# Toward a Context-Aware Methodology for Information Security Governance Assessment Validation *

Marco Angelini[0000−0001−9051−6972], Silvia Bonomi[0000−0001−9928−5357], Claudio Ciccotelli[0000−0003−4687−8241], and Alessandro Palma

Department of Computer, Control, and Management Engineering "Antonio Ruberti"
Sapienza University of Rome
Via Ariosto 25, 00185 Rome, Italy
{angelini, bonomi, ciccotelli}@diag.uniroma1.it,
palma.1871556@studenti.uniroma1.it

**Abstract.** Conducting a cybersecurity assessment is a central activity in protecting a generic organization from cyber-attacks. Several methods exist in research and industry to assess the security level of an organization, from manual activities to automated attack graphs. Unfortunately, automated approaches fail in taking into account the governance aspect that still need to be evaluated manually by the assessor, introducing possible biases or problems deriving from the level of expertise. In this paper, we provide a methodology to support the assessor in the task of evaluating the coverage of cybersecurity controls coming from technical standards, regulations, internal practices. This is done by providing him/her with a multi-layer model that takes into account several organizational layers, a mapping procedure to tie the security controls to the multi-layer model, and the definition of a validation factor that can be used to possibly refine the level of coverage and to suggest possible layers where evidences should be collected to verify and assess the coverage of a security control. A usage scenario provides an initial validation of our approach based on ISO 27001. Developments of this methodology are on-going toward its application to the support of broader cyber-risk assessment activities through discounting risk factors.

**Keywords:** Information Security Governance · Risk Assessment · ISO 27001 · Multi-layer model.

## 1 Introduction

According to NIST SP 800-100 [13], Information Security Governance (ISG) can be defined as *"the process of establishing and maintaining a framework and supporting management structure and processes to provide assurance that information security strategies (i) are aligned with and support business objectives,*

* This work has been partially supported by the EU H2020 PANACEA project under the Grant Agreement n. 826293.

*(ii) are consistent with applicable laws and regulations through adherence to policies and internal controls and (iii) provide assignment of responsibility, all in an effort to manage risk".*

Said differently, ISG is the global effort needed to ensure the well-being of the company's electronic resources and it should be supported by effective and efficient processes. In order to evaluate company's cyber security posture and to measure the adequacy and maturity of its ISG processes, periodic risk assessment must be carried out in order to estimate (quantitatively or qualitatively) the risk related to possible cyber incidents by evaluating (i) the level of exposure to specific threats analyzing existing vulnerabilities and the maturity level of the governance environment and (ii) the impact of the incidents generated by threats materialization.

How to correctly structure the risk management process is suggested by multiple standards and best practices. However, how to support, through automatic tools, assessors in this complex and delicate task is still an open problem, especially when assessing the governance aspects.

Currently, risk assessment is still mainly a manual process carried on by expert assessors through interviews with relevant stakeholders, analysis of technical elements (e.g., data extracted from network and vulnerability scans) and collection of evidences that allow them to properly evaluate the maturity of the governance process. The result is a non-completely objective and time-consuming activity where the result may be deeply influenced by the expertise of the assessor and his/her personal sensitivity.

Recently, automated approaches to dynamic risk estimation have been proposed to support risk assessors with a fast technical assessment (e.g., [7]). These approaches are based on attack graphs to estimate the likelihood of possible attack patterns and on mission impact models to estimate the consequences of a successful attack. Unfortunately, they fail in taking into account the governance aspects that still need to be evaluated manually by the assessor.

In this paper, we take a step to close this gap by providing a methodology to support the assessor in the task of evaluating the coverage of cyber security controls (e.g., the coverage of NIST Cybersecurity framework controls). This is done by providing him/her with the definition of a validation factor that can be used to possibly refine the level of coverage and to suggest possible layers where evidences should be collected to verify and assess the coverage of a security control.

This is done basically in two steps: (i) we map security controls performed by the assessor to a multi-layer model representing the most relevant components of an enterprise (i.e., we create a contextual map between controls and the company under analysis) and (ii) based on the obtained mapping, we compute a validation factor that provides the assessor with an indication about the need of a possible further analysis of the control due to its incidence on multiple organizational layers.

We also provide a usage scenario where we discuss how to apply our methodology to an assessment performed against ISO 27001.

The rest of the paper is organized as follows: Section 2 provides an overview of the main risk assessment methodologies and on attack graph models, Section 3 introduces the multi-layer model we used to identify relevant organizational layers used by the proposed methodology to map security controls, Section 4 presents our methodology, Section 5 introduces the definition of the validation factor for coverage, Section 6 discusses a usage scenario and finally Section 7 concludes the paper.

## 2 Related Work and Background

**Risk assessment methodologies.** The most common risk model is based on two factors: *likelihood* and *impact*. Currently, there exist different methodologies and tools supporting the two-factor risk assessment. OWASP [10] includes a risk assessment framework based on the two-factor evaluation. In the OWASP risk rating methodology the likelihood is estimated by assessing parameters related to threat agents (skill level, motive, opportunity and size) and vulnerabilities (ease of discovery and exploit, awareness and intrusion detection) while, for the impact, it takes into account technical impact (loss of confidentiality, integrity, availability and accountability) and business impact (financial and reputation damage, non compliance and privacy violation).

