

# Security Triage: An Industrial Case Study on the Effectiveness of a Lean Methodology to Identify Security Requirements

Matteo Giacalone  
Security and Safety, Poste  
Italiane SpA  
giacal69@posteitaliane.it

Federica Paci  
DISI, University of Trento  
federica.paci@unitn.it

Rocco Mammoliti  
Security and Safety, Poste  
Italiane SpA  
mammoliti.rocco@posteitaliane.it

Rodolfo Perugino  
Security and Safety, Poste  
Italiane SpA  
peruginor@posteitaliane.it

Fabio Massacci  
DISI, University of Trento  
Fabio.Massacci@unitn.it

Claudio Selli  
Security and Safety, Poste  
Italiane SpA  
c.selli@posteitaliane.it

## ABSTRACT

**Context:** Poste Italiane is a large corporation offering integrated services in banking and savings, postal services, and mobile communication. Every year, it receives thousands of change requests for its ICT services. Applying to each and every request a security assessment “by the book” is simply not possible. **Goal:** We report the experience by Poste Italiane of a lean methodology to identify security requirements that can be inserted in the production cycle of a normal company. **Method:** The process is based on surveying the overall IT architectures (*Security Survey*) and then a lean dynamic process (*Security Triage*) to evaluate individual change requests, so that important changes get the attention they need, minor changes can be quickly implemented, and compliance and security obligations are met. **Results:** The empirical evaluation conducted for over an year at Poste Italiane shows that the process significantly reduces the time to identify security requirements at the pace of change. **Conclusions:** The Security Survey and Triage process should thus be embedded in a company’s production cycle as mandatory step to manage change requests so that security initiatives are prioritized based on the relevance of the assets and of the business objectives of the company.

## Categories and Subject Descriptors

[Software and its Engineering]; [Security and Privacy]

## General Terms

Theory, Experimentation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ESEM '14, September 18-19, 2014, Torino, Italy  
Copyright 2014 ACM 978-1-4503-2774-9/14/09 ...\$15.00.

## Keywords

Security Requirements Elicitation, Change Requests

## 1. HOW TO SECURE EVOLVING ICT SYSTEMS?

Modern corporations must evolve at an increasing pace to offer advanced and innovative services. They must do it securely, fending off old and emerging threats by cyber-criminals. Further, they must comply with a slate of regulations at different levels of abstraction and all the above must be achieved with a lean and cost-effective budget.

For Poste Italiane - the largest Italian employer offering integrated services in finance, logistics, and mobile communication with a turnaround of around 24 billion Euro - balancing security and change means identifying security requirements for over 150 change requests/month and over 2000/year.

A simple solution would be to mandate the most difficult security requirements across the board. Yet, very high security often bring severe performance or usability penalties. For example, strong authentication (e.g. by biometric or a hardware token) is used for appropriate financial transactions but could be mandated for all services as well. This would yield a significant drop in performance, huge deployment costs and would be rightly perceived by many users as a ridiculous burden if they just need to check whether grandma’s birthday parcel has arrived.

The alternative of sloppy-security-for-all is not an option either. Individual changes of one application may have domino effects on other services. Many intermediate IT components are used by different top level services and are subject to different compliance requirements. Touching one application without careful analysis of its implication may lead to severe fines or even criminal prosecution.

Company’s management wants that every change request goes through a security gate and the simple solution is to just follow the books: many security risk assessment standards and methodologies can be used (e.g. ISO 27005 [13], USA’s NIST 800-30 [25], CoBIT [11] Germany’s BSI [5], France’s EBIOS [2], Spain’s Magerit [8], UK’s IAS [7], etc.). At lower abstraction levels one can also follow company-based methodologies such as Cigital’s BSIMM or Microsoft’s

STRIDE [10]. Academic methods are also available like SI\* [9], CORAS [15], SQUARE, [18], and SREP [19].

Both industry and academic approaches alike look simple and straightforward on academic papers and in glossy magazines of security consultants. Yet, they very rarely report the *actual effort* needed to perform a security analysis “by-the-book” in an industrial setting. They have been designed in order to support the identification of security requirements “at the early phases”. Therefore, they are appropriate when an entirely new service is designed, implemented and deployed, but require a huge effort to manage change request.

The earliest publication in the open literature mentioning the actual effort for the identification of security requirements [3] reported that “The CommerceNet requirements analysis [Re-engineering the web server for taking electronic payments] was conducted by 4 analysts, the authors, and various stakeholders, for approximately 30 hours a week over a period of four months”. A back of the envelope calculation sum up to over a thousand person/hour, essentially a full-time employee for a year.

