International Journal on Advances in Security

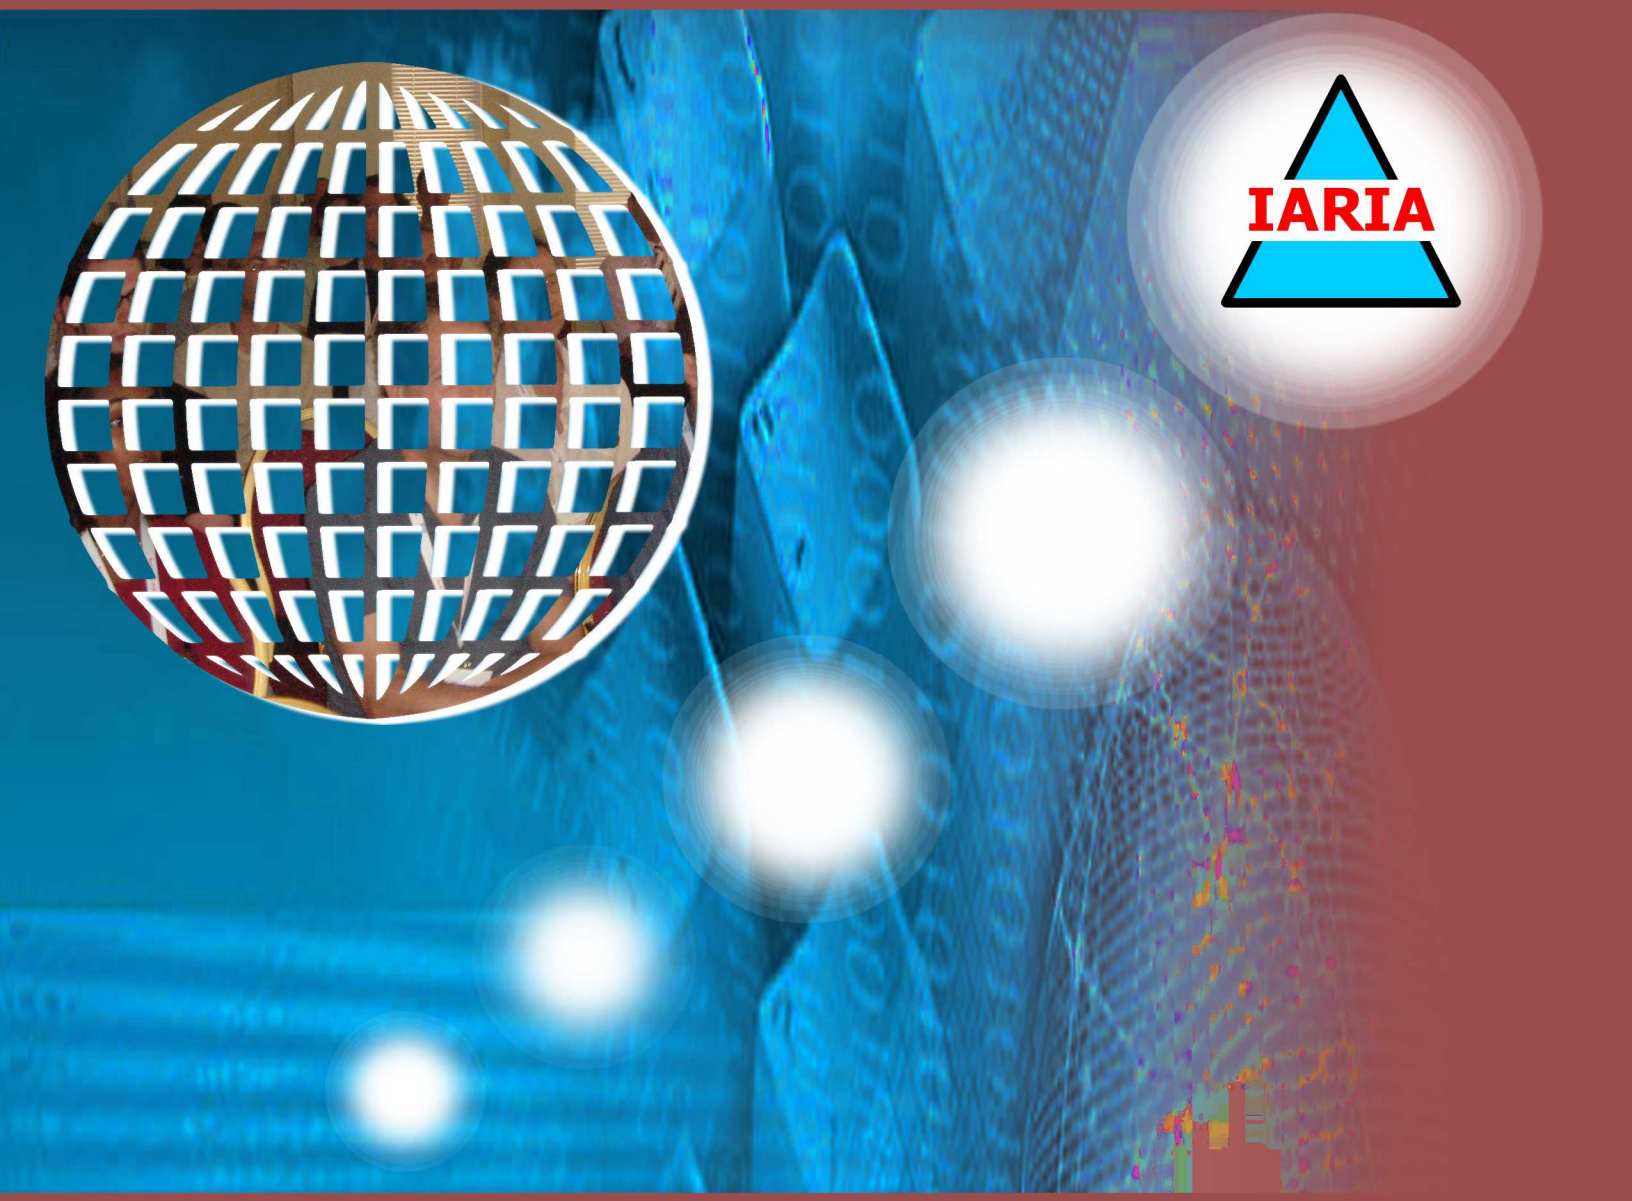

2018 vol. 11 nr. 3&4

The *International Journal on Advances in Security* is published by IARIA.
ISSN: 1942-2636
journals site: http://www.iariajournals.org
contact: petre@iaria.org

Sponsored by IARIA
www.iaria.org

International Journal on Advances in Security
Volume 11, Number 3 & 4, 2018

Catalin V. Birjoveanu, "Al.I.Cuza" University of Iasi, Romania
Wolfgang Boehmer, Technische Universitaet Darmstadt, Germany
Alexis Bonnecaze, Université d'Aix-Marseille, France
Carlos T. Calafate, Universitat Politècnica de València, Spain
Juan-Vicente Capella-Hernández, Universitat Politècnica de València, Spain
Zhixiong Chen, Mercy College, USA
Clelia Colombo Vilarrasa, Autonomous University of Barcelona, Spain
Peter Cruickshank, Edinburgh Napier University Edinburgh, UK
Nora Cuppens, Institut Telecom / Telecom Bretagne, France
Glenn S. Dardick, Longwood University, USA
Vincenzo De Florio, University of Antwerp & IBBT, Belgium
Paul De Hert, Vrije Universiteit Brussels (LSTS) - Tilburg University (TILT), Belgium
Pierre de Leusse, AGH-UST, Poland
William Dougherty, Secern Consulting - Charlotte, USA
Raimund K. Ege, Northern Illinois University, USA
Laila El Aimani, Technicolor, Security & Content Protection Labs., Germany
El-Sayed M. El-Alfy, King Fahd University of Petroleum and Minerals, Saudi Arabia
Rainer Falk, Siemens AG - Corporate Technology, Germany
Shao-Ming Fei, Capital Normal University, Beijing, China
Eduardo B. Fernandez, Florida Atlantic University, USA
Anders Fongen, Norwegian Defense Research Establishment, Norway
Somchart Fugkeaw, Thai Digital ID Co., Ltd., Thailand
Steven Furnell, University of Plymouth, UK
Clemente Galdi, Universita' di Napoli "Federico II", Italy
Emiliano Garcia-Palacios, ECIT Institute at Queens University Belfast - Belfast, UK
Birgit Gersbeck-Schierholz, Leibniz Universität Hannover, Germany
Manuel Gil Pérez, University of Murcia, Spain
Karl M. Goeschka, Vienna University of Technology, Austria
Stefanos Gritzalis, University of the Aegean, Greece
Michael Grottke, University of Erlangen-Nuremberg, Germany
Ehud Gudes, Ben-Gurion University - Beer-Sheva, Israel
Indira R. Guzman, Trident University International, USA
Huong Ha, University of Newcastle, Singapore
Petr Hanáček, Brno University of Technology, Czech Republic
Gerhard Hancke, Royal Holloway / University of London, UK
Sami Harari, Institut des Sciences de l'Ingénieur de Toulon et du Var / Université du Sud Toulon Var, France
Daniel Harkins , Hewlett Packard Enterprise, USA
Ragib Hasan, University of Alabama at Birmingham, USA
Masahito Hayashi, Nagoya University, Japan
Michael Hobbs, Deakin University, Australia
Hans-Joachim Hof, Munich University of Applied Sciences, Germany
Neminath Hubballi, Infosys Labs Bangalore, India
Mariusz Jakubowski, Microsoft Research, USA
Ángel Jesús Varela Vaca, University of Seville, Spain
Ravi Jhawar, Università degli Studi di Milano, Italy
Dan Jiang, Philips Research Asia Shanghai, China

Georgios Kambourakis, University of the Aegean, Greece
Florian Kammueller, Middlesex University - London, UK
Sokratis K. Katsikas, University of Piraeus, Greece
Seah Boon Keong, MIMOS Berhad, Malaysia
Sylvia Kierkegaard, IAITL-International Association of IT Lawyers, Denmark
Hyunsung Kim, Kyungil University, Korea
Geir M. Køien, University of Agder, Norway
Ah-Lian Kor, Leeds Metropolitan University, UK
Evangelos Kranakis, Carleton University - Ottawa, Canada
Lam-for Kwok, City University of Hong Kong, Hong Kong
Jean-Francois Lalande, ENSI de Bourges, France
Gyungho Lee, Korea University, South Korea
Clement Leung, Hong Kong Baptist University, Kowloon, Hong Kong
Diego Liberati, Italian National Research Council, Italy
Giovanni Livraga, Università degli Studi di Milano, Italy
Gui Lu Long, Tsinghua University, China
Jia-Ning Luo, Ming Chuan University, Taiwan
Thomas Margoni, University of Western Ontario, Canada
Rivalino Matias Jr ., Federal University of Uberlandia, Brazil
Manuel Mazzara, UNU-IIST, Macau / Newcastle University, UK
Catherine Meadows, Naval Research Laboratory - Washington DC, USA
Carla Merkle Westphall, Federal University of Santa Catarina (UFSC), Brazil
Ajaz H. Mir, National Institute of Technology, Srinagar, India
Jose Manuel Moya, Technical University of Madrid, Spain
Leonardo Mostarda, Middlesex University, UK
Jogesh K. Muppala, The Hong Kong University of Science and Technology, Hong Kong
Syed Naqvi, CETIC (Centre d'Excellence en Technologies de l'Information et de la Communication),Belgium
Sarmistha Neogy, Jadavpur University, India
Mats Neovius, Åbo Akademi University, Finland
Jason R.C. Nurse, University of Oxford, UK
Peter Parycek, Donau-Universität Krems, Austria
Konstantinos Patsakis, Rovira i Virgili University, Spain
João Paulo Barraca, University of Aveiro, Portugal
Sergio Pozo Hidalgo, University of Seville, Spain
Yong Man Ro, KAIST (Korea advanced Institute of Science and Technology), Korea
Rodrigo Roman Castro, University of Malaga, Spain
Heiko Roßnagel, Fraunhofer Institute for Industrial Engineering IAO, Germany
Claus-Peter Rückemann, Leibniz Universität Hannover / Westfälische Wilhelms-Universität Münster / North-German Supercomputing Alliance, Germany
Antonio Ruiz Martinez, University of Murcia, Spain
Paul Sant, University of Bedfordshire, UK
Peter Schartner, University of Klagenfurt, Austria
Alireza Shameli Sendi, Ecole Polytechnique de Montreal, Canada
Dimitrios Serpanos, Univ. of Patras and ISI/RC ATHENA, Greece
Pedro Sousa, University of Minho, Portugal
George Spanoudakis, City University London, UK

International Journal on Advances in Security
Volume 11, Numbers 3 & 4, 2018

CONTENTS

Applying Soft Systems Methodology to Complex Problem Situations in Critical Infrastructures: The CS-AWARE Case Study

Veronika Kupfersberger*, Thomas Schaberreiter*, Chris Wills†, Gerald Quirchmayr* and Juha Röning‡
*Faculty of Computer Science
University of Vienna (Vienna, Austria)
e-mail: veronika.kupfersberger@univie.ac.at
e-mail: thomas.schaberreiter@univie.ac.at
e-mail: gerald.quirchmayr@univie.ac.at
†CARIS Research Ltd. (Fowey, United Kingdom)
e-mail: ccwills@carisresearch.co.uk
‡Faculty of Information Technology and Electrical Engineering
University of Oulu (Oulu, Finland)
e-mail: juha.roning@oulu.fi

***Abstract*—Modern technology, in addition to all its benefits, creates new threats and attack vectors to individuals and organisations. In the past years, the number of cyber attacks has increased drastically as has the extent of their effects. These circumstances clearly show that a different approach to cybersecurity is required: a holistic, collaborative strategy to improve the security situation for society and the economy as a whole. In the European Union, the legal framework that is currently developing (like the network and information security (NIS) directive), recognises the increasing need for cooperation and collaboration among individual actors to improve cybersecurity. Information sharing is therefore one of the key elements of the NIS directive. In this paper, we present and demonstrate a system and dependency analysis based on soft systems thinking. This approach is able to capture the relations between assets and their internal and external dependencies in the complex systems of organisations. It is applicable to critical infrastructures or other organisations that base their operations on complex systems and interactions. The analysis approach introduced is done in a socio-technological manner; the human aspect of the systems is considered as important as the technical or organisational aspects. The case study presented in this paper, covering the first steps towards the development of a holistic cybersecurity awareness solution, is based on three focus points: an initial threat assessment for local public administrations (LPAs), an analysis of external information sources and an analysis of the piloting scenarios based on the first round of soft systems analysis workshops. The results of which are essential to the development of the solutions implementation framework and further software development.**

Keywords–Cybersecurity; Critical Infrastructures; System Analysis; Soft Systems Methodology; Socio-technological Analysis; Cyber Situational Awareness; Information Sharing.

I. INTRODUCTION

Cybersecurity is one of today's most challenging societal security problems, affecting both individuals and organisations, such as strategic/critical infrastructures, large commercial enterprises, SMEs, non-governmental organisations (NGOs) or governmental institutions. The extensive variety of these attacks is one of the issues, as is the lack of communication between organisations and administrations that have been the target of an attack. Deliberate or accidental threats and attacks threaten digitally administered data and digitally handled processes. Sensitive data leaks can ruin the reputation of companies and individuals, and the interruption of digital processes that organisations rely upon in their daily work flow can cause severe economic disadvantages. This work builds on the paper on how to adress complex situations in critical infrastructure published in SECURWARE 2017 [1]. Reaching beyond the technology-focused boundaries of classical information technology (IT) security, cybersecurity strongly interrelates with organisational and behavioural aspects of IT operations, and the need to comply with the current and actively developing legal and regulatory framework for cybersecurity. For example, the European Union (EU) recently passed the NIS directive that obliges member states to get in line with the EU cybersecurity efforts [2]. Most EU member states and the EU itself have a cybersecurity strategy in place which will eventually lead to the introduction of laws and regulations that fulfil cybersecurity requirements. One of the main aspects of the NIS directive, as well as the European cybersecurity strategies, is cooperation and collaboration among relevant actors in cybersecurity. Enabling technologies for coordination and cooperation efforts are situational awareness and information sharing. Situational awareness in this context is a runtime mechanism to gather cybersecurity relevant data from an IT infrastructure and visualise the current situation for a user or operator. Information sharing refers to the ability to share this information with cybersecurity information sharing communities, like the NIS relevant authorities. In the long term, it is expected that due to the awareness generated information sharing can improve cybersecurity sustainably and benefit society and economy as a whole.

One of the major aspects of information sharing to facilitate collaboration and cooperation, is a proper understanding of the cybersecurity relevant aspects within an organisation's systems. This is a complex and often neglected task that will, as we argue in this paper, greatly improve the cy-

bersecurity of organisations in the context of cybersecurity situational awareness and cooperative/collaborative strategies towards cybersecurity. We introduce and demonstrate a system and dependency analysis methodology to analyse the environment and: (a) Identify the assets and dependencies within the system and how to monitor them; (b) capture not only technological aspects, but the socio-technical relations within the organisation; (c) identify external information sources that could either be provided by official and cybersecurity specific sources (for example, legal/regulatory framework, standardisation, cybersecurity information sharing communities), or more general publicly available information relating to cybersecurity (for example, social networks or twitter); (d) provide the results in a form that can be utilised by support tools.

We base our work around established and well proven methods related to systems thinking, the soft systems methodology (SSM) and PROTOS-MATINE/GraphingWiki. The case study presented in this paper tests the idea of using these methods to analyse complex domains and derive a coherent analysis. The results of the case study will be critical in assuring a high quality software development of a cybersecurity awareness solution for local public administrations. As of now, the first round of user workshops, the initial threat assessment and the analysis of external information sources have yielded essential information for defining an implementation framework. The upcoming second and third round user workshops held in the pilot municipalities will work mainly with information collected during the before-mentioned analysis of threats, external sources and the first workshops.

The paper is organised as follows: Section II discusses background and related work, Section III details our system and dependency analysis approach. In Section IV, an example in the context of CS-AWARE, a European H2020 project which uses the presented system and dependency analysis as a core part of its cyber security solution, is introduced and followed by the summary of the first round of workshops in the pilot scenarios in Section V. Section VI discusses the results of approach and Section VII concludes the paper.

II. Related Work

In December 2015, the European Parliament, the European Council and the European Commission agreed on the European NIS directive as the first EU wide legislation on cybersecurity [2]. The directive lays down the obligations of member states concerning NIS. Most notably for this work, it requires the implementation of proper national mechanisms for incident prevention and response, in addition to information sharing and cooperation mechanisms. The NIS directive is the main action stemming from the EU cybersecurity strategy [3], which emphasises the need for a decentralised prevention and response to cyber incidents and attacks. By now, most EU countries have put a national cybersecurity strategy in place [4] that is in line with many actions proposed by the NIS directive. Coordination and information sharing are key elements of the strategy, with the requirement for national NIS authorities, national law enforcement and defence authorities to interact with each other, as well as their EU counterparts. International cooperation and coordination is envisioned at the EU level. On the standardisation front, the ISO/IEC 27000 [5] standard is the first in a series of standards on information security management that have provided organisations with a best practice framework for assessing security risks and implementing security controls as countermeasures. Similarly, the privacy focused ISO/IEC 29100 [6] standard provides a framework to help organisations to manage and protect personally identifiable information. In 2011 the European standardisation organisations CEN, CENELEC and ETSI have formed the cybersecurity coordination group (CSCG), which was converted to the focus group on cybersecurity in 2016 [7], in order to undertake the strategic evaluation of IT security, cybersecurity and NIS standardisation.

The systems analysis methodology which will be mainly used in this work is the Soft Systems Methodology developed by Peter Checkland [8][9]. Cognitive mapping, casual loop diagrams [10] or a combination of stakeholder analysis and cognitive mapping as suggested by Ferretti [11], would have been alternatives. Generally, the key thought behind the soft systems methodology is that it is hard to completely analyse and describe a complex system, especially if human interaction plays a key role. SSM represents an analysis methodology that aims to achieve an holistic understanding of the system while at the same time only focusing on the actual problems at hand. Soft Systems Methodology has been used in an extraordinarily wide variety of problem domains as diverse as knowledge management in the building industry [12], to evaluating government policy to promote technological innovation in the electricity sector [13]. In the case of the building industry example, the tacit knowledge held by staff involved in the tendering process was made explicit by the application of SSM. In the case of the electricity supply industry, SSM was used understand how better to to promote and foster technological innovation in the sector.

The PROTOS-MATINE methodology [14] is another approach that relates to systems thinking. While the SSM focuses on understanding complex systems and processes by interviewing its users, PROTOS-MATINE takes the standpoint that a truly holistic view on complex situations can only be achieved if as many relevant information sources as possible (e.g., technical, organisational, human on all organisational levels as well as external and publicly available information), are combined to create a complete picture and eliminate discrepancies between information from different sources. The key to PROTOS-MATINE is that collected information from different sources is set in context to each other and graphically processed and visualised to make it simple for domain experts to identify discrepancies in information coming from different sources. For this purpose, GraphingWiki [15], a graphical extension to the MoinMoin Wiki, was developed to visualise dependencies between semantic data collected in Wiki pages in the context of PROTOS-MATINE. The methodology was used in many case studies, for example for highlighting vulnerabilities in anti-virus software [16] and for a socio-technological analysis of a VoIP (voice over IP) provider [17]. In [18], the methodology was extended for analysing complex systems in the critical infrastructure context, where the analysis goal is to achieve a dependency graph of critical infrastructure assets, dependencies between the assets and measures to observe those assets (base measurements).

III. Soft Systems Analysis in the Context of Cybersecurity for Complex Systems

The system and dependency analysis proposed in this paper is seen as the basis for the automatic incident detection and cybersecurity situational awareness efforts of future cybersecurity initiatives, as discussed in the related work. The objective is to identify in the specific organisational context what needs cybersecurity protection and what are the main threats it needs protection from. More specifically, this means that the challenge for system and dependency analysis is to identify the assets within an organisation and their internal and external dependencies in order to be able to protect them from cybersecurity threats. Observable information sources that can be used to determine the on-line state of those assets need to be identified to allow for monitoring and detecting abnormal behaviour, thus describing the security state. Furthermore, the goal of the system and dependency analysis is to identify external information sources that can provide information to help detect and classify security threats correctly. Those information sources can be dedicated cybersecurity information providers like, for example, computer emergency response teams (CERTs) or other threat and vulnerability databases, or they can be publicly available information sources via, for example, platforms like Twitter, Facebook or Google+. The usage of open source intelligence (OSINT) has been proved to be valuable before in other contexts like disaster management. Sail Labs Media Mining System is an example of a system which makes use of freely available information. It aims to allow accurate situational analysis of crisis locations by analysing different relevant data feeds. It gathers information from multiple sources including television, radio and various Internet sources and uses data mining techniques to extract information about the content [19].

Since technology is only one factor in cybersecurity, the system and dependency analysis is designed to capture and monitor the socio-technical nature of an IT infrastructure. It takes into account the human, organisational and technological factors, as well as other legal/regulatory and business related factors that may contribute to the cybersecurity in a specific context. As can be seen in Figure 1, systems thinking is a way of looking at some part of the world, by choosing to regard it as a system, using a framework of perspectives to understand its complexity and undertake some process of change. The key concepts are holism - looking at things as a whole and not as isolated components and systemic - treating things as systems, using systems ideas and adopting a systems perspective.

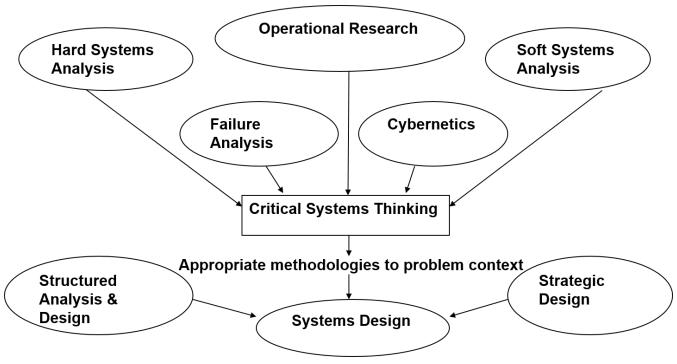

Figure 1. Systems thinking - The systems approach

Two concepts of systems thinking are hard systems thinking and soft systems thinking. Hard systems design is based on systems analysis and systems engineering. It assumes that the world is comprised of systems that we can describe and that these systems can be understood through rational analysis. It is based on the assumption that it is possible to identify a "technically optimal" engineering solution for any system and that we can then write software to create the "solution". Hard systems design assumes that there is a clear consensus as to the nature of the problem that is to be solved. It is unable to depict, understand or make provisions for "soft" variables such as people, culture, politics or aesthetics. It is based on the assumption that it is possible to identify a "technically optimal" engineering solution for any system. It assumes that those commissioning the system have the ability and power to implement the system. While hard systems design is highly appropriate for domains involving engineering systems structures that require little input from people, the complex systems and interactions in critical infrastructures or other organisations - especially with cybersecurity in mind - usually do not allow this type of analysis. Hard systems design is inappropriate and unsuitable for analysing human activity systems that require constant interaction with, and intervention from people. Such systems are complicated, fuzzy, messy and ill defined and are typified by unclear situations, differing viewpoints and unclear objectives, containing politics, emotion and social drama. This is the type of system domain for which an SSM design approach is highly appropriate and to which it should be applied. That is not to say that the SSM approach cannot or should not be used in the design of engineering systems and structures, indeed one of the authors has used this approach very successfully in many complex and diverse problem domains. For example, SSM has been used by one of the authors in the design of naval command and control systems for the British Navy and in the design of system architectures for automated fare collection in very large light railway and mass transit operations.

