CAESAR8: an Agile Enterprise Architecture Approach to Managing Information Security Risks

Paul Loft[a,*], Ying He[b], Iryna Yevseyeva[a] and Isabel Wagner[a]

[a] De Montfort University, The Gateway, Leicester, LE1 9BH, UK
[b] University of Nottingham, University Park, Nottingham, NG7 2RD, UK

ARTICLE INFO

Keywords:
Information Security
Operational Security
Enterprise Architecture
Risk Management
Cyber Security Risk Management
Cyber Security Risk Assessment
Stakeholders
Agile
Lean
Change Management

ABSTRACT

In theory, implementing an Enterprise Architecture (EA) should enable organizations to increase the accuracy of information security risk assessments. In reality, however, organizations struggle to fully implement EA frameworks because the requirements for implementing an EA and the benefits of commercial frameworks are unclear, and the overhead of maintaining EA artifacts is unacceptable, especially for smaller organizations. In this paper, we describe a novel approach called CAESAR8 (Continuous Agile Enterprise Security Architecture Review in 8 domains) that supports dynamic and holistic reviews of information security risks in IT projects. CAESAR8's nonlinear design supports continuous reassessment of information security risks, based on a checklist that assesses the maturity of security considerations in eight domains that often cause information security failures. CAESAR8 assessments can be completed by multiple stakeholders independently, thus ensuring consideration of their tacit knowledge while preventing groupthink. Our evaluation with experienced industry professionals showed that CAESAR8 successfully addresses real-world problems in information security risk management, with significant benefits particularly for smaller organizations.

1. Introduction

When designing new information systems, an Enterprise Architecture (EA) approach can help organizations understand the true effects that a change will have on their business strategy and its operations (Spears, 2004; Li and Hongyan, 2010; Andrews, Monk and Johnston, 2014). EA artifacts document how the business operates and describe how business assets and processes depend on information technology. Likewise, adopting the practice of EA in the design and implementation of security strategies can help companies manage complex business processes and support business strategies (Goudalo and Seret, 2009; Wang, Xu, Lu and Shen, 2009). Information security benefits from an architectural approach, because it integrates security in all aspects of the design of information systems (Loft, He, Janicke and Wagner, 2019). Adopting an Enterprise Information Security Architecture (EISA) approach for Information Security Risk Management (ISRM) therefore provides a security strategy that is focused on business requirements (Li and Hongyan, 2010; Andrews et al., 2014).

In reality, however, there are barriers to achieving the benefits of EA in general (Löhe and Legner, 2014) and EISAs in particular. For example, EAs are often implemented to support technical projects (Bahmani, Shariati and Shams, 2010; Bischoff, Aier and Winter, 2014) to deliver a technical architecture. However, because technology is only one component in the overall company strategy (Sherwood, Clark and Lynas, 2005), this approach can fail to address the wider business context and result in inaccurate impact assessments. In addition, EAs struggle to capture the tacit knowledge within a business (Kotusev and Kurnia, 2020). When this tacit knowledge is not present in EA artifacts, the EA documentation cannot provide the explicit knowledge required to make judgments on enterprise information security risks.

Commercial EA and EISA frameworks include Zachman (Zachman, 1987), TOGAF (The Open Group, 2011) and SABSA (Sherwood et al., 2005). However, following these frameworks requires substantial resource and commitment (Kaisler, Armour and Valivullah, 2005) because they prescribe specific documentation to be created and step-wise methodologies to be followed for integrating EA activities with the business (Kotusev, Singh and Storey, 2015). As a result, most real-world EA implementations do not resemble the theoretical EA frameworks (Kotusev and Kurnia, 2020).

In addition, our previous analysis of security failures (Loft et al., 2019) showed that the structure of commercial EA frameworks is not a good fit for addressing the root causes of security failures. For example, SABSA uses a matrix that is arranged under the interrogative clauses of What, Why, How, Who, Where and When. However, we found very few security failures associated with *Where*, but very complex and varied failures associated with *Who*, e.g., pertaining to end users, supervisory matters and governance issues. Therefore, the SABSA structure did not fit with our findings of holistic security risk assessments. This is supported by calls in the literature for a new model that is both agile and holistic and helps smaller enterprises build out an architecture in a repeatable process (Ross, Weill and Robertson, 2006).

Modern businesses need to change and adapt quickly to remain competitive, and their information security strategies should adapt to ensure continual alignment (Soomro, Shah and Ahmed, 2016). However, there is little research on

✉ P15250056@my365.dmu.ac.uk (P. Loft); Ying.He@nottingham.ac.uk (Y. He); iryna@dmu.ac.uk (I. Yevseyeva); iw@ieee.org (I. Wagner)
ORCID(s): 0000-0002-7449-6585 (P. Loft); 0000-0003-2023-5547 (Y. He); 0000-0002-1627-7624 (I. Yevseyeva); 0000-0003-0242-6278 (I. Wagner)

how EAs can be aligned to changes in business (Kaisler and Armour, 2017). Korhonen, Lapalme, McDavid and Gill (2016) called for a "radical re-conceptualization to inform a more adaptive EA practice", suggesting that EA needs to become a shared competency to support a continuous evolution of the EA with the business environment.

In addition, there is a gap between complex commercial EA frameworks and the agile needs and resource constraints of small-to-medium size enterprises (SMEs) (Bernaert, Poels, Snoeck and Backer, 2014). The original aim of our research, therefore, was to identify a way to help SMEs implement existing commercial EA/EISA frameworks in agile environments.

To find existing solutions, we conducted a systematic literature review to identify the root causes of security failures as well as similar artifacts (Loft et al., 2019). While we found that risk management and enterprise architecture are beneficial for security, we did not find specific frameworks, models or tools that provide organizations with practical agile EA solutions. In addition, we found that commercial EA frameworks struggle to track changes to cyber security risks because existing EA models focus on describing how the enterprise architecture is constructed in layers or phases. Finally, we did not find evidence that commercial frameworks deliver the expected benefits, or that they can be implemented consistently (Kotusev, 2019).

Contributions. In this paper, we aim to move EA from the theoretical realm into a practical solution, especially for smaller organizations. We present a new model called CAESAR8 (Continuous Agile Enterprise Security Architecture Review in 8 domains), which supports ISRM in projects in an agile and holistic way. To ensure the effectiveness of CAESAR8, we follow five novel design principles:

1. **A practical, holistic design.** CAESAR8's nonlinear design encourages continual reassessments of information security risks, which makes it easy to align with agile processes. Progression through five *levels* indicates the maturity of security considerations in each of the eight *domains* (such as assets, business processes, or human factors).

2. **Multiple stakeholder perspectives.** Different stakeholders use CAESAR8 to complete independent assessments of projects from their own perspectives, using their own knowledge of the impact of a project. After completion, all assessments are consolidated into a single result that represents the collective judgment of all stakeholders.

3. **A tractable checklist.** CAESAR8 uses a structured checklist that can be completed in minutes (as opposed to days or weeks with traditional EA approaches), even with limited expertise in information security.

4. **EA process rather than EA artifacts.** CAESAR8 ensures that architecture is considered as part of the assessment process, instead of focusing on creating extensive architecture documentation.

5. **Visualization.** The results are summarized in a radial format. Colors (red, amber, green) indicate the status of each domain and each level, clearly showing where and when intervention of business executives is needed.

Our evaluation with industry experts showed that CAESAR8 succeeds in addressing common problems at the intersection of information security and business. In particular, the ability to include different stakeholders as well as the speed and ease of the assessment process were considered valuable. Thus, the evaluation confirmed that CAESAR8 provides a pragmatic solution that organizations can use to improve how they manage information security risks.

The remainder of this paper is organized as follows: we present related work in Section 2, followed by the design of CAESAR8 and its design principles in Section 3. We show how CAESAR8 can be applied in practice using a case study in Section 4, and present the results of our external evaluation of CAESAR8 in Section 5. Finally, we conclude in Section 6.

2. Related Work

2.1. Enterprise Architecture

An Enterprise Architecture (EA) provides theoretical benefits for Information Security Risk Management (ISRM) because it can increase the accuracy of information security risk assessments. In practice, however, many organizations have been unable to fully implement commercial EA frameworks such as TOGAF, SABSA, or Zachman because of the overheads of creating and maintaining EA artifacts, especially for organizations following agile programs or having limited resource.

In addition, EA frameworks may not deliver their theoretical benefits (Kotusev, 2019), and many commercial frameworks have been shown to have little or no proven benefit (Kotusev et al., 2015; Löhe and Legner, 2014; McClintock, Falkner, Szabo and Yarom, 2020).

2.2. Agile

Our research is directed at finding an agile approach to benefiting from EA-style approaches to ISRM. Agile is a method of project management originally used in software development (Beck, Beedle, Van Bennekum, Cockburn, Cunningham, Fowler, Grenning, Highsmith, Hunt, Jeffries et al., 2001). It involves dividing tasks into short phases of work called *iterations*. Each iteration is reviewed with the business stakeholders and changes are agreed for the next iteration (bottom-up). This method contrasts with more traditional waterfall developments, where management fully agrees the design in advance of development (top-down). Compared to a top-down approach, agile is better suited to meet the rapid demands of the business. However, EA is often seen as a top-down process, making agile and EA seemingly incompatible (Chivers, Paige and Ge, 2005).

Information management and information security can be integrated into agile software development review cycles, for example by capturing security actions for agile user stories on agile dashboards (Madison, 2010; Dorca, Munteanu, Popescu, Chioreanu and Peleskei, 2016). However, this integration alone does not achieve faster review cycles (Kaisler and Armour, 2017) or a move away from the heavy-weight assurance processes required by traditional security standards (Beznosov and Kruchten, 2004).

