  - *Involvement:* particularly in the risk decision-making process – giving stakeholders an opportunity to take part in risk and concern assessment and partake in conflict resolution.

Security culture and awareness

- Dekker’s principles on *Just Culture* [25] aim to balance accountability with learning in the context of security. He proposes the need to change the way in which we think about accountability so that it becomes compatible with learning and improving the security posture of an organisation.
- It is important that people feel able to report issues and concerns, particularly if they think they may be at fault. Accountability needs to be intrinsically linked to *helping the organisation*, without concern of being stigmatised and penalised.
- There is often an issue where those responsible for security governance have limited awareness and understanding of what it means to practise it in the operational world.
- In these cases there needs to be an awareness that there is possibly no clear right or wrong, and that poorly thought-out processes and practices are likely to have been behind the security breach, as opposed to malicious human behaviour.
- If this is the case, these need to be addressed and the person at fault needs to feel supported by their peers and free of anxiety.

Security culture and awareness

- One suggestion Dekker makes is to have an independent team to handle security breach reports so people do not have to go through their line manager.
- If people are aware of the pathways and outcomes following security breaches it will reduce anxiety about what will happen and, therefore, lead to a more open security culture.

Security culture and awareness

- Given that security awareness and education is such an important factor in effective governance, Jaquith [26] links security awareness with security metrics through a range of questions that may be considered as useful pointers for improving security culture:
  - Are employees acknowledging their security responsibilities as users of information systems? (Metric: % new employees completing security awareness training).
  - Are employees receiving training at intervals consistent with company policies?(Metric: % existing employees completing refresher training per policy).
  - Do security staff members possess sufficient skills and professional certifications? (Metric: % security staff with professional security certifications).
  - Are security staff members acquiring new skills at rates consistent with management objectives? (Metrics: # security skill mastered, average per employee and per security team member, fulfilment rate of target external security training workshops and class- room seminars).
  - Are security awareness and training efforts leading to measurable results? (Metrics: By business unit or office, correlation of password strength with the elapsed time since training classes; by business unit or office, correlation of tailgating rate with training latency).

Security culture and awareness

- Metrics may be a crude way to capture adherence to security policy, but when linked to questions that are related to the initial risk assessment, they can provide an objective and measurable way to continually monitor and report on the security of a system to the decision makers, as well as those responsible for its governance in an understandable and meaningful way.
- However, it is worth noting the complexity of metrics here with the use of the term ‘acknowledging’ in the first bullet point. It does not necessarily mean the person will acknowledge their responsibilities merely by completing awareness training. This reinforces the points already made about the importance of human factors and security culture, and the following section on enacting security policy.

Enacting Security Policy

- Overall, effective cyber risk governance will be underpinned by a clear and enactable security policy.
- From the initial phase of the risk assessment there should be a clear focus on the purpose and scope of the risk assessment exercise. During this phase, for more complex systems or whole system security, there should be a focus on identifying the objectives and goals of the system.
- These should be achievable with clear links from objectives to the processes that underpin them. Risks should be articulated as clear statements that capture the interdependencies between the vulnerabilities, threats, likelihoods and outcomes (e.g., causes and effects) that comprise the risk.
- Risk management decisions will be taken to mitigate threats identified for these processes, and these should be linked to the security policy, which will clearly articulate the required actions and activities taken (and by whom), often along with a clear timeline, to mitigate the risks.
- This should also include what is expected to happen as a consequence of this risk becoming a reality.

Enacting Security Policy

- Presentation of risk assessment information in this context is important. Often heat maps and risk matrices are used to visualise the risks.
- However, research has identified limitations in the concept of combining multiple risk measurements (such as likelihood and impact) into a single matrix and heat map [30].
- Attention should, therefore, be paid to the purpose of the visualisation and the accuracy of the evidence it represents for the goal of developing security policy decisions.