With thousands of requests per year, applying to each and every request a security requirements methodology “by the book” is just not possible. Using the numbers from [3] would require a workforce of around 400 people working full time for the whole year long, just to identify the requirements! Practice may shrink both months and analysts by a factor, but will not largely change the picture.

We need to identify security requirements in matter of weeks, or even days for minor applications, while giving the right attention to important change requests that affect critical assets of the company.

In this report we present the results of a year long project at Poste Italiane where a lean, innovative security requirements methodology has been experimented *in vivo* and successfully deployed at a large scale. First, we introduce the business objectives and provide an overview of the solution. The core of the report presents the high level issues behind managing a change request within Poste Italiane (§3), and the process to identify security requirements “by the book” (§4). Then, we introduce the two key components of our solution - Security Survey (§5), and Security Triage (§6) and we discuss their efficacy (§7). We conclude with threats to validity (§8), related work (§9) and our conclusions (§10).

## 2. OVERVIEW OF THE SOLUTION

The high level business objective is to streamline the security requirements identification process so that it can process thousands of requests per year. This can be broken down into a number of sub goals:

- *Q1* Can we identify the criticality level of a change requests (and its related security requirements) more quickly than with the standard process?
- *Q2* Can we use resources in proportion to change requests’ relevance for the company?
- *Q3* Can we still be sure that the appropriate security requirements are identified?

Our solution to these challenges combines two key ideas from two very related disciplines: architecture and medicine. We summarize them in Figure 1.

The first notion is that of *Security Triage* [1]. In medicine, the Triage is the process where “medical personnel systematically categorize victims of a disaster into three groups: those who will die whether treated or not; those who will resume normal lives whether treated or not, and those for whom medical treatment may make a significant difference. Each group requires a different strategy. The first group receives palliative care, the second group waits for treatment, and the third requires some ranking in light of available resources. As new victims appear, personnel must repeat the categorization”. In our methodology, the Security Triage is performed directly by the proposer of the change request, the “owner” of the business service, along the guidelines of the Security Team. The Service Owner classifies change requests based on their relevance for the company as described in Section 6. Requests with a “red code” are subject to a full fledged analysis “by the book”. Requests with “white code” proceed directly to implementation of baseline security requirements and deliver quickly value to internal users and services. This addresses directly *Q1* and *Q2*. Still, we must be sure that a Security Triage is not just a politically correct term for a sloppy security assessment, which would fail *Q3*. It would also fail *Q1* if it didn’t help us to identify the *right* level of security importance. The second instrument, the *Enterprise (security) Survey* is our solution to the problem. In architecture, a land surveyor builds a detailed map of an area by observations, measurements in the field, research of legal instruments, and data analysis in order to establish property boundaries, identify buildings and support planning (of new buildings). The Survey provides the identification of the components of the IT architecture, the breakdown of those into compliance and security perimeter against which a Triage (for the new component) can be successfully performed. Notice that a survey is *not* just an architectural diagram, no more than a map is the only result of a land survey. Attaching business values, identifying owners, drawing legal boundaries, etc. are all essential part of both a land and security survey. The combination of Survey and Triage makes sure that also *Q3* is met on the new system.

## 3. REQUESTS FOR SECURITY ASSESSMENT

When a change request or a new project is proposed within the company the *Security Department* is responsible for the security analysis and the identification of the appropriate security requirements while the *ICT Department* is in charge of the actual implementation and deployment (or its outsourcing) of the concrete solutions (including the security services and security monitors).

The *Service Owner* is the department using or managing the service; this typically includes “functionalists”, people that manage the services, who are experts in the domain and its functional, legal, and business requirements;

The *System Owner* manages the technological chain that actually deliver the service to the end customers. This role can be further classified into Development Manager who takes care of the development of the application and the System Manager who takes care of the administration and operations of the systems involved. Other people (e.g. from the legal or marketing departments) may also be involved, depending on the complexity of the change requests.

A number of issues must be addressed by the analysis. At first the proposed requirements must address all direct or

**triage:** *noun*,

1. (in medicine) the assignment of degrees of urgency to wounds or illnesses to decide the order of treatment of a large number of patients or casualties.
2. (in cybersecurity) the assignment of degrees of security criticality to change requests or new projects to decide the order of security treatment of a large number of ICT services

**survey** *verb*,

1. (in architecture) examine and record the area and features of (an area of land) so as to construct a map, plan, or description.
2. (in cyber security) examine and record the components, features and interactions of (a business service) so as to construct a map, plan, or description of the IT architecture.

