

Central Lancashire Online Knowledge (CLoK)

Title	Cyber threat modeling for protecting the crown jewels in the Financial Services Sector (FSS)
Type	Article
URL	https://clock.uclan.ac.uk/43356/
DOI	https://doi.org/10.1080/19393555.2022.2104766
Date	2022
Citation	Alevizos, Charalampos and Stavrou, Eliana (2022) Cyber threat modeling for protecting the crown jewels in the Financial Services Sector (FSS). Information Security Journal: A Global Perspective, 32 (2). pp. 134-161. ISSN 1939-3555
Creators	Alevizos, Charalampos and Stavrou, Eliana

It is advisable to refer to the publisher's version if you intend to cite from the work.
<https://doi.org/10.1080/19393555.2022.2104766>

For information about Research at UCLan please go to <http://www.uclan.ac.uk/research/>

All outputs in CLoK are protected by Intellectual Property Rights law, including Copyright law. Copyright, IPR and Moral Rights for the works on this site are retained by the individual authors and/or other copyright owners. Terms and conditions for use of this material are defined in the <http://clock.uclan.ac.uk/policies/>

Cyber Threat Modeling for Protecting the Crown Jewels in the Financial Services Sector (FSS)

Lampis Alevizos¹, Eliana Stavrou²

Applied Cyber Security Research Laboratory

University of Central Lancashire Cyprus

12 -14 University Avenue, Pyla

7080 Larnaka, Cyprus

{ 1- (lampis@redisni.org, calevizos@uclan.ac.uk), 2-estavrou@uclan.ac.uk }

Abstract

Financial institutions are undergoing the so-called “de-perimeterization”. The security model up to today, is heavily dependent on “border patrols” focusing mostly on providing a secure perimeter while the internal network is inherently trusted. In the upcoming borderless networks, the focus is shifting to protection of the data itself, considering the full lifecycle or switching towards context aware defensive strategies also known as zero trust networks.

The focus of this work is to critically discuss existing threat modelling methodologies, available and used in the financial services sector (FSS). The objective is to investigate the extent at which existing methodologies cover the different threat actors & events and if they reflect the current threat landscape in the FSS. The investigations are supported by a real-world case study to uncover if any process can reflect the current threat landscape without any customizations or special know-how, and whether the final outcome helps in reaching a secure or compliance state. Through the case study, it is evidenced that by utilising the IRAM2 methodology resulted in a high ratio of compliance, however, considering the Crown Jewels of a Financial Institution (FI), a secure, as much as possible, state should be the desired outcome.

Keywords

Threat profiling, threat modelling, de-perimeterization, threat landscape, attack surface, crown jewels

1. Introduction

Cyber threat modelling nowadays is considered as a recognized discipline in driving the development of end-to-end IT security policies. Threat modelling helps in identifying, prioritising and proactively preventing threats by sinking the likelihood of probable cyber breaches. The threat landscape is changing rapidly, having attackers searching and discovering for new methods to bypass security defences, therefore, continuous improvement is a vital part of threat modelling posing a strong weapon in the arsenal of defenders. When a Financial Institution (FI) decides to add a new device, system, or even better an IT or OT asset, its attack surface is immediately increased, therefore it can influence a FI's security governance design. That said, new threats will automatically arise causing IT professionals, security professionals and executive leaders to become increasingly concerned regarding the safety and security of their data and assets (ThreatModeler, 2016). When it comes to defining the abovementioned "new threats", internal and external threats are the most common categories to specify, however, this is not sufficient when it comes to mission critical assets.

For centuries FIs have been utilising all kinds of assets e.g., assets participating in financial transactions or providing the basis to support such. Considering the ecosystem these assets are operating and the actual low-level functions they are performing, one can conclude that they are inherently vulnerable against theft and fraud. It comes of high significance that FIs should emphasise on such assets that are of the utmost value and risk, frequently indicated by decision makers and business leaders as "Crown Jewels" (CJ). Old-style risk assessment methodologies and threat modelling tactics that classify assets based on a simple classification chart e.g., negligible, low, medium, and high, would not ascertain CJ data distinctly, therefore mission-critical assets characteristically need a diverse context towards proper identification (Holt, 2016). In addition, it is imperative to properly categorise and differentiate between critical and mission-critical assets, as highlighted in Figure 1.

[Figure 1- Asset's classification]

This context is essential to be considered as CJ within banks, are subject to a broader threat landscape, heavily influenced by the threat actor's motivation, determination, funding, tools, skill level etc., all of which are vastly increased due to the nature & value of the mission critical asset being targeted. Moreover, financial services are dominated by a compliance-driven

notion, which frequently results in controls-first mindset (Michael Muckin, 2019). As such, the security architecture is driven by a known set of security controls. Compliance against a set of controls, though mandated via central banks or any other industry authority in general, does not assure a secure system mostly due to the way of measuring controls' effectiveness, often treated as a binary condition. As a consequence, CJs within FIs require comprehensive threat mapping to identify the relevant threat landscape. The output of such methodology should serve as a valuable input in risk management, business continuity & disaster recovery.

This research work investigates at which extend existing threat modelling methodologies reflect the current threat landscape in the context of FIs and identifies the most effective method to map all possible threats against a CJ through a case study analysis. Furthermore, through the case study analysis, it is investigated whether compliance with a set of controls, though mandated by multiple banking authorities, provides assurance for a secure system, or if the compliant state is propagating a false sense of security.

2. Background Work and Literature Review

There are already quite a few researches performed, either using literature reviews, case studies, comparison or a combination of all the above against different threat modelling frameworks, methodologies and tools. Special attention was given to MITRE's research paper studying a system-of-systems threat model (Deborah J. Bodeau, 2018) due to the fact that it is discussing a very specialised threat modelling approach, accurately outlining today's complex mission critical systems. Equally important is the work done from NGCI Cyber Apex Project (HSSEDI, 2019) publishing a suite of documents studying cyber threat modelling, cyber war gaming, advanced cyber risk management, cyber risk metrics survey & an augmented cyber threat model for FSS institutions. Michael Muckin & Scott C. Fitch from Lockheed Martin Corporation examined Microsoft's STRIDE in-depth and managed to enhance it by adding the lateral movement (LM) concept (Michael Muckin, 2019). The CEO of VerSprite and author of multiple books, Tony UcedaVelez has contributed many studies and comparisons against multiple threat models, while he is one of the original authors of the PASTA threat model (UcedaVélez & Morana, 2015). Adam Shostack, Microsoft's threat modelling expert has a vast contribution into cyber threat modelling as well. In one of his books named "threat modelling designing for security" a huge amount of STRIDE details can be found along with innovative ideas and explanations concerning the threat modelling as of today. Finally, the European Union Agency for Network and Information Security (ENISA) has been publishing threat

landscape reports for the last couple of years, highlighting the developments done by the defenders and the latest significant changes brought by the cyber criminals and state sponsored actors against mission critical assets.

2.1 Threat Modelling Approaches

Four approaches to threat modelling arise, software-centric, asset-centric, attacker-centric, data-centric & threat-centric. Based on this categorization, multiple methodologies and frameworks have been developed. First-rate modelling methodologies are very comprehensive while others have a higher conceptual level and some will just concentrate on a specific domain, although with far superior granularity (Shevchenko, 2018). Threat modelling methodologies are not static and cannot be considered as isolated individual frameworks. Threat methodologies and tools along with threat taxonomies and even risk management processes, can and should be integrated, in order to create a customised implementation for the special needs of a financial institution, that will ultimately provide security assurance for CJs (ThreatModeler, 2016).

2.2 Existing Threat Modelling Methodologies

Within the context of information security, threat modelling seeks to identify, understand and communicate threat information to security decision makers. Threat modelling can be used to secure system networks, applications, mobile, web, Internet of Things (IoT) embedded devices and more. In this work, we will investigate how existing methodologies and tools can be utilised in order to create an effective threat model for the highly complex and demanding environment of a financial institution. Threat modelling is a practice that shifts security as far to the left as possible in reference to the three lines of defence concept 3LoD (Veltos, 2017). Threat modelling can occur as early as the planning stages within the software development life cycle (SDLC), however, it can also occur when the product/asset is already operational. This work will examine the reactive or adversarial threat modelling approach, since the Crown Jewels or mission critical assets are already in operational state, meaning they are already in production.

2.2.1 NIST Special publication 800-154

NIST, which stands for National Institute of Standards and Technology, in March 2016 published a data centric system threat modelling guide. There are four core steps presented in the publication from Murugiah Souppaya & Karen Scarfone, however, the aforementioned document is not anticipated to be a novel methodology rather than an outline to data centric system threat modelling.