Enacting Security Policy

- Human factors (see the Human Factors CyBOK Knowledge Area [27]), and security culture are fundamental to the enactment of the security policy.
- As discussed, people fail to follow the required security behaviour because they are unable to behave as required, or they do not want to behave in the way required [22].
- A set of rules dictating how security risk management should operate will almost certainly fail unless the necessary actions are seen as linked to broader organisational governance, and therefore security policy, in the same way HR and finance policy requires.
- People must be enabled to operate in a secure way and not be the subject of a blame culture when things fail. It is highly likely that there will be security breaches, but the majority of these will not be intentional.
- Therefore, the security policy must be reflective and reactive to issues, responding to the *Just Culture* agenda and creating a policy of accountability for learning, and using mistakes to refine the security policy and underpinning processes – not blame and penalise people.

Enacting Security Policy

- Security education should be a formal part of all employees' continual professional development, with reinforced messaging around why cyber security is important to the organisation, and the employee's role and duties within this.
- Principles of risk communication are an important aspect of the human factor in security education. Frequent communication, tailoring the message to the audience, pretesting the message and considering existing risk perceptions that should be part of the planning around security education.
- Extensive discussion of such risk communication principles that are particularly relevant for messaging regarding risk can be found in [29].

Enacting Security Policy

- Part of the final risk assessment and management outcomes should be a list of accepted risks with associated owners who have oversight for the organisational goals and assets underpinning the processes at risk.
- These individuals should be tightly coupled with review activity and should be clearly identifiable as responsible and accountable for risk management.

Risk Assessment & Management Principles

- The UK NCSC guidance [14] breaks down risk management into:
  - *Component-driven risk management*, which focuses on technical components, and the threats and vulnerabilities they face (also known as bottom up); and
  - *System-driven risk management*, which analyses systems as a whole (also known as top down).
- A major difference between the two is that component-driven approaches tend to focus on the specific risk to an individual component (e.g., hardware, software, data, staff), while system-driven approaches focus more on the goals of an entire system – requiring the definition of a higher level purpose and subsequent understanding of sub-systems and how various parts interact.

Risk Assessment & Management Principles

- Rasmussen’s work [31] enables us to consider a hierarchy of abstraction and show how systems-driven and component-driven risk assessment techniques are complementary.
- Goals and purposes of the system at the higher level. Focus on dependencies between sub-goals and also what the system must not do (pre-defined failure states).
- These are important to design into the system and, if omitted, lead to having to retrofit cyber security into a system that has already been deployed.
- The lower levels then consider capabilities and functionality needed to achieve the overarching goals. At this level component-driven risk assessments of real-world artefacts (e.g., hardware, software, data, staff) consider how these may be impacted by adverse actions or events.

Risk Assessment & Management Principles

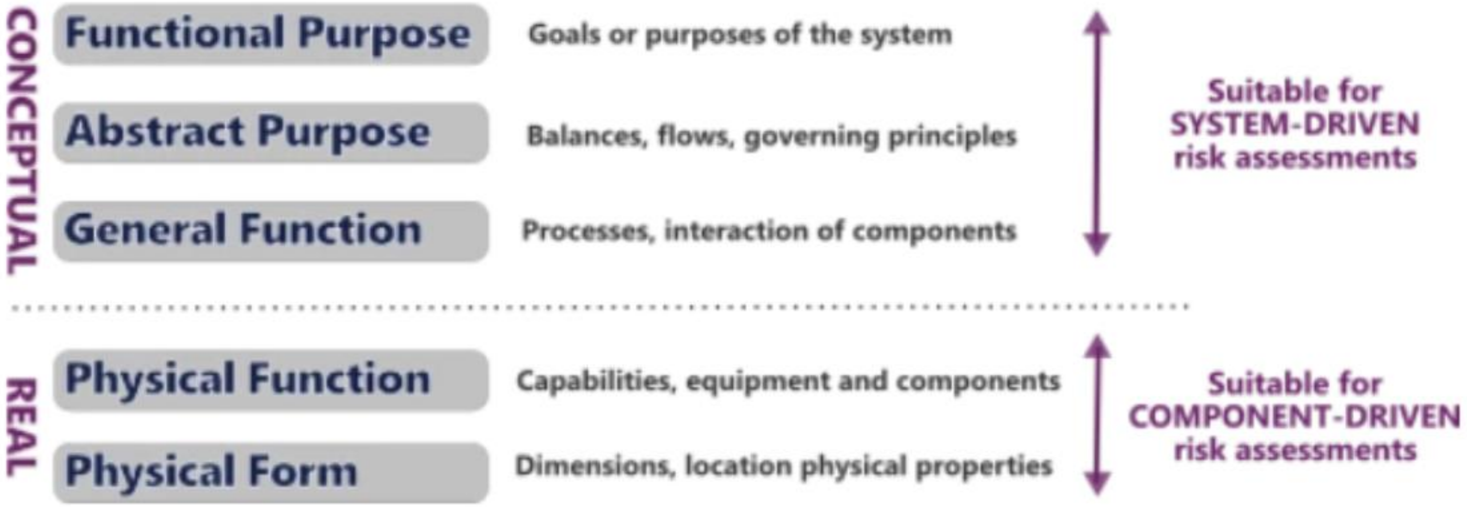
CONCEPTUAL
Functional Purpose
Goals or purposes of the system
Abstract Purpose
Balances, flows, governing principles
General Function
Processes, interaction of components
Suitable for SYSTEM-DRIVEN risk assessments
REAL
Physical Function
Capabilities, equipment and components
Physical Form
Dimensions, location physical properties
Suitable for COMPONENT-DRIVEN risk assessments