MEHARI (MEthod for Harmonized Analysis of RIsk) [6] is a free, open-source risk management methodology where the risk assessment task is decomposed in three main activities: (i) *risk identification* i.e., identification of assets, vulnerabilities and threats, (ii) *risk estimation* in terms of seriousness and (iii) *risk evaluation* in terms of its acceptability. As for OWASP, also MEHARI considers the two-factor risk but, in this case, both likelihood and impact are considered *intrinsic* (i.e., with no consideration of security measures) and then *reduction factors* may be applied (i.e., dissuasion and prevention for likelihood, and protection and palliation for impact).

EBIOS [3] is a risk assessment tool supporting the two-factor risk model. Differently from the others, it stresses the importance of the impact generated from different sources as humans, services, financial, legal and reputation. The assessment criteria used in EBIOS deal with exposure (dependency and penetration) and cyber reliability (maturity and trust).

When considering cybersecurity frameworks, we can relate the mentioned risk-assessment methodologies as follows. The mapping of OWASP risk rating methodology to the security framework NIST is such that some NIST functions are covered (Identify, Protect, Detect). However, considering the Respond and Recover NIST functions, the methodology suggests the general rule to fix first the most severe risks, but it does not offer a detailed approach to do it. Instead, MEHARI and EBIOS are strictly compliant with ISO 27000 family, with direct references to standard ISO 27005. Moreover, EBIOS offers details about security principles (e.g., anticipation, protection, defense, resilience) very similar to NIST functions, meaning that such tool can be applied to NIST framework.

**Attack Graph Model and Risk Estimation.** An attack graph represents possible ways via which a potential attacker can intrude into the target network by exploiting a series of vulnerabilities on various network hosts and gaining certain privileges at each step. A huge body of literature exists about attack graph generation and analysis and such models can be used both on-line (e.g. [2]) and off-line (e.g. [1]) to support security operators in their decision making process. More in details, focusing on the off-line usage, attack graphs can be used to

- determine optimal locations for the firewalls and intrusion detection/prevention systems ([9], [14]),
- compute network security evaluation metrics ([12], [18], [21]),
- perform network security risk analysis ([4], [7], [2]) and
- compute near-optimal proactive defense measures ([22], [7]).

Depending on the way information are represented, we may have two main categories of graphs:

- *State-based representations* [19] depict the whole state of the network for each node in the graph. The main advantage of this representation is its completeness (given the set of vulnerabilities in the network, the Attack Graph is able to represent all the possible attack scenarios). However, this is also its main limitation as it brings to an exponential cost (computation, size of the graph) with respect to the size of the network and the number of vulnerabilities.
- *Logical Attack Graphs* [16] are bipartite graphs representing the dependencies between vulnerabilities and security conditions. In this representation, duplicate paths are eliminated and a more compact representation is provided that scales polynomially with the number of vulnerabilities.

There are a number of attack graph generating tools and techniques, i.e., TVA (Topological Analysis of Network Attack Vulnerability) [15], NETSPA (A Network Security Planning Architecture) [8] and MULVAL (Multihost, multistage, Vulnerability Analysis) [17], that starting from a description of the environment (mainly from topology, routing restrictions and vulnerability scans) are able to generate the resulting attack graph. An alternative approach is to apply correlation techniques on network datasets to create attack graphs.

It is interesting to note that almost all the attack graph models and tools described here support the risk estimation. However, the risk is computed by considering only technical aspects, with governance aspects not taken into account.

## 3 Multi-layer Model

In this section, we briefly describe the multi-layer model defined in the context of the PANACEA Project[1] [11, 5] that will be used as reference to identify relevant layers considered by our methodology to map security controls. This model

[1] https://www.panacearesearch.eu

extends the classical concept of attack graph by including multiple dimensions where vulnerabilities can be identified and exploited to generate attack paths. We can distinguish two main different but interconnected components in the model:

- *Multi-layer Attack Graph*: modeling possible multi-step attacks exploiting the organization's vulnerabilities (both of the assets and the personnel) to reach a target.
- *Business Dependency Model*: modeling the business processes of the organization and their dependencies to other business processes, services and assets.

**Multi-layer Attack Graph.** As pointed out in Section 2, existing attack graph models are relatively easy to build, e.g., by scanning the corporate network for technical vulnerabilities (i.e. CVEs) through automatic tools (e.g. [16, 8, 15]) but unfortunately they do not consider other sources of vulnerabilities like, for example, the human being. Indeed, an organization might install the most advanced technical solutions to protect its assets, still an attacker may circumvent them by exploiting a poorly trained employee which, e.g., leaves its credentials unprotected or is prone to provide sensible information through a social engineering attack or a phishing campaign.

The multi-layer attack graph model [5] is based on three interconnected layers (cfr. Fig. 1): (i) the *human layer* aiming at modeling employees, their relationships and personal vulnerabilities., (ii) the *network layer* modeling the ICT part of the company and (iii) the *access layer* modeling the credentials that humans (represented in the human layer) may use to access devices (residing in the network layer).

