



UNIVERSITY
OF TWENTE.

Department
of Information Engineering and Computer Science

Master's Degree in

Cybersecurity

FINAL DISSERTATION

**STUDY ON THE CYBER RISK
INSIDE A MUSEAL ENTITY**

Supervisor at University of Trento

Paolo Giorgini

Supervisor at University of Twente

Florian W. Hahn

Student

Alessandro Gnutti

ACADEMIC YEAR 2020/2021

Contents

1. Introduction	3
2. Background	5
3. Methodology	7
3.1 Cyber risk evaluation	7
3.1.1 Terminology	8
3.1.2 Risk	11
3.2 The process	12
3.3 Study on employees' behaviours in regard to cybersecurity	12
3.3.1 The questionnaire	12
3.3.2 The phishing test.....	15
4. The results	18
4.1 The risk evaluation - Set up	18
4.2 The risk evaluation - Frequency of a loss event	27
4.3 The risk evaluation - Loss magnitude.....	34
4.4 The risk evaluation - Risk	42
4.4.1 Risk for the Data asset.....	42
4.4.2 Risk for the Web Applications asset	43
4.4.3 Risk for the IT Infrastructure asset.....	44
4.5 The questionnaire and phishing test	45
4.5.1 Policy knowledge and hacker perception	46
4.5.2 Phishing, emails and malicious attachments	48
4.5.3 Malware attacks and induced malfunctions.....	49
4.5.4 Web security.....	49
4.5.5 Password.....	50
4.5.6 Working station security	51
4.5.7 The phishing test.....	52
4.5.8 Results conclusion.....	53
5. Evaluation.....	54
6. Conclusion	64
7. Acknowledgements	65
References	66

1. Introduction

“A building in which objects of historical, scientific, artistic, or cultural interest are stored and exhibited.”
[Oxford Dictionary 2021]

A museum is a place which stores and exhibits artifacts from the past and the present of some kind of interest to society; they hold the traces and remains of how we came to the here and now, they hold our history. History. The story of our places, our societies... of our race as a whole. It is often said we can learn from the mistakes of the past, from what we have done before. In science this is always true: research builds on previous research to go forward, to find the something new that had not been found before. However, the same concept should apply to all other aspects of our societies, such as politics and sociality. Museums can help in spreading knowledge and awareness on many aspects of our world and societies; by presenting our history to the wider public, they can help the wider, not technical population to understand what had been done before, what is being done now and hopefully help in building a better future.

If museums had historically taken a passive approach to showing our history, now, we often see a more proactive approach on their part, to be more present in the social fabric of the cities they are located in. However, museums are fragile structures, not just from a physical point of view as they can be at risk from the environmental changes, but from social phenomena's too; the more active they become in the fabric of society, the more their overall risk grows. Risk in the wider sense of the term. Not just risk of being broken into, or the risk to their physical collections, but risk to their IT infrastructure, to their data and the information they store too.

We live in a world that is growing more and more interconnected: the internet is pervading every single aspect of our lives, from the smallest of things like IoT sensors, to the greater distributed systems that help running complex systems key to our lives and economies. As museums update themselves and grow more present in the fabric of society, it is obvious how they will, and they are, embracing information technologies with their upsides and downsides: employing information technologies can smoothen many operational aspects on one side, but it can provide a greater interactivity and appetability to the end users too. However, this opening to the technologies shaping our world comes with a price: an increased risk and a new set of threats and vulnerabilities.

These new threats and vulnerabilities are an integral part of the cyberspace, the notional environment in which communication over computer networks occurs, the same way the weather is an integral part of the physical world. They are ever present and they have to be dealt with much in the same way weather events are dealt with: either by accepting them or mitigating them.

However, these threats and vulnerabilities have one big difference compared to all the threats and vulnerabilities museums have worked with till now: they are not visible to the naked eye. You cannot see with your naked eye the vulnerability lying in a website in the same way you can see with the naked eye the fact that a door has not a proper lock. You have to inspect the code building the website to try see such vulnerability and as often has happened, you might not even see it till you test it alive, till someone tries to take “a different path” than the one you encoded and expected them to take. And the damage is often visible only afterwards, with a time that goes from minutes and hours up to days or at worse weeks or months, after the damage is done and when only “damage control” is possible.

Museums are vulnerable like every other entity connected to the internet, like every entity open to the outside world. Certain actors that arose to prominence in the cyberspace landscape do not make any difference between museums, businesses and people: if an entity has something they are interested in, they attack it. It had reached the headlines before that there had been cyber attacks against museums, such as the Asian Art Museum of San Francisco [Cascone 2019] or the Smithsonian Institution and the United Kingdom's National Trust [Cascone 2020].

Museums hold our past, hold the mistakes and the greatness of it, and as such they must be protected from these new threats and vulnerabilities that have arisen to prominence. They must build cyber security programs much in the same way business do, to protect their systems, their employees and users; they have to make their

employees aware of such threats, vulnerabilities and programs. In addition, they can play a key role in making the wider public more aware of the cyberspace with its many facets.

The goal of this thesis is to take a dive into the cyber security of museums. This dive will be led by the discussion on the cyber risk analysis I conducted in collaboration with the MUSE – Museo delle Scienze of Trento, where we analysed the state of things and evaluated how much cyber risk the museum is subject to, and it will end in a discussion on a possible approach that could help museums be more cyber secure.

The Background section will present the background from which the internship and the thesis were born from. The Methodology section will present the “how we did it”: the methods we followed, the choices we made. The Result section will present the results we got from our work. The Evaluation section will put into perspective our work and, abstracting toward a more generalized approach, will present an approach for building cyber security inside a museal entity.

For the sections relating to the work conducted during the internship, I will prefer the “we” pronoun as the work had been done together by me and the MUSE staff and their inputs had been key to the results. For the other sections, I will prefer the “I” pronoun as it is more about my own research and evaluation and less about external inputs.

2. Background

The work we will be going to present started in the Autumn/Winter 2021 when, amid the ongoing restrictions for the lingering Covid-19 crisis, I started searching for an internship as part of my Master Degree in Cybersecurity. While I was open to many aspects of cyber security, I was particularly interested in the cyber risk: how it is evaluated and how it is managed in real cases.

I have been an eager museum goer for a long time and I had wondered about how cybersecurity was done inside a museum; for the way I saw it, it was something different from “the usual business”: you have artefacts, most often than not priceless, you have a flow of guests coming and going on which you have limited control on. And then, how is the IT infrastructure into such particular environment? Is the old paradigm of “everything locally and self hosted” still present or the new, overwhelming paradigm of “Something as a Service” has grown? Do they do things “in-house” or rely entirely on outside businesses for their IT infrastructure? Do internal security programs even exist? Last but not least, I wondered about their cyber risk. Are they aware? Do they conduct evaluation of it? These are some of the questions I was asking myself as, through common acquaintances, I was put in contact with the head of IT at MUSE - Museo delle Scienze, the science museum of Trento.

The MUSE – Museo delle Scienze is an auxiliary body of the Autonomous Province of Trento which has as task *“to interpret nature, starting from the mountains, using the eyes, tools, and applications of scientific research, taking advantage of the challenges of the contemporary world, stimulating scientific curiosity and the pleasure of knowledge, giving value to science, innovation, and sustainability.”* [MUSE 2021].

The main part of the museum is located in the building designed by the internationally famous architect Renzo Piano [MUSE 2021] and there are two remote locations at Ledro and Predazzo. While being primarily a museum, the MUSE is an important research center for the Trentino region too: they investigate many aspects of the world, such as biology and geology. As part of their research and museal status, they participate in both local and European scientific projects and they are part of the Italian network of museums and research centers.

As I said in the introduction, museums are at risk of cyber vulnerabilities and cyber attacks like any other entity connected to the cyberspace. However, I could not find much academic material on the cyber security and cyber risk of museums which are entities with specific needs and constraints, different from the ones of business. This lack of materials prompted the idea that lies at the core of the internship’s work and of this thesis: conduct an evaluation of the cyber risk a real museum is subject to. Then, from the empirical results, try to abstract and generalize how cyber security should be done in such a particular setting.

The MUSE had always tried to keep a proactive, research and experimentation focused approach to many aspects of being a museum and the head of IT was more than happy to work with me on this project. Cybersecurity was something they were doing already but with a pretty hands-on approach “We need this and this to be secure, so we do this and this.”; for example, they never did a cyber risk evaluation before, but they were already following a set of recommended security postures as per the Italian cyber security legislation for public administrations [AGID 2017].

My interactions were mainly with the head of IT services and cyber security of the MUSE, dr. Vittorio Cozzio, and their fellow staff member from IT services, dr. Luca Zucchelli. Our initial objectives were the evaluation of cyber risk, as said before, and the draft of a set of recommendations to improve the cybersecurity posture of the museum. To these, we added the execution of a study on how the employees perceived and behaved in regard to cybersecurity. This last objective had been an idea the IT staff of MUSE had in mind for some time but, due to many factors, they had not been able to conduct.

I gladly accepted to work on this additional objective for two reasons. Firstly, the way employees behave can change how the cyber risk of the museum is evaluated: for example, a great percentage of employees ignoring the basic of cyber security means that the risk deriving from them is going to be higher than the same employees set but with the majority doing cyber security properly. Secondly, employees are the first line of defence of a business as well as of a museum [Callahan 2021]. Breaches originating from employees’ mistakes had made the headlines before, even in public administrations [Colin 2020].

You can secure the IT system with the best technology available but if employees still do certain mistakes, like guessable passwords, the best technology might not be enough to protect the system from intrusions and bad actors. Understanding how your first line of defence is doing becomes paramount to secure the whole organization.

The whole work was gathered into three main documents, each covering one part of the work, which were very well received by the MUSE and were added to their documentation pool; specifically, the report on the questionnaire and phishing test, given its educational purpose, was added to the training resources for the employees.

3. Methodology

In this section we will go through the “how we did it”. We will present the tools and methodologies we employed to conduct our work, what choices with did and which contexts prompted them.

3.1 Cyber risk evaluation

Cyber risk evaluation and assessment is a process where the risk that results from the use of information systems is identified, estimated and prioritized. In other words, during a cyber risk evaluation, the threats to the subject entity are defined, the possible damages they can do is estimated and how much the subject entity should worry about them is evaluated.

Such process does not tell the subject entity they will never be breached or that they will: such question is practically impossible to answer. However, it can tell the state of things from the point of view of the risk coming from information systems an entity is subject to. Such point of view helps in building awareness on the state of things and in prioritizing the operational choices in regard to IT and cybersecurity every entity faces.

It is not a one-shot operation, quite the opposite: it should be a continuous and long-term operation building on previous results that gets done again and again as time passes. The cyberspace is an ever changing landscape and something relevant now might not be relevant in five, ten or fifteen years. A glaring example of this changing landscape are the social media, such as Facebook and Twitter, or e-commerce giants like Amazon: fifteen years ago, they were nowhere near the size and importance they are now.

The evaluation of cyber risk can be done with a wide range of techniques and methodologies, each and every one with their pros and cons, with different ways of looking at the same concept: the cyber risk. With all these methodologies, I wanted to employ a methodology that could result in a relatively easy understanding from non-technical people while still having a structured and clear approach that could cover the intrinsic technicalities of cyber risk and be accountable.

The choice was the FAIR methodology, acronym for Factor Analysis of Information Risk [FAIR Institute 2021], a proprietary methodology available for non-commercial purposes under a Creative Common License. This methodology of risk analysis, originally created in early 2000s, takes a structured approach to the risk itself, defining it as the composition of various, well defined factors. It results in a taxonomy like the one shown below, in figure 1.




Figure 1: The FAIR methodology taxonomy [Freund 2015]

The original FAIR methodology has a quantitative approach: to each and every element in the taxonomy is associated a number derived from probabilities and monetary values. However, we decided to take a slightly different approach and modify the methodology: we kept the taxonomy but we choose a qualitative approach where, instead of numbers and probabilities, we employed scales, such as “Very Low-Low-Moderate-High-Very High”.

This choice was dictated by a few factors. First, the result in the form of “Annualized Loss Exposure”, shown as a graph and number, of the quantitative approach resulted of difficult comprehension; second, it would have required a greater quantity of data compared to the ones readily available and they would have often been of harder retrieval, give the organizational context. Last but not least, the FAIR methodology is at its core a proprietary standard and when using numbers and probabilities, it would have required proprietary software which usage we choose to avoid.

Changing from quantitative to qualitative let us keep the structured, easily explainable, taxonomy while simplifying some of the work and having results of easier understanding for non-technical people, which made the most of who would read our reports at MUSE.

Instead of referencing the various definitions that will be needed to fully understand our cyber risk evaluation for the MUSE in the Result section, I will present these definitions in the following subsections, detailing them and explaining where we took specific approaches.

3.1.1 Terminology

A few definitions of terms used by the FAIR methodology that will be recurring in the rest of the thesis will follow.

With the term “*Asset*” we consider all “*the things we value*”: the things an organization value. They can be physical assets, like computers and buildings, as well as incorporeal assets, like data and reputation.

With the term “*Threat*” we consider an entity that threatens an organization and its assets. It can be a person as much as a group, but it has to be “people”: a threat is a terrorist or a terrorist organization, but things like social engineering or the cloud are not threats, the first is a technique which a threat employs to further its goals, the second is just a technology that is neither good or bad.

A threat is defined by a profile detailing key aspects of said threat:

- Motivation: what is the motivation behind the threat’s actions.
- Intent: what is the intent for attacking the organization.
- Sponsor: which third party supports the threat and their actions.
- Targets of choice: which are the targets favoured by the threat.
- Personal risk tolerance: how much personal risk the threat can tolerate.
- Concern for collateral damage: how much concern for collateral damage originating from their actions the threat can tolerate.
- Threatened assets: which assets of the organization are threatened by the threat.
- Damage: what kind of damage can the threat do to the organization’s assets.
- Contact frequency: how often the threat comes into contact with the organization and its assets.
- Threat Capability: how skilled is the threat and the amount of damage it can do.

Threats with similar profiles can be grouped into Threat Communities as an excessive granularity can be counterproductive for the analysis.

With the term “*Event*” we consider all the events that can generate a loss, be it material or immaterial, by the organization. A data breach is a loss event, bad publicity resulting in a diminishing of reputation is a loss event.

With the term “*Loss*” we consider a loss an organization had, it can be anything that the company had and has not anymore; some examples are a monetary loss where money had been lost, a reputation loss where the reputation had been diminished, a data loss where data had been stolen, a technological loss where some kind of technology is not available anymore to the organization.

A loss event can be a “*Primary Loss Event*” where the loss directly impacts the organization or it can be a “*Secondary Loss Event*” where the loss that impacts the organization is generated by the reaction that third parties have following a Primary Loss Event. An example of primary loss is a data breach where sensitive data had been exfiltrated, a secondary loss event generated by our example’s primary loss is the sanction that comes

from a ruling in accordance with the European General Data Protection Regulation (GDPR) law which fines businesses that does not protect data properly.

Threat Event Frequency

“The probable frequency, within a given time-frame, that threat agents will act in a manner that may result in loss.” [Freund 2015]

In other words, how often a threat can act against one of the assets of the organization under evaluation. It can be made by composing how often a contact happens between the organization and the threat, defined by the Contact Frequency, and how probable is that the threat acts against the organization, defined by the Probability of Action. However, in this analysis we decided to define it right away, without employing its decomposition.

Threat Capability

“The capability of a threat agent.” [Freund 2015]

This element defines how good is a threat, evaluating their skills and knowledge. The value is defined by the following scale:

Threat capability	Relative ability
Very high	The top 2% of the entire threats' group
High	The top 16% of the entire threats group
Moderate	Average knowledge and skills (between the bottom 16% and the upper 16%)
Low	The bottom 16% of the entire threats group
Very low	The bottom 2% of the entire threats group

Table 1: The threat capability and the threat's relative ability

Difficulty

“The level of difficulty that a threat agent must overcome.” [Freund 2015]

The difficulty is evaluated depending on the Controls against which a threat directly collides; in other words, a control is that structure, either software or hardware, that puts a check or limitation on conducting an action or accessing a resource. An example of control is asking for a specific password when accessing a specific resource.

Controls that make a threat's life harder make the difficulty's level grow while the controls that lower the contact frequency, discourage the action in the first place or limit the magnitude of the loss do not impact the difficulty value.