**Figure 1: Survey and Triage: the Key Components for Managing Security Requirements**

**Table 1: Regulatory Compliance Obligations**

Perimeter	Description	Baseline
Privacy	Protection of personal data, sensitive and judicial data	National Law, Technical Annex to Law, Internal Guidelines
Financial Data	Protection and tracking of financial transactions, money transfers and financial information	National Authority Regulation, Technical Annex to Law, Internal Guidelines
Bank of Italy Regulation	Compliance with the provisions of direction and control issued by the Bank of Italy	National Authority Regulation, National Regulator Terms of Reference
Traffic Data	Communication Directive, Traffic (Phone and internet) Data Management	National Authority Regulation, Technical Annex to Law, Internal Guidelines
Financial Relevant	Adaptation of organizational models and systems of control provided by the so-called "Savings Protection Law"	National Laws

**Table 2: Security Compliance Obligations**

Perimeter	Description
Credit Card Data	PCI DSS standard to enhance payment card data security
Electronic Payments	Security guidelines for the management of electronic payment systems
Internet Web Sites	Security guidelines for the management of Internet Web Sites
Critical Infrastructures	Guidelines for the management of Critical Infrastructures

indirect security regulations. Table 1 summarizes *some* of the the regulatory issues faced by Poste Italiane.

The *Baseline* describes the norms that must be addressed to achieve a minimum compliance while the *Perimeter* is the set of Services and Applications which have to comply with the Baseline. We will see later how such classification will be used by *Security Survey* to identify some of the components of the minimum security requirements. The table has to be intended as illustrative because many other obligations (e.g. workers' regulation) are not included.

Beside external regulations, the team must also make sure that internal guidelines adopted by the company are correctly implemented. Table 2 lists some of the external security standards and internal security guidelines that are applied to the applications within the appropriate perimeter.

As we mentioned change requests tally over 2000/year with an average around 150/month. They have different

levels of complexity in terms of implementation. For example, the first months of 2014 featured 17 highly complex initiatives, 35 of medium complexity and more than 200 of low complexity. They must be processed quickly and the security assessment is a mandatory quality gate.

#### 4. SECURITY ANALYSIS "BY THE BOOK"

The default application of an *Information Security Risk Management Process* (ISRM) for establishing, implementing, maintaining and continuously improving an Information Security Management System (ISMS) follows the ISO 27001 standards. We briefly sketch its key steps (see [13] for details):

1. *Asset and Process Identification* captures and describes the overall enterprise architecture of the process to be identified;
2. *Business Impact Analysis* focuses on the information used by each service and the impacts of possible compromise of the confidentiality, integrity, and availability of that information;
3. *Risk Assessment* at the level of *Process, People, Application, Infrastructure, Facilities* is then performed in order to identify gaps and current risk levels;
4. *Security Requirements Identification* addresses the gaps determined in the previous phase and produces a plateau of security measures that can be implemented by the Service Owner;

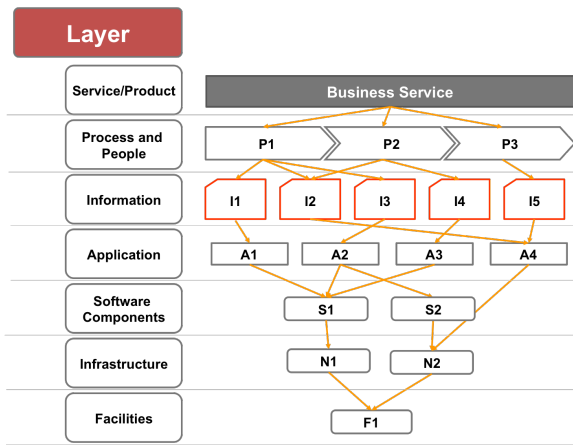


Figure 2: An Enterprise Architecture of a Service

5. *Risk Treatment*, including acceptance of residual risks, is performed by the Service Owner on the basis of the analysis and the business consideration with the support of the ICT Department who actually implements the technical solutions proposed and accepted by the Service Owner.

The first step of the process generates an enterprise architecture. Figure 2 sketches a generic enterprise architecture for a service provided by Poste Italiane that could be identified during the Asset and Process Identification step. The architecture spans different layers: Service/Products, Process and People, Information, Application, Software components, Infrastructure, and Facilities. They will be the subject of the detailed Risk Assessment at step 3 above.

The Business Impact Analysis can take many forms which depend on one's favorite standard. For the Service and Information layers this step has been progressively streamlined in order to include it in the Triage process. We will therefore discuss it later in Section 6.

During the Risk Assessment step, for each of the architectural layers a significant number of interviews is conducted with the Service Owner and the System Owners in order to perform a detailed gap analysis. Table 3 summarizes the effort just in terms of interviews and controls for the example in Fig. 2. So, for example, for each of the three process composing the service one needs to consider 300 questions about the presence of controls and fill a questionnaire which requires a least 3 hours. This does not include the time to actually acquire the knowledge to correctly answer the questions. Once the risk has been assessed, the missing controls and some additional compliance measures are delivered to the Service Owner in order to decide the appropriate risk treatment.