The aim of the *human layer* is to model how an attacker can compromise individual identities by exploiting human vulnerabilities of the personnel and their relationships. As shown in Figure 1, the human layer is a subgraph of the multi-layer attack graph. Each node represents a possible level of use that an individual may get on digital identities. Edges are associated to exploitable human vulnerabilities. A directed edge from a node $x_i$ to a node $x_j$, associated

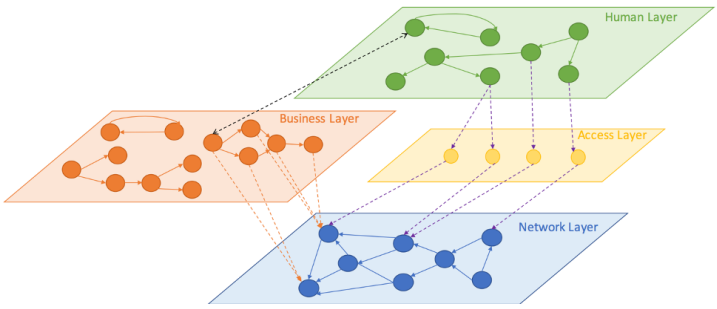


**Fig. 1.** Overview of the Multi-Layer Model.

to a human vulnerability $v$, represents the fact that the human $h_i$ having level of usage $x_i$ on its own digital identity (or an attacker which has gained such usage privilege by exploiting a previous vulnerability) can get level of usage $x_j$ on the digital identity of another individual $h_j$ by exploiting the human vulnerability $v$ on $h_j$ (e.g., a social engineering attack). Notice that the human layer is a multi-graph, where multiple edges between the same pair of nodes are characterized by different human vulnerabilities.

The *network layer* has the aim to represent the ICT network infrastructure serving the organization mission and used by individuals represented in the human layer. Devices expose (technical) vulnerabilities that can be exploited by an attacker to gain access on them. An important concept to model cyber-attacks over IT devices is the *privilege level* that an attacker can gain on such assets. For instance, an intruder might start an attack from the Internet, i.e., with no privilege on the internal IT infrastructure of the organization, or in the case of an insider threat, they might have an initial given privilege on a machine. As a consequence of attack steps that involve vulnerability exploits, they might raise their privilege level on the current machine (privilege escalation) or gain privileges on other machines (remote privilege gain). To model this, each node of the network layer subgraph is associated to a given privilege level on a given device. A directed edge from a node $z_i$ to a node $z_j$, associated to a (technical) vulnerability $v$, represents the fact that an attacker having privilege level $z_i$ on a device $d_i$ can get privilege level $z_j$ on a device $d_j$ by exploiting vulnerability $v$ on $d_j$. Notice that also the network layer is a multi-graph, meaning that there can be multiple edges between each pair of nodes associated to different vulnerabilities.

Individuals are authorized to use assets via various kinds of access credentials, such as badges, tokens, or user accounts, which provide, to various extents, authorization/authentication mechanisms to the network assets. A credential represents the ability of an individual, having a given digital identity, to access a particular asset.

The aim of the *access layer* is to represent such credentials. Thus it is the layer that connects the human layer with the network layer (cfr. Fig. 1). In particular, nodes of the access layer subgraph represent credentials (e.g., a pair $\langle$ username, password $\rangle$, a badge, a biometric key, etc.) or a two-factor authentication that is composed by multiple credentials. Nodes are characterized by a type (e.g., user/password pair, badge, token, etc.) and a level of robustness that can be used in order to weight associated risks when computing attack paths.

The *inter-layer edges* (cfr. Figure 1), are those that connect the subgraphs of the three layers to form the multi-layer attack graph. An edge from a node $x$ in the human layer to a node $y$ in the access layer, represents the fact that an individual with digital identity $x$ has a credential $y$, while an edge between $y$ and a node $z$ in the network layer represents the fact that credential $y$ allows to get privilege $z$ on the associated device.

**Business Dependency Model.** The business dependency model describes the business-level entities and their interdependencies. The business-level entities are partitioned into three disjoint sets: (i) *business processes*, (ii) *services* and (iii)

*assets*. The assets of the business layer are the direct counterparts of the assets of the network layer (e.g., physical/virtual hosts, network equipment, hardware devices, etc.) The services are the direct counterparts of the services and applications running on network layer assets (e.g., software components, applications, etc.). Therefore, there is an interconnection between the network layer of the multi-layer attack graph and these two classes of nodes of the business dependency model (cfr. Figure 1). Conversely, business processes have no direct counterparts in the other layers of the model and represent the business processes of the organization. However, there may be an interconnection between a business node and a node in the human layer, e.g., modeling the fact that a particular member of the organization is fundamental to run a given business process (cfr. Figure 1). The business dependency graph is a directed graph where there is a node for each business-level entity. A directed edge from a business-level entity $be$ to a business-level entity $be'$ models the fact that $be$ depends on $be'$. In other words, in order for $be$ to function properly, $be'$ has to function properly.

## 4 An Assessment Validation Methodology

In this section we present our methodology to support the assessor in the security control coverage validation task and how it is supported by the multi-layer model. The proposed methodology is composed mainly of three sequential steps:

1. performing a control-based security assessment
2. weighted mapping of each security control on the multi-layer model
3. computation of validation factors confirming or suggesting review for the coverage level of the security controls.

In the first step, the security assessor conducts interviews with stakeholders to evaluate the compliance with respect to a set of security controls, collected by applicable technical standards, regulations, best practices. During this activity, the assessor collects, for each relevant security control, a set of evidences that will be used to finally score its coverage and the maturity of implementation. Let us note that, in real scenarios, this activity is prominently conducted manually by the assessor and his/her team, and is heavily based on the expertise of the assessor and on an inherent bias that can influence the final coverage estimation. Most of these activities tend to overweight technical aspects over other considerations (e.g., business processes, human component). Finally, the results of the evaluation can be biased also by the specific set of collected evidences.