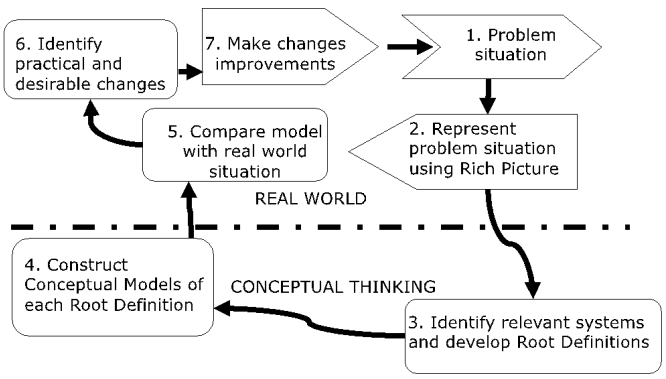

Figure 2. Soft systems design

An overview of the stages of SSM is set out in Figure 2. The SSM methodology has 7 steps: (1) Enter the problem situation; (2) Express the problem situation; (3) Formulate root definitions of systems behaviour; (4) Build conceptual models of systems in root definitions; (5) Compare models with real-world situations; (6) Define possible and feasible changes; (7) Take action to improve the problem situation. A detailed description of the approach is beyond the scope of this paper,

however, reader may wish to refer to Checkland's work [8][9]. In this work, we will focus on the earlier steps of the SSM that deal with the system analysis and problem definition (specifically, steps 1-4). One key element of this phase is that systems stakeholders (users, managers, administrators, etc.) are engaged in workshops to define the problems they are facing, since those who are using systems on a daily basis are the ones that have the most information about it. Since this is not explicit knowledge, but tacit knowledge, it is important to create an environment that facilitates information sharing. The SSM utilises rich pictures for this purpose, and depicting the problem in a rich picture is a key stage early in the process. Rich pictures are a representation of the problem domain. They utilise "cartoon-style" techniques to portray a complex situation and concentrate on:

- Structure - Key individuals, organisations etc.
- Process - What could be or is happening?
- Climate - Pressures, attitudes, cultures, threats etc.

An example of a Rich Picture depicting a malfunctioning airline passenger check-in system appears in Figure 3, outlining different viewpoints in case the system goes off-line.

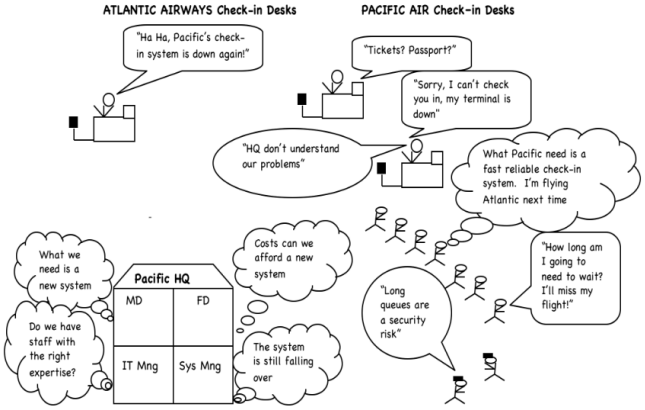

Figure 3. Rich picture of an airline check-in system

Rich pictures are a tool for understanding where we are and are a mix of drawings, pictures, symbols and text. They represent a particular situation or issue and they are depicted from viewpoint(s) of the person or people who drew them. They can both record and evoke insight into a situation. Rich pictures are pictorial 'summaries' of a situation, embracing both the physical, conceptual and emotional aspects of a problem situation. They can depict complicated situations or issues, and relevant systems are identified from the rich picture. These systems are described in Root Definitions, which are then used in conjunction with the rich pictures to develop Conceptual Models. These are formed from the actions stated or implied in the Root Definition(s). Of course, each rich picture may be interpreted from quite differing 'world view points'. A Conceptual Model is like an activity sequence diagram, but is aimed at representing a conceptual system as defined by the logic of the Root Definition and not just a set of activities.

The role of PROTOS-MATINE and GraphingWiki in this proposed analysis method is to complement the information gathering effort in the user workshops with information from other sources, and provide a solid base for discussion in those workshops through visualisation. The main additional sources are expected to be legal requirements and regulatory efforts like the NIS directive; cybersecurity relevant standardisation like the ISO/IEC 27000 family of standards and information about relevant and current risks and threats via official sources like CERTs, or more dynamic information sources like social media. Where relevant, the information received via rich pictures from the workshop participants can easily be complemented by more detailed information available such as, for example, technical manuals, business continuity plans or disaster recovery plans. One of the capabilities of GraphingWiki is to instantly link gathered information to other relevant information and thus allowing to update the graphical representation of the analysed system as soon as new information arrives. We hope to utilise this feature in the user workshops to create more dynamic discussions and give even more incentive to the participants to create a system model that is as close to reality as possible.

The expected result of the proposed system and dependency analysis will be a dependency graph containing an organisations security relevant or critical assets and the dependencies among them. Furthermore, observable measurements that are able to determine the security state of those assets are identified and associated to them. Through GraphingWiki this dependency graph is in digital form and can be further utilised as the basis for advanced cybersecurity situational awareness and monitoring services. One example of such a service will be given in the next section.

IV. The CS-AWARE Approach

CS-AWARE is a European H2020 project that was funded by the European Union under the project number 740723. The aim of the project is to improve the cybersecurity situation in local public administrations (LPAs). While the project is focused on LPAs, the ideas and methods developed in this project are applicable to any organisations that rely on complex systems, interactions and procedures (like strategic/critical infrastructures, large organisations or SMEs).

As can be seen in Figure 4, the main building blocks of the CS-AWARE solution are the system and dependency analysis, data collection and data analysis to achieve the project's goals of cybersecurity situational awareness, cybersecurity information exchange and system self-healing. The proposed solution aims at improving automated situational awareness in small- to medium-sized IT infrastructures, however it is expected that the same principals would also apply to large organisations or critical infrastructures. The system and dependency analysis presented in the previous section is an integral part of two project phases. Besides the actual system and dependency analysis, which will be conducted according to the methodology presented in Section III (Steps 1-4 of the SSM as well as PROTOS-MATINE/GraphingWiki related aspects), it will provide the main input for the self-healing component, based on steps 5-7 of the SSM.

The core idea of the CS-AWARE project is to automate the cybersecurity effort of organisations as much as possible, and provide an on-line situational awareness tool that aims to base its recommendations on a holistic view of an organisation's IT systems and dependencies, but also on the cybersecurity

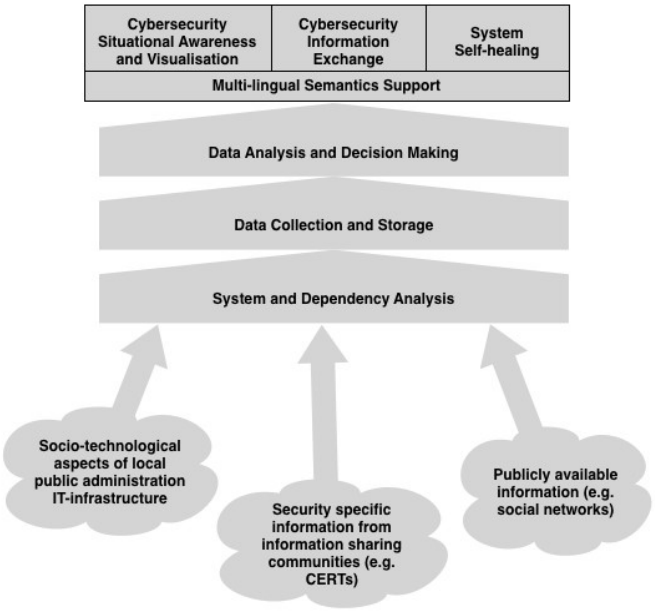

Figure 4. CS-AWARE overall concept

situation in general (for example by observing the risk and threat landscape). The end users of the CS-AWARE solution are expected to be the people responsible for cybersecurity in an organisation, such as the chief security officer (CSO), or system administrators. CS-AWARE is a decision support system that will allow its users to detect cybersecurity incidents quickly and identify the affected systems, since the key assets and security relevant dependencies have been identified during system and dependency analysis. Countermeasures can be initiated by the people responsible for cybersecurity in a timely manner. Besides manual countermeasures, CS-AWARE includes a self-healing component that is closely tied to the system and dependency analysis. The later steps of the SSM (especially steps 5-7) are concerned with defining solutions to the problems identified during analysis. In CS-AWARE one focus point will be to identify and develop possible countermeasures to cybersecurity threats and define policies and procedures that can be invoked if such a threat materialises. Those policies and procedures will be utilised by the self-healing component and can be configured to be invoked automatically if a threat materialises. This will allow the system, depending on the scenario, to prevent or mitigate the damage and/or recover from the incident.

The intelligent and fully automated part of the CS-AWARE project are the *data collection and storage* and *the analysis and decision making* components. Based on the system and dependency analysis results, the base measurements from internal and external sources are observed and when relevant data points are collected, pre-processed and stored. The data analysis component is capable of detecting suspicious behaviour like threat and attack patterns in the data sets it receives and will classify and rank them accordingly, as an input to the decision support in the situational awareness and visualisation component. The accuracy of the decision making component will depend on the cooperation and collaboration efforts and the quality of data that is provided by information sharing authorities. It is envisaged that threat detection can achieve highly accurate unsupervised results once cybersecurity information exchange is an established concept and can provide accurate information relating to cybersecurity threats and attack patterns.

The *cybersecurity situational awareness and visualisation* component is the user interface to the CS-AWARE solution. It will visualise the security relevant aspects of an organisations socio-technological systems, based on the dependency graph received during system and dependency analysis. State changes triggered by the decision making component will cause a visualisation of the affected components and its dependencies. Possible countermeasures will be suggested and self-healing procedures can be configured and invoked, where relevant.

The *cybersecurity information exchange* is the connection point to the cybersecurity information sharing authorities, for example NIS competent authorities like national or EU CERTs. While cybersecurity information sharing is currently still in its infancy, it is seen as one of the major building blocks to a safer cyberspace in future. The CS-AWARE solution will on the one hand, benefit from the information provided by those authorities and on the other hand, provide information about newly detected and unmatched incidents (like threat or attack patterns). It is assumed that with more and more tools that provide capabilities for organisations to participate in security related information sharing, the benefit of sharing information for the common good will become evident and encourage organisations to engage in cybersecurity related information sharing. Cybersecurity information exchange would in that case become one of the most important information sources for cybersecurity awareness and threat detection.

In order to deal with the expected language barriers and usability concerns in the context of European local public administrations, the main focus of the CS-AWARE project, *multi-lingual semantics support* will be part of this project's solution. Where relevant, security related information coming from within the end user organisations, or information from external information sources, will be automatically translated to benefit from the information of different cultural contexts.

The project includes two pilot scenarios in the LPA context: the municipalities of Larissa (Greece) and Rome (Italy). This set-up will allow us to develop tailored system and dependency analysis procedures for the LPA context. The project will commence with workshops in both municipalities. A representative cross section of the LPA's staffs will be formed in each LPA and will use SSM in a workshop setting, where the LPA's staff, facilitated by the project team can help create a detailed understanding of the problem domain and the system dependency analysis, together with security experts, legal experts and CERT representatives.

V. Case Study

The first step in applying the introduced approach was to determine the largest threats to LPA's based on expert knowledge and state-of-the-art research on the topics. This analysis was followed by evaluating the most valuable external, preferentially publicly available, information sources for cyber crime related data. This analysis was also based on expert knowledge and collected in a detailed report. Finally, with the specific information on potential threats and available external sources, the SSM workshops in the pilot cities were conducted.

A. Initial Threat Assessment

During the threat assessment we have determined that the main asset to be threatened from the cyber domain for local public administrations will most likely be the data that is managed by the administrations, including personal citizen and employee data. The main cybersecurity challenge in local public administrations is assumed to be the prevention of unauthorised data access, modification and destruction of those data. It was assessed that local public administrations are not a high valued target for potential threat actors, as for example critical infrastructures (potential large-scale disruption of economy) or financial institutions (potential high financial gain) are. However, there is a certain level of risk associated, since there are relevant threat actors that may have a vested interest in gaining unauthorised access to data managed in LPAs. We assume a low to medium level of risk against LPA managed data from the cyber domain. Additionally, we have identified that the most valued asset in LPAs is the potentially sensitive and/or private citizen and employee data that is managed by LPA systems, and that unauthorised data access, modification and destruction as well as data theft are the most relevant threats towards LPAs.

Table I shows the results of the initial analysis of potential threats and their risk level, based on the expert analysis and internationally acclaimed cybercrime threat reports [20] [21].

TABLE I. LPA RISKS GROUPED BY THREAT

Threat	*Risk level*		
	High	*Medium*	*Low*
Unauthorised data access, modification, destruction		X	
Data Theft		X	
Extortion	X		
Advanced Persistent Threat (APT)			X
Ransomware (untargeted)	X		
Ransomware (LPA specific)			X
Distributed Denial of Service - DDoS (untargeted)	X		
Distribtued Denial of Service - DDoS (LPA specific)			X
Web page defacement / shaming			X
Malware infection		X	

In Table II the most likely threat actors and their corresponding risk levels have been summarised. We assess that untargeted large-scale attacks with the goal of extortion, like Ransomware or Distributed Denial of Service (DDoS) attacks carry a higher risk for LPAs. We have identified the cybercriminal (high) as well as the malicious insider (medium) as the most relevant threat actors. Furthermore, disgruntled citizens, script kiddies and hacktivists are also seen as relevant threat actors, but we assess the risk from those actors to be low due to low potential pay-off for those actors as well as the low expected damages for LPAs.

TABLE II. LPA RISKS GROUPED BY THREAT ACTOR

Threat	*Risk level*		
	High	*Medium*	*Low*
Cyber criminal		X	
Malicious insider	X		
Disgruntled citizen / script kiddie			X
Hacktivist			X

B. Analysis of External Information Sources

As part of this initial analysis of cybersecurity relevant information sources, the main categories and respective sources which were identified can be seen in Table III.

The first possible information sources are related to organisations that the European Cybersecurity Strategy (European Commission and High Representative of the European Union, 2013) classifies as one pillar of coordination and information sharing efforts. While the CS-AWARE project does not expect as much cooperation with law enforcement as with NIS competent authorities due to the higher requirements for protecting information relating to cybercrime, we have identified several organisations that may be able to provide relevant information for CS-AWARE. For the open source data providers, we focused on sources that provide loosely structured information without a dedicated feed or data format, or if they provide a feed the provided data is usually utilised by aggregated data providers. The information sharing tools discussed are mainly community efforts to provide mechanisms for data aggregation. While data aggregation is already covered in the CS-AWARE solution, it is worth looking at those tools to see if it would help us to further simplify the data aggregation effort in CS-AWARE.

CS-AWARE will try to rely solely on free and open source data, it is however worth investigating which commercial data sources exist in case the free and open source data is not available. We may try to ask some of those companies for free access to their data in the context of this European research project. Overall, it seems that the sources listed provide up-to-date information, at least in the cases where the refresh time interval was stated explicitly or was easily verifiable (e.g., by accompanying timestamps). For retrieving data from the sources listed in this section some demo prototypes are available (mainly implemented in Python and Java), which were used by CS-AWARE partners for evaluation and testing the provided feeds. The idea of malware analysis tools is to be able to get a detailed report, listing the behaviour of a suspicious executable in a controlled environment (e.g., sandbox). In general we expect to collect fast but not in-depth reactions to currently ongoing security incidents from social media sources. While this information may lack the level of depth we expect from more security focused information sources, information collected from social media may help CS-AWARE to react to quickly evolving incidents. While one of the main public data sources provided by NIS competent authorities like CERTs is about vulnerabilities, it is still a good idea to have a look at the most well-known vulnerability trackers. For many years, the CVE list provides a standardised way of enumerating software vulnerabilities.

Our analysis concluded that the most valuable cybersecurity related information (or cybersecurity intelligence) for CS-AWARE can be found from both official organisations, for example NIS competent authorities or law enforcement organisations, as well as private efforts, for example for-profit companies or non-profit communities/ projects. More generalised data, not necessarily provided by the security community, can be found from social media or data visualisation focused data sources. For CS-AWARE we will focus on information that is freely available from either the NIS competent authorities, from companies that provide free data

TABLE III. EXTERNAL INFORMATION SOURCES

Topic	*Source*	*Link*
NIS Competent Authorities	European Union Agency for Network and Information Security (ENISA)	https://www.enisa.europa.eu
	European Public Private Partnership for Resilience (EP3R)	https://www.enisa.europa.eu/topics/national-cyber-security-strategies/ppps/public-private-partnership/european-public-private-partnership-\for-resilience-ep3r
	Computer Emergency Response Teams (CERTs)	https://cert.europa.eu
	Computer Security Incident Response Teams (CSIRTs)	http://www.cert.org/incident-management/national-csirts/national-csirts.cfm
Law Enforcement Agencies	Europol	https://www.europol.europa.eu/activities-services/services-support/intelligence-analysis/cyber-intelligence
	Interpol	https://www.interpol.int/Crime-areas/Cybercrime/Cybercrime
Cyber Intelligence Sources and Information Sharing Tools	Shadowserver	https://www.shadowserver.org/wiki/
	Abuse.ch	https://abuse.ch/
	Spamhaus	https://www.spamhaus.org/
	SANS Internet Storm Center	https://isc.sans.edu
Commercial Providers	Flashpoint	https://www.flashpoint-intel.com/solutions/
	Checkpoint	https://www.checkpoint.com/
	DCU Microsoft	https://news.microsoft.com/presskits/dcu/
	AbuseSA / Clarified Networks	https://www.clarifiednetworks.com
Cybersecurity Intelligence Data Feeds	AlienVault OTX	https://otx.alienvault.com/
	Advanced Cyber Defence Center	https://www.acdc-project.eu/
	Hail a Taxii	http://hailataxii.com
	Facebook Threat Exchange	https://developers.facebook.com/products/threat-exchange/
	Honey DB	https://riskdiscovery.com/honeydb/
Malware Analysis	Hybrid Analysis	https://www.hybrid-analysis.com/
	VirusBay	https://beta.virusbay.io/
	VirusTotal	https://www.virustotal.com/en/
Social Media	Xing	https://www.xing.com/
	Reddit	https://www.reddit.com/
	Facebook	https://www.facebook.com/
	Twitter	https://twitter.com/
	Google+	https://plus.google.com/
Vulnerability Data	CVE List	https://cve.mitre.org/cve/cna.html
	National Vulnerability Database (NVD)	https://nvd.nist.gov/
	CVE-SEARCH	https://www.cve-search.org/

or, probably most relevant, open source intelligence (OSINT) focused communities and projects. However, we will keep the option in mind to ask for-profit companies for access to their cybersecurity intelligence data in the context of this European project, if relevant.

C. Analysis of Pilot Scenarios

A crucial part for designing an effective cybersecurity awareness solution for local public administrations was to gain in-depth knowledge on LPA's, the services they provide, their inner workings and how similar these are across different city sizes and European countries. As of now, we have held the first round of SSM user workshops in the two piloting cities, Roma Capitale (RC) in Italy and Larissa, Greece. The main goal for this round of analysis was to gain an initial understanding of the complexities within LPAs and identify realistic and meaningful piloting scenario that can be managed with the resources available for this project. During the analysis we have met and even exceeded the expectations we set for our first round of analysis. In both piloting scenarios we have now a clear understanding of the critical assets and their dependencies to other critical assets that need to be taken into account, and we have identified how those assets can be monitored. By now we have conducted the first of three rounds of user workshops at our piloting partners. We have seen that if the participants of the user workshops have prepared themselves and have comprehended the added value of system analysis using rich pictures, this method is a powerful tool. It allows the participants to quickly gain a common understanding of the systems and interactions from a high level overview down to more detailed technical specifications. The right composition of participants in the user workshops is crucial. Only if representatives from all relevant organisational levels (such as managers and technicians) are present in the workshops, a complete and holistic understanding of the problem domain will be achieved. It have become clear that it is essential to have stable workshop groups – those who decide to be part of the workshop need to be there for the whole duration of the analysis. We argue that in complex systems good cybersecurity awareness can only be provided if the relevant relations between the mission critical aspects of the system are understood, and relevant case specific monitoring points can be utilised. The first round of analysis has only strengthened our argument. In both municipalities, we were able to achieve good analysis results and were able to identify the most mission critical systems and their dependencies, as well as potential monitoring points for CS-AWARE. The individual set-ups and procedures in the two municipalities differ significantly from each other, especially due to the substantial difference in complexity in the operations of the two very differently sized municipalities. Nevertheless, we were able to draw some generalised conclusions that will allow us to develop guidelines and procedures that will help to further simplify future analysis efforts in LPAs.

1) Municipality of Larissa, Greece: In general, the analyst team was satisfied with the outcome of the workshop, and that the Larissa team could be already released after three days of data gathering, after the required level of analysis detail had been achieved. The two main factors contributing to this result were the manageable complexity of systems in a mid-sized municipality, as well as the excellent preparedness of the team in Larissa. The team had familiarised themselves with the

CS-AWARE project ideas as well as with the analysis methodology, which allowed the analysis team to quickly achieve excellent results. As outcomes of the workshop the analysts concluded that the most interesting connection points for CS-AWARE are monitoring on the service level, the network level, as well as monitoring existing security mechanisms. At the service level, the analysts concluded that it was only necessary to concentrate on two mission critical services. These two systems both store and process personal and sensitive data and both are critical to the operation of the City. Therefore, the proposed CS-AWARE solution will seek to monitor these systems and networks. In both cases, it was established that activity could be recorded and saved to the database via built-in auditing mechanisms, meaning SQL queries could be used to capture audit information about data operations (although any personal data will need to be anonymised at source). Furthermore, similar data can be gathered from built-in database auditing mechanisms.

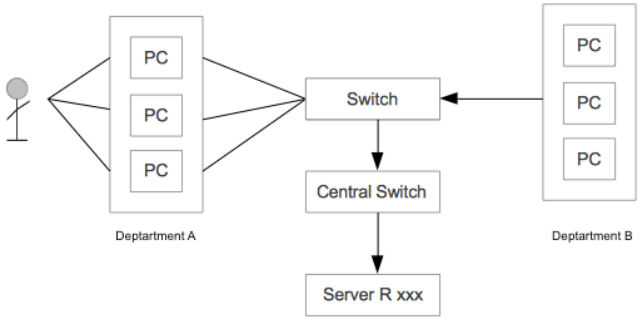

Figure 5. System Rich Picture

Figure 5 shows an anonymised representation of one of the rich pictures created by the team members of the Larissa LPA. As mentioned in Section II they used cartoon-style visualisations of the current situation, in this case part of their network infrastructure.

2) Municipality Roma Capitale, Italy: As expected, the Roma Capitale (RC) systems are much more complex than those that have been seen in the Municipality of Larissa, due to the extraordinary size of Roma and the number of online services that are provided to citizens and employees of RC. The attendees were divided into four groups, which were asked to draw a high level understanding of the systems and dependencies relating to their area of expertise, identifying mission critical systems as well as those parts of the systems that handle sensitive data. This resulted in four initial rich pictures, and while having a unique view on RCs systems, each included many aspects of other parts of the systems that other teams had been investigating in more detail. In the end we were able to identify a piloting scenario that will be possible to manage with the resources available within the CS-AWARE project: It was discussed to focus for now only on one relevant critical service - as well as all systems it depends on. It was identified that the most relevant critical dependencies can be found within the RC data centre (where the application service as well as the relevant application database are running), the web portal together with the identity and access management system (IAM), and several security appliances (like firewalls, proxies and SIEM (Security Information and Event Management) systems) that contain information relevant to service related operations. In the end, we were able to gain a good understanding of the overall architecture of RC systems and dependencies and a more detailed understanding of the system aspects that are the most relevant to CS-AWARE, identifying possible monitoring points for all relevant parts.

VI. Discussion

In Section I, four main points were mentioned by the authors to be essential in creating the introduced strategy of addressing complex situations in large infrastructures by use of a soft system thinking approach.

