LUISS

ROSI as a Tree: A Modular Framework for Context-Aware Return On Security Investment

LAUREA MAGISTRALE IN DATA SCIENCE AND MANAGEMENT
MASTER'S THESIS IN DATA PRIVACY AND SECURITY

Supervisor Paolo Spagnoletti
Co-Supervisor Emilio Coppa
Candidate Elena Tomasella (781321)

Academic Year 2024/2025

With digital transformation now entrenched in both public and private sectors, cybersecurity has emerged as a paramount concern and a key investments priority (Gordon., 2003; Johnson et al. 2024). Currently, while some corporations – especially in the insurance sector – have developed advanced tools to assess the return on security investments and optimize resource allocation, Public Administrations (PAs) lack a structured framework to guide future security investments efficiently. This study seeks to bridge this gap by offering a framework for the modular implementation of ROSI to serve as a tailored tool for individual organizations. Considering ROSI as the trunk of a tree, the structure, development, and potential future evolution of this tree represent the core of this framework.

We follow a Design Science Research (DSR) approach in collaboration with experts from the Italian National Cybersecurity Agency (ACN) to leverage data and integrate knowledge collected through a workshop with key informants. Our artefact consists of a refined and extended Return on Security Investment (ROSI) metric, implemented in Python and enriched with key domain knowledge obtained through a workshop employing the Analytic Hierarchy Process (AHP) technique. We explore complementary methods and incorporate granular data from official reports and secondary data. Using this artefact, PAs – as well as other organizations – can, for example, rank their investment in security based on the assessments generated by our artifact, thereby facilitating the identification of optimal resource allocation strategies and guiding proactive improvements in cyber posture. This aligns with the main objective of our research: by aligning security measures with each organization's unique profile, the proposed framework seeks to enhance resilience, optimize investment decisions, and establish a robust model for continuous cybersecurity improvement.

Introduction

Iniziavo questa Laurea Magistrale con due principali desideri: accorpare alla carriera accademica la creatività, e applicare il pensiero matematico a problemi complessi. Concludo questa Laurea Magistrale con una tesi su un modello generalizzabile e modulare, che cerco di spiegare con gli alberi di Bruno Munari.

Before addressing any technical prerequisites or providing a detailed explanation of the framework, I would like to introduce the essence of this work. To ease intuition and enhance the accessibility of the thesis, a metaphor is used. The only necessary concept to understand about the Return on Security Investment (ROSI) at this stage is that it is a value derived from a formula, assessing the extent to which a security expenditure is not merely a cost for an organization, but rather an investment. Now, please consider ROSI as the trunk of a tree. The structure, development, and potential future evolution of this tree represent the core of our framework. With this in mind, two main areas of discussion emerge. First, the roots of this tree must be analyzed, ensuring their solidity and enabling the right conditions for them to grow further, where possible. Then, there is the canopy to be developed on the top of that trunk. This encompasses the various applications of ROSI, which can be grouped by type and may give rise to a cascade of new branches. The idea of conceptualizing the framework in this way - and of guiding the research through the modular structure typical of trees - emerged when it became clear that the project had both significant extensibility and generalization potential, yet limited resources to implement an advanced version of the artefact. It also became evident that each root and each branch carried its own degree of complexity. It would not have been realistic to pursue an approach that did not allow for the independent development and updating of each component while keeping the rest of the structure intact. This is why the thesis presents chapters dedicated to the roots, for example. Within each, different levels of depth and quality are proposed, depending on the resources available to the organization adopting the ROSI framework. In practice, one may choose the simplest form of each component - whether in the canopy or the roots - and still obtain a functional tree, accepting the simplified assumptions while laying the foundation for a usable prototype. This concept may also resonate with musicians, particularly drummers: there are always core elements - such as the kick drum, hi-hat, snare, floor tom, and a tom-tom. These can be upgraded over time, or new pieces added, and existing ones replaced or enhanced - such as fitting higher-quality drumheads. This is exactly the kind of evolutionary approach an organization might take once the prototype is in place. I'd like to conclude this introduction with a small anecdote: during the course of this thesis research, I spent a few days in Barcelona and I was reminded that Gaudí designed the Sagrada Família by imitating the structural logic of trees, studying their self-supporting architecture. At first, this simply made me smile - but then it gave me confidence. This metaphor might seem imaginative, but it is not entirely original. I am not the first to recognize that nature offers elegant solutions to complex problems. Indeed, high-profile studies now suggest that these are often optimal solutions. Since the tree structure has proven to be an

evolutionary win-win and it is an architecture inherently adaptive to the biome in which it exists, I am confident that this ROSI model can fill a gap by offering a tool that is adaptable to the specific organizational context in which it is deployed.

The work presented in this thesis, as will be explained later, has been guided by the typical dynamics of Action Design Research. This implies that the very nature of the research is grounded in the constructive exchange of knowledge and ideas. For this reason, I am deeply grateful to the people that took part to this Research, each contributing in their own way to the outcomes of this thesis. Every time I exited the metro at Castro Pretorio and walked to the offices of the National Cybersecurity Agency (ACN), I returned to a simple awareness: I am happy with where I am going today, and I truly enjoy what I am doing. That is not a small thing, and it reflects a balance between the research project itself, and the people involved. I sincerely thank my supervisor, Paolo Spagnoletti, for his insightful guidance, which shaped the design and development of this thesis. He gave me the extraordinary opportunity to collaborate with the National Cybersecurity Agency, to which I am especially grateful to Riccardo Zecchinelli e Luigi Addorisio De Feo. Although the collaboration greatly enriched the development of this thesis, I take full responsibility for any misinterpretation or error that may appear in its content. I also wish to thank the authors of the scientific papers I reached out to for clarification, who generously responded to my requests. Finally, I extend my gratitude to Simone Guarino for his valuable contribution to the workshop, which played a key role in shaping the final outcomes and ensuring the accountability of this research.

Chapter 1 – Context

Overview

The National Recovery and Resilience Plan (PNRR) is Italy's strategy for implementing the European Union's Next Generation EU (NGEU) program, designed to support recovery from the economic and social impacts of the COVID-19 pandemic. The PNRR allocates over €190 billion in grants and loans to support investments and reforms across key strategic sectors, including cybersecurity. In this context, the plan aims to strengthen the long-term resilience of the country. The general funding structure for public administrations is exemplified through the case of the National Cybersecurity Agency (ACN). The Agency issues official Calls for Proposals ("Avvisi"), which are public notices inviting eligible institutions to apply for funding. These calls define objectives, eligibility criteria, and submission procedures for specific initiatives. This research originated from the need to evaluate the outcomes of certain investments funded through these calls. Specifically, it focuses on the Return On Security Investment (ROSI) and its strategic relevance, particularly within the context of Italian Public Administrations (PAs), though with broader applicability. The study recognizes the potential of ROSI as a valuable tool to enhance cybersecurity posture by providing a structured approach to assess the effectiveness of security-related investments.

This study is particularly relevant given the disparities in digitalization and cybersecurity advancements across the institutions. With digital transformation now entrenched in both public and private sectors, cybersecurity has emerged as a paramount concern and a key investments priority ((Gordon L. A., 2003); (Jonhson, Maurer, Torres, Guerra, & Mohit, 2024)). Digitalization is reshaping the services provided by nations both horizontally, affecting all sectors, and vertically, driving systematic changes at various levels. While digital transformation progresses at different speeds across sectors, its impact is widespread. However, the increasing reliance on digital systems also amplifies vulnerabilities, both in number and complexity. This first chapter provides an overview of the latest trends in terms of cyber-attacks, along with a referenced analysis of the current cybersecurity landscape, in Italy and in comparison, in the broader European context. It begins by focusing on the distribution of threats among countries and their trends. Subsequently, data is presented on the counterpart as well, examining the various advancement in digitalization, which allow us to assess the extent and investigate the reasons to which Italy has become increasingly more exposed to threats in recent years. At times, these two factors align, while at other times, they do not. Notably, such misalignments can decelerate or even block the initial progress of digital services. Besides the direct effect of the attacks, it is also mindful to consider that cybersecurity is accounted for affecting the success of some initiatives, as exemplifies by the case of e-governments. Additionally, certain issues, such as public trust, have a more significant impact on PAs compared to small and medium-sized enterprises

(SMEs) rather than massive private corporations. Therefore, dedicate attention is given to outlining the evidence that the literature has gathered on this matter. In conclusion, some potential extensions for this study are mentioned, accounting it for being both relevant and necessary. Before delving into the contextual information that explains the usefulness of this Action Design Research, a final clarification must be made. In these initial contextual provisions, the information is focused on the public sector because that is where the research originated, thanks to the fact that it was the National Cybersecurity Agency to make a focus on the topic. However, this should not be interpreted as limiting the scope of the research to a single specific environment, since it was evident from the very beginning that the project had broader potential.

First side: attacks

All data on cyberattacks presented in the referred reports is based on recorded successful incidents that have been publicly disclosed. This has two key implications that must be considered when interpreting the related statistics. First, these rankings rely on publicly known cyberattacks, which poses a limitation in sectoral assessments. Some companies and organizations choose not to disclose comprehensive records due to confidentiality concerns, aiming to maintain public trust and mitigate potential reputational damage. Second, advancements in regulatory frameworks are facilitating both the detection process and its granularity. A presentation of the key statistics from the latest report by Clusit (Clusit, 2025) follows. Over the past five years, the global number of attacks has increased dramatically, from an average of 156 events per day in 2020 to an average of 295 events per day in 2024. While this trend has been evident in recent years, the 2025 report on cyberattacks in 2024 further underscores Italy's status as a preferred target, with a 15% increase in incidents compared to 2023. Despite accounting for only 0.7% of the global population and 1.8% of the world's GDP, Italy experienced 10% of recorded cyberattacks worldwide in 2024. By comparison, France accounted for 4%, while Germany and the United Kingdom each accounted for 3%. Disruptions to services caused by cyberattacks can have severe consequences, ranging from delays in healthcare and emergency response systems to the unavailability of essential administrative functions. In 2024, the Governance, Military, and Law Enforcement (Gov/Mil/LE) sector ranked as the second most targeted by cyberattacks, while the healthcare sector accounted for 13.3% of all incidents. However, Gov/Mil/LE has experienced a 45% increase in cyberattacks compared to 2023, whereas the healthcare sector saw a comparatively lower rise of 18.9%. As illustrated in Figure 1, the sharp increase in cyber incidents within the governance sector is particularly evident on a global scale. Apart from the diverse impact that successful cyberattacks have when targeting PAs as discussed in the dedicated section, another factor distinguishes the private from the public sector. As a demonstration of the different intentions of the attackes and as a further proof that the consequences of the attacks vary significantly, the kind of the attacks are also not similar, al least in their distribution.

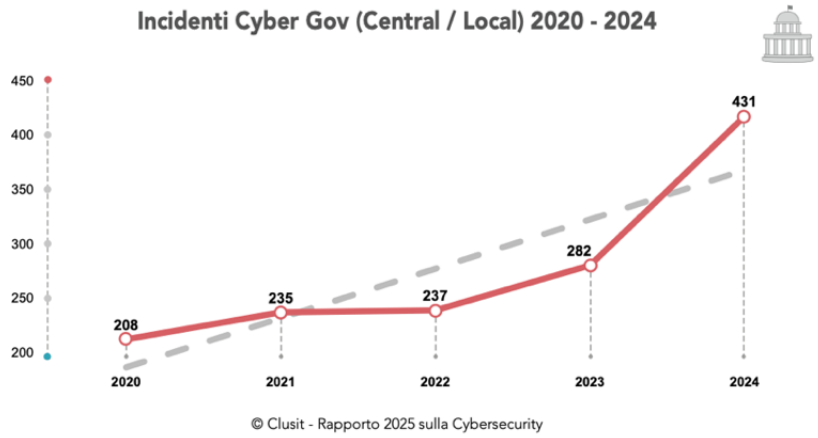

Figure 1 - Cyber Incidents in the Governance Sector

Looking deeper at the kind of attacks that are most occurring in Italy, in 2024 there has been two main kinds which are cybercrime (78%) and hacktivism (22%). The percentage of attacks to PAs is decreasing in respect to the total number, but it is also significant the strong increase in hacktivism as a kind of attack for PAs - from 6 attacks in 2022 up to 35 attacks in 2023. The histogram in Figure 2 is taken from (Rapporto Clusit - Italia e PA, 2024), and it represents the attacks to Italian PAs between 2018 and 2023, grouped by kind of attackers.

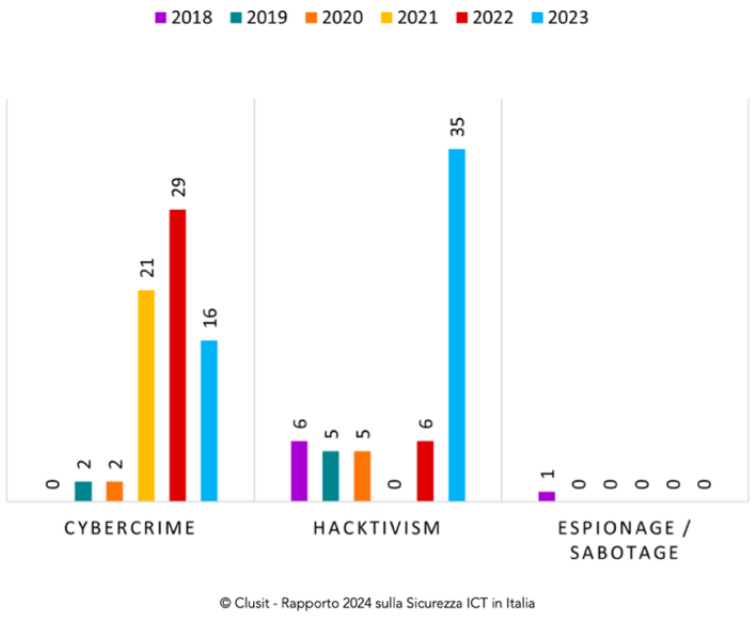

Figure 2 - Attacks to Italian PAs per kind of attacker

To stress the mentioned increase in hacktivism, Table 1 is left. In 2023, a turning point has been registered for the kind of attacks to PAs. Cybercrime increased until 2021, but since then there was a takeoff for hacktivism that had a decreased in cybercrime as a side as well.

	Cybercrime	Hacktivism	Espionage/Sabotage
2018	0%	85,7%	14,3%
2019	28,6%	71,4%	0%
2020	28,6%	71,4%	0%
2021	100%	0%	0%
2022	82,9%	17,1%	0%
2023	31,4%	68,6%	0%

Table 1

From the report Clusit dedicated to Italian PAs in 2024, further information can be integrated to establish a complete understanding of the trend, regarding the nature of the attackers and the severity on the cyberattacks, to PAs. Regarding the nature of attackers, the data highlights the dominance of cybercriminal organizations, which have significantly increased their activities since 2021. The number of attacks attributed to organized crime grew steadily, although a slight decrease was observed in 2023. Meanwhile, politically and ideologically motivated attackers became more prominent in the past year, likely influenced by geopolitical conflicts such as the Russia-Ukraine and Israel-Palestine crises. These attackers primarily engaged in demonstrative cyber operations against government institutions, further contributing to the observed shift in attack severity. Finally, the severity of cyberattacks targeting the Italian PAs showed a concerning trend. Most attacks fell into the "high" and "critical" severity categories, with both experiencing growth over time. However, in 2023, there was an apparent shift, with a decline in critical severity attacks and a corresponding increase in high and medium-severity incidents. This trend aligns with the rising number of ideologically motivated Denial of Service (DDoS) attacks, which, while disruptive, tend to have less severe long-term consequences compared to ransomware or data breaches. From the updated report that has been released in March 2025, it is stood that attacks on Public Administration have seen a significant increase, rising from 560 to 1,430 (+155%), globally. This phenomenon is again linked to growing geopolitical instability and the rise in hacktivist activities, as well as the expansion of the attack surface due to the increasing availability of digital services for citizens, businesses, and public institutions. Hence, data underscores the urgency of adopting structured methodologies to assess risks and prioritize investments effectively in the context of public administrations.

Second side: digitalization and penetration level against global cybersecurity index

As highlighted in the previous paragraph, Italy stands out negatively as one of the most targeted and vulnerable countries in Europe. It has been characterized by a late digitalization, and this is recognized as a reason for which it lacks a strong cyber posture in respect to other European countries (Rapporto Clusit - Italia e PA, 2024).

Consequently, Italy stands as an easier target for cyber-attacks also because it's particularly vulnerable in respect to the rest of the continent. To deep dive into this comparison, three indices are insightful when considered: the level of penetration (PL), the level of digitalization (DL), and the global cybersecurity index (GCI) among European States (Fujs & Bernik, 2024). According to the International Telecommunication Union (ITU), PL is the percentage of the population using a technology, reflecting its adoption and accessibility. DL measures the integration of digital technologies in the society, indicating the shift from analog to digital processes and the adoption of advanced digital solutions. To calculate the GCI, ITU assesses each country based on specific indicators related to the five key pillars on a scale out of twenty: Legal Measures, Technical Measures, Organizational Measures, Capacity Development, and Cooperation. This components are then aggregated into an overall score that is then out of 100 (Global Cybersecuirty Index, 2024). Therefore, if Italy has a GCI score of 96.13, this indicates a high level of commitment to cybersecurity and not a 96% level of protection against cyberattacks. Italy's scores for Digitalization Level, Penetration Level, and the Global Cybersecurity Index (GCI) were 61%, 40%, and 96.13, respectively, in 2024 (Global Cybersecurity Index 2024, 2024). In Figure 3, instead of the country names, abbreviations according to the ISO codes are used and the data is arranged in descending order of DL, and the three scores are reported together. However, GCI is calculated out of 100 as described above but without being a percentage as the other two. For convenience, on the x-axis the % is left as unit of measure.

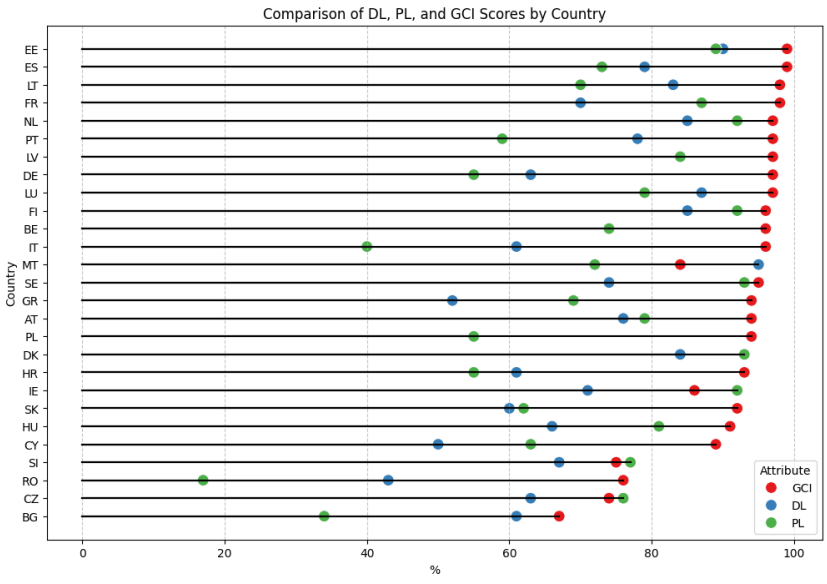

Figure 3 - DL, PL and GCI by Country in 2024

Italy's distinctive position is characterized by a significant low level of penetration compared to the European average, which is 71%, while simultaneously maintaining a GCI score above the European average. In other words, reading the visualization from the left to the right we immediately identify Italy records the third lowest level of penetration and gets the fifth position for digitalization at the same time. While having a high GCI is

promising for Italy, it does not guarantee sustained cybersecurity resilience. As the country works to bridge the digitalization gap and align with the European average, the risk of cyber vulnerabilities will increase due to the rapid pace of technological updates. Thus, digitalization needs to be accelerated cautiously, ensuring that security measures evolve in parallel to mitigate potential risks. Following this wavelength, an additional insight follows. As illustrated in Figure 4 (Global Cybersecurity Index 2024, 2024), Italy's eGovernment is not well-established and remains significantly underdeveloped compared to the broader European landscape. This is particularly noteworthy when considering that, despite these shortcomings, Italy's GCI score surpasses that of several countries positioned in the top-right quadrant of the graph. Specifically, within the group of countries highlighted in the green rectangle - representing Fruitful eGovernance - Italy ranks higher in GCI than Sweden, Austria, Finland, Estonia, and Malta. These five nations constitute nearly half of the twelve countries in this category.

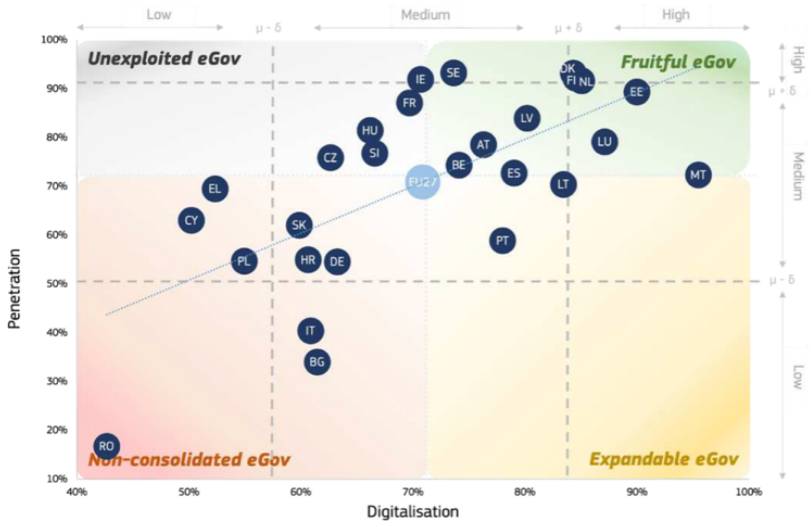

Figure 4 - Levels of Digital Penetration and Digitalisation across Countries

About the actual progress in Italian digitalization, surprisingly the 2024 Digital Decade Country Report for Italy (Commission, 2024) highlights notable advancements in e-government services, particularly in e-Health, where access to electronic health records has surpassed the EU average, reaching 82.7%. In healthcare, cyber threats emerge in several areas. For instance, we highlight the management of sensitive data, which often needs to be transferred between private and public clinics, as well as the malfunctioning of machine learning based tools used for medical examinations. The former pertains to privacy issues, which consistently have a significant impact on public opinion, while the latter affects the accuracy of medical assessments, upon which the health of individual patients depends. Additionally, Italy has made progress in rolling out gigabit networks, with Fiber to the Premises (FTTP) coverage at 59.6%, though still below the EU average of 64%.

Even though Italy has a very low level of digitalization, is it not excluded from the increasing reliance on digital platforms that is a current trend in Western countries. One of the measures of the Italian PNRR is the digitalization, innovation and security in public administration. Below, Table 2 indicates the investments in different (PA digitale 2026, 2025).

Investment Area	**Associated Investment** (million euros)	**Percentage of Total** (%)
Digital Infrastructures	900	10.06
Cloud Migration Enablement and Facilitation	1000	11.18
National Digital Data Platform	556	6.21
Citizen Experience in Public Services	613	6.85
Adoption of PagoPA and App IO	750	8.38
Digital Identity	200	2.23
Digital Services and Digital Citizenship	130	1.45
Digital Skills for Public Administration Staff	500	5.59
Data and Interoperability	140	1.56
Cybersecurity	620	6.93
Technological Infrastructure for Schools	445	4.97
Digital Teaching Integration and Teacher Training	800	8.94
Digital Skills for Students	1100	12.30
Digital Skills for Citizens	195	2.18

Table 2

The investment distribution in the Italian PNRR digitalization plan reflects a strong emphasis on education and cloud infrastructure, with digital skills for students (12.30%) and cloud migration (11.18%) receiving the highest shares of funding. This suggests a long-term vision focused on preparing future generations and modernizing public services, which is explicit in the official documents of national commitment to cybersecurity, and extremely mindful. Indeed, the lack of widespread digital skills implies that the digital services the country provides require parallel advancements in public education, focusing on understanding and the correct use of these services. Also, in the event of cyberattacks targeting public institutions, more than half of the population (reminder: Italian PL is 40%) lacks the knowledge to comprehend the situation. This, in turn, can lead to an underestimation of cybersecurity risks and increased vulnerability to misinformation and manipulation. This issue is further exacerbated by the ease of use of certain AI-based digital tools. The lack of digital skills suggests that many Italians may be using chatbots and similar technologies without fully understanding their associated risks and

functioning. However, in this regard the 2024 report by Clusit stands that the adoption of advanced technologies like artificial intelligence (AI) among enterprises is limited, with only 5% utilizing AI, below the EU average of 8%. The country also faces difficulties in scaling up enterprises, evidenced by having only seven unicorns, accounting for less than 3% of all EU unicorns. To address these issues, Italy has outlined a strategic roadmap aiming to align with EU 2030 ambitions, allocating over €32.5 billion (1.6% of GDP) towards digital transformation initiatives. The report underlines mixed perceptions among citizens regarding digitalization in public services. While Italy has made significant advancements in e-government, citizen satisfaction with these services remains relatively low. Conversely, the literature consistently highlights that concerns related to cybersecurity are hindering progress in the development of e-government services. Only 59% of Italians perceive public administration as effective in providing digital services, compared to the EU average of 68%. Concerns persist regarding accessibility, ease of use, and trust in online government platforms. Additionally, the limited level of digital skills among the population further affects engagement with digital public services.

At first glance, it is surprising that Italy's GCI score significantly exceeds the European average, despite relatively low digitalization and penetration levels. This discrepancy may suggest that while Italy is not among the most advanced nations in terms of digital transformation, the government prioritizes investments in cybersecurity. Notably, through the 'Piano Nazionale per la Ripresa e la Resilienza'[1] (PNRR), additional resources have been allocated to strengthen the country's cyber posture. However, as mentioned before, the annual Clusit Report indicates that Italy is among the most targeted countries for cyberattacks, in Europe but not only. This raises an important question: how can Italy achieve such a high GCI ranking while simultaneously being one of the most affected by successful cyberattacks? As previously explained, the GCI measures a country's commitment to cybersecurity at a global level. However, it does not necessarily reflect the effectiveness of cybersecurity policies or the actual impact of related investments. This highlights a critical issue: the absence of a specific and effective tool for Italian public administrations to enhance their decision-making processes. While there is clear commitment to cybersecurity and an awareness of its importance, the lack of a well-structured organizational framework prevents these efforts from translating into a robust defense against cyber threats. A doubt can arise around the fact that Italian commitment to cybersecurity is strong, but missing effectiveness. Among the five dimensions of Legal Measures, Technical Measures, Organizational Measures, Capacity Development, and Cooperation, Italy performs very well. However, the assessment is practically regarding the commitment ad priori, not the consequent effectiveness in the interventions. This means that there must be an issue in between the commitment and the actualization. Looking at the Market Size per Country 2024 scenario provided in the confidential report (Organisation, 2025) but originally built from the European Cyber Security Organizations in 2024, it is evident that Italy is still among the higher positions. Hence, its problem of being affected by successful

[1] https://www.acn.gov.it/portale/pnrr

cyberattacks is neither a problem of cybersecurity market size. Figure 5 is measuring the billion euros as percentage of the gross domestic product (GDP).

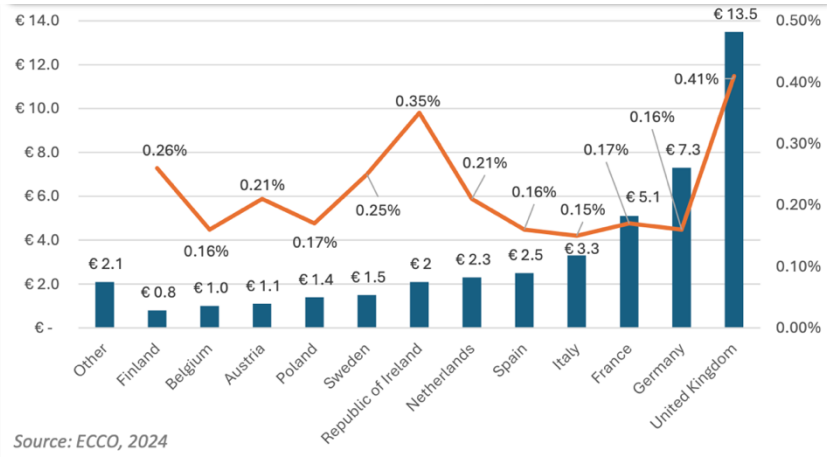

Figure 5 - Market Size per Country 2024

Apart from the United Kingdom, which stands out for its exceptional growth in cybersecurity, the next four leading countries warrant further comparison. Table 3 presents the revenue of the cybersecurity market in these countries alongside their total expenditure in 2024. The data is sourced from Statista[2] and official national reports. These four countries are listed in descending order based on their Global Cybersecurity Index (GCI) rankings.

	revenue of the cybersecurity market in 2024 (billion euros)	Total expenditure in 2024 (billion euros)	Ratio	GCI	PL	DL
Spain	2.85	126.3	2.3%	98.53	73%	79%
France	5.37	166.6	3.2%	97.6	87%	70%
Germany	7	208.8	3,4%	97.41	55%	63%
Italy	2.64	91.2	2.9%	96.13	40%	61%

Table 3

The findings indicate that the size of the cybersecurity market aligns with their high GCI scores. However, no clear correlation emerges between cybersecurity market size and levels of digital penetration or digitalization. This suggests that variations in these factors may account for Italy's distinctive characteristics among the top-ranked countries, potentially offering insights into its relative weaknesses in cyber resilience.