The second step of our methodology provides the instrument to exploit the multi-layer model in the security assessment, in the form of a mapping between the set of security controls that is used during the assessment and the multi-layer model. In this way, the assessor can, for each security control, focus the attention on the specific layers one at a time and check if proper evidences have been collected for it, refining in this way the initial activity. This effort puts the focus on the construction of the mapping. Each security control can be mapped along the following dimensions:

- *Lifetime*: allows to link a security control to a specific part of the security lifetime by distinguishing between controls verified against design aspects of the system (e.g., policy design, network design, system configuration, etc.) and those verified on aspects related to execution time. In most cases these two dimensions are disjoint; however, some controls may be related, to different degrees, to both. An example is represented by security control A.7.2.1 of ISO 27001. This control, which regards "management responsibilities", deals with compliance with policies and procedures, and therefore is mainly run-time. However, it implies also that policies and procedures must be well-designed in order to be applied, so it is linked to design time aspects too.
- *Impacted layer weight*: the controls are mapped to each layer of the multi-layer model (i.e., network, human, access, business) with a certain weight specifying how much they deal with human, access, network and business layers.
- *Management level*: models the impact of the security control over the business operational structure or the security organizational structure. It follows the definition by Von Solms [20] about the Information Security Governance model.

| Control ID | Control name | H | N | A | B | Compl. | Operat. | Design time | Run time |
|---|---|---|---|---|---|---|---|---|---|
| A.5.1.1 | Information security policies | 1 | 0 | 0 | 3 | 4 | 0 | 2 | 0 |
| A.6.2.1 | Mobile device policy | 1 | 3 | 2 | 0 | 4 | 3 | 2 | 0 |
| A.7.2.3 | Disciplinary process | 4 | 0 | 0 | 2 | 4 | 0 | 2 | 0 |
| A.8.1.2 | Ownership of assets | NaN | NaN | NaN | NaN | 4 | 0 | 2 | 0 |
| A.9.2.3 | Management of privileged access rights | 0 | 0 | 4 | 0 | 0 | 4 | 0 | 2 |
| A.10.1.1 | Policy on the use of cryptographic controls | 0 | 4 | 0 | 2 | 4 | 2 | 2 | 2 |
| A.11.2.1 | Equipment siting and protection | 0 | 4 | 4 | 0 | 0 | 4 | 0 | 2 |
| A.12.4.1 | Event logging | 0 | 0 | 0 | 3 | 0 | 4 | 0 | 2 |
| A.13.2.4 | Confidentiality or non disclosure agreements | 4 | 4 | 4 | 4 | 3 | 4 | 2 | 0 |
| A.14.2.2 | System change control procedures | 0 | 0 | 0 | 4 | 4 | 0 | 2 | 1 |
| A.15.1.3 | Information and communication technology supply chain | NaN | NaN | NaN | NaN | 4 | 0 | 2 | 0 |
| A.16.1.7 | Collection of evidence | 4 | 4 | 4 | 4 | 4 | 0 | 2 | 2 |
| A.17.2.1 | Availability of information processing facilities | 0 | 0 | 0 | 4 | 0 | 4 | 0 | 2 |
| A.18.2.3 | Technical compliance review | 0 | 0 | 0 | 4 | 4 | 0 | 2 | 0 |

**Table 1.** A sample of mapping for a subset of controls ISO 27001:2013 (H - Human, N - Network, A - Access, B - Business).

Concerning layers and management levels we map each security control with an integer weight between 0 (control totally not related to the layer/management

level) and 4 (control totally related to the layer/management level). With such scale we can map the many different aspects that are inside layers and management levels. Instead, for the lifetime we introduced an ordinal scale with values LOW-MEDIUM-HIGH (translated in integer scale with values 0-1-2) defining how much the security control is related to the design and/or run-time.

We mapped all the security controls of ISO 27001 to our multi-layer model. Table 1 reports the mapping for a subset of these controls. We provide the full mapping as a supplemental material[2]. Let us note that some controls (e.g. A.9.2.3, A.14.2.2, A.17.2.1, A.18.2.3) are completely mapped in only one of the layers of the multi-layer model while other controls such as A.6.2.1 and A.13.2.4 insist on multiple layers due to their generality.

In addition, there exist few controls for which a meaningful mapping cannot be established. Indeed, for controls A.8.1.2 and A.15.1.3 a "NaN" value is set, meaning that they cannot find a suitable mapping to the multi-layer model.

In the third and final step, we are going to review the assessment, based on the obtained mapping, and we are going to compute the validation factor as described in the following section. The steps of the methodology are then repeated in sequence until all validation factors confirm the coverage levels.