During the first step, system & data need to be acknowledged and categorised. This identification procedure may be fairly limited, meaning that it should only include precise information of a limited group with strictly associated systems or a specific stand-alone host. Categorization within this method, can be defined as the procedure where a system's function & overall utilisation are exhaustively comprehended (Scarfone, 2016).

In continuation, one must identify and select all the potential threat actors that might adversely distress the recognized security goals for the approved data locations. For the overall success of this model, it is of high importance and highly suggested to incorporate all the attack vectors to the model despite the fact that too many resources might be required (Scarfone, 2016).

The third step is related to mitigating strategies and procedures. This is a repeatable process where security controls mitigating the associated risk for every attack vector, and of course are realistic to achieve, are ultimately acknowledged and acknowledged. (Scarfone, 2016).

In the fourth and last step, the threat model is analysed. During analysis we need to verify the individualities and documentation produced from earlier stages and compare all the features composed. Such will lead to proper governance in what ways risks can be lowered crosswise all threat events.

Through this high level four-step process outline, it can be concluded that the NIST special publication 800-154 describes a principally qualitative approach, however, a quantitative approach would produce far more detailed and precise outcomes since it would be based on specific & verifiable data to predict the probability of a risk event outcome. Such would require much more resources and it would not be scalable enough to support the large and complex systems of a financial institution without massive automation of metrics gathering and analysis. On the contrary, the qualitative approach described has other significant benefits (Shuttleworth, 2017).

2.2.2 STRIDE

STRIDE is an acronym, deriving from the six threat elements as per the six letters of the word. Meaning that each letter stands for a “threat category”, in particular, Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service and Elevation of Privilege. STRIDE was first introduced by Loren Kohnfelder and Praerit Garg (Kohnfelder, 1999). The STRIDE framework will ultimately help identifying potential vulnerabilities in any product during a security analysis (Kohnfelder, 1999).

STRIDE utilises a proactive approach to threat modelling, and it is considered to be one of the most mature in that field, henceforth it is profoundly focused on software development. That being said, STRIDE is inherently missing a very important attacking & spreading technique through trusted systems/zones, the lateral movement (Michael Muckin, 2019). Figure 2 - STRIDE Elements, describes the different elements of STRIDE along with a short explanation and the property being violated.

[Figure 2 - STRIDE Elements]

It is of high importance to understand the ultimate goal of STRIDE, it is not meant to simply categorise threats rather than to help find all possible and applicable attacks. Using STRIDE to categorise threats might only be helpful to figure out the right defences and map them back to security controls.

Before deep diving into STRIDE mechanics and start considering threats, one should try to form reasonable questions. For example:

- How could an attacker possibly alter the authentication data?
- What would be the impact, if the adversary could read a user’s profile data?
- What could possibly happen if access is denied to a user’s profile e.g., in a database?
- Can an untrusted user modify data e.g., credit card data, in the database?
- Could somebody deny valid user service from the application?
- Could somebody benefit from a particular system component to elevate their privileges to that of an administrator?
- Can a non authorized user browse the confidential network data?

Of course, the above example questions do not serve as an exhaustive list of questions one might ask, though it provides an indication of the many questions and scenarios that might arise (LeBlanc, 2002).

It is becoming obvious that when architects, security engineers, developments engineers, product owners and all relevant stakeholders modelling e.g., an application utilising STRIDE, each of the above threat elements will need thorough examination and brainstorming in order to figure out all applicable threat scenarios.

2.2.3 *STRIDE & DREAD*

Another acronym initially shaped by Microsoft is DREAD. Similarly, to STRIDE, from each letter derives the relevant category, in detail Discoverability, Reproducibility, Exploitability, Affected users and Damage potential. Starting with discoverability, it is defined how effortlessly a vulnerability can be detected on the subject system/application/environment. Reproducibility assists in determining whether a previously identified attack can be successfully repeated. Exploitability defines the level of ease while exploiting a known vulnerability along with the resources and expertise required. The activity of affected users is trying to predict the impact on diverse kinds of users consuming the subject system or application. Lastly, damage potential assesses the result and the impact of successfully exploiting a vulnerability (UcedaVélez & Morana, 2015). DREAD was created as a threat prioritisation tool initially. However, it serves as a structured scoring system to prioritise and assess threats as identified by STRIDE by assigning every threat a scoring value of the likelihood of occurrence throughout the above mentioned five classes.

An inherent problem of DREAD is that each of the five categories are required to be scored on scale from zero to ten, making the process too subjective, which consequently leads to unusual outcomes in several occurrences and overlooking many risk factors. As a quick example one might think of results produced in multiple teams in large financial institutions or banks not being consistent, which is highly likely. As a result, it is very probable that identified/scored threats will be merged into one backlog item, while the lack of consistency will eventually introduce distortions when it comes to the overall threat prioritisation (Michlin, 2016).

2.2.4 *STRIDE-LM & IDDIL / ATC*

As already discussed during STRIDE model literature review, it is a proactive model heavily focused on software engineering and development. Therefore, an important threat element, that of lateral movement, is inherently not included. Lateral movement can be considered as a system-of-systems (McCollum, <https://www.mitre.org/>, 2018) type of threat. An exceptional paper for system-of-systems type of threats was written by Deborah J. Bodeau & Catherine D. McCollum for the department of homeland security which is operated by the MITRE corporation (McCollum, Mitre.org, 2018). Michael Muckin & Scott C. Fitch, engineers from Lockheed Martin, came up with an enhancement on STRIDE as well as a model named IDDIL / ATC (Michael Muckin, 2019).

The above-mentioned enhancement, lateral movement (LM) to STRIDE, trails the pattern of designating the effect of a threat, which is of exceptional importance considering nowadays threat landscape. For example, a common threat against the very well-guarded perimeter of a bank or financial institution would be to launch spear phishing / phishing campaigns against C-Level executives, secretaries or even regular employees in order to establish a foothold into the internal network, circumventing traditional defences. In continuation attempts to harvest credentials from compromised machines and spread laterally are made, such as accessing / abusing domain controllers and card processing fragments or machines. Authors of STRIDE-LM & IDDIL / ATC (Michael Muckin, 2019) conclusively refined the STRIDE table as shown in Figure 3 - STRIDE Table.

[Figure 3 - STRIDE Table]

We can immediately witness the addition of controls, although only the default, to the STRIDE table, and then the standard definitions & security property. The addition of controls and more specifically, the direct reference from the threat categorization to security controls is a fundamental notion of a threat driven approach (Michael Muckin, 2019), therefore utilising STRIDE-LM is highly recommended as a proactive approach within a financial institution or a bank.

IDDIL / ATC derives from the phrase “There are no idle (IDDIL) threats – they attack (ATC)”, a methodology created by Lockheed Martin, to help memorise the acronym. IDDIL activities are focused on discovery, while ATC activities are focused on the Implementation phase. Throughout the discovery phase, possible threats, attacks, attackers and the applicable assets

are populated. Figure 4 depicts a threat model of a web application incorporating IDDIL elements (Michael Muckin, 2019).

Ultimately, IDDIL / ATC delivers an organised and detailed procedure with the basic principles from Lockheed Martin's cyber kill chain and stresses or even urges the use of threat intelligence. The procedure utilises attack trees and a variation of STRIDE, as discussed above, named STRIDE-LM, which contains the attacker's tactic recognized as "lateral movement" to the original STRIDE model. The final outcome is a summary of threats, threat events and associated characteristics, drafted in a threat profile, which is shaped accordingly to interconnect the outcomes of the final threat model (Selin, 2019).

[Figure 4 - Web application threat model incorporating IDDIL elements]

2.2.5 PASTA

Process for Attack Simulation & Threat Analysis, PASTA, developed by Tony UcedaVélez in 2012, is a risk/threat-centric modelling methodology. It stipulates a seven-stage process (UcedaVélez & Morana, 2015) for associating business goals and technical necessities, considering compliance concerns along with business analysis. Ultimately the goal of this methodology would be to deliver a proper framework for risk mitigation, grounded on practical threat patterns alongside multiple varieties of erratic & erudite threats & motivations.

The inherent risk profile is being calculated during the first stage, while at the same time business impact considerations should be addressed as quickly as possible in the second phase. The thinking behind these stages is better described and understood by the motto "You cannot protect what you do not know". The third phase is focused on thoroughly comprehending the information flows between application or system components and the associated services. In essence, enumeration of all use cases should be performed, in order to capture the related functions of the subject system/service/application in sufficient detail. The output of the enumeration is then used to build the data flow diagrams (DFDs) (Visual-Paradigm.com, 2019) that are easily consumed by the relevant teams including all the essential functions and interactions. A sample data flow diagram for wire transfers of a financial institution is depicted in Figure 5 (Delzer, 2018).