The controls we considered as relevant to the Difficulty value had been defined and grouped following the CIS Controls document, from the Center for Internet Security, version 7, April 2020 [CIS Center for Internet Security 2019]. Specifically, we considered the following ones:

- CIS 1: Inventory and Control of Hardware Assets (ABSC 1)
- CIS 2: Inventory and Control of Software Assets (ABSC 2)
- CIS 3: Continuous Vulnerability Management (ABSC 4)
- CIS 4: Controlled Use of Administrative Privileges (ABSC 5)
- CIS 5: Secure Configuration for Hardware and Software on Mobile Devices, Laptops, Workstations and Servers (ABSC 3)
- CIS 7: Email and Web Browser Protections
- CIS 8: Malware Defences (ABSC 8)
- CIS 9: Limitation and Control of Network Ports, Protocols and Services
- CIS 11: Secure Configuration for Network Devices, such as Firewalls, Routers and Switches (ABSC 3)
- CIS 14: Controlled Access Based on the Need to Know
- CIS 15: Wireless Access Control

Other controls in the CIS list had not been considered relevant to the Difficulty. The strings between parentheses reference the AGID – Agenzia per l’Italia digitale directive “Misure minime di sicurezza ICT per le pubbliche amministrazioni” [AGID 2017] to which the MUSE, as part of the Italian public administration, has to adhere in regard to cyber security. The AGID directive refers to an older version of the CIS controls documentation but we choose to refer to the newest version available at the time of the internship which was the April 2020 version: between the two there had not been extreme changes and we believed it was better to refer to up to date documentation.

The difficulty each control brings is evaluated in relation to the scale measuring the threat’s capability; the Threat’s Capability scale is shown in table 1. Thus, a control able to impact the top 2% of threats is considered a control giving a “Very hard” difficulty.

Ability	Control’s difficulty
Very high	Very high difficulty
High	High difficulty
Moderate	Moderate difficulty
Low	Low difficulty
Very low	Very low difficulty

Table 2: A threat’s ability and the control’s difficulty

Vulnerability

“The probability that a threat agent’s actions will result in loss.” [Freund 2015]

This value is defined by the composition of the how much the threat is skilled through the Threat’s Capability value and how good the controls are with their Difficulty.

Frequency of a Loss Event

“The probable frequency, within a given time-frame, that loss will materialize from a threat agent’s action.” [Freund 2015]

This value is defined as the product of how much our asset is vulnerable, defined by the Vulnerability, and how often the threat is going to act against the asset, defined by the Threat Event Frequency.

Primary loss

“Primary stakeholder loss that materializes directly as a result of the event.” [Freund 2015]

The first typology of primary loss we looked at is about the data loss; we defined it through the well-known CIA triad: Confidentiality, Integrity, Availability. This triad can work well for other assets too.

To the data-related losses, we add other primary losses related to the IT infrastructure, the employees and the organization’s operations. These primary losses are:

- Substitution of parts of the IT infrastructure. For example, substitution of computer or other tools compromised by a malware attack.
- Monetary expenses incurred for the management of an event and the following restoration of normal operations. Examples of these expenses are:
 - The salaries paid to employees unable to work during the event
 - The additional wages for who operates during the event and the following restoration process
 - The loss of revenue caused by the event itself
- The time, measured in hours of work, spent during the event and the following restoration of normal operations.
- The possible damage to reputation and image generated by the event.

Secondary loss event frequency

“The percentage of primary events that have secondary effects.” [Freund 2015]

This value defines how often a primary loss event generates a secondary loss event.

Secondary loss magnitude

“Loss associated with secondary stakeholder reactions.” [Freund 2015]

When a secondary loss happens, this value defines how big is the secondary loss happening.

Secondary Risk

“Primary stakeholder loss-exposure that exists due to the potential for secondary stakeholder reactions to the primary event.” [Freund 2015]

The Secondary Risk is the risk that is generated, by composition of the Secondary loss event frequency and the Secondary loss magnitude, following an event generating a primary loss.

A list of secondary losses relevant to our organization are:

- Expenses generated from sanctions, for example GDPR induced sanctions.
- Expenses generated from legal procedures, for example by a trial for bad conduct or one against a caught malicious actor.
- Expenses generated from the public relationship efforts.

These kinds of expenses can be important but they vary on a case by case basis.

Loss magnitude

“The probable magnitude of primary and secondary loss resulting from an event.” [Freund 2015]

How big is the loss the organization faces when we cumulate the Primary Loss and the Secondary Risk generated by it.

3.1.2 Risk

“The probable frequency and probable magnitude of future loss.” [Freund 2015]

In other words, how big and often a loss can be and happen. It is a composition of how often a loss event can happen, defined by the Frequency of a Loss Event, and how big such loss can be, defined by the Loss Magnitude. As said before, it is a probability, it does not say specifically a loss will happen but it only quantifies how often and how big it might happen.

3.2 The process

The application of the FAIR methodology to a real case is made of a series of consecutive steps with the intent of building an incremental knowledge base leading to final definition and evaluation of risk itself.

The first step we took was the definition of the Assets, the “things we value”, the things the organization values; by taking a look at the assets of the museum, we selected the ones that are relevant to our focus on the cyber risk.

The second step was the definition of a set of threats: what poses a threat to the organization under exam and to its assets. In our cases, we started by doing some research followed by an evaluation of what threats applied to our organization.

The third step took us into scenarios' building. Scenarios are moments in space and time which let us connect threats to assets and the effects assets suffers at the hand of a threat. Scenarios help us scoping our analysis and we built them by connecting assets and threats to answers few key questions targeting our analysis.

At the fourth and final step, we brought things together by giving a value to the elements in the taxonomy, starting from the bottom and moving up, till we reached the root of the tree and we defined the risk itself.

In our specific case, we were working during the Covid-19 pandemic and thus, we relied on remote work tools to carry on our analysis. Each step was written in a self-standing document shared in between me and the MUSE's staff and reviewed through real time calls. At the end, we gathered all these documents into a single and final document that carried the whole process start to end.




Figure 2: Evaluation's process diagram

3.3 Study on employees' behaviours in regard to cybersecurity

Our study on how employees behave in regard to cybersecurity was structured in two parts. The first part was a questionnaire with the purpose of evaluating the “conscious” behaviours and thinking of employees. The second part was an attempt of phishing we conducted internally against the employees; much like penetration testing tries to uncover vulnerabilities in a system by trying to penetrate it, we tried to phish the employees to evaluate their “unconscious” behaviours in the most realistic way possible.

3.3.1 The questionnaire

The questionnaire composed itself of fifteen (15) questions, either multi or single answer, which tried to cover various aspects of cyber security such as phishing and bad attachments, malware attacks, web security, password security and workstation security. Our objective was to get a general coverage on various topic relevant to the organization and the employees while not diving into too technical aspects which would have resulted difficult for non-technical employees.

Security policies on the usage of the IT infrastructure are at the base of a security posture and are, by law, required to be present. How much our employees know about the regulations in place? Then, hackers had changed from the 80s and their skills had grown as much as the internet had; how is the perception of hackers' abilities of our employees? Would they be able to copy a website and make of it a fake or they still are perceived as simple kids messing up with the internet?

Are you aware of the company's security policies on the use of IT equipment?	• Completely. • Enough. • Sufficiently. • None.
How much do you think possible for a hacker to duplicate a business' website in such a way that, to a not accurate control, they look the same?	• A lot. • Enough. • None. • I don't know.

Phishing is the most widely known and visible form of attack and it is still considered relevant and dangerous. Do our employees know about it? Recognize it? Avoid it?

In your opinion, phishing is:	• Another term to say we are going fishing. • The attempt by third parties to obtain information through fake emails and fake websites. • The gathering of sensitive data in a vulnerable computer. • The use by third parties of other people's passwords.
Which are the characteristics to recognize if an email is a fake?	• The text is without grammar. • The sender's address is unreliable or incorrect. • Both the subject and the content push for the submission of data from the reader. • All the characteristics above. • None of the characteristics above. • I don't know.
In your opinion, who is more at risk of being targeted by phishing emails?	• The member of the administration area. • The member of the management area. • Every member of the organization. • The member of the human resources area. • The member of the research area. • None of the above.
Files on a computer have different extensions. Which one of the following could hide a malicious element?	• .zip • .exe • .doc • .bat • All of the above. • Neither of the above.
You receive an unexpected mail saying it is coming from the administration. They attach a pdf file which has a .exe extension (es.: invoice.pdf.exe). What do you do? (multiple choice)	• I open the file. • I ask information to the sender before opening the file. • I inform who is in charge of the cyber security and wait for instruction on how to proceed. • I delete the mail.

You receive a mail from an address that seems to be the assistance of Facebook; this email notify you that your account had been compromised and asks you to click on a link to fix it; what do you do? (multiple choice)	• I click the link. • I search for information on the Facebook website or through Google. • I delete the mail. • I log in into Facebook and change the password.
You receive a mail from an unknown sender; this email notifies you that you just won a new 6 th generation iPad and you just need to send them your data. What you do?	• I reply with all the necessary data. • I reply to the mail asking for additional information. • I delete the mail. • I advertise the link of the win on the socials.

Malwares had been on the constant raise for the past few years becoming one of the most dangerous types of attacks. Knowing what to do when seeing a malware in action can be key to how fast an organization can react and limit the damage spread. How do our employees behave in such case?

Your computer starts behaving in an anomalous and unexpected way, for example the file change names or disappears. What do you do?	• Nothing, it's part of the normal operations, I do not worry. • I call the IT assistance and hope they know what is happening. • I power off the computer and call for the IT assistance. • I immediately pull the power cable off and call for the IT assistance.
Your computer suddenly freezes: you can't access its resources but you only see a screen asking for a ransom to be paid with bitcoins. In your opinion, what is happening?	• A simple malfunctioning of the computer. • A ransomware attack. • An overload of the mail. • The computer reached the end of its operative life.

Fake websites and unsecure websites had been a vehicle for malicious actors for a long time. How do our employees check if a website is secure?

What is the difference between http:// and https://?	• No difference. • The website had a better visual quality. • The website supports the data encryption to and from it. • The website shows the most recent version available. • I don't know.
You're searching for a website selling domestic appliances and you find one with very good prices. What do you look at to find if the website is safe?	• I check for the presence in the address bar of the lock icon. • I check for the absence of any icon in the address bar. • I search the web. • None of the above.

Passwords are a key element of any organization's security. How our employees create their passwords?

It is notoriously hard to memorize passwords. What you do in this regard?	• I create simple passwords so I can remember them by memory. • I leave a post-it under the keyboard. • I gather them in a excel file hidden in the computer's folders. • I put them in a password manager.
---	--

An unsecured workstation can be a door into the IT infrastructure of the organization. How do our employees behave in regard to workstation security when in a rush?

You receive an urgent call that makes you leave your working station. What do you do?	• I leave the working station as it is and I leave. • I lock the pc and cover any document before leaving. • I power off the monitor. • I ask the colleague not to let anyone approach my working station.
---	---

3.3.2 The phishing test

The second part built on a similar test the IT company GitLab run internally [Manzuik 2020] and it was composed by a fake mail that asked the target for their work credentials to avoid the account deactivation; this mail was sent from an internal account to simulate its compromise by a malicious actor.

As we prepared this part, the first step had been to decide which subject our fake email would have. Were we going to declaim the win of some kind of object or service? Were we going to lure our subjects with some benefits in exchange of their credentials? Were we going to fake an account deactivation? Were we going to fake a service's credentials reset?

In the end, we choose to fake the deactivation of the work account as we considered it as the option that would get the most attention from the employees. On the opposite, we evaluated that internal emails offering third party benefits as a member of the organization are often ignored and the emails declaiming the win of some object had been regarded unanimously as phishing when presented in the questionnaire.

The mail itself was built to vaguely resemble a true internal mail, with font, headers and logos of the MUSE but we made sure to include some elements that, to a careful observer, could represent a warning sign that the mail was not what it looked like it was. Specifically, these elements were the addresses, some were fake built in the shape of older, no more used addresses, the tone of the mail, it was not the tone the sender usually used in similar communications, and the procedure it was calling for which is an anomalous one.

The mail directed the user to a fake website we built on the internal network by stripping and copying a real login form used by the IT services. If the credentials were inserted, a simple success page resembling real ones would be returned; no sensitive data was gathered but basic access data used for statistical purposes.

By hosting our website internally, our web certificates were in order, thus simulating a certificate spoofing attempt while the internal origin of the mail, as said before, simulated the compromise of an employee's business account.

The phishing e-mail, as well as the fake login page and success page can be seen in figure 3, 4 and 5.

Sblocco account Microsoft 365



Vittorio Cozzio

mer 21/04/2021 18:03

A: mario.rossi@mtsn.tn.it; scuole@mtsn.tn.it; luca.zucchelli@mtsn.tn.it



Gentile utente,
nella verifica mensile della sicurezza delle policy del servizio Microsoft 365
si è evidenziato il pericolo di attacco al vostro account.
Per evitare che qualcuno si introduca nella vostra mail o nei vostri documenti
il sistema ha avviato una procedura di blocco del vostro account.

Se vuoi continuare ad usare il tuo account devi confermare la tua identità al link qui sotto

[CONFERMA IDENTITA'](#)

ATTENZIONE!!!!

Se entro 48 ore non sarà fatta la tua identificazione, il sistema eliminerà il tuo account.
I tuoi documenti e le tue mail saranno eliminate definitivamente
e non saremo responsabili di questa perdita.

Questo messaggio di posta elettronica e i file trasmessi con esso sono confidenziali e destinati esclusivamente
all'uso dei singoli utenti o enti ai quali sono indirizzati.

Nel caso in cui il messaggio sia stato ricevuto per errore, si prega di informare il gestore di sistema.

Figure 3: The fake email




Figure 4: The fake login page




Figure 5: The fake success page

4. The results

The work with the MUSE – Museo delle Scienze resulted in three documents which covered various aspects of the work we did; in this section we will concentrate on the document about the risk evaluation, titled “Valutazione del Rischio con Metodologia FAIR dell’organizzazione MUSE”, and the one reporting in regard to the questionnaire and the phishing test, titled “Report questionario e test”.

4.1 The risk evaluation - Set up

As mentioned in the Methodology section, we choose to use the FAIR methodology with qualitative measurements for the risk evaluation of the MUSE – Museo delle Scienze.

The first step we took was to define the assets, the “things the organization” valued; we divided them into two broad categories: one for physical assets, assets that can be physically touched, and one for incorporeal assets, assets that cannot be physically touched but which can be accessed by IT means. With the scope of our analysis in mind, specifically the cyber security of the museum, we choose to exclude all the assets that did not relate directly to the IT systems of the museum, such as the collections themselves, the buildings or the tools not related to the IT systems as we considered them to be out of our scope.

In the physical assets category, we included the IT infrastructure as a whole: the machines, such as personal computers, laptops and servers, the network devices, such as routers and switches, and other physical devices. All these assets are assets the museum owns directly and on which has a direct control.

In the incorporeal assets we included all the other IT assets that are not physically touchable: all the data managed by the museum, such as employees’ and guests’ data, the web applications and the reputation/image of the museum itself. Data had been the first asset to be included in this category as they represent the most important and yet incorporeal asset any organization owns and make use during their operations. We decided to include the web applications in this category as they are something that an employee cannot touch without the aid of a device like a computer or a smartphone; these web applications can be either self-hosted and thus under the direct control of the organization or hosted by a third parties and provided as part of a Software-As-A-Service agreement. The reputation/image was included as it is incorporeal but very important as it represents what the customers think of the organization. A bad image can weigh heavily on revenues while a good image can help smoothen the public perception of incidents like a data breach.

Physical assets	Incorporeal assets
• IT infrastructure ◦ Servers ◦ Computers and laptops ◦ Network devices • Other hardware	• Data ◦ User’s data ◦ Management data ◦ Research and scientific data • Web applications • Reputation/image

Table 3: The list of relevant assets

Once we decided what were the things the museum values, we could think about who could be a threat to the museum and its assets: the threats. To define a set of relevant threats, we first looked at the cyber security history of the museum to see what kind of threats had come in contact with it.

Then, we searched through the web for well known threats, such as the Cyber Criminal threat, to compile them into a list. However, not all the threats we found were relevant and thus we asked ourself two key questions: what would this threat want from us? What is their motivation for attacking us? With these questions, we could narrow down the threats that were relevant to our organization; for example, threats like a Nation State sponsored group or a terrorist group with a religious background were delisted because they did not represent a relevant threat to the organization. It is true that a terrorist attack could happen to the museum the same way a Nation State sponsored attack could be conducted against it but their probabilities are considered to be so small to make them not relevant for our analysis.