Beside the effort required to the Security Team, a significant effort is also asked to Service and System Owners. They are the only members of the company that can provide to the Security Team the appropriate information to perform the asset identification, and the business, process, and application analysis. A simple back-of-the-envelope calculation shows that the questionnaire by itself takes 2 working days and with 2000 requests per year it is more than 10 persons working full time for a year just to answer the questions, let alone understanding the systems, understand-

Table 3: Example of Effort for ISRM Analysis

Level	Questions	Time	Unit
Process/People	300	3hrs	Process
Information	16	1hrs	Data
Applications	250	3hs	Application
Software components	200	2hrs	Type of Asset
Infrastructure	200	2hrs	Type of Asset
Facilities	100	1hrs	Facility

Traditional ISRM requires to answer almost 1000 questions, for more than a full day of work, without mentioning the time necessary for actually finding the answer for each question.

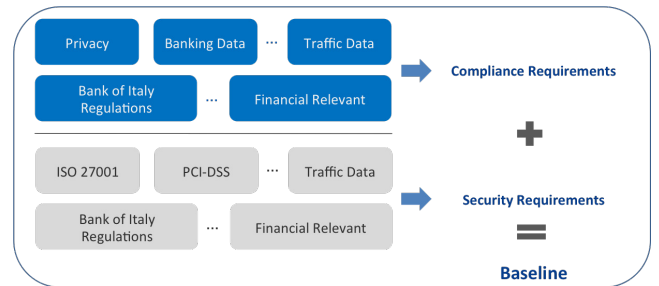


Figure 3: Perimeters determine baseline of security requirements

ing the risks, doing any mitigation, and finally delivering the desired change requests to the end customers.

## 5. ICT SECURITY SURVEY

The high level purpose of the *ICT Security Survey* process is to provide a comprehensive characterization of the ICT and business services of the company. They are registered in a comprehensive catalog (CIRM - Catalog Information Security Risk Management) that is then periodically updated.

The first step of this process can be seen as an ongoing refinement of the enterprise architecture assessment steps of the standard security requirements identification process. It brings together the Business View and the IT View of the systems. To keep the comparison with land surveying this is the measurements of the fields and the drawing of the maps.

Each "logical" component of the business and ICT landscape (services, macro-products, products, applications) is then categorized as belonging to a number of *Perimeters* that determine the baseline in terms of security requirements that must be implemented. We called this part of the process "Mapping". In land surveying this would correspond to draw the boundaries and identifying the ownership of the various tract of land that have been measured in the first phase. Each perimeter would be a layer in the physical map (e.g. rainfall, ownership, vegetation etc.). Overall there are more than 20 different perimeters that can be cross-combined in a variety of ways. Table 1 and Section 3 already described some of them.

The combinations of the compliance and security perimeters determines the security baseline of each service as shown in Figure 3.

A further classification is then performed on services, applications and components to assess the relevance of each

asset for the company. This assessment takes into account a variety of sources and in particular: a) the type of data that is processed, b) the relevance of the various perimeters to the particular application or service, c) the impact that a security compromise might have on the service or the application, and d) the economic relevance of the service.

A suitable function transforms these assessments in a 0-1000 scale and this is then used to cluster them in five macro categories from C1 (lowest level) to C5 (highest level). The aggregation function cannot be disclosed for obvious reasons, but we can provide a indication of what is included in each level:

- **Level C1:** services that do not manage personal data and are not associated with security perimeters;
- **Level C2:** services that handle personal data which are not associated with any of security and compliance perimeters;
- **Level C3:** services that manage personal data which are bind to security and compliance perimeters;
- **Level C4:** services that manage personal and sensitive data or services that have medium economic relevance;
- **Level C5:** services that are fundamental for the company from a business perspective and that are bound to relevant security and compliance perimeters.

Each time a request for a change or a new IT initiative is accepted a *consolidation of assessment levels in the survey's catalog* is performed:

1. at first an initial level of the request  $C_{req}$  is determined by the Triage;
2. then levels of services affected by requests are updated

$$C_{srv} = F\{C_{req}|request\ affects\ service\} \cup \{C_{srv}\}$$

The definition of  $F$  and  $G$  functions is currently under revision: several alternatives are possible like maximum or weight by costs or revenues.

## 6. SECURITY TRIAGE

The ISRM process considers the whole stack of layers of a service while the Security Triage process is centered on two of the upper layers: Services and Information. The main purpose of this simplification is to make the assessment doable by the Service Owner herself because she might not have an inkling on the internal IT plumbing (nor should she be required to have it). Knowledge of the latter should be indeed the responsibility of the System Owner, whose knowledge have been already captured by the survey's process.