- *Identifying assets and their dependencies*
  Based on the results of the two SSM workshops, general assets and their dependencies could be identified and were grouped into four main categories: Network, Database, Service and Security-appliance level.
  The first question we asked the participants in both workshops was: "Which systems are mission critical and/or handle sensitive data?". Mission critical systems were different between the two LPAs, but shared common characteristics such as complementing infrastructure could be identified.
- *Identifying technological and socio-technical relations in the organisation*
  Next to identifying mission critical systems, the technical infrastructure and organisational structure in which these systems are used were determined. During the workshops, the socio-technical characteristics of the assets and the processes they are used in were determined.
- *Identifying external information sources*
  The external information sources used to complement the internal data collected by CS-AWARE were analysed extensively by experts and will be selected according to their relevance and quality of input they offer. Next to Social Media sites, such as Twitter and Reddit, Open Source Intelligence platforms and Commercial Providers, many other potential sources were identified and summarised in Table III.
- *Providing results in reusable form*
  Besides compiling a detailed deliverable on the results of the SSM analysis for external and internal information sources, the GraphingWiki mentioned in Section IV was used to produce a visual representation of the dependencies in the systems. Additionally to the graph, the tool produces a JSON formatted summary of the features of the system, which will be used for configuration purposes by the other components to specify the individual settings of each LPA implementation.

The exemplary dependency graph in Figure 6, depicting part of the LPA's system, shows how the components are linked to each other. Each of the components as well as the different relations between them have a Wiki-page where all relevant information is summarised. These include knowledge obtained in all aforementioned steps of the analysis - the pilot workshops as well as the external sources. An example for

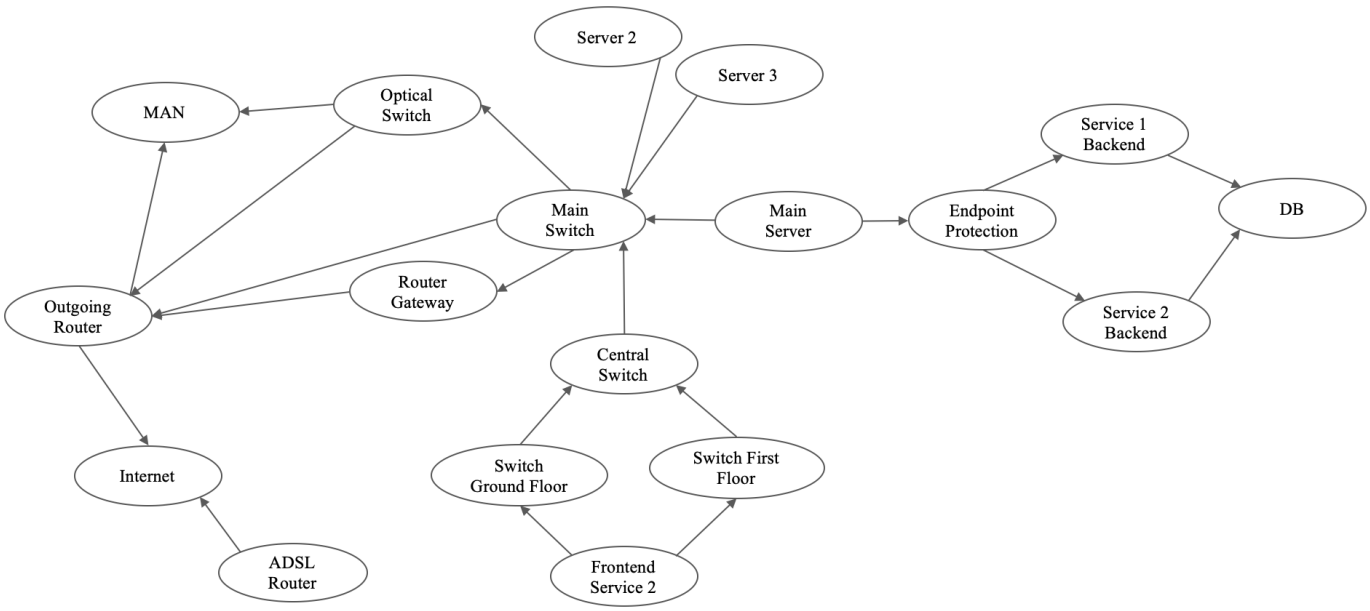

Figure 6. Larissa Dependency Graph

how such a Wiki-page can be structured is shown in Figure 7, including the individual categories selected for this project as well as the semantic text relevant for this component.

Router Gateway

Although now redundant, gateway router was still connected to the main Linksys switch and could be used as a monitoring point for network traffic. After further investigation it was established that almost all of the network traffic could be monitored through the gateway, and that most likely built-in monitoring mechanisms can be used to gather this data.

Observed

Outgoing-Router

?CategoryService1 ?CategoryService2 ?CategoryLarissa

Figure 7. Router Gateway Wikipage

All information collected during these workshops was summarised in the dependency graph and can be extracted in a JSON file to use in external applications. For the purpose of CS-AWARE, this will function as a basis for implementing and configuring the other components in the system. CS-AWARE combines multiple existing tool providers to a single, holistic cyber security awareness solution as can be seen in Figure 4. The System Dependency Analysis described and demonstrated in this paper builds the foundation on which the configurations of the other components depend on. It can specify, next to generalised configuration settings applicable for all LPAs, specific parameters for the individual LPA in question. For the deployment of CS-AWARE in any new LPA, generalised configuration settings can be extracted from the GraphingWiki, which then can be manually imported in the other components.

In both municipalities, we were able to achieve good analysis results and were able to identify the most mission critical systems and their dependencies, as well as potential monitoring points for CS-AWARE. While the individual set-ups and procedures in the two municipalities are significantly different from each other, especially due to the substantial difference in complexity in the operations of the two very differently sized municipalities, we were able to draw some generalised conclusions that will allow us to develop guidelines and procedures that will help to further simplify future analysis efforts in LPAs. In line with the initial risk assessment we have identified that the potentially sensitive and/or private data managed by LPAs is their most valuable asset. A cybersecurity awareness solution has to monitor the possible data flows in day-to-day operations. We have investigated potential monitoring points at 4 different levels that allow to identify suspicious behaviour related to data operations: The database level, the application/service level, the network level and the security appliance level.

The first steps of the SSM were applied during the user workshops in the municipalities - *entering the problem situation*, *expressing the problem situation* and *formulating the root definitions of the systems behavior*. The following steps will be undertaken in the upcoming workshop iterations in the pilots: *building conceptual models* and *comparing model to real-life situations*. This will allow for an even better understanding of the internal system and its information flows. Based on the received feedback *possible changes are defined* and the model revised accordingly. The final model will satisfy the last step of the SSM - *improving the problem sitation* by guiding implementation procedures in the respective municipalities. The Soft Systems Methodology approach provided usable results on which the further development of the CS-AWARE solution have been based on. While it was surprisingly easy to obtain relevant results in a short period of time in Larissa, the complexity of the Roman infrastructure required more extensive work time allocation. Nevertheless, both first Soft Systems Workshops were highly successful and we are eager

to continue deepening the knowledge on the respective systems in the upcoming second round.

VII. Conclusion and future work

In this paper, we have presented a system and dependency analysis methodology for complex systems based on soft systems thinking within the context of cybersecurity. The target for the analysis are organisations that rely on complex systems and procedures for their operation, like critical infrastructures, large organisations or SMEs or public institutions. The analysis methodology is focused on providing a holistic socio-technological view of the analysed system, based on the combination and visualisation of different relevant information sources. Since one of the greatest sources of information about a system is coming from its users, workshops where users from all organisational levels and with different backgrounds work together to define the problem situation are a central aspect of this methodology. We have argued that each organisational set-up is different, which makes generalised cybersecurity solutions difficult. We have shown that the presented system and dependency analysis methodology can be seen as an abstraction layer that allows to apply generalised cybersecurity solutions on top of it. As an example, we have presented the EU H2020 project CS-AWARE that utilises the presented system and dependency methodology as a central part of its cybersecurity solution. The goal of CS-AWARE is to develop an automated cybersecurity situational awareness and decision support solution relying on cooperative and collaborative approaches, as laid out by the NIS directive. The case study presented in this paper applied the introduced Soft Systems Methodology to conduct an initial risk assessment, identify potential external sources as well as hold the first round of SSM workshops in the pilot municipalities.

We have been quite happy with the results of the first round of system and dependency analysis workshops. In some aspects we achieved much better results than we had expected, quickly identifying the four main levels requiring our attention: database, application/service, network and security appliance level. In other aspects it took a bit longer than expected to gain a common understanding of the workshop goals, before achieving the expected results. Based on the experiences we have gained so far, we are confident that we have chosen the right approach for CS-AWARE and with some tweaks to accommodate for individual cultural aspects, we expect even better results during the second round of workshops. Based on the analysis results described in Figure 6, more detailed tacit knowledge of the participants will be obtained regarding the socio-technical and infrastructural aspects of the LPA internal systems.

Acknowledgements

We would like to thank the EU H2020 project CS-AWARE ("A cybersecurity situational awareness and information sharing solution for local public administrations based on advanced big data analysis", project number 740723) and the Austrian national KIRAS project CERBERUS ("Cross Sectoral Risk Management for Object Protection of Critical Infrastructures", project number 854766) for supporting this work. The Biomimetics and Intelligent Systems Group (BISG) would like to acknowledge the support of Infotech Oulu.

References

[1] T. Schaberreiter, C. C. Wills, G. Quirchmayr, and J. Röning, "Addressing complex problem situations in critical infrastructures using soft systems analysis: The cs-aware approach," *in Proc. SECURWARE*, pp. 99–105, 2017.

[2] European Commission, "Proposal for a directive of the european parliament and of the council concerning measures to ensure a high common level of network and information security across the union," COM(2013) 48 final, 2013.

[3] European Commission and High Representative of the European Union for Foreign Affairs and Security Policy, "Cybersecurity strategy of the european union: An open, safe and secure cyberspace," JOIN(2013) 1 final, 2013.

[4] ENISA, "National cyber security strategies in the world," accessed: 2018-011-13. [Online]. Available: https://www.enisa.europa.eu/topics/national-cyber-security-strategies/ncss-map

[5] ISO/IEC 27000:2016, "Information technology — security techniques — information security management systems — overview and vocabulary," ISO/IEC, Standard, 2016.

[6] ISO/IEC 29100:2011, "Information technology — security techniques — privacy framework," ISO/IEC, Standard, 2011.

[7] CEN, CENELEC and ETSI, "Focus Group on Cybersecurity (CSCG)," accessed: 2018-011-13. [Online]. Available: http://www.cencenelec.eu/standards/sectors/defencesecurityprivacy/security/pages/cybersecurity.aspx

[8] P. B. Checkland, *Systems Thinking, Systems Practice*. John Wiley & Sons Ltd. 1981, 1998.

[9] P. B. Checkland and J. Scholes, *Systems Thinking, Systems Practice*. John Wiley & Sons Ltd., 1991.

[10] S. Robinson, "Conceptual modelling for simulation part ii: a framework for conceptual modelling," *Journal of the Operational Research Society*, vol. 59, no. 3, pp. 291–304, 2008.

[11] V. Ferretti, "From stakeholders analysis to cognitive mapping and multi-attribute value theory: An integrated approach for policy support," *European Journal of Operational Research*, vol. 253, no. 2, pp. 524–541, 2016.

[12] T. Maqsood, A. D. Finegan, and D. H. T. Walker, "Five case studies applying soft systems methodology to knowledge management," in *7th Annual Conference on Systems Engineering Research*, 2009, p. 18.

[13] C. H. Antunes, L. Dias, G. Dantas, J. Mathias, and L. Zamboni, "An application of soft systems methodology in the evaluation of policies and incentive actions to promote technological innovations in the electricity sector," *Energy Procedia*, vol. 106, pp. 258 – 278, 2016.

[14] J. Eronen and M. Laakso, "A case for protocol dependency," in *First IEEE International Workshop on Critical Infrastructure Protection (IWCIP'05)*, Nov 2005, p. 9.

[15] J. Eronen and J. Röning, "Graphingwiki – a semantic wiki extension for visualising and inferring protocol dependency," in *First Workshop on Semantic Wikis – From Wiki To Semantics*, 2006.

[16] J. Eronen et al., "Software vulnerability vs. critical infrastructure – a case study of antivirus software," *International Journal on Advances in Security*, vol. 2, no. 1, pp. 72–89, 2009.

[17] P. Pietikainen, K. Karjalainen, J. Roning, and J. Eronen, "Socio-technical security assessment of a voip system," in *2010 Fourth International Conference on Emerging Security Information, Systems and Technologies*, July 2010, pp. 141–147.

[18] T. Schaberreiter, K. Kittilä, K. Halunen, J. Röning, and D. Khadraoui, *Risk Assessment in Critical Infrastructure Security Modelling Based on Dependency Analysis*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 213–217.

[19] G. Backfried et al., "Open source intelligence in disaster management," in *2012 European Intelligence and Security Informatics Conference*, Aug 2012, pp. 254–258.

[20] European Union Agency for Law Enforcement Cooperation, "Internet organised crime threat assessment 2017," accessed: 2018-011-13. [Online]. Available: https://www.europol.europa.eu/activities-services/main-reports/internet-organised-crime-threat-assessment-iocta-2017

[21] European Union Agency for Network and Information Security, "Threat landscape report 2016," accessed: 2018-011-13. [Online]. Available: https://www.enisa.europa.eu/publications/enisa-threat-landscape-report-2016

A Survey on Microservice Security–Trends in Architecture, Privacy and Standardization on Cloud Computing Environments

Luciano de Aguiar Monteiro[1], Washington Henrique Carvalho Almeida[1], Raphael Rodrigues Hazin[1], Anderson Cavalcanti de Lima[1], Sahra Karolina Gomes e Silva[2] and Felipe Silva Ferraz[1]
[1]Center of Advanced Studies and Systems of Recife
Recife, Brazil
[2]Nassau Mauritius College
Teresina, Brazil
E-mail: {lucianoaguiarthe, washington.hc.almeida, raphaelhazin, andclima, sahrask}@gmail.com
E-mail: fsf@cesar.org.br

***Abstract* — Microservices have been adopted as a natural solution for the replacement of monolithic systems. Some technologies and standards have been adopted for the development of microservices in the cloud environment. Application Programming Interface and Representational State Transfer were used on a large scale for the implementation. The purpose of the present work is to carry out a bibliographic survey on the microservice security trends focusing mainly on architecture, privacy and standardization aspects in Cloud Computing environments. This paper presents a bundle of elements that must be considered for the construction of solutions based on microservices.**

Keywords- Microservice; Security; Cloud; Architecture; API; Monolithic

I. INTRODUCTION

Migration of the monolithic architecture to the cloud has been a major problem. In this paper a research was carried out on the topic of microservices that have been adopted as a natural solution in the replacement of monolithic systems. The main question lies in how its architecture has been used and issues of security and privacy keys on a Cloud Computing environment. Cloud Computing provides a centralized pool of configurable computing resources and computing outsourcing mechanisms that enable different computing services to different people in a way similar to utility-based systems such as electricity, water, and sewage.

The motivation for this collection was the fact that more and more microservices have been found as a solution for cloud applications. This paper analysis in further details aspects related to Survey on Microservice Architecture [1]. Due to its architecture, a concern about security issues is fundamental, unlike a monolithic architecture where security is implemented in physical barriers and limiting access to resources, the microservice architecture has its main characteristic in interoperability, reuse and scalability. The purpose of this paper is to compile security issues in microservices, as shown in the following sections.

For the recent advances of Cloud Computing technologies, the use of microservices on applications has been more widely addressed due to the rich set of features in such architecture. These are cloud-based applications that make users use it at low cost, threshold, and risk. Therefore, their practical use in business can be expected as a trend for the next generation of business applications [2].

Scaling monolithic applications is a challenge because they commonly offer a lot of services. Some of them are more popular than others. If popular services need to be scaled because they are highly demanded, the whole set of services will also be scaled at the same time, which implies that unpopular services will consume a large amount of server resources even when they are not going to be used [3].

The microservice-based architecture has emerged to simplify this reality and is a natural evolution to application models.

Microservices are a software oriented entity, which have the following features [4]:

Isolation from other microservices, as well as from the execution environment based on a virtualized container;

Autonomy – microservices can be deployed, destroyed, moved or duplicated independently. Thus, microservices cannot be bound to any local resource because microservice environment can create more than one instance of the same microservice;

Open and standardized interface that describes all specific goals with effectiveness, efficiency and available communication methods (either Application Programming Interface (API) or Graphical User Interface (GUI));

Microservice is fine-grained – each microservice should handle its own task.

Microservice architecture does not make an application any simpler, it only distributes the application logic into multiple smaller components, resulting in a much more complex network interaction model between components. When a real-world application is decomposed, it can easily create hundreds of microservices [5]. For this reason, this paper presents basic principles for the implementation of microservices aimed at classic security aspects for commercial applications. Organization of applications based on these standards mitigates common security issues.

The microservice architecture is a cloud application design pattern that implies that the application be divided into a number of small independent services, each of which is responsible for implementing a certain feature, as noted in Figure 1.

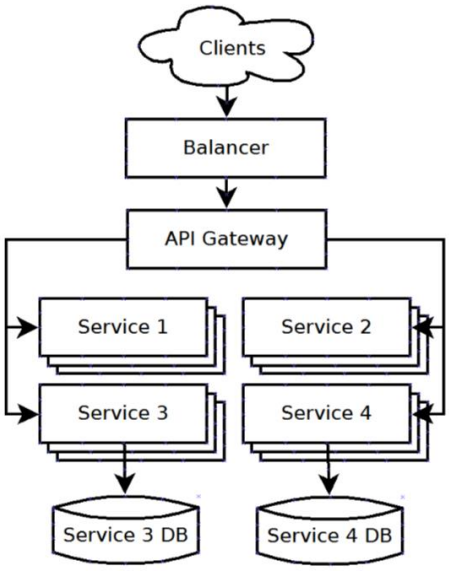

Figure 1. Microservice system architecture [4].

Microservices can be considered meta-processes in a Meta Operating System; they are independent, they can communicate with each other using messages and they can be duplicated, suspended or moved to any computational resource and so on [4]. Meta-modeling process is a type of modeling for analysis and modeling applicable to some known problems. Meta process modeling supports the effort of creating flexible process models.

The adopted methodology for this paper included a research in IEEE Xplore Digital Library, ACM Digital Library and Web of Science sources to provide all necessary information through published works. The strings Microservice AND security; Microservice AND Privacy; Microservice AND Cloud Computing was used to identify these works.

The remainder of this article is structured as follows: Section II is an overview of microservice in research topics; Section III presents security in Cloud Computing environment and Section IV shows the privacy model adopted in cloud applications for microservices and we present the standards of cloud environment and then conclude and summarize all results of that exercise in Section V.

II. MICROSERVICE

A. State of the Art

The microservice architecture was first approached in May 2011 at the workshop of software architecture [6], and since then it has been evolving and being adopted and implemented in Cloud Computing servers like Amazon AWS, Google Cloud and Azure.

From the technological perspective, early microservice applications were strongly influenced by a new generation of software development, deployment, and management tools [6]. Figure 2 show timeline with the technologies that drove the microservice architecture. The use of Linux Containers (LXC) was the first widely used container technique. It uses kernel namespaces to provide resource isolation [7], until the service mesh which build on sidecar technologies to provide a fully integrated service-to-service communication monitoring [6].

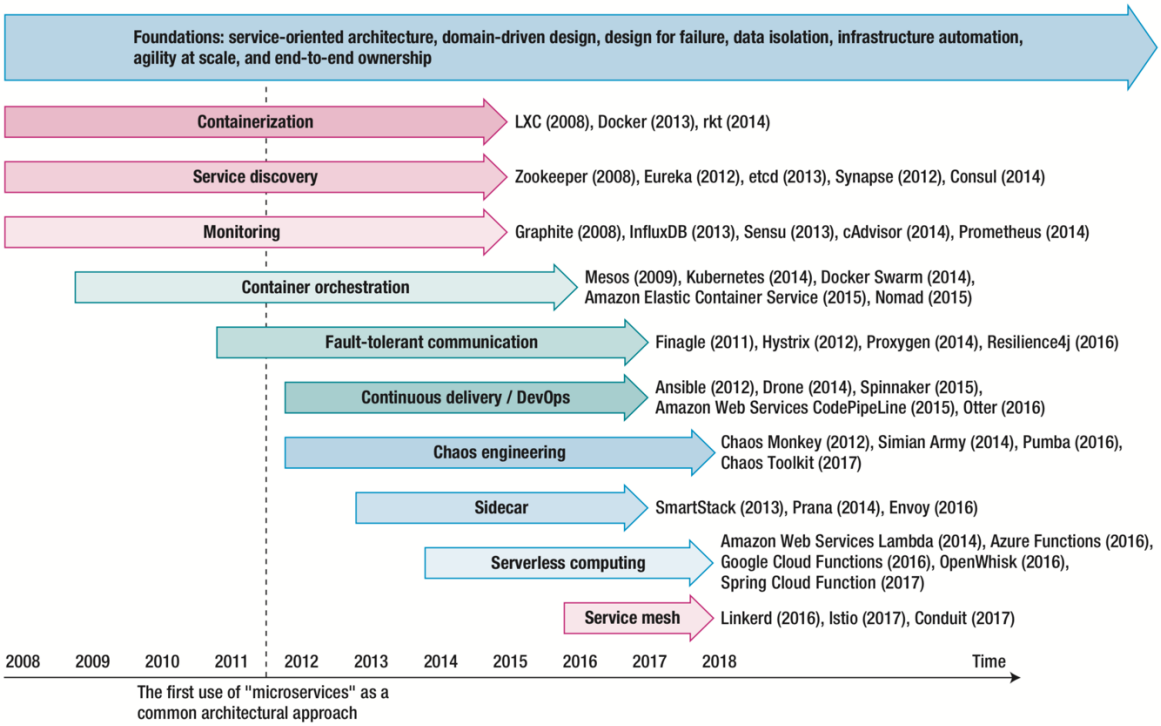

Figure 2. A microservice technologies timeline [6].

State of the art in microservices focuses its searches on the question of architecture analysis, performance, maintenance and security [5]. With the migration to continuous delivery culture, the organization of companies has changed from long processes which sometimes went through different areas to smaller teams, where each is responsible for developing their own microservice and for providing an API that will be used by other teams. On this point, the DevOps culture ends up being the most commonly used.

Information security according to ISO/IEC 27001 standards is based on three principles of confidentiality, integrity and availability. The microservices implementation is heavily bogged down in the irrevocability guarantee of these principles, for that reason some measures must be taken due to its complexity, the architecture proposed in this work treats numerous benefits for the guarantee of information security.

The development of solutions based on microservices has naturally used Cloud Computing environments so as to make the most of the best characteristics and functionalities provided by various solutions in the market. In this study we identified 3 relevant topics: the question of granularity, the deployment process and the resulting patterns.

A Microservice Architecture is a way of architecting software applications as independently deployable services. Based on Fowler, microservices can be characterized by a number of principles [8]:

- organization around business capability
- evolutionary design
- deployment / infrastructure automation
- intelligence in the endpoints
- heterogeneity and decentralized control
- decentralized control of data
- design for failure

The aforementioned principles are fundamental in the architecture that will be better described in the next section. The implementation of microservices is based on trade-offs between security and performance. This research found that the implementation of microservices uses the most advanced resources from Cloud Computing. The main characteristics of Cloud Computing can be summarized in the following points [9]:

- *Multi-Tenancy*

Refers to having more than one occupant of the cloud living and sharing with other occupants of the provider's infrastructures, including computational resources, storage, services, and applications. Through multi-tenancy, clouds provide simultaneous, secure hosting of services for various customers using the same cloud infrastructure resources. It is an exclusive characteristic to resource sharing on clouds.

- *Elasticity*

Another important aspect of Cloud Computing implies that the user is able to scale up or down resources assigned to services or resources based on the current demand. For providers, scaling up and down of a tenant's resources gives a prospect to other tenants to use the tenant's previously assigned resources.

- *Availability of Information*

Service Level Agreement (SLA) is a trust bond between cloud provider and customer. It defines a maximum time for which the network resources or applications may not be available for use by the customer. Due to the complex nature of customer demands, a simple measure and trigger process may not work for SLA enforcement.

- *Multiple Stakeholders*

In a Cloud Computing model, there are different stakeholders involved: cloud provider (an entity that delivers infrastructures to the cloud's customers), service provider (an entity that uses the cloud infrastructure to deliver applications/services to end users), and customer (an entity that uses services hosted on the cloud infrastructure).

Another important characteristic is the deployment of microservices. A cloud deployment model signifies a specific type of Cloud Computing environment, renowned by ownership, size, and access. There are three common cloud deployment models, namely private cloud, public cloud, and hybrid cloud [9].

Figure 3 shows differences between two architectures and demonstrates the microservices implantation. Their independence and granularity can be provided in several infrastructures.