This incremental process led to the set of threats show in the following table:

Internal Threats	External Threats
• Privileged Insiders • Not privileged Insiders	• Cyber Criminals • Amateur Hackers • Activists/hacktivists

Table 4: The list of threats relevant to the organization

We made the choice of grouping up single threats into communities, as suggested by the FAIR methodology itself [Freund 2015]: a too much fine grained selection of threats would have given us a lot of threats with many similarities but alone by themselves not too relevant. A clear example of this reasoning is the Cyber Criminals threat community which covers a wide range of possible actors but which all shares similar ways of operating, similar objectives and motivations and similar tools of choice.

The grouping under internal and external threats was decided as we believe that the point of origin can change a lot the way we approach and manage the two groups of threats; managing an internal threat which we know having already access to our infrastructure is different from managing a threat that has to go through our external defences before being able to operate inside our infrastructure.

We will continue by presenting the profile of each threat we mentioned in Table 4.

Privileged Insider

This category gathers all the employees and internal collaborator whom, during their work, have a “privileged” access to the organization’s assets, such as the IT infrastructure and the data contained inside.

The logic of choice is pretty straightforward: the employees are the first threat to any organization. Listing employees as a threat does not mean that one of them will conduct an attack with malicious intent from the inside with a 100% of accuracy; in fact, we believed to be more probable that the harm to the organization would originate from human errors than for a malicious intent. However, knowing such a threat exists is the first step to mitigate it.

Motivation	Ignorance, vengeance or personal gain
Intent	Ignorance, take a vengeance following an injustice perceived as received or due to external factors such as a lack of money
Sponsorship	None, in rare case a collusion with external agents is possible
Targets of choice	Parts of the organization of which they are part of or which alternatively offer a satisfaction to the vengeance or a monetary return
Skills	Variables, depending on the position inside the organization
Personal risk tolerance	From low to high, depending on the motivation
Concern for collateral damage	Variable, high if part of a cohesive group
Threatened assets	• IT Infrastructure • Data • Web application • Reputation and image
Damage	• Confidentiality • Integrity
Contact frequency	Very low
Threat Capability	Variable depending on the role and skills

Table 5: Privileged Insider's profile table

The main objective of this threat is to harm the organization they are part of; having access to wide parts of the IT infrastructure and keeping in mind their intent, the damage focus on a loss of confidentiality, through the disclosure of information, and on a loss of integrity, through the modification from the inside of sensitive data.

This category has a very low Contact Frequency as it did not seem to be any preconditions for a greater internal risk.

The Threat Capability remains variable as it depends on the professional figure: for example, a botanist without specific IT knowledge can indeed damage equipment and research close to him or her but cannot, in principle, have the ability to do a high degree of damage toward the whole IT infrastructure. An internal employee, on the other hand, part, for example, of the administration has a greater ability to do damage given the privileged access to various types of data and tools. Extreme cases such as a botanist who sets the entire building on fire are not cases considered in the analysis.

Non-Privileged Insider

This category gathers all the external collaborator and partners with which the organization operates but which are not granted particular privileges toward the organization's assets, such as the IT infrastructure and the data in it.

The reasoning behind this threat is once again straightforward: every organization has a set of people at various levels and skills that are not full fledged employees but still work with the organization and its assets; as such, they can represent a threat.

Motivation	Vengeance or personal gain
Intent	Take a vengeance following an injustice perceived as received or due to external factors such as a lack of money
Sponsorship	None, in rare case a collusion with external agents is possible
Targets of choice	Parts of the organization of which they are part of or which alternatively offer a satisfaction to the vengeance or a monetary return
Skills	The skills varies but they are limited by the absence of privileges
Personal risk tolerance	Variable depending on the person's situation
Concern for collateral damage	Variable depending on the person's situation
Threatened assets	- IT infrastructure - Reputation and image - Data, only with high skills
Damage	- Confidentiality - Integrity
Contact frequency	Very low
Threat Capability	Variable depending on the person

Table 6: Non-Privileged Insider profile table

The motivations of a Non-Privileged Insider are similar to the ones of a Privileged Insider and they share the threatened assets and the possible damages; however, the Threat Capabilities are heavily limited by the absence of privileges and access to the IT infrastructure of the target organization.

This category has a very low Contact Frequency since the probability that an external collaborator or a partner try to operate against the organization's infrastructure should be pretty limited.

Cyber Criminals

This Threat Community gathers a set of people that conducts cybercrimes, not necessarily connected to the organized crime, but which nonetheless operate illegally and without an honest purpose.

This Threat Community could not be avoided: they are the number one perpetrator of cybercrimes together with Nation State sponsored actors. It is not a matter of if but a matter of when such Threat Community will get in contact with the organization.

Motivation	Financial
Intent	Monetary return
Sponsorship	Autonomous or sponsored by criminal organizations or other criminal agents
Targets of choice	Any service or organization through which they can get a monetary return, often following the rule of the maximum return with the minimum effort
Skills	In majority high: professional hackers with the abilities, the knowledge and the support to conduct cyber crimes
Personal risk tolerance	High; when the economic return is no longer sufficient or the risk of discovery is excessive, it is possible they abandon the attack
Concern for collateral damage	Low toward third parties as long as the collateral damages do not damage their own operations

Threatened assets	• Data ◦ User's data ◦ Management data • Web Applications • IT infrastructure • Reputation and image
Damage	• Confidentiality • Integrity • Availability
Contact frequency	Very high
Threat Capability	Very high

Table 7: Cyber Criminals profile table

Cyber criminals are a threat that threatens every asset of the organization; their primary objective is an economical return and this leads them to search for valuable data to exfiltrate or damage, and/or to limit the availability of the Web Applications and IT infrastructure in exchange of remuneration, such as with ransomware. Their damage covers the entirety of the CIA triad.

In this analysis, they are considered as the most probable and dangerous threat with a Contact Frequency and a Threat Capability rated Very High.

Amateur Hackers

This Threat Community gathers the people that test the IT infrastructure of an organization without a specific objective and without a particularly malicious intent; however, they might not follow the right procedures for reporting and disclosure of vulnerabilities. For example, instead of reporting the presence of a vulnerability to the organization concerned, it is possible for them to prefer getting in contact with news outlets.

The reasoning behind this Threat Community might not be as obvious as for the previous ones. Amateur Hackers are those people that hacks "for fun" and the reader will wonder why such "hacking for fun" matters

to an organization like the MUSE. First, there had been a data breach in the past generated by a hacker that can be considered part of this threat community; second, the MUSE is a very visible organization and, as such, it might attract such individuals, in practice by representing a “high profile” target.

Motivation	Personal interest, recognition of their work
Intent	Curiosity
Sponsorship	None
Targets of choice	Any service or organization that intrigues them or that can represent an image or economic return
Skills	Medium, they are not professional hackers such as the ones employed by criminal organizations but still have enough knowledge to be moderately dangerous
Personal risk tolerance	Low, they do not do it as part of their work but for curiosity, one might think that there is a limit to what can be done out of curiosity
Concern for collateral damage	High
Threatened assets	• Data • IT infrastructure • Web Applications • Reputation and image
Damage	Confidentiality
Contact frequency	Moderate
Threat Capability	Moderate

Table 8: Amateur Hackers profile table

The Amateur Hacker defined here does not have the same knowledge and skills of a cybercriminal, despite their interest covers the whole IT infrastructure. This reasoning leads their Threat Capability and Contact Frequency to Moderate since the organization under exam is very visible and active in the community it is located in.

Hacktivists

This Threat Community gathers the people who identify themselves as activists and/or are hackers engaged in actions of protest or civil disobedience via the Internet. For example, environmental and animal rights groups are considered part of this category.

The reasoning for this Threat Community starts once again in the history of the museum. A few years back, the MUSE had a few problems with some activists that physically protested against the museum’s involvement in a project in which there was as a partner a well know oil and gas company. The MUSE, as for the Amateur Hackers, can represent a “high profile” threat to activists. The reader will be wondering where is the connection between activists and hacktivists: hacktivists are activists that makes a greater use of the IT technologies to further their goals. While as of the time of writing, they are still a very limited phenomena, we decide to include them as they might take on more prominent role in the years to come, in a similar way to criminals that moved from physically entering industrial buildings to steal products and information to attacking the industrial networks remotely.

Motivation	Ideology
Intent	Demonstrate the ideological inaccuracy and misconduct of their targets
Sponsorship	None
Targets of choice	Any service or organization perceived as in error or dangerous to their ideology

Skills	Variable, between low and moderate; these groups do not yet appear to have developed threatening capabilities similar to those of other actors such as professional hackers
Personal risk tolerance	Variable, depending on the person and group
Concern for collateral damage	Moderate
Threatened assets	- IT infrastructure - Data - Management's data - Research data - Reputation and image
Damage	- Confidentiality - Availability
Contact frequency	Low
Threat Capability	Very Low/Low

Table 9: Hacktivists profile table

The main objective of activists and of their “computerized” versions, hacktivists, is to demonstrate the correctness of their ideology and the errors in which their adversaries incur: this leads them to protest actions in order to prevent the normal operations of their targets and to demonstrate the incorrectness of those operations. From this point of view, the assets at risk are the IT infrastructure and the Data themselves. Their harms focus on confidentiality and availability of the assets.

Given the analysis in which this threat definition had been done, the frequency of this threat is evaluated at Low as their Threat Capabilities. This threat had been included as the organization can be a target of opportunity for them.

Just by reading at the profiles it is hard to understand where they place between each other. Who is the one more dangerous or the one happening with more frequency. To our aid, I thought about organizing these Threat Communities into a table with two axes: on the x axis the Threat's Capacity, in other words how dangerous a Threat Community can be, on the y axis, the Threat Frequency, in other words how often we expect the Threat Community to act against an asset. The result can be seen in Table 10.

Threat Frequency	Very High					Cyber criminals
	High					
	Moderate			Amateur Hackers		
	Low	Hacktivists	Hacktivists			
	Very Low	Hacktivists, Not Privileged Insiders	Hacktivists, Not Privileged Insiders, Privileged Insider	Not Privileged Insider, Privileged Insider	Privileged Insider	Privileged Insider
		Very low	Low	Moderate	High	Very high
		Threat's Capacity				

Table 10: Correlating Threat Frequency and Threat's Capacity

It was kind of expected how the cyber criminals ranked at the top of the list in both categories as much as the hacktivists ranked at the bottom of the list. Cyber criminals are second only to Nation State sponsored actors

in term of resources and abilities; however, we decided to exclude Nation State sponsored actors as we did not believe the MUSE fits into their target list. Nation State sponsored actors have a preference for specific field of research, public offices and critical infrastructures; however, the MUSE is not doing research in any of the research fields' that might of interest such actors, they are not a public office in itself nor they are a critical infrastructure.

Activists/Hacktivists had been included because the MUSE can represent a target of opportunity for such groups. It had happened before that the MUSE had to deal with some activists protesting against a big energy company with which the museum was having a partnership. Hacktivists are the "branch", if we can call it, of activism that focus on protest actions through the internet; as of now, they had not shown the technical abilities and reach of threats like cyber criminals but we thought them to be relevant nonetheless and thus to be included.

Amateur hackers are another category that we defined building on past experiences of the museum. While not as dangerous as cyber criminals, we consider them to still be relevant: not as skilled, they could attack the IT infrastructure for reasons that does not fully ascribe to the ones of other threats. For the recognition from their peers and others, for curiosity and for a test of their skills are the main reasons for which this threat operates; we considered them kind of a "neutral" threat as there are not intrinsically "evil" like cyber criminals but they are not all "white hackers" as they might not follow the proper procedures when it comes to vulnerabilities disclosure.

Insiders, either with or without privileges, could not be excluded by the list. Edward Snowden is a prime example of a privileged insider but not all insiders operate with an ill intent. In the insiders category we consider insiders that operates against the organization, yes, but we consider the insider that damage the organization by mistake too. Being insiders, however, they can have a widely varying set of skills and threat values: a simple researcher have limited access to valuable data while an administration's clerk works with sensitive and confidential data every day.

Threat Communities by themselves are not enough, we need more to fully bridge them to the assets, the type of damage they can generate and to the risk they truly pose to the organization. Scenarios come into play at this point of analysis: each threat community is linked to the assets it threatens and to the intent with which it operates.

Scenarios organized into tables are the step that will let us start give a value to elements of the FAIR taxonomy that will come into play later, for example the Frequency of a Loss Event or the Magnitude of the Loss. They are going to help us replying to a set of questions that defines the scope and are at the core of our analysis. These questions are:

- To which risk is the IT infrastructure subject to?
- To which risk are the organization's Web Applications subject to?
- To which risk are the Data managed by the organization subject to?

As we were doing a risk evaluation for the whole organization, we choose to keep these core questions at a higher level and focused on the three main asset typologies the organization owns. Nothing would have stopped us to go in depth and analyse the risk of a specific asset but, for our objectives, it was deemed not necessary.

The oncoming tables, from Table K to Table K, will illustrates the scenarios organized by Threat Community.

Privileged Insider

A Privileged Insider can threaten every asset of the organization; for this reason, we considered all three main assets: Data, Web Applications and IT infrastructure. The reputation, while being a primary asset itself, is threatened in a more "indirect" way through the actions against the other three primary assets.

A Privileged Insider can act with a malicious intent or it can threaten an asset unknowingly or due to an error.

Asset	Intent	Effect
Data	Malicious	• Confidentiality • Integrity • Monetary expenses
Data	Error	• Confidentiality • Integrity • Monetary expenses
Web Applications	Malicious	• Availability • Monetary expenses
Web Applications	Error	• Availability • Monetary expenses
IT Infrastructure	Malicious	• Availability • Monetary expenses • Substitution
IT Infrastructure	Error	• Availability • Monetary expenses • Substitution

Table 11: Scenario table for the Privileged Insider threat

Non-Privileged Insider

The Non-Privileged Insider can threaten any of the organization's assets but, due to the absence of privileges, the threat's capability is limited; the assets considered to be vulnerable are the IT Infrastructure and the Data assets. The reputation, while being a primary asset itself, is threatened in a more "indirect" way through the actions against the other two primary assets.

A Non-Privileged Insider can act with a malicious intent or it can threaten an asset unknowingly or due to error.

Asset	Intent	Effect
Data	Malicious	• Confidentiality • Integrity • Monetary expenses
Data	Error	• Confidentiality • Integrity • Monetary expenses
IT Infrastructure	Malicious	• Availability • Monetary expenses • Substitution
IT Infrastructure	Error	• Availability • Monetary expenses • Substitution

Table 12: Scenario table for the Non-Privileged Insider threat

Cyber Criminals

The Cyber Criminals can threaten every asset of the organization; for the data, there is a focus on Management Data and User's Data as they can have a greater value in comparison to Scientific Data. The reputation, while being a primary asset itself, is threatened in a more "indirect" way through the actions against the other primary assets.

The Cyber Criminals only act with malicious intent.

Asset	Intent	Effect
User's Data	Malicious	- Confidentiality - Integrity - Availability - Monetary expenses
Management Data	Malicious	- Confidentiality - Integrity - Availability - Monetary expenses
Web Applications	Malicious	- Confidentiality - Integrity - Availability - Monetary expenses
IT Infrastructure	Malicious	- Availability - Monetary expenses - Substitution

Table 13: Scenario table for the Cyber Criminals threat

Amateur Hackers

An Amateur Hacker mainly targets the IT Infrastructure and the Web Applications, the asset Data is considered a target of opportunity, with a focus on User's Data and Management Data. The reputation, while being a primary asset itself, is threatened in a more "indirect" way through the actions against the other primary assets.

An Amateur Hacker does not act right away with a malicious intent the same way a Cyber Criminal does but, at the same time, it cannot be said it act by error or unknowingly.

Asset	Intent	Effect
User's Data	Malicious / Neutral	- Confidentiality - Monetary expenses
Management Data	Malicious / Neutral	- Confidentiality - Monetary expenses
Web Applications	Malicious / Neutral	- Confidentiality - Monetary expenses
IT Infrastructure	Malicious / Neutral	- Confidentiality - Monetary expenses

Table 14: Scenario table for the Amateur Hacker threat

Hacktivists

Hacktivists threatens every primary asset, with a focus on Management Data and Research and Scientific Data. The reputation, while being a primary asset itself, is threatened in a more "indirect" way through the actions against the other primary assets.

Hacktivists only act with a malicious intent.