[Figure 5 - DFD sample of wire transfers of a FI]

During threat analysis, the fourth stage of PASTA, we need to review the threat statements from the related data throughout the ecosystem while also taking into account FI's industry related threat intelligence that are applicable to the specific service and deployment model. The high-level threat scenario is drafted and analysed by ascertaining probable threat scenarios previously targeting applications and/or systems with analogous architectural design or technology involved.

Stage five detects the vulnerabilities against an application design and/or code and associates them to check the threat assertions from the prior stage four are supported. During this stage, one must identify and analyse possible weaknesses and vulnerabilities throughout the system's/application environment through specified activities.

The sixth phase focuses on imitating threat events (actual attacks) that could potentially exploit acknowledged vulnerabilities from the preceding phases. It is also very helpful in determining the threat viability via attack patterns. Core objective within the sixth stage of PASTA is to compliment or reinforce the attack trees with attack modelling and simulation.

Lastly, the residual risk calculation & analysis, seventh phase, emphasizes vulnerability remediation in actual code or system/application design that could enable threat scenarios and underlying attack patterns.

2.2.6 OCTAVE Allegro

OCTAVE is yet another acronym, meaning Operationally Critical Threat, Asset, and Vulnerability Evaluation. It was published in 1999 by Carnegie Mellon Software Engineering Institute that introduced OCTAVE Framework v1.0. Since then, four versions have been released during 2001-2005 presenting OCTAVE Framework v2.0, Criteria V2.0, OCTAVE-S v0.9 & OCTAVE-S v1.0. In June 2007 the last refinement of OCTAVE Allegro v1.0 was introduced and that is subject to discussion in this paper (Eric Lachapelle, 2015). The primary goal of OCTAVE is to develop an improved, rationalised, optimised and structured procedure to assess information security associated risks by investing only a limited number of resources, such as people and time. OCTAVE approach is applied through three phases.

Applying the original OCTAVE starts with identification of the financial institution's most valuable information-related assets, asset prioritisation, identification of most critical assets and documentation of their security requirements. Then an asset-based threat profile is created and all identified threats interfering with the security requirements are identified. The second

phase starts with the evaluation of the subject infrastructure and identification of all possible vulnerabilities deriving from the preceding threat profile(s). During the third and final stage, a security tactic is being developed by the team and then plans need to be drafted for risk mitigation applicable to critical assets. OCTAVE is mainly aimed at big organisations, e.g., 300+ employees with self-maintained IT ecosystems and multi-layered hierarchy. However, as described above, variations such as OCTAVE-S exist for small organisations, hence out of scope for this paper.

OCTAVE Allegro is an information-centric risk centred tactical assessment and planning methodology, aiming to allow more vigorous risk assessment outcomes with fewer information regarding extensive risk assessment. Its primary focus is on IT assets, though most prominently providing an answer on how and where those assets are stored & utilised, transported & handled but also how and where those same assets are wide-open to threat and vulnerabilities. OCTAVE Allegro, which is the refined method of OCTAVE, has eight steps in four stages, depicted in Figure 6 (Brunschwiler, 2013).

[Figure 6 - OCTAVE stages & steps]

During stage one, “Establishing drivers”, certain qualitative events that shape risk measurement conditions dependable with business needs are established. Continuing in the second stage, “Profile assets”, two steps are necessary, the development of profiles of critical information assets in order to ascertain sharp borders neighbouring the assets, and then adequately classify asset-related security necessities. Third stage known as “Identify threats”, identifies likely circumstances that may impede the information assets within a FI’s boundaries are acknowledged in association with the ecosystem of the subject asset or in general wherever the asset operates. The ultimate goal is to successfully identify and document unsolicited conditions that ascend logically from brainstorming meetings, leading to a limitation when it comes to protecting a financial institution’s CJs, rather than identifying and documenting all possible threat scenarios. The last stage, fourth, comprises steps for classifying & analysing risks while also choosing mitigation tactics. The threat and risk landscape are then finalised with the impact against a financial institution of an identified threat being materialised while all probable effects are represented (Richard A. Caralli, 2007).

Conclusively, two limitations of OCTAVE Allegro have been identified. The first one refers to the operation and maintenance phases of the SDLC, in which OCTAVE Allegro is usually applied, eventually leading to the development and implementation of mitigation strategies for

assets currently in operation being enormously expensive (Richard A. Caralli, 2007). The second limitation is the absence of an integrated, or at least a proposed, threat taxonomy. Therefore, during Step 6 – risk identification, activity 1 is heavily dependent on professional judgement of the relevant stakeholders, their expertise & experience, which is insufficient when it comes to protecting a financial institution’s Crown Jewels.

2.2.7 TRIKE

TRIKE is an open-source threat modelling methodology which adopts a defensive standpoint plus it encourages the automation of tedious tasks during threat modelling and as a result provides more resources and focus to risk management activities. When applying TRIKE methodology, one must first define the subject environment and build a requirements model that will contain all assets, threat actors, threat events & rules of the environment in a AAA (Actor-Asset-Action) matrix. Next, every cell of the matrix is separated based on what is known as CRUD(XF) actions and stands for creating – reading – updating – deleting – executing – configuring (Zalewski, 2012).

Afterwards one must build the data flow diagrams (DFDs) in order to draft a blueprint implementation model that will be used in turn to create the real threat model. Each of the DFDs will have to be iterated through so as to identify all the related threats. Every unique threat that will be discovered will eventually represent a root node in an attack tree. At this point it needs to be noted that TRIKE includes merely two threat categories, one of them being elevation of privilege and second, denial of service.

During the last stage, the threat model will facilitate the creation of a risk model. Though the risk model is stated in the methodology, it has not been properly applied in the TRIKE tool, therefore this is one of the downsides of the methodology. Risk assessment is completed by utilising a scale with five points, which is based on the likelihood of the risk (Zalewski, 2012). Overall TRIKE is using a very analytical and risk-based approach; therefore, this is the main reason the methodology is well known and referenced in multiple sources, yet again it has to be mentioned that it appears to be no longer maintained.

2.2.8 VAST

VAST, another acronym, derives from Visual-Agile-Simple-Threat modelling and it was originally created by Anurag Agarwal. The methodology is the foundation of a commercial

threat modelling platform known as ThreatModeler which is based on heavy automation. VAST claims to be the only highly scalable and operational methodology up to this date due to the fact that its core idea is that threat modelling is useful if it unfolds the whole software development life cycle (SDLC) throughout the whole enterprise (ThreatModeler, 2016). That being said, VAST can be adopted by large financial institutions or banks with a highly complex IT/OT landscape in order to include both their SDLC and the full IT/OT infrastructure in a systematic way.

In practice, VAST methodology covers both the proactive and reactive approaches with two different deliverables. The proactive approach, named application threat modelling as per the ThreatModeler platform, provides developers insights to help them in prioritising and mitigating threats very early in the design phase of applications, though before rolling them out to production. Application threat modelling as per OWASP's threat dragon open-source tool is illustrated in Figure 7 (OWASP, OWASP Threat Dragon, 2020).

[Figure 7 - Application threat model example]

On the contrary, the operational threat model is focused on infrastructure level including but not limited to servers, databases, firewalls, switches, routers, load balancers and other relevant infrastructure components within a financial institution. The ultimate goal of the operational threat model is to bring together relevant teams and provide them with insights to mitigate the inherent risks (Optiv, 2011) in the infrastructure of the FI while at the same time enable them to align with business strategy goals and overall budgeting. An example illustration of an operational threat model, of an operational Citrix environment is shown in Figure 8.

[Figure 8 - Operational threat model example]

2.2.9 IRAM2 (ISF)

ISF, which stands for Information Security Forum is a well-known in the security industry independent & not-for-profit association of leading organisations from all over the world, which was founded in 1989. It serves as the world's leading authority on cyber, information security and risk management (ISF, 2019). ISF's guidance, practical tools & research material are commercially available for its members. From a pure risk assessment perspective, while

many other methodologies and frameworks will most likely finish at the stage of risk evaluation, IRAM2 covers a wider extent of the global risk management lifespan by delivering realistic direction on risk treatment. IRAM2 is the ISF's risk assessment methodology, which is laid out in six phases. Every phase features stages and core actions essential to accomplish preceding phase's goals, in addition to recognizing the crucial information risk aspects and outcomes. As this paper's main focus is on threat modelling on financial institutions, we will deep dive on how IRAM2's third phase "threat modelling" along with the relevant upper and subphases can be utilised during a case study later on.