Figure 2: Jens Rasmussen's Hierarchy

Risk Assessment & Management Principles

- System-driven approaches can help to better understand the complexity between sub-components and their components.
- These may include people, technology, and organisational processes for which the interactions and dependencies are non-trivial.
- Taking such an approach (which may perhaps prove more resource intensive than component based approaches, due to identification of inter-dependencies) is only necessary where complexity actually exists.
- If interactions and dependencies are clear and the system is less complex (e.g., a simple office-based IT infrastructure) then a component-driven approach may be more appropriate.

Risk Assessment & Management Principles

- These discussions are crucial in finding the balance between component-level and system-level failure and how best to manage the risk.
- Component-risk is likely to be more important to operational employees who need the component to be functioning in order for their part of a bigger system to perform (e.g., staff, data, devices).
- Systems-level risk is likely to be more important to higher-level managers who have a vested interest in the strategic direction of the system. For them a component further down the value/supply chain may not be perceived to be important, while for the operational employee it's the number one risk.
- The challenge is to work with both perspectives to develop a representation of risk and an associated risk management policy enacted by all parties.

Elements of Risk

- There are four concepts that are core to a risk assessment in most models – vulnerability, threat, likelihood and impact.
  - A *Vulnerability* is something open to attack or misuse that could lead to an undesirable outcome. If the vulnerability were to be exploited it could lead to an impact on a process or system. Vulnerabilities can be diverse and include technology (e.g., a software interface being vulnerable to invalid input), people (e.g., a business is vulnerable to a lack of human resources), legal (e.g., databases being vulnerable and linked to large legal fines if data is mishandled and exposed) etc.
  - A *Threat* is an individual, event, or action that has the capability to exploit a vulnerability. Threats are also socio-technical and could include hackers, disgruntled or poorly trained employees, poorly designed software, a poorly articulated or understood operational process etc. To give a concrete example that differentiates vulnerabilities from threats – a software interface has a vulnerability in that malicious input could cause the software to behave in an undesirable manner (e.g., delete tables from a database on the system), while the threat is an action or event that exploits the vulnerability (e.g., the hacker who introduces the malicious input to the system).

Elements of Risk

- *Likelihood* represents a measure capturing the degree of possibility that a threat will exploit a vulnerability, and therefore produce an undesirable outcome affecting the values at the core of the system. This can be a qualitative indicator (e.g., low, medium, high), or a quantitative value (e.g., a scale of 1-10 or a percentage).

- *Impact* is the result of a threat exploiting a vulnerability, which has a negative effect on the success of the objectives for which we are assessing the risk. From a systems view this could be the failure to manufacture a new product on time, while from a component view it could be the failure of a specific manufacturing production component.