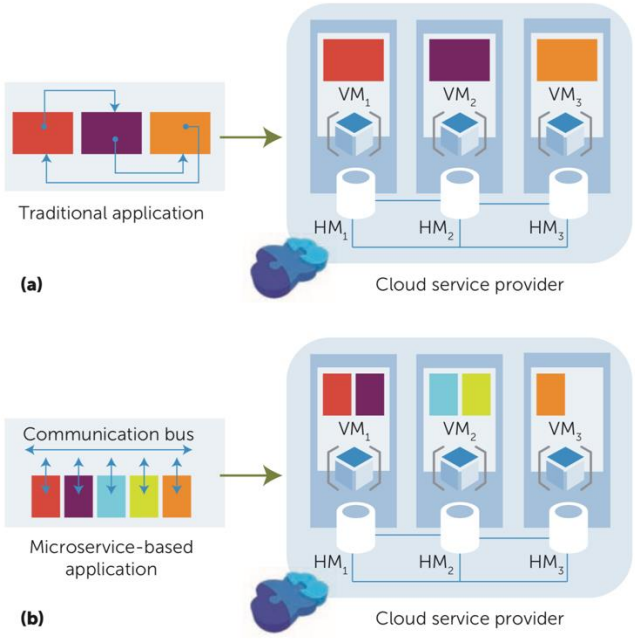

Figure 3. Software deployment in a cloud platform using (a) conventional and (b) microservice-based software [10].

Microservices are relatively small and autonomous services deployed independently, with a single and clearly-defined purpose. Because of their independent deployment, they have a lot of advantages. They can be developed in different programming languages, they can scale independently from other services and they can be deployed in the hardware that best suits their needs [10].

Moreover, because of their size, they are easier to maintain and more fault tolerant since the failure of one service will not break the whole system the way it could happen in a monolithic system [10].

Another characteristic of microservices is cloud native applications, the support of the IDEAL properties: Isolation of state, Distribution, Elasticity, Automated management and Loose Coupling. Microservices propose to vertically decompose the applications into a subset of business-driven services. Every service can be developed, deployed and tested independently by different development teams, and by means of different technology stacks. The responsibility of the development of a microservice belongs only to one team, who is in charge of the whole development process, including deploying, operating and upgrading the service when needed [10]. Figure 4 shows the complexity related.

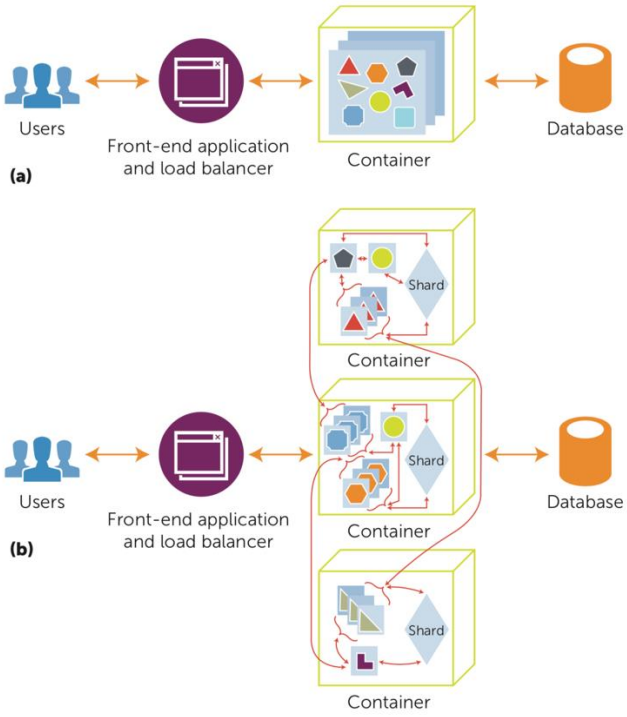

Figure 4. Architectural complexity of (a) monolithic and (b) microservice-based software [10].

Decoupling applications in this manner yields several benefits: it simplifies scaling (each service can be scaled independently), provides greater flexibility in resource allocation and scheduling, allows greater code reuse, enables new fault tolerant mechanisms, provides better modularity, and allows application developers to take advantage of services from other provides e.g., Amazon S3. As a result, this architecture has been widely adopted by both startups and large established companies (e.g., Uber and Netflix), and is being deployed at significant scale (e.g., Uber's application is composed of over 1000 microservices) [12].

The use of microservices can reduce the operational costs, as shown in the study [10]. The comparison was made in a cloud and monolithic solutions environment.

1) Cost comparison

In the study carried out in paper [13], it is shown a cost comparison in the various commercially used architectures of software development. In summary, use of microservices brings lower infrastructure spending by allowing scalability as well as scalability since the measurement of operating cost is done by use. In the old monolithic architecture many resources end up being loaded to memory even without being used, this is one of the great differences for the strong diffusion of this new architecture in software industry.

Given that each architecture was deployed in different infrastructures, we defined and calculated the metric Cost per Million of Requests (CMR) for each architecture in the three scenarios, in order to easily compare their execution costs. For each scenario and architecture, this metric was calculated by dividing the monthly infrastructure costs by the number of requests supported per month, which is calculated by multiplying the number of requests supported per minute by 43,200 —the number of minutes per month (60*24*30)—. We assumed a constant throughput per minute during a month [13]. The CMR metric for each architecture is shown in Figure 5.

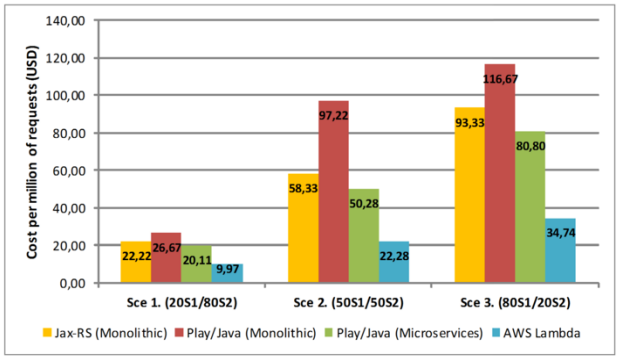

Figure 5. Cost Comparison of The Three Architectures per Million of Requests [13].

2) Granularity

Microservices can be declared with varying levels of capability, and the size of this functionality is typically referred to as its granularity, that is, the functional complexity coded in a service or number of use cases implemented by a microservice. Since microservices are discrete and must be composed into greater functional entities to support business workflows, it follows that message passing between microservices (as a result of method invocation) increases as the microservices become finer-grained.

The 'building-block' approach to service composition is attractive from an architectural perspective; arguments for service reuse can be made, and the gap between application design and the user requirements documentation can be reduced. However, the increase in communication between services (manifesting as out-of-process calls and the number of service calls made) also increases the response time of an application, particularly when many small increases in latency are compounded together [7].

Container-based technologies, and in particular its best known implementation Docker, made deployment of our new application possible through several characteristics [12]:

- A Docker image contains all of its dependencies, which means a given service can be treated as a black box, only exposing its API in exchange for resources.
- The containers are by default sealed from one to another, which results in guaranteed low coupling, without the high cost associated with virtual machines.
- Docker Compose made it possible to easily deploy any number of services, by composing in a text file an application made of several services.
- Docker Swarm mode allows for complete decoupling of the containers and the machines supporting them. In its recent version 3, Docker Compose allows for Distributed Application Bundles, which define applications made of several services without any dependence other than the presence of a Docker host IP address and access credentials.

Recommended patterns on how to compose microservices together [8]:

1. *Aggregator Microservice Design Pattern* – e.g., a service invoking others to retrieve / process data.
2. *Proxy Microservice Design Pattern* – a variation of the Aggregator with no aggregation.
3. *Chained Microservice Design Pattern* – produces a single consolidated response to a request.
4. *Branch Microservice Design Pattern* – extends the Aggregator and allows simultaneous response processing from possibly mutually exclusive chains of microservices.
5. *Shared Data Microservice Design Pattern* – towards autonomy through full-stack services with control of all components.
6. *Asynchronous Messaging Microservice Design Pattern* – use message queues instead of Representational State Transfer (REST) request/response pattern.

Integration is another important feature. The architecture of microservices allows for better integration of corporations where there are areas that handle a number of business activities. As this pattern is based on the independence of technologies, the services made available can be developed without a change in technology, which is usually expensive. Throughout this study we present data that show how this pattern brought about significant improvements in the development of solutions for the software industry.

B. Architecture

Microservice architecture has become a dominant architectural style choice in the service-oriented software industry. Microservice is a style of architecture that puts the emphasis on dividing the system into small and lightweight services that are purposely built to perform a very cohesive business function, and is an evolution of the traditional service oriented architecture style [14], in which what is presents a scenario of microservice architecture (Figure 6) in which five services working independently provide requests of a mobile app through an API [15].

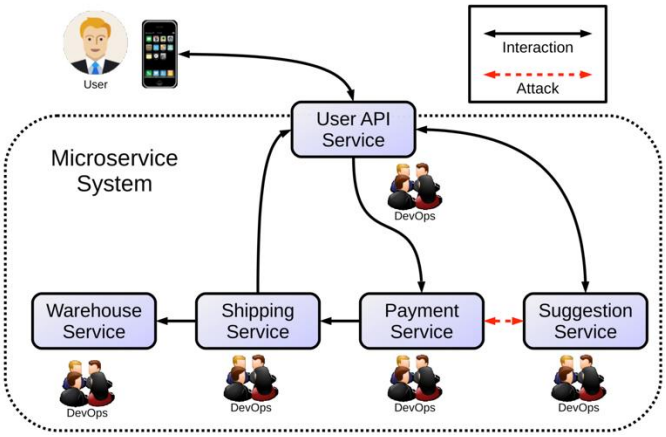

Figure 6. Example scenario of a microservice system[15].

The idea of splitting an application into a set of smaller and interconnected services (microservice) is currently attracting the interest of many application developers and service providers (e.g., Amazon [16][17], Netflix [12][18][19] and eBay [20][5]).

A Microservice based architecture has a pattern for development of distributed applications, where the application is composed of a number of smaller "independent" components; these components are small applications in themselves [21].

A microservice normally comprises three layers as a typical 3-tiered application[22], consisting of an interface layer [23], a business logic layer [20] and a data persistence layer, but within a much smaller bounded context. This sets a broad scope of the technical capabilities that a microservice could possess. However, not every microservice provides all capabilities. This would vary depending on how the function provided is meant to be consumed. For example, a microservice used primarily by providers of APIs would have a communications interface layer, business logic and data persistence layers but not necessarily have user interfaces [21].

We are considering a reference architecture model of microservices, demonstrating the main components and elements of this standard [21]. Table I presents a comparison between monolithic architecture and microservice architecture.

Table I. Comparing monolithic and microservice architecture [21].

Category	Monolithic Architecture	Microservice Architecture
Code	A single code base for the entire application.	Multiple code bases. Each microservice has its own code base.
Understandability	Often confusing and hard to maintain.	Much better readability and much easier to maintain.
Deployment	Complex deployments with maintenance windows and schedule downtimes.	Simple deployment as each microservice can be deployed individually, with minimal or zero downtime.

Category	Monolithic Architecture	Microservice Architecture
Language	Typically, entirely developed in one programming language.	Each microservice can be developed in a different programming language.
Scaling	Requires you to scale the entire application even though bottlenecks are localized.	Enables you to scale bottle-necked services without scaling the entire application.

In this paper, we will cover the following main elements:

1) API Proxy

To "de-couple" the microservice from its consumers, this proxy pattern is applied at the microservice interface level, regardless of the "API proxy" component. Organizations will provide APIs to different consumers, some of whom are within and others outside of the enterprise. These microservices would differ in SLA, security requirements, access levels, etc. [21].

2) Enterprise API Registry

The "discovery" requirements of the microservices are met through the use of the API registry service. Its purpose is to make the interfaces exposed by the microservice visible to consumers of the services both within and outside of the enterprise. An "Enterprise API registry" is a shared component across the enterprise, whose location must be well known and accessible. Its information content is published in a standard format, information should be in consistent and human readable format, and must have controlled access. It must have search and retrieval capabilities to allow users to look up details on available API specifications at design time [21].

3) Enterprise Microservice Repository

The "enterprise microservice repository" would be a shared repository for storing information about microservices. It provides information such as microservice lifecycle status, versions, business and development ownership, detailed information like its purpose, how it achieves the purpose, tools, technologies, architecture, the service it provides, any APIs it consumes, data persisted and queried and any specific non-functional requirements. In the absence of well-defined repository standards, the enterprise must define its own standard specification artefacts for microservices [5].

These elements are fundamental to the organized implementation of microservices and have been considered in this survey.

C. Microservice Standards and Solutions

In the centralized structure, the standardization becomes almost a natural way, but in microservices implementation this philosophy changes.

Traditional enterprise applications are divided into the front-end User Interface (UI), service-side logic components, and database. Front-end UI components run on user devices, such as web pages or mobile-side interfaces. Server-side logic components run on a server or in the cloud. The back-end database hosts application data. Server-side components work in conjunction with the database to handle requests issued by users [24].

Teams building microservices prefer a different approach to standards too. Rather than using a set of defined standards, written down somewhere on paper, they prefer the idea of producing useful tools that other developers can use to solve problems similar to the ones they are facing. These tools are usually harvested from implementations and shared with a wider group, sometimes, but not exclusively. Using a git and github has become the *de facto* version control system of choice. Open source practices are becoming more and more common in-house [25].

A microservice is an application on its own to perform the functions required. It evolves independently and can choose its own architecture, technology, platform, and can be managed, deployed and scaled independently with its own release lifecycle and development methodology. This approach takes away the construct of the Service-Oriented Architecture (SOA) and Enterprise Service Bus (ESB)and the accompanying challenges by making "smart endpoints" and treating the intermediate layers as network resources whose function is that of data transfer [21].

Unlike SOA, microservices do not have integration components responsible for service orchestration and prefer choreography. Business processes are embedded in services and there is no logic in the integration. Thus, Microservices themselves are responsible for the interaction with others. This gives limited flexibility to design or adjust business processes over the company's IT. It is a payoff for microservice independent service management. However, Netflix considers even the option to orchestrate microservices, which is not a mainstream path [26].

The applications that expose interfaces that can be used by other applications to interact with are defined as API [5]. Microservice APIs which are built using internet communication protocols like HTTP adhere to open standards like REST [27][28] and SOAP [3] and use data exchange technologies like XML [29] and JavaScript Object Notation (JSON) [5].

Applications developed in a monolithic architecture perform multiple functions such as providing address validation, product catalogue, customer credit check, etc. When using the microservice based architecture pattern, applications are created for specific functions, such as address validation, customer credit check and online ordering; these applications are cobbled together to provide the entire capability for the proposed service. The approach to application development based on microservice architecture addresses the challenges of "monolithic" application and services [21].

In the research undertaken in this paper, the microservices are implemented and documented as follows [14]:

1) Architectural views/diagrams

- UML
- Standard modeling languages, e.g., RAML and YAML.
- Specifically designed modeling languages, e.g., CAMLE.

- Standard specification languages, e.g., Javascript (Node.js), JSON and Ruby.
- Specifically designed specification languages, e.g., Jolie.
- Pseudocode for algorithms.

2) REST

Representational State Transfer (REST) consisting of a set of architectural principles that, when followed, allows a well-defined interface design to be created. Applications that use REST principles are called RESTFul. REST [5][29][28][30] is often applied to provide services to other services (web services) and to the same full use of messages. To better understand the architectural style, it is important to highlight three important concepts: (i) feature; (ii) operations and (iii) representations. Resource is any information that is made available to customers through a Unique Identifier (URI). We can also define resource as being the source of representations. The representations are a set of data that explains the state of the requested resource. URIs must have a notation pattern, be descriptive, and have a previously defined hierarchy. The same resource can be identified by one or more URIs, but a URI [31], [32] identifies only one resource.

3) API

API is a basic authentication, including API user registration with strong password protection, (b) modern security mechanisms such as message level security, web signature and web encryption, and (c) security mechanisms within API and its backend services as a third security factor such as token based API for backend authentication, public key infrastructure and transport layer handshake protocol [23].

REST APIs [18] are developed in many technologies and microservices developed using different types of programming languages (Java, .NET, PHP, Ruby, Phyton, Scala, NodeJs etc.) and persistent technologies (SQL, No-SQL, etc.) [3][8][33]. They can be managed and exposed to web clients, who can then access the microservices and receive their responses through a "*livequery*" mechanism whereby updates to database data are instantly communicated to subscribing clients [29]. Figure 7 best presents categories of practices for designing REST-based web services.

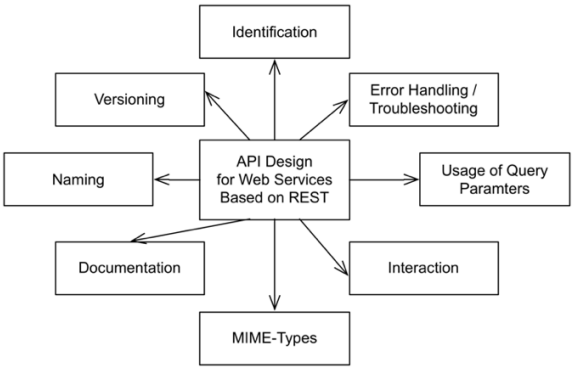

Figure 7. Categories of best practices for designing REST [34].

NoSQL databases are used in these implementations [29][35][36][37]. The NoSQL nature of the database is essential for providing the scaling, sharing and replication functionality expected from modern architectures, as well as to better support hierarchical data required for collaborative document editing [29].

The popularity of microservice-based architecture is evident from the report by the popular jobs portal *indeed.com*, in which the number of job openings on microservices-related technologies such as JSON [5][38][32] and REST [3][29][28] has grown more than 100 times in the last six years, whereas jobs in similar technology areas like SOAP and XML have remained nearly identical [5].

Solutions for microservices seek to implement simple algorithms that meet specific needs with the elements presented in this section. security on Cloud Computing

III. SECURITY ON CLOUD COMPUTING MICROSERVICES

Switching from a monolithic or centralized architecture to a decentralized architecture requires some care. In the past, security was focused on a single point [15], responsible for receiving all service requests. In the microservice-based architecture, the resources are offered through several points of access that interconnect each other, forming a unique solution.

Microservices combined with secure containers can facilitate new ways of building critical applications. These applications will benefit from tools and services built for less critical software. Secure containers and compiler extensions can help address more stringent requirements of critical applications. Although this approach is sufficient for implementing fail-stop applications, there are still several open research questions regarding whether and how it might support fail-operational applications [40].

Monolithic security services are relatively easier to implement than microservices. Monolithic services have a clear boundary and encapsulate their intercommunications. This will obscure security vulnerabilities [41][42] within the inner layers of the system. A microservice also encapsulates its communications. Both microservices and services are based upon clear requirements.

In a microservice-based system a simple routine completion requires the microservices to communicate with each other over network, for example. This will expose more data and information (endpoints) about the system and thus it expands the attack surface [19]. Some care must be taken in the communication between other services in the same network, and this is one of the major challenges [23][43][29] in this approach.

Monolithic applications, as previously explained, have a single and shared code base where all the developers work together. This development methodology has a few downsides, as it needs to struggle with handling cases in which the number of users exceeds the capacity of the server and it is hard to manage and maintain due to the lack of mechanisms aimed at modularization [26]. The evolution of the development of an application in monolithic architecture

becomes quite complex, considering that in order to add new functionalities, one must change the source code, and still considering the same reasons, making the software hard to maintain. Monolithic architectures are typically difficult to deploy, difficult to upgrade and maintain and difficult to understand [27].

The deployment of monolithic architecture applications in Cloud Computing environments causes a very negative impact: services need to be scaled because they are highly demanding. The whole set of services will also be scale at the same time, which generates unpopular services that consume a large amount of server resources even when they are not going to be used [3].

The organization of teams for the development of a system based on microservices is generally subdivided into teams and services, and these teams are generally responsible for the implementation and delivery of services. For this type of implementation, the teams have to be aligned in the purposes of the microservices and the interconnection between them, thus also synchronizing the protocol [44] used to carry out the communication, thus respecting a standard for access protection or improper interception. Defining the way services are interconnected and interacting is the key point of security [38].

The security challenge brought by such network complexity is the ever-increasing difficulty in debugging, monitoring, auditing and forensic analysis of the entire application [45]. Since microservices are often deployed in a cloud that the application owners do not control, it is difficult for them to construct a global view of the entire application [5].

In microservice architecture, an application is essentially a collection of workflows. These workflows can compose many levels of services, each processing and modifying the data before its final destination. What we need is a way to certify the metadata related to a data stream and manage its validity during time and re-elaboration [46].

Security is a major challenge that must be carefully thought of in microservices architecture. Services communicate with each other in various ways creating a trust relationship. For some systems, it is vital that a user is identified in all the chains of a service communication happening between microservices.

Microservices predominant execution environments are containers, that remove dependencies on the underlying infrastructure services, which reduces the complexity of dealing with those platforms [47], microservices need high availability and scalability characteristics provided by providers of Cloud Computing, environment preferably used by the developers. In this architecture the four security aspects that should be considered are: containers, data, permission, and network [48], as noted in Figure 8.

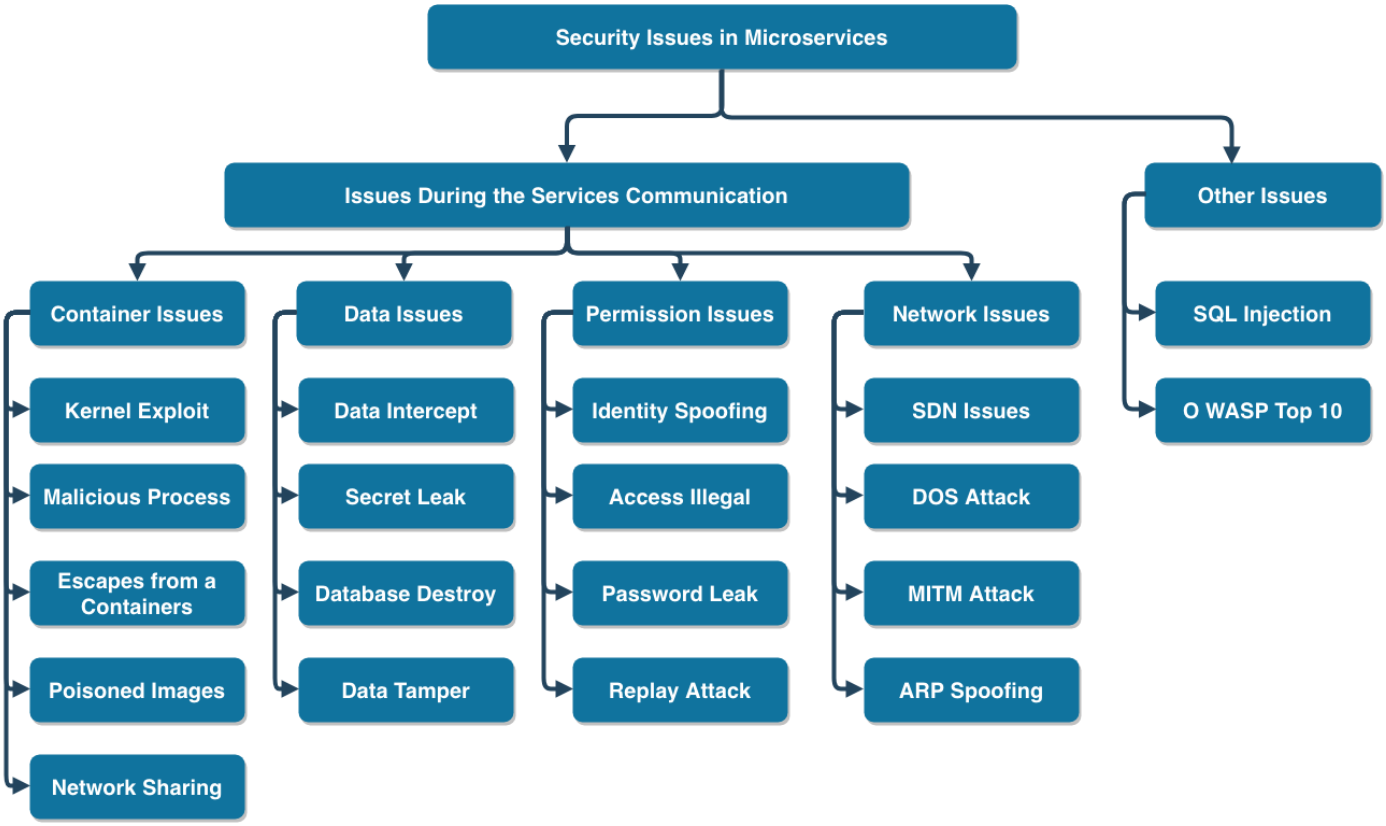

Figure 8. A taxonomy of security issues of Microservices [48]

Based on main security aspects of microservices discussed in Figure 6, the following main safety mechanisms will be presented to prevent safety deficiencies in microservices:

- mutual authentication of Services Using Mutual Transport Layer Security - with a self-hosted Public Key Infrastructure as a method to protect all internal service-to-service communication [49];
- host-authenticated TLS with in-band authentication are well-known solutions that are employed by designers to handle security challenges [14][15];
- principal propagation via Security Tokens: after a user has been authenticated by the gateway, the microservices behind it will be processing user's requests, a security token is created on the server side upon the successful validation of the clients credentials and given to the client for subsequent use [49].