The starting point for IRAM2 is to develop an environmental profile and define and agree the scope for assessment. This step helps the assessor and stakeholders to acquire a thorough understanding of the subject ecosystem. That said, deep comprehension of the subject business & technology and how they interconnect & interdependent within the context of a FI is required. Understanding gained from the previous profiling process provides the basis from which to discuss and agree the scope of assessment. An agreement between the key stakeholders should be drafted including the Business service or process(es) in scope for assessment, the technology service(s) supporting the business service or process(es) in scope for assessment and the parties responsible for assessing the various risk domains (e.g., business process and technology/information risk) for the environment.

Once the environmental profile is concluded and the scope of the assessment is decided, the subsequent phase is to classify information assets within the ecosystem and lastly assess the business impact. The first step in performing a business impact assessment (BIA) is to find every relevant information asset associated with the ecosystem being assessed (i.e., in-scope). Throughout this step, one should identify and document the complete information lifecycle for each information asset e.g., where it is generated/input, processed, stored/archived, transmitted and destroyed, as well as the associated business or technology component(s). This helps similar information assets to be grouped and provides the basis for aggregating business impacts. Figure 9 (ISF, 2019) is an example of using these activities to identify information assets and their associated lifecycles:

[Figure 9 - Asset identification activities example]

Once the information assets within the environment being assessed have been identified, the next step is to assess the business impact. The value of information assets to an organisation is

assessed by evaluating properties of the information assets known as information attributes. In IRAM2, business impact is measured using a business impact reference table (BIRT). It is important to note that IRAM2 differs from other methodologies in assessing business impact as it considers both realistic and worst-case inherent impacts. At the end of the Business Impact Assessment phase, one should have acquired a concrete understanding of the information assets in the subject ecosystem, along with their business impact ratings.

The next phase involves identification and prioritization of the associated threats to the subject ecosystem, and determination of how they could potentially cause harm to that ecosystem. As per IRAM2, a threat is anything that is capable by its action or inaction, of causing harm to an information asset. For each threat, IRAM2 defines specific characteristics known as threat attributes. These attributes are used by security architects, engineering teams and overall stakeholders to help model the behaviour of each threat, a process known within IRAM2 as threat profiling and referring to step two of phase one. IRAM2 groups threats at the highest level by the intent threat attribute, which results in the creation of three standard threat groups:

1. **Adversarial:** This group includes threats that perform one or more deliberate actions against an FI's information systems or assets, with the intention of causing harm.
2. **Accidental:** This group includes threats that, as a result of error or unintentional action (or inaction) cause harm to a FI's information systems or assets.
3. **Environmental:** This group includes threats which are outside the control of the financial institution, such as natural or man-made hazards, or failures of critical infrastructure, that act to cause harm to FI's information systems or assets.

This grouping of threats enables the relevant stakeholders to more readily compare and contrast similar threats as part of the threat profiling process. If one gets these fundamental concepts in place, it is now possible to proceed with the first step of this phase, populating the threat landscape. The first step in threat profiling contains decisive threat calculation pertinent to the subject ecosystem, therefore facilitating stakeholders to draw the threat landscape.

Next, if all threats have been profiled, one should create a prioritised threat landscape. This can be accomplished by for example sorting the threat landscape firstly by Likelihood of Initiation (LoI), and then by Threat Strength (TS). This list should then be numbered sequentially to provide the threat priority rating. The next immediate step in this phase is to identify the threat events associated with each threat.

At the end of the Threat Profiling phase, one should have acquired concrete comprehension of the threat actors applicable to the subject ecosystem, their associated threat events, and how they might distress the numerous assets (via components) in the ecosystem. A record of the prioritised threat landscape and agreements with the relevant stakeholders should be in place, along with the in-scope threat events and impacted assets.

Lastly and once all risks have been assessed and the residual risk has been calculated for every risk, the resulting phase is determining risk treatment for each identified risk, and the first step of this phase is to determine whether each identified risk exceeds the organisation's risk appetite (LogicManager, 2019). During this step, one should determine the risk category or categories to which each risk belongs, and then compare it to the organisation's appetite in each relevant risk category.

3.3 Threat taxonomies & threat intelligence

The complexity, persistence and capabilities of the adversaries in the current threat landscape result in a speed contest amongst threat actors, incident responders and security analysts. Hacking groups, individual hackers, APTs, criminals, activists/hacktivists, nation-state and other adversaries are increasingly and dangerously successful in attacking the Crown Jewels and accomplishing their objectives. Banks and financial institutions are adopting cyber threat intelligence (CTI) (Doerr, 2018) to address the rapid increase in adversarial threats and understand the current threat landscape. The assignment of collecting evidence-based knowledge considering context, mechanisms, indicators, repercussions and actionable insights, in response to prevailing or developing threats to assets that can be cast-off to enlighten verdicts in reference to the subject's reaction to that danger, is referred to as threat intelligence (Bromander, 2017). One of the core activities of CTI is to share an adversary's activities, therefore several taxonomies exist for maintaining a common lexicon within and between FI's, while it is of high importance and great value when these taxonomies are incorporated during threat modelling and risk assessment (Launius, 2018).

Threat taxonomies can improve threat modelling and risk assessments in two ways. First, the language barriers between technical, non-technical security & business audiences with different expertise can be broken down into clear definitions for IT/OT threats. Proper communication while modelling throughout the stakeholders will ultimately provide decision

makers with a clear understanding of active threats. Second, a well-structured, up-to-date threat taxonomy enables modelling, analysis and assessment at various granularities. In-depth analysis taking into consideration the current threat landscape and the thorough understanding of adversary behaviour will eventually lead to results beyond the compliant state usually achieved during risk assessments & threat modelling exercises, which is common within FI's, leading at least one step closer to secure state.

This work's scope is not an evaluation of threat taxonomies; however, it recognizes and makes use of during the case study, therefore within this segment a synopsis of taxonomies is laid out. MITRE's common weakness enumeration, also known as CWE, is a community developed list of common software security weaknesses. The core of CWE is to act as a common language for measuring software security tools and additionally as a baseline for weakness identification, mitigation & prevention efforts (MITRE, mitre.org, 2019). Associated with CWE, CAPEC, Common Attack Pattern Enumeration and Classification is heavily focused on application security describing the common attributes, tactics, techniques and procedures engaged by aggressors to exploit known vulnerabilities in cyber enabled capabilities such as SQL Injections, session hijacking / session fixation, XSS, clickjacking (MITRE, MITRE CAPEC, 2019), etc. Another taxonomy from MITRE, ATT&CK, Adversarial Tactics, Techniques & Common Knowledge is focused on network defence describing mostly the operational phases during and adversary's lifecycle pre & post exploitation e.g., persistence, lateral movement & data exfiltration. Additionally, it provides details on the specific tactics, techniques and procedures (TTPs) that APTs utilise to accomplish their goals whilst targeting, compromising and operating after breaking the perimeter (MITRE, MITRE ATT&CK, 2019).

In a notional scenario of a clustered internet banking application, it becomes evident that by utilising and combining MITRE's CAPEC & ATT&CK threat taxonomies, one can reflect the full spectrum of an adversary's "lifecycle". Meaning that web related TTPs would be applied through CAPEC, covering the internet facing part of the internet banking cluster, while ATT&CK, would apply TTPs as if an attacker had already bypassed the perimeter moving laterally towards the cluster.

TAL, Threat Agent Library, developed by Intel delivers a reliable up-to-date reference defining threat agent activities against IT/OT systems and other information assets (Casey, 2007). MITRE's Common Vulnerabilities and Exposures, CVE, is a dictionary that provides common identifiers for publicly known vulnerabilities mostly focused on software packages (MITRE,

MITRE CVE, 2019). There are more taxonomies, sharing standards such as OpenIOC (FireEye, FireEye, 2013), STIX (MITRE, STIXproject, 2018), MAEC (MITRE, MAECproject, 2019) and one of the most complete ontologies, the unified cybersecurity ontology, UCO (Zareen Syed, 2016).

4. Pros/cons summary and suitability table

Table 1 - Threat modelling methodologies summary & suitability summarises the pros and cons of existing threat modelling methodologies.

[Table 1 - Threat modelling methodologies summary & suitability]

It becomes evident that multiple methodologies or frameworks come along with their pros and cons, while the suitability varies. Another important differentiation factor is the threat domain coverage, which is something to take into consideration prior selecting a methodology or framework. Most methodologies and/or frameworks do not offer default threat catalogues. However, the methodologies/frameworks that do offer a default table of threats is generic and not directly applicable to the highly regulated and highly complex environments of FIs.