Asset	Intent	Effect
Research and Scientific Data	Malicious	• Confidentiality • Availability • Monetary expenses
Management Data	Malicious	• Confidentiality • Availability • Monetary expenses
Web Applications	Malicious	• Confidentiality • Availability • Monetary expenses
IT Infrastructure	Malicious	• Availability • Monetary expenses • Substitution

Table 15: Scenario table for the Hacktivists threat

4.2 The risk evaluation - Frequency of a loss event

The Frequency of a loss event is the root value of the left side of the FAIR topology as it is the left child of the root Risk, opposed to the Loss Magnitude which is the right child of Risk.

To define and evaluate this value, we build on the scenarios we prepared in the previous step and we will organize the tables by threat.

Privileged Insider

The Privileged Insider is a Threat Community representing the organization's personnel itself; when we did our evaluation, we kept in mind the actual state of the organizations and the results we got from our questionnaire and phishing test. Our small study empirically demonstrated how there are the conditions for practical vulnerabilities as much as there are dangerous behaviours that could be employed especially by the Cyber Criminals threat to attack the organization and its assets.

In general, we evaluated the frequency that a privileged insider acts with a malicious intent very low: while the perfect work environment hardly exists, the atmosphere and people did not seem to warrant a higher level of probability for such event.

The capability of the threat itself and the assets' vulnerability to it are rated "High": the privileged insider has, as its name implies, a privileged access to many, if not all, of the organization's assets and we assumed it knows about the protections in place to defend them if not completely at least partially.

In the specific case of the Web Application asset, this threat's capacity and the asset's vulnerability are rated "Moderate". This asset category is made for the greater part of the software offered with a Software-As-A-Service approach, implemented and managed by a third party and thus not under the direct control of the organization. Such situation led us to evaluate the capacity and vulnerability "Moderate" because a privileged insider does not have a complete access to these assets the same way it would with on-premise ones.

In opposition to the malicious intent, an insider generating a threat by error is considered to be much more probable, especially in regard to the Data asset and the IT infrastructure asset. We believe it is important to recognize that human errors happen and they cannot be completely avoided.

If the frequency is higher, the Threat's Capability and the asset's Vulnerability with this intent result lower, rated as "Moderate". This choice came as we balanced what a human error could generate opposed to the controls and security solutions the organization has. The controls should be able to prevent most errors against our assets but there is still a probability that they are not enough or are not able to prevent a loss event, for example an employee attaching the wrong file to an email is something we can hardly predict or stop with any of the controls the organization has implemented.

In conclusion, the Frequency of a loss event for the Privileged Insider varies between “Very High” and “Very Low” depending on the scenarios. We consider more probable for a loss event to happen in regard to the Data and IT Infrastructure assets in a scenario where we had a human error than for the same scenarios but with a malicious intent. The scenarios centred on the Web Applications asset, as we said before, are considered with either intent having a limited probability to generate a loss event.

The table gathering the Frequency of a loss event for the threat Privileged Insider can be seen at Table 16: The Frequency of a Loss event for the Privileged Insider threat.

Non-Privileged Insider

The Non-Privileged Insider Threat Community represents all the people, such as external collaborators and partners, that work with the organization and have a certain degree of access to the organization’s assets but are not full-fledged employees. In the specific case of the organization under exam, this Threat Community is made by the museum guides as well as other external contractors and collaborators.

The frequency of this threat’s action is generally rated “Very Low” for all scenarios but the ones centred around the IT Infrastructure; while there had been some problems and a strained relationship with certain groups making up this threat community, we still evaluated “Very Low” that they could operate against the organization’s IT assets. The scenarios centred on the IT Infrastructure have a slightly higher probability because it is the asset that comes the most in contact with this Threat Community and thus it might be subject to a loss event with a higher probability.

The Threat’s Capacity and the asset’s Vulnerability are limited in every scenario because we believe that the controls in place and how the Non-Privileged Insiders are managed, for examples by limiting the access to certain asset subsets, are enough to prevent most of the malicious or error actions; we are aware that a very motivated attacker could, even with a non-privileged access to the assets, be a threat but we evaluated such situation to have a minimal and not relevant probability to happen.

Following the above mentioned reasoning, the Frequency of a loss event by the action of this Threat Community is evaluated to be either “Low” or “Very Low.”: based on the organization experience, we do not believe such events to be very probable. Of course, being not very probable does not mean it will not happen.

The table gathering the Frequency of a loss event for the threat Non-Privileged Insider can be seen at Table 17: The Frequency of a loss event for the Non-Privileged threat.

Cyber Criminals

As I mentioned before, the Cyber Criminals threat is the Threat Community covering all who commit cyber crimes. It is generally considered as the most dangerous and frequent of many threat listings, sometimes second only to the Nation State sponsored actors, groups that operates with the backing of nations. For this reason, we evaluated their capacity to be “Very High”: they have the knowledge and the skills to operate against the assets of the organization.

Another key characteristic of this Threat Community is that they tend to be ubiquitous: just think about how phishing campaigns fill the spam folders of many email inboxes again and again [Enisa 2021]. Being so present and active, it led us to evaluate their frequency as “High”: we believe that it is not a matter of if the organization will get into contact with this Threat Community but of when. You will be wondering why we rated their frequency as “High” and not “Very High”: the reason lies in the organization itself. The MUSE is a museum, it is not a research center or a financial institution that are very palatable for this Threat Community. We expect them to attack the museum, yes, but not as much as if we were a different kind of organization: for this reason, we evaluated the frequency of this Threat Community to be “High”.

In regard to the assets’ vulnerability, the evaluation takes the shape of a balancing between their skills and knowledge and the organization’s controls. On one side we have a very skilled adversary, on the other we have our controls: the strength of such controls had been considered to be enough to stop most of the possible attacks. Of course, the possibility that a very motivated adversary works around the controls or that a human error leads to a loss event are both there; however, as in other cases, we evaluated them to not be as high as to warrant a higher rating of the assets’ vulnerability.

At the end of the evaluation, the Frequency of a loss event turned out to be “Moderate” for very scenario as we balanced the high threat’s capacity and its frequency against the assets’ vulnerability. The scenarios for this Threat Community all have the malicious intent as this Threat Community does not operate with any other intent.

The table gathering the Frequency of a loss event for the threat Cyber Criminals can be seen at Table 18: The Frequency of a Loss event for the Cyber Criminals threat.

Amateur Hackers

The Amateur Hackers were defined during the set up for the analysis as we analysed which Threat Communities to be included as a group of people that prod and attack an IT system for their own curiosity and recognition by their peers and wider public.

For the same reasons we included them in the first place, which were a past case for the organization as well the organization being a target of opportunity, we evaluated the frequency with which the organization gets in contact with them “Moderate”: while not happening as often as the Cyber Criminals Threat Community, we believe it still happens with a reasonable frequency. However, we choose to distinguish depending on the assets: as their primary, and easier, targets are the Web Applications and IT Infrastructure assets, we evaluated the frequency for these assets a bit higher than the Data assets which, the way we see it, gets threatened in a secondary form; being a “secondary target” of sort, leads to the evaluation of the frequency of contact for the Data asset as “Low”.

In the definition, this Threat Community was a moderately skilled Threat Community, less than the Cyber Criminals Threat Community but still dangerous enough to have its Threat Capability rated “Moderate”. In a similar fashion, the controls in place should be able to greatly limit the operativity of the Amateur Hacker Threat Community and thus, their capacity to harm an asset is rated “Low”.

In conclusion, the Frequency of a loss event for the scenarios centred around the Amateur Hacker Threat Community was rated either “Low” or “Moderate” depending on the asset under analysis. The Web Applications and IT Infrastructure assets, being their primary targets, had been rated more at risk of being attacked and thus gained the “Moderate” evaluation. The Data asset, on the other hand, as we said is a “collateral target” and thus had been rated with a lower value.

The table gathering the Frequency of a loss event for the threat Amateur Hackers can be seen at Table 19: The Frequency of a loss event for the Amateur Hackers threat.

Hacktivists

This Threat Community had been defined as the “cyber arm” of the well-known activists which, instead of banners and spray paint, engage in protest action via the internet and with the use of IT tools.

This Threat Community is one that, at the moment of the writing, does not represent a particular dangerous threat; their frequency of contact is very limited and had been rated to “Very Low” as the MUSE represents a target of opportunity for them but not a primary target. Their ability to threaten the assets had been rated “Very Low” as for now: these groups had not shown any particular skill or action that would have lead to a much higher rating.

In conclusion, the Frequency of a loss event had been rated, for all scenarios and assets taken into consideration, to be “Very Low”. We believe this threat could arise and grow to prominence in the oncoming years.

The table gathering the Frequency of a loss event for the threat Hacktivists can be seen at Table 20: The Frequency of a loss event for the Hacktivists threat.

Asset	Intent	Effect	Threat Frequency	Threat Capability	Vulnerability	Frequency of a loss event
Data	Malicious	- Confidentiality - Integrity - Monetary expenses	Very Low	High	High	Very Low
Data	Error	- Confidentiality - Integrity - Monetary expenses	Moderate	Moderate	Moderate	Moderate
Web Applications	Malicious	- Availability - Monetary expenses	Very Low	Moderate	Moderate	Very Low
Web Applications	Error	- Availability - Monetary expenses	Low	Low	Moderate	Low
IT Infrastructure	Malicious	- Integrity - Availability - Monetary expenses - Substitution	Very Low	High	High	Low
IT Infrastructure	Error	- Integrity - Availability - Monetary expenses - Substitution	Moderate	Moderate	Moderate	Moderate

Table 16: The Frequency of a loss event for the Privileged threat

Asset	Intent	Effect	Threat Frequency	Threat Capability	Vulnerability	Frequency of a loss event
Data	Malicious	- Confidentiality - Integrity - Monetary expenses	Very Low	Very Low	Low	Very Low
Data	Error	- Confidentiality - Integrity - Monetary expenses	Very Low	Very Low	Very Low	Very Low
Web Applications	Malicious	- Availability - Monetary expenses	Very Low	Low	Low	Low
Web Applications	Error	- Confidentiality - Integrity - Availability - Monetary expenses	Very Low	Low	Low	Low
IT Infrastructure	Malicious	- Integrity - Availability - Monetary expenses - Substitution	Very Low	Very Low	Low	Very Low
IT Infrastructure	Error	- Integrity - Availability - Monetary expenses - Substitution	Low	Very Low	Very Low	Very Low

Table 17: The Frequency of a loss event for the Non-Privileged threat

Asset	Intent	Effect	Threat Frequency	Threat Capability	Vulnerability	Frequency of a loss event
User's Data	Malicious	- Confidentiality - Integrity - Availability - Monetary expenses	High	Very High	Moderate	Moderate
Managements Data	Malicious	- Confidentiality - Integrity - Availability - Monetary expenses	High	Very High	Moderate	Moderate
Web Applications	Malicious	- Confidentiality - Integrity - Availability - Monetary expenses	High	Very High	Moderate	Moderate
IT Infrastructure	Malicious	- Integrity - Availability - Monetary expenses - Substitution	High	Very High	Moderate	Moderate

Table 18: The Frequency of a loss event for the Cyber Criminals threat

Asset	Intent	Effect	Threat Frequency	Threat Capability	Vulnerability	Frequency of a loss event
User's Data	Malicious / Neutral	- Confidentiality - Monetary expenses	Low	Moderate	Low	Low
Management Data	Malicious / Neutral	- Confidentiality - Monetary expenses	Low	Moderate	Low	Low
Web Applications	Malicious / Neutral	- Confidentiality - Monetary expenses	Moderate	Moderate	Low	Moderate
IT Infrastructure	Malicious / Neutral	- Confidentiality - Monetary expenses	Moderate	Moderate	Low	Moderate

Table 19: The Frequency of a loss event for the Amateur Hackers threat

Asset	Intent	Effect	Threat Frequency	Threat Capability	Vulnerability	Frequency of a loss event
Research and Scientific Data	Malicious	- Confidentiality - Availability - Monetary expenses	Very Low	Very Low	Very Low	Very Low
Management Data	Malicious	- Confidentiality - Integrity - Availability - Monetary expenses	Very Low	Very Low	Very Low	Very Low
Web Applications	Malicious	- Confidentiality - Availability - Monetary expenses	Very Low	Very Low	Very Low	Very Low
IT Infrastructure	Malicious	- Availability - Monetary expenses - Substitution	Low	Low	Very Low	Very Low

Table 20: The Frequency of a loss event for the Hacktivists threat

4.3 The risk evaluation - Loss magnitude

The Loss magnitude is the root value for the right side of the FAIR topology as it is the right child of the root Risk, opposed to the Frequency of a Loss Event which is the left child.

With Loss magnitude we include immaterial losses, such as the data themselves, reputation or the time spent recovering during a loss event and restoring the normal operations, as well as material losses, such as the change of compromised hardware and software, the monetary costs that the organization incurs during and following a loss event. The types of losses had been defined in a similar manner with the threats: we first searched the web for typologies that could apply to our case and then we evaluate them one by one against the organization under analysis.

The Loss magnitude is once again organized by threat gathering the relevant scenarios together.

Privileged Insider

Once again, as for the previous section, we build on the Privileged Insider profile to evaluate the magnitude of the loss this Threat Community could cause.

A Privileged Insider is the threat with the best and most direct access to any of the organization's assets: for this reason, if such threat operates with a malicious intent, meaning the threat is operating with the purpose to do some kind of damage, then the magnitude of the loss could be great. This reasoning led us into evaluating the Loss magnitude for every scenario having a malicious intent as "Very high".

It is quite clear how a privileged insider operating with malicious intent would with near certainty generate a secondary loss and thus, the Secondary Risk would be sizeable. This secondary loss would take the shape, in the majority of cases, of regulatory sanctions, legal procedures expenditures and public relationship expenses.

When the intent is due to an error, the situation is slightly different. For the scenario centred on the Data asset, the probability that a secondary loss event happens following a primary loss event is, once again, quite possible; only if not sensitive data are disclosed there is a possibility of not incurring into a secondary loss or at least the loss being of limited size. However, the majority of the data the organization manage are sensitive one way or another and thus, the frequency of a secondary loss was rated to be "High", the Secondary Risk followed suit and the Loss magnitude was rated "High".

On the scenario with error as intent and Web Application as asset, we evaluated that a secondary loss could happen with less frequency than in previous cases: while a disclosure of data would have a much greater impact, a damage due to error to the web applications alone could, in many cases, be limited and thus lower the probability of a secondary loss to "Moderate". The resulting Secondary Risk follows suit and the Magnitude of the loss for the Web Application asset is set to "Moderate". The Web Applications asset are for the most part software offered on a "Software-As-A-Service" premise and what a privileged insider can do with an error is limited.

For the IT Infrastructure asset alone, an error by a privileged insider can generate a loss that, in many cases, is an internal type of damage that can be limited and the external exposure that generates a secondary loss is itself limited. For example, if a part of the internal network goes down due to a routing mistake, the one knowing it will be the employees but unless there are critical operations with the outside being done in that part of the network, it is something that gets a limited visibility outside. At the same time, the loss of availability of the IT Infrastructure serving clients or partners during a particularly important event may well lead to a secondary loss of varying degree. With a probability and magnitude less than the scenarios with a malicious intent, it still could not be dismissed completely. With this reasoning, the frequency of a secondary loss event and the entity of the Secondary Risk are set to "Moderate".

The effects this Threat Community can generate vary depending on the scenario with the only constant of the monetary expenses, a type of effect that has to be taken for granted the moment we think of a loss event. For the Data assets, the effect of a loss event are a loss of confidentiality and integrity in both scenarios. When an error happens during day to day operations, usually the problem is about releasing sensitive data or modifying

them the wrong way; for the malicious intent, we believe the target would be to modify or make public sensitive data in exchange of a return of some sort. For the Web application asset, we consider the effect of a loss event to be a loss of availability beside the monetary expenses: while the application themselves cannot be touched, the asset can be made unavailable to others. Last but not least, we believe the IT Infrastructure asset could suffer from a loss of integrity and availability as well as requiring substitution: being a physical asset, the availability of the service could be compromised as well as the whole asset integrity when the threat operates against it and, at the extreme, requiring the substitution of some of its parts.

The table gathering the Loss magnitude for the threat Privileged Insider can be seen at Table 21: The Loss magnitude event for the Privileged Insider threat.

Non-Privileged Insider

In a similar fashion to the Privileged Insider Threat Community, for the Non-Privileged Insider Threat Community we have a distinction between scenarios having a malicious intent and ones having error as intent.

With a similar reasoning, a non-privileged insider operating with a malicious insider can do more damage than one with an error intent and, as such, it can lead to a greater probability of generating a secondary loss as well as a greater overall Secondary Risk. However, we believe the probability to generate a secondary loss is limited, it can happen yes, but it is not as probable as with a privileged insider or a cyber criminal; the frequency had thus been rated “Moderate”. Following suit, the Secondary Risk had been rated to be “Moderate” as we could not completely dismiss the possibility of a secondary loss.