Although the microservices are independent and do not cause dependencies among the modules, the biggest challenge nowadays is to guarantee availability [50]. The DevOps movement (set of practices to integrate the software development to IT operations) is currently collaborating with cloud environments and microservice architecture, providing continuous integration from the code compilation to the availability of the test and production environment, making it a facilitator for systems implementation utilizing microservices.

Ensuring the availability of services is presented as a security requirement facilitated by the use of the microservice architecture. This approach usually works by fragmenting the entire solution in smaller pieces [51]. Considering that these fragments are parts of the code with specific functions (microservices), in the event of a fragment failure, it would not result in the unavailability of all system resources. Availability has some critical points as they are bound to be observed, such as: implementing software versions, software crash recovery, invasions, unavailability of infra features beyond points.

In a microservice architecture, it is typical for many instances of a particular service to be running at any one time and for these instances to stop and start over time [52]. The problem of service discovery is to enable service consumers to locate service providers in real time to facilitate communication [53]. Docker Containers have been gaining a lot of hard work because of their agility and ease of making new services available [50]. The containers allow the microservices to be packaged [54] and available next to their dependencies in a single image, thus facilitating the availability of the service in a timely manner, minimizing downtime. This mode is called code portability [33]. In the context of microservices, the use of Docker containers for service delivery has resulted in benefits under various aspects such as automation, independence, portability and security, especially when considering ease of management, creation and continuous integration of environments systems offered by the Docker platform. In Docker, each container consists of only the application and the dependencies that the application needs to run, ideally no more and no less [33].

Another security concern involves the trust among the distributed microservices. An individual microservice may be compromised and controlled by an adversary. For example, the adversary may exploit vulnerability in a public facing microservice and escalate privilege on the virtual machine that the microservice runs in. As another example, insiders may abuse their privileges to control some microservices. As a result, individual microservices may not be trustworthy [5].

IV. PRIVACY ISSUES

Privacy has been a barrier for the adoption of Cloud Computing [51][55]. The migration to microservices has helped overcome this obstacle due to the scale gains proposed in this architecture.

In general, privacy refers to the condition or state of hiding the presence or view [56]. There is a need to attain this state in the places where confidential things are used such as data and files. In cloud data storage privacy is needed to attain the data, user identity and controls [57].

Trust is a crucial factor in Cloud Computing environments in current practice. It depends mostly on observation of characteristics, and self-evaluation of cloud service vendors. Existing trust mechanisms in the cloud are characteristics-based trust, SLA confirmation-based trust, Cloud transparency techniques, Trust as a service, Formal endorsement, audit and standards. In order to attain the service, it requires to be used in blend with social and technological mechanisms for providing persistent trust [58].

The exchange of sensitive data is intense in large-scale scenarios of Cloud Computing, with several federations, where multiple Identity Providers and Service Providers work together to provide services. Therefore, identity management should provide models and privacy mechanisms in order to manage the sensitive data of its users [43].

For service provider's standards in Cloud Computing environment, the contract is usually based on good deeds of that company, and hence the users need to pay attention to the security requirements, contract terms and other credentials. The users must have clear understanding in detailed terms and conditions of service providers and also the risk involved in signing the service provider's contract before moving to cloud [58].

Cloud service provides various options to the business customers to choose the level of protection needed for their data. The most common of these approaches is encryption. The customer chooses the type of encryption that they prefer and store the encryption key in a safe place under their control [44].

To ensure privacy, a well referenced model is used. This model is presented in Figure 9.

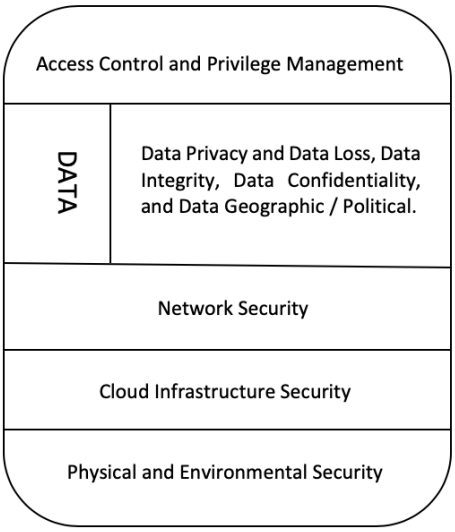

Figure 9. Cloud security and privacy model [55].

According to the proposed model in [55], a secure and private cloud model is divided into five layers: Physical and Environmental Security, Cloud Infrastructure Security, Network Security, Data and Access Control and Privilege Management:

A. Physical and Environmental Security

Layer of policies adopted with the objective of protecting physical access to the cloud provider [16]. Another benefit of cloud service is the ability to meet the elasticity of demand. Business processes should consider the availability of open-ended resources at an affordable cost. Use of services such as Software as Service (SaaS) enables the business to focus more on their core strengths. Since availability of computing resource is no longer a constraint, the business should take advantage of computing power to experiment with new ideas to serve the customers better. Since the cost is usage based, changing business processes to take advantage of newer technologies is advantageous to a business. Cloud service addresses an important business process for every business, namely backup and recovery. Many businesses do not pay enough attention to data backup and recovery because it is time consuming and does not provide immediate benefit until some disaster strikes, which is rare. With the cloud taking care of all the management aspects of data backup and recovery, businesses tend to focus on their strengths and a cloud provides the essential service of backup and recovery when needed. A common perception is that in order to provide security the user must have control over the devices. This usually applies to physical security. Given the elastic nature of demand for service and the centralization of service, the cloud environment is in a better position to provide greater physical security to the hardware [44].

B. Cloud Infrastructure Security

Addresses issues with cloud infrastructure security, but specifically with the virtualization environment [59]. Above this, the combination of software layers, the virtualization layer and the management layer allow for the effective management of servers. Virtualization is a critical element of cloud implementations and is used to provide the essential cloud characteristics of location independence, resource pooling and rapid elasticity. Differing from traditional network topologies such as client–server, Cloud Computing is able to offer robustness and alleviate traffic congestion issues. The management layer is able to monitor traffic and respond to peaks or drops with the creation of new servers or the destruction of unnecessary ones. The management layer has the additional ability of being able to implement security monitoring and rules throughout the cloud [60].

C. Network Security

Specifies the medium to which the end user connects to the cloud, comprising browsers and their connection [20]. The client-server in the cloud is accomplished by a client sending a request to the server and waiting for the result, where the server performs the computational process. A connection medium between a client and cloud service provider is the Web browser which relates to the cloud system. As discussed before, a client sends a request and needs to validate it on its own to check the authority of the user on the cloud system. Client credentials are signed by using Extensible Markup Language signature to authenticate and Extensible Markup Language (XML) encryption to encrypt the Simple Object Access Protocol (SOAP) messages [58].

D. Data

Layers cover data privacy, integrity, confidentiality, and geographic location [46]. To prevent data loss in cloud different security measures can be adopted. One of the most important measures is to maintain backup of all data in cloud which can be accessed in case of data loss. However, data backup must also be protected to maintain the security properties of data such as integrity and confidentiality. Different data loss prevention mechanisms have been proposed in research and academics for the prevention of data loss in network, processing, and storage. Many companies, including Symantec, McAfee, and Cisco, have also developed solutions to implement data loss prevention across storage systems, networks and end points [61].

E. Access Control and Privilege Management

Policies and processes used by cloud services provider to ensure that only the users granted appropriate privileges can use or modify data. It includes identification, authentication [62] and authorization issues [55]. The access control and privilege management are policies and processes used by cloud providers to ensure that only the consumers granted appropriate privileges can accede, use or modify data. Lately, researchers have proposed many models (such as Attribute Based Encryption (ABE), Key Policy Attribute Based Encryption, Cipher Text Policy Attribute Based Encryption, etc.) that are useful to provider security and access control. The majority of these proposed models are the modified form of the classical ABE model [63].

The implementation of the architecture proposed in this paper and the use of API brings some issues that must be identified to avoid problems.

API is used by the developers which act as an interface between the cloud service providers and the client. It allows users to manage and get the information from service providers. API and the related software need to be highly secured as it is used by the cloud users to access their data. API is the public front door entry to the data and accessible externally, thus it incorporates many threats in it [61].

V. CONCLUSIONS AND FUTURE WORK

Microservice-based architectures have become increasingly popular as an architectural style for software development. In this architectural style, the services provided by software solutions are divided into smaller parts and focused on the specific service of some functionalities. The approach of developing microservices with the construction of smaller software components has a number of advantages over the traditional monolithic architecture such as increasing the resilience of the software implemented as a microservice and the ease of scaling the solution implemented through the microservices.

Security aspects are critical in this architecture because the widespread use of Cloud Computing services, as demonstrated by the complexity of implementation, requires care with the privacy and security of information that is handled by those services.

The development of software using the microservice-based architecture comprises important aspects that must be observed in order to obtain good results. The objective of this article was to present the elements that should be considered for the development of solutions based on microservices and describing how the microservice-based architecture is defined. In addition to identifying the elements related to their implementation in Cloud Computing environment and explaining the privacy model applicable and relating the elements that intergrade the standards and solutions linked to the microservice-based architecture.

Future work will be developed to present case studies demonstrating the implementation of the microservice architecture in a Cloud Computing environment with the use of Docker containers for its construction and summarization of security troubles.

REFERENCES

[1] W. H. C. Almeida, L. D. A. Monteiro, R. R. Hazin, C. De Lima, and F. S. Ferraz, "Survey on Microservice Architecture - Security , Privacy and Standardization on Cloud Computing Environment," *ICSEA 2017*, no. c, pp. 199–205, 2017.

[2] S. H. Jyhjong Lin, Lendy Chaoyu, "Migrating Web Applications to Clouds with Microservices Arcitectures," *Int. Conf. Appl. Syst. Innov.*, pp. 1–4, 2016.

[3] M. Villamizar and et al, "Evaluating the Monolithic and the Microservice Architecture Pattern to Deploy Web Applications in the Cloud Evaluando el Patrón de Arquitectura Monolítica y de Micro Servicios Para Desplegar Aplicaciones en la Nube," *10th Comput. Colomb. Conf.*, pp. 583–590, 2015.

[4] D. I. Savchenko, G. I. Radchenko, and O. Taipale, "Microservices validation: Mjolnirr platform case study," *2015 38th Int. Conv. Inf. Commun. Technol. Electron. Microelectron. MIPRO 2015 - Proc.*, no. May, pp. 235–240, 2015.

[5] Y. Sun, S. Nanda, and T. Jaeger, "Security-as-a-service for microservices-based cloud applications," *Proc. - IEEE 7th Int. Conf. Cloud Comput. Technol. Sci. CloudCom 2015*, pp. 50–57, 2016.

[6] P. Jamshidi, C. Pahl, N. C. Mendonca, J. Lewis, and S. Tilkov, "Microservices: The journey so far and challenges ahead," *IEEE Softw.*, vol. 35, no. 3, pp. 24–35, 2018.

[7] Á. Kovács, "Comparison of different linux containers," *2017 40th Int. Conf. Telecommun. Signal Process. TSP 2017*, vol. 2017–Janua, pp. 47–51, 2017.

[8] C. Pahl and P. Jamshidi, "Microservices: A Systematic Mapping Study," *Proc. 6th Int. Conf. Cloud Comput. Serv. Sci.*, vol. 1, no. Closer, pp. 137–146, 2016.

[9] H. Bennasar, M. Essaaidi, A. Bendahmane, and J. Ben-Othman, "State-of-The-Art of cloud computing cyber-security," *Proc. 2015 IEEE World Conf. Complex Syst. WCCS 2015*, 2016.

[10] C. Esposito, "Challenges in Delivering Software in the Cloud as Microservices," *IEEE Cloud Comput.*, vol. 3, no. 5, pp. 10–14, 2016.

[11] D. Taibi, V. Lenarduzzi, C. Pahl, and A. Janes, "Microservices in agile software development: a workshop-based study into issues, advantages, and disadvantages," *Proc. XP2017 Sci. Work.*, p. 23, 2017.

[12] A. Panda, M. Sagiv, and S. Shenker, "Verification in the Age of Microservices," *Proc. 16th Work. Hot Top. Oper. Syst. - HotOS '17*, pp. 30–36, 2017.

[13] M. Villamizar *et al.*, "Infrastructure Cost Comparison of Running Web Applications in the Cloud Using AWS Lambda and Monolithic and Microservice Architectures," *Proc. - 2016 16th IEEE/ACM Int. Symp. Clust. Cloud, Grid Comput. CCGrid 2016*, pp. 179–182, 2016.

[14] N. Alshuqayran, N. Ali, and R. Evans, "A systematic mapping study in microservice architecture," *Proc. - 2016 IEEE 9th Int. Conf. Serv. Comput. Appl. SOCA 2016*, pp. 44–51, 2016.

[15] K. Jander, L. Braubach, and A. Pokahr, "Defense-in-depth and Role Authentication for Microservice Systems," *Procedia Comput. Sci.*, vol. 130, pp. 456–463, 2018.

[16] A. Krylovskiy, M. Jahn, and E. Patti, "Designing a Smart City Internet of Things Platform with Microservice Architecture," *Proc. - 2015 Int. Conf. Futur. Internet Things Cloud, FiCloud 2015 2015 Int. Conf. Open Big Data, OBD 2015*, pp. 25–30, 2015.

[17] H. Khazaei, C. Barna, N. Beigi-Mohammadi, and M. Litoiu, "Efficiency analysis of provisioning microservices," *Proc. Int. Conf. Cloud Comput. Technol. Sci. CloudCom*, pp. 261–268, 2017.

[18] R. Heinrich *et al.*, "Performance Engineering for Microservices: Research Challenges and Directions," *Proc. 8th ACM/SPEC Int. Conf. Perform. Eng. Companion*, pp. 223–226, 2017.

[19] M. Ahmadvand and A. Ibrahim, "Requirements reconciliation for scalable and secure microservice (de)composition," *Proc. - 2016 IEEE 24th Int. Requir. Eng. Conf. Work. REW 2016*, pp. 68–73, 2017.

[20] T. Q. Thanh, S. Covaci, T. Magedanz, P. Gouvas, and A. Zafeiropoulos, "Embedding security and privacy into the development and operation of cloud applications and services," *2016 17th Int. Telecommun. Netw. Strateg. Plan. Symp.*, pp. 31–36, 2016.

[21] Yale Yu, H. Silveira, and M. Sundaram, "A microservice based reference architecture model in the context of enterprise architecture," *2016 IEEE Adv. Inf. Manag. Commun. Electron. Autom. Control Conf.*, pp. 1856–1860, 2016.

[22] J. Rufino, M. Alam, J. Ferreira, A. Rehman, and K. F. Tsang, "Orchestration of Containerized Microservices for IIoT using Docker," pp. 1532–1536, 2017.

[23] M. B. Mollah, M. A. K. Azad, and A. Vasilakos, "Security and

privacy challenges in mobile cloud computing: Survey and way ahead," *J. Netw. Comput. Appl.*, vol. 84, pp. 38–54, 2017.

[24] C. Y. Fan and S. P. Ma, "Migrating Monolithic Mobile Application to Microservice Architecture: An Experiment Report," *Proc. - 2017 IEEE 6th Int. Conf. AI Mob. Serv. AIMS 2017*, pp. 109–112, 2017.

[25] J. Fowler, Marthin; Lewis, "Microservices: a definition of this new architectural term," *Microservices:a definition of this new architectural term*, 2014. .

[26] S. Systems, T. Cerny, and M. J. Donahoo, "Disambiguation and Comparison of SOA, Microservices and Self-Contained Systems," pp. 228–235.

[27] S. Yamamoto, S. Matsumoto, and M. Nakamura, "Using cloud technologies for large-scale house data in smart city," *CloudCom 2012 - Proc. 2012 4th IEEE Int. Conf. Cloud Comput. Technol. Sci.*, pp. 141–148, 2012.

[28] J. Bogner and A. Zimmermann, "Towards Integrating Microservices with Adaptable Enterprise Architecture," *Proc. - IEEE Int. Enterp. Distrib. Object Comput. Work. EDOCW*, vol. 2016–Septe, pp. 158–163, 2016.

[29] C. Gadea, M. Trifan, D. Ionescu, and B. Ionescu, "A reference architecture for real-time microservice API consumption," *Proc. 3rd Work. CrossCloud Infrastructures Platforms - CrossCloud '16*, pp. 1–6, 2016.

[30] D. Guo, W. Wang, G. Zeng, and Z. Wei, "Microservices architecture based cloudware deployment platform for service computing," *Proc. - 2016 IEEE Symp. Serv. Syst. Eng. SOSE 2016*, pp. 358–364, 2016.

[31] P. Marchetta, E. Natale, A. Pescape, A. Salvi, and S. Santini, "A map-based platform for smart mobility services," *Proc. - IEEE Symp. Comput. Commun.*, vol. 2016–Febru, pp. 19–24, 2016.

[32] A. de Camargo, I. Salvadori, R. dos S. Mello, and F. Siqueira, "An architecture to automate performance tests on microservices," *Proc. 18th Int. Conf. Inf. Integr. Web-based Appl. Serv. - iiWAS '16*, pp. 422–429, 2016.

[33] D. Jaramillo, D. V. Nguyen, and R. Smart, "Leveraging microservices architecture by using Docker technology," *Conf. Proc. - IEEE SOUTHEASTCON*, vol. 2016–July, pp. 0–4, 2016.

[34] P. Giessler, R. Steinegger, S. Abeck, and M. Gebhart, "Checklist for the API Design of Web Services based on REST," vol. 9, no. 3, pp. 41–51, 2016.

[35] A. Gueidi, H. Gharsellaoui, and S. Ben Ahmed, "A NoSQL-based Approach for Real-Time Managing of Embedded Data Bases," *Proc. - 2016 World Symp. Comput. Appl. Res. WSCAR 2016*, pp. 110–115, 2016.

[36] T. I. Damaiyanti, A. Imawan, and J. Kwon, "Extracting trends of traffic congestion using a NoSQL database," *Proc. - 4th IEEE Int. Conf. Big Data Cloud Comput. BDCloud 2014 with 7th IEEE Int. Conf. Soc. Comput. Networking, Soc. 2014 4th Int. Conf. Sustain. Comput. C*, pp. 209–213, 2015.

[37] R. Simmonds, P. Watson, and J. Halliday, "Antares: A Scalable, Real-Time, Fault Tolerant Data Store for Spatial Analysis," *Proc. - 2015 IEEE World Congr. Serv. Serv. 2015*, pp. 105–112, 2015.

[38] A. Ciuffoletti, "Automated Deployment of a Microservice-based Monitoring Infrastructure," *Procedia Comput. Sci.*, vol. 68, pp. 163–172, 2015.

[39] T. Combe, T. Paris-tech, A. Martin, R. Di Pietro, and N. B. Labs, "To Docker or Not to Docker : A Security Perspective," *IEEE Cloud Comput.*, vol. 3, no. 5, pp. 54–62, 2016.

[40] C. Fetzer, "Building critical applications using microservices," *IEEE Secur. Priv.*, vol. 14, no. 6, pp. 86–89, 2016.

[41] I. Khalil, A. Khreishah, and M. Azeem, "Cloud Computing Security: A Survey," *Computers*, vol. 3, no. 1, pp. 1–35, 2014.

[42] C. Saravanakumar and C. Arun, "Survey on interoperability, security, trust, privacy standardization of cloud computing," *Proc. 2014 Int. Conf. Contemp. Comput. Informatics, IC3I 2014*, pp. 977–982, 2014.

[43] J. Werner, C. M. Westphall, and C. B. Westphall, "Cloud identity management: A survey on privacy strategies," *Comput. Networks*, vol. 122, pp. 29–42, 2017.

[44] S. Srinivasan, "Data privacy concerns involving cloud," *2016 11th Int. Conf. Internet Technol. Secur. Trans. ICITST 2016*, pp. 53–56, 2017.

[45] M. Fazio, A. Celesti, R. Ranjan, C. Liu, L. Chen, and M. Villari, "Open Issues in Scheduling Microservices in the Cloud," *IEEE Cloud Comput.*, vol. 3, no. 5, pp. 81–88, 2016.

[46] F. Callegati, S. Giallorenzo, A. Melis, and M. Prandini, "Data security issues in MaaS-enabling platforms," *2016 IEEE 2nd Int. Forum Res. Technol. Soc. Ind. Leveraging a Better Tomorrow, RTSI 2016*, pp. 0–4, 2016.

[47] D. S. Linthicum, "Practical Use of Microservices in Moving Workloads to the Cloud," *IEEE Cloud Comput.*, vol. 3, no. 5, pp. 6–9, 2016.

[48] D. Yu, Y. Jin, Y. Zhang, and X. Zheng, "A survey on security issues in services communication of Microservices-enabled fog applications," *Concurr. Comput.*, no. September 2017, pp. 1–19, 2017.

[49] T. Yarygina and A. H. Bagge, "Overcoming Security Challenges in Microservice Architectures," *Proc. - 12th IEEE Int. Symp. Serv. Syst. Eng. SOSE 2018 9th Int. Work. Jt. Cloud Comput. JCC 2018*, pp. 11–20, 2018.

[50] H. Kang, M. Le, and S. Tao, "Container and microservice driven design for cloud infrastructure DevOps," *Proc. - 2016 IEEE Int. Conf. Cloud Eng. IC2E 2016 Co-located with 1st IEEE Int. Conf. Internet-of-Things Des. Implementation, IoTDI 2016*, pp. 202–211, 2016.

[51] K. Bao, I. Mauser, S. Kochanneck, H. Xu, and H. Schmeck, "A Microservice Architecture for the Intranet of Things and Energy in Smart Buildings," *Proc. 1st Int. Work. Mashups Things APIs - MOTA '16*, pp. 1–6, 2016.

[52] D. Escobar *et al.*, "Towards the understanding and evolution of monolithic applications as microservices," *Proc. 2016 42nd Lat. Am. Comput. Conf. CLEI 2016*, 2017.

[53] J. Stubbs, W. Moreira, and R. Dooley, "Distributed Systems of Microservices Using Docker and Serfnode," *Proc. - 7th Int. Work. Sci. Gateways, IWSG 2015*, pp. 34–39, 2015.

[54] R. Roostaei and Z. Movahedi, "Mobility and Context-Aware Offloading in Mobile Cloud Computing," *Proc. - 13th IEEE Int. Conf. Ubiquitous Intell. Comput. 13th IEEE Int. Conf. Adv. Trust. Comput. 16th IEEE Int. Conf. Scalable Comput. Commun. IEEE Int.*, pp. 1144–1148, 2017.

[55] K. El Makkaoui, A. Ezzati, A. Beni-Hssane, and C. Motamed, "Data confidentiality in the world of cloud," *J. Theor. Appl. Inf. Technol.*, vol. 84, no. 3, pp. 305–314, 2016.

[56] C. Perra and S. Member, "A Framework for the Development of Sustainable Urban Mobility Applications," 2016.

[57] M. Thangavel, P. Varalakshmi, and S. Sridhar, "An analysis of privacy preservation schemes in cloud computing," *Proc. 2nd IEEE Int. Conf. Eng. Technol. ICETECH 2016*, no. March, pp. 146–151, 2016.

[58] G. Shanmugasundaram and A. P. Cloud, "A COMPREHENSIVE REVIEW ON CLOUD COMPUTING SECURITY," 2017.

[59] H. Gebre-amlak, S. Lee, A. M. A. Jabbari, Y. Chen, and B. Choi, "MIST : Mobility-Inspired SofTware-Defined Fog System," 2017.

[60] D. Zissis and D. Lekkas, "Addressing cloud computing security issues," *Futur. Gener. Comput. Syst.*, vol. 28, no. 3, pp. 583–592, 2012.

[61] G. N. Dev, "A Survey on Security Threats in Cloud Computing Technology," *Int. J. Res.*, vol. 1, no. 8, pp. 1071–1081, 2014.

[62] R. H. Steinegger, D. Deckers, P. Giessler, and S. Abeck, "Risk-based authenticator for web applications," *Proc. 21st Eur. Conf. Pattern Lang. Programs - Eur. '16*, no. February 2017, pp. 1–11,

2016.

[63] K. E. Makkaoui*, A. Ezzati, A. Beni-Hssane, and C. Motamed, "Cloud Security and Privacy Model for Providing Secure Cloud Services," 2016.

Crime Forecasting in Small City Blocks Using a General Additive Spatio-temporal Model

Maria Mahfoud
CWI,
Stochastics,
Amsterdam, The Netherlands
Email: M.Mahfoud@cwi.nl

Sandjai Bhulai
Vrije Universiteit Amsterdam,
Faculty of Science,
Amsterdam, The Netherlands
Email: s.bhulai@vu.nl

Rob van der Mei
CWI,
Stochastics,
Amsterdam, The Netherlands
Email: R.D.van.der.Mei@cwi.nl

***Abstract*—Spatio-temporal modeling is widely recognized as a promising means for predicting crime patterns. Despite their enormous potential, the available methods are still in their infancy. A lot of research focuses on crime hotspot detection and geographic crime clusters, while a systematic approach to include the temporal component of the underlying crime distributions is still under-researched. In this paper, we gain further insight in predictive crime modeling by including a spatio-temporal interaction component in the prediction of residential burglaries. Based on an extensive dataset, we show that including additive space-time interactions leads to significantly better predictions.**

Keywords–Predictive analytics; forecasting; spatio-temporal modeling; residential burglary.