Another aspect that needs to be taken into consideration when choosing a threat modelling methodology is the threat actors and their motivations, which can provide insights as to the threats that FI might have to face, including potential impact if the threat is materialised. This knowledge can assist in understanding the malicious strategy taken by specific threat actors, and identifying their tactics, techniques and tools utilised. Having a better understanding of the threat actors, defenders can identify and prioritise activities to address high impact threats and minimise, even eliminate, potential impact on CJ. Myriam Dunn Cavelty (Dunn Cavelty, 2010) defines a fivefold classification model of threat actors based on their motivational factors. Namely, the first level of motivational factors comprises cyber vandalism which includes hacking and hacktivism threat actors. Their essence of motivation being egoism or anarchy, and therefore their actions can have high impact against an FI, that can lead to financial theft, business disruption or reputational damage. The second level consists of cybercrime with their motive being anarchy or money. Threat actors with such motivational factors will cause high impact reputational damage, financial theft & fraud and even regulatory issues. The third level consists of cyber espionage. Actors with these motivational factors in the context of FSS/FIs will try to have financial gains by either direct financial theft or fraud, or indirectly by obtaining sensitive information and trying to sell it onto the dark web. The fourth level consists of cyber terrorism, which aims to cause destruction, and therefore this threat actor is unlikely to target FSS/FIs. Fifth and last level is cyber warfare, which entails nation state actors. They seek primarily money and power hence a likely threat actor to consider against FSS/FIs. A summary of the threat actors, their motivation and impact are shown in Figure 10 (Deloitte, 2014)

[Figure 10 - Cyber threat actors and motivational factors against FSS/FIs]

It is evident that one of the main applicable threat actors to a FI, is the organised criminal group(s) or even advanced persistent threats (APTs), therefore it is of utmost importance to

include the latest tools, techniques, tactics and procedures followed by such actors. That said, it becomes imperative that a financial institution searching to implement a threat modelling methodology or framework, should invest heavily upfront in terms of time and resources and draft a detailed tailor-made threat event catalogue in order to achieve the secure state, rather than a compliant state. Moreover, the threat catalogue should be maintained and kept up to date by frequent reviews, taking into account FI's industry specific threat intelligence along with advanced and current threat taxonomies such as MITRE's ATT&CK (MITRE, MITRE ATT&CK, 2019).

5. Case study

5.1 Rationale & goals

The goal of this real case study is to answer the first research question (quoted text)

"at which extend do existing threat modelling methodologies cover all the different threat actors and threat events, and do they reflect the current threat landscape?"

by applying the specific part of threat modelling of IRAM2 framework. IRAM2 phases have been described in section [3.3.9](#). This is a widely accepted and applied risk assessment framework including a threat modelling methodology in phase three, which will be the focus and the main topic of this case study. The modelling approach was selected due to the fact that it is currently, continuously updated, adopted by the 400+ global member organisations and it is very well documented though some resources are available to members only, hence require subscription. The second research question (quoted text)

"does compliance with a set of controls, though mandated by multiple banking authorities, provides assurance for a secure system, application or environment, or the compliant state is propagating a false sense of security? "

will also be discussed, since it is a consequence and outcome of research question one, that will be highlighted through the first part of the case study.

5.2 Assumptions and limitations

The case study derives from a real-world threat modelling exercise against the online banking environment carried out within a known bank. It has to be made clear that details of the network, infrastructure and subject to model environment components have been obfuscated or simplified for obvious reasons while their confidentiality is protected accordingly. It is, however, representative of the architecture and the basic business functions of the online banking infrastructure subject to model, therefore it provides a reasonable basis for examining our two initial research questions.

The aspects of the bank relevant to the online banking environment are explained in some detail, hiding the bank's IT architecture, risk profile, business functions and cyber defence capabilities. Whilst the bank's intention is to have a vigorous and highly proficient set of cyber defence capabilities, it is not flawless. Again, this is a real-world case study, therefore it needs to be understood that many of the details of the bank's security configurations and defensive practices reflect verdicts concluded based on a risk-aware manner, thus inherently including trade-offs of cost, benefit and usability.

5.3 IT landscape overview

The bank's WAN entails numerous & diverse locations, sites and branches, spread geographically throughout the world, interconnected through a series of different types of leased lines and virtual connections via multiple telecommunications providers. The diverse kinds of sites/branches comprising the WAN comprise main & minor datacentres, major & minor offices along with the automated teller machines (ATMs) network, remote employees, vendor supplier's connections and cloud services can be seen in Figure 11.

[Figure 11 – Major data centre high level diagram]

The WAN has multiple forms of packet-blocking and packet/traffic inspection technologies in place; however, very little efforts have been made to properly segment and restrict the traffic flow from the internal and between the user endpoint farm towards the server's farm. However, the financial transactions processing Mainframe (IBM, 2017) is segregated and access is allowed only through specific servers. This is a conscious decision taken within the bank

primarily driven by a cost-benefit analysis that concluded the risk was not sufficient to warrant the capital or operational expenses.

In reference to external connectivity, there are four major interfaces that can be listed. The first one is the internet interface. Second, the dedicated/trusted third parties interface including B2B connections, links to European Central Bank (ECB) & other authorities. Third, the mergers & acquisitions interface consisted of partly trusted companies or subsidiaries in the process of being acquired. The fourth interface includes the daughter companies where there is a high level of trust established and basic security controls applicable. Every connection to and from an external interface is strictly controlled and a different set of combined security controls has been considered suitable for the identified risk. Internet connections are considered to be of high risk; therefore, they have the most security controls implemented.

5.4 Threat modelling

IRAM2 risk assessment methodology requires a threat profiling phase to take place, hence a structured approach to model threats is documented. The focus of this work is on the threat modelling methodologies, therefore we will focus on the relevant phase of IRAM2, Phase C, along with the related activities, as per the description in section 3.3.9 and depicted in Figure 12 (ISF, 2019):

[Figure 12 - IRAM2 Phase C activities]

The main goal of Phase C is to model and prioritise all possible threats applicable to the subject ecosystem & identify the potential ways that the highest priority threats could manifest to cause harm to the subject to model. The subject of the model is the basic payments application, part of the online-banking infrastructure which is a cluster of servers located throughout the bank's demilitarised zone (DMZ), internal trusted zone server farm, and the restricted mainframe network. The basic payments application is hosted on the bank's mainframe; therefore, all the above-described principles are applicable. All of the online-banking components are considered to be mission-critical (CJs), part of the WAN described above and the overall bank's IT landscape. A lightweight version of the architectural structure of the online-banking environment illustrating the mainframe zone along with the data flow towards the subject basic

payments application is shown in the data flow diagram (DFD) presented in Figure 13 (OWASP, 2016):

[Figure 13 - Online banking environment DFD & core components]

5.4.1 Populate the threat landscape

The methodology provides a common threat list (CTL) in order to assist in the selection process of the relevant threats for the environment being assessed. According to ISF, a threat is anything that is capable, by its action or inaction, of causing harm to an information asset (ISF, 2019) The default CTL consists of 32 threats across three threat groups (adversarial, accidental, environmental), and it has been utilised in the real-world case study, filtering out all the non-relevant threat actors considering the restricted zone (mainframe) the basic payment application resides. Moreover, the environmental threats although in-scope during the original assessment, during the outline of this case study they considered as out of scope in order to focus more on the research questions of this paper. Environmental threats do not impact the research questions, the overall implementation of the modelling approach or the conclusions of this paper in any case, hence considered out of scope. The threat landscape was drafted after multiple workshops including all relevant stakeholders supported by development and run/change teams, security architects, engineers and risk assessors. The outcome of the threat landscape population is illustrated in Figure 14 - Threat landscape population (ISF, 2019).

[Figure 14 - Threat landscape population]

Due to the fact that the basic payments application, is located within a restricted & trusted zone (mainframe) with no internet interfaces, along with the inherent origin distinction (internal/external), it was concluded that some threat actors are not applicable, or even more precise, their threat scenarios were highly unlikely to be materialised. The origin of a threat refers to whether the threat originates internally within the organisation or is external to the organisation. Determining whether the threat is internal or external is a vital aspect since it will determine, later on, the threat events that a threat is able to initiate. As a result, some threat actors have been greyed out (removed) as shown in the above Figure 14 (ISF, 2019). At this step, it is imperative to note the differentiation in threat actor's context and origin. Namely, the

main threat scenarios in scope are employees, either privileged or unprivileged while the threat group is internal only. This was a conscious decision, based on two reasons:

- 1) considering that mainframe is a highly sensitive banking ecosystem and has zero internet exposed interfaces, the stakeholders considered that it cannot be directly attacked through internet by a hacking group for instance, hence only an internal employee would / could attack mainframe
- 2) If any threat actor (e.g., nation-state, organised crime) is able to compromise the bank's perimeter via phishing or spear phishing attacks, then the same scenarios of an internal employee would be applicable. This is already in scope; hence nation-state, organised crime and all external threat actors were de-scoped.