In agile projects, the *Product Owner* is a member of the agile team who is responsible for translating business requirements into project tasks. If the Product Owner recognizes that good (security) architecture practice contributes significantly to the business value of the project, agile project can make incremental steps in a good architectural direction. However, even though agile principles can be applied even to front-loaded, feature-heavy architectural designs (Madison, 2010), an agile approach for enterprise information security architecture does not exist yet.

2.3. Root Causes of ISRM Failures

The root causes of information security failures can be grouped into eight domains: Human Factors (HF); Management Influence (MI); Enterprise Architecture (EA); External Factors (EF); Business Process (BP); Information Assets (IA); Security Governance (SG); and Technology Infrastructure (TI) (Loft et al., 2019). Each of these domains groups together key factors (*nodes*) that are related to ISRM and are often the root causes of information security incidents, such as internal threats, information sharing, and the social environment. These nodes have differing levels of influence on information security successes and failures, where *influence* can be estimated based on how strongly a node is correlated with success or failure and how many other nodes it has correlations with (Loft et al., 2019).

Nodes also affect each other, forming *chains of influence*, which can be estimated based on pairwise correlations between nodes. Figure 1 illustrates this concept. For example, on the top left of the figure, the green arrows show how an *Internal Threat* is closely associated with the security *Controls Operation*. However, the *Controls Operation* is itself correlated with *Human Factors* and there is a strong correlation (r=0.73) between the *Internal Threat* and *Human Factors* nodes. This correlation supports the observations that human behavior often creates an intentional or unintentional *Internal Threat*. But this can be mitigated. For example, the *Human Factor* is affected by a strong security *Culture*, which requires *Monitoring* to ensure that the culture is being properly *Managed*. In addition, the *Social Environment* influences the security *Culture*, so close *Supervision* within peer networks can ensure that the correct training and work attitude are being applied.

However, analysis of the correlations also uncovers important factors that can lead to security failures. For example, following the red arrows on the right of Figure 1, if the correct selection of security controls is not made, this can lead to *Performance Degradation*, such as making user authentication a difficult process for users. This in turn creates an *Internal Threat* as users try to circumvent the drop in their productivity that this causes. In addition, when *Information Sharing* is conducted within a *Third Party Relationship*, separate data pools are often created to enable the operation of new third party systems or processes. This can lead to *Fragmentation* and the removal of data from the organization's own systems, thereby losing control of company data and increasing its attack surface as a result.

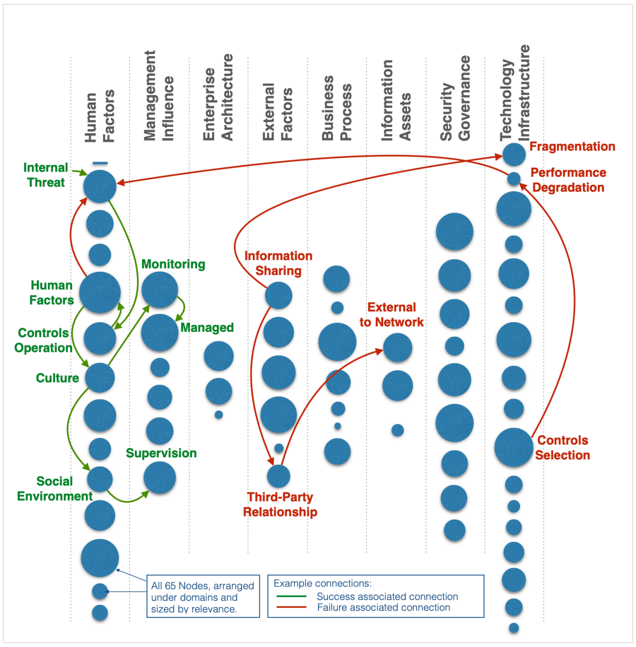

Figure 1: Influence of individual nodes on security successes and failures (denoted by the sizes of the circles), and examples of chains of influence determined by pairwise correlation analysis.

Further analysis of these chains of influence informed our design of CAESAR8, in particular the checklist and structure of the model (presented in Section 3).

2.4. Cognitive Diversity

A diverse group of knowledgeable people are likely to make better risk decisions than individual experts (Page, 2019). *Cognitive diversity* refers to the differences in people's knowledge and experience, including information, knowledge, heuristics, representation, and mental models. While security experts play a vital role in defining security strategies and reviewing changes, relying on a single expert is dangerous because one individual may not have complete awareness of all relevant corporate issues. Business stakeholders, on the other hand, are likely to have a practical understanding of how their respective parts of the business are affected by a change, but may not be able to judge all security aspects correctly. For this reason, "diversity trumps ability" (Hong and Page, 2004), as solutions to complex problems often require multiple points of view. More diverse perspectives deliver a better overall solution, as long as each individual brings relevant knowledge and perspectives.

Individuals, or homologous groups, can show an under-reaction or over-reaction to new information when considering threats (Atanasov, 2020). Therefore, organizations that rely on a single expert's judgment may accept a level of subjectivity in decisions as a result of untreated biases (Skjong and Wentworth, 2001). Experts are also more prone than lay people to miss alternatives and to treat their preferred models as infallible (Skjong and Wentworth, 2001). Also, when posed with a question, experts have been shown to substitute similar alternative questions for which they already have a familiar response (Kahneman, 2011).

An expert's intuitive judgment is therefore prone to errors and biases, meaning that their judgments may be made with a misplaced level of confidence (Kahneman and Tversky, 1977). Moreover, the quality of an expert's decision making rarely improves with experience, partly due to limited opportunities for feedback (Kirkebøen, 2009). Further, experts are not always consistent in their judgments. This can be because of venality, where experts take positions that serve their immediate self-interests (Mumpower and Stewart, 1996), but also because of other biases such as *cognitive dissonance*, where an expert's overarching belief may overpower their concerns with the specific issues under review. As a result, even the same expert can make a different judgment on a different day.

Therefore, the opinions of a knowledgeable, diverse group of experts who offer different knowledge, perspectives and heuristics (Hong and Page, 2004) are important for information security risk assessments. Ideally, these stakeholders are people who share a common goal in the project but who also offer different perspectives or skill sets and are likely to employ a different mode of thinking. However, existing approaches to enterprise architecture struggle with obtaining and maintaining stakeholder involvement (Kurnia, Kotusev, Shanks, Dilnutt and Milton, 2021).

2.5. Groupthink

The term *groupthink* refers to the mode of thinking that persons engage in when concurrence-seeking becomes so dominant in a cohesive ingroup that it overrides realistic appraisal of alternative courses of action (Janis, 1971). In group settings, the sense of conformity (this is a *request*, and is different from obedience, which is an *order*) can be so strong and powerful for an individual that it subconsciously overrides their own thoughts to the contrary. This overriding sense of conformity has led to disasters across industries. For example, in the aviation industry, co-pilots have chosen to risk death, rather than contradict their captains (Milanovich, Driskell, Stout and Salas, 1998). In medicine, junior health care staff have failed to speak up about concerns with their supervisors' patient management plans (Sur, Schindler, Singh, Angelos and Langerman, 2016).

Working in groups can also lead individuals to increase their risk taking. This *Illusion of Invulnerability* is a consequence of groups making judgments together (Janis, 1971; Hart, 1991). Therefore, it is vital that security risk assessments are conducted independently, so that stakeholders can inform the overall decision without undue influence.

The process of stakeholder involvement in CAESAR8 has been designed to avoid groupthink scenarios by making sure that stakeholders retain their independence of judgment.

2.6. Security Standards and Checklists

Many organizations are familiar with tracking compliance using maturity models and security standards. Compliance strategies are an important requirement for many organizations, however, frequent change can result in security standards which have become disconnected from the true risks of the business (Sadki and El Bakkali, 2014; Sen and Borle, 2015). As a result, organizations need to find improved ways of keeping pace with these changes.

Security maturity models are designed to assess practices against standard criteria. However, they do not provide an holistic perspective (Khoshgoftar and Osman, 2009) and often miss the impact of change on human resources. They are also inflexible to rapid change and can be overwhelming in their application (Jugdev and Thomas, 2002), which makes them unsuitable for smaller, agile projects.

Information Security Standards (such as ISO 27001) are valuable references, however, they are very generic and abstract in nature (Diesch, Pfaff and Krcmar, 2020). Smaller organizations rarely adopt the complete standard because they can be too complex to manage. In addition, these standards may not be validated by science and research (Siponen and Willison, 2009).

In general, a tractable set of questions can help individuals check the progress of projects in their known environment, as it provides a valuable tool that all stakeholders can use (Scriven, 2000). Simple, "modest" checklists are best to ensure that concerns highlighted by stakeholders are properly considered and knowledge-loss is avoided (Gawande, 2011). The incorporation of checklists to improve human performance has been proven in many industries, for example, the World Health Organization's surgical checklist improved mortality rates (Weiser and Haynes, 2018).

CAESAR8 uses a checklist to provide a repeatable assessment of the most important issues that businesses need to consider when assessing information security risks in projects, thus allowing the use of the checklist in real-time scenarios.

2.7. Metrics and Visualization

A common rating system in project management is the traffic light system, which uses red, amber, and green colors (RAG) to indicate the status of work: on track (green), started but not complete (amber), and at risk (red). This type of metric, also called *existence metric*, is also useful to provide high-level indications of the security posture of an entity (Axelrod, 2008).

Even though existence metrics enable an easy-to-interpret visualization of potentially complex risk scenarios, especially for senior executives, they can lack granular detail, such as quality and completeness (Axelrod, 2008). However, summarizing project risks in this way, based on accurate assessments and reporting over time, can present management with a valuable prediction about where projects may be running into trouble and when intervention may be required (Hopmere, Crawford and Harré, 2020).

In addition, the use of high-level summary metrics does not preclude more detailed assessments by stakeholders, i.e., the RAG summary simply represents the top of a hierarchy of metrics (Savola and Heinonen, 2011). For example, if a security expert knows that the firewall rules need to be reviewed or updated following a change, they may decide that their assessment for the technical security impact or security strategy is no higher than "amber", thus determining the current status of the overall result.