Risk Assessment & Management Methods

- The purpose of capturing these four elements of risk is for use in dialogue that aims to represent how best to determine the exposure of a system to cyber risk, and how to manage it.
- The US Government NIST [32] guidelines capture the vulnerability, threats, likelihood and impact elements inside the *prepare (pre-assessment), conduct (appraisal and characterise), communicate (cross-cutting), maintain (management)* cycle. A step-by-step detailed guide can be found in the full document, but we summarise the actions here.

Risk Assessment & Management Methods

- **Prepare** involves identifying the *purpose* (e.g., establishing a baseline of risk or identifying vulnerabilities, threats, likelihood and impact) and *scope* (e.g., what parts of a system/organisation are to be included in the risk assessment?; what decisions are the results informing?).
- It also involves defining *assumptions* and *constraints* on elements such as resources required and predisposing conditions that need to inform the risk assessment. The *assessment approach* and tolerances for risk are also defined at this stage along with identifying *sources of information* relating to threats, vulnerabilities and impact.

Risk Assessment & Management Methods

- **Conduct** is the phase where threats, vulnerabilities, likelihood and impact are identified. There are a range of ways that this can be conducted, and this will vary depending on the nature of the system being risk assessed and the results of the *prepare* stage.
- NIST has a very specific set of tasks to be performed. These may not be relevant to all systems, but there are some useful tasks that generalise across multiple system perspectives, including identifying: threat sources and adversary capability, intent and targets; threat events and relevance to the system in question; vulnerabilities and predisposing conditions; likelihood that the threats identified will exploit the vulnerabilities; and the impacts and affected assets.

Risk Assessment & Management Methods

- **Communicate** is one of the most important phases, and one often overlooked. Conducting the risk assessment gives one the data to be able to inform actions that will improve the security of the system.
- However, it is crucial this is communicated using an appropriate method. Executive boards will expect and need information to be presented in a different way to operational team members, and general organisational staff will need educating and guiding in an entirely different way.
- The results and evidence of the risk assessment must be communicated in a manner accessible to each stakeholder and in a way that is likely to engage them in risk management planning and execution.

Risk Assessment & Management Methods

- **Maintain** is an ongoing phase that is essential to continually update the risk assessment in the light of changes to the system environment and configuration.
- Security postures change regularly in digital environments. For instance, IoT units installed from 2014 to 2020 saw a rapid increase in adoption of 2.63 million across the business sector between 2014 and 2018. By 2020 this is projected to grow by a further 3.39 million.
- This kind of rapid integration of devices into corporate IT systems is likely to change the exposure to risk and, therefore, the scope would need to be refined, new risk assessments carried out, and action taken and communicated to all stakeholders to ensure that the new risk is managed.
- This scenario indicates that (i) risk assessment maintenance must be *proactive* and undertaken much more regularly than an annual basis, and (ii) conducting risk assessment for compliance purposes (possibly only once a year) will leave the organisation wide open to new technological threats unless the *maintain* phase is taken seriously.