The next steps are following the standard IRAM2 process, however, this step is crucial in the sense that it scopes or de-scopes threat actors, and therefore scopes the controls to be assessed. Although the above two reasons might seem logical, and they did seem logical to the specific group of stakeholders at that point of time, there needs to be a closer examination. Namely, the second reason (point 2 above) can and should be challenged for the following reason:

- a) If highly skilled threat actors such as nation-state or organised crime, was able to gain foothold onto the bank's network, they would not execute the same range of attacks or they would not follow the same pattern of attacks as an internal (privileged or unprivileged) employee, therefore it must be scoped separately. As a result, the threat profiling exercise automatically is subject to a great loophole, with the reason being the lack of a third context (threat group) applicable to the (privileged or unprivileged) employee, namely, the in-adversarial.

5.4.2 Profile threats

Next step is the profiling of threats. Again, the outcome of this step was a result, and it should be, of multiple workshops with all relevant stakeholder's participation supported by the relevant security architects, engineers, risk assessors and risk managers. Since the environmental threats were de-scoped, only for this case study, the remaining adversarial and accidental threats were assigned attributes such as history, capability, motivation, commitment, rationale etc., as shown in Figure 15 & Figure 16 (ISF, 2019) respectively. As of this step, we are following the standard process which is mandated to map the attributes of the threat actors, however, in the previous step we have already defined the origin, the source of the attack,

external or internal and we have already addressed the context, namely adversarial or accidental. Hence from this point onwards, there is no room for customizations as the process follows and is highly dependent on the previously selected context and threat actors. Our previous selections and scoping exercises refined and defined the controls for the assessment.

[Figure 15 - Adversarial threat attribution]

[Figure 16 - Accidental threats attribution]

5.4.3 Produce a prioritised threat landscape

The next step is to present the results of the threat profiling activities into a prioritised order. The profiled threats from previous adversarial & accidental contexts are transferred into a common list where they are prioritised based on likelihood of initiation (LoI) and then based on threat strength (TS). Part of the produced prioritised threat landscape of basic payments application is shown in Figure 17 (ISF, 2019). Note that the general employee threat actor, with an accidental intent, has the lowest priority. The threat strength is rated as high, while the likelihood scored low, as per IRAM's default scaling. This is the result of forward thinking of a general employee, trying to launch attacks against the internal subject application. Indeed, the strength would be high considering that the attack originated from the internal network, however, the likelihood is set to low because the general employee has inherent trust, as an internal employee. This is the point where we propose a customization in threat profiling exercises that use a default threat catalogue and context. More specifically, point (a) in page above, explains the main reason between a general or privileged employee and a nation-state threat actor having access already in the internal network. The difference in short is context and attack path. Conclusively for the next assessment:

- a) We should add a third threat actor, the internal employee privileged or unprivileged, because both might be targeted and compromised by a spear phishing attack for instance.
- b) And that employee should have a third context, on top of adversarial and accidental, which should be the in-adversarial.

As a result, that would mimic and scope properly the threat scenario along with the reflecting controls of a nation-state actor, compromising an employee, and launching all kinds of internal

attacks with the ultimate goal to spread laterally and compromise internally protected systems. That would greatly increase effectiveness of our assessment as one of the most highly likely scenarios would have correctly scoped, and in turn the applicable controls for such scenarios would have been assessed for operational effectiveness as well.

[Figure 17 - Prioritised threat landscape]

The default values of risk factor scoring tables for the accidental and adversarial contexts related to LoI & TS are illustrated in Figure 18 (ISF, 2019), though they can be modified in case they do not reflect the actual ratings.

[Figure 18 - LOI & TS factors scaling]

5.4.4 Scope and map the threat events

The next step is imperative, as the scope and mapping of the threat events occur. As per ISF, a threat event is an action (or lack thereof), initiated by a threat against an information asset, which is capable of causing harm (ISF, 2019). ISF provides a default threat event catalogue (TEC). The TEC is a list of threat events e.g., session hijacking or theft of information system hardware that can be initiated by a threat e.g., nation state, hacking group or a privileged employee. It needs to be mentioned that the list is editable allowing threat events to be modified, deleted, by deleting rows, and added, by adding a new row at the base of the TEC, in other words customizable. However, if a new threat event is added, one must make sure that it has assigned an origin and a minimum threat strength to initiate threat events. The TEC assigns a threat event ID, depicts the context (threat group adversarial/accidental, etc.), lists the threat event types along with the specific threat events. A threat event description with an assessment guidance is also provided, hence the default TEC is highly usable and consumable throughout different levels of stakeholders & technical and non-technical audiences. Last but not least, the threat origin requirement and the minimum threat strength required to initiate a threat event is set. Part of the produced TEC is shown in Figure 19 (ISF, 2019):

[Figure 19 - Subject threat event catalogue]

Finally, the scoped threat events provided the stakeholders with a mapping of the threat events that can be initiated by one or more threats. Below is the prioritised version mapping illustrated in Figure 20 (ISF, 2019). Since this is a real-world case study, some details have been purposely greyed out due to confidentiality reasons, however the overall notion is the same and does not affect the outcome and conclusions.

[Figure 20 - Threat events mapping]

At this stage an attack tree was also created though it is considered confidential data hence it cannot be incorporated within this case study. However, an example is illustrated in Figure 21 in order to provide the reader sufficient understanding (Edge, 2007). The attack tree was drafted after multiple workshops with the relevant stakeholders, having as a starting question “how can the subject ecosystem be attacked?”. As such, numerous abuse cases and potential attacks were identified, categorised and depicted in the mentioned Figure 21 (ISF, 2019). It can be highlighted that the same attack tree is later on used as the golden source input for use cases with regards to security monitoring, as abuse cases and use cases are highly interconnected in the sense that feedback is unilateral and important.

[Figure 21 - Attack tree example]

As a result, the mapped threat events enable for mapping the in-scoped threats and threat events to the relevant basic payment application components, as shown in Figure 22 (ISF, 2019). The whole mapping exercise was not easy, as multiple threat scenarios correspond to different and multiple threat events, which in turn connect back to numerous threat event types. However, the core logic is to always have the attack tree leading and acting as input, therefore everything else attaches to it and provides clarity as more information comes together.

[Figure 22 - In-scope threats & threat events mapping]

5.4.5 Identify and map assets impacted by threat events

Now that a mapping of the in-scope threats and threat events to the information assets and components is complete, the last step in Phase C (threat profiling), is to incorporate results from the business impact assessment. The Asset Threat Event Map combines key information from the previous steps. During this step, it is crucial to have the business impact assessment (BIA) (Gartner, 2019) as it is a pre-requirement for the correct, complete, and accurate mapping. For obvious reasons the bank's BIA cannot be shared, however, part of the produced output of this step is depicted in Figure 23 (ISF, 2019):

[Figure 23 - BIA sample output]

5.5 Case study inference

The first immediate inference serving as answer to the first research question (quoted text)

“at which extend do existing threat modelling methodologies cover all the different threat actors and threat events and do they reflect the current threat landscape?”

is that by utilising a default or generic threat event catalogue, in our case IRAM2 default threat event catalogue, we managed to identify a pool of applicable threat actors and threat events in high level categories capable of highlighting the pain points. However, in our case study and considering the nature of the business along with the asset status being a CJ, the current threat landscape is not depicted correctly since we missed two things. First, it is the inadvertent context, meaning that an internal employee (privileged or normal) could open up a remote direct/reverse communication channel to threat actors. Second, after that channel is established, lateral movement into the internal network towards the online banking and mainframe environments becomes possible. As already outlined in the introduction section [1.1](#), the threat landscape is changing rapidly, attackers searching and discovering for new methods to bypass security defences, therefore, continuous improvement is a vital part of threat modelling posing a strong weapon in the arsenal of defenders. As a continuous improvement of the default TECs provided by multiple threat models, MITRE's ATT&CK framework can be utilised to ensure that a full threat actor / threat event coverage will be achieved and will eventually reflect the current threat landscape.

After the completion of the threat profiling process, a mapping of all the identified and applicable threat scenarios back to the bank's security controls has been produced. The security controls are subject to policies, compliance e.g. frequent risk assessments, etc., therefore the outcome of this exercise will lead to the input of a risk assessment. The outcome of the risk assessment will eventually provide the C level executives, though not limited to, with oversight to risks, issues and observations supporting an overall risk aware decision for the security posture of the subject to model application, hence the second research question (quoted text) refers to a false sense of security.