CAESAR8 uses this high-level RAG summary to visualize the complete security posture of a project, while allowing experts to drill down to their own specific stakeholder metrics.

3. Design of CAESAR8

CAESAR8 combines different areas of human behavioral-science to solve a practical problem of enterprise information security risk management for small to medium sized organizations. In this section, we describe our research methodology, the design principles for the design of CAESAR8, and the final model.

3.1. Methodology

We designed CAESAR8 following the Design Science Research (DSR) methodology (Hevner, March, Park and Ram, 2004), and specifically the DSR process (DSRP) (Peffers, Tuunanen, Rothenberger and Chatterjee, 2007), as this research paradigm is ideally suited to the iterative development and validation of a practical model that focuses on the understanding of how human performance can be improved in a particular context.

DSR aims at the creation and evaluation of new artifacts, such as algorithms or models. The core of DSR is the *design cycle* which is repeated iteratively. Each design cycle consists of a design phase in which the design of the artifact is improved, and an evaluation phase in which the performance of the artifact is assessed.

We used five design and development iterations, incorporating a new design principle in each iteration. Early iterations were demonstrated using ex ante demonstrations (such as the case study in Section 4), and the final model after the fifth iteration was validated using an extensive ex post external evaluation (see the results in Section 5).

3.2. The CAESAR8 Design Principles

We followed five design principles to ensure that CAESAR8 provides a holistic but agile solution for the continuous assessment of information security risks during projects. We formulated these design principles specifically to address fifteen common problem areas for ISRM in small and medium-sized organizations, consisting of governance problems and solution design problems (Table 1), as well as nine underlying issues which can hinder information security projects (Table 2) (Loft et al., 2019).

Table 1
Common problem areas for information security risk management

Governance	1	Stakeholders not directly engaging with projects
	2	Lack of collaboration across separate teams
	3	Limited understanding of the wider effects of changes
	4	Executive not formally understanding project risks
Solution design	5	Legal compliance reviews not completed for all changes
	6	Security risk management not expressed in a business context
	7	Insufficient rigor applied when working with third parties
	8	Agreed security controls are sometimes omitted
	9	Lack of monitoring of security controls
	10	Project impact on current business processes not fully considered
	11	Lack of clarity over information storage and sharing
	12	Ad hoc deployment of new technology
	13	Not understanding the effect of a new system on all personnel
	14	Testing is not completed adequately
	15	Management unwilling or unable to monitor compliance

Table 2
Underlying issues in information security projects

1	Time-related pressures are a risk to security
2	Budget constraints are a risk to security
3	High workloads are a risk to security
4	Volume of project changes are a risk to security
5	Difficult to recruit skilled security personnel
6	Prioritization of work can be unclear
7	Disparate security and business risk management methods
8	Security documentation sometimes inadequate
9	Lack of adherence to security operating procedures

3.2.1. Principle 1: A practical, holistic design.

A practical, holistic design can help address the common problem areas and underlying issues identified in Tables 1 and 2 by supporting a systematic review of all aspects of a project (holistic) in an agile way (practical), as the project progresses. In this way, changes in key information security risk factors can be captured and, if needed, remedied. The model only encourages the timely production of EA artifacts that benefit the immediate project under review, and does not become sidetracked by compliance with the detached characteristics of international security standards. To realize this principle, we derived twelve design goals:

1. Support continual reassessments of ongoing changes to projects using a non-linear design.

2. Reflects the dependency between security activities via progression through the model.
3. Support integration with existing project processes including DevOps.
4. Support integration with agile working practices.
5. Support the creation of architecture documentation, where required.
6. Focus on the key issues that determine the success of information security in projects.
7. Ensure a clear and easy checklist that all stakeholders can understand.
8. Conducting assessments is easy for all stakeholders.
9. Conducting assessments is a quick process.
10. Assist with the prioritization of work.
11. Combine assessments from different stakeholders to ensure that information security solutions are aligned with the business.
12. Support easy sharing of the overall results of assessments with colleagues and management.

3.2.2. Principle 2: Multiple stakeholder perspectives.

Gathering the perspectives of multiple stakeholders is important to access any applicable tacit knowledge they may have (Page, 2019). *Tacit* knowledge is knowledge which has not yet been articulated or cannot be articulated (Hedesstrom and Whitley, 2000). This is the opposite to *explicit* knowledge which can be articulated, and in the context of EA, is usually articulated in EA artifacts which provide written descriptions of an organization from different perspectives. However, tacit knowledge is information that stakeholders subconsciously *know* from their specific awareness and/or experience. In the context of EA, tacit knowledge could relate to an awareness of how an otherwise unconnected part of an organization's operations might be indirectly affected by a change. Although the interactions of these different parts of the business may not be written down, a stakeholder is likely to instinctively *know*. These matters are hard to capture in EA artifacts. However, it is important to ensure that stakeholders can offer their knowledge independently, in a way that is free from the dangers of *groupthink* (Janis, 1971).

To realize this principle, CAESAR8 requires that all relevant stakeholders for a project are identified, i.e., individuals or groups with full, hands-on knowledge of how the business operates in their particular area. Each stakeholder completes a separate CAESAR8 assessment from their perspective, without guessing or generalizing the assessment for the entire organization. Assessments can be completed remotely and individually to avoid groupthink.

3.2.3. Principle 3: Unify around a tractable checklist.

A checklist provides a valuable tool for project stakeholders because it guides the stakeholder's assessment, ensures that all important areas are considered, and helps to make assessments repeatable (Scriven, 2000). Regardless of a stakeholder's security knowledge, a checklist allows the stakeholder to participate in the abstraction of information security risks that can only be identified from their tacit knowledge of their own business area.

To realize this principle, CAESAR8 uses an ordered checklist that examines a common set of enterprise problems that are at the root cause of security failures, and allows affected business stakeholders to repeatedly check that these problems are being avoided.

3.2.4. Principle 4: Value process over EA artifacts.

A holistic *process* for ISRM is more important than creating EA artifacts. While documentation is important for supporting collaboration on the target architecture or documenting the current state, EA artifacts in general can be difficult to create, difficult to use and difficult to maintain (Kaisler et al., 2005), which causes delays and expense.

To support process over artifacts, CAESAR8 embraces the Agile principles and values (Beck et al., 2001). Specifically, CAESAR8 values individuals and interactions over processes and tools, working software over comprehensive documentation, customer collaboration over contract negotiation, and responding to change over following a plan. These values represent a move to embracing change and having less dependence on the creation of documentation. CAESAR8 assessments are designed to be used on a continuous cycle through all maturity levels, so that the correct levels of maturity can be reestablished at every product iteration.

3.2.5. Principle 5: Provide a collective visualization.

The results of ISRM need to be shared with all those affected by a project and in a format that supports senior management engagement and intervention (Belkadi, Cherti and Bahaj, 2018; Hopmere et al., 2020). In addition, in an agile solution, the process needs to be shared across teams and communicated to corporate decision makers. This is important for achieving a *collective visualization*, which improves human productivity (Belkadi et al., 2018).

To realize this principle, CAESAR8 uses a radial visualization of the checklist, both for individual assessments and the combined overall assessment. Color-coding helps to highlight in which areas information security risks need more consideration. As a result, this visualization allows everyone to see how security risks are emerging and being managed. In particular, the visualization can be shared with other stakeholders, serve as a basis for discussions, and be used by senior management as part of their decision making process.

3.3. Final CAESAR8 Model

3.3.1. Checklist

Based on our work on the root causes of information security failures, including the pairwise analysis to discover *chains of influence*, we have identified a common set of 40 performance markers that can guide an holistic information security risk management process in a way that helps to avoid common causes of information security failures.

These performance markers, shown in Figure 2, are arranged in a matrix of eight domains (rows) and five maturity levels (columns). The eight domains are broad areas which can cause information security failures if they are not considered during a project: Human Factors (HF), Management Influence (MI), Enterprise Architecture (EA), External Factors (EF), Business Process (BP), Information Assets (IA), Security Governance (SG), and Technology Infrastructure (TI). The five levels correspond to increasingly detailed considerations of security within a project: Level 1 reviews aspects of the current business that are impacted by a project; Level 2 reviews the nature of the planned changes; Level 3 considers the information security risks in relation to the change; Level 4 analyzes the security approach to mitigating the risks of the change; and level 5 looks at actions that improve the organization's future resilience to information threats.

The order for reviewing these performance markers is critical to this process. For example, the current business processes that are affected by a change (domain BP, level 1) and the proposed changes to this information processing (domain BP, level 2) must be fully understood so that a reliable understanding of the dependence on a third party is achieved (domain EF, levels 2/3) and, subsequently, which changes to security controls need to be agreed with that third party (domain EF, level 4). An assessment may be unreliable if it determines that a change to a third party contact is or is not required without fully understanding changes to all information processing.

3.3.2. Visualization

The results of a stakeholder's assessment, as well as the combined assessment for all stakeholders, are presented in radial format as shown in Figure 3. Each track sector represents one performance marker. Five tracks correspond to the five maturity levels, starting at level 1 on the outside, and each circular sector represents one of the eight domains. Valid responses for each performance marker form an extended RAG status. In addition to the RAG responses of *Yes* (green), *No* (red), and *Partial* (amber), CAESAR8 allows additional responses to avoid distorting a stakeholder's assessment: *Trust* (blue), *N/A* (black), and *Unknown (?)* (grey). Using the *trust* response, a stakeholder can indicate that they trust another stakeholder to address this performance marker. The *N/A* response indicates that this performance marker does not apply to the stakeholder's area, and the *Unknown* response indicates that the stakeholder currently lacks the necessary information to answer the performance marker.