I. INTRODUCTION

How the police should respond to crime is a constant source of discussion and debate among scholars and practitioners. Over time, new strategies have been developed that use data to influence decision making and direct crime control [1]. This data was first used to indicate the underlying problems within a community by identifying clusters of repeating crime incidents. This was followed by using data to map crime to allow for rapid response to emerging crime problems and hotspots. The most recent development is intelligence-led policing, an objective method for formulating strategic policing priorities by using data analysis and crime intelligence for strategic planning and resource allocation in order to reduce, disrupt and prevent crime. The better integration of the available information systems allows the police to create a picture of the criminal environment and to predict the emerging areas of criminality [2].

Within an intelligence-led framework, proactive policing corresponds with an initial response of the law enforcement agencies to prevent crimes before being committed rather than reacting to criminal acts. Proactive policing requires the ability to predict crime hotspots and concentrations to identify likely targets for police intervention. The identification of these targets is one of the main goals of predictive policing [3].

Although the use of statistical analysis for predicting crimes has been around for decades, the Geographical Information System (GIS) revolution, in the recent years, has led to a surge of analytical techniques to identify likely targets in order to prevent criminal activities. Perry [3] organizes these techniques around six analytic categories: hot spot analysis, regression methods, data mining techniques, near-repeat methods, spatio-temporal analysis and risk terrain analysis. As stated by [4], "the most under-researched area of spatial criminology is that of spatio-temporal crime patterns". The same point has been made by Law et al. [5] who discusses spatio-temporal approaches in past crime research proposing a Bayesian spatio-temporal approach for modeling crime trends. Bernasco and Elffers [6] also address this issue of integrating the spatial and the temporal dimension of crime in order to advance the analysis of crime data. They mentioned that crime varies spatio-temporally illustrating this by an example from [7] on residential burglaries. Especially for residential burglaries, a body of research has shown the repeat and the near-repeat victimization effects [8]–[11]. Therefore, modeling the space-time interactions of residential burglaries are important to capture these effects.

Displaying statistical information on a map allows for conveying information in a format which is ideal for operational decision making. Spatio-temporal information can ideally be understood when displayed on a map, however, there are a number of issues related to the mapping of information in the policing domain. Among these is the use of choropleth maps. As noted by [4], "one particular problem among crime analysis is the incorrect tendency to map real values with choropleth (thematic) maps, resulting in the misleading impression that is often given by larger or unequal areas (Harries, 1999)". Chainey et al. [12] also mention the need of a threshold specification to identify hotspots. In their paper, they indicate also the influence of the parameter setting on the ability to predict future crimes using hotspot maps. The same problem was discussed by [13] who addresses the problem of hotspot identification and the variation of maps that can be obtained using the same data. They state that the choice of a thematic range represents a problem in itself.

An additional problem related to crime mapping is the varying sizes and shapes of geographic administrative boundary areas. Eck et al. [13] propose the use of small uniform grids as a solution to this problem. This results in a high-resolution model. This type of models provides a more realistic forecast in terms of structure and spatial variability [14]. However, it does not necessarily improve the forecast accuracy [15]. Roberts [16] highlights the necessity of evaluating the spatial

and temporal variation in the skill of the model in order to define the scales at which the model forecast should be believed.

This research focuses on residential burglaries and attempts to provide more clarity in predictive crime modeling and mapping by addressing the limitations discussed above. The major aims of this study are to find an accurate probability distribution of residential burglaries taking account of the space-time interactions, and to identify a cut-off value to classify areas as high-risk areas. Wang and Brown [17] model criminal incidents in Charlottesville using a spatio-temporal generalized additive model (ST-GAM) and extend it to a local spatio-temporal generalized additive model (LST-GAM). They applied the ST-GAM to predict the probability distribution of criminal incidents. In the ST-GAM, the temporal information of previous criminal incidents is modeled as a dummy variable indicating the time of the last committed criminal incident. They show that the ST-GAM and the LST-GAM outperform their previous spatial generalized linear model (GLM) and the hot spot model. This research extends the model proposed by [17] by allowing for more complicated space-time interactions.

Inspired by [18], we propose a generalized additive model (GAM) for modeling the probability distribution of residential burglaries in one district of Amsterdam based on regular lattice data (grid boxes of 125 $\times$ 125 meters). The model extends the base model similar to the one discussed in [17] by allowing for additive space-time interactions. We show that the model provides a useful forecast from a radius of 312.5 meters from the centroid of the grid. However, a clear improvement in the forecast accuracy is observed from the first neighborhood (187.5 meters from the centroid of the grid).

The remainder of this paper is organized as follows. Section II describes the used data set and the data analysis. Section III provides the methodological framework underlying this research. Section IV illustrates the results of the analysis. Section V concludes this research.

II. DATA

A. Data description

The data used for this research was provided by the Dutch Police. It contains all recorded incidents of residential burglaries that happened in one district of Amsterdam, with the highest burglary rate, between January 1, 2008 and April 30, 2014. The data was recorded at a monthly level and grouped into grids of 125 $\times$ 125 meters. The data is thus regular lattice data. Only the grids that correspond to urban areas were selected resulting in 1,812 grid locations. In total, there were 115,968 records with a total number of 11,450 incidents.

In addition, each crime incident recorded contained the latitude/longitude coordinates on the grid level, the time of occurrence (month, year) and different covariates that correspond to the demographic factors and the socio-economic factors that are associated with this grid. Next, to these covariates, the Dutch police also use some spatio-temporal indicators that specify when the last incident happened in a specific grid or combination of grids (neighborhood) using different time intervals. These spatio-temporal indicators are crime specific, for example, the number of residential burglaries in a specific grid one month before the reference date. The covariates that correspond to the demographic and the socio-economic factors are location-specific covariates and are constant over time. These covariates count 44 attributes, including population, average values of houses in the postal code area of the corresponding grid, percentage low incomes in the postal code area of the corresponding grid, and so on. Next, to these covariates, we also used some covariates that correspond with the geographic information of the city, such as the distance to the nearest highway access. In total, there were 61 covariates. The description of the discussed covariates is given in Table III.

Table I. Covariates including missing values and the corresponding percentage of the observed missing values.

Covariate	Missing (%)	Covariate	Missing (%)
POP	23	TPH	26
MP	23	ND	31
FP	23	AVH	46
NH	23	NLI	46
AHS	26	NHI	73
NWI	27	NPI	28
SH	26	PB	84
SPH	26	NE	94
MPH	26	AMI	28

B. Data exploration

A first analysis of the recorded incidents shows that only 1.2% of the total records had a higher number of residential burglaries than 1, while 91.61% of the records was equal to 0. For this reason, the occurrence of residential burglaries (binary) was considered as the response variable.

1) Missing values: The first problem encountered using the above-described data was a large number of missing values. The response variable contains no missing values. However, 113,408 of the 115,968 records contain at least one missing value. It is clear that removing every row that contains a missing value is not the best option as it will reduce the sample size by 97.8%.

Further analysis of the missing values shows that all missing values were observed for the location-specific covariates. Moreover, when a covariate contains missing values, at least 23% of the data was missing. Due to a large number of covariates and the high percentage of missing values we decided to remove the corresponding covariates. This concerns 18 of the 44 location-specific covariates. Table I shows these covariates with the corresponding percentage of missing values.

A deeper analysis of the covariates shows that the covariates that correspond to age categories were not complete (they did not sum up to 100%) and at least 25% of the observed values for each variable was equal to zero, which is not likely. For these reasons, these variables were also removed from the data set. Furthermore, the variable TSLI (the number of months since the last incident in the grid) was not always consistent with the corresponding spatio-temporal indicators and based on common sense, this variable is expected to be highly correlated with the other spatio-temporal indicators. For this reason, this variable was also removed from the data set.

2) Near zero-variance covariates: Further analysis of the data shows that many covariates have only some unique values

Table II. Near zero-variance covariates.

Covariate	Covariate	Covariate	Covariate
NA	NS	ACCOM	GI
BANK	SMKT	CS	SCS
NNC	LS	PFS	YC
HOSP	HFE	GH	TO

with low frequencies. These variables, also called near zero-variance variables, can cause numerical problems. Kuhn [19] considered a variable as a near zero-variance variable if two conditions were approved. The first one is that the percentage of unique values should be less than 20%. The second one is that the ratio of the most frequent to the second most frequent value should be greater than 20. The analysis of the near zero-variance covariates in our data set was performed using the `nearZeroVar` function from the caret package [20]. This analysis reveals that 16 covariates have a near zero-variance, which were removed from the data set. These covariates are illustrated in Table II.

The final data exploration was, mainly, performed following the protocol described in [21].

3) Outliers: First, a Cleveland dotplot was drawn for each covariate to identify potential outliers. The plots (see Figures 1-2) show that some covariates have potential outliers indicated by the isolated points. These outliers were replaced by the maximum values observed after removing the outliers from the data set. Moreover, the covariates CB (number of cafes and bars in the grid), REST (number of restaurants in the grid), and SHOP (number of shops in the grid) are highly unbalanced, as illustrated in Figure 1. To avoid problems due to a large number of zeros and to reduce the dimensionality of the data, these covariates were grouped into one covariate called public places (PP). This covariate has 19 unique values but is highly unbalanced. PP was divided into three categories. The first category is when no public places were observed in the grid. The second category is when there are at most five public places in the grid, and the last category is when there are more than five public places. This to distinguish between the grids in terms of crowdedness. Furthermore, EI (the number of educational institutions in the grid) is also highly unbalanced and has only three unique values, this covariate was used as a binary covariate (fPP).

4) Collinearity: Ignoring collinearity increases type II errors and leads to serious problems with forward and backward selection procedures [22]. As we are, among others, interested in the covariates that drive residential burglaries, we should be very careful about collinear covariates. To assess collinearity between the covariates, the variance inflation factor (VIF) was used. The VIF measures the amount by which the variance of a parameter estimator is increased due to collinearity with other covariates rather than being orthogonal [23]. First, the VIF was calculated using all covariates. The covariate with the highest VIF was removed and the VIFs have calculated again. This process was repeated until all VIF values were smaller than two. Note that the use of this threshold is subjective as there is no true VIF threshold. In the literature, different VIF values were suggested. Kennedy [24], among other authors, recommends a threshold of ten. A threshold of five was recommended by [25]. However, as mentioned in [22], the use of a VIF threshold of ten or even five is too high [26]. By using a threshold of two, we aim to be more conservative about collinearity. The VIF analysis shows that L6MN (number of incidents in the direct neighborhood in the sixth month and earlier before the reference time), L6MG (number of incidents in the grid in the sixth month and earlier before the reference time), and ADFS (average distance from the centroids of the grid to the nearest known 10 burglars) are collinear with other covariates and were removed from the data set.

Residential burglaries are known to have the repeat and near-repeat victimization effect where residential burglaries cluster over time and space [8] [27] [10]. Due to this effect, collinearity is expected between the spatio-temporal indicators. To provide more insight into the relationships between these covariates, the principal component analysis (PCA) biplot was used. The left panel of Figure 3 shows higher correlations between the number of incidents observed in the grid and in its direct neighborhood within the same time unit. The spatio-temporal indicators that correspond to the same time unit were aggregated resulting in three covariates TL1M, TL2M, and TL3M where TLxM is the total number of incidents observed in the grid and its direct neighborhood x months before the reference time. A PCA biplot was drawn using these covariates. As can be seen from the right panel of Figure 3, higher collinearity is observed between TL1M, TL2M, and TL3M. Again, to avoid loss of information, these covariates were grouped together into a new covariate, TL3, which is the total observed incidents in the grid and its direct neighborhood in the last three months. To check for outliers in TL3, a Cleveland dotplot was drawn and this plot shows no extreme observations. A PCA biplot was drawn again using TL3, MDFS (distance from the center of the grid to the nearest known burglar) and DTNHA (distance from the center of the grid to the nearest highway access), which shows that MDFS is negatively correlated with TL3 (this plot is not shown here but the same result can be concluded from Figure 3). We decided to use TL3 and leave MDFS out of the analysis.

Furthermore, conditional boxplots were used to assess collinearity between continuous and categorical covariates. This reveals that collinearity between SD and DTNHA exists. The covariate sub-district (SD) also shows some collinearity with TL3. To avoid problems due to collinearity, SD was omitted from the analysis.

The final set of covariates includes eight covariates, namely the space covariates X and Y; the temporal covariates YEAR and MONTH; the categorical covariates public places (fPP) and educational institutions (fEI); the total observed incidents in the grid and its direct neighborhood in the last three months (TL3) and finally, the distance to the nearest highway access (DTNHA).

5) Relationships between the response and the covariates: The relationship between the response variable and the nominal variables was assessed graphically by a design plot (Figure 4). As illustrated in Figure 4, higher mean values of the residential burglaries were observed between October and February, with the highest mean in December. This period is characterized by a short daylight period, while occupancy times of the residents remain the same. Due to the cover of darkness and the absence of the residents, burglars have a lower risk of being spotted. The highest value observed in December can be explained by the Christmas days and New

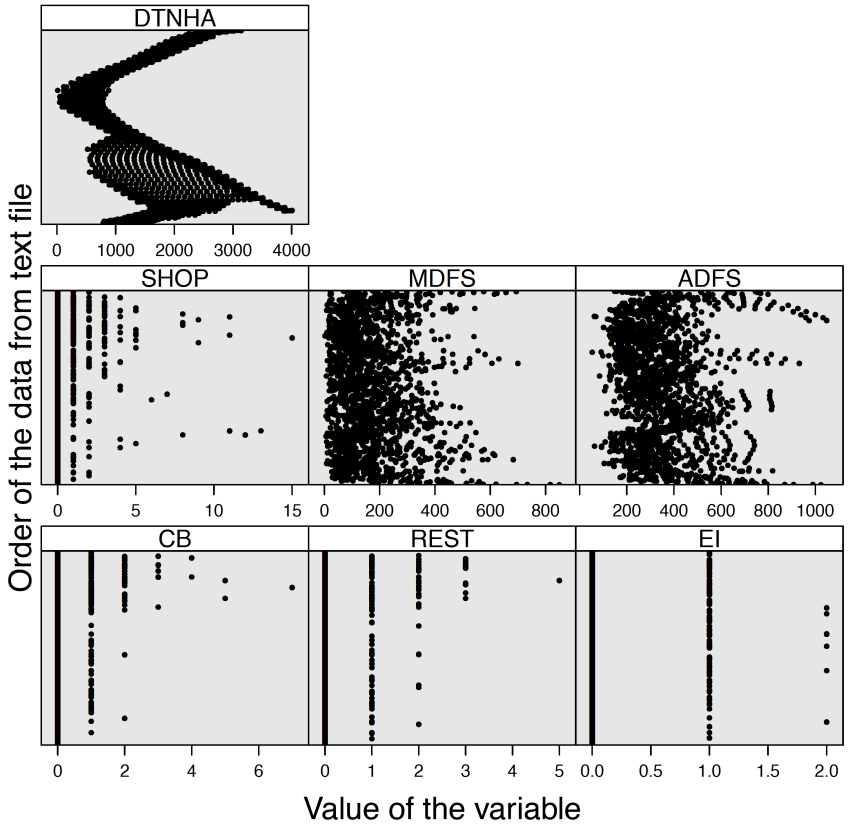

Figure 1. Multi-panel Cleveland dot plot for the location specific covariates. The horizontal axes represent the values of the covariates, and the vertical axes represent the order of the data as imported from the data file. Note the data is sorted on X, Y, YEAR and MONTH, respectively. This figure indicates the existence of some outliers in the most covariates. These are given by the isolated points in the right-hand side of the panels. This figure also shows that the discrete covariates (CB, REST, EI and SHOP) are highly unbalanced with some unique values.

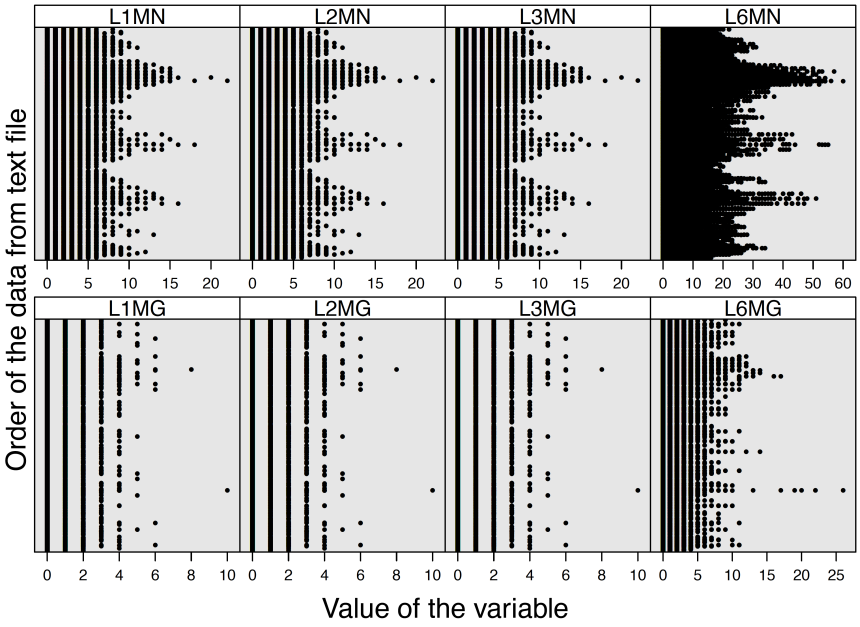

Figure 2. Multi-panel Cleveland dot plot for the spatio-temporal indicators. The horizontal axes represent the values of the covariates, and the vertical axes represent the order of the data as imported from the data file. Note the data is sorted on X, Y, YEAR and MONTH, respectively. This figure indicates the existence of some outliers in these covariates. The panels show roughly the same pattern indicating some collinearity between these covariates.

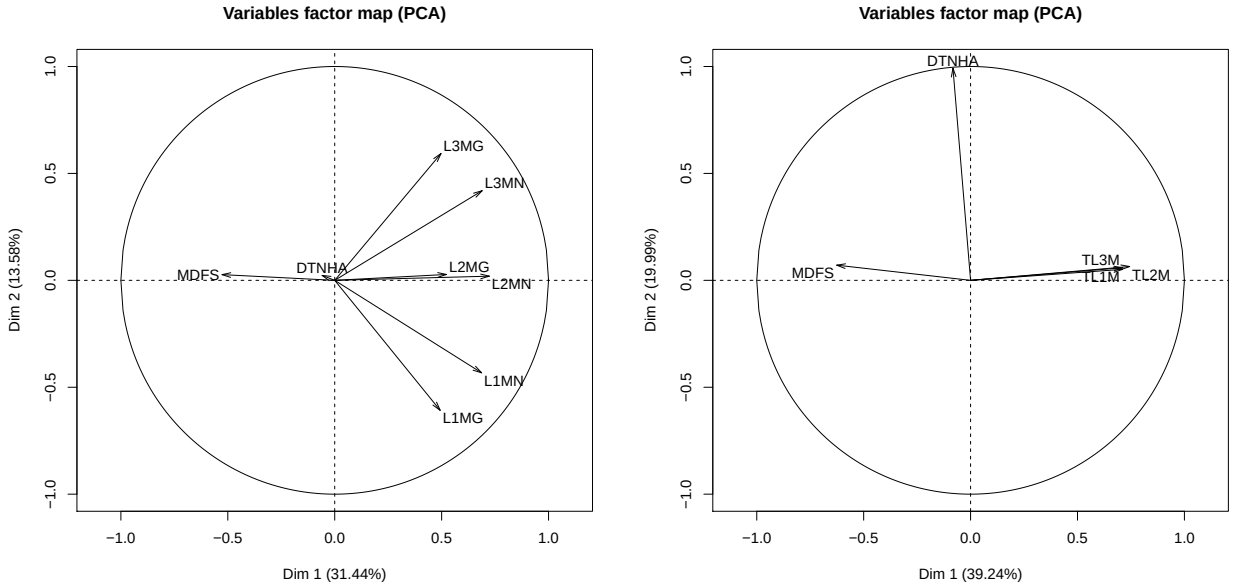

(a) (b)

Figure 3. PCA biplot of the covariates. The left panel indicates higher correlation between the number of residential burglaries observed in the grid and its direct neighborhood within the same time unit. The right panel shows the PCA biplot after aggregating the spatio-temporal indicators that correspond to the same time-unit. As can be seen from this panel, TL1M, TL2M and TL3M are highly correlated.

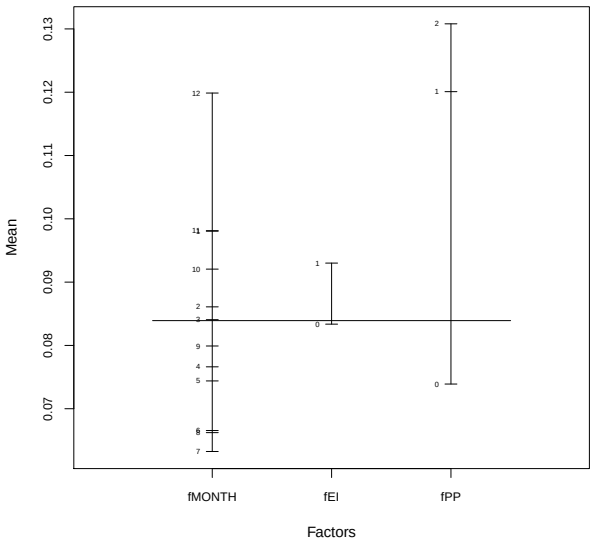

Figure 4. Design plot showing the average incidents per class for each factor variable.

Years Eve that are attractive days for burglars. Furthermore, a higher mean was observed in grids containing educational institutions (fEI) or public places (fPP). Moreover, crowded areas have a higher mean compared to quiet areas.

Finally, histograms of the TL3 and DTNHA for areas with residential burglary were plotted. A deeper analysis on TL3 shows that 93.14% of the incidents has occurred within grids with TL3 higher than zero. For this reason, the histogram of TL3 was drawn considering only TL3 values that are higher than zero. This shows a highly skewed distribution with peaks for TL3 values between two and four. Moreover, the distribution of DTNHA reveals a high peak of residential burglaries for distances between 875 and 1,000 meters.

In the next section, we introduce our generalized additive model (GAM) for modeling the probability distribution of residential burglaries. The model extends a base model by allowing for additive space-time interactions.

III. METHODOLOGY

Given the covariates discussed in Section II, the occurrence of residential burglaries in a certain grid i, and in a certain month t, was modeled using a GAM using the binomial distribution and the logistic link function (see, e.g., [28], [29]). To be more precise, the model is not a GAM with the binomial distribution but rather one with a Bernoulli distribution. The use of the logit link is to ensure that the fitted values are bounded in $(0, 1)$.

The choice of GAM is based on the expected non-linear relationships between the covariates and the response. A non-linear relationship is expected between the response and the distance to the nearest highway access (DTNHA). This can be explained by the two types of burglars identified by [30], the first being the opportunity burglar that prefers to operate within its own neighborhood and the second being the professional burglar who selects its targets based on the highest expected loot and operates mostly in suburban areas and areas that are near highways, because they are unaware of the local situation

and escape routes. A non-linear relationship is also expected for TL3 due to the repeat and near-repeat victimization effects. The covariate MONTH is also expected to have a non-linear effect on the residential burglaries. This is due to the repeat victimization effect and the daylight-darkness effect [31]. For these reasons, smoothers will be used to model these covariates.

We use a GAM model that allows for space-time interactions as follows:

$$\begin{aligned}\text{logit}(\mu_{i,t}) = \text{fEI}_i + \text{fPP}_i + \text{YEAR}_t + f_1(\text{TL3}_{i,t}) + \\ f_2(\text{DTNHA}_{i,t}) + f_3(\text{MONTH}_t) + f_4(\text{X}_i, \text{Y}_i),\end{aligned} \quad (1)$$

where $\mu_{i,t} = \mathbb{E}(y_{i,t})$, $y_{i,t}$ follows a Bernouilli distribution, $i \in \{1, \ldots, 1812\}$, $t \in \{1, \ldots, 60\}$. The functions f_1 and f_2 are one-dimensional smoother functions of the covariates represented by a cubic regression spline (CRS). f_3 is a one-dimensional smoother represented by a cyclic cubic regression spline (CCRS). This is to avoid big jumps between the January and the December value of the smoother [32]. The function f_4 is a two-dimensional isotropic smoother for space represented by thin plate regression splines (TPRS). The TPRS was used for smoothing the spatial co-ordinates because they are measured on the same unit [29].