[2] https://www.statista.com/topics/12924/cybersecurity-in-europe/#topicOverview

Considerations for PAs to be different

As in private corporations, particularly larger organizations or those heavily involved in information systems, dedicated teams and managers, often referred to as Security Operations Centers (SOC), oversee cybersecurity, public institutions must select individuals or teams responsible for cybersecurity, decisions on interventions, and the overall management of resources. Despite this formal similarity, there are factors that are most relevant to the public sector, distinguishing it from private corporations due to its unique characteristics. A research by (Wirtz & Weyerer, 2017) identified five risk-related factors in the literature associated with the state of cybersecurity in the public sector: the relevance of cyber-attacks, the degree and the sources of potential danger, the security concerns, and the lack of risk awareness by superior authorities. These points are covered across the sections in this first chapter, since some statistics about the attackers for PAs has been presented already, for example. Considering the attacks in terms of kinds and frequency, some differences are detected as well. From the (Rapporto Clusit - Italia e PA, 2024), the insightful graphs below are taken. In

Figure 6 the data is shown for the public sectors in light blue and for all the sectors together in red.

Figure 6 - Annual increase in number of attacks

In Figure 7, the ratio between the number of attacks to PAs and the total number of attacks is visualized. The peak in 2021 can be intuitively related to Covid19, although this is not the most interesting takeaway. Indeed, it is evident that PAs seems to be less targeted when looking at the graph on the right, although it must be considered that the number of attacks is just a counting that is not taking into account their entity, nor their consequences. Attacks to the private sector, especially for smaller enterprises, are easier to monetize, and furthermore more frequent. Although from the peak in 2021 the percentage of attacks to PAs in respect to the total is diminishing, the rate to which is it doing so it is also going down, but slower.

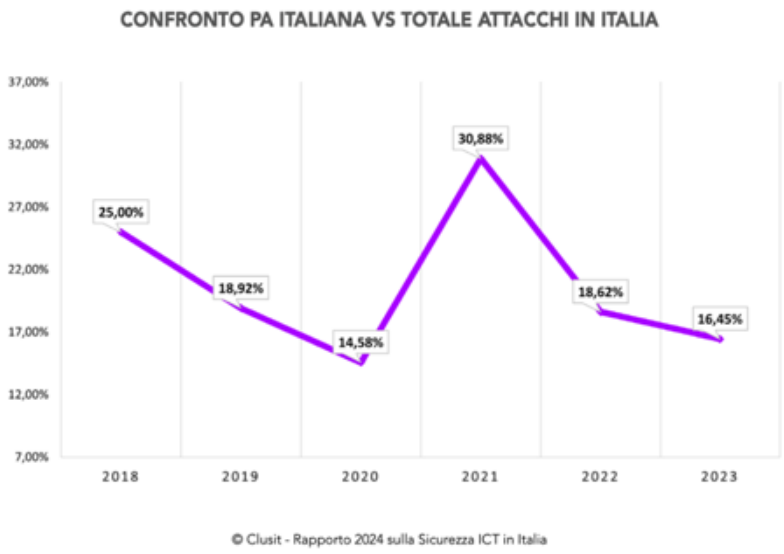

Figure 7 - Rate of attacks to PA over total attacks in Italy

In practice, a public administrator responsible for cybersecurity investments must determine how much to allocate to protect critical assets from evolving cyber threats. Because of what have been highlighted about the Italian context regarding digitalization in the previous section, this decision-making process is not just a matter of financial optimization but a fundamental issue of public trust and governance. If we look beyond the lack of enduring physical consequences, cyber-attacks can still levy tremendous damage by undermining societal cohesion and trust in government institutions, traumatizing civilians, and dividing communities (Shandler & Gomez, 2023). In this regard research into how emotions such as anger, anxiety, and fear affect public confidence in the government's ability to prevent future attacks reveals that while anxiety and fear stemming from exposure to an incident generally reduce trust in authorities, anger has the opposite effect, enhancing confidence (Shandler & Gomez, 2023). This study was validated through an analysis of an incident in Düsseldorf, Germany. When discussing public trust as a critical factor in cybersecurity within the public sector, these emotional responses must also be considered. The relationship between public institutions and citizens operates within an interdependent and dynamic system. Given the sensitivity and the wide-ranging impact of potential breaches in public administrations, providing public administrators with a practical, quantitative tool to guide investment decisions is particularly beneficial.

In addition, public administrations provide essential services that are increasingly digitalized and automated, making them attractive targets for cyberattacks. Unlike in the private sector, where customers can choose between service providers, citizens have no alternative but to rely on government institutions, entrusting them with sensitive personal data. Moreover, PAs operate in a highly regulated environment, bound by stringent legal and ethical obligations. Cyber incidents in the public sector not only lead to financial losses but can disrupt essential services, compromise national security, and undermine public confidence. A distinction can be made between

actions that explicitly seek to undermine public confidence and social cohesion (i.e., misinformation operations), and actions where the primary aim is something else (data manipulation, physical destruction, etc.) (Shandler & Gomez, 2023). While the latter occurs in both the private and public sectors, the former predominantly affects the public sector, further highlighting its particular vulnerability and the associated risks. The field of cybersecurity is currently considered to be in its third era, as identified by (Ganapati, Ahn, & Reddick, 2011). This indicates that cybersecurity is no longer solely a technical or legal issue but also a managerial concern. The article emphasizes this shift by noting that 'literature on cybersecurity in public administration blossomed from 2006 onwards'. However, the most significant takeaway from this study is the recognition of the need for continued research to develop a more adaptive, agile, and effective policy framework to address cybersecurity challenges at the macro level. In this context, the macro level refers to national cybersecurity efforts, highlighting the necessity of fostering widespread awareness to ensure consistent investment in cybersecurity. Such investments would enable law enforcement agencies to operate cohesively, aligning with and reinforcing the priorities identified through security initiatives.

Finally, some trends are particularly critical for the public sector in terms of cybersecurity. One of them is the European aim of transforming cities - and more broadly, infrastructures - into smart systems. Fiber To The Premises (FTTP) is just one example among others, with the widespread use of the Internet of Things (IoT) demanding an extensive and reliable network. With the proliferation of IoT devices, nations are collecting vast amounts of data, which in turn introduces significant vulnerabilities into their systems. In general, hackers exploit both basic protocols and application vulnerabilities depending on the situation (Šimec, 2019). Other emerging threats stem from the increasing reliance on cloud-based services, the implementation of tools to facilitate e-voting (Ganapati, Ahn, & Reddick, 2011). These scenarios highlight that certain European transformation strategies aim to be data-driven rather than relying solely on massive data collection. However, this shift undoubtedly introduces new vulnerabilities, emphasizing the need for diverse investments to ensure nations are adequately prepared, responsive, and resilient against emerging cybersecurity threats.

Potential extensions

As stated from the outset of this research, the intention to generalize the framework and its associated tool beyond the context of public administration was clear. This project has the potential to serve as a pivotal development in the management of cybersecurity concerns within the broader domain of IT management. In this regard, Figure 8 is particularly relevant, as it presents the most pressing technical and operational challenges facing IT departments, based on a widely recognized industry survey sector (Johnson, et al., 2024). Notably, 'Security / Cybersecurity / Privacy' emerges prominently among the top-ranked concerns, underscoring the urgency and relevance of advancing effective managerial tools in this domain.

Organizational IT Issue	2023 (n= 436)	2022 (n= 540)	2021 (n= 454)	2020 (n= 624)	2019 (n= 618)	2018 (n= 793)	2017 (n= 769)	2016 (n= 801)	2015 (n= 785)	2014 (n= 717)	2013 (n= 483)
Alignment of IT with the Business	1 (44.7%)	2 (33.9%)	2 (33.3%)	2 (35.1%)	2 (33.2%)	2 (32.8%)	2 (37.3%)	1 (41.7%)	1 (42.4%)	1 (26.2%)	1 (43.7%)
Security/ Cybersecurity/ Privacy	2 (41.7%)	1 (51.1%)	1 (42.5%)	1 (36.1%)	1 (35.9%)	1 (38.3%)	1 (41.9%)	2 (36.0%)	2 (31.5%)	2 (17.6%)	7 (11.2%)
Data Analytics/ Data Management	3 (27.1%)	3 (28.7%)	3 (24.7%)	3 (25.3%)	3 (25.7%)	3 (26.9%)	3 (23.4%)				
Digital Transformation	4 (25.9%)	5 (22.2%)	4 (24.4%)	4 (24.4%)	4 (22.2%)	7 (19.5%)	8 (18.7%)				
Compliance and Regulations (e.g., HIPAA, SarBox, SAS70, PCI etc.)	5 (21.8%)	4 (28.3%)	5 (23.6%)	4 (24.4%)	5 (20.6%)	6 (19.9%)	4 (20.7%)	12 (13.5%)	11 (16.2%)	12 (9.1%)	16 (6.0%)
AI/Expert Systems/Machine Learning	6 (20.4%)	22 (9.6%)	23 (7.5%)	16 (10.1%)	13 (12.8%)	16 (11.6%)					
Cost Reduction/ Control - IT	7 (19.0%)	17 (11.5%)	10 (15.6%)	6 (23.7%)	8 (18.6%)	9 (17.8%)	5 (20.0%)	7 (19.0%)	8 (17.3%)	17 (8.2%)	5 (16.8%)
Business Continuity	8 (18.3%)	7 (17.8%)	7 (19.2%)	7 (22.8%)	16 (12.0%)	12 (14.0%)	18 (10.8%)	11 (13.7%)	15 (12.4%)	22 (5.0%)	
Cost Reduction/ Control - Business	9 (17.7%)	14 (13.7%)	13 (13.2%)	8 (19.4%)	10 (16.7%)	10 (14.5%)	6 (19.9%)	6 (19.7%)	10 (16.3%)	9 (12.3%)	4 (18.6%)
Cloud/Cloud Computing	10 (15.8%)	8 (16.9%)	6 (19.4%)	9 (18.3%)	6 (19.7%)	13 (13.7%)	14 (12.2%)				

Figure 8 - Organizations' Most Important IT Management Issues, 2013-2023

Organizational IT Issue	2023 (n= 436)	2022 (n= 540)	2021 (n= 454)	2020 (n= 624)	2019 (n= 618)	2018 (n= 793)	2017 (n= 769)	2016 (n= 801)	2015 (n= 785)	2014 (n= 717)	2013 (n= 480)
Security/ Cybersecurity/ Privacy	1 (44.5%)	1 (55.2%)	1 (46.9%)	1 (40.9%)	1 (46.3%)	1 (46.4%)	1 (47.7%)	1 (46.4%)	1 (36.8%)	1 (25.5%)	2 (19.8%)
IT Talent/Skill Shortage/Retention	2 (27.1%)	2 (38.1%)	3 (22.2%)	5 (18.4%)	3 (21.5%)	2 (25.6%)	3 (23.5%)	2 (28.3%)	3 (28.3%)	2 (20.9%)	3 (19.6%)
Alignment of IT and/with the Business	3 (26.8%)	3 (25.2%)	2 (24.4%)	2 (23.4%)	2 (25.1%)	4 (19.8%)	4 (21.8%)	3 (24.0%)	2 (29.7%)	3 (19.9%)	1 (32.5%)
AI/Expert Systems/ Machine Learning	4 (24.3%)	27 (6.5%)	32 (5.3%)	27 (7.7%)	31 (6.3%)	26 (7.3%)					
Credibility of IT/ Perception of IT Leadership	5 (20.9%)	6 (17.8%)	4 (21.6%)	4 (20.4%)	4 (20.4%)	3 (22.1%)	2 (24.4%)	4 (20.3%)	6 (16.4%)	18 (7.1%)	
Compliance and Regulations (e.g., HIPAA, SarBox, SAS70, PCI etc.)	6 (19.3%)	4 (19.4%)	6 (16.5%)	7 (15.9%)	5 (15.7%)	6 (16.3%)	5 (16.9%)	11 (12.1%)	13 (12.2%)	14 (7.5%)	16 (7.5%)
Business Continuity	7 (17.7%)	5 (19.3%)	5 (17.4%)	3 (21.8%)	8 (14.1%)	5 (17.4%)	8 (14.0%)	5 (17.4%)	7 (16.2%)	13 (7.8%)	
Improving IT Communications and Relationships with the Business	8 (15.8%)	11 (13.7%)	11 (11.7%)	14 (11.7%)	6 (15.2%)	11 (12.4%)	9 (13.4%)	10 (12.6%)			
Cost Reduction/ Control - IT	9 (15.6%)	25 (8.3%)	16 (10.4%)	9 (14.4%)	17 (10.2%)	15 (11.3%)	20 (10.1%)	13 (11.1%)	21 (9.0%)	30 (3.9%)	13 (8.3%)
Data Analytics/Data Management	10 (15.4%)	7 (15.2%)	9 (14.1%)	12 (12.3%)	7 (14.2%)	8 (14.1%)	7 (14.2%)				

Figure 9 - IT Leaders' Most Important or Worrisome IT Management Issue, 2013-2023

It is particularly noteworthy that similar evidence is presented in Figure 9, which offers a ranking of IT management issues from the valuable perspective of IT leaders. This highlights a critical gap - one in which the generalization of the proposed framework could be effectively positioned. The findings underscore the need for

advanced tools capable of addressing cybersecurity challenges at the managerial level, where strategic decision-making and risk mitigation are most impactful.

A potential extension of this work could also involve a focused analysis of cybersecurity investments for specific PA assets. In this regard, artificial intelligence (AI) emerges as a particularly relevant topic. According to (Jonhson, Maurer, Torres, Guerra, & Mohit, 2024), artificial intelligence has risen significantly in IT risk management rankings, moving from 22nd place to 5th place within a year. This shift highlights the increasing recognition of AI-related risks in IT governance. Additionally, the report emphasizes the persistent challenges of legacy systems maintenance, particularly with government organizations which continue to struggle with updating and re-platforming outdated applications. These insights emphasize the broader importance of cybersecurity investments, extending beyond the public sector. It must be underscored the accelerating pace at which AI is becoming a key driver of the global economy and the associated concerns that emerge alongside its adoption, especially since Generative AI (GenAI) come into place. This trend further supports and reinforces the relevance of this Action Design Research (ADR) approach in addressing these evolving challenges. A possible next step of this research project could develop an artefact capable of guiding cybersecurity investments in AI-enabled environments, by optimizing the prioritization and type of interventions based on the specific characteristics of the organization under analysis. The tailored version of the artefact would be applicable both in contexts already characterized by AI applications - where it can be used to assess the effectiveness of past security investments and to optimize future ones - and in environments where AI adoption is still under evaluation, where it can provide a quantitative basis for estimating the expected benefits of security-related investments. The growing adoption of artificial intelligence as a driver of economic development is often accompanied by a perception of high risk, particularly among small and medium-sized enterprises (SMEs). The investments required for implementing AI-enabled solutions are indeed significant, while quantitative evidence of their return on investment is often limited or difficult to compare. This uncertainty leads to increased caution in the adoption of such technologies, slowing their diffusion and limiting the competitive potential of SMEs. Considering these challenges, it is crucial to develop tools that enable structured and data-driven evaluations of the value of security investments in AI-enabled environments. This research project could address this gap by offering a model that allows organizations to optimize their cybersecurity strategies based on their unique characteristics and the specificities of their operational context. The project proposes to combine the economic evaluation technique of Return on Security Investment (ROSI) with the methodological rigor of Capability Maturity Models (CMMs). In particular, the ROSI approach might be enhanced in its quantitative dimension by integrating the Monte Carlo method to simulate attack costs, and by refining other parameters through dedicated functions. Focusing the artefact on AI-enabled environments would allow for a more targeted approach, enabling both the adaptation of the Capability Maturity Model and the refinement of the ROSI framework according to the specific vulnerabilities of such systems. The

output of the artefact, developed through the Action Design Research (ADR) methodology, has dual value: thanks to the interaction between the Capability Maturity Model and the investment evaluation tool, organizations will be able to obtain increasingly accurate analyses and tailored recommendations. The combined output can contribute to encouraging secure investment expansion in AI-enabled environments.

Apart from AI, which is a topic tackling almost all different sectors, a brief instance of application in one specific sector is discussed to account further the research applicability and potential extension. Healthcare could warrant peculiar examination as well, for example. As previously discusses, there are several ways to which its vulnerabilities serve as entry points for cyberattacks, and at the same time healthcare ranked second among the top targeted sectors worldwide and in Italy. A distinctive aspect of the Italian healthcare system is that it primarily operates as a public service, although several private clinics are emerging. The growing competition driven by waiting times for medical services places patients in a difficult position, often forcing them to choose between paying for private care or enduring long wait times for public services. Given this scenario and considering that cybersecurity is equally critical for private clinics to protect their business, a successful cyberattack on the Local Health Authorities (ASL) would have severe consequences, not only disrupting essential services but also eroding public trust and further complicating healthcare management. Therefore, investing in cybersecurity to enhance resilience and strengthen cyber defenses should be an absolute priority. Among various public administrations, some possess distinct structural characteristics and engage in extensive interactions with the population, necessitating tailored cybersecurity measures. This further supports our ADR research, demonstrating the significant value of a PA-specific cybersecurity framework.

Chapter 2 - Overview

Modularity

Modularity represents the key value and core strength of this artifact. It manifests in two primary ways. First, the modular approach to the various components of ROSI (e.g., Economic Benefit, Cost of Investment) allows for their independent development, enabling prioritization of those elements that can be implemented or upgraded immediately. Second, the artifact's modularity also lies in the capacity to specialize each of these components independently, thereby enhancing the granularity of resources and the precision of potential recommendations. Furthermore, individual components can be replaced or updated as needed. In this sense, the artifact is composed of "building blocks" that can be addressed one at a time, according to convenience and strategic importance. From a structural standpoint, this modular approach is essential to ensuring a model that is both exponentially scalable and easily adjustable at every level - without compromising coherence or meaning. Indeed, the research holds the potential for further extension from the very beginning, and the intention was then to structure it in a way that ensures scalability in a subsequent phase. For this reason, the methodology is strongly focused on enabling the broadest possible applicability and flexibility of the artifact as explained. In practice, this approach entails avoiding assumptions that cannot be disregarded in future adaptations of the tool. Moreover, data that has been used so far is almost taken from secondary sources for most of the factors, but the idea of having primary sources at same point in future developments of this initial prototype has been clarified from the beginning. As soon as that data would become available, the corresponding module of the tool would be updated with it, not only as inputs but also as resources to readapt the definitions and the algorithms defining the related factor. Thanks to this approach, each time this kind of updates are experienced, the model registers an improve, since new data concerning an area is not impacting negatively any other but just enhancing the overall appropriateness of the result. Moreover, the concept of modularity offers a significant advantage to this model compared to others. Beyond enabling upgrades and the extension of foundational elements and applications, modularity allows the artefact to be effectively adapted to diverse contexts without compromising its usefulness. This is achievable by selecting context-specific modules that serve as appropriate profilers, and by designing algorithms that compute factors based on these modules. Ultimately, the modular nature of the artefact unlocks a high degree of flexibility. This point is emphasized at the outset of the work because it was crucial to maintain this perspective from the initial design phase. Indeed, creating a framework with such a level of generalizability - despite limited data availability - required a prolonged and thoughtful reflection on the architectural structure. This process involved considering all potentially valuable primary resources that could be integrated and effectively utilized.

To facilitate understanding of the research essence and future potential, as well as the strength of its methodological approach – modularity -, the metaphor of a tree is used. Inspired by Bruno Munari's "Disegnare

un albero" (literally: ‘Drawing a tree’) (Munari, 1978), the framework grows in the same way one learns to draw a tree: starting from a solid, modular root structure and developing branches organically and progressively. Each root or branch can be deepened or expanded independently, without compromising the overall balance and coherence of the system. In this perspective, the ADR method works on the ROSI components (the roots), unlocking the possibility of building, above them, a series of strategic paths, interfaces, and asset-focused developments (the canopy). Both roots and canopy are modular: they can grow vertically in one direction at a time, maintaining the robustness of the model. Organizations can choose from which of the three core ROSI factors (root bodies) to start expanding, depending on their context and strategic priorities. Just as in a modularly designed tree, the framework roots itself in the solidity of ROSI components and grows upward and outward through targeted strategies and assets, ensuring flexibility, adaptability, and structural coherence.

Figure 10 - The canopy

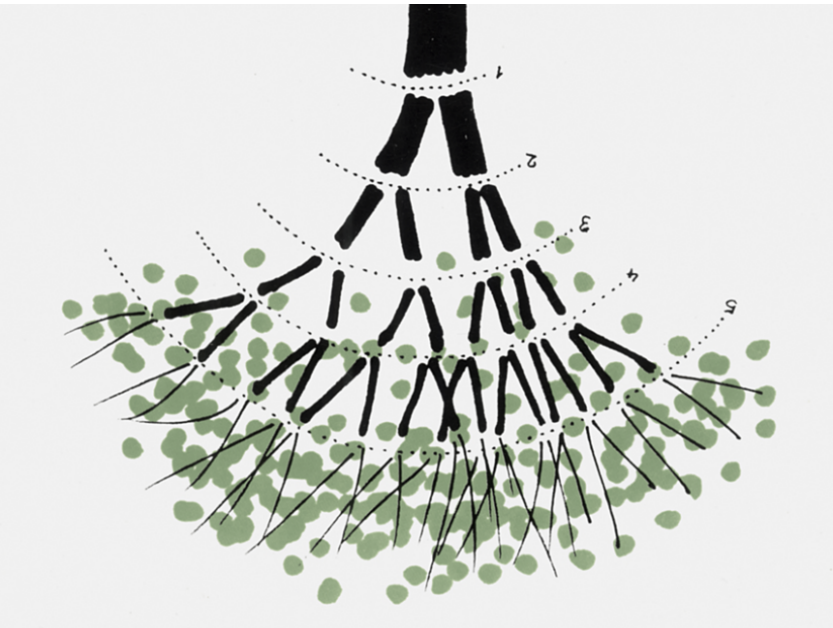

Figure 11 - The roots

Now, it is time to start talking about the framework and the first thing that is pointed out is the core of the tree, which is the ROSI, metaphorically represented in the trunk. The Return on Security Investment (ROSI) is a widely recognized metric used to assess the financial benefits of cybersecurity measures relative to their costs. Originally developed for the private sector through the adaptation of the Return on Investment (ROI) concept to security-related topics, ROSI provides a structured approach to evaluating whether security investments yield a positive return by preventing potential losses due to security breaches. However, its applicability to public administration contexts requires significant revision to account for the unique operational, economic, and organizational conditions of the public sector. For example, the conventional ROSI model often relies on financial risk reduction as the primary measure of effectiveness. In PAs, the impact of security investments extends beyond financial returns to include aspects such as public trust, and the preservation of critical infrastructure, for example. Since the introduction of ROSI, various studies have been conducted, leading to the development of multiple frameworks that are now applied globally. Each approach is distinguished by specific characteristics, and in this chapter, we aim to identify and analyze the most significant ones. Our objective is twofold: first, to extract the beneficial aspects that may contribute to the development of PA-specific framework, and second, to critically assess the limitations of these existing models in the context of public sector requirements.

In response to the need of a guidance in managing investments to mitigate the cyber risk, “numerous methodologies of Return On Security Investment (ROSI) exist to help decision makers but they pose great challenges in the domain of cyber security” (Arshad, Abbas, Faisal Amjad, Shafqat, & Yaqoob, 2019). Therefore, the literature has proposed various refinements to the ROSI model aimed at enhancing its quantitative rigor and reducing subjectivity in investment evaluations. While these refinements are not specifically thought for the public

sector, several studies have highlighted that "National cyber programmes have different objectives to organization-led cyber programmes" (Onwubiko & Onwubiko, Cyber KPI for Return on Security Investment, 2019), emphasizing the need for tailored methodologies that account for sector-specific challenges such as regulatory constraints, public accountability, and the non-discretionary nature of service provision. Hence, we structure our literature review as follows. We first discuss the models that we consider significant adaptations of the ROSI framework. Next, we examine the contributions of the literature regarding the appropriate weighting of various assets when addressing public organizations. Additionally, we may expand our literature review to include models for assessing organizational maturity in cybersecurity. This aligns with our objective of developing a framework to guide security investments for organization. The structure of the literature review is informed by the approaches proposed by (Webster & Watson, 2002) and (Sein, Henfridsson, Purao, Rossi, & Lindgren, 2011). In particular, a concept matrix is employed as a tool to concisely visualize the current state of the art in the relevant research areas.

Definition of ROSI

A first dive into the presentation and explanation of the standard formula is provided, followed by a discussion of its various factors. Following the previously introduced metaphor, this corresponds to beginning with the trunk before examining its roots. It is important to note that a critical challenge concerns the estimation of Economic Benefits, which often relies on approximations, historical data, and industry benchmarks. Although some sources refer to the Return on Security Investment using different terms (e.g., cyber-ROI), there is general consensus on its definition, which is presented as follows.

$$ROSI = \frac{Economic\ Benefit - Cost\ of\ Investment}{Cost\ of\ Investment}$$

The *Economic Benefit* represents the monetary value of the losses that are prevented thanks to your security measures. This is often complex to estimate because - briefly summarizing the issue - it requires to associate a price on something that did not happen and that it dependent on many variables. On the other hand, the *Cost of Investment* is the actual cost of the interventions. Based on experience and as emphasized in the literature, the investment for the interventions is not a primary concern when evaluating the ROSI. However, two aspects are introduced here and will be addressed in its dedicated chapter. Upon completing the evaluations of specific single investments, the results can be aggregated to determine the overall ROSI for each PA, but comparing ROSI of different PAs doesn't have significance before having decided the management of indirect benefits. Another critical factor is the extent to which different expenditures should be included in the calculation of monetary investments for a single intervention *i*. It is essential to establish a rigorous methodology for determining the scope of investment inclusion: this is vital for enabling meaningful comparisons across different assessments. Without a standardized approach, two security investment scenarios may appear to differ significantly in their ROSI merely

because one has included a broader range of associated costs. A structured classification system, as explained in the dedicated chapter, would mitigate such misleading results, ensuring a more objective and consistent evaluation of security investments. Overall, a detailed and extensive framework can be discussed around the calculation of ROSI itself, but also about its roots and its applicability. This thesis is addressing all of these, knowing that this is however a first step for a wider research project to be developed.

Literature Review

The adaptations of the ROSI tool and some other relevant studies identified in this literature review are reported next, focusing on those most relevant to this research. These adaptations vary mostly in their approaches to assessing and calculating economic benefits. While some methods adopt a more qualitative perspective, others emphasize quantitative analysis. Each adaptation is outlined individually, with a comparative discussion provided in the final paragraph. Generally, the Return on Security Investment has been widely recognized as a decision-support tool for cybersecurity investments, often grounded in formal models for the calculation of costs and benefits ((Barik, Misra, Fernandez-Sanz, & Koyuncu, 2023), (Butler, 2002), (Gheorge, 2012), (Liu, Zhang, & Chen, 19), (Tsiakis & Stephanides, 2005), (Arshad, Abbas, Faisal Amjad, Shafqat, & Yaqoob, 2019) (Arshad, Abbas, Faisal Amjad, Shafqat, & Yaqoob, 2019), (Böhme & Moore, 2010), (Collier, Briglia, Slutzky, & Lambert, 2023) (Marican, Othman, Selamat, & Razak, 2024), (Pontes, Guelfi, Silva, & Kofuji, 2011)). Several frameworks have been proposed to support the adoption of the ROSI approach, aiming to provide structured methodologies for its implementation across different organizational contexts ((Arshad, Abbas, Faisal Amjad, Shafqat, & Yaqoob, 2019) (Arshad, Abbas, Faisal Amjad, Shafqat, & Yaqoob, 2019), (Butler, 2002), (Erolaa, et al., 2021), (Tsiakis & Stephanides, 2005), (Gheorge, 2012), (Sonnenreich, Albanese, & Stout, 2006)). However, persistent challenges remain in defining standardized inclusion criteria for cost components, as well as in accurately mapping defense mechanisms and system vulnerabilities for reliable ROSI calculations ((Arshad, Abbas, Faisal Amjad, Shafqat, & Yaqoob, 2019) (Arshad, Abbas, Faisal Amjad, Shafqat, & Yaqoob, 2019), (Butler, 2002), (Onwubiko & Onwubiko, Cyber KPI for Return on Security Investment, 2019), (Erolaa, et al., 2021), (Böhme & Moore, 2010), (Liu, Zhang, & Chen, 19), (Sonnenreich, Albanese, & Stout, 2006)). Other research has been driven on these topics. For instance, the University of Zurich appears to be at the forefront of advanced and specialized research in this field, particularly in aligning the applicative potential of machine learning with complex systems such as Cyber Value at Risk. Notably, one of their contributions includes the development of an interface tool through which organizations are profiled based on predefined drivers, allowing for the calculation of Return on Security Investment (ROSI). Interestingly, such an interface has also been made available as a commercial service offered by certain insurance companies. In this regard, this ADR's artefact draws inspiration from these existing solutions, envisioning a similar trajectory in terms of future applicability and potential service integration.