The radial format of CAESAR8 offers a number of advantages for reading the results of the assessment. Of particular relevance is the centroid (center point) of the radial visualization (Draper, Livnat and Riesenfeld, 2009). For a *"Go / No Go"* decision, where senior executives use the graph to decide whether to launch a new system, mature projects should have turned the figure to a *green* status by progressing to the point in the center of the figure. The centroid logically denotes the end of the assessment. For the CAESAR8 model, the center is level 5, which considers optimizing the security strategy to improve future resilience. This is significant because the center ring is the smallest and will therefore be regarded as the least significant (Diehl, Beck and Burch, 2010) – which is indeed the case for CAESAR8.

Progress across the eight domains is more critical than achieving levels because the levels mainly ensure that the domain requirements are achieved in the correct order. Studies have shown that radial diagrams offer the best format for presenting progress in primarily one dimension (the domains in CAESAR8) (Diehl et al., 2010) and that this is read more accurately if this dimension is displayed in sectors as opposed to the rings (Diehl et al., 2010; Goldberg and Helfman, 2011). This quality of radial diagrams to highlight the symmetry across sectors (the domains) as the assessment moves between the levels is an important benefit (Goldberg and Helfman, 2011). While it may be easier to read two dimensions in a matrix style diagram, the prominence of a second dimension in the context of CAESAR8 assessments would detract from the more important first dimension of the CAESAR8 domain results.

3.3.3. Combination of stakeholder assessments

Multiple stakeholders complete independent assessments following the extended RAG status. All n independent assessments are consolidated to provide a final result. The collective, or consolidated, response is calculated using a worst-case formula. This ensures that all concerns are highlighted and outlier assessments are not suppressed due to groupthink or consensus-seeking mechanisms. The formula, shown in Equations 1 and 2, calculates the final result P_v for each performance marker P by examining the corresponding assessments P_s from all stakeholders s, where the set of performance markers M = Domains × Levels = $\{EF1, SG1, \dots, MI5, EA5\}$.

$$\forall P \in M : P_v = \begin{cases} \text{“No”}, & \text{if } \sum_{s=1}^{n} P_{sr} > 0, \\ \text{“?”}, & \text{if } \sum_{s=1}^{n} P_{su} > 0, \\ \text{“Trust”}, & \text{if } \sum_{s=1}^{n} P_{sb} > 0, \\ \text{“N/A”}, & \text{if } \sum_{s=1}^{n} P_{sv} = n, \\ \text{“Yes”}, & \text{if } \sum_{s=1}^{n} P_{sg} + \sum_{s=1}^{n} P_{sv} = n, \\ \text{“Partial”}, & \text{otherwise,} \end{cases} \tag{1}$$

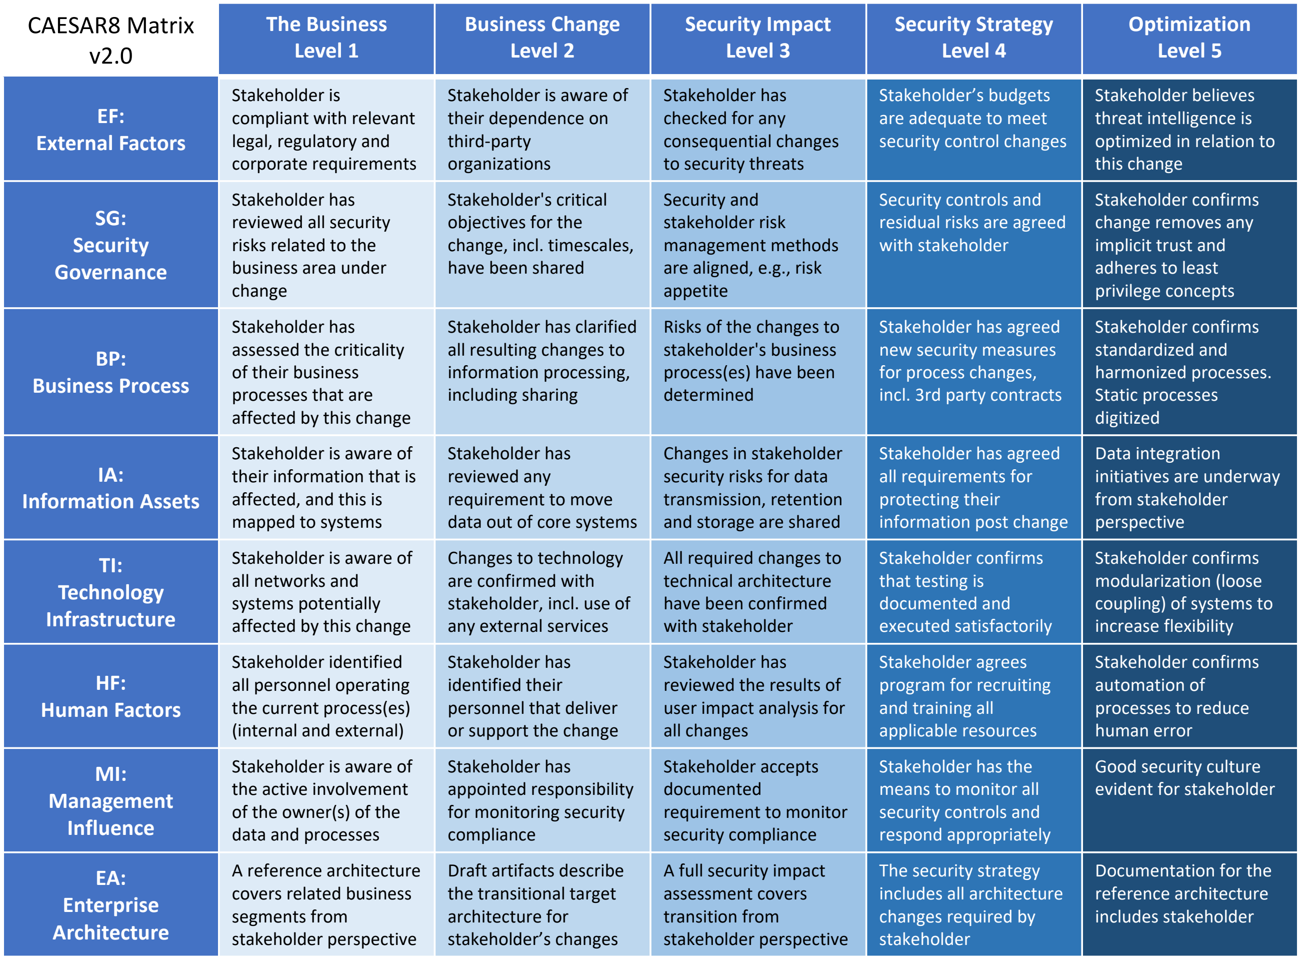

CAESAR8 Matrix v2.0	The Business Level 1	Business Change Level 2	Security Impact Level 3	Security Strategy Level 4	Optimization Level 5
EF: External Factors	Stakeholder is compliant with relevant legal, regulatory and corporate requirements	Stakeholder is aware of their dependence on third-party organizations	Stakeholder has checked for any consequential changes to security threats	Stakeholder's budgets are adequate to meet security control changes	Stakeholder believes threat intelligence is optimized in relation to this change
SG: Security Governance	Stakeholder has reviewed all security risks related to the business area under change	Stakeholder's critical objectives for the change, incl. timescales, have been shared	Security and stakeholder risk management methods are aligned, e.g., risk appetite	Security controls and residual risks are agreed with stakeholder	Stakeholder confirms change removes any implicit trust and adheres to least privilege concepts
BP: Business Process	Stakeholder has assessed the criticality of their business processes that are affected by this change	Stakeholder has clarified all resulting changes to information processing, including sharing	Risks of the changes to stakeholder's business process(es) have been determined	Stakeholder has agreed new security measures for process changes, incl. 3rd party contracts	Stakeholder confirms standardized and harmonized processes. Static processes digitized
IA: Information Assets	Stakeholder is aware of their information that is affected, and this is mapped to systems	Stakeholder has reviewed any requirement to move data out of core systems	Changes in stakeholder security risks for data transmission, retention and storage are shared	Stakeholder has agreed all requirements for protecting their information post change	Data integration initiatives are underway from stakeholder perspective
TI: Technology Infrastructure	Stakeholder is aware of all networks and systems potentially affected by this change	Changes to technology are confirmed with stakeholder, incl. use of any external services	All required changes to technical architecture have been confirmed with stakeholder	Stakeholder confirms that testing is documented and executed satisfactorily	Stakeholder confirms modularization (loose coupling) of systems to increase flexibility
HF: Human Factors	Stakeholder identified all personnel operating the current process(es) (internal and external)	Stakeholder has identified their personnel that deliver or support the change	Stakeholder has reviewed the results of user impact analysis for all changes	Stakeholder agrees program for recruiting and training all applicable resources	Stakeholder confirms automation of processes to reduce human error
MI: Management Influence	Stakeholder is aware of the active involvement of the owner(s) of the data and processes	Stakeholder has appointed responsibility for monitoring security compliance	Stakeholder accepts documented requirement to monitor security compliance	Stakeholder has the means to monitor all security controls and respond appropriately	Good security culture evident for stakeholder
EA: Enterprise Architecture	A reference architecture covers related business segments from stakeholder perspective	Draft artifacts describe the transitional target architecture for stakeholder's changes	A full security impact assessment covers transition from stakeholder perspective	The security strategy includes all architecture changes required by stakeholder	Documentation for the reference architecture includes stakeholder

Figure 2: CAESAR8 Matrix with 40 performance markers across eight domains (rows) and five levels (columns).