The model was fitted using the penalized iteratively re-weighted least squares $(\text{P} - \text{IRLS})$, while the optimal amount of smoothing was estimated using the UnBiased Risk Estimator (UBRE) [29]. All analyses were conducted using the `mgcv` package [29] from the R statistical and programming environment [33].

IV. Results

Now that we can generate the probability function of residential burglaries through the GAM model, which cut-off value θ should be used to classify high-risk areas and which spatial scale provides a useful forecast? In practice, the choice of the cut-off value is mostly left to law enforcement agencies who choose a cut-off value based on the available resources and their risk preferences. Some of them choose a cut-off value of 0.8, others select areas based on the top 3% or the top 5% percentiles to classify areas as high-risk areas. However, the use of a hard cut-off value as 0.8 strongly depends on the estimated probabilities. In our case, this will result in a clear under-estimation of risk areas. If one decided to use a fraction of top percentiles, then this should be at least equal to the expected percentage of incidents. Elsewhere, the risk areas will be undoubtedly under-estimated.

Considering our training set, the average incidents (binary) over the five years, ranging between 2008 and 2012, was about 8.3%. This means that on average 151 grids, from the total grids of 1,812, should be considered as risky grids. Using the 97% percentile results in considering only 55 grids as high-risk areas. Doing this, we know apriori that we are under-estimating the risk areas. Some people will argue that the given resources do not allow to cover this high number of grids. In our point of view, from a safety perspective, the grids that should be flagged as high-risk areas should at least match the expected grids with incidents and should be independent of the available resources. After classifying the areas as high-risk areas, smart allocation methods can be used to cover the risk areas using the available resources.

Given the estimated probability distribution, the optimal cut-off (the average) considering the different neighborhoods ($\theta_1 = 0.171$) and the optimal cut-off at the grid level ($\theta_2 = 0.126$) were further used to classify areas as high-risk areas. The reason of using both cut-off values is because the optimal cut-off on grid level was quite different from the optimal cut-offs that correspond to the other neighborhoods.

The generated heat maps of January and April are given in Figure 5. From this figure, a clear difference is observed in the number of grids that are flagged as high-risk grids. In fact, more incidents are expected in January compared to April. Therefore, the predicted high-risk area in January is larger compared to the one in April. The heat maps also show that most realizations were located within the high-risk area or within their lower bounds.

In January, more incidents are expected compared to April, this is in agreement with historical data (see Figure 4). The heat maps also show that most realizations are located within the high-risk area or within its lower bound.

V. Conclusion

In this research, we developed a GAM model to predict the probability distribution of residential burglaries. The results show that the covariate TL3, the total incidents in the grid and its neighborhood in the last three months, has a dominant effect in the model. Apparently, this covariate captures a large part of the spatio-temporal effect in residential burglaries. Moreover, a small part of the variation in the data was captured by the model. The low power of the model may be due to the high resolution of the data used.

Finally, θ_1 and θ_2 were used to assess the performance of the model and these cut-offs were compared with the cut-off obtained for the maximum performance. Results show that both values provide similar results to the maximum performance observed, while the cut-offs that correspond to the maximum performance considering the different metrics cover a wide range, which can be difficult to interpret from a decision-making point of view.

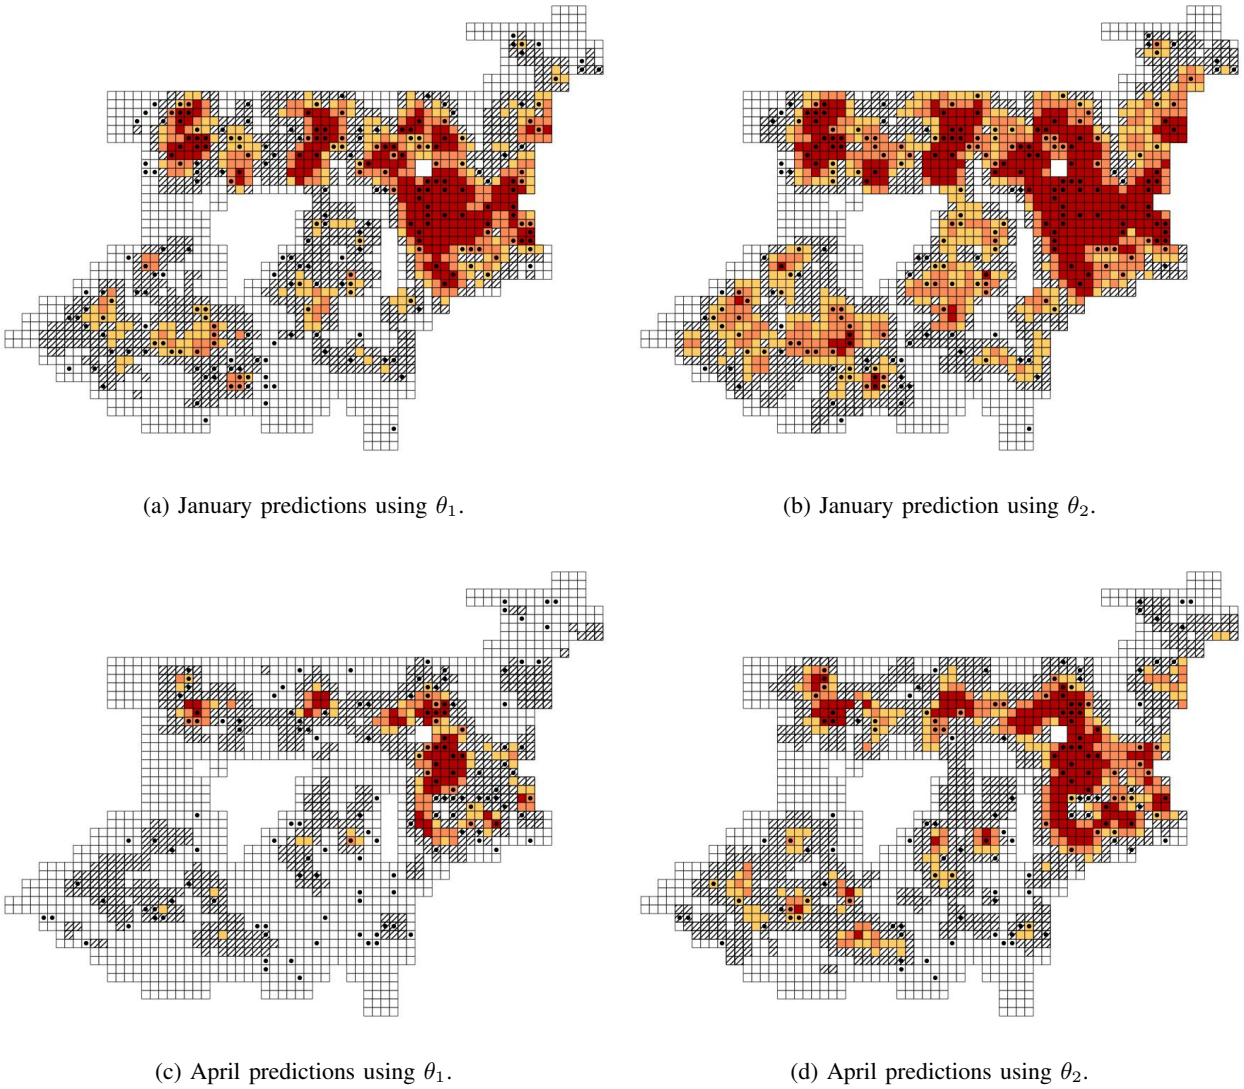

(a) January predictions using θ_1.

(b) January prediction using θ_2.

(c) April predictions using θ_1.

(d) April predictions using θ_2.

Figure 5. Heat maps of January and April using θ_1 and θ_2 and including the estimated lower bounds. The heat maps show that almost all incidents are located within the estimated high-risk area or within their lower bounds. It can also be seen that the estimated high-risk area of January is larger than the one of April. The maps obtained using θ_1 show that almost all incidents are located within the high-risk area or within their lower bound. However, the total high-risk area is smaller compared to a high-risk area obtained using θ_2. This result is very appealing for the resource allocation.

REFERENCES

[1] M. Mahfoud, S. Bhulai, and R. v. d. Mei, "Spatio-temporal modeling for residential burglary," in Proceedings of the 6th International Conference on Data Analytics. IARIA, 2017, pp. 59–64.

[2] J. H. Ratcliffe, Intelligence-Led Policing. Willan publishing, 2008.

[3] W. L. Perry, Predictive policing: The role of crime forecasting in law enforcement operations. Rand Corporation, 2013.

[4] J. H. Ratcliffe, "Crime mapping: spatial and temporal challenges," in Handbook of quantitative criminology. Springer, 2010, pp. 5–24.

[5] J. Law, M. Quick, and P. Chan, "Bayesian spatio-temporal modeling for analysing local patterns of crime over time at the small-area level," Journal of quantitative criminology, vol. 30, no. 1, 2014, pp. 57–78.

[6] W. Bernasco and H. Elffers, "Statistical analysis of spatial crime data," in Handbook of quantitative criminology. Springer, 2010, pp. 699–724.

[7] J. H. Ratcliffe, "Residential burglars ad urban barriers: a quantitative spatial study of the impact of canberra's unique geography on residential burglary offenders," http://crg.aic.gov.au/reports/ratcliffe.html, 2001, last access date: 31 October, 2017.

[8] W. Bernasco and P. Nieuwbeerta, "How do residential burglars select target areas? a new approach to the analysis of criminal location choice," British Journal of Criminology, vol. 45, no. 3, 2005, pp. 296–315.

[9] S. D. Johnson, W. Bernasco, K. J. Bowers, H. Elffers, J. Ratcliffe, G. Rengert, and M. Townsley, "Space–time patterns of risk: a cross national assessment of residential burglary victimization," Journal of Quantitative Criminology, vol. 23, no. 3, 2007, pp. 201–219.

[10] M. Short, M. D'Orsogna, P. Brantingham, and G. Tita, "Measuring and modeling repeat and near-repeat burglary effects," Journal of Quantitative Criminology, vol. 25, no. 3, 2009, pp. 325–339.

[11] W. Bernasco, S. D. Johnson, and S. Ruiter, "Learning where to offend: Effects of past on future burglary locations," Applied Geography, vol. 60, 2015, pp. 120–129.

[12] S. Chainey, L. Tompson, and S. Uhlig, "The utility of hotspot mapping for predicting spatial patterns of crime," Security Journal, vol. 21, no. 1, 2008, pp. 4–28.

[13] J. Eck, S. Chainey, J. Cameron, and R. Wilson, "Mapping crime: Understanding hotspots," http://discovery.ucl.ac.uk/11291/, 2005, last access date: 31 October, 2017.

[14] E. E. Ebert, "Neighborhood verification: A strategy for rewarding close forecasts," Weather and Forecasting, vol. 24, no. 6, 2009, pp. 1498–1510.

[15] C. F. Mass, D. Ovens, K. Westrick, and B. A. Colle, "Does increasing horizontal resolution produce more skillful forecasts?" Bulletin of the American Meteorological Society, vol. 83, no. 3, 2002, pp. 407–430.

[16] N. Roberts, "Assessing the spatial and temporal variation in the skill of precipitation forecasts from an nwp model," Meteorological Applications, vol. 15, no. 1, 2008, pp. 163–169.

[17] X. Wang and D. E. Brown, "The spatio-temporal generalized additive model for criminal incidents," in Intelligence and Security Informatics (ISI), 2011 IEEE International Conference on. IEEE, 2011, pp. 42–47.

[18] N. H. Augustin, V. M. Trenkel, S. N. Wood, and P. Lorance, "Space-time modelling of blue ling for fisheries stock management," Environmetrics, vol. 24, no. 2, 2013, pp. 109–119.

[19] M. Kuhn, "Building predictive models in r using the caret package," Journal of Statistical Software, vol. 28, no. 5, 2008, pp. 1–26.

[20] M. Kuhn, J. Wing, S. Weston, A. Williams, C. Keefer, A. Engelhardt, T. Cooper, Z. Mayer, B. Kenkel, the R Core Team, and M. Benesty., caret: Classification and Regression Training, 2014, r package version 6.0-37, Last access date: 31 October, 2017. [Online]. Available: http://CRAN.R-project.org/package=caret

[21] A. F. Zuur, E. N. Ieno, and C. S. Elphick, "A protocol for data exploration to avoid common statistical problems," Methods in Ecology and Evolution, vol. 1, no. 1, 2010, pp. 3–14.

[22] A. F. Zuur, E. N. Ieno, and G. M. Smith, Analysing ecological data. Springer New York, 2007, vol. 680.

[23] D. Liao and R. Valliant, "Variance inflation factors in the analysis of complex survey data," Survey Methodology, vol. 38, no. 1, 2012, pp. 53–62.

[24] P. Kennedy, A guide to econometrics, 6th ed. Willey-Blackwell, 2008.

[25] P. Rogerson, Statistical methods for geography. Sage, 2001.

[26] D. C. Montgomery, E. A. Peck, and G. G. Vining, Introduction to linear regression analysis. Wiley, 1992.

[27] S. D. Johnson, "Repeat burglary victimisation: a tale of two theories," Journal of Experimental Criminology, vol. 4, no. 3, 2008, pp. 215–240.

[28] T. J. Hastie and R. J. Tibshirani, Generalized additive models. CRC Press, 1990, vol. 43.

[29] S. Wood, Generalized additive models: an introduction with R. CRC press, 2006.

[30] C. van den Handel, O. Nauta, P. van Soomeren, and P. van Amersfoort, "Hoe doen ze het toch? modus operandi woninginbraak," https://hetccv.nl/onderwerpen/woninginbraak/documenten/hoe-doen-ze-het-toch-modus-operandi-woninginbraak/, 2009, last access date: 31 October, 2017.

[31] T. Coupe and L. Blake, "Daylight and darkness targeting strategies and the risks of being seen at residential burglaries," Criminology, vol. 44, no. 2, 2006, pp. 431–464.

[32] A. Zuur, E. N. Ieno, N. Walker, A. A. Saveliev, and G. M. Smith, Mixed effects models and extensions in ecology with R. Springer, 2009.

[33] R Core Team, R: A Language and Environment for Statistical Computing, 2013, last access date: 31 October, 2017. [Online]. Available: http://www.R-project.org/

Appendix

Table III. List of covariates with a short description.

Covariate	Description
X	X coordinate of grid
Y	Y coordinate of grid
YEAR	Year of reference time
MONTH	Month of reference time
DISTRICT	District
SD	Sub police district
POP	Number of residents in postal code area of the grid
MP	Number of male residents in postal code area of the grid
FP	Number of female residents in postal code area of the grid
NH	Number of households in postal code area of the grid
AHS	Average household size in postal code area of the grid
AC1	Percentage residents between 0 and 14 years old in postal code area of the grid
AC2	Percentage residents between 15 and 24 years old in postal code area of the grid
AC3	Percentage residents between 25 and 44 years old in postal code area of the grid
AC4	Percentage residents between 45 and 64 years old in postal code area of the grid
AC5	Percentage residents between 65 and 74 years old in postal code area of the grid
AC6	Percentage residents 75 years and older in postal code area of the grid
NWI	Percentage non-western immigrants in postal code area of the grid
SH	Percentage of single-person household in postal code area of the grid
SPH	Percentage of single-parent household in postal code area of the grid.
MPH	Percentage of multiple households without children in postal code area of the grid
TPH	Percentage two-parent households in postal code area of the grid
ND	Number of dwellings in postal code area of the grid
AVH	Average value of the houses in postal code area of the grid
NLI	Percentage low income households in postal code area of the grid
NHI	Percentage high income households in postal code area of the grid
NPI	Number of persons that generate income in postal code area of the grid
PB	Percentage of persons that receive social benefits in postal code area of the grid
NE	Percentage of entrepreneurs in postal code area of the grid
AMI	Average monthly income in postal code area of the grid
CB	Number of cafes and bars in the grid
REST	Number of restaurants in the grid
EI	Number of educational institutions in the grid
NA	Number of associations in the grid
NS	Number of snack bars in the grid
ACCOM	Number of hotels in the grid
GI	Number of government institutions in the grid
BANK	Number of banks in the grid
SMKT	Number of supermarkets in the grid
CS	Number of coffee shops in postal code area of the grid
SCS	Number of sex shops, clubs and shows in the grid
LS	Number of liquor stores in the grid
PFS	Number of petrol filling stations in the grid
NNC	Number of nightclubs in the grid
YC	Number of youth centres in the grid
HOSP	Number of hospitals in the grid
HFE	Number of nursing home for the elderly
GH	Number of gambling houses in the grid
TO	Number of tourist offices in the grid
SHOP	Number of shops in the grid
TSLI	Number of months since the last incident in the grid
L1MG	Number of incidents in the grid in the first month before the reference time
L1MN	Number of incidents in the direct neighbourhood of the grid in the first month before the reference time
L2MG	Number of incidents in the grid in the second months before the reference time
L2MN	Number of incidents in the direct neighbourhood of the grid in the second months before the reference time
L3MG	Number of incidents in the grid in the third month before the reference time
L3MN	Number of incidents in the direct neighbourhood of the grid in the third month before the reference time
L6MG	Number of incidents in the grid in the sixth month and earlier before the reference time
L6MN	Number of incidents in the direct neighbourhood in the sixth month and earlier before the reference time
MDFS	Distance (m) from the center of the grid to the nearest known burglar
ADFS	Average distance (m) from the centroids of the grid to the nearest known 10 burglars
DTNHA	Distance (m) from the center of the grid to the nearest highway access

Achieving GDPR Compliance with Unikernels

Bob Duncan
Computing Science
University of Aberdeen, UK
Aberdeen, UK
Email: bobduncan@abdn.ac.uk

Andreas Happe
Dept. Digital Safety & Security
Austrian Inst. of Tech. GmbH
Vienna, Austria
Email: andreas.happe@ait.ac.at

Alfred Bratterud
Dept. of Computer Science
Oslo and Akershus University
Oslo, Norway
Email: alfred.bratterud@hioa.no

***Abstract*—IT security and privacy has always been a challenging problem to address, but with cloud, there is an exponential increase to the challenge. Once an attacker successfully breaches a cloud system, the intruder will seek to escalate privileges in order to delete the forensic trail, thus covering their tracks. There is little to prevent this from happening in cloud, and this is known as the Cloud Forensic Problem. Under the new European Union General Data Protection Regulation, following a cyber breach, it is necessary for the breached company to report the impact of the breach within 72 hours of becoming aware of the breach. Where the forensic trail has been compromised, this will present a serious compliance challenge. We address this problem through the use of Unikernel based monitoring systems which can ensure both full forensic and audit trails can be maintained. Our early results are very promising. We are continuing our work with a larger pilot study.**

Keywords–Cloud Forensic Problem; unikernels; EU GDPR, compliance.

I. INTRODUCTION

All business is the subject of cyber attacks, no matter whether it is a public corporation, a private firm, a financial institution, a government agency, a non-government agency or a charity. In previous work [1], we proposed the use of a unikernel based system to help defend against such attacks. No matter what type of organisation is involved, all those who will be subject to the rules of the European Union (EU) General Data Protection Regulation (GDPR) [2], will need to comply fully with the regulation. No matter where the company is located in the world, should they hold personally identifiable information (PII) belonging to any EU resident, they will fall under the jurisdiction of the EU GDPR regulator. In a post-Brexit world, the UK Government has indicated that the GDPR will still apply in the UK. Indeed, the UK Government has indicated that the UK GDPR will be enforced with greater rigour, and will accord greater rights to private individuals.

In order to achieve compliance with the rules of the GDPR, companies who fall under the scope of the GDPR will necessarily require to undertake considerable extra work, and expense, in order to be able to achieve compliance. Each organisation will require to appoint a data controller, who either must have the requisite technical skills, or must be assisted by a person with such technical skills. This will likely be an unwelcome additional expense. They must also have a data processor and a data protection officer, meaning further costs. In addition, they will have to take all necessary technical steps to ensure the security and privacy of all PII belonging to data subjects of the organisation, again at yet more expense.

Many companies are likely to be unprepared for achieving compliance. Many (erroneously) believe that because the reporting requirement has been changed from "within 72 hours of a breach occurring" to "within 72 hours of discovering a breach", they will have no problem being compliant [3]. The reality is that they will be wrong! They must also be able to report which records were accessed, modified, deleted or exfiltrated from the system. However, once an attacker breaches a system and becomes resident as an intruder, one of the first tasks they seek to carry out is to delete the forensic trail which recorded their incursion into the enterprise systems, so that their presence becomes more covert, allowing them to remain hidden inside the system. This allows them to harvest whatever information or secrets they desire for as long as they remain hidden in the system.

Without a complete forensic trail in any system, compliance will be a challenge, if not impossible. This will particularly be the case with cloud systems, since there is nothing to prevent such an intruder from deleting not only the forensic trail, but anything else they desire, including the very running cloud instance that they are hiding within. If there is no record of the trail of events relating to the database contents, then the company is unlikely to be able to identify which records have been accessed, modified, or deleted, resulting in a failure to be compliant with the GDPR. Since failure to comply can result in fines which can rise to the greater of €20 million or 4% of global turnover, then this will certainly have a substantial impact, although there are many who still fail to grasp this important point.

We start by considering the cloud forensic problem in Section II, and discuss why this is such a challenge for GDPR compliance in cloud systems. We are concerned with achieving both good security and good privacy. While it is possible to have security without privacy, it is not possible to have privacy without security. Thus our approach will be to first ensure a good level of security can be achieved, and to that end, we start by listing the specific security issues we seek to address and discuss how we propose to tackle them in Section III. The remainder of the paper is organized as follows: in Section IV, we consider how we might go about finding a cloud based solution, in Section V, we discuss the outline technical details of our proposed approach; In Section VI, we consider possible attack vectors. In Section VII, we consider just how robust a unikernel approach might be. In Section IX, we discuss our conclusions.

II. THE CLOUD FORENSIC PROBLEM AND THE GDPR

Cloud computing has been around now for over 10 years, and a great deal of good quality research has been carried out,

particularly regarding matters of security and privacy. Cloud systems have become highly popular with companies due to the flexibility of cloud offerings. The speed of starting a cloud instance, the removal of long lead times to access compute and storage resources, the ability to scale up, as well as down, to match demand presents a particularly good incentive to use cloud computing. The fact that companies can write costs off entirely against revenue provides a further attractive incentive for their use. Kratzke [4] has long warned of the dangers of thinking that conventional software is just the same as cloud-native software. Kratzke et al. [5] do suggest the possibility of using existing Container technologies to improve cloud-native programming.

There have been many good papers produced on both security [6]–[17] and privacy [14], [18]–[32], and we laud the efforts of countless researchers who have tried to provide this area with better security and privacy, which speaking generally, has been successfully achieved over the years. A number of others have looked at better accountability as a means to meeting these ends [10], [11], [15], [20], [27], [30], [33]–[52] But there remains one fundamental weakness that has not been resolved, namely the "cloud forensic problem". All computer systems are the subject of attack, and cloud systems are no exception. Unfortunately, no system is immune to attack, and that is certainly true for cloud systems.

No computer system is immune to attack. It is the primary goal of an attacker to breach a system. This can involve quite a considerable amount of work on the part of a serious attacker. They are very likely to perform extensive surveillance and compile many analyses of how the company systems and their architectures are organised. Many will carry out considerable amounts of social engineering work to attempt to fully understand the people of the organisation, since people are frequently the weakest link. But understanding the organisational structure can also provide vital intelligence to understand how company procedures operate, all of which can help them achieve their goals. This intelligence gathering will be very comprehensive and thorough, usually covering every possible aspect of all the systems of the company in order to discover everything they can about the business architecture before they start their attacks. By understanding fully how the company is structured and how it operates, they are far less likely to make any errors when they start the process of penetrating the systems.

Other attackers, are much less organised. They will simply try to hack in to company systems, without any regard or thought of the overview of the company concerned. They will merely look for known software vulnerabilities and try their best to successfully attack them. They care little about whether they are discovered while attempting to penetrate the system. Theirs is a short term view, rather than the long view held by others. They want to get in, and out, quickly with whatever they can lay their hands on. For them, time is money, and if they are unable to get in within a reasonable amount of time, they will move on to the next prospect.

Yet other potential intruders will perpetrate their attack through the people of the business using a variety of other attack methods: such as using social engineering attacks, email attacks that might use malicious links and malware payloads, attempting to use web based drive by attacks, or the use of phishing, vishing and many other approaches. These attackers are much less concerned with purely technical attacks, but are often extremely talented in the use of these methods, and in particular social engineering.

No matter which type of attacker they are, they all share one fundamental goal — and that is to penetrate the system in order to become an intruder. The aim here is not just to get in, and out, as quickly as possible, but to develop a long term foothold inside company systems which will allow them to return, time and again, to help themselves to whatever they wish, as they escalate privileges more and more, the longer they remain inside the system. This will necessarily involve some serious attempts to escalate privileges to allow them to mdify the forensic trail.