## 5 Computing Coverage Validation Factors

Validation factors play a central role in the methodology described in this paper because they support the understanding of how, in what way and to what degree we can better fit the security controls coverage with respect to assessment results. Once the validation factor has been figured out, then the cybersecurity assessor has the information about how reliable the assessed coverage level is for each security control. We can generally identify three main cases, identified by two reliability thresholds $T1$ and $T2$:

- $validation\ factor > T2$: the coverage level of the security control resulting from the assessment can be considered reliable enough (*Confirmed - OK*);
- $T1 < validation\ factor < T2$: the results of the assessment are partially reliable, but a supplement of analysis could be necessary. At this point the mapping described in the previous section suggests the assessor which are the layers in which applying such procedures (*Recommended Review - RR*);
- $validation\ factor < T1$: the procedures considered in the assessment are not enough for covering the security control, therefore a revision of the security control and evidences is critically necessary (*Absolutely Review - AR*).

Reasonable values for such thresholds are $T1 = 0.3$ and $T2 = 0.6$, but other assignments are possible depending on the context. The computation of the validation factor is based on the assessment information collected in step 1 and the mapping produced in step 2. In the following we report how each of these pieces of information is interpreted in the computation of the validation factor:

[2] The complete mapping can be found at the following link: https://drive.google.com/file/d/1PHEbU38H4NtyzLiqHrZ-YczN-4NhBe5z/view

- Coverage ($cv$): is the coverage level of the security control (i.e., the extent to which the provisions of the control are implemented). The coverage level is derived directly from the assessment and can assume the values C (Fully Covered), PC (Partially Covered) or NC (Not covered).
- Lifetime ($lt$): whether the control can be implemented at run-time or design-time; the rationale for this parameter is that if a control is implementable at run-time, it is generally easier to remedy and study. The parameter $lt$ should be put at 0 if the control is totally design-time, 2 if it is totally run-time and 1 in all cases in which the distinction between the two is not evident.
- Management Level ($ml$): represents the extent to which the control is related to security operational management rather than compliance management. We account for this information as, generally, operational actions are more explicitly defined than compliance ones, thus, the coverage level of a security operational level control has more significance with respect to that of a compliance level control. Therefore, the confidence in a coverage level will be higher if the control concerns operational actions rather than compliance ones. We assign to this parameter a value between 0 and 4 (0: totally not operational, 4: totally operational).
- $g_M$: is the maximum gap between the level of mapping on the multi-layer model. We assign to each security control a level of mapping to the layers (human, access, network and business) with an integer scale $[0, 4]$ (the same we used for the $ml$ value). The maximum gap between all the layers for a specific control models the uncertainty of the mapping of the control to the multi-layer model: the greater is the gap, the more the layers are disjoint, meaning that the security control is more defined.
- $p_{RC}$: is the reliability-coverage factor, interpreted as the degree of reliability to which a security control coverage contributes to the analysis of the security coverage of the assessment. We found three degrees of reliability:
  - HIGH: the security control is implemented at run-time and it is mapped in the multi-layer model;
  - MEDIUM: the security control is implemented at design-time and it is mapped in the multi-layer model;
  - LOW: the security control is not mappable into the multi-layer model.

  With such information and taking into account the coverage level, we produced the following table containing the values of $p_{RC}$ parameter for every case:

| Reliability Coverage | HIGH | MEDIUM | LOW |
|---|---|---|---|
| C | 4 | 3 | 1 |
| PC | 3 | 2 | 1 |
| NC | 1 | 1 | 1 |

Considering the parameters described above the validation factor can be computed for each security control mapped in the multi-layer model with the following formula:

$$validation factor = \frac{lt + ml + g_M + p_{RC}}{lt_{max} + ml_{max} + g_{M,max} + p_{RC,max}}$$

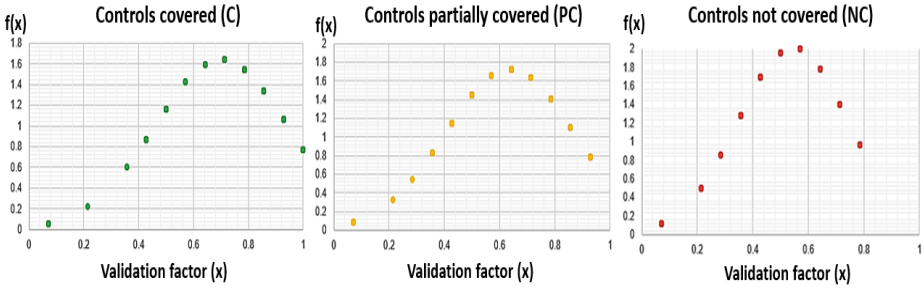


**Fig. 2.** Validation factor distribution in different cases of controls coverage.

Let us remark that if the lifetime is run-time instead of design-time and the management level is operational instead of compliance, then the validation factor increases due to the fact that the coverage of the controls is more precise. The more the gap between layers increases, the more the control is well positioned into the multi-layer model and the aspects to take under observation are more precise. Finally, the more the mapping is reliable, the more the formula is precise and the validation factor is high. The validation factor is normalized in the interval $[0, 1]$ by having in the denominator the sum of maximum values for each parameter ( $lt_{max} = 2$ and $ml_{max} = g_{M,max} = p_{RC,max} = 4$). Once calculated, the validation factor represents the confidence of how much the coverage level of the security control fits. As an example, if a security control coverage is assessed as "completely covered" (C), and the computed validation factor is 0.4, then it means that the security controls should be reviewed ($RR$). The additional inspection of the relative mapping to the multi-layer model suggests which layers and associated security elements should be reviewed.