When a change request (or a new IT initiative) is placed, the Service Owner performs the steps below:

- Identify the service related to the change request and the relevant information handled by the processes and applications supporting the service under analysis;
- Identify the compliance and security perimeters (if changed from the service already described in the survey's catalog);

- Provide additional information on applicability of Privacy regulations to the data.

After this initial analysis, the *Business Impact Analysis* (BIA) is performed. The Service Owner is faced with a number of categories of potential losses. Some of these categories are for example a) Economic Operating Loss is evaluated directly by Service Owner as she can identify the actual monetary amount that a breach to the service will imply; b) Loss of Reputation as could be perceived by suppliers, end users, and national regulators in case of breaches of the service's security; c) Loss of Competitive advantage includes the possibility for competitors to exploit the security breach to gain market share or even directly exploit the leaked information for direct purposes; and d) Legal Liabilities include issues such as fines or criminal prosecution related to security breaches.

For each category a Service Owner is faced by 16 questions grouped by impact type: *loss to confidentiality*, *loss to integrity and loss to availability*. Answer to each category have been streamlined to "make sense" for a Service Owner (as opposed to a Security Expert). For example, for the economic losses, a Service Owner should assess whether a loss to availability for a certain number of hours might lead to minor or major economic loss. She will have to provide the information for a range of hours of downtime. Another example: for the legal liability category, she might be asked whether a violation to integrity might lead to an administrative offense with monetary fine or a criminal offense with minimum jail terms.

Out of experience, we discovered that losses to availability can be easily assessed in degrees (e.g. 1 hour vs 2 hours vs 1 day etc.) whereas for integrity and confidentiality it is better to provide a on/off state (the confidentiality is either compromised or it is not compromised). Typical notions used by Security Expert such as session compromise, forward compromise, root control etc. are difficult to grasp. They would be investigated for the change requests that have the highest level C5.

At this point the information provided by the Service Owner during the Triage questions are combined with the information from the survey's catalog to obtain a final value for the level  $C_{req}$  aggregating together the value of *Perimeters*, *BIA*, *Economics* and additional information. This combination is company specific. Different companies may appreciate differently the amount of a regulatory fine for non-compliance needed to move from C3 to C4. Change requests or new initiatives that have a high level (e.g. C5 and possibly C4) are then subject to a complete information security risk assessment (ISRM) "by the book" as detailed in Section 4.

As final outcomes, the Security Triage process produces: a) Baseline of Security Requirements to which the project has to comply, b) Security relevant Level for each correlated service, c) Business Impacts deriving from a loss of confidentiality, integrity, and availability, d) Business Continuity Objectives for the services and and e) an updated survey's catalog.

The results of the assessment and the changes are then fed back to the survey's catalog as we mentioned at the end of Section 5. Therefore the updated service will be tagged with the new level as determined by the Triage (and the eventual ISRM follow-up for the critical initiatives).

Table 4 shows the difference between the outcomes of the main steps of "by the book" approach (ISRM) and the pro-

**Table 4: Comparison of Security Survey and Triage with ISRM**

ISRM Activity	Security Survey & Triage	ISRM
Asset Identification	Service and information	All levels
Business Impact	Critical information	All
R.A. Process	Main process	All
R.A. Application	Main application	All
R.A. Infrastructure	Derived from baseline	Specific
R.A. Facilities	Derived from baseline	Specific
Security Reqs	By perimeters assignment	Gap analysis
Risk Treatment	Derived from baseline	Gap analysis
Service Owner Effort	Low	High/medium

posed one by Security Triage in terms of risk assessment step 3. The assignments of baselines of security requirements depending on perimeters is where the need of both Survey and Triage is most apparent. Without a survey that pre-determine the perimeters it would be possible to misplace the security requirements services. For example a Service Owner could modify a components used for C1 service without being aware that this component also delivers vital information to a C5 service. This might have led the company to severely under-estimate the impact of an apparently innocuous change request and a likely violation of compliance obligations.