Practical Quantitative Approach

The Practical Quantitative Model in (Sonnenreich, Albanese, & Stout, 2006) emerges from a strong practical background, aiming to bridge the gap between theoretical models and real-world business needs. The study underscores the importance of accurate security metrics, acknowledging that while absolute precision may be unattainable, consistent and repeatable measurements can still provide meaningful insights for decision-makers. One critical issue raised is the "ostrich response", where organizations fail to track security incidents properly due to internal embarrassment or concerns about public perception. This lack of reliable incident data weakens actuarial estimates, leading to flawed risk exposure assessments. Moreover, the model emphasizes the economic impact of lost productivity, noting that even minor daily security disruptions can accumulate into significant financial losses. By incorporating productivity loss into ROSI calculations, the model provides a tangible link between security investments and business efficiency. To operationalize these concepts, the authors developed SecureMark, a real-world benchmarking system designed to provide a standardized method for measuring ROSI. SecureMark functions by conducting structured security assessments, where organizations are evaluated on key security parameters derived from industry standards such as ISO 17799, NIST, and ISF recommendations. The system quantifies risk exposure in monetary terms by calculating lost productivity due to security vulnerabilities. It also assesses the effectiveness of security measures, providing a scoring system that ranks security investments based on their ability to mitigate risks and improve operational efficiency. Unlike traditional ROSI models that struggle with uncertainty regarding the likelihood and impact of security incidents, SecureMark shifts the focus to measurable, recurring productivity losses, which can be systematically reduced through targeted security investments. By leveraging structured surveys, SecureMark collects data on employees' perceived downtime, disruptions, and inefficiencies caused by security-related issues. SecureMark allows organizations to compare their security performance against industry benchmarks, facilitating data-driven decision-making regarding security investments.

ROA (Return On Attack)

This paper poses a quite different approach, considering not only the return on investment economically speaking, but also the relevance of attackers being discouraged to perceive their goal rather than their being incentivize in finding a more efficient solution. We report the definition of their proposal, the Return-On-Attack (ROA), where S stands for Security measure: $ROA = \frac{gain\ from\ successful\ attack}{cost\ before\ S - loss\ caused\ by\ S}$. There remains significant complexity in determining how to quantify the factors in the denominator, as for the others approaches. Moreover, this one aligns with game theory, as the authors explicitly acknowledge at the conclusion of their paper. However, within the context of PAs, this perspective appears to be less suitable for addressing investment security. Game theory is particularly relevant when all players have access to predictable strategic outcomes; however, this assumption

does not hold in the case of attackers, whose actions are inherently uncertain. PAs can vary considerably in their attractiveness to attackers depending on their operational functions, more than the private sector. Furthermore, they are more susceptible to political events and dynamic shifts, making a game-theoretic approach less effective in this domain.

A Cost-Benefit Approach

The Security Attribute Evaluation Method (SAEM) proposed by (Butler, 2002) follows a four-step process to evaluate security design choices. The paper presents the overall framework and its results in comparison with alternative security designs in the context of a non-profit organization's financial and accounting system. This necessitates a focus primarily on the theoretical aspects, as the case study itself is not directly applicable as a model for PAs. However, the author specifies that the framework is also being tested to evaluate a government e-commerce system, suggesting its potential relevance to our target context as well. A critical consideration highlighted in the paper is the scarcity of statistical threat data, which leads security managers to rely on experience, judgment, and the best available knowledge to assess the likelihood of attacks and their associated impact. This presents a key limitation of the framework, as it depends heavily on the perspectives and expertise of security managers. While this reliance may be advantageous for a private corporation, where the eventual security operations center (SOC) team possesses detailed and context-specific knowledge that can enhance the cost-benefit approach, it raises significant challenges in the case of PAs. In this context, a comparative approach across different organizations becomes impractical if it depends on the assessments of individual security managers or teams within each institution, as a standardized, cross-institutional evaluation would require a consistent and centralized source of expertise - something that is neither feasible nor scalable in the public sector.

ROSI for Security-Oriented Organizations

The proposed structured framework in (Arshad, Abbas, Faisal Amjad, Shafqat, & Yaqoob, 2019) for calculating Return on Security Investment (ROSI) addresses limitations in traditional approaches. The model follows a six-phase methodology to systematically assess security investments. A little explanation of each of them is given, exposing step by step the related limitations for the eventual application of the framework in the public sector.

I. The initial phase involves asset identification and analysis, during which assets are categorized and prioritized according to their criticality. The classification of these assets follows the guidelines outlined in ISO-27001. Given that the objective is to establish a unified framework for diverse PAs, it is important to recognize the inherent differences in the specific types of assets they hold. Moreover, a suitable number of assets to be considered should be determined, complementing the

broader framework. To deep dive into the evaluation of the asset value, we would like to state the two formulas that the paper suggests using. The first aims to evaluate the Criticality of an asset; all the single factors take values for 1 to 5, C stands for Confidentiality, I for Integrity and A for Availability: $Criticality = C + I + A$. The second formula associates a monetary value to an asset, based on its criticality level in respect to the corporation we consider: $Asset\ Value = Physical\ Cost * Criticality\ Value$.

II. The subsequent phase involves the identification of vulnerabilities and threats, utilizing tools and threat modeling techniques to detect weaknesses in critical assets. This process relies on publicly available data, the expertise of security professionals, and the application of security scanning tools. In the context of PAs, the effectiveness of this phase is heavily contingent upon the level of advancement achieved by the government in terms of information security. In Italy, for instance, security experts leverage both their experiential and implicit knowledge, as well as explicit data, including a list of known vulnerabilities, to facilitate risk identification. Notably, an important resource that can be seamlessly integrated into this approach is provided by the Computer Emergency Response Team (CERT-AGID – Agenzia per l'Italia Digitale). Indeed, CERT-AGID offers a valuable service to PAs, supplying a stream of textual Indicators of Compromise (IoCs) related to Italian malware and phishing campaigns detected through Open-Source Intelligence (OSINT) and Closed Source Intelligence (CLOSINT).

III. The third phase, likelihood and impact determination, employs a Bayesian theorem-based statistical model to estimate the probability of security breaches, addressing the subjectivity inherent in previous approaches. Bayes' theorem is expressed as follows: $P(A|B) = \frac{P(B|A)P(A)}{P(B)}$, where $P(A|B)$ represents the posterior probability of event A given that B has occurred, $P(B|A)$ is the conditional probability of B occurring given that A has taken place, $P(A)$ and $P(B)$ are the marginal probabilities of A and B, respectively. In this context, the model leverages Bayes' theorem to estimate the probability of a security breach by incorporating multiple conditional probability values. These values are more practical to collect and more reliable than generic probability estimations, as they pertain to specific cases rather than broad assumptions. A critical aspect to consider is the accuracy of these conditional probabilities. In the referenced study, such values were derived from the Common Vulnerability Scoring System (CVSS) dataset, which provides a generalized perspective. To enhance the model's precision, a possible approach would be to improve data fitting by collaborating with internal organizations specializing in cybersecurity incident analysis, such as CERT-AGID in the Italian context. This would allow for a more context-aware and refined probabilistic assessment of security threats.

IV. The framework then advances to the countermeasure analysis phase, in which potential security investments are assessed based on their effectiveness and associated costs. A critical challenge in both the public and private sectors is addressing the overlap that may arise when a single security investment impacts multiple assets. To ensure accuracy in cost estimation and avoid redundancy, it is essential to establish a systematic approach for managing this overlap, thereby preventing the duplication of investment calculations.

V. The ROSI calculation phase integrates a cost-benefit analysis by comparing estimated annual losses before and after the implementation of security measures.

VI. Finally, the framework provides strategic recommendations, allowing decision-makers to evaluate whether a security investment is justified based on its anticipated financial and security benefits. If the Return on Security Investment (ROSI) is negative, the investment is generally not recommended, whereas a neutral ROSI suggests the need for further investigation before proceeding. However, in the specific context of PAs, additional factors such as public interest and data privacy may influence the final recommendation. These considerations could justify an investment even if its immediate financial return appears neutral or negative, emphasizing the broader societal and regulatory implications of security measures.

Gordon - Loeb

In addition to ROSI, another widely recognized economic framework for assessing the quality of security investments is the Gordon-Loeb model. Like ROSI, it emphasizes cost-effectiveness, though its quantitative aspects can still be further refined. A prevailing trend in the field is the adaptation of financial quantitative methods to calculate Value at Risk in the context of cybersecurity. However, it is crucial to acknowledge that, among the various financial methodologies available, the dynamic and unpredictable nature of cyber risks necessitates considering results derived from assumptions that may not always be entirely restrictive. The Gordon-Loeb model, as presented here, does not incorporate financial tools but adopts a significantly more quantitative approach compared to the other frameworks discussed above. Back in 2002, from their research on the economics of security investments, Gordon and Loeb have released three interesting propositions regarding the relation between the eventual change in vulnerabilities and the optimal investments to be made (Gordon & Loeb, 2002). From their work, the quantitative approach is strongly inspirational in relation of this research as well, for the fact they introduced a function to express the probability that an information set vulnerability would be breached. Some quantities must be given to consider the pipelines that are described to calculate the expected benefits of the information security investments:

1. λ: the loss conditioned on a breach occurring,
2. τ: the probability of the threat to occur,

3. υ: the vulnerability, intended as the probability that a threat once realized would be successful).

Going on, some derived measures are important to be defined as well:

4. L: the loss or potential loss for an information set, which is given by the product of the probability of a threat occurring and the loss conditioned on a breach occurring,
5. S(z, υ): the probability that an information set with vulnerability υ will be breached, conditional on the realization of a threat and given that z is the security investment that has been made.
6. EBIS(z) = [υ - S(z, υ)]L: the expected benefit from the investment in information security, which can be easily extended as the net benefit if we subtract the value of z afterwords. An important note here: EBIS is depending on z only, since the monetary investment is the only thing that is up to the organization to determine its cyber posture.

Representing the EBIS(z) on a graph in which the monetary investment is set on the x-axis, we get that the actual best investment can be identified as the z* for which the distance between the benefits and the costs is maximized. Hence, we can take the first derivative of S(z, υ) in respect to z to investigate the best investment. Also, considering the dynamism of the vulnerabilities landscape, we might take into consideration the second derivative of S(z, υ), the mixed one. Let me list here the three propositions, taking into mind that we particularly leverage the first and the third one for our purposes. Please consider that along with the proposition, some considerations are reported about the ways each of them would be useful for our case study.

1. *Proposition 1. Under assumptions […], for L and for a range of υ, an increase in υ would be followed by an increase in z*.* This one is just making sense of what we get from intuition, so basically the fact that if the vulnerability increase, also the related loss would do the same. While it's clear that this first proposition is not tackling the way this relation is possibly linear rather that following other trends.
2. *Proposition 2. Under assumptions […], it's not always necessary the case that the optimal investment in information security z*(υ) is weakly increasing in vulnerability.* This one is mostly targeting an eventual misinterpretation of the first proposition, from which we might think that the loss is linear with the vulnerability. This is not necessary the case and in the case of public administrations we might also find a solution to characterize the kind of vulnerabilities, since for sure some threat would affect loss in a more consistent way than others.
3. Without stating the proposition itself, we would just report the results of it got under the same assumptions of the other two and considering class I and II of security breach probability functions. In simple terms, it's proved that the best security investment is always at maximum 36, 79% of the loss that would be expected in absence of any investments. This percentage is considerably valuable for the second phase on this research, since it seems to set a range and an advise for the investments the public administrations have to make to enhance their resilience.

In estimating the probability of a threat affecting a public administration, at the current version of the artifact the probability for each case is treated as a fixed value from ISTAT (2023) and Bankitalia (2024). In contrast, the approach proposed in this paper aims to adopt a more flexible function for estimating probability, introducing a dependency on both the level of security investment and the associated vulnerability that is indeed very much appetible for future updates of the model.

Cyber Value-at-Risk

As mentioned when introducing Gordon-Loeb as one of the alternative quantitative approaches to ROSI, financial tools are also commonly considered to be useful. A relevant financial metric in this context is the Risk-Adjusted Value at Risk (RVaR), which presents both advantages and drawbacks. It is stood in (Franco, Künzler, Von der Assen, Feng, & Stiller, 2024) that cybersecurity costs cannot be assumed to follow a normal distribution due to their intrinsic characteristics and behavior. Specifically, these costs are never zero and typically exhibit a heavy-tailed distribution, reflecting the potential for extreme financial losses. Since RVaR is historically data-driven, it inherits both the strengths and limitations of financial risk models. While it can be effectively employed in cybersecurity investment assessments, its reliance on past data makes it less suitable for quantifying the risk associated with rare but high-impact events. Moreover, it fails to account for the fact that an incident that has already occurred may have a nonzero probability of recurring, making it a limited predictor of emerging cyber threats. A widely recognized challenge in this domain is the underreporting of cyberattacks by corporations – the 'ostrich reaction' was already mentioned, which compromises the reliability of datasets used to develop risk assessment models. At the core of this issue lies a misalignment of incentives - corporations are often reluctant to disclose security breaches due to concerns that doing so could negatively affect their ability to raise capital by increasing perceived risk among investors. Another noteworthy aspect is the approach taken to classify corporations based on time and size scalers in cybersecurity risk assessments. For the factor scaler, the methodology employs a weighted sum that aggregates all available reports. Additionally, it incorporates rankings across different factors based on their relative influence, ensuring a more nuanced evaluation of cybersecurity risks across diverse corporate structures.

Cyber KPI for Return on Security Investment

The article (Onwubiko & Onwubiko, Cyber KPI for Return on Security Investment) is dated back in 2019, but its discussion on the challenges of measuring the actual cost of a cybersecurity incident remains highly relevant. We analyze the two distinct sets of cyber KPIs proposed for organizational and national ROSI calculations. By comparing their differences and similarities, we seek to understand the extent to which ROSI calculations differ across these two scenarios and how these differences impact the assessment process. To facilitate the discussion,

the terminology of graph theory is adopted, referring to nodes, branches, and their respective degrees and depths in relation to the central node of each schema - namely, the Organizational Cyber KPI and the National Cyber KPI. A key distinction between these two structures lies in the depth of their leaf nodes: in the case of organizational KPIs, the leaves have a depth of two, whereas for national cyber KPIs, the depth is only one. This difference arises from the authors' decision to further categorize each organizational metric into three subcategories, while national cyber KPIs were limited to a single level. The selection of national KPIs was based on a common ground detected from a review of national cybersecurity strategies in various countries. The comparison between National and Organizational Cyber KPIs reveals both similarities and differences in their structure and objectives. While organizational cyber KPIs are primarily focused on protecting business assets, infrastructure, and data, national cyber KPIs address broader societal and strategic objectives, such as national security, economic resilience, and citizen protection. Structurally, organizational KPIs tend to be more granular for the reasons we already mentioned. In contrast, national KPIs emphasize aspects like countering cybercrime, protecting critical national infrastructure, and fostering cybersecurity education. Despite these distinctions, both frameworks share common principles, such as the need for continuous monitoring, the importance of incident detection, and the emphasis on mitigating vulnerabilities. Ultimately, the differences in depth and scope between national and organizational cyber KPIs highlight the varying scales and priorities of cybersecurity investments in these two contexts. In qualitative approaches to ROSI, the assessments of an organization's cyber posture are typically conducted through structured questionnaires. As a result, there is a clear need for a tailored assessment artifact specifically designed for PAs. Relying on frameworks developed for private organizations would likely produce misleading results, as these models prioritize factors that may not be critical for public sector cybersecurity while overlooking key elements essential for safeguarding governmental and societal digital infrastructures.

Concept Matrix

Following the guidelines outlined by (Webster & Watson, 2002), a concept matrix Table 4 is employed to synthesize the key contributions identified in the literature regarding adaptations of the ROSI model. In the subsequent analysis, we discuss the frequency and prominence of the recurring themes observed across the reviewed studies.

A. Definition of ROSI (or cyber ROI)
B. Issues in calculating ROSI
C. Economic benefits evaluations in cybersecurity
D. Threat assessment and prioritization and identification
E. Cost of investment in cybersecurity
F. Alternatives to ROSI

G. Monetary cost of breaches

H. Applicability of ROSI in public sector

	A	B	C	D	E	F	G	H
Security Attribute Evaluation Model: a Cost-Benefit Approach (Butler, 2002)	x	x		x	x	x		
Evaluating Information System Investments from Attackers' Perspective: the Return On Attackers (ROA) (Cremonini & Martini, 2005)	x	x			x	x		
A System to Calculate Cyber Value-at-Risk (Erolaa, et al., 2021)					x	x		
Security Investment and Information Sharing under an Alternative Security Breach Probability Function (Gao, Zhong, & Mei, 2013)					x	x		
Framework for Calculating Return on Security Investment for Security-Oriented Organizations (Arshad, Abbas, Faisal Amjad, Shafqat, & Yaqoob, 2019)	x	x	x	x				x
Cyber KPI for Return On Security Investment (Onwubiko & Onwubiko, Cyber KPI for Return on Security Investment)	x	x	x	x				
A Comprehensive Risk Management Framework for Approaching the Return on Security Investment (ROSI) (Pontes, Guelfi, Silva, & Kofuji, 2011)				x	x			x
Information Security Assessment in Public Administration (Szczepaniuk, Szczepaniuk, Rokicki, & Klepacki, 2019)				x	x			
Estimating Benefits from Investing in Software Development (Arora, Telang, & Frank, 2007)			x		x			
Return on Security Investment (ROSI) – A Practical Quantitative Model (Sonnenreich, Albanese, & Stout, 2006)	x	x		x		x		
Cybersecurity Investments Metrics using FAIR-ROSI (He, Xin, & Luo, 2024)		x		x		x	x	

Table 4 - Concept Matrix

We acknowledge that it is straightforward to observe that concept G, which pertains to the monetary cost of breaches, is the most underexplored aspect in the existing literature. Similarly, concept H, which concerns the applicability of ROSI in the public sector, has not been widely investigated. We emphasize this not as a justification for positioning the research within this specific domain- since we have already provided a detailed rationale for this Action Design Research (ADR) approach in the first chapter - but rather to highlight the extent to which the field of ROSI adaptations remains fragmented and presents significant gaps. These gaps indicate a

need for further empirical studies that systematically address the economic implications of security breaches and evaluate the applicability of ROSI frameworks beyond the private sector, particularly in public institutions where security investments follow distinct decision-making processes and budgetary constraints. Considering the opposite perspective, an examination of the concepts that are most frequently addressed in the literature is reviewed. Concept B is unsurprising, as nearly all sources discussing ROSI promptly acknowledge its limitations and challenges. Furthermore, we argue that Concept D is particularly significant due to its implications for risk management and investment strategies. Finally, we highlight the relevance of Concept E, given the inherent complexity in estimating and assessing the costs associated with cybersecurity investments.

Chapter 3 - Methodology

This brief chapter outlines the research methodology, structured into three sections that reflect the distinctive integration that characterizes this study. Given the complexity of the problem addressed, the flexibility sought in the design of the artefact was mirrored in the researchers' methodological approach. Various strategies were explored and adapted throughout the research process to align with the available resources and the evolving demands of different phases and dimensions of the study.

Action Design Research

As soon as the topic for the Master's Thesis was chosen, the Action Design Research (ADR) approach (Sein, Henfridsson, Purao, Rossi, & Lindgren, 2011) was considered convenient and wise to structure the work. Action design research (ADR) has been proposed as a tool for conducting an engaged form of research for advancing theory while producing useful knowledge (Spagnoletti, Resca, & Sæbø, 14). It is a research method that aims to generate prescriptive design knowledge through the iterative process of building, intervening, and evaluating IT artifacts within their organizational context. It is characterized by four distinct stages, each underpinned by specific principles, that are summarized in Figure 12 which is taken from the original paper.

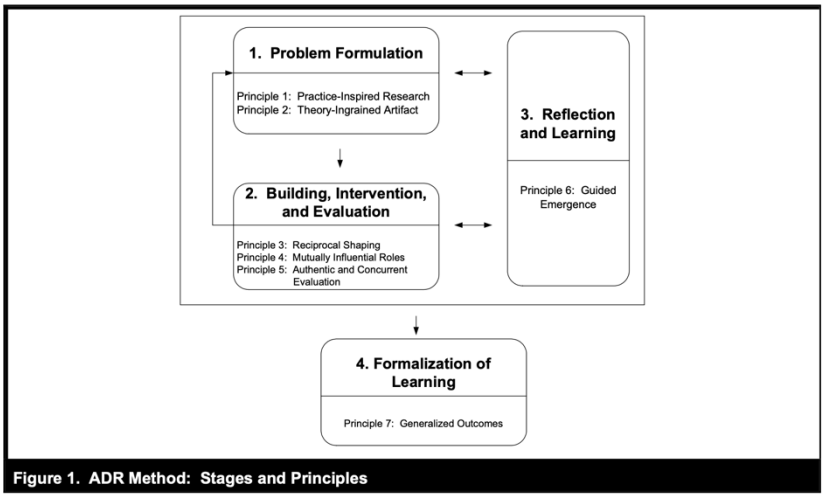

Figure 12 - Action Design Research

The process begins with Problem Formulation, which is guided by the principles of Practice-Inspired Research and Theory-Ingrained Artifact, emphasizing the importance of addressing real-world problems with theoretically informed artifacts. In the case study, the theoretical foundation was initially developed through an extensive literature review, while the practical problem emerged directly from the field, specifically at the Italian National Cybersecurity Agency. A dialogue between academia and the domain field was initiated concerning the need to evaluate the PNRR 1.5 investments: this constituted the context in which the problem was properly formulated, serving as a bridge between theory and practice. The second stage, Building, Intervention, and Evaluation, is where the artifact is iteratively developed and implemented, shaped by the principles of Reciprocal Shaping,

Mutually Influential Roles, and Authentic and Concurrent Evaluation. This stage stresses both the interactive dynamic of this kind of research, supported by evidence from the case study and input from domain experts, and the fact that the development of the artifact is validated and refined through iterative interactions. Instead of reporting again a schema from the original paper, an adaptation of the BIE phase that is specifically naming the parts for this case study follows, for further clarifying the roles.

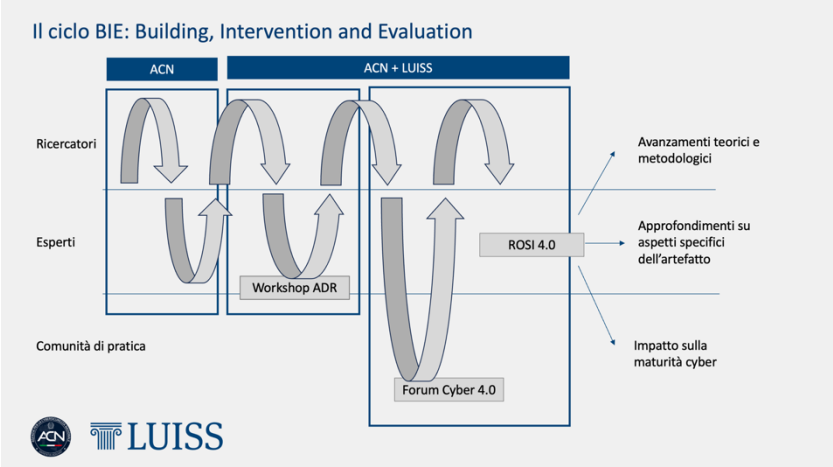

Figure 13- BIE Cycle in this ADR

The first iteration sees ACN as the main actor, and this is because, prior to engaging with academia, the Agency conducted a desk-based investigation and developed a pilot version of the artefact. Although this process remained internal during the first iteration, it still included expert input, as the team involved was hybrid in nature and personnel from various technical departments were consulted on specific topics to advance the artefact. Once this iteration was concluded, Luiss formally entered the ADR process. In the second iteration, the highest level of expert consultation was reached during the workshop held in May 2025, which is presented and discussed in the dedicated chapter. On that occasion, a diverse group of domain experts was assembled and consulted on key topics and issues aimed at improving the artefact developed in the interim, and to validate the relevance and usefulness of the ongoing research. The third iteration of the BIE phase involved a third stakeholder, which is the Competence Center Cyber 4.0. While the workshop was hosted at their premises, the most significant event in this regard is yet to come. In early June 2025, this ADR will be presented to a broader audience at the annual Forum Cyber 4.0, providing an opportunity to receive feedback from the professional community. The artefact is continuously refined and potentially enhanced throughout all iterations. In the end, each of the clustered stakeholders namely researchers, experts and community of practice, benefits from specific returns generated during this BIE phase. From the point of view of the researchers, technological and theoretical advances were achieved. At the same time, the experts got some deep dives about specific aspects of the artefact; even if the artefact itself is not available for an extent use, the research that was developed around its factors generated independent valuable insights for the field. Last but not least, the community of practice could, in the near future, rely on tools similar to the artefact or be generally inspired by its data-driven nature. Moving forward, the third stage of ADR,

Reflection and Learning, is guided by the principle of Guided Emergence, which emphasizes understanding how the artefact evolves through intervention. This should be sufficiently clear after the extensive discussion in the previous phase. Finally, the Formalization of Learning stage, driven by the principle of Generalized Outcomes, aims to produce design knowledge that can be applied beyond the specific context of the study. In this regard, a few words are necessary. Both the workshop and the forum serve as opportunities to present the artefact to an audience primarily composed of domain experts from the field, particularly from industrial contexts rather than public administrations. Although the artefact was developed to be ready for the use of data from public authorities, the research design itself is context-independent, and the intention to provide Italian companies with a strategic investment guidance tool was present from the very beginning of the project. Therefore, the generalization phase is of high importance within the overall process, as one of its key objectives is to enable the artefact's application beyond the initial scope of the National Cybersecurity Agency.

Analytic Hierarchy Process (AHP) Technique

Leveraging the collaborative networks of both the University and the National Cybersecurity Agency, the concept of a workshop emerged as a suitable platform for discourse and the generation of valuable qualitative and quantitative rankings to be incorporated into the model. The workshop represents an exceptional opportunity to bring together experts from different fields within the framework of this research. Participants from both private and public organizations shared information and perspectives on two main issues, guided by the Analytic Hierarchy Process (AHP) technique. The choice of AHP was inspired by the work presented in "A multi-criteria model for the security assessment of large-infrastructure construction sites" (Oliva, Faramondi, Setola, Tesei, & Zio, 2021), which illustrates a method for translating qualitative domain knowledge into quantitative values, an insight particularly relevant for enriching the artefact. Now, specific details about the workshop are provided.

Due to confidentiality, the names of the organizations cannot be disclosed; however, classification by sector and by monetary income, rather than by number of employees, can be presented. The range of the domain experts varies both in terms of the contexts they come from and the expertise for which they are referenced. The strength of this workshop lies in the fact that academia, as well as private and public organizations from various sectors, were represented at the table. This composition enables the creation of a suitable setting for generating generalizable knowledge and for discussing differing perceptions and viewpoints grounded in the participants' respective experiences. Some of the professionals participating in the workshop are active operatives in the field, such as certain Chief Information Security Officers rather than Heads of Cybersecurity. Others are engaged in cybersecurity from a legal perspective, including professors of Luiss Business School specializing in cybercrime, digital transformation, cyber defense, and risk management. Since insurance products could be tailored based on the artefact's outcomes, banks were also involved in the discussion. Beyond the diversity of domain experts, the

range of industrial organizations involved is noteworthy as well, as stressed by Table 5. The sectors represented by the organizations are notably diverse, encompassing banking and finance, education, energy, information technology, and shipbuilding. In addition to these, seven persons from the National Cybersecurity Agency were there as well. A final note: the expert panel comprises representatives from both academia and private corporations, although the artefact is designed primarily for public administration. This choice is not only influenced by expert availability but also aligns with the modular nature of the research approach, which supports future extension and generalization. Again, the workshop brings together a diverse group of experts to capture a broad range of issues and vulnerabilities within cybersecurity systems. While variables specific to public administration are considered in the artefact currently in use, the group is not dependent on them, thereby ensuring that the workshop's outcomes remain broad and adaptable.

Participant Cluster	Role/Institution	Sector	Organizational Size	Perspective Offered
Corporate ICT Experts	Chief Information Security Officers and cybersecurity managers	ICT / Cloud / Digital Services	Medium to large enterprises	Operational cybersecurity and infrastructure management
Infrastructure & Utility Leaders	Directors of Security and Cyber Defense	Energy / Critical Infrastructure	Large corporations	Systemic risk management and protection of essential services
Public Sector & Government	Heads of cybersecurity in public institutions	Public Administration / ICT	Large government agency	Institutional cybersecurity and data protection
Academic Experts	University professors with expertise in finance, law, and cybercrime	Academia / Finance / Legal Studies		Risk analysis, regulatory frameworks, and legal implications
Strategic Industry Representatives	Cybersecurity leads in industrial and manufacturing sectors	Defense / Shipbuilding / Industry	Very large enterprise (over 10,000 employees)	Protection of industrial critical infrastructures

Table 5 - Partecipants' Clusters

The workshop was held in Rome, at the Competence Center Cyber 4.0, Italy's National Cybersecurity Competence Center. A brief introduction about the Center to make the context clear: it was established by the Ministry of Enterprises and Made in Italy in 2019 and operationalized in 2021. Headquartered in Rome, the Center operates as a public-private association, drawing strength from a wide-ranging constituency that includes Italian universities, leading cybersecurity firms, public research organizations, specialized SMEs and government institutions. This ecosystem enables Cyber 4.0 to act as a neutral, multi-stakeholder platform that bridges academia, industry, and public policy. Its mission is to strengthen Italy's cybersecurity posture through fostering innovation, developing training, providing advisory, co-funding applied research and technological transfer, and implementing capacity building initiatives at national and international level, with a commitment to sharing knowledge and practices across borders. In this context, the center is actively involved in EU-funded programs (e.g., Horizon Europe, Digital Europe), national strategic projects, and emerging international collaborations aimed at building sustainable cybersecurity capabilities. With this overview, it should be clear that the Competence Center 4.0 matched the ideal setting for the aim of the workshop. During the event and for its preparation, a significant contribution was also provided by 'Università Campus Bio-Medico di Roma', a recognized academic and research institution. In particular, it was responsible for managing the AHP technique during the event, as experts in this methodology were specifically dispatched by the institution to support the process.