### 5.1 Measurement accuracy

In this section we illustrate the accuracy and reliability of the validation factor. For this purpose, we evaluate all possible cases of assessment for each security control of ISO 27001:2013, and we use such information to analyze the statistical distribution of data. The results are reported in Figure 2: the behavior of the three cases (C/PC/NC) is the normal distribution evaluated through mean ($\mu$) and standard deviation ($\sigma$) in each case of coverage level (all C, all PC and all NC), and then applying the normal distribution function:

$$f(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-} \frac{(x-\mu)^2}{2\sigma^2}$$

The mean and standard deviation in the three cases are respectively:

- $\mu = 0.701754$ and $\sigma = 0.243285$ if controls are all covered;
- $\mu = 0.637845$ and $\sigma = 0.230917$ if controls are all partially covered;
- $\mu = 0.545739$ and $\sigma = 0.198709$ is controls are not covered.

We use such analysis for the evaluation of the accuracy and reliability of the validation factor because it expresses how much the reliability of coverage level can vary in the different cases, considering the average value. What the data shows is that the more the security controls are covered, the more the dispersion increases; in other words, the more a security control has a high coverage level, the more the variability of the validation factor increases (but still remains limited). This is due to the fact that when a security control is assessed near the full coverage, it means that many aspects of the control are assessed, and on each of them, depending on the mapping, the multi-layer model can express their validity or not. Instead, when a security control is assessed as not covered (or near low values for coverage), it means that many aspects of the security control has been found as not covered, or not considered, and the resulting validation factor from the multi-layer model is more clear-cut, resulting in less variability.

The mean expresses the expectation of the impact over different layers: if the validation factor has a very low value (e.g., lower than 0.3), it means that the control under analysis impacts on many different layers with different weights and thus the suggestion for the assessor is to verify that all the layers impacted have been considered in the determination of the coverage level. The verification can be done by using the mapping in order to identify layers that may be included in the evaluation. Contrarily, if the validation factor has a high value (i.e., between 0.7 and 1), it means that the control is clearly related to a specific layer and thus the expectation is that the assessor already verified the relevant elements in its analysis.

## 6 Usage Scenario

This section presents the application of the assessment validation methodology to a realistic scenario. The context in which it is applied is related to the healthcare domain, that is showing an increased number of cyber-attacks in the last two years (even exacerbated by COVID-19) and an attention to cybersecurity that only lately is starting to become more deeply considered. As a resultant, this domain represents a good scenario in which the multi-layer model can be instantiated and provides benefits.

We applied our methodology considering as reference a unit of a hospital. The scenario of application is based on the security controls defined in ISO 27001, with an assessment based on interviews and evidences collection limited only to the technological layer. We denote this procedure as "classic assessment". While in principle this could be seen as a limitation, in practice many cybersecurity assessments are based on these assumptions. We take this situation as a reference for security assessment initiatives. ISO 27001 presents 18 categories for 114 security controls. The usage scenario covers all the security control categories, while not considering all the security controls for each category. Rationale behind this choice is that some of the controls are not applicable in the scope of this usage scenario and others are still under assessment or modeling. At the

| Control ID | Control name | H | N | A | B | Compl. | Operat. | Design time | Run time | Cov. | Valid. factor | Valid. result |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| A.5.1.2 | Review of the policies for information security | 1 | 0 | 0 | 3 | 4 | 0 | 2 | 0 | PC | 0.425 | RR |
| A.6.1.1 | Information security roles and responsibilities | 4 | 0 | 0 | 1 | 4 | 0 | 2 | 0 | PC | 0.429 | RR |
| A.7.2.1 | Management responsibilities | 3 | 0 | 0 | 3 | 3 | 4 | 1 | 2 | NC | 0.714 | OK |
| A.8.1.3 | Acceptable use of assets | 3 | 3 | 3 | 4 | 1 | 4 | 2 | 0 | C | 0.571 | RR |
| A.9.2.3 | Management of privileged access rights | 0 | 0 | 4 | 0 | 4 | 0 | 0 | 2 | C | 0.714 | OK |
| A.9.4.3 | Password management system | 2 | 0 | 4 | 0 | 0 | 4 | 0 | 2 | PC | 0.929 | OK |
| A.10.1.1 | Policy on the use of cryptographic controls | 0 | 4 | 0 | 2 | 4 | 2 | 2 | 2 | C | 0.857 | OK |
| A.11.1.2 | Physical entry controls | 3 | 2 | 4 | 1 | 0 | 4 | 1 | 2 | C | 0.928 | OK |
| A.12.4.1 | Event logging | 0 | 0 | 0 | 3 | 0 | 4 | 0 | 2 | PC | 0.857 | OK |
| A.13.2.4 | Confidentiality or non disclosure agreements | 4 | 4 | 4 | 4 | 3 | 4 | 2 | 0 | NC | 0.357 | RR |
| A.14.3.1 | Protection of test data | NaN | NaN | NaN | NaN | 0 | 4 | 2 | 0 | NC | 0.357 | RR |
| A.15.1.3 | Information and communication technology supply chain | NaN | NaN | NaN | NaN | 4 | 0 | 2 | 0 | PC | 0.07 | AR |
| A.16.1.7 | Collection of evidence | 4 | 4 | 4 | 4 | 4 | 0 | 2 | 2 | NC | 0.214 | AR |
| A.17.2.1 | Availability of information processing facilities | 0 | 0 | 0 | 4 | 0 | 4 | 0 | 2 | PC | 0.928 | OK |
| A.18.1.4 | Privacy and protection of personally identifiable information | 0 | 0 | 0 | 4 | 4 | 3 | 2 | 0 | C | 0.714 | OK |

**Table 2.** A sample of mapping with assessment for a subset of controls and the related validation (H - Human, N - Network, A - Access, B - Business).

same time all the security control categories are covered with at least 1 security control bringing the total to 15 security controls.