However, the controls and how non-privileged insiders are managed still limits the ability of a non-privileged insider to do damage: the Loss magnitude generated by such malicious non-privileged insider remains “Moderate”.

The scenarios with error as intent have a much lower probability to generate a secondary loss event since the access to the assets by the non-privileged threat is limited: while not to be completely dismissed, we set the frequency by which a secondary loss can happen to “Low”. Consequentially, the Secondary Risk had been evaluated to “Low” and the overall Loss magnitude is limited.

The typology of secondary loss is the same for the Privileged Insider Threat Community: regulatory sanctions, expenses for legal procedures and expenses for the public relationship effort following the loss event.

The effect this Threat Community can generate by operating against the organization’s assets varies. The monetary expenses are a fixed effect while the others vary depending on the scenario. In regard to the Data assets, the main effect is a loss of confidentiality and integrity of the asset; in either case we did not consider a loss of availability because we believe the main objective of this threat is to either make sensitive data public or to modify such data. In regard to the Web Application asset, the access this Threat Community has is very limited and with malicious intent the most it could probably do is to try make the software unavailable to others; due to error, however, we regarded a loss of confidentiality and one of availability too as a possible effect. Last but not least, in regard to the IT Infrastructure asset we believe the main effects are a loss of integrity, a loss of the availability of the asset and a possible substitution, the reason is simple: the IT Infrastructure asset is a physical asset that can be touched and physically damaged. Due to error or a malicious intent this threat could physically damage the asset thus limiting the availability of the service, its integrity and in extreme cases, requiring the substitution of physical parts.

The table gathering the Loss magnitude for the threat Non-Privileged Insider can be seen at Table 21: The Loss magnitude event for the Non-Privileged Insider threat.

Asset	Intent	Effect	Frequency of a secondary loss event	Secondary loss (expenses for ...)	Secondary Risk	Magnitude of the loss
Data	Malicious	- Confidentiality - Integrity - Monetary expenses	Very High	- Sanctions - Legal procedures - PR	Very High	Very High
Data	Error	- Confidentiality - Integrity - Monetary expenses	Alto	- Sanctions - Legal procedures - PR	Alto	Alta
Web Application	Malicious	- Availability - Monetary expenses	Very High	- Sanctions - Legal procedures - PR	Very High	Very High
Web Application	Error	- Availability - Monetary expenses	Moderate	- Sanctions - Legal procedures - PR	Moderate	Moderate
IT Infrastructure	Malicious	- Integrity - Availability - Monetary expenses - Substitution	Very High	- Sanctions - Legal procedures - PR	Very High	Very High
IT Infrastructure	Error	- Integrity - Availability - Monetary expenses - Substitution	Moderate	- Sanctions - Legal procedures - PR	Moderate	Moderate

Table 21: The Magnitude of the Loss for the Privileged Insider threat

Asset	Intent	Effect	Frequency of a secondary loss event	Secondary loss (expenses for ...)	Secondary Risk	Magnitude of the loss
Data	Malicious	- Confidentiality - Integrity - Monetary expenses	Moderate	- Sanctions - Legal procedures - PR	Moderate	Moderate
Data	Error	- Confidentiality - Integrity - Monetary expenses	Low	- Sanctions - Legal procedures - PR	Low	Low
Web Application	Malicious	- Availability - Monetary expenses	Moderate	- Sanctions - Legal procedures - PR	Moderate	Moderate
Web Application	Error	- Confidentiality - Integrity - Availability - Monetary expenses	Low	- Sanctions - Legal procedures - PR	Low	Low
IT Infrastructure	Malicious	- Integrity - Availability - Monetary expenses - Substitution	Moderate	- Sanctions - Legal procedures - PR	Moderate	Moderate
IT Infrastructure	Error	- Integrity - Availability - Monetary expenses - Substitution	Low	- Sanctions - Legal procedures - PR	Low	Low

Table 22: The Magnitude of the Loss for the Non-Privileged Insider threat

Cyber Criminals

In the case of this threat community, a Secondary Loss event is to be taken nearly for granted for every scenario. A breach due to a cyber criminal attack means, first of all, the risk of being sanctioned due to the current legislations, an example of such legislations is the European General Data Protection Regulation (GDPR); then, it would mean a nearly certain image fallout leading to an increased expenses in the public relationship effort to make up for the damage. Last but not least, at the extreme case, expenses for legal procedures can happen.

For these reasons, we rated the frequency of a secondary loss event “Very High”, meaning they will happen nearly every time there is a primary loss event. The Secondary Risk follows suit and it is rated as “Very high”. Overall, the magnitude of the loss for every scenario was considered to be “Very High” because this Threat Community is the most dangerous and knowledgeable of all: if they succeed in breaking the defences in place, then a primary loss is going to happen and a secondary loss would follow.

Of course, the secondary loss might not always be sizeable: a good incident response and damage control can limit the expanse of a secondary loss as well as the location of the primary loss can. For example, if no real data was stolen, the secondary loss could be limited while, on the opposite, a theft of clients’ sensitive and confidential data would generate a much bigger secondary loss.

The effect of this Threat Community operations is a loss of the CIA triad (Confidentiality, Integrity and Availability) as well as monetary expenses; in regard to the IT infrastructure asset, a Cyber Criminal attack could entail the physical substitution of hardware and software, for example after a ransomware attack.

The table gathering the Loss magnitude event for the threat Cyber Criminals can be seen at Table 23: The Loss magnitude event for the Cyber Criminals threat.

Amateur Hackers

This Threat Community is similar to the Cyber Criminals Threat Community in the fact that a secondary loss is quite probable if the threat’s attack succeeds: this is due to the fact that there is a high probability the threat will make the vulnerability or the positive outcome of the attack public. Being a less frequent threat community than the previous one, however, together with the fact that there is still the possibility they work with the organization before going public, makes the frequency of a secondary loss event to be evaluated “Moderate”. The Secondary Risk follow suit and it is evaluated “Moderate”.

The secondary loss takes the same shape of the secondary loss for the Cyber Criminal Threat Community with regulatory sanctions, legal procedures expenses and public relationship expenses. The effect of this Threat Community operations, however, is more limited: we expect a loss of confidentiality and the ever present monetary expenses. This decision had been made by reviewing the profile of the Threat Community: their target is to gain exposure toward their peers and the wider public and the easier way to do this is by making public confidential information. Furthermore, they do not have the skills and resources to conduct attacks that could, seriously and directly, impact the availability and integrity of the assets.

The overall magnitude of the loss it is not expected to be as high as with the Cyber Criminals Threat Community. Compared to that Threat Community, the Amateur Hackers Threat Community is less skilled, knowledgeable and motivated, leading us to evaluate the magnitude of a loss from their action as “Moderate”.

The table gathering the Loss magnitude for the threat Amateur Hackers can be seen at Table 24: The Loss magnitude event for the Amateur Hacker threat.

Hacktivists

As explained before, the Hactivist Threat Community is, at the moment of writing, the Threat Community evaluated to be the less dangerous of all the Threat Communities included in the analysis.

The frequency of a secondary loss generated by their action is considered to be “Low” for all scenarios but the ones focused on the Web Applications and IT Infrastructure assets. The frequency for those two assets had

been rated “Moderate” for the following reason. The Web Applications asset includes the website and social media while the IT Infrastructure represents a physical target relatively easy to be reached: they both represent a high interest target for this Threat Community with the main intent of defacing them.

However, albeit the probability of a secondary loss with those two specific assets is higher, the overall Secondary Risk for all scenarios remains “Low” as the Threat Community is considered the less dangerous and less probable. The Loss magnitude follows the reasoning and is rated to be “Low”. While we could have set it to “Very Low” we believe there is a risk with this Threat Community and lowering the rating too much might let it “slip under the radar”.

The main goal of this Threat Community is to demonstrate the wrongness of their target and because of this, we believe the main effect of their actions would be centred on the loss of confidentiality of the assets together with the usual monetary expenses.

The table gathering the Loss magnitude for the threat Hacktivists can be seen at Table 25: The Loss magnitude event for the Hacktivists threat.

Asset	Intent	Effect	Frequency of a secondary loss event	Secondary loss (expenses for ...)	Secondary Risk	Magnitude of the loss
User Data	Malicious	- Confidentiality - Integrity - Availability - Monetary expenses	Very High	- Sanctions - Legal procedures - PR	Very High	Very High
Management Data	Malicious	- Confidentiality - Availability - Integrity - Monetary expenses	Very High	- Sanctions - Legal procedures - PR	Very High	Very High
Web Application	Malicious	- Integrity - Availability - Monetary expenses	Very High	- Sanctions - Legal procedures - PR	Very High	Very High
IT Infrastructure	Malicious	- Integrity - Availability - Monetary expenses - Substitution	Very High	- Sanctions - Legal procedures - PR	Very High	Very High

Table 23: The Magnitude of the Loss for the Cyber Criminals threat

Asset	Intent	Effect	Frequency of a secondary loss event	Secondary loss (expenses for ...)	Secondary Risk	Magnitude of the loss
User Data	Malicious / Neutral	- Confidentiality - Monetary expenses	Moderate	- Sanctions - Legal procedures - PR	Moderate	Moderate
Management Data	Malicious / Neutral	- Confidentiality - Monetary expenses	Moderate	- Sanctions - Legal procedures - PR	Moderate	Moderate
Web Application	Malicious / Neutral	- Confidentiality - Monetary expenses	Moderate	- Sanctions - Legal procedures - PR	Moderate	Moderate
IT Infrastructure	Malicious / Neutral	- Confidentiality - Monetary expenses	Moderate	- Sanctions - Legal procedures - PR	Moderate	Moderate

Table 24: The Magnitude of the Loss for the Amateur Hackers threat

Asset	Intent	Effect	Frequency of a secondary loss event	Secondary loss (expenses for ...)	Secondary Risk	Magnitude of the loss
User Data	Malicious	• Confidentiality • Monetary expenses	Low	• Sanctions • Legal procedures • PR	Low	Low
Management Data	Malicious	• Confidentiality • Monetary expenses	Low	• Sanctions • Legal procedures • PR	Low	Low
Web Applications	Malicious	• Confidentiality • Monetary expenses	Moderate	• Sanctions • Legal procedures • PR	Low	Low
IT Infrastructure	Malicious	• Confidentiality • Monetary expenses	Moderate	• Sanctions • Legal procedures • PR	Low	Low

Table 25: The Magnitude of the Loss for the Hacktivists threat

4.4 The risk evaluation - Risk

Risk is the root of the FAIR taxonomy, evaluated as the product of the Frequency of a primary loss event and the Loss magnitude.

The two children of the Risk were organized by threat as it made easier explaining the reasoning of the choices we did; however, it was decided to present the Risk by asset. At this point in the analysis, we want to know the overall risk our assets are subject to and thus, it resulted easier to present the results ordered by asset instead of by threat. Table 26 to 28 presents the three main assets with their scenarios and define the risk they are subject to.

To evaluate the Risk, we did not have an exact number the same way we would have if we took a pure quantitative approach: we tried to balance the qualitative results we got from the evaluation of the Frequency of a primary loss event and the Loss magnitude and output a value that balanced the two while keeping in mind the scenario under review.

4.4.1 Risk for the Data asset

The Risk to the Data asset varies greatly depending on the Threat Community we look at.

The Insider, either Privileged or Non-Privileged, Threat Community has a range of risk, depending on the privileges granted; the Privileged Insider scenarios obviously presents a much higher risk than the Non-Privileged ones. Going against the obvious reasoning, we evaluated that the risk posed to the asset by an error of a privileged insider to be higher than the same threat with malicious intent: errors happen, they are part of the human nature and, as such, they are much more probable. In addition, they are more often than not exploited by third parties, such as the Cyber Criminals Threat Community, and thus they bring a greater risk than a malicious insider. We are not saying that the risk by a privileged insider with malicious intent is negligible, we are simply saying that we evaluated it lower on the fact that it happens with a much smaller frequency.

The Cyber Criminals Threat Community have some of highest risk scenarios for this asset. Even with the controls in place, lowering the possibility of a loss event, the possible Loss Magnitude can be extremely high since the Data are the primary target for this Threat Community. Balancing these two values, the overall risk results slightly lower but still being ranked “High” for both scenarios.

While the possible Loss magnitude had been evaluated to “Moderate”, the scenarios centred on the Amateur Hackers Threat Community had been considered to happen not that often and consequentially the overall Risk is rated to “Low”.

The Hacktivists Threat Community follows their earlier definitions and the risk they pose to the Data asset had been evaluated to “Very Low”, the lowest for this asset.

Asset	Threat	Intent	Frequency of a primary loss event	Loss magnitude	Risk
Data	Privileged Insider	Malicious	Very Low	Very High	Moderate
Data	Privileged Insider	Error	Moderate	High	High
Data	Non-Privileged Insider	Malicious	Very Low	Moderate	Low
Data	Non-Privileged Insider	Error	Very Low	Low	Low
User's data	Cyber criminals	Malicious	Moderate	Very High	High
Management Data	Cyber criminals	Malicious	Moderate	Very High	High

User Data	Amateur Hackers	Malicious / Neutral	Low	Moderate	Low
Management Data	Amateur Hackers	Malicious / Neutral	Low	Moderate	Low
Research and Scientific Data	Hacktivists	Malicious	Very Low	Low	Very Low
Management Data	Hacktivists	Malicious	Very Low	Low	Very Low

Table 26: The Risk for the Data asset

4.4.2 Risk for the Web Applications asset

The risk to the Web Applications asset varies at a smaller degree than the Data asset but it still depends on the Threat Community we look at.

With either intent, the Privilege Insider presents a “Moderate” risk to this asset: the malicious intent with a very high magnitude of the loss is balanced by the very low frequency while the error intent present a moderate magnitude and a low frequency, leading in both cases to the “Moderate” outcome.

The Non-Privileged Insider Threat Community has such a limited access to this asset that in either scenario the risk it poses is considered to be “Low”.

The Cyber Criminals Threat Community once again has one of the highest possible Loss magnitude for this asset; as the frequency of such loss is rated “Moderate”, we evaluated the overall Risk posed by this Threat Community to the asset in question to be “High”.

The Amateur Hackers Threat Community is “Moderate” in both the Frequency of a primary loss event as well as in the possible Loss magnitude; the Risk posed by this Threat Community had been evaluated “Moderate”.

The Hacktivists Threat Community poses the lowest threat to the Web Applications asset as we rated the Risk as “Low”.

Asset	Threat	Intent	Frequency of a primary loss event	Loss magnitude	Risk
Web Applications	Privileged Insider	Malicious	Very Low	Very High	Moderate
Web Applications	Privileged Insider	Error	Low	Moderate	Moderate
Web Applications	Non-Privileged Insider	Malicious	Low	Moderate	Low
Web Applications	Non-Privileged Insider	Error	Low	Low	Low
Web Applications	Cyber criminals	Malicious	Moderate	Very High	High
Web Applications	Amateur Hackers	Malicious / Neutral	Moderate	Moderate	Moderate
Web Applications	Hacktivists	Malicious	Low	Moderate	Low

Table 27: The Risk for the Web Applications asset

4.4.3 Risk for the IT Infrastructure asset

The risk to the IT Infrastructure asset varies at a smaller degree than the Data asset but it still depends on the Threat Community we look at; overall, it is lower than the two previous assets.

The Privileged Insider Threat Community poses with either intent, only a “Moderate” risk to this asset. The risk that such Threat Community operates against this asset is present and non negligible: for the malicious intent the very high Loss magnitude is balanced out by the “Low” frequency, for the error intent we consider both values to be “Moderate”.

For the Non-Privileged Insider Threat Community, once again there is a possibility that a loss happens but we believe it is minimal: we do not think a non-privileged insider going thrashing the IT equipment to be a very frequent happening. While the possible Magnitude with a malicious intent could be higher, the “Very Low” frequency balances it out and in both scenarios the risk had been evaluated “Low”.

The Cyber Criminals Threat Community poses the highest risk, as in previous case, to this asset: their “High” risk is due to the “Very High” possible magnitude of the loss and the “Moderate” frequency by which this Threat Community can act against this asset.

For the Amateur Hacker Threat Community, the IT Infrastructure is one of what we believe be their favourite target and thus, we reasoned the frequency of a loss by their action to be “Moderate”; being somewhat in the middle in regard to skills and knowledge between Cyber Criminals above and Hacktivists below, we consider the possible Magnitude of the loss to be “Moderate”. The Risk this Threat Community poses to this asset had been evaluated “Moderate”.