Once the idea of the workshop was approved by the National Cybersecurity Agency, the approach to guide the workshop needed to be established. Given the objectives of the workshop and the key informants that were targeted to participate, the method for a multi-criteria analysis was immediately proposed and subsequently selected. The Analytic Hierarchy Process (AHP) is a widely used decision-making framework that systematically addresses complex problems by structuring them into a hierarchy, where preferences are elicited through questionnaires involving pairwise comparisons of alternatives, and these relative preferences are then mathematically transformed into absolute weights or priorities. Participants are asked to respond anonymously to multiple rounds of questionnaires - two in this case - and a facilitator provides a synthesized summary of the group's feedback after each round to moderate a discussion around the results. The Analytic Hierarchy Process (AHP) offers significant advantages for complex decision-making scenarios like this one. Specifically, it leads to an improvement of results accuracy by systematically structuring the problem and allowing for the quantitative evaluation of subjective judgments through pairwise comparisons. The method also provides flexibility in evaluating relative preferences, enabling decision-makers to weigh criteria and alternatives according to their specific importance and how they contribute to the overall goal. Furthermore, a crucial strength of AHP is the verification of consistency of preferences assigned by experts, which helps to identify and mitigate inconsistencies in judgments, thereby increasing the reliability and validity of the final decision.

In the case of the workshop held in May 2025, there were thirteen people answering the questionnaires. A first round of the technique, because of the settings of the software in use, was a question regarding general feedback of the presented framework. However, in subsequent phases, the questions are closed, as they aim to guide the analysis and evaluation of the themes in the most quantitative manner possible. The results from each iteration are shared with the experts. This facilitates a process of mediated interaction between individuals and the group throughout the workshop. The technique preserves the anonymity of individual contributions - an essential feature, as it encourages the disclosure of more sensitive information that experts might otherwise be reluctant to share, given the vulnerability associated with cybersecurity topics. Among its various advantages, this method is particularly well-suited for evaluating the weighting of information in decision-making processes. It is also valuable in contexts where differing perspectives may not directly conflict but could nonetheless offer significantly divergent and contrasting insights. For example, in this specific case a Chief Information Security Officer (CISO) and a legal expert may approach the evaluation of a cyberattacks' impact on an organization from fundamentally different perspectives.

Coding

In parallel with the development of the ADR, the artefact was implemented in Python and subsequently integrated with the outcomes of the AHP technique. Due to the confidential nature of the data, the source code cannot be explicitly disclosed in this thesis. However, it is appropriate to describe the logical structures and procedural steps adopted, as well as the strategies employed to address the challenge of data unavailability within the required timeframe. Again, the development of the code for the artefact progressed in parallel with the preparation of the workshop as well. The pilot project was informed by the proposed framework and incorporated outputs from both the workshop and the forum as structured. Nevertheless, it aimed to improve upon the initial model by experimenting with methods that, although ultimately discarded in later iterations, represented a more advanced approach compared to earlier versions. The code was entirely developed in Python and organized within a Jupyter Notebook using Visual Studio Code. Python was particularly well-suited for this project for several reasons. First, its extensive use of functions aligns closely with the modular design of the artefact, which was conceived and refined through multiple iterations. Second, Python's ability to incorporate interactive elements - such as through the use of functions like *input()* enabled the development of a preliminary version of the interactive features envisioned for the final tool. Additionally, Python's robust capabilities for data visualization proved especially advantageous, supporting the strategic management goals underlying the long-term applications of the artefact. The notebook was organized in seven sections, that are: SetUp/Libraries, Preprocessing, Functions, Interactive Calculation of ROSI, Generation of ROSI Datasets, Visualizations, and Workshop's Graphics. The first two regarded the setting of the python environment and the adjustments that the data needed to be used in the model appropriately, comprehending the data cleaning. In the third section, all the functions are written, so that the

readiness of the code is smooth. For the following two sections, the content is kind of repeated but the way it is presented makes the difference. Running the first, the user could get a ROSI and the related insight via a step by step question and answer about the organization considered, based on the iterative use of the *input()* function. Then, a section is dedicated instead to the generation of datasets comparing ROSI of different organizations rather that related to distinct interventions, and so forth. This one is also fundamental to enable the final section of visualization and graphics. That is where the side of ROSI that is most useful from a managerial point of view comes out: some easily-understandable barplots make possible to grasp take-aways and insights from the ROSI calculations. Among the libraries utilized, Pandas (*import pandas as pd*) was essential for data manipulation and tabular processing, while Matplotlib (*import matplotlib.pyplot as plt*) and Seaborn (*import seaborn as sns*) were used to generate static visualizations to support the analysis. Additional libraries such as NetworkX and Plotly were considered (but not actively used in the baseline), offering potential for advanced graph processing, interactive visualization, and numerical operations, respectively. The *re* library provided support for regular expressions and was essential for pattern matching and text parsing tasks throughout the data processing phase. Finally, the warnings module was used to manage warning messages - such as suppressing or displaying them - to keep the program's output clean or handle deprecated features gracefully.

Now, an extract of the functions used in the implementation is described. For some of them, additional comments are provided to clarify the reasoning behind specific logical choices in the code. Overall, the use of functions proves highly advantageous for the interactive nature of this ADR. In fact, for future developments of the artefact, the same functions currently in use can remain relevant, allowing the work to focus more on adaptation rather than on rethinking the entire structure.

- The function *ask_for_area()* is designed to interactively prompt the user to select an area of intervention from a predefined list, which is specifically the following and it is taken from the official Call: ['Analisi della postura di sicurezza e piano di potenziamento', 'Miglioramento dei processi e dell'organizzazione di gestione della cybersecurity', 'Miglioramento della consapevolezza delle persone']. It copies the list of available areas and displays up to three options at a time, allowing the user to choose by typing 1, 2, or 3. The function includes basic input validation to ensure a correct selection is made and returns the selected area once a valid choice is confirmed. This interactive approach reflects the user-oriented nature of the baseline model, leveraging Python's input() function to guide decision-making in a simple and accessible way.
- The function *mappa_interventi_dimensioni(area)* defines a static mapping between selected intervention areas and the corresponding cybersecurity dimensions they influence. Given an input area, the function returns a list of related dimensions based on a predefined dictionary. If the input area is not found in the mapping, it returns an empty list. This function supports logical consistency in the

model by linking high-level interventions to specific operational domains within the cybersecurity framework.

- The function *interventi_per_area(org, area)* is designed to extract specific intervention data for a given organization and area of intervention from the investments DataFrame. It filters the rows by matching the organization and area criteria using two conditions: the column which identifies the organization and the column which specifies the area. The function then returns the values from the column corresponding to the filtered rows, using the *tolist()* method to convert the result into a list.
- The *cost_attack_org(org)* function estimates the monetary impact of a cyberattack based on general contextual characteristics related to the targeted organization. The function prompts the user to answer a series of binary (yes/no) questions concerning the geographical location of the attack, namely if it is targeting an Italian organization or not, the sector in which it occurred (public or private), and how the attack was identified or disclosed. Each response incrementally adjusts the estimated cost based on predefined values taken from the (Security, 2024). The final cost reflects the highest applicable scenario, ensuring that the estimation captures the most significant financial consequence among the considered factors. This function is interactive and relies on user input, simulating an expert-informed cost attribution process.
- The *match_drivers_call_and_IBM(cost_df, investments, org)* function performs a customized filtering of a cybersecurity cost dataset built on the top of the IBM 2023 report (Security, 2024), by aligning it with the specific technological and organizational characteristics of a selected organization. The function dynamically identifies which cost-related drivers of 'Avviso 2' are applicable to the organization in question and classifies them based on defined thresholds (e.g., number of applications, development practices, security systems). The cost types that are not specifically entailing the case of the selected organization are excluded from the analysis to better reflect its actual risk and mitigation profile. This tailored filtering supports a more precise estimation of cyberattack-related costs by integrating structural and behavioral indicators into the cost attribution logic.
- The function *aggiustamento()* is designed to apply an adjustment factor specifically to the cost data reported in 2023, in light of 2024 updates. It prompts the user to confirm whether a standard 10% cost adjustment is acceptable; if so, it returns a clarifying note. Otherwise, the user can input a custom percentage, and the function returns the corresponding multiplier to be applied to the original price. This mechanism ensures flexibility and traceability over time.
- The functions *cp_aritmetica(cp_dim)* and *cp_pesata(cp_dim)* are mutually exclusive. They get as input the list of cyber postures for the six dimentions, and they return the total cyber posture, which is evaluated using throught the arithmetical mean in the first case and a weighted mean in the second.

- The functions *eff_lineare(cp)* and *eff_workshop (cp)* are mutually exclusive. They get as input the total cyber posture of the organization, and they return the effectiveness that is related to that level, the first assuming a linear relation between the two quantities and the second one relying on the interpolated function that was got as output from the second question of the workshop.

The data used in this case study is confidential. For this reason, a Non-Disclosure Agreement (NDA) was required for the candidate to access the data. Consequently, it is not possible to show specific examples from the datasets in this thesis; however, their structure will be explained. All datasets were initially provided in .xlsx format and were then exported sheet by sheet into .csv files. Not all data at disposal was made available to the candidate, for confidentiality reasons. This introduced complexity, as more assumptions had to be accepted in the model.

Chapter 4 – The Root of Cost of Investment

Case Study

An initial overview of the cost of investment's factor is provided through a specific case study in this section, but a framework to generalize its quantification is then discussed. It is highly useful to remind the context in which the Agency is working before starting the overview, which concerns some PNRR 1.5 funds. Organizations and Public Administrations could apply to some Calls, that are essential instruments for allocating financial resources transparently and efficiently. The eligible institutions differ among the various Calls. For the case study of this research, two of them are of particular relevance, respectively the second and the seventh. Once the selection process for determining the winners among the applicants for each Avviso is completed, the grant amount is based on the interventions that the selected beneficiaries have declared for implementation within their respective systems. For each winner of the PNRR funding, the related amounts for the selected interventions can be taken as the cost of intervention as input for the ROSI calculation. A theoretical pipeline to level up this process is discussed in this chapter, with the intention to make it compatible with other factors of the artefact. The exact volume to allocate for each intervention available in the Call is decided in advance appropriately, ideally with a bottom-up approach. Before delving into the specific details of individual cases, it is valuable to first provide additional information regarding the potential eligibility criteria and general structure of the Calls. Each Avviso may target either central or local public administrations and can be classified into two categories: "a titolarità" or "a regia". In the first case, the Agency is directly responsible for implementing the interventions requested by the PA and subsequently approved and financed by the ACN. Conversely, in the latter case, the ACN solely provides the funding, while the PA retains full autonomy in executing the interventions. Table 6[3] summarizing the category to which each Avviso belongs to is provided below.

	Title	Period		Central PA	Local PA	A titolarità	A regia	Grant (euro)
Avviso 1	Interventi di potenziamento della resilienza cyber - PA Centrale	03/03/22	07/04/22	X			X	15.100.000,00
Avviso 2	Interventi di potenziamento della	03/03/22	23/03/22	X		X		7.800.000,00

[3]Data from https://www.acn.gov.it/portale/documents/d/guest/acn_relazione_annuale_al_parlamento_2024

	resilienza cyber - PA Centrale							
Avviso 3	Interventi di potenziamento della resilienza cyber - PA Locale	02/08/22	17/10/22		X		X	63.700.000,00[4]
Avviso 4	Interventi di potenziamento delle capacità di analisi e scrutinio software nella PA Centrale	05/09/22	23/09/22	X			X	900.000,00[5]
Avviso 5	Attivazione di laboratori di prova per l'area di accreditamento Software e Network	20/11/22	30/11/22				X	3.000.000,00
Avviso 6	Attivazione e potenziamento CSIRT regionali	11/08/23	25/09/23		X		X	28.500.000,00
Avviso 7	Interventi di potenziamento della resilienza cyber - PA Centrale	11/10/23	05/12/23	X		X		12.600.000,00
Avviso 8	Interventi di potenziamento della resilienza cyber - PA	17/06/24	12/07/24		X		X	108.000.000,00[6]

Table 6 - Calls of the National Cybersecurity Agency

As an exemplification, a particular analysis of 'Avviso 2'[7] is done. It is the call that has been highlighted in light grey in Table 6. The available interventions for this call are grouped into three clusters:

[4] Initially, the grant amounted to 45.000.000,00 euros. Additional funding has since been released for this Call.
[5] One PA only (MEF – Ministero dell'Economia e delle Finanze) received the grant for this call. For this reason, the initial budget has not been used fully.
[6] https://www.acn.gov.it/portale/documents/20119/88395/PNRR_Inv1.5_Avviso_2_InterventiCyber1.pdf/66b56c83-9e83-727c-af0b-ac7e4045da17?t=1704713847368

1. assessment of the cyber posture and a plan to enhance it,
2. upgrades in processes and management of cybersecurity,
3. situational awareness.

Each entity can ask for up to a fixed number of different activities to be implemented in its system, among all the ones proposed throughout the three areas. In the case of the second call, the PAs could ask for five interventions maximum. Each applicant had to literally check the intervention that it required and the related dimensional drivers on the piece of sheet they submitted in the application. In the following table, all the drivers that can be encountered in the second call are listed, with the number of classes that each one of them identified to allocate an appropriate budget consequently in case the PA would win the funding. Before reading the table, it is recommended to go through the following glossary, so that a complete understanding of the rows is available.

- Processes: structured and monitored sequence of activities that are thought to provide a specific output
- Application: a single or o a group of software running on a server
- Security systems: systems and platforms with security purposes
- Identity: identities that are managed by IAM (Identity and Access Management)
- Services (management of digital identities): IAM services available for employees
- Policies (management of digital identities): policies of access and profiling of systems based on elements such as organizational position/level
- Strategic documents: operative instruction define how an activity that is described in the current flow and how plans regarding the configurations to apply for specific systems are defined
- Programming languages: e.g., Java, Python, PHP
- Technologies: kind of platforms in use and hosting
- Applications on premise: installed applications that are managed by internal IT and servers
- Applications in cloud: installed applications that are managed by the use of cloud services such as IaaS, platform, and SaaS
- VPN, MLPS, SD-WAN, Dark Fiber: kind of connection to define and manage datacenters and decentralized offices

The two columns on the right in Table 7Table 7 - Drivers of the Call contain preliminary notes, which should be revised and expanded in the future using both secondary and selected primary sources. The aim is to establish a connection between the quantified information on attack-related costs and other ROSI factors, and the internal profile of the organization, as determined by the class of drivers to which it belongs.

	# classes		From IBM report 2023	From IBM report 2024
Number of processes	3			
Number of applications	3			

Number of security systems	3		Most common investment types among those increasing security investments following a data breach	
Identities	3		Cost of a data breach by head count, per-record cost of a data breach by type of record compromised	Type of data compromised by percentage (2023 against 2024)
Number of services	3		Cost of a data breach by head count	
Number of policies	3		Distribution over time of data breach costs in low-data versus high-data regulation environments	
Number of applications in the field	3			
Number of servers	3			
Number of clients	3			
Kind of output	2/3	*		
Number of programming languages that has been used	3			
Number of technologies	3			
Number of applications on premise	3		Cost of a data breach by storage location of breached data	Cost of a breach by storage location
Number of applications in cloud	3		Cost of a data breach by storage location of breached data, Migration to the cloud as cost amplifier for a data breach	Migration to the cloud as cost amplifier for a data breach, Cost of a breach by storage location
Number of VLAN	3			
VPN, MLPS, SD-WAN, Dark Fiber	3			

Number of web proxy users	3			
Number of training sessions	3		Employee training as cost mitigator for a data breach, most common investment types among those increasing security investments following a data breach	Employee training as cost mitigator for a data breach, cost of a data breach based on the level of security skills shortage

Table 7 - Drivers of the Call

As already mentioned, the investment for an intervention i could be easily defined in this case study as the amount of funding awarded to the winner of a call for proposals related to the requested intervention: $M(i) := grant(i)$. M is used as notation to avoid confusion with C, which is later used to denote the cost of an attack, and with I, which represents the set of interventions. Moreover, M, standing for money, serves as a reminder that this amount corresponds to the monetary investment. This amount $M(i)$ is based on the specific drivers selected by the winner in the application form. The response to each driver is associated with a budget, in the official call[8]; if the driver concerns the size of the PA and it stands in the biggest cluster, the highest budget is then associated, for example. This is done for each driver independently, so there is no differentiation ad priori in the budget given for all the possible commutations in the responses to the drivers. For clarity, if a public administration requests an assessment of its cybersecurity posture (this is indeed one of the interventions that can be requested), various combinations of drivers' can be given. For 'Avviso 2', for instance, the number of classes of budget for each intervention is equal to the maximum number of categories the drivers for that single intervention have at most. This means, for instance, that if an intervention has two drivers, one made of three categories and one made of two as in the example above with D_1 and D_2, there would be three possible budgets for both drivers, but for the second one D_2 the middle budget is not to be considered as an option.

d_1^1	High Budget - $M(d_1^1)$	High Budget - $M(d_2^1)$	d_2^1
d_1^2	Middle Budget - $M(d_1^2)$	Middle Budget	
d_1^3	Low Budget - $M(d_1^3,)$	Low Budget - $M(d_2^2)$	d_2^2

Basically, a team of experts could look at the two or more suggested budgets that are associated to the different drivers for a single intervention and decide whether to go for a mean among the values (Middle Budget and Low Budget in the case above), rather than giving the middle budget, or the lowest, or even the highest. This first practical approach can be implemented quite briefly. In first place, in case the system was run in Excel, it could be moved to Python, and a series of conditional loops could be designed ad priori, not only standardizing the

[8] https://www.acn.gov.it/portale/documents/20119/88395/PNRR_Inv1.5_Avviso_2_InterventiCyber1.pdf/66b56c83-9e83-727c-af0b-ac7e4045da17?t=1704713847368

estimation of budgets but also introducing a confidential interval. Ideally, each combination of drivers should have a predefined cost estimation for the corresponding scenario. Let D_1 and D_2 respectively be the first and the second driver for an activity that has been selected by a PA. If, for example, $|D_1| = 3$ and $|D_2| = 2$, then there should be $|D_1| * |D_2| = 6$ distinct budget for each of the possible combination of response for the two drivers that are identified by the following table.

		D_1		
		d_1^1	d_1^2	d_1^3
D_2	d_2^1	$M(d_1^1, d_2^1)$	$M(d_1^2, d_2^1)$	$M(d_1^3, d_2^1)$
	d_2^2	$M(d_1^1, d_2^2)$	$M(d_1^2, d_2^2)$	$M(d_1^3, d_2^2)$

For completeness: if there are two drivers for an activity and each of them has three possible classes, there should be nine potential budgets corresponding to each of the nine possible combinations derived from those two drivers. The first step towards enhancing the cost estimation system involves generating these case-specific budgets, keeping the cost of interventions still derived from the questionnaire responses. The scenario in which the highest budget is selected when the two drivers' responses indicate the same different budget levels as in the first table can still be validated and is clearly marked in green in the second table above. The distinction lies in the fact that the budget selection for each combination is predetermined and does not necessitate human intervention a posteriori.

Generalization

For the cost of the investment in the case study there is no need for estimation, since the exact value of the allocated budget should be readily available. This represents a peculiar case, which makes it suitable for an output for ACN, but also makes it difficult to apply the same setting in different contexts. Moreover, the potential use of the artefact for strategic comparisons among various sets of investments is limited by the absence of cost estimations. For instance, a board may decide to allocate a fixed amount of money to cybersecurity and need support for its management. On the other hand, a board may decide to proceed with a specific intervention and requires an estimated cost. This second scenario reinforces the value of this ADR not only for the ROSI output but also for enhancing the evaluation and refinement of its contributing factors. A structured and rigorous framework for defining and quantifying such cost of investments is discussed, both for future calls and for broader applicability of the tool in other contexts. To keep the description as concrete as possible, the following proposal addresses the definition of budgets for all combinations of drivers presented in the last table above. Currently, these budgets lack specificity and a clear process for their determination. The method outlined aims to introduce greater precision and coherence to this process. First of all, a consistent currency must be established to ensure uniformity in calculations. While seemingly trivial, this is fundamental. For the case study, as well as for the general purpose

of the research and its potential extensions, the Euro has been selected. There are two additional aspects to immediately address regarding costs. The first concerns the classification of interventions and, consequently, their associated costs. This classification turns to be particularly useful when analyzing insights derived from comparing different types of interventions rather than focusing solely on different public administrations. In other words, by specializing the ROSI both by public administration and by investment type, results can be grouped according to either dimension. If ROSI by type of investment is considered, pre-classification facilitates the identification of trends and further supports investment prioritization, beyond merely distinguishing between sectors within public administration. Then, a precise decision must be made regarding the extent to which associated costs should be included in the main expenditure when calculating the cost of each individual investment. This consideration arises because, following a single intervention, a series of subsequent updates is often required for integration into the existing system. Each of these updates may trigger further modifications, potentially leading to a cascade effect. Therefore, it is essential to define a threshold at which additional costs are capped in relation to the main expenditure. Following this overview, the proposal is presented.

Each cost is divided into three components: personnel, software, and hardware. The first category includes all expenses related to employee training, while the second and third are directly recorded as costs of tools. This tripartite classification applies not only to the main intervention but also to the associated costs that enable the intervention and those that arise consequently. The schema Figure 14 illustrates the full range of cost interactions across different interventions, from which a framework for evaluating the cost of the intervention under consideration is derived. The sketch is symmetrical when centered on the main intervention, with two levels of related expenditures extending backward and two forward. The first distinction concerns the costs to be included in the evaluation: greater weight is assigned to the left side of the flow, as it represents the expenditures necessary to enable the actualization of the investment is made. Within the green area, additional considerations for reducing the costs to be included remain. Starting from the top left, an orange block is marked by two outgoing arrows - one exiting the green diagram and the other remaining within it. The orange color indicates cases where the cost enables the system to benefit from the investment not only for the specific intervention under consideration but also for others. In such cases, it is intuitive to allocate the cost among all the investments it supports. Although this allocation may not be equally distributed, further refinements are left for future research extensions. As a preliminary approach, a simple division of the cost by the number of arrows exiting the block is suggested. Remaining on the left side of the schema, the red blocks also require explanation. These costs are still associated with interventions that enable the implementation of other interventions; however, they would still be incurred by the organization as maintenance costs for those other interventions. Whenever such a cost enters the schema, the entire left-side branch is disregarded. On the right side of the green area, the light blue color remains to be explained. This represents costs that consistently enable the use of multiple interventions. Similar to the orange

blocks, and acknowledging this as a limitation of the model, the simplest way to account for this cost overlap across different investments is to divide the cost of the block by the number of arrows leading into it and include only the resulting fraction as the cost, rather than the entire amount.

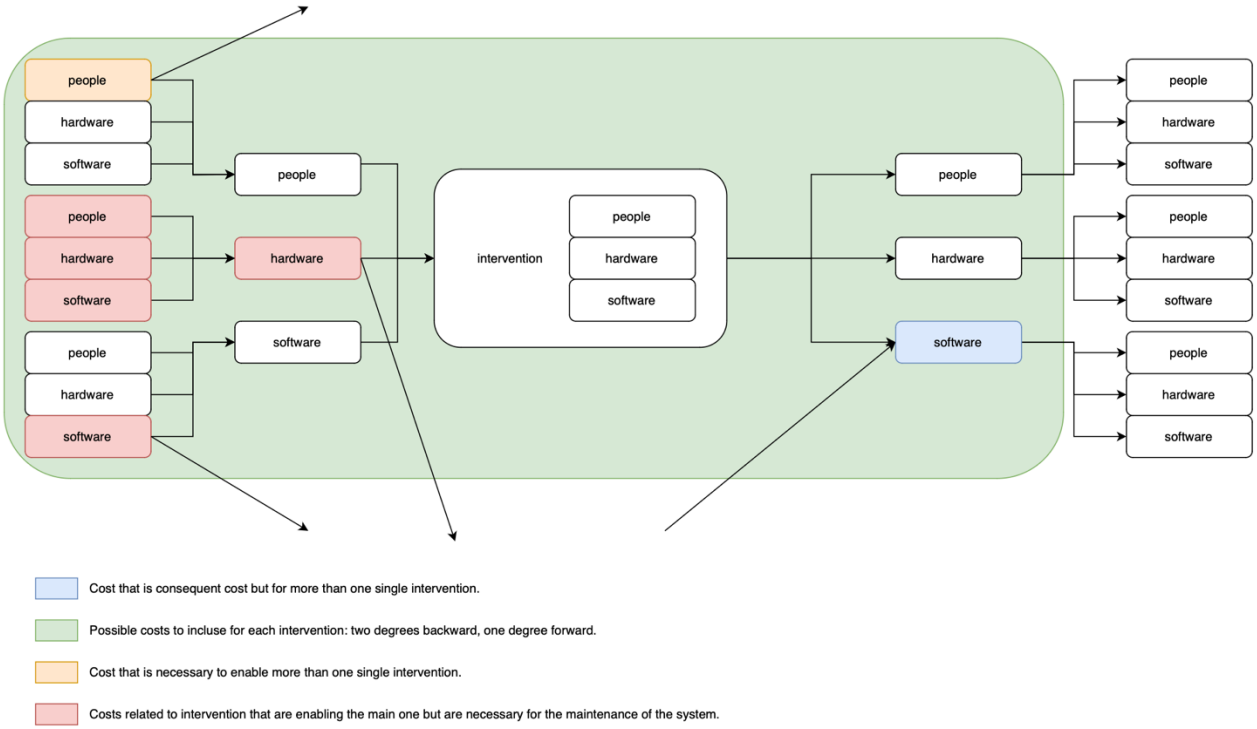

Figure 14 - Framework to calculate the cost of investment

Having seen this example in the graph above, it should be easier now to progress with the description and understanding of a pipeline to picture a similar schema for each activity. Once this is done, it is smooth to calculate via addition the cost of the intervention as defined, not having any kinds of interpretation about what to include. A more mathematical notation is now introduced to help in conceptualizing the workflow. Let $i \in I$ be an intervention, where I is the set of all the interventions made. For convenience, a clarification about the appendixes is briefly provided: I_j^k is the set of intervention for the public administration j (the public administrations are numbered appropriately), and it is impacting the area $k \in \{P, H, S\}$, where P is the set of investments regarding the personnel, H for the hardware and S for the software. M_j^k the related investment. Whenever an intervention (or an investment) is missing one of the two explicit specifications, this means that the entire set is considered. For instance, the investment $M_1$is the total investment for the first public administration of the list, considering the three areas personnel, hardware and software in their entirely. On the other side, $M^P$is the total investment entailing personnel among all the public administrations for which data is provided. It is then smooth to measure some quantities that are giving direct insights at managerial level, such as, given that there are J public administration in total (as for saying they are numbered 1, 2, …, J-1, J):

$$mean(M^P) = \frac{\sum_{j=1}^{J} M_j^P}{J}$$

is the average investment in personnel for the public administrations that have been considered. After this brief discussion on the importance of using a precise and rigorous notation, a clear outline is provided on how to construct and complete the schema as shown above. For each investment, there is a one-to-one correspondence with a single intervention, and vice versa. However, the two terms are not synonymous: an investment inherently includes the monetary amount spent, whereas for an intervention, it must be explicitly stated that its cost is the relevant factor when discussing financial matters. Otherwise, the term refers to the implementation of security measures or similar actions.

The main investment for which the cost is to be evaluated must be placed at the center of the diagram. A dichotomous analysis then divides the total cost of the main intervention into three categories: personnel, software, and hardware. These three categories are subsequently mirrored on both the left and right sides of the main block. On the left side, each of the three categories is further subdivided in a second iteration moving backward. For the left side of the block, representing the upstream flow of costs for the main investment, all additional tools or skills necessary for the implementation of the main intervention must be accounted for. However, any expenditures incurred before the intervention that facilitate its introduction into the system but are not strictly necessary should be excluded. This pipeline is repeated to fill out the second level of breaches on the left side of the main block as well. On the right side of the main block, the additional expenditures that follow the main investment to enable its benefits are represented. Any subsequent costs aimed at enhancing the use of the intervention but exceeding what is strictly necessary are not included. After this first round of designing the schema, the arrows must be added accordingly to the interactions among different investment for a single public administration. Disclaimer: although it can be considered the mean of cost for PAs as outlined above, the interactions that are considered are not going beyond a single public administration. Whenever an upstream intervention is necessary for the implementation of other main investments, an arrow is drawn from its block to outside the green rectangle. In the upstream flow, an arrow must be added whenever an expenditure is required to enable the implementation of other investments at the same subsequent level. Conversely, in the downstream flow, arrows enter the blocks when an additional expenditure plays a role in ensuring the proper functioning of the intervention. Summing up, while looking back at table that could be got from a first step in the implementation, that table should take the following updated format. Each budget is derived from a cost estimation encompassing three areas: personnel, hardware, and software, based on the scenario indicated by the responses to the drivers. While this approach necessitates greater effort initially, it provides a robust foundation for standardizing budget definitions. Moreover, it generates valuable insights through the potential grouping of costs according to their nature.