Table 2 presents the results of this activity. The first two columns (Control ID, Control Name) reports the used security controls, where the following eight columns report the mapping to the multi-layer model valued for each of the security controls on the usage scenario. Finally:

1. "Coverage" column reports the coverage level related to security controls coming from the classic assessment;
2. "Validation Factor" column reports for each security control the resulting validation factor computed from the mapping;
3. "Validation Result" column reports the validity of each security control, coming from the interpretation of the validation factor, in terms of three possible results: assessment reliable (OK), assessment partially reliable (RR: Recommended Review) and assessment not reliable (AR: Absolute Review).

Looking at the results, the classic assessment presents five security controls fully covered, six partially covered and the remaining four not covered. We observe that security controls A.5.1.2 (Security policies review) and A.6.1.1 (security roles and responsibilities) represent similar situations in which the original assessment set a partial coverage, and that both controls are mapped to two layers

with one layer prominent with respect to the other (Business layer for the first and Human layer for the second). This brings to a moderate validation factor that potentially could lower the coverage of the relative security control: for this reason it is recommended to review the assessment looking at the relative mapped layers. Focusing on security control A.8.1.3 (Acceptable use of assets), the classic assessment produced a full coverage for this control (C). The multi-layer model identifies Human, Access and Business layers as important, on top of the technical layer (Network). This mapping brings to a moderate validation factor, meaning that the assessment should be verified for elements concerning Human layer (e.g., misuse of organizations devices) or Access layer (e.g., verification of permissions policies). For this reason the final outcome is a recommended revision, that could probably lead to a PC coverage level. Security control A.9.2.3 (Privileged access rights) is assessed as fully covered (C). The multi-layer model identifies the Access layer as the only one interested by this control, with some contributions coming from lifetime (Design time) and management level. Overall the validation factor is above the $T2$ threshold and it confirms the coverage level as fully covered. Finally, we observe that the security controls A.14.3.1 (security of test data) and A.15.1.3 (ICT supply chain) have no mapping with the multi-layer model. A.14.3.1 is mapped on the security domain while A.15.1.3 is mapped on the operational domain, and both are labeled as design-time activities. The difference in management level provides a very low validation factor for the latter and a moderate one for the former, that weighted with high contribution for lifetime (Design time for both) brings to recommended review for the first and absolute review for the second security control. Overall, our methodology confirms the coverage level for 8 out of 15 (54%) security controls while asks review for the other half of them, with moderate revisions for 5 controls and strong revisions for other 2.

Concluding, the usage scenario showed where the proposed methodology can help in validating the coverage level of security controls with respect to classic security assessment procedures, supporting a more fit estimation for coverage level, highlighting security controls that needs a supplement of analysis and identifying for them the layers on which additional information should be collected.

## 7 Conclusion

This paper explored the idea of supporting the security assessment processes through a methodology that allows to better estimate the coverage with respect to security controls. The proposed methodology is based on a multi-layer model that captures the cyber-exposure from multiple perspectives, like human component and business inter-relations. It allows the derivation of validation factors that help in better estimating the coverage level of a security control and eventually identifying parts of the original security assessment that needs a supplement of analysis. The identification of additional areas on which to conduct the supplemental analysis can be directly inferred by the layers of the multi-layer model. The proposed usage scenario, even if of limited scope, successfully demonstrated

the validation capabilities of the proposed methodology and the added value it provides with respect to classic security assessment methods, allowing to model the reliability of a security assessment considering (but not limited to) information security governance aspects.

Limitations exist in the proposed approach that we plan to resolve in future works. The first limitation concerns the hypothesis of a purely technological assessment conducted with classic methods. We plan to integrate the methodology in order to be able to assign each collected evidence to one or more of the layers of the multi-layer model, making the methodology able to support mixed initiative assessments. The second limitation concerns the scope of the usage scenario. We are currently working on broadening it, eventually obtaining a controlled environment that better suits a more robust validation of the obtained results. This aspect is linked to future enhancements of this methodology, to fully exploit the multi-layer model in supporting a complete risk assessment process. This future evolution will allow to consider not only the coverage level of a set of security controls, but also a formal definition of cyber-risk. In this scenario we are developing a "risk discount factor" supported by data computed through the multi-layer model (e.g., attack paths, human vulnerabilities) that can positively or negatively affect the risk estimation instead of only the coverage level. Finally we are working on generalizing the mapping to a broader set of security assessment frameworks (e.g. NIST Cybersecurity framework, CIS CSC).