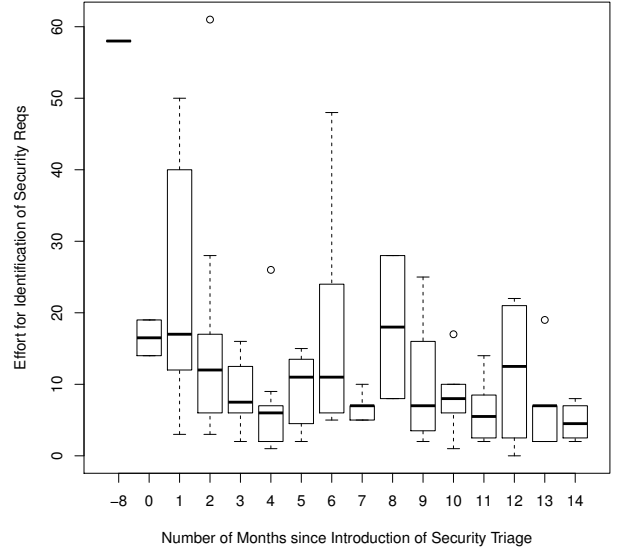
## 7. EMPIRICAL EVALUATION

The process of surveying and triage has been first piloted in 2012, with a product in the e-financial sector in March 2012. In June 2012 over 1600 new change requests have been identified as possible activities to be included in the pilot phase. By the end 2012 other two major pilots for internal processes for financial procurement and supply chain management have been concluded. In 2013 the procedure has been applied to a much larger scale within the company and it is now in full swing.

In the remainder of the section we use the generic term *effort* to indicate the time required for a security assessment. It is not possible to disclose the exact amount in days or hours as this would be a confidential information for Poste Italiane. However, for the purpose of the empirical evaluation, this variable has been measured uniformly across the various requests and makes *relative* comparisons possible.

At first we evaluate the research question *Q1* and namely whether the new process made it possible to identify quickly the C1-C5 level for each new service requests.

Figure 4 shows the boxplot distribution of security assessment requests over the past year and a half that participated into the study (142). It is immediate to see that there has been a sharp drop in the average time needed to identify the security requirements. At the beginning of the study (Month 8), all security research assessments followed the approach “by the book”, with an essentially constant (and very high) effort. At the beginning of the activity the time



*With the progressive adoption of Security Survey & Triage the effort for security analysis progressively decreases in the year.*

**Figure 4: Effort to Identify Requirements by Month**

required to fill the questionnaire is still very variable as the Service Owner must often get back to the Security Team to grasp the questions. As time goes by the process has less variability and the mean is significantly reduced.

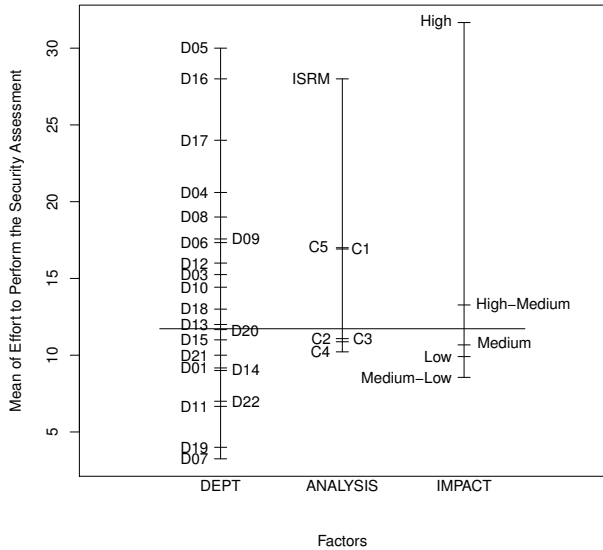
The data shows that the following hypothesis is true and statistical significant by one-sided Mann-Whitney-Wilcoxon test with a p-value of 0.006482

*H<sub>1</sub>: The effort required to identify security requirements in the second semester is lower than in the first semester.*

Figure 5 shows the distribution of the mean of the effort required to identify security requirements grouped by the departments who have placed the change requests, by the process applied to identify the security requirements and by impact that change requests had on business. The effort was measured as number of days required to identify security requirements. The first vertical line shows the mean of the effort grouped by departments placing a change request. The second vertical line shows the mean of the effort required to identify security requirements following the ISRM process and the effort required when using the Security Survey and Triage grouped by the level of relevance of the change requests (C1-C5). The third vertical line represents the mean of the effort required to identify security requirements grouped by the impact that change requests had on the business. ISRM analysis takes almost twice more effort than Security Survey and Triage process (C1...C5). This is also attested by the results of Mann-Whitney test that shows the difference in effort is statistically significant (p-value = 0.00453).

In summary, the following hypothesis is true:

*H<sub>2</sub>: The effort required to identify security requirements with ISRM is higher than with Security Survey and Triage (C1-C5).*



*ISRM analysis takes almost twice more effort than SecurityTriage-based analysis (C1...C5). As a control, notice effort is distributed evenly across all departments and high impact cases take a significant effort as expected.*

**Figure 5: Plot Design of Effort by Category**

We have also tried to test whether C5 assessments took more time than other assessments (a refinement of  $Q2$ ) but this is only weakly confirmed with p-value of the MW one-sided test of 0.0739 (just above the threshold). This may be a random fluctuation and is the subject of further investigation because the C1-C5 level of the request is the final result of the analysis which is therefore not affected by it.