The Hacktivists Threat Community poses a limited risk but, albeit their frequency is “Very Low”, we believe this asset could be an easy target for this kind of Threat Community and thus, the possible magnitude of the loss had been rated “Moderate”. Balancing these two results, we evaluated the risk to be “Low”.

Asset	Threat	Intent	Frequency of a primary loss event	Loss magnitude	Risk
IT Infrastructure	Privileged Insider	Malicious	Low	Very High	Moderate
IT Infrastructure	Privileged Insider	Error	Moderate	Moderate	Moderate
IT Infrastructure	Non-Privileged Insider	Malicious	Very Low	Moderate	Low
IT Infrastructure	Non-Privileged Insider	Error	Very Low	Low	Low
IT Infrastructure	Cyber criminals	Malicious	Moderate	Very High	High
IT Infrastructure	Amateur Hackers	Malicious / Neutral	Moderate	Moderate	Moderate
IT Infrastructure	Hacktivists	Malicious	Very Low	Moderate	Low

Table 28: The Risk for the IT Infrastructure asset

In conclusion and in overall, the Cyber Criminals Threat Community is the one we considered to be the more frequent and the more dangerous: we evaluated the risk they pose to the organization's assets "High".

The Privileged Insider Threat Community varies at a greater rate depending on the asset and the intent. We rated to be more frequent scenarios that happens due to errors than the ones due to a malicious intent. Of course, a privileged insider operating with malicious intent can be much more dangerous to any asset than one generating a risk due to error but then, the frequency partly balances out the danger posed by the threat; the scenarios due to error, with a higher frequency, are considered to generally pose a greater risk to the organization.

The Non-Privileged Insider Threat Community is greatly hampered in their possibility to pose a risk to the organization by their "non privileged" status; by the state of the organization, the frequency is very much similar to the Privileged Insider with error scenarios more frequent than malicious ones. Having a limited access and abilities to operate, it lowers their overall risk.

The Amateur Hacker Threat Community is a moderate threat by all means. While not as skilled and knowledgeable as the Cyber Criminals Threat Community and not as "inside" the organization as the Insiders Threat Communities, they can still be a dangerous threat and thus, the risk they pose to the organization had been generally rated "Moderate".

The Hacktivists Threat Community represents the one posing the smallest risk to the organization, either "Low" or "Very Low". The MUSE is a target of opportunity and thus they are not a very frequent threat; then, their skills and knowledge, at the time of writing, is limited and lower, on average, than all the other Threat Communities. While still being a possible threat, they had been considered as the one posing the least risk to the organization.

4.5 The questionnaire and phishing test

The results for the questionnaire and the phishing tests were organized in a report that had two objectives: educating and reporting. We divided it into two main sections: the first about the questionnaire, the second about the phishing test. For each sub section of the questionnaire part, we went on explaining the question's subject, we presented some key results and we proposed the relevant good practices.

When presenting the results to the museum, we choose to show only the relevant percentages; we did not want to clutter the report with an excessive amount of numbers which would have only confused the reader. For example, if in a question we had an 80-15-4-1 distribution of percentages, meaning 80% gave answer A, 15% gave answer B and so on, we would have presented it as "the 80% did right" or "the 80% did it wrong and only the 15 did it right". The possibility to ask for more detailed numbers was available.

The museum has a population of about 180 employees and we gathered 85 responses; if we take a confidence level of the 95% and an error margin around the 8%, we can say the sample we got is statistically relevant for the whole population. Aside from the purely statistical relevance, I and the IT staff at MUSE believe the results we got are relevant and they give us a general overview of the museum's employees and their behaviours.

4.5.1 Policy knowledge and hacker perception

Policies the way they are usually done are often difficult to be read and fully understood; our expectations were to have a minority professing great knowledge of them while the majority knew it in general terms or not at all. The data mostly followed our expectations:

- A minority, amounting at the 5% of the responses, profess perfect knowledge of the organization's policies.
- The majority, amounting at the 34% of the responses, profess a fair knowledge of the organization's policies.
- A sizeable percentage, amounting at the 31% of the responses, profess sufficient knowledge of the organization's policies.
- A sizeable percentage, amounting at the 31% of the responses, profess no knowledge of the organization's policies.




Figure 6: Graph of the questionnaire's answers ordered by policy knowledge

As we noticed the above distribution, we wondered if the knowledge of the organization's policies on IT and cyber security could be a possible indicator of how good or bad people behaved in regard to cyber security. Do people with greater policy knowledge do less mistakes than people with no knowledge? Or there is little difference?

A perfect knowledge of the policies does not mean perfect technical knowledge: the minority that professed a perfect knowledge of the policies still had some wrong answers, for example in recognizing dangerous file extensions or in regard to how to behave during a ransomware attack. In the other three groupings, there had not been a marked difference in the correctness of the answers between greater and lower policy knowledge.

Our conclusion is that the policies the way they are done now matter up to a certain point while the knowledge coming from other sources, such as hands on experiences or training courses, results much more important.

To get a bit of a “feeling” of our employees, we decided to ask about how much do they thought possible that a hacker could replicate a business’ website and that the site would look like the real one at a cursory inspection. The idea behind this question is to try see how much they are technological in deep, how much they believe a hacker’s ability are and how much they expect to be possible.

- The majority with the 55% of responses believes it is moderately possible for a hacker to accomplish such a feat.
- A sizeable percentage, about the 25% of the responses, believes it is something that is greatly possible.
- A sizeable minority with about 14% of the responses admits to not know if such a feat is possible.
- A minority with the 6% of responses believes it is not something possible.




Figure 7: Graph of how the employees perceive a hacker's abilities

This question had neither a right nor wrong answer, much like the previous one on the organization’s IT policies. While we were a bit surprised that a minority did not believe it possible in any way, we expected that the majority would consider it at least partly possible and that a relative minority would be unsure on such a question. The results agreed with us.

4.5.2 Phishing, emails and malicious attachments

For the section on phishing and malicious attachments, we gathered the following results:

- The majority with the 86% of the responses had given the right definition of phishing.
- The majority with the 86% of the responses had correctly identified all the characteristics of a phishing email.
- The majority with the 88% of the responses had correctly identified phishing as a threat to every member of the organization.
- The majority with the 52% of the responses had correctly identified every file extension that can hide a malicious payload. The .exe extension is the second choice with the 21% of the responses.
- The perception if a mail is phishing or not changes depending on the subject and origin:
 - A mail with “You won [something]” is unanimously seen as phishing.
 - A mail in regard to an “Account’s compromise” is unanimously suspect.
 - A mail “unexpected but coming from the inside of the organization carrying a particular attachment” is generally suspected, nearly everyone takes care to inquire with either the senders or whoever is in charge of security.

In general, we can say that phishing is known and can be recognized as such by the museum’s staff. A bit harder is to correctly identify all the file extensions that can carry a malicious payload, but it was kind of expected since in the museum there are many non-technical employees.

An important insight is how the subject/topic and origin of the mail changes the behaviour of the reader. Some, by now, “well known” attempts such as the “You won [something]” are unanimously regarded as phishing; we believe this is due to being one of the oldest and well-known phishing subjects. The responses changes when the external email comes with a “Account’s compromise” subject: these emails are seen as suspect but not outright phishing, many would try to log into their accounts by other means to check if the email is truthful. The responses changes again when the mail comes from the inside of the organization: internal mails are generally considered to be safe or, at most, are seen as suspects; usually the employees prefer to ask the original sender about them.

This situation where origin and subjects change the conscious behaviour of the reader is quite interesting and it could be a follow-up work to study if it happens in other and different settings. An internal account compromised could be the bridgehead not just for lateral movements inside an organization’s network but, at later stages, even a platform from where to launch internal and targeted attacks.

4.5.3 Malware attacks and induced malfunctions

For the section on malware attacks, we got the following results:

- Only a minority, amounting to the 6% of the responses, is able to operate correctly during a malware attack.
- Every response takes care of informing who is in charge of the cyber security.
- All participants had been able to recognize a ransomware attack.

People do not fully know how malware works, how fast and dangerous can be and they are missing such knowledge: only a minority is able to operate properly in case their workstation is under a malware attack. This result was expected as malware are not “an everyday thing to see live” as phishing is; however, every respondent would have taken the care to inform who is in charge of cyber security which is not a bad thing as knowing something is happening is key to react in time.

It was surprising how all participants were able to recognize a ransomware attack: our guess is that they recognized it since it is an attack that keeps getting on the news as more and more organizations fall prey to it and thus they have seen screens similar to the one we described before.

4.5.4 Web security

In the web security section, we got the following results in regard to the difference between the two internet protocol versions http and https:

- The majority with 62% of the responses had been able to recognize correctly the difference between http and https.
- About the 35% responses declares to not know the difference.

When it was asked what was done to understand if a web site is secure, we got the following results:

- The majority with 58% of the responses declares to rely on the web by searching the website through google.
- A minority with 25% of the responses checks the presence of the lock symbol in the address bar.

Http and https are the two version of the same protocol at the front when talking of secure communications via the web; https is by now a well-known protocol and widely employed, even more after Google’s decision to show as insecure and penalize in their search engine’s results the website that did not support https natively in 2018 [Google 2018]. While not perfect, the responses we got were mostly positive and this is a good thing: a sizeable quantity of the employees have some basic knowledge of how the web works.

When checking if a website is secure, it dominates the web search: people rely on the web to see if others had an opinion or had encountered a problem with the site they are looking at. Only a minority takes a “technical” approach and checks if the lock, symbol of the https protocol, is present in the address bar. While searching the web is by itself not a wrong approach, it highlights a lack in basic technical knowledge because the lock is not recognized to be something connected to security.

4.5.5 Password

In the password section, we got the following results:

- The majority with 45% of the responses makes use of a password manager.
- The 32% of responses assert that they create simple passwords as they are easier to be remembered.
- The 18% of the responses mentions the usage of a excel file to gather all their passwords.
- A minority amounting at the 5% of the responses assert that they make use of a post-it under the keyboard.




Figure 8: Graph of how the employees store their passwords

The key insight is how a sizeable amount of the respondents creates weak passwords on purpose: such behaviour is the most dangerous one as a weak password is usually easily guessable and an easily guessable password means an easier work for a motivated threat. Weak passwords had been the trojan horse of many breaches, see the recent [Miller 2020], and it is extremely important to put remedy to this behaviour by raising awareness of the risks and providing adequate solutions to counter the problem.

A minority rely on an excel file to store their password; while it can be somewhat less dangerous than outright easy to remember passwords, it can still present a threat to the organization: a malicious actor with access to the machine storing such file can have unrestricted access to the file and its content. The lower risk is simply because a malicious actor needs access to the physical machine and storage which is usually harder and more complex than guessing a weak password.

We were pleased to register how the 45% percent, which is overall the majority, makes use of a password manager; the IT staff and cyber security head of the MUSE rightfully pushes the use of these tools as they can store passwords securely and they can help in creating passwords without relying on one's memory.

4.5.6 Working station security

In the working station security, we got the following results:

- The majority with 74% of the responses behaves correctly when leaving their working station by locking their computer.
- The 20% of responses leave the working station as it is.
- The 6% of the responses only power off the monitor.




Figure 9: Graph of how the working station is left during office hours when the owner is not present

The working station is the physical door to the IT infrastructure of an organization and its security is as relevant to the whole organization cybersecurity as it is the security of the software they use. If a malicious actor happened upon an unlocked machine, he or she would have complete and nearly unlimited access to the machine itself and the access to the organization's infrastructure would be limited only by the privileges of the logged account: plenty of time and space to see, copy and modify some confidential data or upload a malware inside the computer for example.

There is a good level of education in this regard as the majority behaves correctly by locking the working station before leaving it, even when pushed in a rush.

A sizeable minority, however, leaves their working station as it is, unlocked and without protections. To fully understand the extent of the risk posed by this behaviour we would need to understand where these people are located and work inside the organization: a working station in a back office, after multiple controls poses a relatively smaller risk than one close to the public areas of the museum. However, we believe there should be more emphasis on correct behaviours to lower the risk coming from an unlocked and unattended work station.

A minority limits themselves to powering off the monitor: we like to see this behaviour much like leaving home and closing the door without turning the key, the door looks closed... until you try to open it. If we work to lower the percentage of employees leaving the working station unlocked, we believe we can lower the percentage of people only powering their monitors off as well.

4.5.7 The phishing test

The phishing test results had been compiled by inspecting the access data to the server hosting our fake login in the 36 hours after we launched the attack.

- The 70% of who received the mail did not click the link.
 - The 51% asked about it to who is in charge of IT and its security.
- The 29% of who received the mail clicked the link but did not insert their credentials.
- The 17% of who received the mail clicked the link and inserted their credentials.
 - The 14% of the total had remained completely unaware.
 - The 3% noticed the anomaly after having inserted the credentials and asked about it to who is in charge of cyber security.

Some people fell for it. We expected it. Some stopped short of inserting their credentials while some did and remained completely unaware it was a phishing attack. A minority recognized there might have been something wrong after they inserted their credentials and they got in touch with the IT department; albeit having fell for it, in a real situation, their warning could be key to react in time before their mistake turns into a full blow attack.

We were pleased to register that about half of the recipients got in contact with who oversees the IT infrastructure and the cyber security of the museum either by mail, phone or instant messaging. This means that people tend to prefer asking who is more knowledgeable of the topic: asking before doing means the IT staff can notice an incipit of attack before it develops as well as presenting a possible moment of education.

Another key insight is that internal emails are generally trusted to be safe and they are checked less thoroughly than emails coming from the outside. This behaviour was already visible in the questionnaire where internal mails turned out to be less suspect than external ones, but with the phishing test we proved it empirically.

The way we launched the attack, very late in the afternoon when people were leaving or had already left the working place at the end of the working day, tried to simulate an attack trying to get under the radar. We saw the peak of accesses right at the launch of the attack and on the morning after when they were starting a new working day. Many accesses came from outside the physical perimeters of the museum since many accessed the mail from their homes and because many were, at the time, working remotely due to the Covid-19 pandemic.

By looking at the type of browsers people used, newer versions of Chrome, Firefox and Safari made the most of the accesses and this tells us that people do make updates to their browsers. Older browsers can carry vulnerabilities and as we update software key to the business, the same should be done with browsers and other tools used by employees to access the organization's resources by remote. Operative systems too resulted up to date with newer Windows OSes leading, followed by Apple's Mac iOS, Android and Linux distributions.

4.5.8 Results conclusion

The MUSE – Museo delle Scienze takes cyber security seriously: by relying on enterprise level software, they try to do the best possible working with what they have and what is around them. Given the context, there are many employees that are not properly educated on cyber security but, as I said before, it was to be expected; there is still a sizeable percentage with good behaviours and knowledge.

In truth, when I choose to work with MUSE on their cyber security, I expected to see their cyber security a lot more lacking: I expected it to be less comprehensive, thinner and in general less. I admit I was pleasantly surprised to see it was different. As a cyber security student, it is “painful” when seeing how much cyber security is sometimes undervalued opposed to how much harm could happen. As I said before and will say again, museums should care about cyber security much like any other organization should.

Of course, there is more to be done, there always is: the perfect system does not exist after all and in such a fast changing landscape like the cyber security one, even a perfect system today will not be perfect tomorrow. However, the cyber security stance of the MUSE is already quite the good one on par to many enterprise businesses.

At the end of my internship, I left one final document that gathered a few suggestions to improve their cyber security. These suggestions are:

- Investing in the employees’ training on cyber security. As I will explain in a few pages, employees are the first line of defence and education is key. With this suggestion, I advised to monitor said training in the long run.
- Updating and modernizing documentation and policies, to be repeated in the long run.
- More empirical testing of the IT infrastructure, to be repeated in the long run.
- Some technical suggestions regarding some technical choices they did.

Many of these suggestions will make more sense after reading the Evaluation section of this thesis which will follow, where I will present my take on how cyber security could be inside museums and present some core concepts of cyber security.

5. Evaluation

Our work had been done in collaboration with MUSE – Museo delle Scienze which, albeit being active in the fabric of the city and the region, it is still considered a medium sized museum both for the size of its collections, of the number of employees and for the number of visits.

The results we obtained are not to be taken as a picture of the whole museal sector since the sector sees many variations and changes depending on locations, typology and sizes. An open museum built in a forest or on a historic area, for example the roman Pompei area or the World War I trenches and forts dotting the Alps, has not the same needs and constraints of a museum housed in a historical building in the middle of a bustling city, for example the Musei Vaticani in Rome or the Uffizi in Florence, nor the same of a museum housed in a modern building like the MUSE. The same goes for the traffic they are subject to: a small, local museum has not the same needs and constraints of big, internationally recognized museums.