Risk Assessment & Management Methods

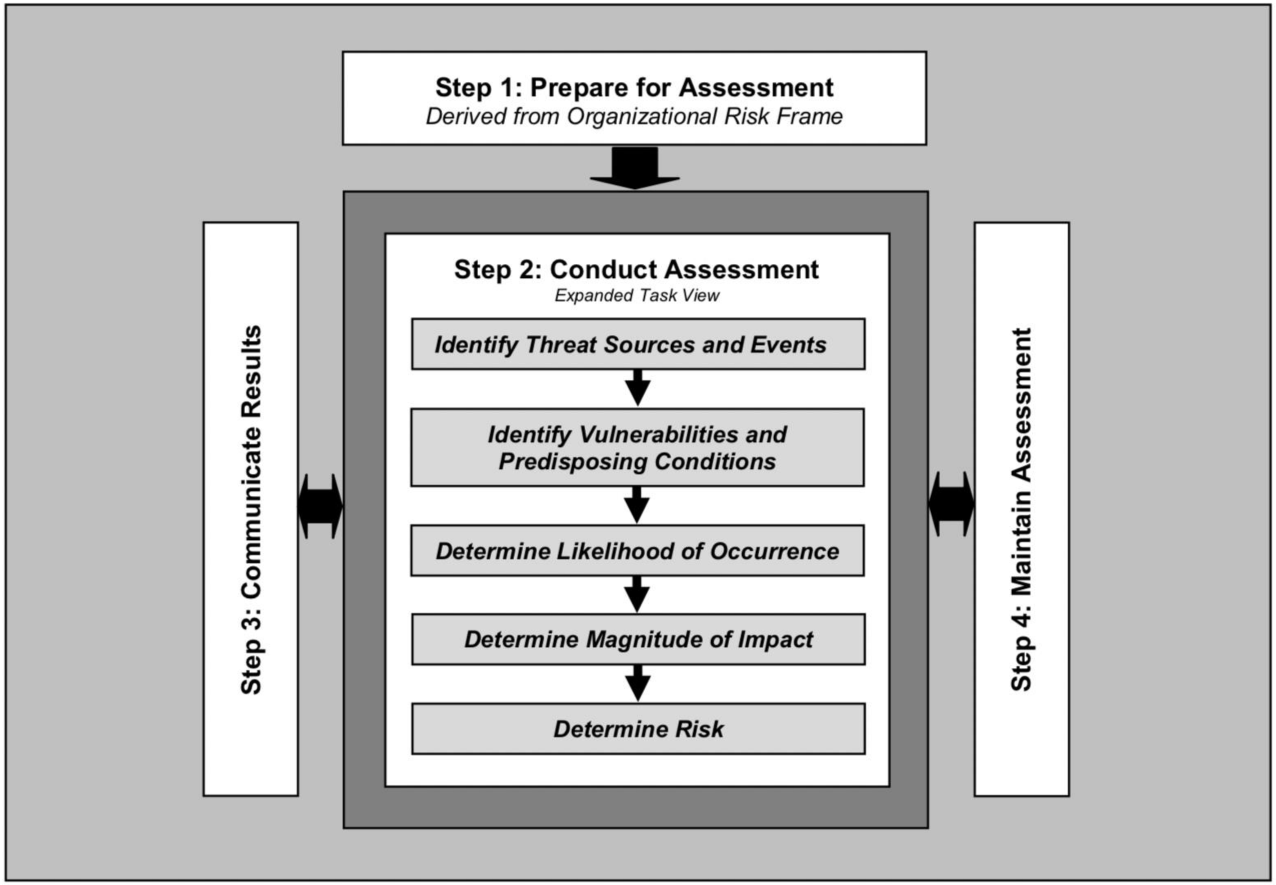
Step 1: Prepare for Assessment
Derived from Organizational Risk Frame
Step 2: Conduct Assessment
Expanded Task View
Identify Threat Sources and Events
Identify Vulnerabilities and Predisposing Conditions
Determine Likelihood of Occurrence
Determine Magnitude of Impact
Determine Risk
Step 3: Communicate Results
Step 4: Maintain Assessment

Risk Assessment & Management Methods

- Reasonably complete list and comparison of component and system level risk assessment methods in the CyBOK Risk Management & Governance KA on pages 18-22

Cyber-physical systems and operational technology

- While traditional IT security (e.g., corporate desktop computers, devices and servers) may generally take a risk assessment perspective focused on minimising access (confidentiality), modification (integrity) and downtime (availability) within components and systems, the world of cyber-physical systems and Operational Technology (OT) typically has a greater focus on *safety*.
- These components and systems, also known as Industrial Control Systems (ICSs) underpin Critical National Infrastructure (CNI) such as energy provision, transportation, and water treatment.
- They also underpin complex manufacturing systems where processes are too heavy-duty, monotonous, or dangerous for human involvement. As a result, OT risks will more often involve a safety or reliability context due to the nature of failure impacting worker and general public safety and livelihood by having a direct impact in the physical world.
- This is perhaps a prime case for the use of systems-driven methods over component-driven, as the former support the abstraction away from components to high-level objectives (e.g., avoiding death, complying with regulation). Taking this view can bridge the security and safety perspective and support discussion on how to best mitigate risk with shared system-level objectives in mind.