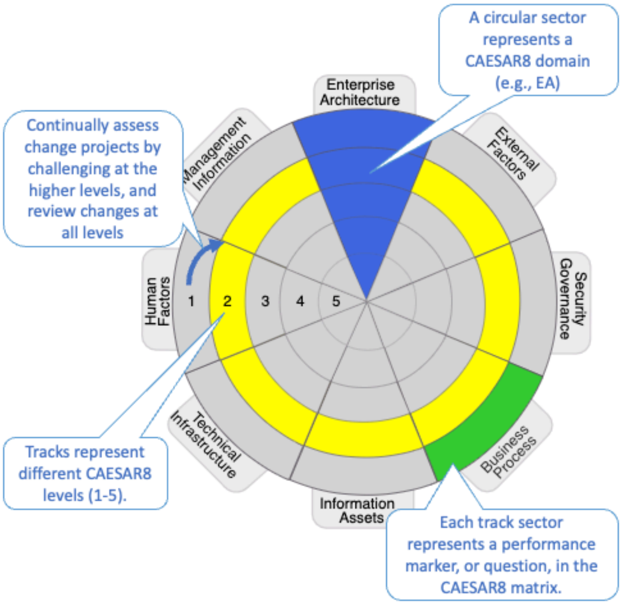
A circular sector represents a CAESAR8 domain (e.g., EA)
Enterprise Architecture
External Factors
Management Information
Continually assess change projects by challenging at the higher levels, and review changes at all levels
Human Factors
1
2
3
4
5
Security Governance
Technical Infrastructure
Business Process
Information Assets
Tracks represent different CAESAR8 levels (1-5).
Each track sector represents a performance marker, or question, in the CAESAR8 matrix.

Figure 3: Radial visualization for CAESAR8.

where

$$
\begin{aligned}
P_{sr} &= \begin{cases} 1, & \text{if } P_s = \text{“No”} \\ 0, & \text{otherwise} \end{cases} \\
P_{su} &= \begin{cases} 1, & \text{if } P_s = \text{“Unknown (?)”} \\ 0, & \text{otherwise} \end{cases} \\
P_{sb} &= \begin{cases} 1, & \text{if } P_s = \text{“Trust”} \\ 0, & \text{otherwise} \end{cases} \\
P_{sv} &= \begin{cases} 1, & \text{if } P_s = \text{“not applicable”} \\ 0, & \text{otherwise} \end{cases} \\
P_{sg} &= \begin{cases} 1, & \text{if } P_s = \text{“Yes”} \\ 0, & \text{otherwise.} \end{cases}
\end{aligned}
\tag{2}
$$

4. Case Study: The Gangs Matrix

We use the case of the *Gangs Matrix*, an IT project pursued by the police in London, UK, to show how CAESAR8 works in the context of a project. This case study shows how CAESAR8 would be used by experts to encode their evaluations, and how the result would be visualized, particularly how CAESAR8 highlights problem areas in individual assessments and emphasizes areas of agreement or conflict between stakeholder assessments.

The Gangs Matrix is particularly suitable as a case study because its security shortcomings were documented in an independent investigation by the UK's data protection regulator ICO (Information Commissioner, 2018), which we used as the basis for the case study. The case study can therefore check whether the CAESAR8 performance markers – identified by scientific analysis of the literature – are able to detect the known information security failings, and if so, how early in the project.

While we conducted the CAESAR8 assessments retrospectively and without inside knowledge of the business, the detailed description of the ICO findings allowed for a high degree of confidence in our assessments.

4.1. Background

In an ongoing effort to reduce the serious crimes committed by gangs in London, UK, the Metropolitan Police Service (MPS) desired to prosecute more offenders and to deter young people from engaging in such crime. The *Gangs Operating Model* was the MPS's central strategy for dealing with gang crime, and the 32 separate local boroughs of the MPS were responsible for implementing the model. The model required that each local borough created its own *Gangs Matrix*, an intelligence database, to record details of gang members. These were then compiled centrally to form a London-wide matrix. However, the ICO investigation found that the Gangs Matrix also contained victims of gang crime. The MPS's operation of the Gangs Matrix resulted in a breach that disclosed the data of gang victims, and the subsequent investigation by the ICO discovered multiple contraventions of the Data Protection Act 1998 (DPA).

Details of the specific contraventions, set out in the ICO's enforcement notice (Information Commissioner, 2018), show that this project suffered significant issues between the implementation of new technology and a change to working practices, and the effect of these issues on information security. Our CAESAR8 assessments are based on seven specific findings:

1. Excessive information sharing with third parties without any agreement regarding the use and protection of information;
2. Lack of protection of the information that was shared with third parties, and no consideration for the sensitivity of the data in question;
3. Routine transfer of Gangs Matrix information without appropriate security, such as encryption;
4. Creation of local copies of unprotected data;
5. No revocation of access privileges when users of the Gangs Matrix moved out of gang-related roles;
6. Lack of governance and central oversight which allowed poor and unlawful processing of data to go unchallenged;
7. Lack of central data protection guidance and failure to monitor compliance with guidance for the operation of the Gangs Matrix. No privacy impact assessment for the system's personal data.

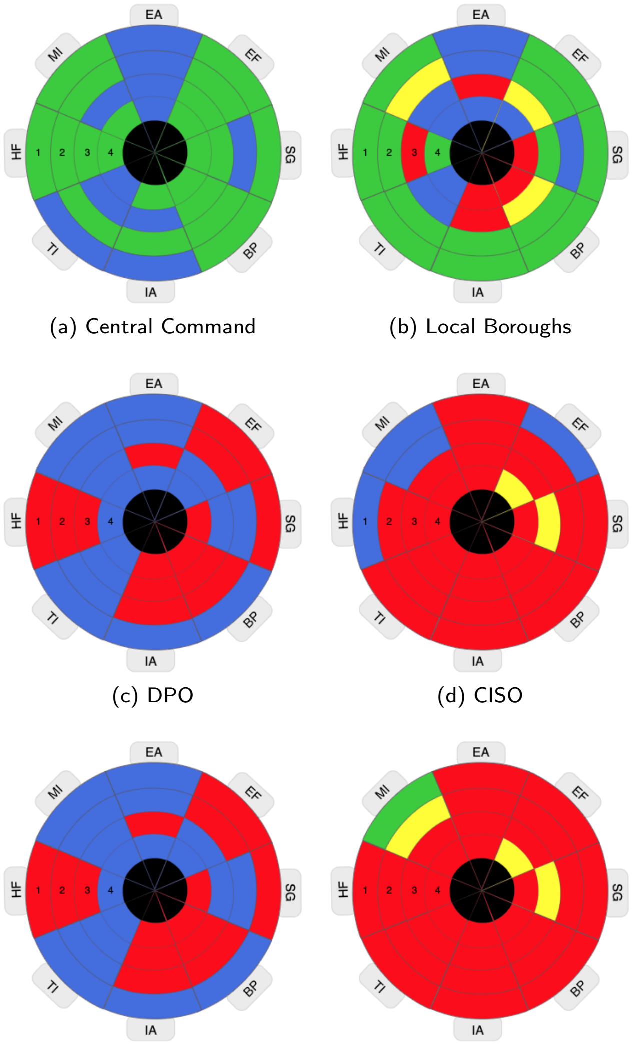

(a) Central Command

(b) Local Boroughs

(c) DPO

(d) CISO

(e) Central and DPO

(f) All stakeholders

Figure 4: CAESAR8 results for Gangs Matrix stakeholders

4.2. Methodology

The case study was conducted by the first author, who has previously worked for a UK police force. The assessments are based solely on the evidence provided in the ICO report. Four separate stakeholder assessments were conducted: i) Central Command; ii) a joint assessment for the 32 Local Boroughs; iii) the Data Protection Officer (DPO); and iv) the Chief Information Security Officer (CISO). In reality, more stakeholders would probably be identified for this project, including a separate assessment for each local borough, but these four already demonstrate the usefulness of CAESAR8. The four separate assessments were then consolidated to produce a final result. We did not assess level 5 (optimization) because the ICO enforcement notice did not have enough information for this level.

4.3. Results

Individual assessments for all four stakeholders are shown in Figure 4, along with their aggregated results. Figure 4a shows that, from Central Command's perspective, the project is on track with no open issues. If Central Command only relied on their own assessment and did not consult a DPO or CISO, this perspective may explain why the project went ahead in reality. However, the assessments by the other three stakeholders (Figures 4b–4d) would have raised issues: the DPO and CISO's assessments early in the project at level 1, and the local boroughs' assessment at level 3. This shows why explicit consideration of multiple stakeholder views in CAESAR8 is valuable.

The combined assessment of all four stakeholders (Figure 4f) shows that issues in the project could have been discovered (and fixed) very early in the project if a systematic approach such as CAESAR8 was in place.

Finally, Figure 4e shows the aggregation of the assessments by Central Command and DPO, which indicates which security risks could have been identified even if the CISO was not aware of the project. It is clear that the DPO is likely to have expressed caution at the start of level 1. In the aggregated assessment, the DPOs concerns have over-ridden the otherwise optimistic assessment by Central Command concerning the readiness to go live with the new information system. This ability to select which stakeholders to combine is a useful feature of CAESAR8, and can be used in *what-if* analysis during risk assessment and treatment.

In summary, this case study highlights how security problems can easily occur without consistent management checks for good security governance. CAESAR8 allows stakeholders to check whether significant issues have been considered in ongoing projects. Aggregating the assessments of different stakeholders allows for the full identification of potential issues, which is not possible when relying on individual assessments alone.

5. Evaluation of CAESAR8

We conducted a formal evaluation of CAESAR8 to determine, (1) whether the final CAESAR8 design meets our design goals, and (2), whether and how real-world users benefit from using CAESAR8.

The evaluation process was a summative ex post process using invited professionals. The professionals were selected for their experience (over 10 years working with information security risks) and their diversity of knowledge, e.g., based on working in different business areas, different roles, different sectors (including the public sector), and organizations of different sizes.