However, starting from the empirical results we got during our work at MUSE, we can try define a generalized approach to cybersecurity for the museal setting that can be at the base of the cybersecurity programs tailored to the specific situations. Some concepts and processes that lies at the core of a good cybersecurity posture are the same whatever you are working with a small system or a big system; what changes is the scalability of the solution employed to follow the chosen processes and concepts, not the core concepts and posture.

I would like to define the set of topics lying at the base of building a cyber security program as “Building blocks”: “building” is about creating something from something else, we build a house, we build a bike, but building is not limited to material constructs, it works for intangible concepts too, for example we build our career, we build our know-how, we build our mindset. In our case, we are building something complex: we are building cyber security, we are building the parts that make up a complex and multi-faceted structure. “Blocks” are self standing pieces of material, multiple blocks together form bigger blocks; in our case we have these parts, for example documentation, people, etc., that are both self-standing, connected to others, and influencing and getting influenced by others. They are not phases, defined as “a distinct period or stage in a series of events or a process of change or development.”, they do not happen in a specific point of time and in separated ways, they are tightly interlinked and happen all together and all together are needed to build the “greater” cyber security block. We can see these blocks much like Lego® bricks which, when put together, can build a complex construct.




Figure 10: The building blocks of a cyber security program

To start the dive into these blocks I have shown in the above figure, I would like to highlight three initial steps which make the first block “Building the mentality” up and which I believe should take place early on as well as at a higher level than all the others: specifically, this higher level is the mentality level, it is in the way museums and the people managing them think.

The first point is that museums, much like many businesses, have to get out of the mindset that “it will not be us”. It will. It is not a matter of if, it is a matter of when. If you are connected to the internet you are a possible target, whatever you are a museum, a business or a non-profit or something else; it can happen to be attacked and it can happen to museums too, in fact it had happened already [Cascone 2019][Cascone 2020].

Then, the second point would be to enter in the mindset that doing cyber security is vital to the well being of the organization. How many businesses, in the rush to go to market, cut on cyber security only to make the headlines because they had been breached or because some researcher found some vulnerabilities in their products, with the resulting damage to revenues and reputation? Doing cyber security should be part of the normal operations and should be budgeted accordingly. Cyber security is not just some kind of security that is not visible and thus not worth, it is a security for the whole organization’s business, for its employees, for its image and overall health. It could be seen akin to insurances: you pay for it, you invest on it, hoping you will never really find out you need it.

The third point is to be agile. Cyber security is not a field that can wait political deliberations of one kind or another that take the time they take. Cyber security is an ever evolving field and new threat and vulnerabilities appear every year; such fast evolution is something which must not be undervalued: museum should become “agile” to be able to adapt to such a fast changing landscape. However, becoming agile is a matter of practices employed during the organization’s operations as much as it is a matter of mentality, of how people think. Agile methodologies, as seen in the software development field, rely on speed and quick feedbacks to build better products in less time with less wastes; I believe the same approach should be taken in regard to cyber security, relying on quicker decisions to adapt faster and better.

I believe these three initial steps, this initial block, are key to get into the right mindset to make cybersecurity truly effective in the short and in the long run; without them, all the other things that “should be done”, such as keeping doing patches or putting state of art firewalls on the network, are not as effective as they could be. If people do not “believe”, they will not put in the effort.

Let’s think on it. If they keep believing it will not be us, they will be under great risk of being breached and, for how things are done now, it will happen: ransomware attacks are often “carpet attacks” trying to encrypt whatever system they find, the same can be said for phishing that simply attacks everyone with an email box. Then, if a museum does not believe in cyber security being important, they will not give it the budget required to defend the museum itself, resulting once again in a greater risk of being breached with the consequential damages if such event happens. Last but not least, if they are not agile enough, by the time they recognize the change in the landscape, it might be already too late for them to defend themselves.

It is a great shift, true, but it is a shift that can save museums from having to deal with worse situations; people go to a museum in their free time, by their own choice, they are not constrained by business choices and contracts: if a museum’s image is tarnished, the damage is real and happens at that moment in time, a museum does not have the time before contracts expires to repair the damage and the image, like a business might have.

As we said, cyber security is more than just putting a nice firewall on the network: cyber security is, at its core, a structured process. In this process, I see three main steps after we got into the right mindset as we highlighted above: starting with the evaluation of the subject’s needs and constraints with the relative planning, down to the implementations of the chosen technical solutions, past the monitoring and evaluation of the results given by the technical solutions and back up to the planning phase. It should be a cycle that does not run once, stopping after a technical solution had been implemented and is running, but that repeats itself continuously.




Figure 11: The cyber security cycle.

The reader might wonder why I first declared cyber security is made of building blocks and now I'm declaring it is a cycle. The reason I made both statements is because they are both true statements. When we go through the cycle, we work with the building blocks I mentioned: the cycle needs the building block to build cyber security and the building blocks need the cycle to be built. It is a lot like looking at the light which is at the same time a wave and photons, it has a dual nature: wave-like and particle-like. Doing cyber security is dual natured too: process-like and building blocks-like.

Let's take a step back and go back inspecting the building blocks we defined in figure 10. The first block, Building the mentality, has happened: people believe cyber security is important and the shift in the organization practices has happened with the process rolling and the cyber security part of the organization's operations. What now? Now, there is more building to be done of course, but the next "building blocks" are blocks that does not necessarily happen one after the other but they can, and I think they should, happen more or less simultaneously. The first of these "building blocks" is about people.

People. People make up an organization, an organization without people is just an empty husk and the people define the organization itself with their beliefs and choices. But what about people in a cyber security context? People are an asset as much as a threat, but they are something else too.

The people, the employees are the true first line of defence of any organization [Callahan 2021]. That is what they are before being a threat, before being an asset, they can be the baseline of a cyber defence. The system can be secured with the best technology around but when the user, the employee, does a mistake, the best technology might not be enough to stop a malicious actor from taking advantage of such mistake. I like to stress a lot on this point as I believe it is extremely important: cyber security had mostly focused on the technical aspects of securing an organization, but I believe they should take a greater interest in "securing" the employees too, the "human element", for the organization's benefit but, at a much larger scale, for the wider internet's benefit too.

If people can be more educated on many aspects compromising cybersecurity, they can become less at risk of doing mistakes that can, and had been, deadly [Miller 2020]. An easy example is about passwords: people often build passwords easy so they can remember them by heart. However, a password to be remembered by heart is usually not secure as it is easily guessable and thus, it makes easier the work of a motivated attacker to gain access to the password protected asset. The password is an easy example which nearly everyone had seen or come in contact with, but it does not stop there.

Obviously, if an employee is less at risk of doing certain mistakes because it had been properly educated, the risk of the organization they are part of diminishes too. The risk at their home is minor too; it is possible to hope they spread the knowledge to who is close to them, such as children and parents, and consequentially they can help making the whole, greater, internet more secure.

Education had always been key to the progress of humanity as a whole and it is no different when we talk about cybersecurity. Education can help the progress of cyber security. A museum, like any other organization, must “build the people” not just their IT systems: awareness and education are key and should be at the core of any cyber program.

There are many ways to raise awareness and conduct trainings, the problem lies in choosing the best solution that fits better with the organization. There are many factors, such as employees’ background and attitude and the organization’s budget and structure, but employees’ training is not “one solution fits all” problem: it requires careful evaluation in both the short and the long term but the goal should always be to strive for the highest knowledge retain rate, in both the short and in the long term, possible.

Much like the whole cyber security, trainings and “building the people” is not a once done phase but should be repeated and done again constantly as new technologies, challenges and solutions arise. The quest for knowledge is a never ending one, after all. For such a quest, it is important to monitor the retention of knowledge as we mentioned before: how useful is a training where the trainee remember little after just few months?

As we said, “building the people” is one of many building blocks that can happen simultaneously. Building the people is helped by providing documentation about the cyber security in place: it is about documentation the block “Building the documentation” which is the next building block will go on inspecting.

Documentation. Verbose, maybe true, not helpful, quite untrue if done right. Documentation is about putting down on paper what is being done, what should be done and who is doing what.

Let’s imagine a scenario where an organization is trying to build a cyber security program but nobody takes care to write down what should be done, what they are doing and who is in charge of what. It is a practical anarchy of sort, where confusion is queen and the efficiency and productivity are at their lowest. What about if there were the right documents to define and scope the work? Confusion would be at its minimum while efficiency and productivity would soar high.

Now, let us imagine a more extreme scenario where the organization is under a ransomware attack. Due to the lack of proper documentation, there is no clear chain of command on who should take the lead to stem the attack as well as no clear reference on what should be done in such a dire case. What if documentation had been prepared beforehand? The obvious outcome might have been different.

Preparing and maintaining the documentation about the cybersecurity posture of the organization is key and it brings multiple benefits:

- It provides a formal guide to approaching cyber security, defining tools and processes to be applied both during normal operations and during emergencies: we know this and write down we are going to use this and that because this and that.
- It can provide a clear chain of command, especially during emergency situations: we know this employee did this and that employee did that, this employee is in charge of this and that employee is in charge of that.
- It can help spotting mistakes and improper behaviours both from the system and the employees: this mistake appeared as the solution had or had not been implemented right, this employee or that system are showing anomalous behaviours.

Such documentations can take the form of a single document or many documents depending on the preferences and sizes of the organizations: a big museum with many branches and departments might want to structure

such documentations hierarchically by having each branch build their own policies from a core general policy, while a smaller museum could just have one document for all the aspects covered by their cyber security policies. As for technical cases, the core idea is the same, what changes is the solution applied.

An important point about documentation is readability. Not everyone is able to understand the technicalities that sometimes are needed to explain some parts of cyber security. I believe that documentation should be divided into two parts: one part should be understandable by non-technical people, sacrificing the technical details to reach a wider audience, while another part, on the opposite, should be technical and go into details as much as needed. However, the documentation should be widely available to the employees: the non-technical documentation should be available to any and every level in the organization and be part of the onboarding processes for any new employee while the technical one, with some obvious restriction on levels and roles for security reasons, should still be available to a non-technical employee with the proper accesses.

Of course, the writing of such documents should be important, and great care should be taken to not end up with documents that are just too long-winded, taking in the wrong details and leaving out something important, or, at worse, ambiguous and, in general, all the mistakes that makes the inspection, the reference and the reading difficult.

To give a bit of reasoning to why how a document is written is important, just think about people reading habits: it is known how a bad writing influences how much care and time is put into reading. How many people do read the EULAs and all that documents that are written with legal languages that result very difficult to comprehend by all the non-legal people?

If we are not satisfied by that, let us look at a scenario: a cyber security breach had happened, some actions had been done that did not went as expected with an increased loss by the organization. In the post-action investigation, it turns up such actions had been done referencing some internal documentation that turns out to be ambiguous; now who is at fault? Who took the decision that led to a greater loss following a documentation that did not provide a clear answer?

Last but not least, documentation should be kept up to date. It is extremely important that documentation follows the technical and organic changes an organization faces. If the people, the processes or the IT solution change, the documentation must be updated and we could go as far as evaluating if new technologies growing to prominence should be included in certain documents. Being up to date is extremely important, even more during emergencies: a documentation that had not been updated does not and cannot provide good, reliable answers when queried and leaves more space for the human error.

Now, we will take a bit of a detour in a block into which the documentation block kind of bleeds: Building the intelligence. Intelligence is about understanding things, about gathering knowledge and learning. Why intelligence in cyber security, the reader will be asking. Knowledge is power, since the dawn of times; not for nothing the intelligence services of nations were born from the attempt of knowing more of their opponents than what was visible and noticeable.

Cyber security can be seen much like a game of two: defender and attacker. The defender tries to secure the systems, tries to stop the attacker from getting in, the attacker on the other hand is, as the name implies, all for the attack, prodding, testing, directly attacking what the defender is defending. Like any game, like any real conflict, it is not something that can be played without information, without knowledge. In our world where information had been constantly gaining value, just think about the big data where their value lies in the information a skilled analysts can extract, information had grown increasingly important in the cyber security world too.

In the cyber security, information is about vulnerabilities and type of attacks, about systems and security solutions, about threats and technologies. Attackers gather all these information in multiple ways: they gain them during their unlawful operations, they trade them, they learn them from lawful channels. The defenders, on the other hand, had historically struggled with building the same knowledge, usually limiting themselves to lawful channels and finding themselves often short compared to their enemy.

The situation had been changing in the past few years as there had been a push on the defender side for a more widespread knowledge sharing between lawful entities together with the usage of the same tools and methodologies attackers had been using, such as penetration testing and white hacking. It is in this context that I believe the “Building the intelligence” block takes shape.

An organization cannot defend blindly: if you do not know who and what you are up against, how can you defend? You only risk to be swinging a fictional sword to the empty air while the attacker sneaks behind you. The same goes for the cyber security field: you need and must know who and what you are up against, hiding the head in the sand is not an option.

In certain settings there is a way of saying: “Sharing is caring”. This, as strange as it can seem, it applies to cyber security too: it is not about sharing physical objects, it is about sharing knowledge. Sharing the knowledge means others can learn, can secure themselves and hopefully not find themselves in the same situation once again. With this idea in mind, frameworks like Common Vulnerabilities and Exposures (CVE) system [CVE 1999] or the OWASP Top 10 [OWASP 2021] had been born. Other projects, like Google’s Project Zero [Google, 2014], had been born with the same idea but with the addition of finding a problem before the problem even happens, by actively researching systems and IT solutions, much like an attacker would.

For an effective cyber security, an organization should build a knowledge pool about what they have at hand and about what is happening outside their perimeter. Such knowledge building takes the shape of two sources: one internal and one external. The internal source is about knowing what you have and you are working with: monitoring is the easiest way to gather knowledge and insights from your tools. External sources are business, open sources initiative and programs that shares insights, warnings and solutions related to cyber security.

Let’s look at a real example: the MUSE, where the work you are reading originated. In regard to cyber security insight, the MUSE has two main external sources and one main internal source. The two external sources are the GARR Consortium [GARR 2021], Consorzio Gestione Ampliamento Rete Ricerca, and the Microsoft Corporation, while the internal source is their own monitoring system.

The GARR Consortium is a nation wide organization connecting research and education centres and providing various kind of services to them such as network accesses, security and more. For the MUSE, one of the services they provide is the internet access and the GARR monitors closely their network, ready to warn any of their members of anomalous traffic running to and from them. Microsoft Corporation is the well know multinational technology corporation which, as part of their work, works together with third parties on security issues and issues security warnings and information. Internally, they rely on the monitoring of their system to gather insights and eventually spot anomalous behaviours.

Building knowledge with internal and external sources has another extremely important characteristic: time. This process should be fast and continuous as certain vulnerabilities do not leave much time before being used by malicious actors. As example of how much time is sometimes the essence, I would like to remember the time of the outbreak of the ransomware Wanna Cry [Wikipedia 2017] which swept the world and created problems to many organizations in May 2017: Microsoft issued a security patch [Misner 2017] to fix the vulnerability [Microsoft 2017] making the infection possible to Windows machines for all their operative systems from Windows Vista onward weeks before the attack while for older, unsupported versions, such as Windows XP [Warren 2017], they released a patch while the attack was still running and making the headlines. How many did not patch and found themselves infected? The fight between attackers is played at a very fast pace and a defender’s knowledge should and must not lag behind.

Until now we discussed about people, about documentation and about intelligence, but what about the technical defences? It is at this point that the next block comes into play: Building the implementation.

Implementation is defined as “the process of putting a decision or plan into effect” and, in software development, it specifically means the code making some kind of abstract concept work “for real”. If you

remember the cyber security cycle we showed in figure 11, the implementation came after we evaluated the needs and constraints of the organization and prepared a plan of action: thus, we are building what we planned, we are building the real implementation of the solutions we thought in the first phase.

Building the implementation is quite a complex block. It is influenced by all the others, the people influence which solutions are chosen, the documentation puts constraints on solution, the intelligence scopes and directs which solutions to be applied, but at the same time solutions influence them back as the solutions influences how people behaves, the documentation is written based on solutions and the intelligence comes in part from the solutions running.

Building the implementations spans both the hardware and the software: in the hardware you have devices such as physical machines and network devices, in the software you have all the software related to the cyber security being run, such as firewalls, antivirus software, access controls and identity management software.