Table 4 shows that research question  $Q3$  is also met. The security requirements are produced by the baseline and are appropriate for achieving the compliance obligations. In this way it is possible to cover all change requests with a baseline security analysis and for C5 services a detailed risk analysis is still performed. At the same time the Service Owner saves significant time as shown by hypothesis  $H_1$  and the Q-Q analysis between the first and second semester.

## 8. THREATS TO VALIDITY

Unknown confounding factors may explain the relationship between the treatment (adoption of Security Survey and Triage) and the outcome (less effort by Service Owners). A relevant confounding factor that could have explained the decrease in time could be the Security Team pushing Service Owner to perform a quicker assessment. An alternative explanatory factor could also have been that the fact of being monitored actually altered the natural behavior of the persons being observed. Both threats do not apply to our study, as incentives are properly aligned: Service Owners have the strongest self-interests to cut the time for their change requests to be shipped to customers. It is well known in the literature that security is always felt as a burden [6]. The anecdotal evidence shows that the communication was also properly aligned: Service Owners usually contacted the Security Team in order to ask for support and advice.

A more important threat to validity is that the Service Owner might understand the Security Triage incorrectly, and the wrong application might lead to incorrect data. This threat is present but we have some mitigations. At first, the Service Owner can ask the support of the Security Team to make sure she correctly understood the issues at stake. As a further measure the questions in the Triage have been set up from a service perspective and not from a security perspective. For example, asking whether a breach to confidentiality may lead to an administrative fine or a competitor gaining unfair advantages is a service question albeit it has security implications.

A final issue concerns the ability to generalize study’s results beyond the study’s settings. The main threat to external validity is that the study has been conducted in one company, and hence in a specific context, which is a threat to generalization. The large number and extremely diverse nature of internal departments of Poste Italiane (ranging from logistics to financial services), give us good counter balance to this threat. In order to control for such factors one would need to perform a randomized assignments of change requests to “placebo” methodology and to the Security Survey and Triage-based methodology. Unfortunately, this is not possible for two reasons: an industrial process must be uniform for all requests, and a company cannot deliberately assign bogus-security to some production software. In fact, whenever conducting industry research in vivo not all variables can be controlled (as is the case for controlled experiments in vitro). This is an intrinsic limitation of any industrial report. Still, the number of data points that we consider for our analysis exceeds those of many in-vitro experiments in security requirements [14, 16, 21].

## 9. RELATED WORK

There are many standards, practices, and methods available for identifying and removing information security risks. The ISO/IEC 27005 [13] and ISO/IEC 31000 [12] are the standards to undertake risk management at the corporate level. The NIST SP 800-30 is another standard for security risk assessment proposed by US National Institute of Standards and Technology [25]. There are also several national-level methodologies for security risk assessment like UK’s HMG Information Assurance Standard [7], France’s EBIOS [2], Spain’s MAGERIT [8] and Germany’s BSI [5].

Other approaches to security risk assessment are the COBIT methodology sponsored by the ISACA institute [11], SABSA [23] and STRIDE [10]. COBIT focuses on identifying business goals first and deriving security controls from those [11]. Similar to COBIT, SABSA methodology focuses on the identification of business requirements and the refinement of them into a set of security controls.

Academic methodologies to identify security requirements are also available like SQUARE [18], SREP [19], CORAS [15], Misuse Cases [24] and Attack Trees [22].

However, only few academic papers [14, 16, 21, 26] have studied the actual effectiveness of these methods. They adopted the Method Evaluation Model (MEM) [20] to compare the applications of different method and evaluate perceived and actual efficacy of different methods. In [26] Opdhal and colleagues have repeated the experiment with industrial practitioners. Both experiments show that attack trees help to identify more threats than misuse cases. More recently, Labunets et al. [14] have conducted a controlled experiment

with MSc students to compare visual methods (CORAS) and textual methods (SREP) for security risk assessment. The experiment shows that visual methods are more effective and better perceived by the participants. In [16], Massacci and Paci report an interesting protocol to perform empirical comparisons of different security and risk assessment methods by using both practitioners and students.

In this paper we presented a novel methodology that facilitates the identification of security requirements at the pace of change and we used a statistical hypotheses testing to show its efficacy in reducing the effort required to manage change requests.

## 10. CONCLUSIONS

In this paper we have reported a lean innovative methodology for the identification of security requirements stemming from a year long project conducted by Poste Italiane. The process is based on an global mapping analysis of the overall ICT landscape (*Security Survey*) and then a lean dynamic process (*Security Triage*) to quickly identify the level of relevance of a request for security assessment and the corresponding security requirements.