		D_1		
		d_1^1	d_1^2	d_1^3
D_2	d_2^1	$\sum_{k\in\{P,H,S\}} M^k(d_1^1, d_2^1)$	$\sum_{k\in\{P,H,S\}} M^k(d_1^2, d_2^1)$	$\sum_{k\in\{P,H,S\}} M^k(d_1^3, d_2^1)$
	d_2^2	$\sum_{k\in\{P,H,S\}} M^k(d_1^1, d_2^2)$	$\sum_{k\in\{P,H,S\}} M^k(d_1^2, d_2^2)$	$\sum_{k\in\{P,H,S\}} M^k(d_1^3, d_2^2)$

Finally, facilitates the seamless introduction of an additional modular implementation, which involves incorporating a confidence interval into budget calculations. This range can still be established using a bottom-up approach; for instance, it may be determined at the management level that all budgets should adhere to a range of $\pm 5\%$. Finally, as the concluding point for this root of the tree, attention is given to the importance of broadening the evaluation of investment costs to include both operational expenditures (OPEX) and capital expenditures (CAPEX). Once these elements are incorporated into the model, a broader discussion will naturally follow - particularly concerning the appropriate time horizon for analysis and the selection of a suitable discount rate.

Chapter 5 – The Root of Economic Benefit

It is now appropriate to turn to another foundational element, or root, of the framework: the assessment of the economic benefits associated with the investment for which the ROSI is calculated. As will be further clarified, this component holds significant value in enabling various applications and generating strategic management insights. This root is further subdivided into three sub-components, which become evident when examining the structure of its formula, composed of the following three factors:

$$EB = P_A * Eff * C_A,$$

in which P_A is the probability on an attack, Eff the effectiveness associated to the organization's cyber posture, and C_A the cost of attack. One of the strengths of this framework lies in its architecture, which is designed to ensure that all three sub-roots are context-specific and tailored to the profile of each organization.

Probability of an Attack

Definition

The probability of an attack is intuitively influenced by various factors, including the type of attack and the industry of the targeted organization. The definition of probability of an attack can be empirically defined as:

$$P_A = \frac{\#\ successful\ attacks}{\#\ attacks}.$$

A long-term recommendation is to refine this definition by retaining an empirically derived probability while integrating theoretical approaches - such as threat probability functions - to lessen dependence on historical data alone and to avoid overly restrictive assumptions, using instead a value tailored to the specific characteristics of the organization in question. However, given the availability of national data (CSIRT) on the frequency of attacks categorized by type and other relevant characteristics, an initial step is proposed to enhance the current framework without introducing a more complex analytical pipeline at this stage. To decide which value to use for the probability of attack, the objective of this ADR is to ensure that the model remains compatible with advanced quantitative tools for future integration. A key challenge in this regard lies in maintaining coherence between the variables underpinning future extensions - particularly those upon which P_A will rely - and the algorithm that adapts the artefact to the contextual information of the organization in the given scenario.

For a first revision of the P_A definition, two key aspects must be balanced as a primary concern. Public administrations operate in different sectors and are therefore exposed to threats in varying ways, both in terms of attack types and frequency. At the same time, a comparative analysis of the Return on Security Investment (ROSI) is crucial for deriving meaningful insights. To ensure consistency, a standardized selection of attack types must be established. A potential solution that addresses both aspects is to adopt a bottom-up approach. This would begin

with identifying and listing the set of attacks to be considered but giving different weights for different kind of attacks depending on the PA, for the calculation of the overall attack probability. These weightings should be tailored to reflect the unique threat landscape of each individual PA. For clarity, let $PA_1$and PA_2be two distinct public administrations for which the probability of attack must be calculated, respectively $P_A(PA_1)$ and $P_A(PA_2)$. Then:

$$P_A(PA_1) = \sum_{a \in A} P_a(PA_1) w_a^1,$$

$$P_A(PA_2) = \sum_{a \in A} P_a(PA_2) w_a^2,$$

where $w_a^1$is the weight (as a percentage) for the attach a for the first public administration, while w_a^2 is the weight the attach a for the second public administration. Enabling this method has another potential return and releases eventual possibilities to modularly extend the accuracy of the methodology. Two advantages come into mind immediately. Firstly, the integration of further information that might affect the weights for the different administrations is easily perceivable having this structure. Moreover, even if internal information is not available from the first day, an average coming from external report can be used temporaneous until better data is collected. This is a practical instance of the way the concept of modularity is ingrained in the model. The second advantage lies in the fact that an interest in investigating the probability of the attacks grouped by kind could naturally arise, and this makes this methodology useful even outside its specific context. For example, being J the set of indexes for the public administrations,

$$P_a = \frac{\sum_{j \in J} P_a(PA_j)}{|J|}$$

gives the average probability of attach for the attacks of kind a among all the different PAs. This could be also insightful when the average is calculated not for the whole set of public administrations but for different identified clusters and compared. A discussion should be driven for the definition of the weights for each PA, and a method to periodically update those is required as well to be outlined. However, this is still one more module that can be delivered in a second iteration of the artifact, and an arithmetic mean can be used for now. To wrap up: the only request would be having the set A with a selection of attacks' type (e.g., top 5 kind of attacks reported by official report from secondary sources) and then for an initial phase the following would be assumed (keep in mind the case below is for PA_1 as for instance):

$$w_{a_1}^1 = w_{a_2}^1 = \cdots = w_{a_n}^1 = \frac{1}{n}.$$

Practically, there are already some secondary sources that are available and useful for this module of implementation. For the kinds of attack to consider in the formula above for the case of public administration (extendible to a generic organization), another strict pipeline is needed to decide the quantity and the selection of them for the assessment of the cyber posture and for calculating the ROSI, eventually. It is important to underline

that further reliable high-quality data is going to be integrated in the model, at least for the case-study of Italian administrations. Indeed, a trend that goes beyond national borders regards the availability of official reports for the kinds of attack that are main characters of the threat landscape. The recommended approach in this regard is to integrate this specific data to fully leverage its potential. Threats will be categorized based on the type of attack and the sector they have impacted. This highlights the importance of developing a strategy that positions the organization within a framework of indicators, which are then crucial for assessing the threat landscape it interacts with. In this context, two main approaches can be considered:

- The probable attacks targeting the organization's sector can be used to define its threat landscape, serving as a basis for evaluating its cyber posture. This means that, a priori, only a few specifications about the organization are required to clearly identify which attacks should be collected from the report, based on the sector to which the institution belongs at that moment. This approach assumes that the report provides data on attacks in a manner that is consistently aligned with the relevant sector, in addition to being categorized by the type of attack.
- Another approach involves identifying the organization's vulnerabilities without relying on historical attack trends for its sector. In this case, only the attacks that specifically exploit these vulnerabilities are considered. This method has the advantage of being implementable from zero level, even though it requires a rigid structure for the identification of the vulnerabilities which must be developed. It is not excluding some assessment focused on the sector of organization.

The first approach is less demanding in terms of preparation. While defining threats solely based on the sector may be considered restrictive rather than merely reductive - and potentially risky in terms of assessment accuracy - this method does not rely on a pre-assessment that could lead to overfitting or the loss of contextual information. On the other hand, the second approach ensures a more robust assessment, as it is not dependent on data but rather on domain knowledge regarding how to evaluate system vulnerabilities. A balance between the two approaches can be achieved: using the second method as a foundation, guidelines can be established for prioritizing vulnerabilities based on the sector in which the organization is presumed to operate. Moreover, given the significant contribution of the literature in supporting the research objectives, particular emphasis is placed on defining a fully secure information system.

In the specific case of PAs, a significant body of research has emerged from the Polish context, offering valuable insights into the assessment and management of information security risks within governmental institutions. That study is explained in the article *Information Security Assessment in Public Administration* (Szczepaniuk, Szczepaniuk, Rokicki, & Klepacki, 2019), in which the implementation of Information Security Management Systems (ISMS) in Polish public institutions from 2016 to 2019 is evaluated. The methodological approach in the paper employs a systemic analysis to assess information security threats in PAs, utilizing a Cartesian-based model

to define interactions between threats and institutional vulnerabilities. This is particularly interesting, and it should be into account for the long-term final aim of proposing a PA-specific framework to enhance resilience. The authors conceptualize PA institutions as complex systems, where security is determined by the relationship between threats and protective measures. They adopt a model based on Clements' security system framework[9], defining a set of threats $Z = \{z_1, z_2, \dots, z_n\}$ and a set of vulnerable objects $O = \{o_1, o_2, \dots, o_n\}$. The relationship between threats and vulnerable objects is expressed as a Cartesian product $R \subseteq Z \times O$, which represents potential penetration paths within the system. This model enables a structured identification of vulnerabilities, allowing for targeted security measures to mitigate risks. The study further extends this approach by introducing a systemic situation model $\Sigma = < S, O, R >$, where S represents the system under threat, O denotes the external environment acting as a source of threats, and R captures the interactions between these components. The methodology incorporates probability-based assessments of security resilience, comparing the defensive potential $P(s)$ of an institution against the destructive potential $P(o)$ of threats, establishing conditions under which an institution remains secure. This structured methodology enables a quantitative analysis of security vulnerabilities, supporting standardized processes for threat assessment. Such standardization is essential for facilitating meaningful comparisons across different organizations. In the long-term perspective of this ADR, the definition of attack probability should ideally be grounded in an analytical framework of this kind. However, at present, it is not feasible to adopt such an advanced model. Consequently, the artefact relies on fixed probabilities assigned to different types of attacks, following the selection of both the relevant attack set and the secondary sources on which the assessment is based.

Sources

In the very first version of the artefact, a constant attack probability of 36% could be assumed, derived from secondary sources such as official reports issued by Istat, CSIRT, ENISA, and the Bank of Italy. Ideally, these sources should be replaced with primary data collected internally from each PA (or general organization). As already mentioned, establishing a standardized and structured approach to data collection is essential to enable the probability of an attack to be determined based on internally sourced information. In the case of public administrations, while the public sector as a whole may face similar threat landscapes on average, there remains significant variability in individual characteristics that should be leveraged to improve the accuracy of attack probability assessments for each PA. In the direction of providing a good tool for the specific case of Italian organizations, it is valuable to mention that the National Cybersecurity Agency published a cyber taxonomy[10],

[9] Hoffman, L. J., Michelman, E. H., & Clements, D. (1977). SECURATE User's Manual (Memorandum No. UCB/ERL M77/49). Electronics Research Laboratory, College of Engineering, University of California, Berkeley.
[10] Agenzia per la Cybersicurezza Nazionale (ACN). (2023). Tassonomia Cyber: Linguaggio comune per la classificazione degli incidenti informatici. https://www.acn.gov.it/portale/documents/20119/552690/ACN_Tassonomia_Cyber_CLEAR.pdf

which is particularly inspiring when it comes to differentiating the model by vector of attack rather than by threat actor.

As a disclaimer, it must be said that official cybersecurity reports often exhibit a certain degree of terminological ambiguity that can lead to confusion between the type of attack and the exploit technique used to leverage a vulnerability. Specifically, there is a tendency to conflate the classification of an attack (such as phishing, ransomware, or DDoS) with the specific method by which a vulnerability is exploited. For instance, a ransomware incident might be labeled as a "phishing attack", even though phishing merely served as the initial vector while the actual attack constitutes a separate and distinct phase with different objectives. This distinction is crucial for accurate incident analysis and the implementation of effective countermeasures, as confusing the entry vector with the nature of the attack itself can result in misguided or insufficient defensive strategies. In light of this consideration, the prototype developed in the case study will, for the purpose of initial testing, focus specifically on the types of attacks rather than the exploit techniques. Accordingly, the selection of the five attack categories for preliminary analysis will be based on their classification as distinct attack types, ensuring conceptual clarity and methodological consistency. For the time being, for the case study of Italian PAs, the monthly operational summaries by CSIRT make a good available source for this objective. These documents provide a monthly overview of key metrics and indicators derived from the operational activities of ACN, aimed at characterizing the state of cyber threats in Italy. Specifically, CSIRT Italia, the Agency's technical-operational division, serves as the national hub for mandatory and voluntary incident notifications required by law. Additionally, it collects information from open sources, commercial platforms, and equivalent national and international entities, shared either proactively or through collaborative agreements. This comprehensive dataset offers the Agency broad visibility into the cyber threat landscape affecting the national system and provides a structured qualitative assessment of threats and the exposure levels of national entities. Two examples of insight that are provided in the operational summaries and relevant for the case study follow: both Figure 15 and Figure 16 represent the distribution in kind of attacks having impact on the constituency. If the dataset from which these graphs are derived would available, for example, a comparative analysis across different months should be conducted, and an historical trend should be continuously monitored. This approach is essential for making informed decisions regarding the definition of the set A of attacks, as outlined in the guidelines for the proposed framework. Although the operational summaries do not provide relevant data for establishing weights among different types of attacks, it is the set of attacks that can be defined by considering the sector 'Pubblica Amministrazione Centrale' in its entirety. It is also important to note the presence of another sector, 'Pubblica Amministrazione Locale'. The same framework applied to the 'Pubblica Amministrazione Centrale' under the call 'Avviso 2' can be similarly employed for all calls pertaining to local administrations. Finally, there is the room for other sources to be integrated as mentioned before.

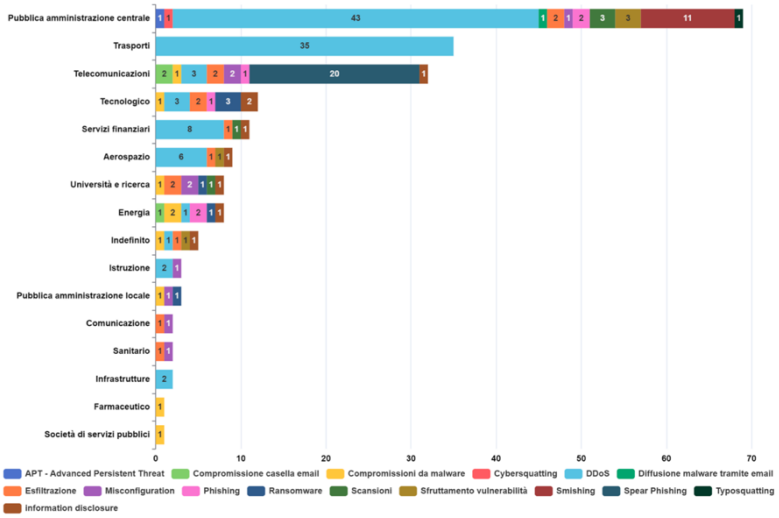

Figure 15 - Kind of attacks per targeted sector May 2024 (CSIRT)

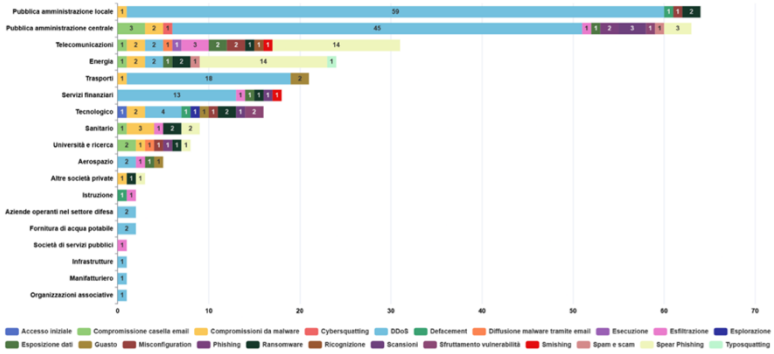

Figure 16 - Kind of attacks per targeted sector Feb 2025 (CSIRT)

Given the need of quantifying accurately the probability of an attack to occur even in the short-term, the horizon of the resources was extended to get full advantage of the Italian national initiatives around cybersecurity in PAs. Since the Computer Emergency Response Team (CERT-AGID – Agenzia per l'Italia Digitale) has been reporting statistics on cyber-attacks in Italy for a long time, it is worthed to mention that it would be extremely valuable to integrate their primary data to the model, so to take care of the assessment of threat probability and to better define the range of threats to be considered. The initial intuition is the potential benefit that could be got from linking the PA-specific vulnerability to a tailored list of Indicators of Compromise (IoC) that are relevant to the targeted one. This approach would also facilitate the identification of critical assets to be prioritized when evaluating the economic benefits of security measures, particularly in the calculation of ROSI. In general, multiple methodologies exist for classifying the probability of an attack, with the choice of approach depending on factors such as the sector in which the organization operates and other contextual variables.

Selection of Attacks

Acknowledging the need to balance a broad and detailed framework with limited data availability, an initial version of the artefact could adopt the previously discussed formula $P_A(PA) = \sum_{a \in A} P_a(PA) w_a^1$, in which the

probability of attack is fixed in advance but differentiated from a set of five selected attack types, meaning that the set A has cardinality of five. In this iteration, the weights assigned to each attack type w_a^1 are equal, employing a simple arithmetic average that assumes identical probability across all five. While this is a strong simplifying assumption, this revision lays the groundwork for future comparisons of ROSI across different organizations, rather than across different attack types, as previously discussed. Considering official national reports (e.g., ACN) and industrial reports (e.g., CISCO, IBM), the following five attacks are considered: Distributed Denial of Service (DDoS), data breach, ransomware, supply chain attack, and trojan-based intrusions. This selection would again be found in the discussion of the Cost of Attack. Indeed, to ensure consistency, the same set of attack types must be used to quantify this factor as well. For example, if a decision-maker responsible for cybersecurity strategy wishes to specifically compare two types of attacks, the ROSI can be calculated for each by using the corresponding probability of attack, P_a, and the associated cost of the attack, C_a as defined by the model for that particular attack type. While recognizing that upgrading this factor - by integrating a mathematical model that maps vulnerabilities to the probability of attack by type as mentioned above - would represent a significant advancement, it is also acknowledged that such an enhancement would require considerable time and effort to implement. As a result, its development is not currently prioritized. Instead, more immediately actionable proposals are considered in this iteration. For instance, integrating real-time data on ongoing attacks into the Python implementation could allow for dynamic updates of attack probabilities, thereby increasing the model's responsiveness and relevance without requiring a full-scale structural overhaul. Alternatively, if real-time data integration is not pursued, it is essential to at least schedule regular updates of the relevant values to ensure that the ROSI calculation incorporates a probability factor that reflects current conditions. In this context, a final consideration should be noted: among all the factors, this one - representing the probability of attack - is the most sensitive to dynamic changes in the real world. For this reason, its timely updating should be given the highest priority. While the cost of an attack may also be influenced by emerging developments, and to some extent the cost of investment as well, these factors tend to evolve more gradually. In contrast, the probability of attack is the factor most directly affected by sudden shifts in the cybersecurity landscape.

Effectiveness

Definition and prerequisites

By definition, the effectiveness is the probability for the organization of being completely protected against cyber-attacks. Keeping this in mind, it must be clear that the way the value is calculated is not considering any calculation of probabilities; instead, for the practical implementation of ROSI, the effectiveness is related to the cyber posture of the organization considered. For this reason, a first overview in regard on this other concept is provided and the relation among the two is then discussed. The cyber posture of an organization is the security status of an

enterprise's information, networks, and systems based on information security resources (e.g., people, hardware, software, policies) and capabilities in place to manage the defense of the enterprise and to react as the situation changes.[11] In the Call 'Avviso 2', which is the one this ADR is taking as instance to exemplifying the artefact's applicability and advantages of use, one of the interventions available[12] for institutions to request involves the assessment of cyber posture. This assessment is conducted by the agency in accordance with national guidelines for the cybersecurity framework. Specifically, six dimensions are evaluated, with each public administration (PA) assigned a level ranging from 0 to 5 to each dimension. This model, known as the Capability Maturity Model (CMM), is employed to assess organizational maturity and is endorsed by the National Institute of Standards and Technology (NIST) as a reference framework. A description for each level of maturity is now left in Table 8.

MATURITY		DESCRIPTION
Uncompleted	0	No formalized cybersecurity processes or capabilities are in place.
Initial/Ad Hoc	1	Cybersecurity measures are implemented in an unstructured and reactive manner, without consistent policies or procedures.
Managed	2	Basic cybersecurity processes are documented and followed, but they remain largely reactive and inconsistently applied.
Defined	3	A well-defined and standardized cybersecurity framework is established, with proactive measures and policies in place.
Quantitatively Managed	4	Cybersecurity processes are continuously monitored and measured for effectiveness, with data-driven improvements.
Optimized	5	Cybersecurity is fully integrated into the organization, with continuous improvement, automation, and strategic alignment with business objectives.

Table 8 - Levels of Maturity in Cyber Posture

These levels serve as the unit of measurement for assessing an organization's overall cybersecurity maturity. However, the evaluation can be more granular: a maturity level can be assigned to each of six distinct dimensions of cyber posture, and the overall cyber posture is then derived from these individual assessments. For completeness, the official description from ACN official pages is left as well in Figure 17.

[11] Johnson, A., Dempsey, K., Ross, R., Gupta, S., & Bailey, D. (2019). Guide for Security-Focused Configuration Management of Information Systems (NIST Special Publication 800-128). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-12

[12] Note: this intervention has been made mandatory for all recipients of funding under the Call 'Avviso 7,' which serves as a continuation of 'Avviso 2.' This requirement was introduced because it was essential for public administrations (PAs) to develop awareness regarding the impact of the interventions. Consequently, a measurement was required to ensure such awareness.

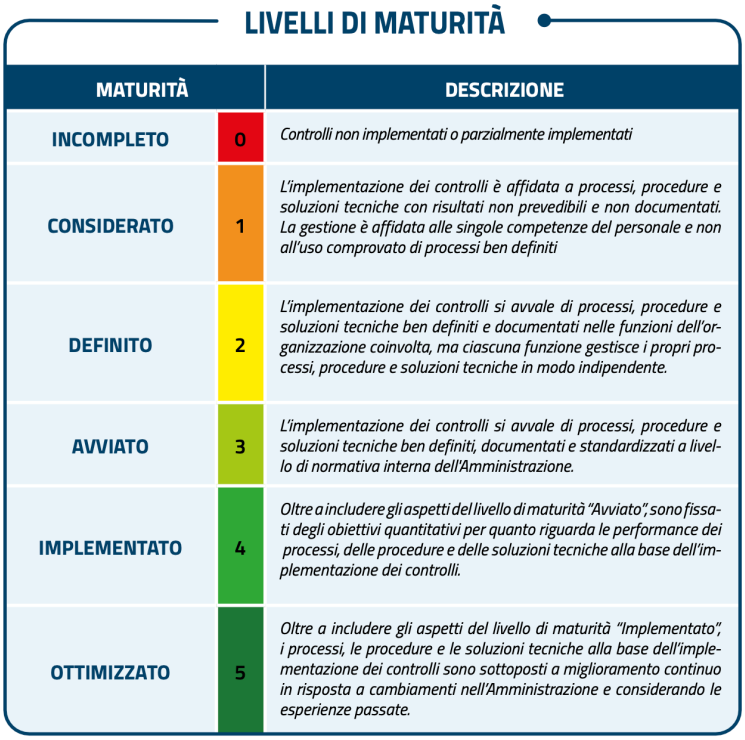

Figure 17 - Levels of Maturity (ACN)

Having already introduced the maturity levels, it is now appropriate to present an overview of the six dimensions.

- Governance and processes

  This dimension encompasses all subcategories of the FNCS[13] related to the effective governance of cybersecurity. It includes the definition of roles, responsibilities, policies, processes, and strategic approaches to cybersecurity. Furthermore, it prescribes the integration of cybersecurity into the organization's overall governance structure and decision-making processes. Externally, it requires the clear identification and communication of roles and responsibilities.

- Management of cyber-risk and operations' continuity

  This dimension includes all FNCS subcategories focused on the systematic management of cyber risk - from identification through mitigation - as well as operational continuity, with the objective of minimizing potential impacts on day-to-day operations.

  Preparatory to risk and impact analyses, it includes asset mapping, threat and vulnerability assessment, definition of risk tolerance, supply chain risk management, and backup procedures. It also includes subcategories addressing the management of risks associated with the potential disclosure of personal data, in accordance with the data protection requirements outlined in the FNCS.

- Security incident management and response

[13] National Cybersecurity Framework

This dimension includes all FNCS subcategories focusing on the organization's capacity to manage security incidents, from detection and analysis through response and recovery. The goal is to restore normal operations in a timely and effective manner. It aligns with the functions of Detect, Respond, and Recover as defined in the FNCS.

- Management of logical access and digital identities

  This dimension includes all FNCS subcategories aimed at ensuring that access to systems, information, and digital resources is managed securely and proportionally to each individual's role and responsibilities. It encompasses the proper administration of digital identities through secure lifecycle management processes, remote access control, implementation of the principles of segregation of duties and least privilege, as well as the adoption of multi-factor authentication mechanisms.

- Training and awareness on cybersecurity

  This dimension incorporates all FNCS subcategories related to awareness-raising, education, and training activities targeted at personnel, considering their respective roles and responsibilities. It addresses both cyber and physical security. The objective is to ensure that personnel are aware of their roles and responsibilities, as well as of the cyber risks associated with their activities. This is achieved through structured training programs and the effective communication of security policies, including the dissemination of fundamental principles (e.g., proper credential usage).

- Security on applications, data and networks

  This dimension covers all FNCS subcategories aimed at ensuring the protection of the IT infrastructure, communication networks, data, and software applications. This includes, but is not limited to: network monitoring, the use of technologies that safeguard the integrity of data and devices, and the proper management of computing environments.

A brief disclaimer is warranted: these six dimensions are specific to the framework proposed by the National Cybersecurity Agency and may not be universally recognized. They were developed as an aggregation and reinterpretation of the NIST Cybersecurity Framework guidelines. Nevertheless, the current model is structured around these categories. Any future modification to the number or nature of the dimensions would not compromise the integrity of the artefact; it would require adjustments to the formulas. For example, the current hexagonal representation that is left in Figure 18 could be adapted into another regular polygon if the number of categories changes, or the labels at each vertex could be revised accordingly. It provides an example of an assessment for these six dimensions (ACN, Relazione Annuale al Parlamento 2023, 2023). It is visualized in Figure 18, in which each vertex corresponds to a dimension and the steps relate to the levels (from 0 to 5) for that dimension. Each PA , and more generally each organization, can have a related spider map like this one, in which the dark blue line gives the caption of the cyber posture before the interventions, then there is the 'quick win' line which concerns the situation for the PA no long after the interventions, and the target to achieve stands in blue.

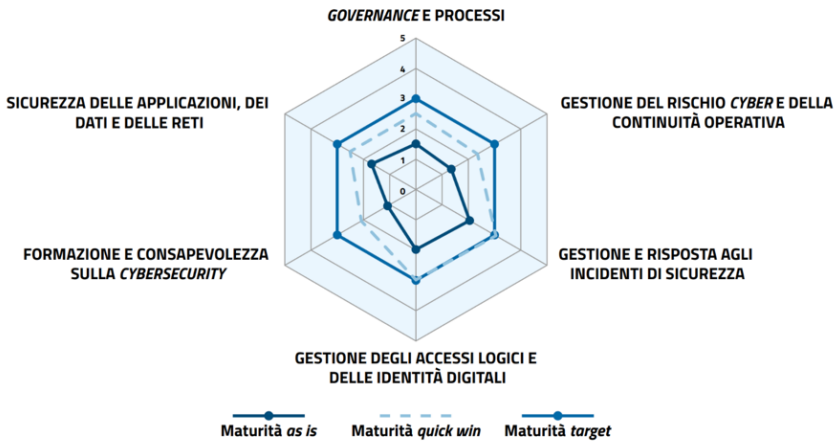

Figure 18 - Dimensions of cyber posture

Before elaborating on these classifications, it is essential to introduce their usage. The total cyber posture of a PA is represented as a numerical value between 0 and 5, that can be calculate in the most basic version as the arithmetic mean of the levels assigned to each of the six dimensions, saying:

$$CP = \sum_{i=1}^{6} \frac{CP(\dim_i)}{6}.$$

Since an initial overview of cyber posture evaluation has been provided, attention can now shift to its relationship with effectiveness. As previously indicated, effectiveness is defined independently of cyber posture. Nevertheless, for the model to be implemented, it is necessary to determine which assumptions should be accepted or rejected regarding how a system's effectiveness varies in relation to its cyber posture. No further requirements need to be introduced at this point, as a detailed discussion of both cyber posture and effectiveness is provided in the following subsection. A final remark before addressing the specifications: the decision to present the update of the total cyber posture evaluation prior to the update concerning its relationship with effectiveness does not imply a required sequence of implementation. This structure once again reflects the modular nature of the framework: in practice, it may be more feasible to implement the cyber posture component first - or vice versa.