Cyber-physical systems and operational technology

- Efforts to continually monitor and control OT remotely have led to increasing convergence of OT with IT, linking the business (and its associated risks) to its safety critical systems.
- Technology such as Supervisory Control and Data Acquisition (SCADA) provides capability to continually monitor and control OT but must be suitably designed to prevent risks from IT impacting OT.
- In Europe the Network and Information Systems (NIS) directive [9] mandates that operators of essential services (such as CNI) follow a set of 14 goal-oriented principles [10], focused on outcomes broadly based around risk assessment, cyber defence, detection and minimising impact.
- Safety critical systems have a history of significant global impacts when failure occurs in the control systems (e.g., Chernobyl, Fukushima), and the addition of connectivity to this environment has the potential to further increase the threat surface, introducing the additional risk elements of global politics and highly-resourced attackers (e.g., Stuxnet, BlackEnergy).
- Recent additions to this debate include the uptake and adoption of IoT devices, including, for example, smart tools on manufacturing shop-floors.
- The cyber security of cyber-physical systems, including vulnerabilities, attacks and countermeasures is beyond the scope of this KA and is discussed in detail in the Cyber-Physical Systems Security CyBOK Knowledge Area [49].

Security Metrics

- Security metrics is a long-standing area of contention within the risk community as there is debate over the value of measuring security.
- It is often difficult to quantify – with confidence – how *secure* an organisation is, or could be. Qualitative representations such as *low, medium, high* or *red, amber, green* are typically used in the absence of trusted quantitative data, but there is often a concern that such values are subjective and mean different things to different stakeholders.
- Open questions include: what features of a system should be measured for risk?, how to measure risk?, and why measure risk at all?
- Some metrics may be related to risk levels, some to system performance, and others related to service provision or reliability. Jaquith provides some useful pointers on what constitutes *good* and *bad* metrics to help select appropriate measures [26].

Security Metrics

- Good metrics should be:
  - Consistently measured, without subjective criteria.
  - Cheap to gather, preferably in an automated way.
  - Expressed as a cardinal number or percentage, not with qualitative labels like ”high”, ”medium”, and ”low”.
  - Expressed using at least one unit of measure, such as ”defects”, ”hours”, or ”dollars”.
  - Contextually specific and relevant enough to decision-makers that they can take action. If the response to a metric is a shrug of the shoulders and ”so what?”, it is not worth gathering. [26]
- Bad metrics:
  - Are inconsistently measured, usually because they rely on subjective judgments that vary from person to person.
  - Cannot be gathered cheaply, as is typical of labour-intensive surveys and one-off spread- sheets.
  - Do not express results with cardinal numbers and units of measure. Instead, they rely on qualitative high/medium/low ratings, traffic lights, and letter grades. [26]

Security Metrics

- The work of Herrmann [36] provides a more pragmatic view based on regulatory compliance, resilience and return on investment.
- The perspective on metrics is grounded in the understanding that we cannot be completely secure, so measuring *actual* security against *necessary* security is arguably a defensible approach, and the metrics described are tailored towards measuring the effectiveness of vulnerability management.
- Essentially, is it possible to quantify whether the risk management plan and associated controls are fit for purpose based on the threats identified, and do the metrics provide evidence that these controls are appropriate?
- Furthermore, are the controls put in place likely to add more value in the savings they produce than the cost of their implementation?

Business Continuity

- Ultimately, despite all best efforts of accountable individuals or boards within a company who have understood and managed the risk they face, it is likely that at some point cyber security defences will be breached.
- An essential part of the risk assessment, management and governance process includes consideration and planning of the process of managing incidents and rapidly responding to cyber attacks.
- The aim is to understand the impact on the system and minimise it, develop and implement a remediation plan, and use this understanding to improve defences to better protect against successful exploitation of vulnerabilities in future (feedback loop).

Business Continuity

- Organisations typically prefer to keep information about cyber security breaches anonymous to prevent reputational damage and cover up lapses in security.
- However, it is likely that other organisations, including competitors will succumb to the same fate in the future, and could benefit from prior knowledge of the incident that occurred.
- At a broad scale, this is something that needs to be addressed, especially given the offensive side of cyber security will communicate and collaborate to share intelligence about opportunities and vulnerabilities for exploiting systems.
- Certain industries such as financial and pharmaceutical sectors have arrangements for sharing such intelligence but it is yet to become commonplace for all types of organisations.
- Large public consortia such as Cyber Defence Alliance Limited (CDA), Cyber Information Sharing Partnership (CISP), and the Open Web Application Security Project (OWASP) are all aiming to support the community in sharing and providing access to intelligence on the latest threats to cyber security.
- For more detailed information on incident management see the Security Operations & Incident Management CyBOK Knowledge Area [53].