We have also provided some preliminary data on its efficacy: the approach significantly reduces the time to identify security requirements at the pace of change.

The Security Survey and Triage process should be embedded in a company's production cycle as mandatory step to manage change requests so that security initiatives are prioritized based on the relevance of the assets and of the business objectives of the company.

## 11. ACKNOWLEDGMENT

The work of University of Trento has been partly supported by the European Union 7th Framework Programme under grant agreements n.256980 (NESSOS) and n.285223 (SECONOMICS) and the Joint Undertaking SESAR WPE under contract 12-120610-C12 (EMFASE). The work of Poste Italiane has been partly funded by MIUR-PON under project PON03PE\_00032\_2\_02 within the framework of the Technological District on Cyber Security.

## 12. REFERENCES

- [1] A. Davis. The art of requirements triage. *Computer*, vol. 36, no. 3, 42–49, 2003.
- [2] ANSSI. EBIOS 2010 - Expression of Needs and Identification of Security Objectives, 2010.
- [3] A. I. Anton and C. Potts. The use of goals to surface requirements for evolving systems. In *Proc. of ICSE '98*, 157–166, 1998.
- [4] D. Baca and K. Petersen. Countermeasure graphs for software security risk assessment: An action research. *JSS*, 86(9):2411 – 2428, 2013.
- [5] BSI. Standard 100-2, The IT-Grundschutz Methodology, 2011.
- [6] H. Cavusoglu, H. Cavusoglu, and J. Zhang. Security Patch Management: Share the Burden or Share the Damage? *Management Science*, 54(4):657-670, 2008.
- [7] CESB. HMG Information Assurance Standard 1, 2009.
- [8] F. L. Crespo, M. A. Amutio Gomez, J. Candau, and J. A. M. Manas. Magerit v2 –Methodology for Information Systems Risk Analysis and Management - Book I - The Method, 2006.
- [9] P. Giorgini, F. Massacci, J. Mylopoulos, N. Zannone. Requirements engineering for trust management: model, methodology, and reasoning. *IJIS* 5(4), 257-274, 2006.
- [10] S. Hernan, S. Lambert, T. Ostwald, and A. Shostack. Threat modeling-uncover security design flaws using the STRIDE approach. *MSDN Magazine*, 68–75, 2006.
- [11] ISACA. COBIT 5: A Business Framework for the Governance and Management of Enterprise IT, 2012.
- [12] ISO/IEC. 31000:2009 – Risk Management, 2009.
- [13] ISO/IEC. 27005:2011–Information Technology–Security Techniques–Information Security Management Systems–Requirements, 2011.
- [14] K. Labunets, F. Massacci, F. Paci, and L. M. Tran. An experimental comparison of two risk-based security methods. In *Proc. of ESEM '13*, 163–172, 2013.
- [15] M. S. Lund, B. Solhaug, and K. Stolen. A guided tour of the CORAS method. In *Model-Driven Risk Analysis*, 23–43, Springer, 2011.
- [16] F. Massacci and F. Paci. How to select a security requirements method? A comparative study with students and practitioners. In *Proc. of NordSec '12*, 89–104. Springer, 2012.
- [17] G. McGraw. Software security: building security in. Addison-Wesley Professional, 2006.
- [18] N. R. Mead and T. Stehney. Security quality requirements engineering (SQUARE) methodology. *SIGSOFT*, 30(4):1–7, May 2005.
- [19] D. Mellado, E. Fernández-Medina, and M. Piattini. Towards security requirements management for software product lines *CSI*, 30(6):361–371, 2008.
- [20] D. L. Moody. The method evaluation model: a theoretical model for validating information systems design methods. In *Proc. of ECIS '03*, 1327–1336, 2003.
- [21] A. L. Opdahl and G. Sindre. Experimental comparison of attack trees and misuse cases for security threat identification. *Inform. Software Tech.*, 51(5):916–932, 2009.
- [22] B. Schneier. Attack trees. *Dr. Dobbs's Journal*, 24(12):21–29, 1999.
- [23] J. Sherwood, A. Clark, and D. Lynas. Enterprise security architecture: a business-driven approach. Backbeat Books, 2005.
- [24] G. Sindre and A. Opdahl. Eliciting security requirements with misuse cases. *REJ*, 10(1):34–44, 2005.
- [25] G. Stoneburner, A. Goguen, and A. Feringa. Risk management guide for information technology systems. NIST, 2002.
- [26] P. Karpati, Y. Redda, A.L. Opdahl, G. Sindre. Comparing attack trees and misuse cases in an industrial setting. *Inf. Soft. Technology*, 56 (3), 294 - 308, 2014.