Weighing the Dimensions of Cyber Posture

For the evaluation of the global level of cyber posture, it is not necessarily the case that each dimension is as important as the other ones to the overall assessment of the cyber posture of a single organization. This can be up to the sector for example, or to other variables. In a long term, these weights should be calculated in respect to some predefined drivers so that the tool of ROSI would fit each peculiar application. Keeping this aim in mind, a first shift to upgrade the artifact can still be made with a general rebalancing the weights for the six dimensions thanks to the contribution of domain experts. Within each dimension, the transition between two consecutive

levels is not always consistent; the rate of improvement in cyber posture is not necessarily linear across the dimensions. This is acknowledged but the issue is postponed to the relation between the effectiveness and the cyber posture. Second, the six dimensions do not necessarily share a uniform scale, as the effort required to progress from one level to the next varies across different areas. The first approach, consisting in assigning different weights w_i to the various dimensions and thus reflecting their specific impact on overall cybersecurity, would turn the formula for the overall cyber posture of an organization into:

$$CP = \sum_{i=1}^{6} w_i CP(\dim_i).$$

At first glance, looking back at the representation of the six dimensions in Figure 18, this update can be represented graphically as a transition from a regular hexagon as the one on the left in Figure 19 to an irregular hexagon, on the right in Figure 19.

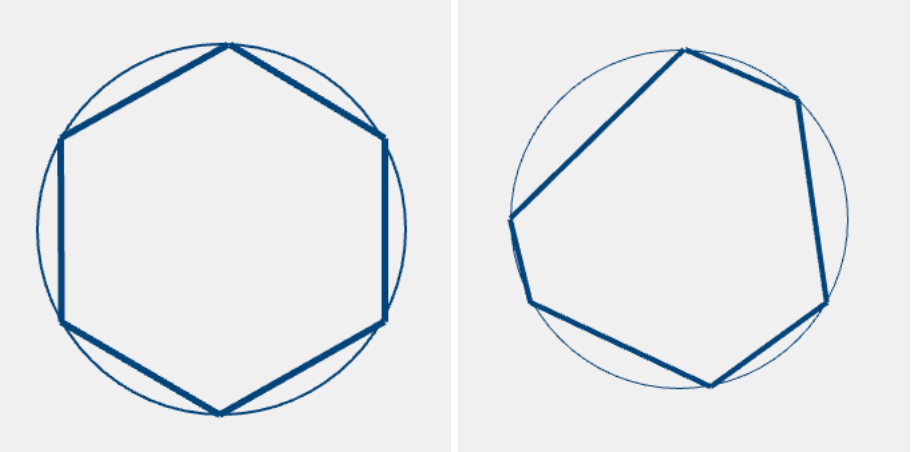

Figure 19

A framework for determining the weight of each dimension in the assessment must be developed to ensure consistent application before and after the interventions. In the initial iteration of the model's enhancement, a single set of weights was established. While this approach abandons the use of an arithmetic mean, the weights remain fixed, resulting in a model that still lacks dynamism. Based on responses to the questionnaire and the interventions carried out, the cyber posture values across different tiers can be calculated relative to the baseline. Although there is no intention to re-evaluate the cyber posture for each individual dimension, the overall cyber posture can be more accurately determined by applying the fixed weights to the six dimension-specific scores. In the first iteration, the artefact relied primarily on data from secondary sources to define these fixed weights. This preliminary version incorporated contextual characteristics of the assessed organization by aligning the drivers recorded in the Call with detailed information on attack costs from IBM's report (Security, 2024). However, the secondary sources used were predominantly international reports, and no primary data was available for integration into the artefact at that stage. In the second iteration, although the weights remained fixed, they were updated based on the context-specific domain knowledge acquired during the research.

Before delving into the discussion of effectiveness, it is important to note that a correlation exists between the cost amplifiers and mitigators identified by IBM (Figure 20) and the various dimensions considered in assessing an organization's cyber posture.

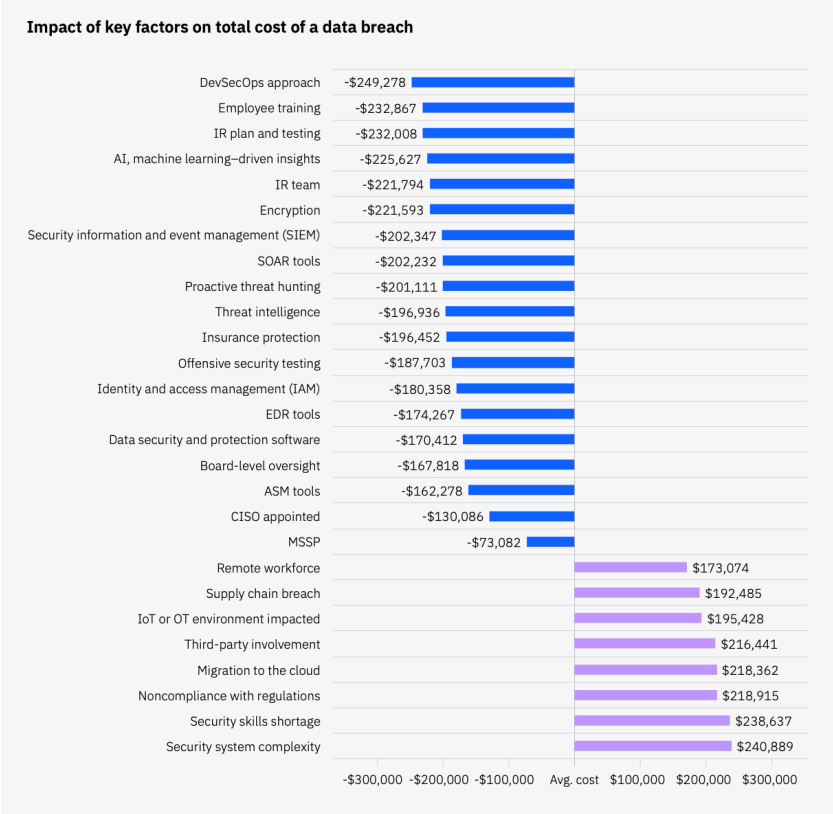

Figure 20 - Impact of key factors on total cost of a data breach in 2023

In an initial attempt to enhance the artefact, this correlation inspired to assign informed weights to the different dimensions of cyber posture, with the aim of producing a more accurate overall assessment. However, this approach was ultimately deemed unreliable, primarily due to the nature of the secondary sources, which were not suitable in form and did not fully align with the intended conceptual framework. Nonetheless, as additional data is expected in the future, this method is briefly outlined here for potential reference. In practice, the costs associated with specific driver classes within the organization are evaluated relative to the theoretical average cost, in order to determine the percentage increase or decrease attributable to each characteristic. Based on the alignment between the drivers and the dimensions of cyber posture, weights for these dimensions - used to calculate the overall cyber posture - are then derived. Greater weight is assigned to dimensions contributing to higher cost amplification, and less to those associated with mitigation. The goal of the weight calculation is to assess how much influence each dimension has on increasing cost. Specifically, for each dimension, the maximum cost increase (when that dimension is active) is divided by the total maximum amplification across all dimensions.

This results in normalized weights that reflect the relative impact each dimension has on cost increases, emphasizing those with the greatest amplification potential, as summarized in Equation I.

$$w_i = \frac{\max amplification\ (\dim i)}{\sum_{j=1}^{6} \max amplification\ (\dim j)}$$

Equation I

This approach relies on a taxonomy that maps the drivers to the dimensions they influence, thereby contributing to the profiling of the associated cyber posture. In case this approach would be adopted in future iterations of the model, special attention must be given to the process of mapping drivers to dimensions. In particular, a rigorous and well-defined method is required. It is also crucial that this mapping process is aligned not only with the six dimensions currently in use, but also with widely recognized frameworks, such as the guidelines provided by NIST.

The relationship between Cyber Posture and Effectiveness

The fundamental question that arose working on the artefact was substantially how cybersecurity posture increases from level 0 to level 5 in relation to X. This question should be posed for each dimension, followed by a comparative analysis to ensure coherence across dimensions after an initial iteration. Other surveys might become more granular via addressing this relation between the cyber posture and the effectiveness of the system dimension per dimension, but this is not the case for the first round. The primary challenge, however, concerns the selection of X, as multiple options exist for this variable. The issue is an inherent difficulty in choosing it but also a lack of sufficient data whatever the choice would be. One possible approach could be to use the effectiveness directly, not introducing other complexities but knowing it would be restrictive and not entirely quantitative. This is exactly what was done, and it means that the domain experts were asked, during a proper workshop, to draw the behavior of the effectiveness in respect to the total cyber posture. A reminder: effectiveness refers to the probability of a system to be completely protected against cyberattacks, which can be else explained as the degree to which a system achieves having zero vulnerabilities. Another possible approach could be to use the volume of Personally Identifiable Information (PII) mapped within the PA for a given dimension. Official reports published by IBM in 2023 and 2024 highlight how the presence of large amounts of PII significantly escalates the average cost of a data breach, taking it completely out of range in respect to the average cost. Nonetheless, mapping this data is complex. A preliminary attempt could involve analyzing the responses from PA questionnaires, where the types of PII handled by the administration are recorded. At the very least, the dependency could be based on the cardinality of the set of PII types. While this approach does not account for volume, it provides a more informed basis than having no data at all. Additionally, this limitation could be addressed in future iterations of the framework, once more data becomes available. With these guidelines in place, PAs and could be encouraged to collect more granular data on their PII to enhance the accuracy and efficacy of this approach. At that point, the

dependency on X could become even more accurate, extending X to $X = (X_1, X_2)$, where the first coordinate indicates the kind of PII and the second one its volume.

One of the objectives of this ADR concerned the refinement of the effectiveness that is associated with the level of cyber posture. As a baseline, the effectiveness was initially assumed to be linear with the percentual increase in cyber posture at two different times, namely before and after the interventions. Quantitively, it means that the effectiveness $Eff(i)$ related to an intervention i was defined as:

$$Eff(PA_1) = \frac{CP_{PA_1}(a) - CP_{PA_1}(b)}{5} = \frac{\Delta CP_{PA_1}}{5},$$

which corresponds to the percentual increase in the cyber posture looking at the cyber posture of a public administration, PA_1, before (time b) and after (time a) the intervention i. While this approach does not present inherent contradictions, it blocks a series of potential application of this factor that would go beyond its use as factor of ROSI. Indeed, relating the effectiveness to a difference in cyber posture instead that relying on a punctual correspondence has the consequence of unavailability of an effectiveness value ad priori. To mitigate this significant limitation, the approach was changed, and the usage of a punctual correspondence was introduced. In terms of formula, this means that the model first took the effectiveness as a function of ΔCP, but it was turned then into a function of CP. Consequently, a discussion arose around the function to use for this. The linearity could be kept as first step, getting the following new applicability-oriented formula:

$$Eff = \frac{CP}{5}.$$

To further clarify the difference from the previous setting, an example is provided. Recognizing that the interest in associating an organization's cyber posture with its effectiveness extends beyond the ROSI tool, consider a company evaluated as having a total cyber posture of 2.5 at a given point in time. While a different positioning of this company at another moment was required in the previous version, the new update allows for a direct interpretation: the company can now be considered 50% protected from cyber-attacks (this is, in fact, the meaning of scoring 50% in effectiveness). After this shift, in the ROSI formula it is not anymore the effectiveness associated to a change in cyber posture to be used, but instead the difference between two effectiveness that are respectively independently associated to two cyber postures.

The next step consists in refining the relationship between Effectiveness (Eff) and Cyber Posture (CP), moving beyond the linear assumption. The ideal solution would involve defining a specific function for each dimension, complementing the weighting customization discussed in the previous subsection. In this initial iteration of the ADR, only a limited advancement toward the long-term objective is undertaken. Non-linearity between the two factors is implemented; however, it is applied only to the total cyber posture, rather than providing a granular differentiation of this non-linearity across all six dimensions. Certain dimensions require greater initial effort to

improve their cyber posture but exhibit diminishing returns as they reach higher levels. In Figure 21 and Figure 22, two example of possible relationships are shown. The colors of the vertical bands are coherent with the colors in the official ACN table presenting the levels of maturity that are unity of measure of the cyber posture (Figure 17). Figure 21, which has a concave shape, may be appropriate in cases where the assessed dimension of cyber posture is more labor-driven, while Figure 22 could correspond to a more capital-driven context. This distinction is reasonable, as investments in employee training tend to yield more immediate improvements in system effectiveness, whereas achieving significant advancements in areas such as network security tends to be more complex and gradual.

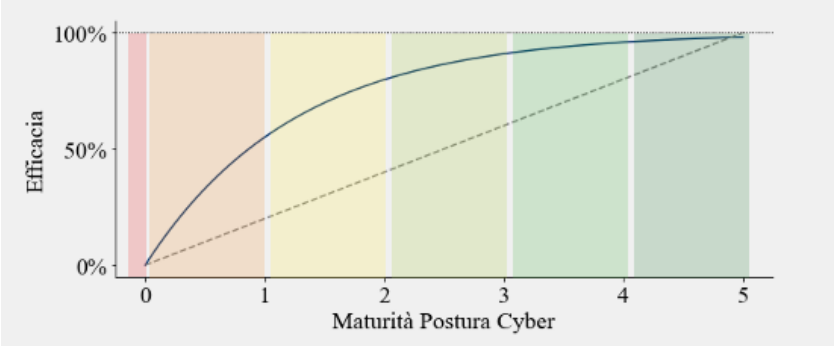

Figure 21 - Example of a labour-driven dimension

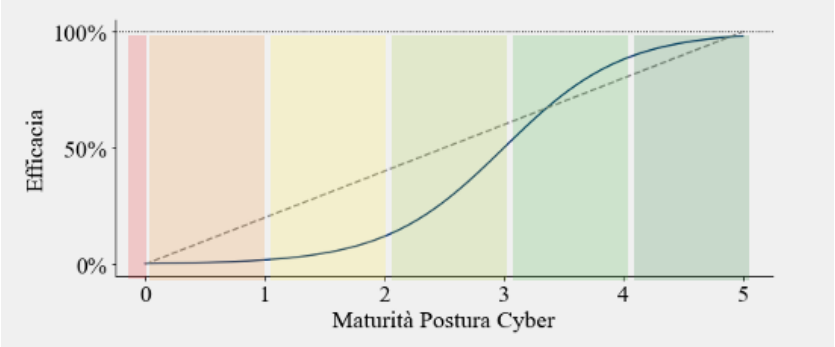

Figure 22 - Example of a capital-driven dimesion

Another visual representation is left and discussed to clarify further the hard limitation in assuming a constant shift among the levels of maturity, whose refinement is one of targets of this approach. Consider for instance the dimensions ‘Formazione e consapevolezza cyber’ and ‘Gestione degli accessi logici e delle identità digitali’, respectively: ‘Cybersecurity Training and Awareness’ and ‘Management of Logical Access and Digital Identities’. Once again, for the first one, a strong upgrade would be intuitively get starting from the maturity zero to the first levels up, meaning that the initial effort is highly impactful on the maturity of this dimension, intuitively. Instead, the Management of Logical Access and Digital Identities requires a long-term approach, and the initial actions enhance the cyber posture slowly. In summary, this comparison can ideally be outlined as follows.

Cybersecurity Training and Awareness	1	2	3	4	5
Management of Logical Access and Digital Identities	1	2	3	4	5

Figure 23 - Comparison of dimensions

Going from light blue to pink, assuming that a PA is assessed at maturity 3 for both dimensions, the associated effectiveness differs significantly between the two cases. The shift from level 3 to level 4 does not correspond to +20% in either case (as it would be with the linearity assumption) and has a considerably greater impact as an upgrade for the second dimension. Unfortunately, this visualization is not based on any kind of reliable methodology for distributing the six maturity levels, but it serves instead just as an example. In any case, following this approach, the effectiveness could be defined as the percentual length of the maturity line corresponding to the cyber posture of the organization, over the total length of five. Referring to the instance in Figure 23, for the first dimension this percentual increase would be the ratio among the length of the dark blue are and the total length of the colorful table, and similarly for the second dimension in purple.

Cybersecurity Training and Awareness	1	2	3	4	5
Management of Logical Access and Digital Identities	1	2	3	4	5

Figure 24 - Comparison of dimensions

As anticipated, in this first iteration of the artifact the upgrade wouldn't be granular enough to get a specific function for each dimension of cyber posture, but a global readjustment of the dilatation of the levels is introduced. At first glance, the transformation can be represented graphically as illustrated in Figure 25.

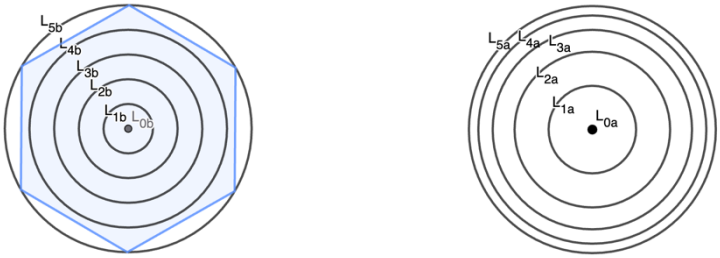

Figure 25

With the upgrade, there is transition from the concentric levels $L_{0a}, L_{1a}, L_{2a}, L_{3a}, L_{4a}, L_{5a}$ where the rays increase uniformly across successive circles, to the concentric levels $L_{0b}, L_{1b}, L_{2b}, L_{3b}, L_{4b}, L_{5b}$, for which the shift must be established. The notation a and b are used consistently to refer to the cyber posture after and before the intervention, respectively. The hexagon remains unaffected throughout this transition (and is therefore not replicated) since the different weights among the dimensions can be (and it is) managed as an independent process in respect to this one. In this regard, balancing the dimensions is modular and compatible with introducing the non-linearity between the cyber posture and the effectiveness of the system.

Coming to one of the goals for this section, regarding the upgrade of the artefact, the time to structure a methodology to define the actual relation between the level of total cyber posture and the effectiveness of the organization arrived. The solution involved adopting a bottom-up approach, by assembling a team of domain experts to conduct interviews and gather insights on the progression curve of effectiveness in respect to the total cyber posture. This is done using the Analytical Hierarchical Process (AHP) technique, on a dedicated workshop organized in support of this research. An entire chapter is dedicated to this workshop, but for convenience its outputs are discusses directly in the subsequent subsection to provide the official update that is got through that contribution.

Enhanced formula

In conclusion, the most advanced formula that the ROSI model can use for the factor of effectiveness, restrictive because of the limited availability of data that is typical of the initial stages, but advanced thanks to the integration with the outcomes from the workshop, can be presented. Let

$$Eff = f(CP) = f\left(\sum_{i=1}^{6} w_i CP(\dim_i)\right),$$

Equation II

Here, f denotes the specific function obtained by interpolating five points on a plane representing cyber posture versus effectiveness. The x-axis consists of natural numbers from 0 to 5, each corresponding to a defined level of cyber posture. These points represent the output of Question 2 from the workshop, as discussed in the dedicated chapter.

A final important note: as mentioned, the two approaches that are proposed – the weighting of the cyber posture's dimensions and the relationship between the cyber posture of a system and its effectiveness - are modular with each other, meaning that they can be integrated and there is no way one is dependent on the other, even though they would be both enhanced by the other's integration. Referring to the Equation II, it means that it can either be chosen to work on w_i or rather defining the function f. For clarity, the schema in Figure 26 is left. Thanks to the

modularity of the method that has been proposed, either following the first (represented in orange) or the second (represented in green) approach would equally take to the optimized final setup.

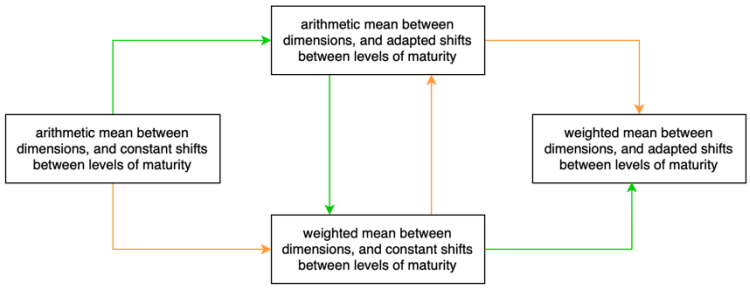

Figure 26 - Modularity within the factor of effectiveness

Cost of an Attack

Definitions

Several issues arise when considering the Cost of an Attack, particularly regarding which consequences should be included in the assessment - both direct and indirect - and how long these effects should be tracked. To address this, it was important to establish a clear time frame for evaluating the impact. For consistency with existing practices, such as the assessment of quick wins in cybersecurity posture, a one-year time horizon is used as the standard. However, where feasible, a second, extended time frame of five years may also be considered. This longer horizon is particularly relevant for evaluating the long-term return on security investments and should be adjusted using an appropriate discount rate in case the investment is fixed at time zero. Considering the proposed implementation for the probability of attacks, again a further opportunity to get a modular approach to make the ROSI calculated more accurate is released in this factor too. Indeed, having already defined a dependency on the type of attack in that section, the exact same can be done here. However, for a first evaluation an average cost of an attack can be used, not considering other specifications nor internal resources, but relying on secondary sources and fixing a standard value. In a longer-term, instead of a general average, the value could be got further considering not only different costs for different attacks but also weighting these costs in respect to the context of the organization. Depending on what it is done in the factor of probability of attack, there are two options to advance. To explain this, the term of the efficiency is just set as it was constant, since it is not the point this time.

$$EB = P_A * Eff * \left(\sum_{a \in A} w_a C_a\right)$$

Equation III

$$EB = \left(\sum_{a \in A} w_a P_a\right) * Eff * C_A$$

Equation IV

$$EB = \left(\sum_{a \in A} w_a P_a * C_a\right) * Eff$$

Equation V

where w_a has the same meaning as in the section dedicated to the probability or cost of an attack. In Equation III, it is the probability of attack to fixed as a constant. In Equation IV, it is instead the cost of an attack, while in the third one is neither of the two and for this reason the indices are granularly referring to the probability and the cost of an attack contemporaneously. For the first implementation of the artifact, the second formula is the one recommended. An important note: for the last version stated above, the weight w_a must be determined as dependent on both the probability and the cost of attack, introducing a multi-variate problem to the formula. However, there is a different step forward which is not enough to get the extent to reach the third version, but it is still a valuable improvement. Specifically, this is extremely important because it is the way the ROSI model is adapted in respect to the different characteristics of the organizations. In practice, a tailored value for C_A is calculated for each organization: it would still be fixed and missing dynamicity in respect to the variate threat landscape of the attack, but at least it would be calculated in respect to a set of drivers that aims to profiling the organization.

Even with these initial updates, the model still relies on secondary sources, which is a limitation to be addressed in the long term. While the ideal objective is to base the model on primary data and develop a tool customizable on a case-by-case basis, several proposals are outlined as future directions. One initial proposal revisits the set of selected attack types discussed in the context of attack probability. A potential enhancement of the framework involves mapping these attack types to specific business functions or services within each organizational context. This would enable a more granular assessment of potential financial losses and operational disruptions. As an additional remark, this structure also facilitates expert elicitation regarding sector-specific vulnerabilities, attack likelihood, and resilience strategies. Subsequently, the impact of the five types of attacks should be calculated based on a set of reliable secondary sources rather than primary data. Assuming that an economic value has been assigned to the affected business function, the expected loss can be computed as the product of this monetary value and the specific impact factor. Finally, summing the expected losses across the five attack types yields the total value of C_A, which can be used in the ROSI formula.

Alignment of Costs with Organizational Profiles

Acknowledging that working on this factor was the easiest way to develop a tool which would be in some way context-specific, the ADR focused on speculating one of the most important and recognized sources, which is the IBM 2023 report (Security, 2024) the cost of attack. The report provides a range of disaggregated estimates - such as costs by country, sector, organization size, breach detection method, and various amplifiers or mitigators - that

offer a valuable basis for contextualizing the economic impact of a breach. Rather than relying on the global average cost of a data breach (approximately USD 4.5 million), the approach proposed here involves selecting a cost estimate that reflects the specific characteristics of the organization under assessment. For instance, costs differ significantly between sectors (e.g., USD 3.86 million in Italy versus USD 2.6 million in the public sector), by headcount (e.g., USD 3.3 million for organizations with fewer than 5,000 employees versus over USD 4.8 million for those with more than 10,000) and based on detection method or security maturity indicators. These detailed cases are intended to be matched with the organizational drivers recorded in the call (or rather via interviews or questionnaires in a generalized applicability of the artefact), enabling a more tailored and realistic assessment of potential financial impact. This matching process reinforces the context-specific nature of the tool, moving beyond generic estimations and aligning more closely with the actual exposure and resilience profile of each organization. To stress these points, it is good to reflect on the fact that the assessment of the cyber posture for an organization (and consequently its effectiveness for what discussed above) and the differentiation between cost of attack among different organizations are derived from the same primary data.

An example is left to clarify the pipeline with which the cost of attack is determined specifically for each organization. The number of employees is an easy information to collect for almost every kind of organization. Figure 27 identifies the different cost that are associated to a data breach in respect to the number of employees in USD dollars. This is just one instance, but the IBM reports are populated by these graphs, being deeply granular but with the counterpart that every time the discussion concerns one factor influencing the cost of an attack, instead of approaching the issue as a multivariate problem.

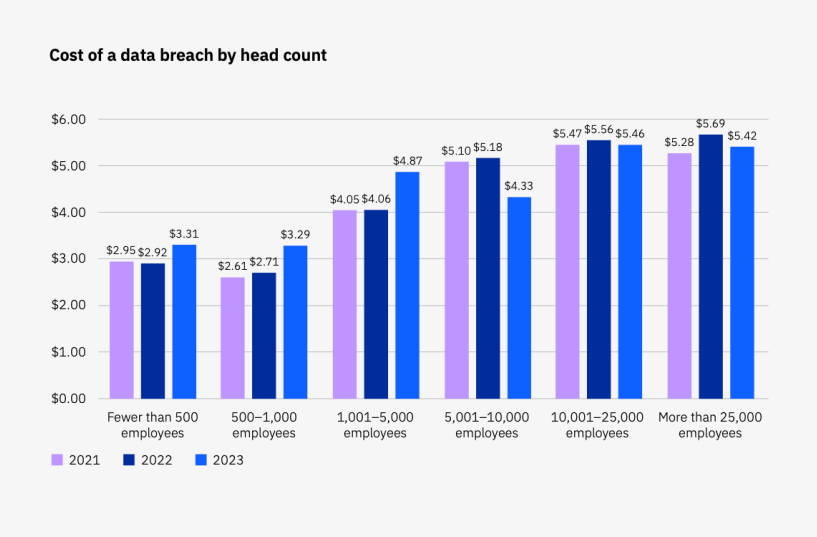

Figure 27 - Cost of a data breach by head count in 2023 (IBM)

The easiest way to explain the framework for the value of cost in use is through an example. Let *PA* be a public administration that participated in 'Avviso 2' and received funding. As a baseline, an initial ROSI can be calculated

using the average cost of an attack. However, the aim is to make this value more tailored to the specific context of PA. An initial value is determined based on the user's responses to two preliminary questions. Specifically, rather than assigning a default cost of 4.5 million USD, the tool asks whether the organization is based in Italy and whether it operates in the public sector. Fixed values are assigned to each combination of responses, and the corresponding value is selected as the starting point before proceeding with the algorithm. Further information about the organization is then matched with the classification provided in the IBM report, so that for each category (or "band") of drivers the organization belongs to, the corresponding cost of an attack estimated by IBM is associated. This process is repeated for all relevant drivers to ensure that the maximum amount of information is considered. The result is a list of costs: it may be the case that the organization is associated with higher costs of attack for some drivers, and lower for others. To determine the cost to use in the final calculation, a conservative approach has been adopted: the maximum value among all associated costs is selected.

Another try was perceived, as an alternative to the one that has just been explained. This one was inspired by one peculiar study reported by IBM, which was shown already in Figure 20. The underlying idea is to identify the factors that influence the cost of a cyberattack, either positively or negatively, and to adjust the default cost accordingly based on whether the organization exhibits any of these amplifying or mitigating characteristics. In this second algorithm, the list of associated costs matching the information regarding the organization and the bands of costs in the IBM records is still used, but the operation among those values is not the maximum this time. Until the match between costs and the organization’s profile, the steps are the same as in the previously explained algorithm. But instead of taking the maximum cost among those, the values are normalized by dividing each value by the fixed benchmark of 4.45 million USD, resulting in the expression of the relative amplification or mitigation effect of each scenario. Finally, the organization-specific potential cost of a cyberattack set as default is scaled by the maximum relative percentual amplification. This approach remains conservative; however, rather than relying directly on the values reported, it seeks to contextualize the scenario to a greater extent by incorporating organization-specific characteristics.

Chapter 6 – Findings

Workshop

The workshop was held in the Conference Room of the Cyber 4.0 Technopole and was moderated by a Luiss professor, who also serves as the president of Cyber 4.0. Following an opening address by the Director of Cyber 4.0, a representative from the National Cybersecurity Agency (ACN) outlined the context that led to the research initiative, emphasizing the need to evaluate the investments made under PNRR 1.5. The speaker also clarified that, from the outset, the Agency recognized the broad application potential of the project, particularly regarding the cost-benefit assessment of cybersecurity-related interventions. A second ACN representative then highlighted the importance of obtaining concrete results from the study before requesting additional financial resources. They stressed the need to frame cybersecurity as an investment rather than a cost. It was then time for the moderator to begin the actual session. Initially, the context for which the research is developed was introduced, as for the importance of cybersecurity in IT management (reporting in particular Figure 8 and Figure 9). Moreover, the basic concepts around ROSI were explained, as for instance its usefulness as a mean to drive strategic decision on investments, the variety of framework around its constituency, and the issues that arise for its definition. The workshop continued with a presentation of the Action Design Research (ADR) project that was entitled "Return On Security Investment (ROSI): A Framework for Measuring and Assessing the Impact of Cybersecurity Investments in Public and Private Organizations" for this occasion. This session also introduced the broader research objective: assessing the return on cybersecurity investments. In addition, a tailored representation of 'Building, Intervention and Evaluation' (BIE) phase was shown to the audience (Figure 13). Thirteen experts participated in the discussions, seated in a semicircle facing the presenters. The first interactive moment was triggered by the completion of a Mentimeter questionnaire proposed by the moderator, as shown in Figure 28. The response scale ranged from 0 (strong disagreement) to 5 (strong agreement) with the statement provided.