5.1. Methodology

The evaluation consisted of three steps. First, a pre-evaluation questionnaire was used to confirm the demographics and experience of the participants, and to check their agreement with the common problem areas (Table 1). Second, participants used a web app implementation of CAESAR8 to conduct CAESAR8 assessments for their own projects from the perspective of different stakeholders. Third, a post-evaluation questionnaire was used to check to what extent the CAESAR8 design goals were met, and how well CAESAR8 helps to address the common problem areas. We obtained ethical approval from our university's IRB prior to conducting the evaluation.

5.1.1. Participants

The evaluation required a substantial time commitment from participants as well as substantial experience on cyber-security related projects. We therefore recruited 14 participants who each had more than ten years experience of working in cyber security projects. Half of the participants were information security professionals, and the other half consisted of project managers, business and operations managers, software engineers, change managers, program managers, and auditors. Three participants dropped out before completing the post-evaluation questionnaire. Their results have therefore been excluded.

5.1.2. Questionnaires

Both pre- and post-evaluation questionnaires were semi-structured. Participants were asked to rate agreement with problem areas, CAESAR8 benefits, and realization of design goals on a 5-level Likert scale. In addition, participants could give free-text responses to explain their scores or make other comments. Participants were given access to the CAESAR8 web app after completing the pre-evaluation questionnaire. Access to the post-evaluation questionnaire was given manually after verifying that participants had interacted with the web app.

5.1.3. CAESAR8 web app

The web app contained two short training videos[1]: a 3-minute video to explain the design of the model and a 10-minute visual tutorial of how to use the web app. These videos ensured that all volunteers received the same level of instruction on the model and app.

Participants were then asked to configure their own set of stakeholders, complete CAESAR8 assessments for each stakeholder, i.e., answer the 40 performance markers from the perspective of each stakeholder, and view individual and aggregate results. Some participants completed their interactions with the web app in a single session, whereas others returned to the web app repeatedly over several days. On average, participants spent 26 minutes completing three separate assessments (excluding training videos and stakeholder configuration).

5.2. Common Problem Areas

Participants rated the importance of the 15 common problem areas (see Table 1) as well as CAESAR8's performance in addressing them. Figure 5 shows agreement with the problem areas and the corresponding CAESAR8 benefits. Each point on the graph indicates the average response for a problem/benefit combination. The figure shows

[1] https://youtu.be/2fLuOi4aC-s

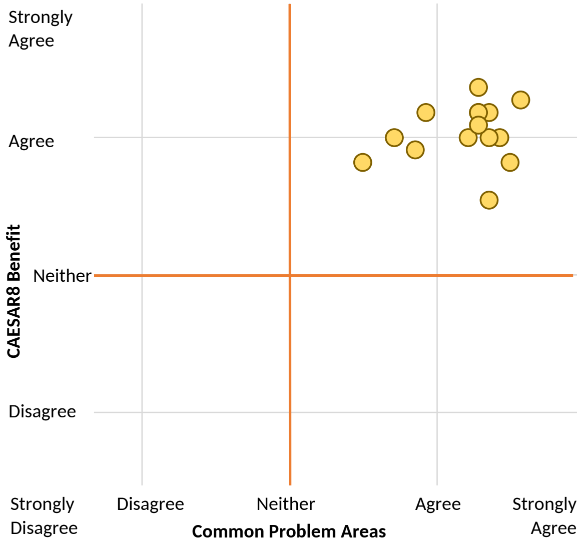

Figure 5: Common problems matched to CAESAR8 benefits

that all problem areas are positioned in the upper-right quadrant, confirming that the problem areas are indeed real information security problems experienced in projects and, importantly, that CAESAR8 helps to address them.

Figure 6 shows detailed results for all problem areas. A novel feature of CAESAR8 is its ability to capture the broad perspectives of all those stakeholders that are involved in a project. These are specifically related to the five problems areas marked with a star in Figure 6. We can see that perceived CAESAR8 benefit in these areas is close to the areas' perceived importance, with average scores between *Agree* and *Strongly Agree*.

Regarding the overall benefit of CAESAR8, participants commented that it is "a really strong concept and definitely a useful way of evaluating security risk," and that the assessment "was able to be populated quickly and yet provided holistic coverage. The tool provides a framework to enable project governance and point towards areas that require greater effort or scrutiny. The product, in my view, would be a positive mechanism to organizations."

For only two problem areas, marked with a triangle in Figure 6, the perceived CAESAR8 benefit was more than half a point[2] lower than the perceived importance of the problem area. The first problem area, *inadequate testing*, indicates that thorough testing of security controls is important to fully understand potential failure modes. However, the corresponding CAESAR8 performance marker (TI4) was worded to only check for a test plan, not that the plan was put into operation. As a result, we updated the wording following the evaluation (Figure 2 already shows the updated version).

The second problem area, *lack of compliance monitoring*, indicates that although CAESAR8 encourages checks for monitoring of system performance, management still needs to enforce this requirement to ensure that compliance monitoring actually happens. However, CAESAR8 helps by providing checks across the maturity levels, such as checking for responsibility, a documented procedure, and the means to carry it out.

5.3. Design Goals

To check whether CAESAR8 indeed provides a practical solution for businesses, the post-evaluation questionnaire asked a series of questions to find out how well the design goals for CAESAR8 had been achieved. The results (Figure 7) show that participants agreed that CAESAR8 achieved its design goals, with a median score of *Agree* for all design goals.

Free-text comments given by participants supported this positive evaluation of CAESAR8. Regarding the design of CAESAR8, participants stated that its "ease of use would assist in continued assessment" (goal 4) and "it is useful in providing a dashboard to highlight areas of concern" (goal 6). Participants also praised the usability, commenting that it is "very easy to navigate the model" (goal 7) and its "questions [...] are easy to understand and apply" (goal 8). The governance focus of CAESAR8 received the most positive comments, stating that CAESAR8 "allows an easy overarching view providing areas of a project that require enhanced attention and effort in order to achieve the desired business assurance" (goal 10), "allow[s] all parts of an organization to be involved in the assessment and ongoing monitoring of cyber security, with significant benefits to the organization" (goal 11), and praising "the results page and how the different stakeholders results can be overlaid to produce an overall picture" (goal 12).

Some comments also explained the reasons for lower agreement ratings. Most significantly, one participant disagreed that the questions were easy to understand, but explained that the level 5 questions (security optimization) needed security expertise and were less suitable for non-technical stakeholders. This is indeed the case, and we believe that level 5 could be removed for non-technical stakeholders. Security-mature projects would only need agreement up to and including level 4, which provides sufficient assurance to support the decision to go live.

5.4. Stakeholder Involvement

To evaluate how well CAESAR8 realizes the involvement of multiple stakeholders and their contributions, we asked additional questions about this aspect of CAESAR8. The results, shown in Figure 8, indicate that the most highly rated aspect of stakeholder involvement in CAESAR8 is the ability to include diverse stakeholders from non-security roles – "This is a great model to ensure the 'buy in' from all stakeholders."

In addition, participants praised the ability for stakeholders to rate only those performance markers in their area of responsibility and to indicate when they *trust* other stakeholders to take responsibility for a performance marker, commenting that "this model makes those gaps and those areas of 'assumed trust' immediately clear. This is something

[2] The >0.5 value was selected because it represents an overall shift in the CAESAR8 benefit score from that of the original problem score on the five-point Likert scale.

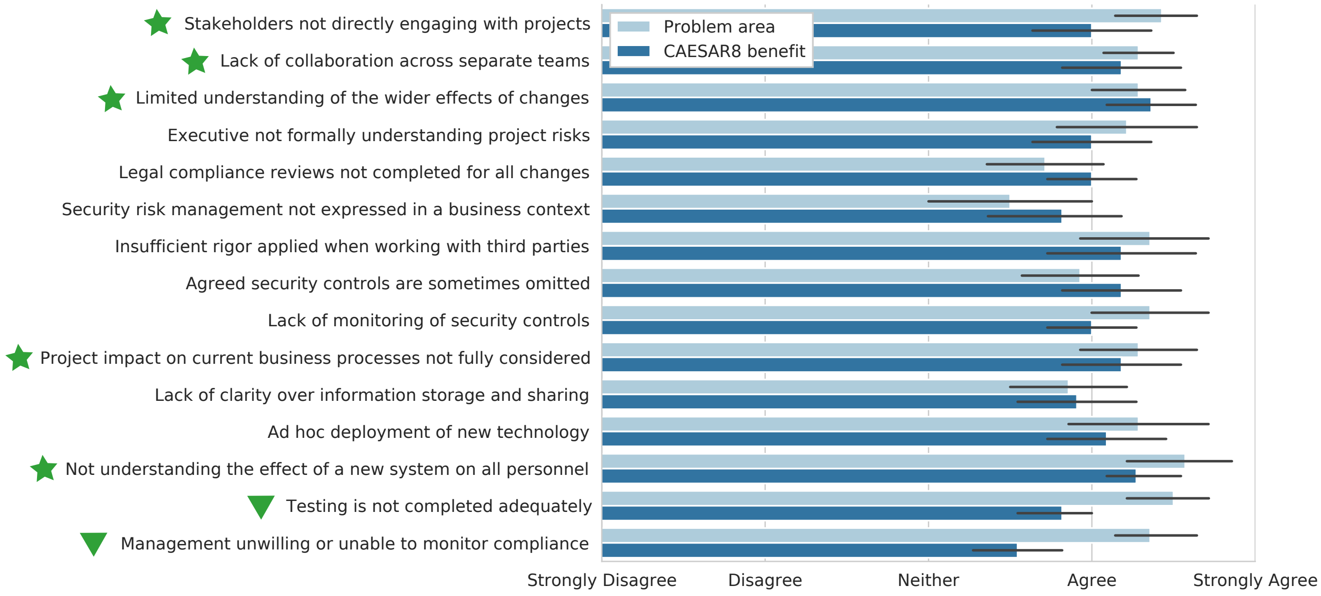

Figure 6: Extent to which participants agreed with 15 common problem areas (light blue), and with CAESAR8 benefiting each area (dark blue). Error bars indicate the 95% confidence interval for the mean.