In the hardware side of the implementation, it is on the organization to inspect carefully, together with external support if there is not enough knowledge internally, which hardware to get. Some evaluation metric to be kept in mind when looking for hardware are:

- The trustiness of the vendor: sometimes it is better to rely on a well known, internationally renowned vendor which can be trusted to deliver a safe and certified product with a known supply chain than on a niche vendor without much history.
- Support availability: support for the hardware should be a key evaluation metric as hardware firmware are often proprietary and a hardware that lags behind in a matter of security can pose a vulnerability to the entire organization.
- Future proof: hardware should be future proof for the longest time possible. A hardware that after a brief span of time is no more supported and cared for by the vendor puts the organization at risk. This metric is tightly linked to the support availability metric.
- Needs and constraint satisfaction: the hardware should satisfy the needs and constraint of the organization.
- Data usage: while less relevant than with software, as we will see, what kind of data the hardware inspects should still be checked to defend against supply chain attacks and hardware trojans.

In regard to software, before running into metrics that are pretty similar to the hardware metrics, we need to look at a dichotomy that had characterized the software development since the seventies: proprietary vs free/open source.

The difference lies in how much rights the software has. A proprietary software heavily restricts the rights of a buyer: a buyer can use it but for example it cannot modify it or share it again. On the opposite, a free/open source, depending on the license, has a much broader set of rights available to the buyer.

Beside the purely legal view, the organization should evaluate carefully whatever they have the skills to deal with free/open source as they often require greater technical knowledge to be managed opposed to a “ready to work” package of many proprietary software.

Both “worlds” have their pros and cons and there is not a clear winner, the goal should be to fulfil at best the needs of the organization and not to have an ideology tower above the everything else. I believe that a middle way approach can often be taken by mixing together in a coherent and functional construct both free/open source software and proprietary software, taking the best from each side.

As I mentioned, the metrics of evaluation for software are, in some cases, similar to the metrics for the hardware but there are some specific aspects that comes into play when speaking of software.

- Software licenses and legal constraint: the legal details of software should be carefully evaluated to not end up with licenses not fitting the usage and other legal problems.
- Longevity: as for hardware, the longevity of a software should be carefully evaluated and kept in mind. As before, a software without updates after a brief spawn of time can represent a risk to the whole organization's system, aside from representing a waste of resources in the first place.
- Support availability: as for hardware, the availability of after-sale support is important as it can be the change between a problem being solved in a limited amount of time opposed to dragging on and impacting the organization's operations. This metric is tightly linked with the longevity metric.
- Data usage: cloud solutions are taking the lead in many software and cyber security is not an exception; organizations should evaluate carefully what data the cloud software under scrutiny gathers and processes as well as where these data are stored. For example, given the different legislation on data protection between the United States of America and the European Union, a European organization could choose a solution storing data in the European area to comply with the legislation in place.
- The trustiness of the vendor: as for hardware, the vendor should be trusted to lower the risk that a malicious attacker uses software to introduce backdoors and malwares inside the organization's system.
- Easy of use: how easy is the software to be used, this metric is particularly relevant for software that gets in contact with non-technical employees. If a software is hard to use, people will find ways to not use it.
- Needs and constraint satisfaction: the software should satisfy the needs and constraint of the organization.

Software and hardware make up the IT infrastructure of an organization but how these two components are interconnected and structured is, from a security perspective, as much important as choosing the right components. Building infrastructures had been subject to many studies and real experimentation and certain concepts had arisen to preeminence. These concepts are:

- Separation of concerns: parts that have nothing to do with each other or that have different needs should be separated. An example of this are the separation of guests' network and internal network as well as the usage of Demilitarization Zones [CISA 2021][ESnet 2021] for servers and applications requiring public external accesses while the internal server and applications lies behind greater defences. Such separations permit to better control the single zones and the links connecting them as well as to isolate more easily zones getting infected by a malicious actor.
- Least privilege possible: software as well as users should have just enough privileges to get their work done, no more no less. Too many privileges are not necessary but provides a broader surface of attack while too little privileges hinder the operation of the organization. An "extreme" approach building on the least privileges concept is the zero trust security model where there is no trust at all and every single action has to be authenticated.
- Security by default: security should be included right away, starting from the designing phase onward, when building an IT infrastructure and not something being added later, on top of all the rest like if it was an optional feature.
- Defence in depth: cyber security should not be just a thin layer at the edge of the IT infrastructure but should permeate all levels of the IT infrastructure. A defence in depth is about designing and implementing multiple layers of security controls in such a way that a malicious actor must compromise multiple levels and multiple controls to be able to freely operate inside the infrastructure. Such approach increases the effort required as well as the time needed to compromise the infrastructure.

Some of these had arisen in software development but they had spread past it as they are concepts that easily apply to people and whole infrastructures; some others had been born from real experiences and research on securing systems that are getting more and widespread outside the old buildings' boundaries in an increasingly interconnected world.

There is one step we still have not inspected from our Cyber Security cycle as shown in figure 11: Monitoring & Results evaluation. This step is the “after action”: like sport matches after action are repeated after each and every match to understand what went well and what did not, this step should be repeated too again and again. We do not have matches and the cycle is a continuous one, the reader might be objecting: this “after action” step should be set that happens on a condition trigger.

This condition trigger should trigger when one of two specific reason evaluates to true:

- Enough time has passed since last run and enough data had been gathered through monitoring and outside inputs to warrant a new run of this step. The exact timespan should be on the organization to decide depending on the subject but I would advise between every few month/a year to at most 2 or 3 years. The longer the wait, the more the cyber security landscape can change and too short might not be relevant and might be too taxing on the organization.
- An event happened:
 - This event can be a loss event where the organization had a loss and this step should be used to try understanding the why, the what and who before feeding the results to a new planning phase that will address them.
 - This event can be an organizational or structural change: the organization or its structure changes considerably so we evaluate what has happened till now and we trigger a new planning phase that will address what will happen from now.

This step is a step that goes all over the Building blocks we talked about previously and the goal is to increase the knowledge pool of the organization. This step will make use of many different tools to reach its goals:

- Real time monitoring for software and hardware solutions
- Questionnaires and tests for the employees
- Theoretical and empirical evaluations such as cyber risk assessments
- Technical testing such as penetration testing

All the raw knowledge gathered with these methods should be structured into a coherent set to be used in a new Needs and Constraint Evaluation & Planning step. This way, we will be starting a new iteration of the cycle with new knowledge to added to the one we had at the end of the previous iteration: incrementally, as we reiterate the cycle, we will be incrementally building a knowledge base that, in the long run, will give us long term insights on how much the process has progressed. Similarly, the end of each iteration brings with it the short term insights on how the process is progressing in a limited time frame. Both insight’s types are valuable to an organization.

The reader might be wondering, whatever I will propose something tangible, something technical, opposed to a theoretical framework like the one I described till now. Well, the problem with proposing technical solutions is that their choice is tightly linked to the needs and constraint of the museum and/or the organization under exam. What I can do is writing down a few technical but still general points that can be relevant to the cyber security of museums but I cannot propose a complete technical solution to securing a museum.

Backups. Broadly defined as “Something that you have arranged in case of failure of something else”, when talking about computers a backup is about “*a copy of information held on a computer that is stored separately from the computer*” [Cambridge University Press 2021]. Backups had become increasingly important as a security against sudden failures of the main IT infrastructure. Museum should take the care to prepare and keep updated backups of their files and programs.

Multi-factor authentication. Multi-factor authentication is “an electronic authentication method in which a user is granted access to a resource only after successfully presenting two or more pieces of evidence (or factors) to an authentication mechanism: knowledge (something only the user knows), possession (something only the user has), and inherence (something only the user is) [Wikipedia 2009][Ping Identity 2021]. Multi-factor authentication had arisen to preeminence as a way to make electronic authentication more secure and less

relying on just alpha numerical passwords [NIST 2016]. An example of multi-factor authentication is the use of an alphanumerical password in tandem with a onetime code.

New advances in multi-factor authentication are studying the use of biometric data which are extremely difficult, if not impossible, to be recreated artificially as well as location data together with well-known methods such as passwords and tokens. Whatever the advance features, the fact is that multi factor authentication works and it is here to stay. Museums should work to enable this kind of protection when accessing their resources, both physical and IT.

The cloud. The cloud has gained traction in the last few years and it is set to dominate the market of IT services: with cheap storage, immense scalability, resource elasticity and accessibility from anywhere in the world, it made easier to do business. Museums can benefit from it the same way businesses and other organizations do but it requires additional knowledge to manage it and secure it: everything is not in the same place anymore but can be spread all over the world, past the physical boundaries of physical locations. The push for the cloud had been aided by the Covid-19 pandemic of 2020 which pushed remote work as a new way to work. It is clear how, from a security point of view, such “remoteness” makes the management of security much more complex.

6. Conclusion

We are at the end of this work, plenty had been said already but I would like to spend a few more words for a conclusion and to recap.

In this work, after I presented the results of my work with MUSE – Museo delle Scienze, I tried to draft a structured approach for doing cyber security inside museums but, in truth, every organization could benefit by a similar approach when wanting to build cyber security. Museums have specific needs and constraints, different from every other type of organization, but, at the core, many cyber security principles apply to them the same way they apply to many other organizations.

I believe that a structured approach like the one I presented with the Cyber Security Cycle and the Building Blocks can help in making it easier to explain to sceptical and not technical people as well as in organizing and tracking the cyber security work inside the organization. The Cyber Security Cycle is a time-based approach, you know you have a phase, following another and so on, so it helps in the long-term organization of work. The Building Blocks is an object-based approach, you have single objects to work with, parts that covers a specific topic that help you focusing the work to be done at some point of time and space. As I mentioned before, it is a dual natured approach to cyber security.

On the state of museums in regard cyber security, I believe the landscape is variegated: like with businesses, for museums that take cyber security seriously there are some that are not, believing in the “It will not be me” thinking. As for businesses, I believe a mental shift is needed: with business the mentality is shifting because cyber security is becoming a business priority. Cyber security let a business avoid damaging breaks of production and operability, while for museums, which are often part of public administration, this shift might be dragging a bit behind.

As with many aspects of business and life, I believe people to be important, I believe in proper education more than in endless and mindless enforcement of rules that are not completely understood. Educating people properly will go a long way in helping a museum, much like any organization, be cyber secure.

In conclusion, I believe and hope cyber security will gain traction becoming a normal part of museums’, much like any organizations’, operations and day to day life.

7. Acknowledgements

I would first like to thank my internship and thesis supervisor at the University of Trento, DISI, professor Paolo Giorgini for the support in this endeavour that started in the Spring of 2021. His feedback had always been clear and on point as I went through the writing of this thesis.

I would then like to thank the MUSE – Museo delle Scienze as a whole for accepting me and being part of the project I had in mind. Specifically, I would like to thank Vittorio Cozzio, head of IT services and security, for the mentoring, their always ready to help attitude and being always mindful of my goals and plans; then, I would like to thank Luca Zucchelli, from IT services, for their support and insights as we worked through the project.

I would like to thank my thesis supervisor at the University of Twente, dr. F.W. Hahn, for the feedback and the support from the Netherlands.

Finally, last but not least, I would like to thank my mother, Margherita, for being the number one in pretty much everything.

References

Listing of all the references found throughout the thesis.

- AGID Agenzia per l'Italia Digitale. 2017. "Misure minime di sicurezza ICT per le pubbliche amministrazioni." *AGID*, April 18, 2017. <https://www.agid.gov.it/it/sicurezza/misure-minime-sicurezza-ict>
- Callahan, Dan. 2021. "The first line of defense: Why employees are the key to stronger cybersecurity" *securitymagazine.com*, January 21, 2021. <https://www.securitymagazine.com/articles/94027-the-first-line-of-defense-why-employees-are-the-key-to-stronger-cybersecurity>
- Cambridge University Press. 2021. "Back-up.", *Cambridge Dictionary*, 2021. <https://dictionary.cambridge.org/it/dizionario/inglese/backup>
- Cascone, Sarah. 2019. "Hackers Saw the Asian Art Museum of San Francisco as Ripe for a Ransom Attack. Are Other Cultural Institutions Next?" *Artnet*, July 22, 2019. <https://news.artnet.com/market/hackers-attack-asian-art-museum-san-francisco-1604188>
- Cascone, Sarah. 2020. "Hackers Have Stolen Private Information From Donor Lists to 200 Institutions, Including the Smithsonian and the UK's National Trust." *Artnet*, September 2, 2020. <https://news.artnet.com/art-world/hackers-hit-smithsonian-parrish-corning-1905256>
- CIS Center for Internet Security. 2019. *CIS Controls v. 7.1*. CIS Center for Internet Security.
- CISA. 2021. "CISA portal" *CISA website*, 2021. https://us-cert.cisa.gov/ics/Control_System_Security_DMZ-Definition.html
- CVE. 1999. "CVE portal" *cve.mitre.org website*, 2021. <http://cve.mitre.org/>
- Enisa. 2021. "Phishing most common Cyber Incident faced by SMEs." *Enisa portal*, June 28, 2021. <https://www.enisa.europa.eu/news/enisa-news/phishing-most-common-cyber-incidents-faced-by-smes>
- FAIR Institute. 2021. "What is FAIR." *FAIR Institute*, 2021. <https://www.fairinstitute.org/what-is-fair>
- Freund, Jack and Jack Jones. 2015. *Measuring and Managing Information Risk - A FAIR Approach*. Elsevier Inc.
- ESnet. 2021. "Science DMZ." *ESnet website*, 2021. <https://fasterdata.es.net/science-dmz/>
- GARR. 2021. GARR website, 2021. <https://www.garr.it/it/>
- Google. 2014. "Google's Project Zero", *Google Project Zero blog*, September 18, 2021. <https://googleprojectzero.blogspot.com/>
- Google. 2018. "A secure web is here to stay." *Google Security Blog*, February 8, 2018. <https://security.googleblog.com/2018/02/a-secure-web-is-here-to-stay.html>
- Manzuik, Steve. 2020. "Red Team Attack Operation RT-011 - Phishing - Fake Laptop Upgrade." *GitLab*, May 20, 2020. <https://gitlab.com/gitlab-com/gl-security/security-operations/gl-redteam/red-team-tech-notes/-/tree/master/RT-011%20-%20Phishing%20Campaign>
- Microsoft. 2017. "CVE-2017-0145." <http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2017-0145>
- Miller, Colin. 2020. "Hackers of the Hof van Twente municipality demand 750,000 euros in ransom." *Netherlands News Live*, December 7, 2020. <https://netherlandsnewslive.com/hackers-of-the-hof-van-twente-municipality-demand-750000-euros-in-ransom/30117/>
- Misner, Phillip. 2017. "Customer Guidance for WannaCrypt attacks." *Microsoft Security Response Center*, May 12, 2017. <https://msrc-blog.microsoft.com/2017/05/12/customer-guidance-for-wannacrypt-attacks/>
- MUSE. 2021. "Mission" *muse.it*, 2021. <https://www.muse.it/it/il-muse/Mission/Pagine/Home.aspx>
- MUSE. 2021. "Edificio" *muse.it*, 2021. <https://www.muse.it/it/il-muse/Edificio/Pagine/Home.aspx>
- NIST. 2016. "Back to basics: Multi-factor authentication (MFA)." *NIST*, April 19, 2021. <https://www.nist.gov/itl/applied-cybersecurity/tig/back-basics-multi-factor-authentication>
- Oxford Dictionary. 2021. Definition of "Museum" as provided to Google by the Oxford Dictionary, September 15, 2021. <https://www.google.com/search?client=firefox-b-d&q=definition+of+museum>
- OWASP. 2021. "OWASP Top Ten", *owasp.org website*, 2021. <https://owasp.org/www-project-top-ten/>
- Ping Identity. 2021. "Multi-factor authentication: Best practices for securing the modern digital enterprise." Ping Identity, 2021. <https://www.pingidentity.com/en/resources/client-library/white-papers/2016/mfa-best-practices-modern-digital-enterprise-3001.html>
- Warren, Tom. 2017. "Microsoft issues 'highly unusual' Windows XP patch to prevent massive ransomware attack" *The Verge*, May 13, 2017. <https://www.theverge.com/2017/5/13/15635006/microsoft-windows-xp-security-patch-wannacry-ransomware-attack>

Wikipedia. 2009. "Multi-factor authentication." Last modified September 5, 2021.
https://en.wikipedia.org/wiki/Multi-factor_authentication
Wikipedia. 2017. "WannaCry ransomware attack." Last modified September 11, 2021.
https://en.wikipedia.org/wiki/WannaCry_ransomware_attack