1. Question 1 – Perceived Challenges of ROSI

   The first question addressed the perceived challenges associated with ROSI (Return on Security Investment). The average score was 4.3 out of 5, with responses clustered toward the higher end of the scale. This indicated broad agreement among participants, prompting a transition to the next topic to stimulate expert discussion.
2. Question 2 – Feasibility of Estimating Economic Benefit

   The second question generated more disagreement, with a mean score of 2.8 out of 5 and a distribution that more closely resembled a Gaussian curve. One of the first widely shared criticisms was the discrepancy between the scores of the first two questions. Since the challenge of ROSI is closely tied to

the estimation of economic benefits, many argued that, logically, Question 2 should have received a significantly higher score. One expert (a CISO) noted that within the professional community, decision-making strategies rarely rely on precise calculations of economic benefits. Instead, organizations often set a target level of cyber maturity and plan interventions accordingly. Another expert, a former CISO in the banking and insurance sectors, explained how economic benefit is typically calculated as the product of event frequency, impact, and a factor of uncertainty. Using this definition, they emphasized that the concept of impact must account for both lost revenue and initial costs. In public sector contexts, additional factors - such as the failure to deliver services to citizens - must also be considered. Subsequently, an actuary expressed skepticism about the strength of statement 2, arguing that an estimation of economic benefit is always feasible. They found the term "impossibility" poorly suited for the topic. They also stressed the importance of using probabilistic models, as outlined by the previous expert. Additional contributions on the topic of economic benefit included references to EBA (European Banking Authority) regulations and the suggestion that penalties should be considered both in terms of impact and likelihood of attack.

3. Question 3 – Cost Drivers and Market Dynamics

   The third question received an average score of 3.4 out of 5, with responses forming two non-extreme peaks - one in the agreement zone and one in the disagreement zone. It elicited less discussion among the experts, with only one participant explicitly commenting that while it is correct to state that cost is determined by the market, in the public sector it is also essential to consider the effects of service disruption.

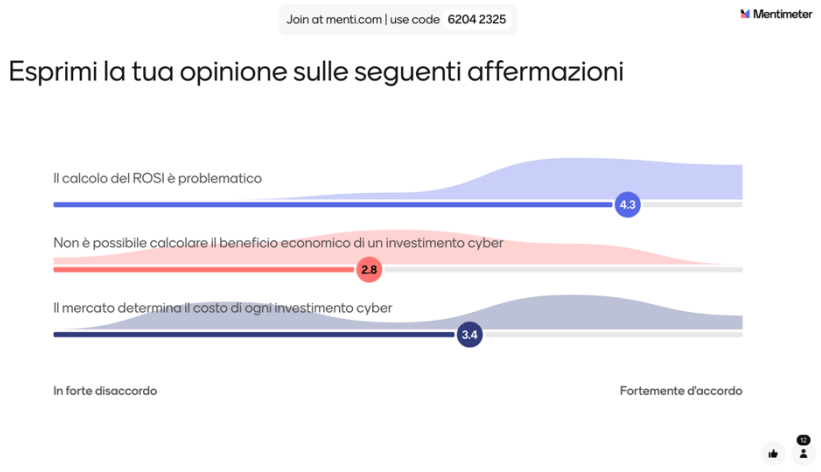

Figure 28 - Mentimeter results

Overall, this initial round of open discussion highlighted a shared recognition of the need to adopt more quantitative approaches - while acknowledging that these must originate from qualitative assessments. Building on this takeaway, the moderator clarified that a primary goal of the next phase of the workshop was to quantify expert perspectives. The Action Design Research methodology was then formally introduced, followed by the presentation of the proposed framework, including methodological aspects and applications. A second feedback session followed, during which the following insights were collected:

- The market capitalization and sector of the organization should be considered when quantifying factors that influence ROSI;
- Given the complexity of the problem, ROSI should be suitable for strategic-level use, particularly for evaluating investment alternatives;
- The industry in which the organization operates must be treated as a key driver in the computational model;
- Qualitative methods should not be excluded from the framework;
- Probabilistic models should be integrated and adapted to contextual information;
- The cost component of the investment lacks a corresponding uncertainty factor, which is present on the benefit side; this discrepancy should be made explicit in the model;
- In cost estimation, both direct and indirect costs must be clearly distinguished and considered;
- Reputational damage should be factored into the adapted ROSI calculation when contextualizing organizational risks;
- The kind of data processed by the organization must be taken into account.

The framework was then presented with the metaphor of the tree. Consequently, a context was introduced to the audience, as case study to consider from the practitioners in responding to the questionnaires. The AHP was conducted within the context of ASL Metropolitana Nord (ASL-MN), a public Local Health Authority located in Northern Italy, serving both urban and peripheral areas. ASL-MN employs approximately 3000 staff members, including doctors, nurses, administrative personnel, and technical staff. It offers a broad range of healthcare services such as hospitals, outpatient clinics, emergency care, occupational medicine, vaccination services, and manages a centralized booking system (CUP). With an annual turnover of around €800 million, the organization operates a centralized IT infrastructure housed in an internal data center, providing remote access capabilities for community-based medical personnel. Its digital ecosystem includes administrative and healthcare management software, PACS systems for medical imaging, and electronic health records. External vendors are engaged for IT maintenance, diagnostic systems, and software support. Given the sensitivity of patient health data protected under the GDPR, the critical nature of life-saving services, and the presence of network-connected biomedical

equipment and hospital servers, ASL-MN represents a complex and high-stakes environment in which to assess decision-making through AHP.

Following this, the experts were introduced to the AHP (Analytic Hierarchy Process) methodology along with its advantages, the context for framing their responses, the two main evaluation questions, and the prerequisites for answering them. Then, some further specification about the six dimensions through which the cyber posture is assessed and a review of the maturity levels were given in preparation to the AHP technique exercise. In particular, the slides about the context, the maturity levels and the details of the dimensions were left in paper format to the key informants during the inquiries phase too. Inspired by the work of (Oliva, Faramondi, Setola, Tesei, & Zio, 2021) the questions were structured in the same format as in that study but tailored to the specific aims of the workshop. The inquiries were written in Italian, as this was the language in which the workshop was conducted. A presentation of each question, the corresponding outputs, and the discussions they provoked follows.

The first question was: "In reference to the ASL context, please provide a comparative assessment of how much each of the following dimensions contributes to Cyber Posture, indicating which are more determinant than others". The practitioners were presented with a framework outlining the six dimensions of cyber posture and were instructed to draw between one and fifteen arrows connecting the various dimensions. For instance, if an expert drew an arrow from 'Governance and Processes' to 'Application, Data, and Network Security' and assigned the value 1 to it, this indicated that, in their view, the two dimensions held equal weight in the overall assessment of cyber posture. Instead, if they assigned value 7, this meant that they considered the first dimension more important than the second one. The overall results were presented at the conclusion of the workshop; however, for clarity and ease of interpretation, the outputs are reported step by step in these sections.

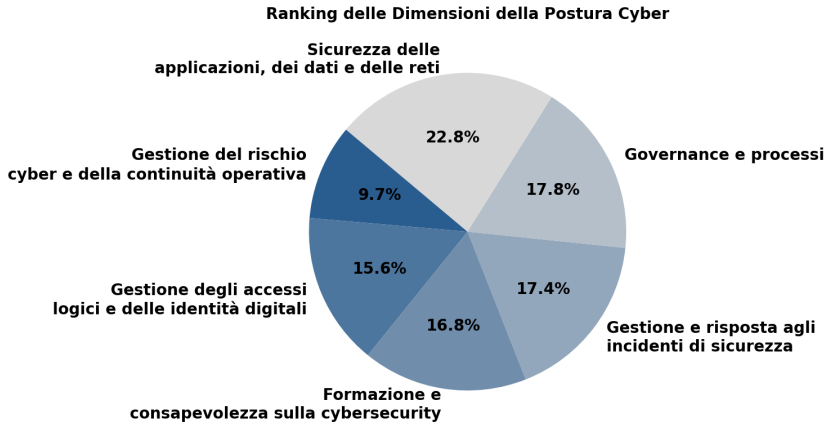

Figure 29 - Output Domanda 1

Experts expressed significant surprise at the results of this first question, especially regarding the weights assigned to "Cyber risk and business continuity management" compared to, for instance, "Application, data, and network

security”. Two participants strongly disagreed with the weights shown in Figure 29, stating that, in their view, the two should have been reversed. During the discussion, it emerged that the responses may have been influenced by the assumed maturity level of the reference organization. For example, the fact that “Cyber risk and business continuity management” ranked last in importance might be attributed to the perception that this is an area already addressed and thus taken for granted. Another interpretation suggested that the low ranking could stem from the assumption that risk management corresponds to a high level of cyber maturity, leading participants to prioritize other dimensions during the survey, under the assumption that the context organization had a low maturity level. Some also suggested that the dimension - particularly the "business continuity" component - may have been misunderstood in relation to its meaning in the ASL context.

The second question was: "Referring to the presented ASL context, please provide a comparative assessment of the increase in overall Cyber Posture effectiveness resulting from transitions between different levels of maturity."

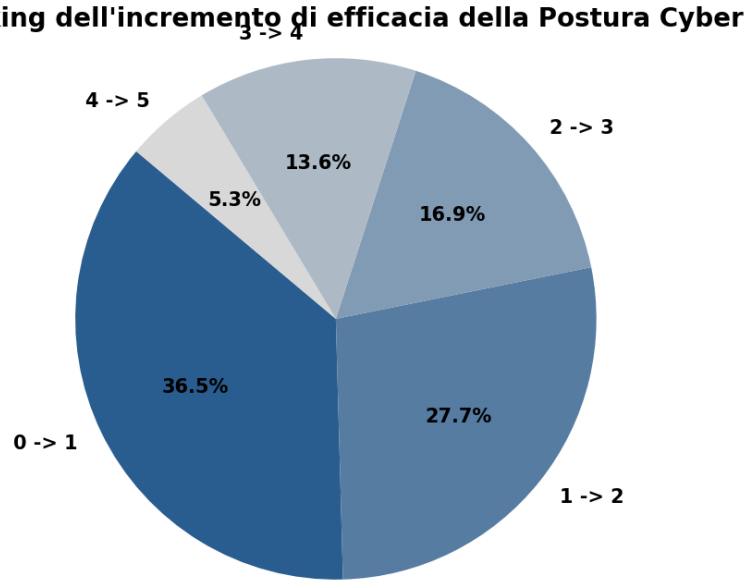

Figure 30 - Output Domanda 2

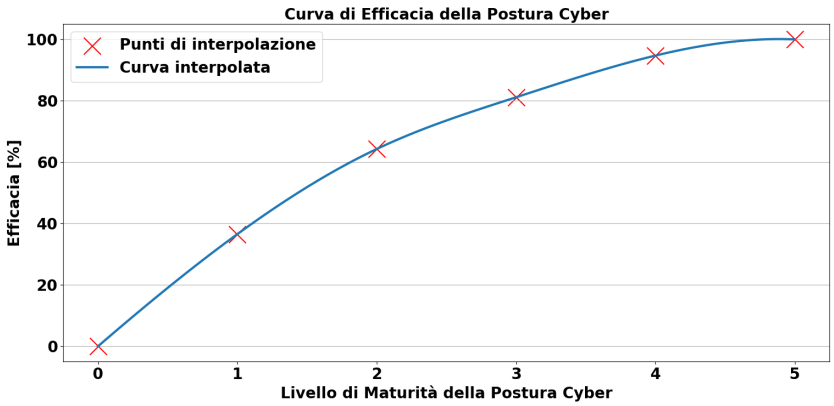

Figure 31 – Processed Output Domanda 2

The second question received unanimous agreement from the experts. In particular, each increase in maturity level was perceived as yielding progressively smaller gains in effectiveness compared to the previous one (Figure 30 and Figure 31). The relationship between cybersecurity maturity and effectiveness was identified as one of the key challenges. There was strong agreement that the initial marginal investments in cybersecurity - such as the transition from level 0 to level 1 - are those with the greatest impact.

Artefact

ROSI per organization

The formula for the Return on Security Investment has been thoroughly explained at this point; however, the actual coding implementation warrants particular attention. Several versions were developed to gradually introduce the complexity of its factors and underlying nature. The initial version is, understandably, considered the baseline. Specifically, this means that no depreciation rates are applied to the benefits. This first ROSI calculation is intended to be performed immediately after the intervention has been completed - at tier 1 - regardless of its timing. The probability of a cyberattack is fixed at 36%, but the effectiveness of the intervention is assumed to be associated to the cyber posture in the way that the workshop outlined. This ROSI is calculated for a single organization. First, the dataset containing the organization's drivers is filtered accordingly. Then, the cyber posture dataset is filtered to extract values measured before and after the intervention. The economic benefits are computed based on the assumptions previously outlined, and the ROSI is then calculated directly using the standard formula. Another version of ROSI, referred to as *ROSI_medio*, was named explicitly to distinguish it from the first formulation reported. In this version, a depreciation rate of 15% was introduced over a five-year time frame to calculate the ROSI from a long-term perspective. This adjustment is particularly relevant, as the previously defined *ROSI_t1* often produced negative values. This outcome is generally expected for most types of investment: initial expenditures tend to be perceived as costs, while their true nature as investments becomes evident over time. The interest rate of 15% and the five-year time horizon are not arbitrarily selected ((Economics, Rates of return to investment in science and innovation, July 2014), (Economics, Rate of Return to Investment in R&D, March 2023)); their justification is provided in the chapter dedicated to economic benefits. It is worth reiterating, however, that the depreciation rate is not applied to the cost of the investments. While this may not align with general accounting practices, it is considered consistent within this case study, as the investment amount corresponds to a pre-allocated budget that remains fixed and is not subjected to inflation. In any case, meaning taking the ROSI at tier 1 or rather the mean of it over a 5-years' time-horizon, other ROSI versions can be tested. These entail the input given to the function for the cost of the attack, from which the benefits are then evaluated. Specifically, three options were identified and used in first place. Intuitively, the first choice was to set a referenced constant as cost of an attack, but then two other paths were considered. A first step

was made using the prepared function cost_attack_org(org), which identifies a more tailored value based on some characterizations of the organization's context, such as the country it is based in rather than whether it is a public sector. Another version is based on the matching between the drivers derived from the funding call and the expected cost of data breaches as reported in the IBM report published in 2024. For the latter, the cost of an attack is defined as the maximum between the predefined default cost (as described in the initial simplified approach) and the cost estimates provided by the IBM report for attacks on organizations with a similar profile to the one under consideration, based on specific characteristics.

ROSI per type of intervention and sub-intervention

Recognizing that calculating ROSI by type of intervention opens the door to further applications and future developments of the prototype, some preliminary attempts were made in this direction - even though the data available may not always be sufficient or entirely appropriate. One issue is due to the fact that the drivers obtained from the funding call are inherently tied to the interventions selected by each participating organization. The call itself was structured around three thematic areas, each offering different types of interventions. Theoretically, it is possible that only one - or even none - of the twelve participating central public administrations opted for a particular intervention, thereby rendering disaggregated evaluations less meaningful or statistically fragile. A more critical challenge, however, lies in the calculation of benefits. This requires determining whether to consider the average increase in cyber posture among organizations that implemented a specific intervention, or to apply a weighted average based on additional criteria. It would also be more meaningful to relate each intervention to the specific dimensions of cyber posture it targets and then assess changes in those particular dimensions before and after the intervention. While the benefit calculation methodology remains an open question, the investment cost component does not present major methodological concerns in this context. The same procedure used to calculate ROSI by intervention type was applied to sub-intervention types, aiming to enhance the granularity and informative value of the ROSI estimates. Additionally, tools were made available to organizations, both at the Tier 1 level and in a version showing the average ROSI over a five-year period, with a 15% annual discount rate applied to the cumulative benefits.

Chapter 7 – Discussion

The roots have been analyzed and discussed, so it is now time to turn to the canopy of the tree. In practice, this involves reflecting on the application of ROSI. Two main branches are identified, from which a wider set of sub-branches can emerge. The distinction between these lies in the fact that the first pertains to the use of ROSI for evaluating past investments, while the second enables applications for the strategic assessment of investments a priori. A dedicated section follows for each, and then the artefact's sources and next steps are also discussed.

The Branch of Ex-Post Economic Evaluation

A disaggregated assessment of the impact of implemented cybersecurity investments enables a more granular understanding of their effectiveness across different contexts. This evaluation can be conducted along three key dimensions. First, by type of organization, recognizing which was the one that operated the best way with the investment made. In a longer term, insights about the sector-specific factors, size, and maturity levels that significantly influence the outcome of security investments should arise with this aggregation. Second, by type of intervention or even sub-intervention, which allows for identifying which categories of cybersecurity measures - such as technological upgrades rather than training programs - worked the best in terms of returns. Lastly, by ROSI bands, to distinguish performance clusters and facilitate benchmarking. This layered analysis supports more informed decision-making and contributes to the development of evidence-based investment strategies.

The creation of historical datasets represents a crucial step for the implementation and further development of the artifact too. In particular, following the investment, a systematic collection of data on an organization's cyber posture at different points in time, as well as disaggregated ROSI values by type of investment, organizational category, or sector, would enable not only more thorough retrospective analyses but also the training of eventual machine learning models. These datasets could serve as valuable input for systems aimed at identifying patterns in organizational features that correlate with high (or low) returns on cybersecurity investments. In this way, beyond fostering a deeper understanding of the key success factors, such efforts would support the development of increasingly tailored and adaptive decision-support tools. Given the context-specific design and customization of the tool, machine learning resources could then be leveraged to identify patterns that influence Return on Security Investment (ROSI) and cybersecurity posture - either positively or negatively. As with all data-driven methodologies, when properly configured, this approach foster greater awareness of the complex, multivariate nature of cybersecurity-related challenges.

Another valuable application of the ROSI factors is the switching point analysis, which refers to identifying the moment in time when the cumulative benefits of an investment equal its initial cost. This analysis can be

conducted even with a basic prototype implementation of the framework, by calculating the time required for an organization to reach this breakeven point. Such insight is particularly useful, as it provides an immediate sense of the average duration needed for a security expenditure to turn into a return-generating investment. For instance, revisiting the PNRR 1.5 case study, it would be straightforward to derive an average switching point across the organizations examined. However, certain methodological specifications must be established. First, there is the definition of a discount rate, which should ideally be grounded in a thorough literature review. For the purposes of this initial implementation, an annual discount rate of 15% ((Economics, Rates of return to investment in science and innovation, July 2014), (Economics, Rate of Return to Investment in R&D, March 2023)) was adopted. This rate is applied progressively to the benefits to calculate their cumulative present value over time, while the investment cost remains fixed in the current model. It should be noted that this analysis will become more nuanced once the cost component of ROSI incorporates indirect and operational expenditure (OPEX), which would require a more dynamic treatment of investment costs over time.

The Branch of Ex-Ante Evaluation of Future Investment Strategies

An economic framework of this kind can serve as a valuable support tool in conducting ex ante evaluations of future investment needs, for the Agency in first place but not only. By providing structured guidance for planning and strategic decision-making that are experience-based, the framework enables the identification and prioritization of areas where cybersecurity investments are most needed. This proactive approach ensures that resources are allocated efficiently and in alignment with broader organizational and national security objectives. Moreover, there are two foundational elements - roots - of the framework that are especially valuable in enabling broader applications. The first is the estimation of attack probability: when derived from a comprehensive mapping of vulnerabilities, it offers meaningful insights into the organization's underlying risk profile. The second is the ability to accurately assess the organization's overall cyber posture at a specific point in time. This capability is crucial, as it enables a picture of the potential return on investment based on the security configuration and maturity at that given moment. Finally, thinking about the utility of this research and its long-term option of being user-oriented at a managerial level, the development of a tool can be run in parallel. By leveraging organization-specific data, such a tool could indeed offer tailored assessments of the Return on Security Investment (ROSI), enhancing decision-making through contextualized insights and promoting more informed cybersecurity investment strategies.

Data Sources

An important disclaimer: this research involves a part of coding to process available data, and indeed a dedicated chapter was also at disposal for getting more information on this side. However, it must be stated that the available data is insufficient for conducting an analysis of historical trends, at the point that also the results got from the

ROSI calculations are limited in terms of quality because of the shortage of appropriate and granular data at disposal. Conscious of this but not losing any trust in the potential valuable application that would be following by a more advanced version of the artefact, this research aims to establish a complete and generalizable framework. Again, this is done in two directions: first, in terms of its applicability to the given context, and second, in terms of scalability. It is regarding scalability that the current lack of extensive data becomes relevant. While historical trends are not yet discernible, this does not hinder the possibility of designing a structured framework that will be ready to accommodate such data when it becomes available, thereby enabling the system to scale up and provide deeper insights. As mentioned, an important concern regards the current lack of historical data to integrate into the model. The ADR process incorporated a workshop to leverage this, as already mentioned, and its outcomes are valuable both independently and as integral components of the model. It served as an opportunity to refine the framework and model by incorporating the reliable insights of practitioners. Along this line, a recommendation for future iterations of the research concerns the sources of ongoing updates. Specifically, it would be highly beneficial to design and manage structured surveys targeted at domain experts, enabling the integration of key informant knowledge into the model. This would reinforce its strength as a practice-inspired and experience-based tool. Furthermore, iterative engagement with the community of practice could help prioritize the most critical roots and branches, while also uncovering potential new extensions to the framework.

Inspired by what has been done for ‘Avviso 2’ but also thinking about giving instructions to have suitable data for the artefact to be upgraded in the future, the availability of questionnaires assessing the organizations’ cyber posture must be stressed as one of the most important data to have. Appropriate multiple-choice inquiries would enable an analysis of their risk exposure in relation to the organizations’ specific characteristics. Considering the case study, a description of the current setup of the questionnaire in use is provided. It is divided into three sections: security management, IT security, and physical security. However, the topics covered in each section are not balanced in terms of quantity: more than half concerns IT, followed in terms of number of questions by security management with sixteen questions, and lastly there are also some about physical security. As emphasized in the literature (Johnson, et al., 2024), cybersecurity is not solely a technical matter but also encompasses managerial aspects. For this reason, the prominence of security management within the questionnaire is justified. This is particularly significant in systems characterized by complex information infrastructures due to their size and/or field of operation. In such cases, the interactions among teams and departments introduce additional vulnerabilities, making a clear framework for securing communications a high priority. Consequently, this imbalance may influence the profiling of organizations discussed later, as the characteristics determining the relative weight of each section should be considered. There is no necessity to increase the number of sections in the questionnaire; rather, it is more beneficial to explore each one in greater depth and to organize them further by structuring subsections. This approach would ensure a more comprehensive assessment of the organization’s

cyber posture. For the future design of questionnaires: each question within each section should be accompanied by two dedicated columns - the first indicating whether the question is applicable to ICT services, and the second to ICT products (including both software and hardware). Two additional columns could then link each question to specific controls or requirements within ISO/IEC 27001 and the National Framework of Cybersecurity (FNCS). This connection would be fundamental to ensure the artefact aligns with both national and international guidelines, enabling compatibility with other tools and facilitating the development of more advanced instruments. Lastly, a series of multiple-choice options - as used in the 'Avviso 2' - could allow stakeholders to select the option that best describes their organizational setup. This form of classification could be particularly useful in the early iterations of the model. This is especially relevant given that the most important secondary source is the IBM report released in 2024, which quantifies the cost of attacks based on certain drivers, classifying the organizations examined through variable-based ranges. As soon as a questionnaire is structured, concerns arise around whether it is appropriate for obtaining a realistic picture of the organization's cyber posture. Given that it is a qualitative approach, a strict method for updating the questions and organizing them systematically must be proposed, aiming to ease the integration of further extensions and the identification of missing points. For example, given that the collection of primary data regarding threats and successful attacks is growing and hopefully its quality is also ensured to be enhanced in the next future thanks to reliable sources (e.g., ENISA and ACN reports about cyberattacks), a direct interaction with that data can be introduce in the recurrent - probably annual - updates of the questionnaire. In practice, this means that whenever a trend is identified among the attackers, such as the specific interest in some areas, related inquiries could be added to the questionnaire, to increase the level of precision of the cyber posture, leaving less space to interpretation.

Next Steps

The next steps of this research are structured around two complementary axes - methodological refinement and practical application, the roots and the canopy respectively - each contributing to the evolution and operationalization of the proposed framework. On the methodological front, future work would entail a more granular analysis of intervention costs, disaggregated by both technology type (e.g., Generative AI, Quantum Computing) and specific cyber solutions. Parallel to this, the economic benefit dimension would be expanded through enhanced estimations of attack probabilities, based on primary data sources and differentiated by attack type and targeted assets. A key advancement would be the incorporation of a more nuanced model for estimating the cost of an attack, including indirect impacts and leveraging data from an upcoming survey rather than mapping the Personal Identifiable Information of the organization considered. Also, a probabilistic approach in quantifying the factor of probability of attack would enable a more robust and context-sensitive evaluation. Finally, the framework would benefit from a reassessment of cyber posture effectiveness, supported again by new surveys, targeting a representative sample of both public and private organizations. Besides the value of ROSI for the

framework, it must be stressed that this ADR is enhancing indirectly the evaluation of the cyber posture of an organization. A more accurate assessment of an organization's cybersecurity maturity enables CISOs and other security leaders to plan more effectively and to target risk mitigation strategies with greater precision. This reality accounts this research for its relevancy, unveiling that each factor has indeed a consistent value itself from a managerial point of view in this complex field, which is cybersecurity. The use of such framework is fundamental for ensuring measurement rigor and enabling comparisons between organizations, thereby allowing iterative assessment of intervention effectiveness. In addition to aiding the initial step of selecting interventions, the artefact would also promote greater uniformity in the collection of data necessary to analyze the cyber posture of organizations, thereby improving the capacity to monitor and strengthen cybersecurity over time. On the application side, as mentioned the framework could be transformed into a practical tool for deployment in both public and private sector organizations. Moreover, the goal is to generalize its use within broader policy frameworks, including national-level strategies[14]. This dual track of operationalization and generalization will significantly enhance the utility and scalability of the artefact. Taken together, these developments underscore a forward-looking research agenda that not only deepens the academic contribution but also ensures real-world applicability and policy relevance.

With the aim of underling the extent to which the framework is designed beyond the code of the artefact that have been written by the candidate independently from the case study, Figure 32 is left. This comprehensive schema wraps up the framework for enabling the calculation of the Return on Security Investment (ROSI), meticulously organized into the Probability of Attack, the Cost of Attack, and Effectiveness. The boxes highlighted in green represent the candidate's implemented coding (independent from the case study), demonstrating a concrete application of specific methodologies within the broader framework. The white sections, conversely, illustrate the candidate's proposed extensions and future enhancements for the framework, encompassing additional feedback alternative methodologies (e.g., "Probabilistic Approach (dynamic probability)" for the probability of attack, and surveys for total cyber posture in the factor of effectiveness). Finally, loops allow continuous refinement and adaptation based on enhanced and more accurate contextual information, aimed at tailoring the model as much as possible. The two dotted arrows crossing different factors express the need for alignment between them; in contrast, the dashed line refers to optional paths that substitute the standard setting or can simply be added on top of it. A final disclaimer: considering the proportion between the green and white sections, the overall volume of the proposed extensions is not fully appreciated relative to the potential impact that some of these changes could have as turning points for enhancing the quality of the artefact.