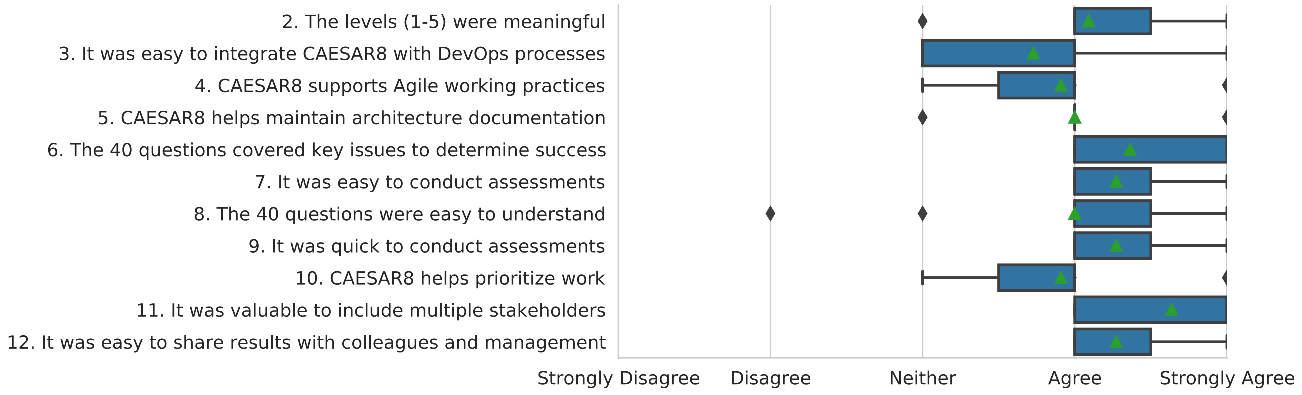

Figure 7: Extent to which participants agreed that CAESAR8 met our 12 design goals (note that we did not evaluate agreement with the first design goal because it described the intrinsic design of the model). Boxes indicate the lower and upper quartiles, and the green triangle indicates the average score.

that is enormously powerful," and that "identifying where one stakeholder is trusting another is a strong feature."

5.5. Time Needed for CAESAR8 Assessments

We instrumented the web app to record timestamps for all actions performed by the participants. This data allows us to analyze how long the CAESAR8 assessments took for each performance marker. Note that this data includes large outlier values that occurred when participants left the web app open while taking a break or working on other unrelated tasks. For this reason, we plot the results in a box plot which focuses on the quartiles and median, which are not sensitive to outliers. We also omit outliers in the plot.

Figure 9 shows the time taken per performance marker. We observe that participants took longest for EF1 (median of 16s), which is the first performance marker in the assessment. This is likely because they opened the assessment page and looked over all performance markers before beginning the assessment.

On average, the median time taken is 10.05s per performance marker, which means that a stakeholder can complete a full assessment in as little as 7 minutes, assuming that the stakeholder is familiar with the project and how it relates to their own area of responsibility.

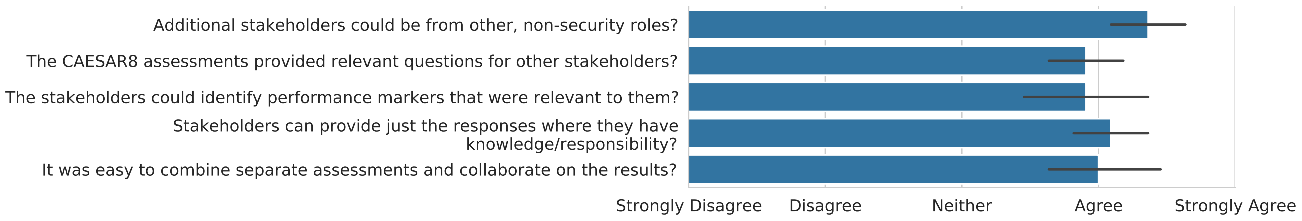

Figure 8: Extent to which participants agreed with aspects of stakeholder involvement in CAESAR8. Error bars indicate the 95% confidence interval for the mean.

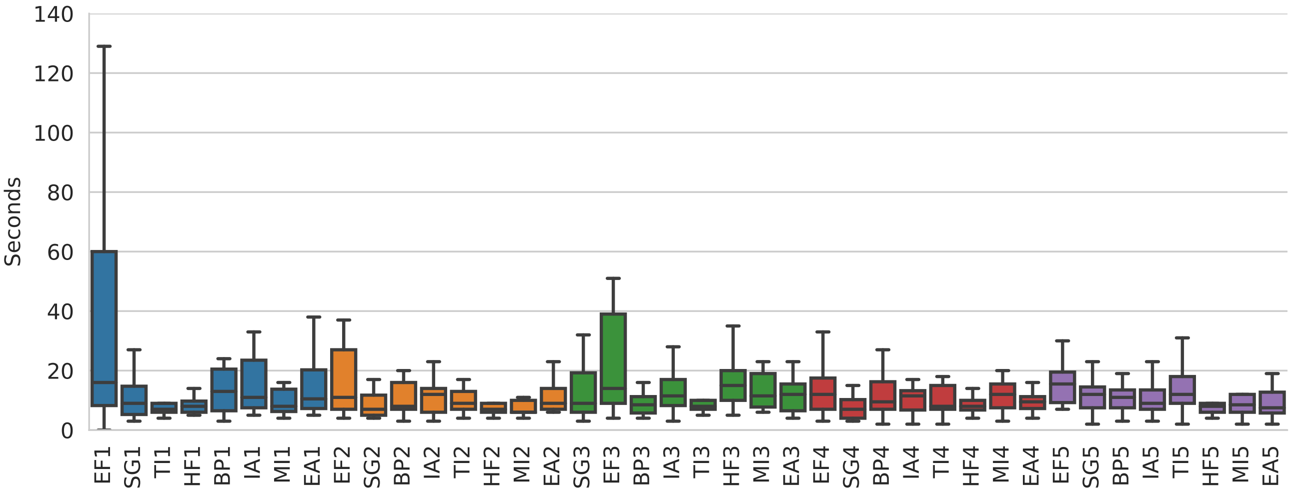

Figure 9: Time taken per performance marker.

6. Conclusion

> “The assessments contain some powerful questions. The impact each set of assessment questions has on the result is really clear and very telling. I believe that it will make the process of initial evaluation of cyber security, and ongoing monitoring of cyber security related processes, understandable and easily accessible, even to small organizations with, potentially, a lesser basic understanding of cyber.”
>
> *Senior technical auditor, 25+ years experience in impact of information security risks on safety.*

We have presented a novel model, CAESAR8, that helps organizations manage information security risks under the constraints experienced in agile projects.

Contributions. CAESAR8 supports an holistic approach to assessing the information security risks by (1) reviewing projects across eight domains that are common causes for information security failures, (2) ensuring systematic progress across five maturity levels, and (3) involving multiple independent stakeholders in the assessment to access tacit stakeholder knowledge while avoiding concurrence-seeking and groupthink.

Contributions for practice. In this way, CAESAR8 helps to avoid many of the common pitfalls for information security without the need for prior assurance activities that can hinder agile development. Through its incremental reviews, CAESAR8 can uncover key issues in early project iterations. If used continuously, CAESAR8 supports the development of an enterprise information security architecture within agile teams. CAESAR8 can also be used to assess safety in Operational Technology (OT) projects and provide an integrated perspective of safety and security. Our evaluation with experienced industry professionals confirmed that CAESAR8 is easy to use, supports continued assessment and monitoring of cyber security, and provides holistic coverage, with significant benefits especially for smaller organizations.

Limitations. CAESAR8 is a promising model for holistic reviews of information security risks in real-time. However, while CAESAR8 allows evaluating the current cyber security posture, it does not provide a method to move to the next maturity level in the best or easiest possible way.

CAESAR8 provides a generic assessment in relation to the overall status of information security risks, but it does not attempt to define what the individual risks are in detail. It still requires the skills of individual stakeholders and subject matter experts, using their own standards and tools, to determine how to respond to specific performance markers. This process may look very different for different stakeholders.

Future research. Engaging the correct stakeholders when conducting CAESAR8 assessments for a given project is key to the accuracy of the final CAESAR8 result. However, more research is needed into methods for automating the selection process and the best methods to engage stakeholders in the continual CAESAR8 assessment process.

In addition, a useful extension of CAESAR8 would be as a tool to compare before and after, for example, comparing the status of security risks before/after a project, which could be used in modeling to highlight the potential security impact of projects.

References

Andrews, C., Monk, C., Johnston, R., 2014. Integrated architecture framework and security risk management for complex systems IET.

Atanasov, P., 2020. Small steps to accuracy_ Incremental belief updaters are better forecasters. Organizational Behavior and Human Decision Processes , 17.

Axelrod, C.W., 2008. Accounting for value and uncertainty in security metrics. Information Systems Control Journal 6, 1–6.

Bahmani, F., Shariati, M., Shams, F., 2010. A survey of interoperability in Enterprise Information Security Architecture frameworks, in: Information Science and Engineering (ICISE), 2010 2nd International Conference on, IEEE. pp. 1794–1797.

Beck, K., Beedle, M., Van Bennekum, A., Cockburn, A., Cunningham, W., Fowler, M., Grenning, J., Highsmith, J., Hunt, A., Jeffries, R., et al., 2001. Manifesto for agile software development .

Belkadi, S., Cherti, I., Bahaj, M., 2018. Lean in information technology: Produce the human before the software, in: International Conference on Advanced Intelligent Systems for Sustainable Development, Springer. pp. 203–213.

Bernaert, M., Poels, G., Snoeck, M., Backer, M.D., 2014. Enterprise architecture for small and medium-sized enterprises: a starting point for bringing ea to smes, based on adoption models, in: Information systems for small and medium-sized enterprises. Springer, pp. 67–96.