“does compliance with a set of controls, though mandated by multiple banking authorities, provides assurance for a secure system, application or environment, or the compliant state is propagating a false sense of security? “

Consider that we implement every mandated control as per the policies, national/international authorities, regulations, etc., as well as our internal threat modelling exercise on the web banking application. However, the ecosystem that hosts the web banking was designed to include administrative and management interfaces that ultimately allow lateral movement throughout the internal network of the bank, if any of the assets are compromised. It becomes clear that the threat actors and threat events were not identified during the threat modelling, yet the environment deemed 100% compliant, however, it remains exceptionally vulnerable to high-impact targeted attacks. Appropriate utilisation of e.g., MITRE's ATT&CK, framework during the threat modelling workshops phase, may have acknowledged such matters and allocated the suitable number of controls, incorporating the compensating controls, so as to reach above and beyond compliance.

6. Key takeaways of threat modeling in FI

Performing threat modeling exercises, even by utilising the default threat event catalogues, as per the industry standards or provided/embedded by the subject methodology is far better than not performing threat modeling at all. The responsible team should be aware of the methodologies and emerging trends on the field; however, many methodologies share common steps and ideologies, which makes it easier to follow up. Based on the literature review and the observations from the real-world case study the following key takeaways can be drafted.

- Start by defining the depth level. Meaning that you need to define if the objective of your exercise is referring to a Crown Jewel or simply an asset within your organisation. By defining the depth, one is able to start identifying all possible threats considering the current threat landscape or simply identify the most likely and harmful threats.
- Conduct multiple workshops and make sure you are engaging all kinds of stakeholders.
- Gather, utilise and ultimately embed threat intelligence feeds into your threat profiling lifecycle.
- Draft a detailed enough blueprint of the trust borders, either technical, logical or human associated to the system subject to modeling.
- Always include the administrators, or privileged users in general, in your models and make sure the least privilege principle is in scope.
- When it comes to the context of a threat event, do not rely on adversarial and accidental split only, take into account the inadvertent context.
- Profile as much accurately as possible the most likely threat actors along with their motives and skills, always with the context of your subject to the model system in mind.
- Do not assume that part/s of the system to model is another's responsibility. One should always look at the big picture and test all possible layers of defence as the threat actors will most likely do.
- Draft a list of Crown Jewels and categorise them according to your business needs but make sure you include the core components as these are the basis for the mission critical assets to rely upon.
- Be thorough and consistent with documentation.
- Be thorough and consistent throughout the model process so as to reach coherent and comparable outcomes between iterations.

7. Conclusions

Some FIs implement a compliance-based (or audit-based) approach to protecting particular information, which is unlikely to single out mission-critical information assets for specialised protection. This can result in significant gaps that remain undiscovered until a security incident occurs.

Other FIs apply a risk-based approach, although research indicates that in many instances these efforts are not focused on the risks specific to mission-critical information assets. Consequently, important activities such as detailed analysis of the threats or accounting for the complete footprint can be overlooked.

For many FIs the skills shortage, combined with investment constraints, inhibits their ability to build on existing approaches to provide balanced and comprehensive protection. A common side effect is an over-reliance on fundamental controls and a lack of enhanced and specialised controls. The standard / default use of a single threat modelling methodology, within an FI looking would indeed raise the security posture level but up to a certain level. However, based on the literature review and the answers provided by the real-world case study, when the scope is explicitly narrowed down to protecting the Crown Jewels, then the FI must invest & build on the selected approach in order to reach an effective, secure end state of CJ's, beyond compliance.

References

- VerSprite. (2019, 10 8). *VerSprite*. Retrieved from VerSprite: <https://versprite.com/>
- David, G. (2017, 6 30). *assist-software.net*. Retrieved from assist-software.net: <https://assist-software.net/blog/scrum-framework-roles-activities-and-artifacts>
- Techopedia. (2019, 10 8). *techopedia.com*. Retrieved from techopedia.com: <https://www.techopedia.com/definition/22193/software-development-life-cycle-sdlc>
- Delzer, C. (2018, 7 11). *SBS CyberSecurity LLC*. Retrieved from SBS CyberSecurity LLC: <https://sbscyber.com/resources/data-flow-diagrams-101>
- Visual-Paradigm.com. (2019, 10 8). *Visual-Paradigm.com*. Retrieved from Visual-Paradigm.com: <https://www.visual-paradigm.com/guide/data-flow-diagram/what-is-data-flow-diagram/>
- TrendMicro. (2019, 10 8). *TrendMicro*. Retrieved from TrendMicro: <https://www.trendmicro.com/vinfo/us/security/definition/indicators-of-compromise>
- Petters, J. (2018, 10 23). *Varonis*. Retrieved from Varonis: <https://www.varonis.com/blog/ids-vs-ips/>
- Brook, C. (2018, 12 5). *digitalguardian.com*. Retrieved from digitalguardian.com: <https://digitalguardian.com/blog/what-user-and-entity-behavior-analytics-definition-ueba-benefits-how-it-works-and-more>
- McAfee. (2019, 10 8). *McAfee*. Retrieved from McAfee: <https://www.mcafee.com/enterprise/en-us/security-awareness/operations/what-is-soc.html>
- UcedaVélez, T., & Morana, M. M. (2015). Risk Centric Threat Modeling: Process for Attack Simulation and Threat Analysis. In T. UcedaVélez, & M. M. Morana, *Risk Centric Threat Modeling: Process for Attack Simulation and Threat Analysis* (p. 675). Wiley.
- UcedaVelez, T. (2015, 8 31). *GSA.gov*. Retrieved from GSA.gov: <https://interact.gsa.gov/sites/default/files/Mon%20AM2-SW%20Assurance%20Fall%20SSCA%20Forum-Sept%202015.pdf>

- ISACA. (2014, 6 4). *ISACA.org*. Retrieved from ISACA.org:
<http://www.isaca.org/chapters5/Ireland/Documents/Forms/DispForm.aspx?ID=11>
- Lefkowitz, J. (2019, 7 1). *SecurityWeek.com*. Retrieved from SecurityWeek.com:
<https://www.securityweek.com/risk-based-vulnerability-management-must-security-compliance>
- PenTestuniversity. (2017, 5 23). *pentestuniversity.org*. Retrieved from pentestuniversity.org:
<https://www.pentesteruniversity.org/escaping-scope-creep/>
- Agrawal, G. (2019, 1 18). *mrcissp.com*. Retrieved from mrcissp.com:
<https://mrcissp.com/2019/01/18/threat-modeling-a-step-by-step-guide/>
- Shostack, A. (2014). *Threat modeling: designing for security*. John Wiley & Sons.
- ThreatModeler. (2016, 4 15). Retrieved from <https://threatmodeler.com>
- ThreatModeler. (2016, 4 15). Retrieved from <https://threatmodeler.com>
- Shevchenko, N. (2018, 12 3). *insights.sei.cmu.edu*. Retrieved from Carnegie Mellon University:
https://insights.sei.cmu.edu/sei_blog/2018/12/threat-modeling-12-available-methods.html
- Veltsos, C. (2017, 11 20). *SecurityIntelligence.com*. Retrieved from SecurityIntelligence:
<https://securityintelligence.com/take-a-load-off-delegate-cyber-risk-management-using-the-three-lines-of-defense-model/>
- CyberEdu, F. (2019, 10 10). *ForcePoint*. Retrieved from ForcePoint:
<https://www.forcepoint.com/cyber-edu/ot-operational-technology-security>
- Kohnfelder, P. G. (1999). *The threats to our products*.
- LeBlanc, M. H. (2002). *Writing Secure Code*. Microsoft Press.
- Wilson, T. (2012, 5 16). *pluralsight.com*. Retrieved from pluralsight.com:
<https://www.pluralsight.com/blog/it-ops/access-control-list-concepts>
- Microsoft. (2009, 12 11). *Microsoft*. Retrieved from [https://docs.microsoft.com/en-us/previous-versions/commerce-server/ee823878\(v=cs.20\)?redirectedfrom=MSDN](https://docs.microsoft.com/en-us/previous-versions/commerce-server/ee823878(v=cs.20)?redirectedfrom=MSDN)
- Imperva. (2019, 10 10). *Imperva.com*. Retrieved from <https://www.imperva.com/learn/application-security/botnet-ddos/>
- Michael Muckin, S. C. (2019). *Lockheed Martin*. Retrieved from Lockheed Martin:
<https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/LM-White-Paper-Threat-Driven-Approach.pdf>
- McCollum, D. J. (2018, 6 28). *Mitre.org*. Retrieved from Mitre.org:
https://www.mitre.org/sites/default/files/publications/pr_18-1631-ngci-system-of-systems-threat-model.pdf
- Deborah J. Bodeau, C. D. (2018, 4 7). *MITRE*. Retrieved from mitre.org:
https://www.mitre.org/sites/default/files/publications/pr_18-1174-ngci-cyber-threat-modeling.pdf
- Michlin, I. (2016, 3 16). *NCCgroup*. Retrieved from <https://www.nccgroup.trust/uk/about-us/newsroom-and-events/blogs/2016/march/threat-prioritisation-dread-is-dead-baby/>