Business Continuity

- ISO/IEC 27035-1:2016 [54] is an international standard defining principles for incident management. It expands on the aforementioned ISO/IEC 27005 model and includes steps for incident response, including:
  - *Plan and Prepare:* including the definition of an incident management policy and establishing a team to deal with incidents.
  - *Detection and Reporting:* observing, monitoring detecting and reporting of security incidents
  - *Assessment and Decision:* determining the presence (or otherwise) and associated severity of the incident and taking decisive action on steps to handle it.
  - *Response:* this may include forensic analysis, system patching, or containment and remediation of the incident.
  - *Learning:* a key part of incident management is learning–making improvements to the system defences to reduce the likelihood of future breaches.

Conclusion

- We have explained the fundamental concepts of risk, using a working definition of *the possibility that human actions or events may lead to consequences that have an impact on what humans value*, and placed this in the context of cyber risk management and governance.
- Using academic foundations that have been widely adopted in international practice, we have explained the links between pre-assessment and context setting, risk and concern assessment, characterisation and evaluation, management, and governance
- Risk governance is the overarching set of ongoing processes and principles that underpin collective decision- making and encompasses both risk assessment and management, including consideration of the legal, social, organisational and economic contexts in which risk is evaluated.
- We have defined some of the core terminology used as part of the structured processes that capture information, perceptions and evidence relating to what is at stake, the potential for desirable and undesirable events, and measures of likely outcomes and impact – whether they be qualitative or quantitative.

Conclusion

- A major aspect of risk is human perception and tolerance of risk and we have framed these in the extant literature to argue their significance in risk governance aligned with varying types of risk – routine, complex, uncertain and ambiguous.
- We have particularly drawn on factors that influence the perception of risk and discussed how these link to the human factors of cyber security in the context of security culture.
- Training, behaviour change, creation of confidence and trust, and stakeholder involvement in the risk governance process have been highlighted as crucial success factors.
- This is based on well-established literature that people’s intuition and bias will often outweigh evidence about risk likelihood if they believe the management of the risk is not trustworthy, does not apply to them, or is beyond their control.
- We need people to buy into risk governance rather than impose it upon them.
- Accordingly, we introduced the concept of balancing accountability with learning, proposing that failures in the risk governance process should lead to feedback and improvement where individuals that may have breached risk management policies should feel able to bring this to the attention of risk managers without fear of stigmatisation.

Conclusion

- We differentiated between system-level risk management that analyses the risk of a system as a whole and considers inter-dependencies between sub-systems; and component-level risk management that focuses on risk to individual elements.
- A number of well-established risk management methods from the systems and component perspectives are analysed in the CyBOK KA, with core strengths of each highlighted and some insights into how the methods function.
- While the core principles of risk – based around vulnerability, threat and impact – exist across all methods, there are individual attributes (we refer to as strengths) of each method that may make them a better fit to an organisation depending on what the risk stakeholders require as evidence of exposure.

Conclusion

- Measuring security and the limitations of metrics were discussed in the context of possible options for security metrics, as well as differing views in the community on the benefits and limitations of metricised risk.
- Finally, we linked to incident response and recovery, which should provide a feedback loop to risk management planning within the risk governance process.
- Risk governance is a cyclical and iterative process, and not something that can be performed once.
- The crosscutting aspects of communication, stakeholder engagement and context bind the risk assessment and management processes and are core to the continual reflection and review of risk governance practices.
- Incidents, when they occur, must inform risk management policy to improve cyber security in future – and we must accept that we will likely never be completely secure.
- In line with this, human factors and security culture must respond to the ever changing need to manage cyber risk, enabling and instilling continual professional development through education and *Just Culture* where lessons can be learned and governance methods improved.