## References

1. Angelini, M., Blasilli, G., Catarci, T., Lenti, S., Santucci, G.: Vulnus: Visual vulnerability analysis for network security. IEEE Transactions on Visualization and Computer Graphics **25**(1), 183–192 (2019)
2. Angelini, M., Bonomi, S., Borzi, E., Pozzo, A.D., Lenti, S., Santucci, G.: An attack graph-based on-line multi-step attack detector. In: Proceedings of the 19th International Conference on Distributed Computing and Networking. ICDCN 18, Association for Computing Machinery, New York, NY, USA (2018), https://doi.org/10.1145/3154273.3154311
3. ANSSI: EBIOS Risk Manager. https://www.ssi.gouv.fr/en/guide/ebios-risk-manager-the-method/, online; accessed 12 July 2020
4. Beckers, K., Heisel, M., Krautsevich, L., Martinelli, F., Meis, R., Yautsiukhin, A.: Determining the probability of smart grid attacks by combining attack tree and attack graph analysis. In: Cuellar, J. (ed.) Smart Grid Security. pp. 30–47. Springer International Publishing, Cham (2014)
5. Bonomi, S., Ciccotelli, C., Laurenza, G., Lenti, S., Palleschi, A., Santucci, G., Sorella, M., Tanasache, F.D.: Understanding human impact on cyber security trough multilayer attack graphs. Tech. rep., Department of Computer, Control and Management Engineering, Sapienza University of Rome. (2020), https://bonomi.diag.uniroma1.it/research/publications
6. CLUSIF: MEHARI (MEthod for Harmonized Analysis of RIsk). http://meharipedia.x10host.com/wp/, online; accessed 12 July 2020
7. Gonzalez Granadillo, G., Dubus, S., Motzek, A., García, J., Alvarez, E., Merialdo, M., Papillon, S., Debar, H.: Dynamic risk management response system

to handle cyber threats. Future Gener. Comput. Syst. **83**, 535–552 (2018), https://doi.org/10.1016/j.future.2017.05.043
8. Ingols, K., Lippmann, R., Piwowarski, K.: Practical attack graph generation for network defense. In: Proceedings of the 22nd Annual Computer Security Applications Conference. p. 121130. ACSAC 06, IEEE Computer Society, USA (2006), https://doi.org/10.1109/ACSAC.2006.39
9. Jajodia, S., Noel, S.: Topological Vulnerability Analysis, pp. 139–154. Springer US, Boston, MA (2010)
10. Jeff Williams: OWASP Risk Rating Methodology. https://owasp.org/www-community/OWASP_Risk_Rating_Methodology, online; accessed 12 July 2020
11. L. Coventry and D. Branley-Bell and E. Sillence and S. Bonomi and C. Ciccotelli and S. Lenti and A. Palleschi and L. Querzoni and G. Santucci and M. Sorella, and F. D. Tanasache: D2.2 - Human Factors, Threat Models Analysis and Risk Quantification. PANACEA Project https://www.panacearesearch.eu
12. LeMay, E., Ford, M.D., Keefe, K., Sanders, W.H., Muehrcke, C.: Model-based security metrics using adversary view security evaluation (advise). In: 2011 Eighth International Conference on Quantitative Evaluation of SysTems. pp. 191–200 (2011)
13. Nist, Aroms, E.: NIST SP 800-100 Information Security Handbook: A Guide for Managers. CreateSpace, Scotts Valley, CA (2012)
14. Noel, S., Elder, M., Jajodia, S., Kalapa, P., O'Hare, S., Prole, K.: Advances in topological vulnerability analysis. In: 2009 Cybersecurity Applications Technology Conference for Homeland Security. pp. 124–129 (2009)
15. Noel, S., Wang, L., Singhal, A., Jajodia, S.: Measuring security risk of networks using attack graphs. IJNGC **1**(1) (2010)
16. Ou, X., Boyer, W.F., McQueen, M.A.: A scalable approach to attack graph generation. In: Proceedings of the 13th ACM Conference on Computer and Communications Security. p. 336345. CCS 06, Association for Computing Machinery, New York, NY, USA (2006), https://doi.org/10.1145/1180405.1180446
17. Ou, X., Govindavajhala, S., Appel, A.W.: Mulval: A logic-based network security analyzer. In: Proceedings of the 14th Conference on USENIX Security Symposium - Volume 14. p. 8. SSYM05, USENIX Association, USA (2005)
18. Pamula, J., Jajodia, S., Ammann, P., Swarup, V.: A weakest-adversary security metric for network configuration security analysis. In: Proceedings of the 2nd ACM Workshop on Quality of Protection. p. 3138. QoP 06, Association for Computing Machinery, New York, NY, USA (2006), https://doi.org/10.1145/1179494.1179502
19. Sheyner, O., Wing, J.: Tools for generating and analyzing attack graphs. In: de Boer, F.S., Bonsangue, M.M., Graf, S., de Roever, W.P. (eds.) Formal Methods for Components and Objects. pp. 344–371. Springer Berlin Heidelberg, Berlin, Heidelberg (2004)
20. Solms, S.V., Solms, R.V.: Information security governance. Springer Science and Business Media, Springer (2009). https://doi.org/10.1007/978-0-387-79984-1
21. Wang, L., Jajodia, S., Singhal, A., Cheng, P., Noel, S.: k-zero day safety: A network security metric for measuring the risk of unknown vulnerabilities. IEEE Transactions on Dependable and Secure Computing **11**(1), 30–44 (2014)
22. Wang, L., Albanese, M., Jajodia, S.: Network Hardening - An Automated Approach to Improving Network Security. Springer Briefs in Computer Science, Springer (2014), https://doi.org/10.1007/978-3-319-04612-9