- Seller, D. (2006, 4 12). *Microsoft Blogs*. Retrieved from <https://blogs.msdn.microsoft.com/dansellers/>
- Eric Lachapelle, F. R. (2015, 14 10). *PECB*. Retrieved 10 12, 2019, from Professional Evaluation and Certification Board : <https://pecb.com/whitepaper/risk-assessment-with-octave>
- BRUNSCHWILER, C. (2013). *compass-security.com*. Retrieved 10 12, 2019, from <https://blog.compass-security.com/2013/04/lean-risk-assessment-based-on-octave-allegro/>
- Richard A. Caralli, J. F. (2007, 5 4). *Carnegie Mellon*. Retrieved 10 12, 2019, from https://resources.sei.cmu.edu/asset_files/TechnicalReport/2007_005_001_14885.pdf
- Scarfone, M. S. (2016, 3). *National Institute of Standards and Technology*. Retrieved 10 14, 2019, from https://csrc.nist.gov/CSRC/media/Publications/sp/800-154/draft/documents/sp800_154_draft.pdf
- Shuttleworth, M. (2017). *Project Risk Manager*. Retrieved 10 14, 2019, from <https://www.project-risk-manager.com/blog/qualitative-and-quantitative-risk-analysis/>
- Optiv. (2011). *Optiv*. Retrieved 10 15, 2019, from <https://www.optiv.com/blog/inherent-and-residual-risk-how-both-scores-drive-enterprise-risk-decisions>
- Zalewski, A. H. (2012). *semanticscholar.org*. Retrieved 10 15, 2019, from https://pdfs.semanticscholar.org/d3a8/8f79f3baf7c1f3ad75fada8ec2b71b27ca99.pdf?_ga=2.167142267.456045756.1571166544-1103033312.1571166544
- ISF. (2019). *ISF - SecurityForum*. Retrieved 10 18, 2019, from <https://www.securityforum.org/>
- LogicManager. (2019). *LogicManager*. Retrieved 10 18, 2019, from <https://www.logicmanager.com/erm-software/knowledge-center/best-practice-articles/risk-appetite-risk-tolerance-residual-risk/>
- FireEye. (2019). *FireEye*. Retrieved 10 18, 2019, from <https://www.fireeye.com/current-threats/apt-groups.html>
- Doerr, C. (2018, 11 16). *Enisa*. Retrieved 10 19, 2019, from <https://www.enisa.europa.eu/events/2018-cti-eu-event/cti-eu-2018-presentations/cyber-threat-intelligence-standardization.pdf>
- Launius, S. (2018, 3 1). *sans.org*. Retrieved 10 19, 2019, from <https://www.sans.org/reading-room/whitepapers/threatintelligence/evaluation-comprehensive-taxonomies-information-technology-threats-38360>
- Bromander, V. M. (2017, 9). *Cyber Threat Intelligence Model: An Evaluation of*. Retrieved 10 19, 2019, from <https://www.researchgate.net/publication/319701970>
- MITRE. (2019). *mitre.org*. Retrieved 10 19, 2019, from <https://cwe.mitre.org/>
- MITRE. (2019). *MITRE ATT&CK*. Retrieved 10 20, 2019, from <https://attack.mitre.org/>
- MITRE. (2019). *MITRE CAPEC*. Retrieved 10 20, 2019, from <https://capec.mitre.org/data/index.html>
- Casey, T. (2007, 11). *www.sbs.ox.ac.uk*. Retrieved 10 20, 2019, from <https://www.sbs.ox.ac.uk/cybersecurity-capacity/system/files/Intel%20->

%20Threat%20Agent%20Library%20Helps%20Identify%20Information%20Security%20Risks.
pdf

- MITRE. (2019, 10 16). *MITRE CVE*. Retrieved 10 20, 2019, from <https://cve.mitre.org/>
- FireEye. (2013, 10 1). *FireEye*. Retrieved 10 20, 2019, from <https://www.fireeye.com/blog/threat-research/2013/10/openioc-basics.html>
- MITRE. (2018, 6 1). *STIXproject*. Retrieved 10 20, 2019, from <https://stixproject.github.io/about/>
- MITRE. (2019, 9 7). *MAECproject*. Retrieved 10 20, 2019, from <https://maecproject.github.io/>
- Zareen Syed, A. P. (2016). <https://www.semanticscholar.org>. Retrieved 10 20, 2019, from https://pdfs.semanticscholar.org/67b3/c0893013cbdcc9f35ec9359fa4466df7360e.pdf?_ga=2.108917250.125639893.1571524259-1103033312.1571166544
- HSSEDI. (2019). *U.S. Department of Homeland Security (DHS)*. Retrieved 10 20, 2019, from <https://www.dhs.gov/science-and-technology/apex-ngci>
- Matthew B. Miles, A. M. (2014). *Qualitative Data Analysis: A Methods Sourcebook and The Coding Manual for Qualitative Researchers*.
- Holt, M. (2016, 8 16). *How to secure mission-critical*. Retrieved 10 21, 2019, from <http://www.securityforum.org>
- David B. Fox, E. I. (2018, 5 2). *MITRE*. Retrieved 10 24, 2019, from https://www.mitre.org/sites/default/files/publications/pr_18-1613-ngci-enterprise-threat-model-technical-report.pdf
- IBM. (2017, 9 8). *IBM*. Retrieved 10 24, 2019, from <https://www.ibm.com/downloads/cas/8MGOLOB7>
- OWASP. (2016, 9 28). *Wikimedia Commons*. Retrieved 10 24, 2019, from https://commons.wikimedia.org/wiki/File:Data_Flow_Diagram_-_Online_Banking_Application.jpg
- Gartner. (2019). *Gartner*. Retrieved 10 24, 2019, from <https://www.gartner.com/en/information-technology/glossary/bia-business-impact-analysis>
- Edge, K. R. (2007). *semanticscholar*. Retrieved 10 24, 2019, from <https://www.semanticscholar.org/paper/The-Use-of-Attack-and-Protection-Trees-to-Analyze-Edge-Raines/eaf9e4bc88ce24f85dae46aa45d918171445ce49>
- McCollum, D. J. (2018, 6 28). <https://www.mitre.org/>. Retrieved 1 1, 2020, from https://www.mitre.org/sites/default/files/publications/pr_18-1631-ngci-system-of-systems-threat-model.pdf
- Selin, J. (2019). *Evaluation of threat modeling methodologies*. JAMK University of applied sciences.
- Dunn Cavelty, M. (2010). The Reality and Future of Cyberwar. *Parliamentary Brief*.

List of Tables

List of Figures

FIGURE 1 - ASSET'S CLASSIFICATION	41
FIGURE 2 - STRIDE ELEMENTS	41
FIGURE 3 - STRIDE TABLE	42
FIGURE 4 - WEB APPLICATION THREAT MODEL INCORPORATING IDDIL ELEMENTS	42
FIGURE 5 - DFD SAMPLE OF WIRE TRANSFERS OF A FI	43
FIGURE 6 - OCTAVE STAGES & STEPS	43
FIGURE 7 - APPLICATION THREAT MODEL EXAMPLE	44
FIGURE 8 - OPERATIONAL THREAT MODEL EXAMPLE	44
FIGURE 9 - ASSET IDENTIFICATION ACTIVITIES EXAMPLE	45
FIGURE 10 - CYBER THREAT ACTORS AND MOTIVATIONAL FACTORS AGAINST FSS/FIS	45
FIGURE 11 – MAJOR DATACENTRE HIGH LEVEL DIAGRAM	46
FIGURE 12 - IRAM2 PHASE C ACTIVITIES	46
FIGURE 13 - ONLINE BANKING ENVIRONMENT DFD & CORE COMPONENTS	46
FIGURE 14 - THREAT LANDSCAPE POPULATION	47
FIGURE 15 - ADVERSARIAL THREAT ATTRIBUTION	47
FIGURE 16 - ACCIDENTAL THREATS ATTRIBUTION	47
FIGURE 17 - PRIORITIZED THREAT LANDSCAPE	47
FIGURE 18 - LOI & TS FACTORS SCALING	48
FIGURE 19 - SUBJECT THREAT EVENT CATALOGUE	48
FIGURE 20 - THREAT EVENTS MAPPING	49
FIGURE 21 - ATTACK TREE EXAMPLE	49
FIGURE 22 - IN-SCOPE THREATS & THREAT EVENTS MAPPING	50
FIGURE 23 - BIA SAMPLE OUTPUT	50