Beznosov, K., Kruchten, P., 2004. Towards agile security assurance, in: Proceedings of the 2004 workshop on New security paradigms, ACM. pp. 47–54.

Bischoff, S., Aier, S., Winter, R., 2014. Use it or lose it? the role of pressure for use and utility of enterprise architecture artifacts, in: 2014 IEEE 16th Conference on Business Informatics, IEEE. pp. 133–140.

Chivers, H., Paige, R.F., Ge, X., 2005. Agile Security Using an Incremental Security Architecture, in: Baumeister, H., Marchesi, M., Holcombe, M. (Eds.), Extreme Programming and Agile Processes in Software Engineering, Springer Berlin Heidelberg. pp. 57–65.

Diehl, S., Beck, F., Burch, M., 2010. Uncovering strengths and weaknesses of radial visualizations—an empirical approach. IEEE Transactions on Visualization and Computer Graphics 16, 935–942.

Diesch, R., Pfaff, M., Krcmar, H., 2020. A comprehensive model of information security factors for decision-makers. Computers & Security 92, 101747.

Dorca, V., Munteanu, R., Popescu, S., Chioreanu, A., Peleskei, C., 2016. Agile approach with Kanban in information security risk management, in: 2016 IEEE International Conference on Automation, Quality and Testing, Robotics (AQTR), pp. 1–6. doi:10.1109/AQTR.2016.7501278.

Draper, G.M., Livnat, Y., Riesenfeld, R.F., 2009. A survey of radial methods for information visualization. IEEE transactions on visualization and computer graphics 15, 759–776.

Gawande, A., 2011. The checklist manifesto: How to get things right. Journal of Nursing Regulation 1, 64.

Goldberg, J., Helfman, J., 2011. Eye tracking for visualization evaluation: Reading values on linear versus radial graphs. Information visualization 10, 182–195.

Goudalo, W., Seret, D., 2009. The Process of Engineering of Security of Information Systems (ESIS): The Formalism of Business Processes, IEEE. pp. 105–113. doi:10.1109/SECURWARE.2009.24.

Hart, P., 1991. Irving l. janis' victims of groupthink. Political Psychology , 247–278.

Hedesstrom, T., Whitley, E.A., 2000. What is meant by tacit knowledge? towards a better understanding of the shape of actions., in: ECIS, pp. 46–51.

Hevner, A.R., March, S.T., Park, J., Ram, S., 2004. Design science in information systems research. MIS quarterly , 75–105.

Hong, L., Page, S.E., 2004. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. Proceedings of the National Academy of Sciences 101, 16385–16389. doi:10.1073/pnas.0403723101.

Hopmere, M., Crawford, L., Harré, M.S., 2020. Proactively monitoring large project portfolios. Project Management Journal 51, 656–669.

Information Commissioner, 2018. Metropolitan police service enforcement notice. URL: https://ico.org.uk/media/action-weve-taken/enforcement-notices/2260336/metropolitan-police-service-20181113.pdf.

Janis, I.L., 1971. Groupthink. Psychology today 5, 43–46.

Jugdev, K., Thomas, J., 2002. 2002 student paper award winner: Project management maturity models: The silver bullets of competitive advantage? Project management journal 33, 4–14.

Kahneman, D., 2011. Thinking, fast and slow. Macmillan.

Kahneman, D., Tversky, A., 1977. Intuitive prediction: Biases and corrective procedures. Technical Report. Decisions and Designs Inc Mclean Va.

Kaisler, S., Armour, F., 2017. 15 years of enterprise architecting at hicss: Revisiting the critical problems, in: Proceedings of the 50th Hawaii International Conference on System Sciences.

Kaisler, S.H., Armour, F., Valivullah, M., 2005. Enterprise architecting: Critical problems, in: System Sciences, 2005. HICSS'05. Proceedings of the 38th Annual Hawaii International Conference on, IEEE. pp. 224b–224b.

Khoshgoftar, M., Osman, O., 2009. Comparison of maturity models, in: 2009 2nd IEEE International Conference on Computer Science and Information Technology, IEEE, Beijing, China. pp. 297–301. doi:10.1109/ICCSIT.2009.5234402.

Kirkebøen, G., 2009. Decision Behaviour- Improving Expert Judgement, in: Williams, T.M., Samset, K., Sunnevåg, K.J. (Eds.), Making Essential Choices with Scant Information. Palgrave Macmillan UK, London, pp. 169–194. doi:10.1057/9780230236837_9.

Korhonen, J.J., Lapalme, J., McDavid, D., Gill, A.Q., 2016. Adaptive Enterprise Architecture for the Future: Towards a Reconceptualization of EA, in: 2016 IEEE 18th Conference on Business Informatics (CBI), IEEE, Paris, France. pp. 272–281. doi:10.1109/CBI.2016.38.

Kotusev, S., 2019. Fake and real tools for enterprise architecture: The zachman framework and business capability model. Enterprise Architecture Professional Journal , 1–14.

Kotusev, S., Kurnia, S., 2020. The theoretical basis of enterprise architecture: A critical review and taxonomy of relevant theories. Journal of Information Technology doi:10.1177/0268396220977873.

Kotusev, S., Singh, M., Storey, I., 2015. Investigating the usage of enterprise architecture artifacts, in: ECIS 2015 Research-in-Progress Papers.

Kurnia, S., Kotusev, S., Shanks, G., Dilnutt, R., Milton, S., 2021. Stakeholder engagement in enterprise architecture practice: What inhibitors are there? Information and Software Technology 134, 106536.

Li, X., Hongyan, L., 2010. Proposal for information security architecture based on a company, in: Communication Systems, Networks and Applications (ICCSNA), 2010 Second International Conference on, IEEE. pp. 17–20.

Loft, P., He, Y., Janicke, H., Wagner, I., 2019. Dying of a hundred good symptoms: Why good security can still fail - a literature review and analysis. Enterprise Information Systems 15, 1–26. doi:10.1080/17517575.2019.1605000.

Löhe, J., Legner, C., 2014. Overcoming implementation challenges in enterprise architecture management: a design theory for architecture-driven it management (adrima). Information Systems and e-Business

Management 12, 101–137.

Madison, J., 2010. Agile Architecture Interactions. IEEE Software 27, 41–48. doi:10.1109/MS.2010.35.

McClintock, M., Falkner, K., Szabo, C., Yarom, Y., 2020. Enterprise security architecture: Mythology or methodology?, in: Proceedings of the 22nd International Conference on Enterprise Information Systems (ICEIS 2020) – Volume 2, pp. 679–689.

Milanovich, D.M., Driskell, J.E., Stout, R.J., Salas, E., 1998. Status and cockpit dynamics: A review and empirical study. Group Dynamics: Theory, Research, and Practice 2, 155–167. doi:10.1037/1089-2699.2.3.155.

Mumpower, J.L., Stewart, T.R., 1996. Expert Judgement and Expert Disagreement. Thinking & Reasoning 2, 191–212. doi:10.1080/135467896394500.

Page, S., 2019. The diversity bonus. Princeton University Press.

Peffers, K., Tuunanen, T., Rothenberger, M.A., Chatterjee, S., 2007. A design science research methodology for information systems research. Journal of management information systems 24, 45–77.

Ross, J.W., Weill, P., Robertson, D.C., 2006. Enterprise Architecture As Strategy: Creating a Foundation for Business Execution. Harvard Business School Press, Boston, Mass.

Sadki, S., El Bakkali, H., 2014. Towards controlled-privacy in e-health: A comparative study, in: Multimedia Computing and Systems (ICMCS), 2014 International Conference on, IEEE. pp. 674–679.

Savola, R.M., Heinonen, P., 2011. A visualization and modeling tool for security metrics and measurements management, in: 2011 Information Security for South Africa, IEEE. pp. 1–8.

Scriven, M., 2000. The logic and methodology of checklists. Technical Report. Western Michigan University.

Sen, R., Borle, S., 2015. Estimating the Contextual Risk of Data Breach: An Empirical Approach. Journal of Management Information Systems 32, 314–341. doi:10.1080/07421222.2015.1063315.

Sherwood, J., Clark, A., Lynas, D., 2005. Enterprise Security Architecture: A Business-Driven Approach. 1 edition ed., CRC Press, San Francisco.

Siponen, M., Willison, R., 2009. Information security management standards: Problems and solutions. Information & management 46, 267–270.

Skjong, R., Wentworth, B.H., 2001. Expert Judgment and Risk Perception, in: Proceedings of the Eleventh International Offshore and Polar Engineering Conference, pp. 537–544.

Soomro, Z.A., Shah, M.H., Ahmed, J., 2016. Information security management needs more holistic approach: A literature review. International Journal of Information Management 36, 215–225. doi:10.1016/j.ijinfomgt.2015.11.009.

Spears, J.L., 2004. A holistic risk analysis method for identifying information security risks, in: Working Conference on Integrity and Internal Control in Information Systems, Springer. pp. 185–202.

Sur, M.D., Schindler, N., Singh, P., Angelos, P., Langerman, A., 2016. Young surgeons on speaking up: when and how surgical trainees voice concerns about supervisors' clinical decisions. The American Journal of Surgery 211, 437–444. doi:10.1016/j.amjsurg.2015.10.006.

The Open Group, 2011. TOGAF Version 9.1. 10th ed., van Haren Publishing, Zaltbommel.

Wang, H., Xu, H., Lu, B., Shen, Z., 2009. Research on security architecture for defending insider threat, in: 2009 Fifth International Conference on Information Assurance and Security, IEEE. pp. 30–33.

Weiser, T., Haynes, A., 2018. Ten years of the surgical safety checklist. Journal of British Surgery 105, 927–929.

Zachman, J.A., 1987. A framework for information systems architecture. IBM systems journal 26, 276–292.