[14] https://www.cybersecurityframework.it

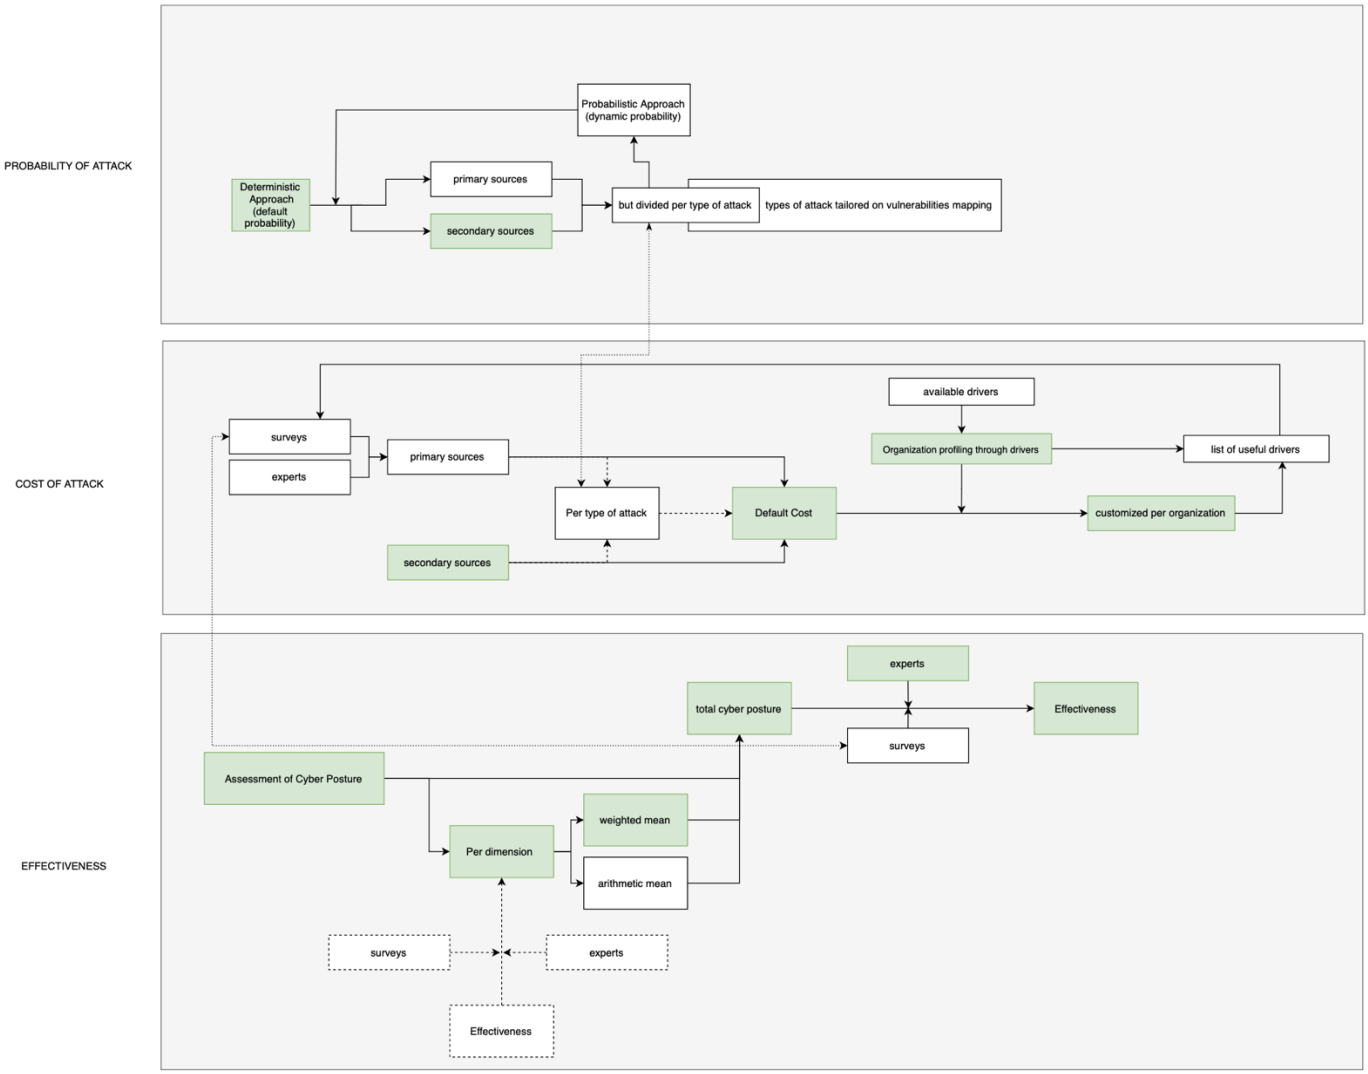

Figure 32 - Extension of the artefact

In conclusion, returning to the starting point to reconnect with the metaphor of the tree, a similarity that does not concern the architecture itself but rather the timing of growth is highlighted. Just as trees are sensitive to seasonal cycles in their growth, the development of this framework should also be continuous over time - not necessarily constant, but iterative - gradually expanding both in scope and structural solidity, step by step. There is indeed room for further development. One of the key takeaways from the workshop - which served as a highly credible opportunity to aggregate informed knowledge - was the identification of a prevailing gap in the sector. Currently, there are too few models that employ ROSI as a tool at the managerial level, and the community of practice largely relies on methods that lack sufficient quantitative rigor. The experts expressed a clear need for a framework that begins qualitatively but evolves through quantitative development, with the ultimate goal of enabling decision-makers to adopt a more data-driven approach. This, ultimately, is what this master's thesis set out to explore.

The Return on Security Investment (ROSI) framework can be effectively utilized for scenario analysis[15], enabling organizations to evaluate the potential economic outcomes of different cybersecurity investment strategies under varying conditions. Theoretically, by modeling for example diverse threat landscapes, investment levels, and organizational responses, scenario analysis helps decision-makers understand how changes in external and internal factors impact the expected return on security expenditures. This approach fits perfectly with this framework: it facilitates more informed strategic planning by highlighting which investments yield the greatest resilience or cost-benefit balance across possible alternatives. Furthermore, scenario analysis with ROSI supports risk management by identifying vulnerabilities and opportunities within specific contexts, ultimately guiding organizations to allocate resources more efficiently and proactively anticipate evolving cyber risks.

[15] https://www.ncsc.gov.uk/collection/risk-management/using-cyber-security-scenarios

Chapter 8 – Conclusions

This Action Design Research (ADR) project was initiated in response to a need identified within the community of practice, specifically when the National Cybersecurity Agency (ACN) began developing an advanced methodology for evaluating investments under PNRR 1.5. From the outset, the potential for broader applicability of the artefact was evident, and the development was guided by a mindset aimed at ensuring its generalizability from the very beginning. The workshop held in May 2025 served as an initial, positive validation that the theoretical direction of the research aligned with practical needs. The results of its questionnaires reflected the researchers' initial intuitions about the key issues, which were those same issues that the model was designed to address and refine. The ADR is currently in its third iteration, and we hope it can serve as a catalyst for further inquiry and deeper understanding of the topic. From an ethical and societal perspective, this project has the potential to offer the country a strategic advantage in technological development. By enabling a more tangible economic evaluation of security investments, the tool may help align technological advancement with effective risk mitigation. In conclusion, beyond the modularity that has been extensively discussed, we wish to highlight a key strength of this research: its foundational commitment to methodological integration. It is grounded in the recognition that while qualitative methods offer valuable contextual insights, they are insufficient on their own, and the same holds true for quantitative approaches. The aim, therefore, is to bridge these two paradigms by translating qualitative, context-specific information into quantitative values, enabling the calculations to be tailored to the organization under consideration. Moreover, this objective is pursued through multiple methodological pathways, balancing and integrating diverse techniques. While this may appear to complicate the process, it does not. In mathematics, the principle of "less is more" often holds true, and we generally find this notion both intellectually compelling and inspirational. However, when addressing complex socio-technical problems, excluding an entire methodological perspective may lead to incomplete or biased outcomes. In this context, mediation emerges as a particularly valuable approach. The goal is not to minimize the number of variables or data sources, but rather to understand what each source, be it a key informant, a data stream, or a method, can best contribute. Simplicity is indeed vital, especially in formulating questions; but mediation serves as the mechanism through which multiple valid perspectives can be reconciled into a coherent and actionable solution.

Acknowledgements (Ringraziamenti)

Per voi sono *amore*, *la nostra piccola*, *sorellina, Lelle*, *la mia amica*, *stelin*, *svampi*, *Elenina*, *Ele, heartquake*, *Elenuccia*. Voi siete la mia geografia.

In chiasmo con il tuo esserci dall'inizio, ecco alla fine qualcosa solo per te. Con te, la parola 'sempre' non è utopia. Avremo gli occhi lucidi mentre in silenzio mi sistemerai la corona d'alloro sulla nuca.

Bibliografia

Szczepaniuk, E., Szczepaniuk, H., Rokicki, T., & Klepacki, B. (2019). Information Security Assessment in Public Administration. *Computers & Security*.

ACN. (2023). *Relazione Annuale al Parlamento 2023.*

ACN. (2024). *La minaccia cibernetica nel settore sanitario.*

ACN. (2024). *Operational Summary - Maggio 2024.*

ACN. (2024). *Relazione Annuale al Parlamento.*

Šimec, A. (2019). Cyber-attacks and Internet of Things as a threat to critical infrastructure. *ECONOMIC AND SOCIAL DEVELOPMENT*, 108-111.

Apruzzese, A., Chiantore, L., Cicognini, M., Leoncin, M., Molinari, G., Marchetti, M., & Andreolini, M. (2024). *Rapporto CLUSIT 2024 sulla Sicurezza ICT Italiana.* Milano.

Armin, J., Thompson, B., & Kijewski, P. (2016). Cybercrime Economic Costs: No Measure No Solution. *Combatting Cybercrime and Cyberterrorism: Challenges, Trends, and Priorities*, 135-155.

Arora, A., Telang, R., & Frank, S. (2007). Estimating Benefits from Investing in Secure Software Development. *Software Engineering Carnegie. Mellon University Institute and US Department of Homeland Security*.

Arshad, A., Abbas, H., Faisal Amjad, M., Shafqat, N., & Yaqoob, T. (2019). Framework for Calculating Return on Security Investment (ROSI) for Security-Oriented Organizations. *Futur Generation Computer System*, 754-763.

Böhme, R., & Moore, T. (2010). The Iterated Weakest Link. *IEEE Security and Privacy Magazine*, pp. 53-55.

Bahugunaa, A., Bishtb, R. K., & Pandec, J. (2020). Country-level cybersecurity posture assessment:Study and analysis of practices . *NFORMATION SECURITY JOURNAL: A GLOBAL PERSPECTIVE*, 250-266.

Barik, K., Misra, S., Fernandez-Sanz, L., & Koyuncu, M. (2023, 10 14). RONSI: a framework for calculating return on network security investment. *Telecommunication Systems*, pp. Vol. 84; p. 533–548.

Butler, S. A. (2002). Security Attribute Evaluation Method: A Cost-Benefit Approach. *Proceedings of the 24th international conference on Software engineering, ACM*, 232–240.

Clusit. (2025). *Rapporto 2025 Clusit sulla Cybersicurezza in Italia e nel mondo.*

Collier, Z. A., Briglia, B., Slutzky, D. L., & Lambert, J. H. (2023). On metrics and prioritization of investments in hardware security. *Systems Engineering*, pp. 425-437.

Commission, E. (2024). *Italy 2024 Digital Decade Country Report.*

Cremonini, M., & Martini, P. (2005). Evaluating Information System Investments from Attackers Perspective: the Return-On-Attack (ROA).

Damjan Fujs, I. B. (2024). Analyzing Cybersecurity Strategies of the European Union: Challenges and Opportunities for Public Administration. *ELEKTROTEHNIŠKI VESTNIK* , Vol. 91(1-2), pages 8-20.

David Kroodsma, B. S. (2021, May 11). Global Fishing Watch.

Economics, F. (July 2014). *Rates of return to investment in science and innovation.* London: Frontier Economics Ltd.

Economics, F. (March 2023). *Rate of Return to Investment in R&D.* London: Frontier Economic Ltd.

Ekelund, S., & Iskoujina, Z. (2019). Cybersecurity economics – balancing operational security spending. *Information Technology & People*, pp. Vol. 32 No. 5, pp. 1318-1342.

Erolaa, A., Agrafiotisa, I., Nurseb, J. R., Axona, L., Goldsmitha, M., & Creese, S. (2021). A system to calculate Cyber Value-at-Risk. *Computers & Security*.

Fan, J., Zhanga, P., & Yen, D. C. (2013). G2G information sharing among government agencies. *Information & Management*.

Franco, M. F., Künzler, F., Von der Assen, J., Feng, C., & Stiller, B. (2024). RCVaR: An economic approach to estimate cyberattacks costs using data from industry reports. *Computers & Security*.

Fujs, D., & Bernik, I. (2024). Analyzing Cybersecurity Strategies of the European Union: Challenges and Opportunities for Public Administration. *Elektrotehniski Vestnik*, Vol. 91(1-2), pages 8-20.

Ganapati, S., Ahn, M., & Reddick, C. (2011). Evolution of Cybersecurity Concerns: A Systematic Literature Review. *Proceedings of the 24th Annual International Conference on Digital Government Research*, 90-97.

Gao, X., Zhong, W., & Mei, S. (2013). Security investment and information sharing under an alternative security breach probability function. *Information Systems Frontiers*, Vol. 17, pp. 423-438.

Gheorge, M. (2012). Investment Decision Analysis In Information Security. *Révista Economica - Information Security Management* , pp. Vol. 13; p.85-93.

Global Cybersecuirty Index. (2024). Retrieved from ITU: https://www.itu.int/en/ITU-D/Cybersecurity/pages/global-cybersecurity-index.aspx

(2024). *Global Cybersecurity Index 2024.* ITUPublications.

Gordon, L. A. (2003). A framework for using insurance for cyber-risk management. *Communications of the ACM*.

Gordon, L. A., & Loeb, M. P. (2002). The Economics of Information Security Investment. *ACM Transactions on Information and System Security* .

Hausken, K. (2006). Returns to information security investment: The effect of alternative information security breach functions on optimal investment and sensitivity to vulnerability. *Information System Frontiers*, Vol. 8, pp. 338–349.

He, Y., Xin, T., & Luo, C. (2024). Cybersecurity Investments Metrics using FAIR-ROSI Cybersecurity Investments Metrics using FAIR-ROSI. *UK Academy for Information Systems Conf erence Proceedings 2024*.

Hiscox. (2022). *Cyber Readiness Report 2022.*

Johnson, V., Maurer, C., Torres, R., Guerra, K., Mohit, H., Srivastava, S., & Chatterjee, S. (2024). The 2023 SIM IT Issues and Trends Study The 2023 SIM IT Issues and Trends Study. *MIS Quarterly Executive*, pp. Vol. 23: Iss. 1, Article 7.

Jonhson, V., Maurer, C., Torres, R., Guerra, K., & Mohit, H. (2024). The 2023 SIM IT Issues and Trends Study The 2023 SIM IT Issues and Trends Study. *MIS Quaterly Executive*.

Kryshtanovych, M., Andriyash, V., Bondar, H., Kushnir, Y., & Ozarko, K. (2022). Public Administration Mechanisms for Ensuring Cybersecurity in Modern Conditions of Socio-Economic Development. *IJCSNS International Journal of Computer Science and Network Security*, VOL.22 No.3.

Liu, G., Zhang, J., & Chen, G. (19, 09 2014). An approach to finding the cost-effective immunization targets for information assurance. *Decision Support Systems*, pp. 40-52.

Manley, M. (2015). Cyberspace's Dynamic Duo: Forging a Cybersecurity Public-Private Partnership. *Journal of Strategic Security*, Vol.8.

Marican, M. N., Othman, S. H., Selamat, A., & Razak, S. A. (2024). Quantifying the Return of Security Investments for Technology Startups. *Baghdad Science Journal*, pp. 2449-2461.

Mattiolli, R., Malatras, A., Hunter, E. N., Biasibetti Penso, M. G., Bertram, D., & Neubert, I. (2023). Identifying Emerging Cyber Security Threats and Challenges for 2030.

Munari, B. (1978). *Disegnare un albero.* Edizioni Corraini.

Okoli, C. (2015). A Guide to Conducting a Standalone Systematic Literature Review. *Communications of the Association for Information Systems*, 37(43), 879-910.

Oliva, G., Faramondi, L., Setola, R., Tesei, M., & Zio, E. (2021). A multi-criteria model for the security assessment of large-infrastructure construction sites. *International Journal of Critical Infrastructure Protection*.

Onwubiko, C., & Onwubiko, A. (2019). Cyber KPI for Return on Security Investment. *International Conference on Cyber Situational Awareness, Data Analytics And Assessment (Cyber SA)*.

Onwubiko, C., & Onwubiko, A. (n.d.). Cyber KPI for Return on Security Investment.

Organisation, E. C. (2025). *ECSO 2024 Investment and M&A Report.*

PA digitale 2026. (2025). Retrieved from Dipartimento per la trasformazione digitale : https://padigitale2026.gov.it/misure/

Pontes, E., Guelfi, A. E., Silva, A. A., & Kofuji, S. T. (2011). A Comprehensive Risk Management Framework for Approaching the Return on Security Investment (ROSI). In *Risk Management in Environment, Production and Economy* (p. chapter 7).

(2024). *Rapporto Clusit - Italia e PA.* Milano.

Schatz, D., & Bashroush, R. (2017). systematic review of security investment evaluations method. *Information Systems Frontiers*.

Security, I. (2024). *Cost of a Data Breach Report 2023.* IBM.

Sein, M. K., Henfridsson, O., Purao, S., Rossi, M., & Lindgren, R. (2011). Action Design Research. *MIS Quarterly* , Vol. 35 No. 1 pp. 37-56.

Shandler, R., & Gomez, M. A. (2023). The hidden threat of cyber-attacks – undermining public confidence in government. *Journal of Information Technology & Politics*, 20:4, 359-374.

Sienkiewicz, P. (2010). Systems analysis of security management. *Scientific Journals, Maritime University of Szczecin*.

Sonnenreich, W., Albanese, J., & Stout, B. (2006). Return On Security Investment (ROSI) – A Practical Quantitative Model. *Journal of research and practice in information technology.*, Vol.38(1), p.45-56.

Spagnoletti, P., Resca, A., & Sæbø, Ø. (14, May 2015). Design for social media engagement: Insights from elderly care assistance. *Journal of Strategic Information Systems*, pp. 128-145.

Sukumar Ganapati, M. A. (2011). Evolution of Cybersecurity Concerns: A Systematic Literature Review. *Proceedings of the 24th Annual International Conference on Digital Government Research* , 90-97.

Tsiakis, T., & Stephanides, G. (2005). The economic approach of information security . *Computers & Security*, pp. 105-108.

Webster, J., & Watson, R. T. (2002). Analyzing the past to prepare for the future: Writing a literature review. *MIS Quarterly*, 26(2), xiii-xxiii.

Wirtz, B. W., & Weyerer, J. C. (2017). Cyberterrorism and Cyber Attacks in the Public Sector: How Public Administration Copes with Digital Threats. *International Journal of Public Administration*, 1085-1100.

Zhang, A., Collins, R., & O'Connor-Close, C. (2020). Cyber incident cost estimates and the importance of building resilience. *Reserve Bank of New Zealand Bulletin*, Vol. 84, No. 2.

Appendix A

Agenda

1. Investimenti Cyber: contesto e criticità
2. Action Design Research
3. Proposta di framework di valutazione economica
4. Workshop
5. Conclusioni

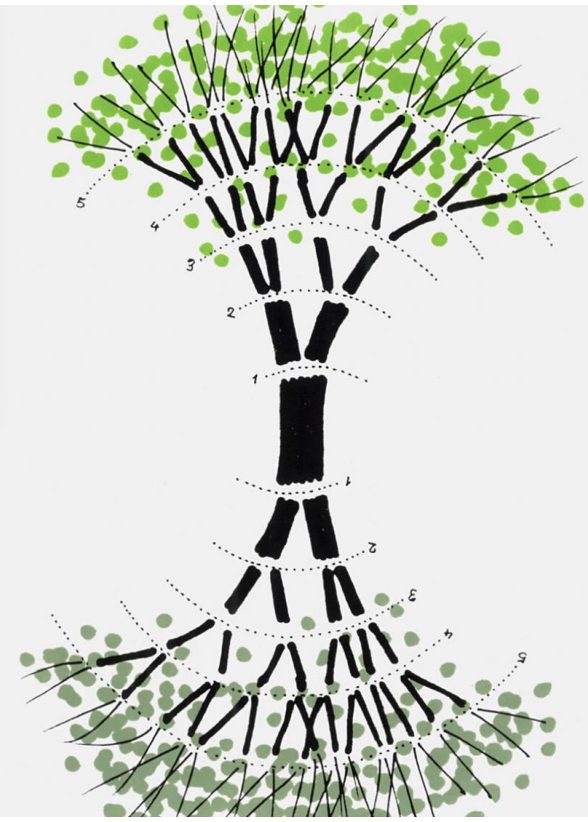

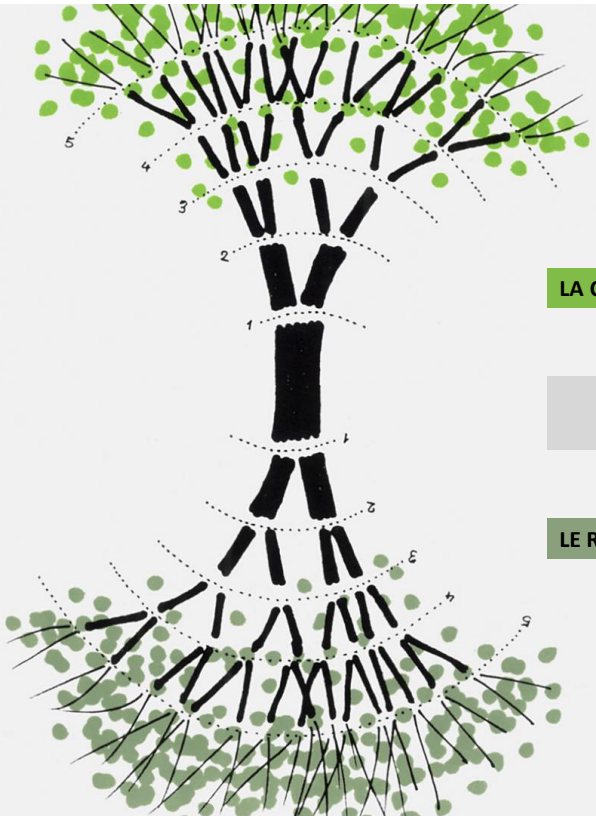

LA CHIOMA – le applicazioni del framework

$$ROSI = \frac{Beneficio\ Economico - Costo\ dell'Investimento}{Costo\ dell'Investimento}$$

LE RADICI – la metodologia del framework

LA CHIOMA – le applicazioni del framework

Valutazione disaggregata dell'impatto degli investimenti effettuati:
- Per tipologia di intervento
- Per tipologia di organizzazione
- Per fasce di ROSI

Creazione di **dataset** storici e di **stime** utili per implementare e aggiornare l'artefatto

VALUTAZIONE ECONOMICA A POSTERIORI DEGLI INVESTIMENTI

Framework economico a supporto dell'Agenzia nella valutazione dei fabbisogni di investimento futuri, nella pianificazione e nella definizione delle strategie

Tool utilizzabile dagli utenti per una valutazione economica basata sulla propria Postura Cyber

VALUTAZIONE ECONOMICA DI FUTURE STRATEGIE D'INVESTIMENTO

ACN LUISS

LE RADICI – la metodologia del framework

COSTO DEGLI INVESTIMENTI

- Capex
- Opex
- Costi diretti
- Costi indiretti

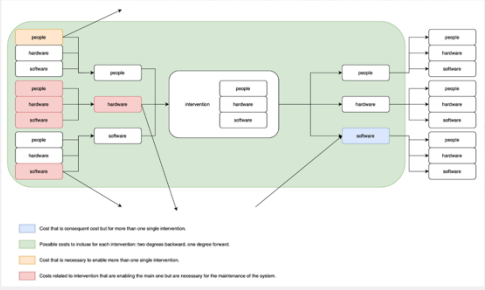

$$EB = P_A * Eff * C_A$$

BENEFICIO ECONOMICO (EB) DELL'INVESTIMENTO

P_A

Probabilità di un attacco cyber

- Fonti secondarie
- Specifico per tipo di attacco

Eff

Efficacia del livello di Postura Cyber

- Probabilità di proteggersi completamente dagli attacchi cyber
- Associata ad un livello di maturità

Output workshop

C_A

Costo economico di un attacco cyber

- Fonti secondarie
- Specifico per organizzazione

ACN LUISS

Maturità della Postura Cyber

DOMANDA 2

Postura Cyber

Lo stato di sicurezza delle informazioni, dei sistemi e delle reti di un'organizzazione, basato sulle risorse di Information Assurance (IA) (es. persone, hardware, software, politiche) e sulle capacità messe in atto per gestire la difesa dell'organizzazione e reagire al mutare della situazione.

Tradotto da: National Institute of Standards and Technology. (2011). Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations (NIST SP 800-137). U.S. Department of Commerce. https://nvlpubs.nist.gov/nistpubs/legacy/sp/nistspecialpublication800-137.pdf

LIVELLI DI MATURITÀ

MATURITÀ		DESCRIZIONE
INCOMPLETO	0	*Controlli non implementati o parzialmente implementati*
CONSIDERATO	1	*L'implementazione dei controlli è affidata a processi, procedure e soluzioni tecniche con risultati non prevedibili e non documentati. La gestione è affidata alle singole competenze del personale e non all'uso comprovato di processi ben definiti*
DEFINITO	2	*L'implementazione dei controlli si avvale di processi, procedure e soluzioni tecniche ben definiti e documentati nelle funzioni dell'organizzazione coinvolta, ma ciascuna funzione gestisce i propri processi, procedure e soluzioni tecniche in modo indipendente.*
AVVIATO	3	*L'implementazione dei controlli si avvale di processi, procedure e soluzioni tecniche ben definiti, documentati e standardizzati a livello di normativa interna dell'Amministrazione.*
IMPLEMENTATO	4	*Oltre a includere gli aspetti del livello di maturità "Avviato", sono fissati degli obiettivi quantitativi per quanto riguarda le performance dei processi, delle procedure e delle soluzioni tecniche alla base dell'implementazione dei controlli.*
OTTIMIZZATO	5	*Oltre a includere gli aspetti del livello di maturità "Implementato", i processi, le procedure e le soluzioni tecniche alla base dell'implementazione dei controlli sono sottoposti a miglioramento continuo in risposta a cambiamenti nell'Amministrazione e considerando le esperienze passate.*

Agenzia per la Cybersicurezza Nazionale. (2023). Relazione annuale al Parlamento 2023. ACN. https://www.acn.gov.it/portale/documents/20119/446882/ACN_Relazione_2023.pdf

Efficacia della Postura Cyber

DOMANDA 2

Relazione tra efficacia* e livello di maturità di Postura Cyber

*probabilità di proteggersi completamente dal rischio cyber

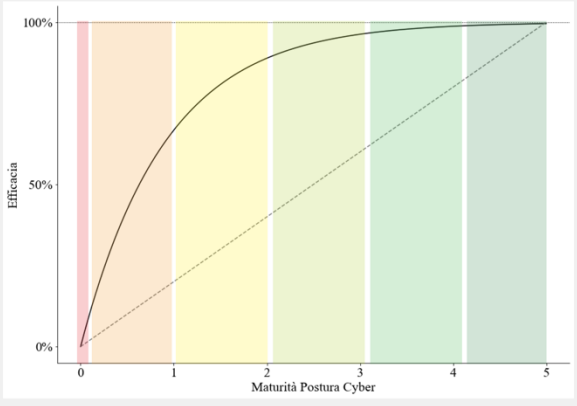

Esempi illustrativi

1 Dimensione di Postura Cyber di tipo *labour-driven*

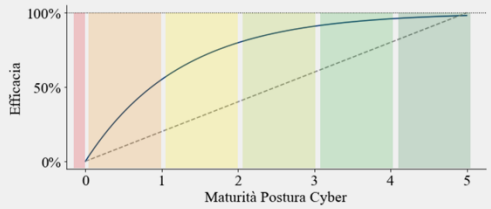

2 Dimensione di Postura Cyber di tipo *capital-driven*

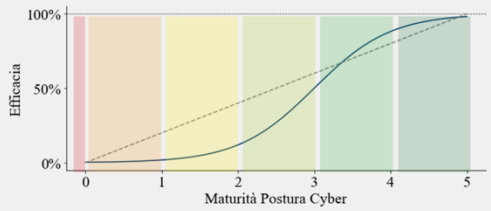

Contesto - ASL

APPLICAZIONE AD UN CONTESTO

Si immagini di assumere il ruolo di Chief Information Security Officer (CISO) presso una Azienda Sanitaria Locale (ASL).

Profilazione

Nome: ASL Metropolitana Nord (ASL-MN)
Tipo di ente: Azienda Sanitaria Locale pubblica
Localizzazione: Nord Italia (area urbana e periferica)
Dipendenti: 3.000 (tra medici, infermieri, amministrativi, tecnici)
Servizi: Ospedali, ambulatori, pronto soccorso, medicina del lavoro, vaccinazioni, sistema CUP (prenotazioni)
Fatturato medio: €800 milioni annui

Infrastruttura

- Infrastruttura centralizzata in un data center interno
- Accessi remoti per medici e infermieri territoriali
- Software gestionali (sanitari e amministrativi), sistemi PACS (immagini mediche), cartelle cliniche elettroniche
- Fornitori esterni per manutenzione IT, sistemi diagnostici, e software gestionali

Obiettivi sensibili

- Dati sanitari dei pazienti (GDPR)
- Servizi salvavita (emergenze, pronto soccorso)
- Infrastruttura critica (server ospedalieri, apparecchiature biomedicali connesse in rete)
- Portali di prenotazione online

Esercizio AHP

DOMANDA 1

«Relativamente al contesto ASL, La invitiamo a esprimere una valutazione comparativa dell'incidenza delle seguenti dimensioni sulla Postura Cyber, indicando quali risultano più determinanti rispetto alle altre.»

- Le preferenze sono definite attraverso frecce orientate
- È possibile inserire da **1 a 15 preferenze**
- Il valore della preferenza va assegnato secondo la tabella della **scala di Saaty**
- Il valore di incertezza può essere omesso

Scala di Saaty	
1	Uguale importanza
3	Moderata importanza
5	Importanza forte
7	Importanza molto forte
9	Estrema importanza