It is rather unfortunate that they are often greatly aided in this quest by the companies themselves. Usually, this occurs through a degree of laziness whereby the companies are clearly failing to monitor server logs properly. Looking at previous cyber breach reports [53], at which time there was a global average of 6 months between breach and discovery, it is clear that had these companies been more rigorous about reading and analysing their server logs, they would have been in a better position to discover intruders rather more quickly. Even last year, where the time between breach and discovery has dropped to a number of weeks rather than months [54], this is still not good enough. Some companies contribute further by refusing or failing to update security patches to both operating systems and software systems, usually under the guise of "last time I did a security update, all the systems crashed".

This all leads towards the, as yet unresolved, cloud forensic problem — namely, that once an intruder is in the system, and has escalated sufficient privileges, there is absolutely nothing to prevent them from deleting the forensic trail, which allows them to hide all evidence of their successful penetration. Worse, by this stage they will also have control of all the system logs and audit trails, and there is nothing to prevent them from deleting every last trace of their intrusion and ongoing ex-filtration of private data.

Surely that has nothing to do with the GDPR you might ask? Sadly, that is not the case. In the event of a breach, you are required to report the breach within 72 hours of discovering the breach. You must be able to report how many relevant records have been compromised, whether by having been read, amended, deleted or ex-filtrated from the system. Given that many system logs are not even turned on by default, this means identifying which records have been compromised, whether by having been read, amended, deleted or ex-filtrated from the system, will present a serious enough challenge in the first place. However, given that the intruder will likely have thoroughly worked through all forensic trails in the system, the likelihood of being able to realise that a breach has occurred at all will likely be very slim, let alone having the ability to properly identify which records have been compromised.

From a holistic perspective, it would have been helpful if these matters might have been addressed by the Cloud implementation itself. However, no such attempt has taken place during the past decade, no doubt due to the massive challenge involved. Consequently, all organizations subject to the provisions of the GDPR are required to safeguard their own systems and therefore take such steps as are necessary to ensure adequate privacy is achieved.

This will mean non-compliance with the GDPR, which can then trigger fines which can rise to the greater of €20,000,000 or 4% of global turnover. This will certainly catch the attention of top management within organisations. Considering that these fines can be levied for every single breach, and that the upper limit is based on turnover rather than profit, that should be sufficiently concerning to get some serious attention. Of course, all sensible Cloud users should have been thinking about this long before now, and we are aware of many who on hearing that notification 'within 72 hours of discovery of a breach', rather than 'within 72 hours of occurrence of a breach', heaved a collective sigh of relief and stopped worrying about implementing a solution. This is what motivates our work.

III. How do we tackle the problem?

At this time, no system is fully secure. Operating systems, transport protocols, software applications — all of this software has evolved during previous decades. Any relevant standards were defined decades ago. The primary goal at that time was functionality. Security and privacy were very much an afterthought, which has remained the case for decades. Security and privacy has very much been a case of "Let us bolt something on to tackle that". Default settings are geared for ease of setting up, not for security and privacy. This means proper security and privacy presents a massive challenge, which increases exponentially for cloud, Internet of Things and Big Data.

Since the primary goal of the successful intruder is to delete or obfuscate the forensic trail which could expose their presence, then we must consider protection of this data a priority. However, before becoming a successful intruder, the attacker has first to get into the system. This process will be capable of triggering certain alarms, if activated. At the very least, proper scrutiny of server logs would be a big help here. It is not necessary to have human eyes on all these logs, but it would be sensible to use some automated means to detect anomalous behaviour and to flag this up before the attacker can gain a permanent foothold within the system. Thus, there are two specific needs to fulfil here. One is the proper protection of all forensic data and audit trails, and the other is to analyse the system traffic in a timely manner to detect potential anomalous behaviour.

One might imagine that it would logically be more efficient to deal with the second need first before considering the first. However, as we have already stated, retaining a full and proper record is not only vital for GDPR compliance requirements, but with compromised forensic and audit trails, there will not be a full picture to analyse for anomalous behaviours, rendering the task less than useful. We therefore suggest the protection of the forensic and audit data has to be the priority, meaning that the subsequent analysis of this data will at least be run on a full set of data.

We therefore address the security of the forensic and audit trail data as our first priority, returning to the analysis of log data to detect anomalous behaviour in Section VIII. We therefore seek a suitable mechanism that will be fit for our purposes, and consider here the advantages and disadvantages of a number of possible alternatives.

Conventional algorithms running on the server could potentially work well, but their weakness lies in running on the server instance where they are vulnerable to attack. They would also present a considerable overhead to the smooth running of the main web application on the cloud instance.

We could opt to use Containers, such as Docker, LXD or Rocket. However, Bratterud et al. [55] warn of some security issues with this approach, and Kratzke [56] also warns of the unexpected, and unwelcome overhead these solutions can bring.

In previous work, [57], we considered how well unikernels might be used to improve on dealing with our target list of security goals, and found the potential for an improved approach. In [58], we developed a suitable framework, providing detailed definitions of how this might be tackled. In [59], we demonstrated how a unikernel based solution could reduce complexity, while improving security and privacy. We also considered in [60], whether unikernels could help address some of the key weaknesses introduced by use of the Internet of Things (IoT). In each case, we build on the work of the previous papers, in order to ensure we do not miss anything important as we develop the system.

Unikernels run natively on cloud, they have an exceptionally small footprint, they run without many of the conventional "toys" associated with normal web based cloud instances. This means a seriously minimal attack surface. They are lightweight, can be activated "on demand", and are extremely difficult to attack. Virtually every single conventional attack fails due to there being a heavily restricted means of accessing the running unikernels. A typical cloud instance will be over 150MB in size. Even Docker containers will be a minimum of 24MB in size, whereas a unikernel instance can be as little as 2MB in size. This approach is therefore of interest to us in working towards a good solution to the problem.

Given the limitation we face in terms of most software being insecure, how can we approach developing a potential solution for this problem? In [61] [62] Duncan and Whittington proposed that all cloud based systems which would be subject to compliance under the GDPR, should have ALL audit trails and forensic logs captured and stored off-site in a highly secure immutable database running on a dedicated and highly secure server. These proposals also suggested the immutable database be set up off-site from the cloud instance. This solution has the advantage that the data is not available on the running cloud instance for an attacker to try to compromise, leading to a more secure approach.

While we accept that advice might be highly appropriate given the pervasive extent of the cloud forensic problem, could there be any other way that we might be able to find a cloud based solution? As we shall see in the next section, there may be a way to achieve just that objective.

IV. Finding a Cloud Based Solution

We certainly do accept the sensible logic proposed by Duncan and Whittington [61] [62] to keep the immutable database separated from all running cloud instances. While that makes perfect sense, there is no reason, other than the cloud forensic problem, why the immutable database should not run on a cloud system. However, we do agree that it should not run on the same system as the company system it is trying to protect.

We are keen to explore the idea of running a system on cloud, since that will have the attraction of having all the characteristics that make cloud an interesting proposition for enterprises to use. It provides an agile way to match demand needs to the supply of resources, which can be acquired on demand. It is highly flexible and infinitely scalable. When provisioned by a serious CSP, it is likely to be much more secure than a conventional distributed network system that has been poorly configured. It is also a revenue expense, which can be advantageous for fiscal reasons.

So the question we must now address is how we might go about solving this particular problem. This is where the unikernel based system might be able to help.

Let us first consider the advantages from a security point of view of unikernels:

- The larger a piece of software, the more vulnerabilities are usually present. As we already stated, a unikernel instance can be as little as 2MB;
- The smaller an instance is, the faster a new instance loads;
- The smaller instances are, the cheaper they are to run;
- There is no terminal to log into. The terminal presents one of the easiest attack routes into any system and is usually not well protected from attack. If the attacker cannot log in, achieving a successful attack will be rather difficult to perpetrate;
- The running instance of any unikernel runs with immutable code, meaning no user may inject code into the running unikernel instance.

And now, let us look at any potential disadvantages of unikernels:

- No terminal to log into — a disadvantage for most sys admins. We view this as a huge advantage. If the sys admin cannot login, the attacker has no chance of doing so;
- The running instance is immutable, so it cannot handle changes. We view this as a positive. We are particularly keen to be able to log all changes, system, forensic and audit trail data in a persistent and immutable storage medium off-site. If we cannot change anything, neither can the intruder.

In our view, every item in the above list of advantages and disadvantages all present positive attributes. From a performance, cost, reduced latency and minimised attack surface perspective, all the attributes are highly beneficial for our purposes. This provides us with a degree of confidence that we might be moving on the right track to find a workable solution.

In the next section, we will look at how we might set about developing a system to deploy these instances in a suitable manner that might help us to solve our security challenge.

V. Outline Technical Solution Proposed

We have seen that our unikernel instances can be extremely lightweight, are immutable in operation, have a very small attack surface, perform well, are cheap to run with reduced latency. Because of these advantages, we can use a number of these instances to build a much more robust system.

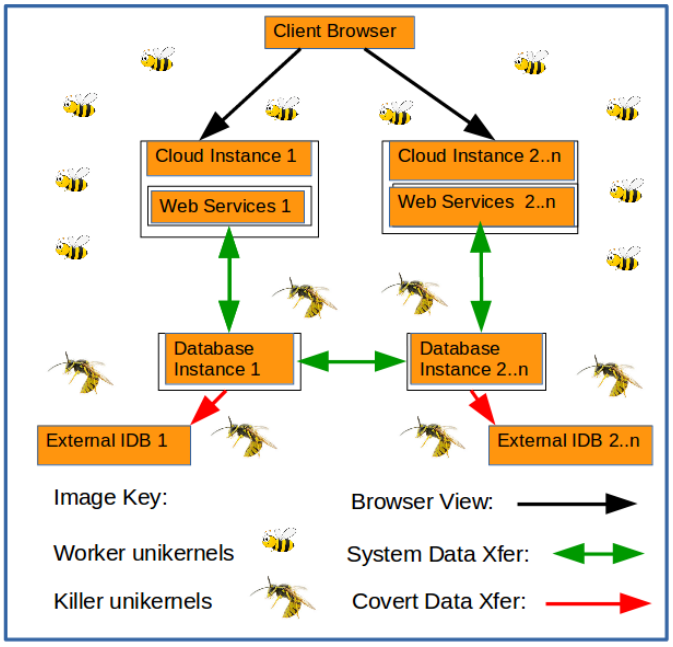

Figure 1: A Unikernel Based Solution to the Cloud Forensic Problem.

If we use the analogy of a bee hive, we can apply this approach as part of our solution. In a bee hive, there are a number of specialised bees — there is a single queen bee, hundreds of male drones (whose responsibility is to mate with a queen, after which they die), anything up to 80,000 female worker bees, who look after developing eggs, larvae and pupae, as well as the whole hive, gathering food from flowers outside the hive and defence duties, which they perform to the death, if needed. Each bee performs a specialised function depending on its age. And in the event a queen leaves, gets lost, or dies accidentally, the colony is capable of generating new queens, either full queens, or temporary queens. The ultimate in sustainability.

Our main company system will have a presence on a cloud platform, using one or more cloud instances as needed, which will be running on a conventional cloud setup. The cloud instance will have the capability to replicate at scale as demand increases and also to shut down instances when demand falls. The main cloud instance system will not be able to be shut down from within. We shall call this the front end Cloud Instance 1.

A conventional database management system will be included in all cloud instances in the normal way except they will instead be removed from within these instances and will run inside a single instance with every non-required function removed from that running instance in order to reduce the attack surface. Should database replication be later required, this can be accommodated through setting up similar database instances. We shall call this original Database Instance 1.

Thus Database Instance 1 will only accept input from the known running front end Cloud Instance 1. There will be no direct access allowed from outside the cloud environment. In the event that replication is required, Cloud Instance 1 will setup as many replicated instances as needed, including

Database Instance 2..n, which will all be replicated, expanding to deliver the required resources.

Worker unikernels will be assigned to each Cloud Instance as they are spooled up, and shut down as no longer needed. They will have specific tasks to perform, such as policing, audit, or whatever. Killer unikernels will be assigned to the task of protecting database systems. Their primary goal will be to ensure the safety of both the forensic trail and the audit trail for all database components, which will be safely stored in the immutable database. These records cannot be deleted. If required, these killer unikernels can turn on attackers trying to breach the systems. All unikernel instances will be tracked, with forensic data collected also for them.

As we can see, each different type of instance is specialised, sticking to its own designated tasks. So what is special about this, apart from splitting up the functions? When a cloud instance runs with a variety of different types of software running on it, this can present a big challenge to configure the overall package in a secure way. By specialising each instance, it becomes much easier to configure securely, because every single unused port can be shut down. Security controls can focus on only what they have to, thus cutting down the potential attack surface.

Any new front end instance, if not registered with the control instance, will not be allowed access to any database instance. Likewise where any new database instance is not registered with the control instance, the front end instances will refuse to connect with it.

The secure immutable database for storing system logs, forensic and audit trail data should not be directly visible to the client browser. Each running instance will send a copy of all system logs, forensic and audit trail data to the immutable database instance as it is generated. The source and timing of all events will be logged by the immutable database.

With the unikernel instances, because they are so lightweight, we can deploy as many of them as we need to carry out very specific tasks. We can have some to police various events, some to carry out audit tasks, some to keep a track of what is live within the system. Each of the components of the main system can be looked after by a number of dedicated unikernel instances, whose sole function will be dedicated to looking after the one conventional cloud instance. Since these unikernels are self sufficient, there is unlikely to be any real adverse impact on the running main instances.

Figure 1 shows a cross-section of the proposed solution. The Client browser will see the front end which provides conventional running cloud instances, with controllers hidden behind the scenes. These controllers can be protected by 'killer bee' unikernels. The external Immutable Database instances will be hosted elsewhere, and can also be protected by 'killer bee' unikernels. The 'worker bee' unikernels clustering around the conventional cloud instances will carry out normal policing and other required tasks. Additional 'bee workers' of whatever kind needed can be spooled up as required. They are fast to provision, take little space and will carry out their programmed task.

As to the question of how many of each type of unikernel we should aim to use, we believe that it would be pointless to speculate at this stage until we can test what will be optimal after we carry out some live experimentation to establish what works well in various loading scenarios. With the use of proper control systems, we can ensure that each new instance is properly registered, constantly and properly monitored, with the control system having the capability to spool up new instances as needed qhickly and efficiently, as well as shutting down those which are no longer required. We expect that such flexibility will allow a highly scalable system to be developed, which can offer minimal running cost, in conjunction with a low latency approach to dealing with attacks. This testing will form part of our future work.

VI. Possible Attack Vectors to Consider

Since we are mostly working with web services, we will look at the Open Web Application Security Project (OWASP) 2017 Top 10 Web Vulnerabilities [63]. We choose these, because they represent the top 10 vulnerabilities with the biggest financial impact on web user systems.

A1:2017-Injection Vulnerability: Injection flaws, such as Structured Query Language (SQL), Not Only SQL (NoSQL), Operating System (OS) injection and Lightweight Directory Access Protocol (LDAP) injection, occur when untrusted data is sent to an interpreter as part of a command or query. The attacker's hostile data can trick the interpreter into executing unintended commands or accessing data without proper authorization. **Solution:** Use a strong Application Programming Interface (API), separate content from commands in the database, and sanitise **ALL** user input.

A2:2017-Broken Authentication Vulnerability: Application functions related to authentication and session management are often implemented incorrectly, allowing attackers to compromise passwords, keys, or session tokens, or to exploit other implementation flaws to assume other users' identities temporarily or permanently. **Solution** Implement multi-factor authentication; no default passwords, especially from admins; reject all top 10,000 worst passwords.

A3:2017-Sensitive Data Exposure Vulnerability: Many web applications and APIs do not properly protect sensitive data, such as financial, healthcare, and PII. Attackers may steal or modify such weakly protected data to conduct credit card fraud, identity theft, or other crimes. Sensitive data may be compromised without extra protection, such as encryption at rest or in transit, and requires special precautions when exchanged with the browser. **Solution:** Encrypt all PII.

A4:2017-XML External Entities (XXE) Vulnerability: Many older or poorly configured eXtensible Markup Language (XML) processors evaluate external entity references within XML documents. External entities can be used to disclose internal files using the file Uniform Resource Identifier (URI) handler, internal file shares, internal port scanning, remote code execution, and denial of service attacks. **Solution:** Whenever possible, use less complex data formats such as JavaScript Object Notation (JSON), and avoiding serialization of sensitive data.

A5:2017-Broken Access Control Vulnerability: Restrictions on what authenticated users are allowed to do are often not properly enforced. Attackers can exploit these flaws to access unauthorized functionality and/or data, such as access other users' accounts, view sensitive files, modify other users' data, change access rights, etc. **Solution:** With the exception of public resources, deny by default; no unrestricted access to users; log all failures.

A6:2017-Security Misconfiguration Vulnerability: Security misconfiguration is the most commonly seen issue. This is commonly a result of insecure default configurations, incomplete or ad hoc configurations, open cloud storage, misconfigured Hypertext Transfer Protocol (HTTP) headers, and verbose error messages containing sensitive information. Not only must all operating systems, frameworks, libraries, and applications be securely configured, but they must be patched/upgraded in a timely fashion. **Solution:** Secure installation processes should be implemented. Keep it simple and log all errors.

A7:2017-Cross-Site Scripting (XSS) Vulnerability: XSS flaws occur whenever an application includes untrusted data in a new web page without proper validation or escaping, or updates an existing web page with user-supplied data using a browser API that can create Hyper Text Markup Language (HTML) or JavaScript. XSS allows attackers to execute scripts in the victim's browser which can hijack user sessions, deface web sites, or redirect the user to malicious sites. **Solution:** Preventing XSS requires separation of untrusted data from active browser content.

A8:2017-Insecure Deserialization Vulnerability: Insecure de-serialization often leads to remote code execution. Even if de-serialization flaws do not result in remote code execution, they can be used to perform attacks, including replay attacks, injection attacks, and privilege escalation attacks. **Solution:** The only safe architectural pattern is not to accept serialized objects from untrusted sources or to use serialization mediums that only permit primitive data types.

A9:2017-Using Components with Known Vulnerabilities Vulnerability: Components, such as libraries, frameworks, and other software modules, run with the same privileges as the application. If a vulnerable component is exploited, such an attack can facilitate serious data loss or server takeover. Applications and APIs using components with known vulnerabilities may undermine application defenses and enable various attacks and impacts. **Solution:** There should be a patch management process in place to ensure known vulnerabilities are never used.

A10:2017-Insufficient Logging & Monitoring Vulnerability: Insufficient logging and monitoring, coupled with missing or ineffective integration with incident response, allows attackers to further attack systems, maintain persistence, pivot to more systems, and tamper, extract, or destroy data. Most breach studies show time to detect a breach is over 200 days, typically detected by external parties rather than internal processes or monitoring. **Solution:** This paper is all about solving this problem!

And for no 11 of 10, go check out your site and make sure your system is not vulnerable.

There are, of course, many more vulnerabilities you can check out, and you should. The more you eliminate, the stronger and more robust your system becomes. You can be sure the attacker already knows all the potential vulnerabilities, so you need to make sure you do too, and plug them.

VII. Discussion on Outline Technical Solution

We strongly believe that a unikernel based system would have a positive and robust impact because of the extra muscle offered to check and log everything that is happening within the system. Given that unikernel instances have a very low attack surface, no conventional attacker 'toys', are immutable in operation, and highly compact, as well as everything being logged to the immutable database - we are cutting out a huge range of vulnerabilities from existing cloud systems. By ensuring the cloud instance running can also withstand the OWASP top ten web vulnerability test, we are in a very strong position to resist a great many attacks.

Some experimentation will be required to identify what the optimal setup of the 'unikernel hive' instances will be in order to obtain the most effective approach. We need to ensure the controller instances are efficiently organised to allow scalability of the overall cloud installation, while at the same time ensuring maximum security and privacy can be achieved. At this time, the Cloud Forensic Problem means that conventional cloud systems cannot guarantee GDPR compliance for cloud users. Container based solutions are likely to be subject to the same issues as conventional cloud instances. While they may very well offer some improvement, it is likely that improvement will come at an overhead cost.

Using the unikernel approach, it is likely that it will certainly be possible to be compliant with the GDPR, that the overhead of running the unikernel instances will be minimal, and that the system can be highly responsive to the need for scalability. Not only that, but the ability to provide a means for compliance for cloud systems has to be big improvement on the status quo.

While we have carried out a number of minor tests on various aspects of our proposal, we have yet to carry out any serious testing, which forms the main thrust of our next stage of the work. We have built a 'cloud in a box' with which to carry out extensive testing of our proposed system. The hardware comprises a Xeon server, running a fast Xeon processor, 16GB of fast RAM and a 525GB SSD drive, together with a 4TB fast storage drive. On this, we have loaded an HP Eucalyptus full cloud management system, which is also Amazon Web Services (AWS) compatible.

This will initially host a conventional web based system to use as a control. We will then run a system based on the proposals contained within this article. Then, we will conduct a series of typical attacks on each of the systems, and will log and analyse the results. We believe that this testing will amply support our belief that this approach will not only prove feasible, but also highly robust against attack.

Having considered how an outline technical solution might be developed, and assessing its feasibility, we now turn to the second need, namely detection of anomalous behaviour, which we address in the next section.

VIII. Detection of Anomalous Behaviour

Following the successful implementation of the solution to retaining full and proper details of the forensic and audit trails, we can now consider how we might go about detecting anomalous behaviour. Since we will now be dealing with a complete data set, then we will have a worthwhile task that we can now set about performing. Obviously, without a full forensic and audit trail available to us, it would seem rather a pointless exercise to analyse incomplete logs to attempt to detect anomalous behaviour. However, with a complete data set, this will prove to be much more worthwhile and meaningful exercise.

The common approach on this problem is often by using technical means alone. This is frequently expressed as policies authorising some action or other. However, the business architecture of an enterprise comprises a combination of people, process and technology [64], not technology alone. Such solutions are generally doomed to failure, as suggested by Duncan and Whittington in [65]–[68], who note such approaches ignore the impact of people and process on security. Both people and process are generally considered to be the weakest link in the business architecture of any enterprise.

However, in this case, we believe that to introduce people and process to the mix at this stage would be counterproductive. First, the scale of the transactional volume can be potentially enormous. Second, the work of analysis would be exceptionally boring, leading to the possibility of mistakes. Third, the introduction of people and process at this stage could lead to both errors and potential corruption, which we must consider as a large potential weakness to the system. Thanks to the robust nature of our proposed solution, we believe in this particular case, we can leave out the intervention of people and process. Naturally, the output from the system would be passed to humans for consideration and investigation, but we are confident that the analysis work on detection of anomalous behaviour could properly be performed without human intervention at this point.

We favour a straightforward approach, such as the soft security approach proposed by Neovius and Duncan [69]. In this approach, they proposed a theoretical framework that could address the highly complex challenge of securing cloud based accounting systems, which are notoriously difficult to secure properly. This would work in conjunction with an immutable database to ensure there could be no loss of audit trail or forensic records.

There is no doubt that inspecting and analysing server logs would present a very effective way to monitor what is happening with any system. Equally, there is no doubt that many companies fail to perform this rather mundane task. Usually, this comes down to a question of huge volume of transactional data, the boredom of manually analysing this data and the opportunity for errors and possible corruption due to the human input.

We suggest that leaving humans out of the main loop here would allow the work to be performed by a suitable algorithm, without the potential corrupting influence of the human input, leading to a better quality of output, performed more accurately and far more quickly. This could potentially lead to faster identification of a breach being perpetrated, thus leading to catching the culprits far more quickly and eliminating their presence from the system. Providing that appropriate forensic and audit trail data has been properly preserved, then it may be possible to ensure sufficient data is collected to assist a possible prosecution of the culprit.

There is no doubt that these tasks could also be provisioned to run on unikernels, leading to a more efficient use of resources. There is also no doubt that this is a task that cannot be left out. Analysis of server logs is one of the key ways to determine whether a breach has occurred, hopefully accompanied by sufficient forensic records to be able to do something about it. At the very least, there will be a very early warning about the possibility of an intrusion, and also the prospect of identifying what damage has been done in respect of GDPR compliance.

Providing a means of being able to identify what data has been compromised is a vital part of the armoury in mitigating the level of fines for non compliance of the GDPR. Anything we can do to ensure this can be achieved will be a good thing. Anything that can be done efficiently will be a bonus.

IX. Conclusion and Future Work

As we have already stated, the Cloud Forensic Problem presents a very serious challenge for all cloud users, especially in light of the forthcoming GDPR. We have proposed a possible solution for this problem, which is a little different from conventional approaches. However, it offers a highly robust solution to a major challenge for all organisations who will be subject to compliance with the GDPR.

We believe this solution offers such merit that we plan to run a pilot test to establish just how well it will be able to cope with a system under serious attack. Initially, it will run on a private network, under attack from professional penetration testers. Once we are sure of how well the solution is likely to perform, we will set up a real live cloud instance to see just how well it might perform.

When the GDPR comes on stream, there will not be time for organisations to mess about. If they cannot comply with the regulation, and they are breached, resulting in a loss of PII, then they can expect huge fines, the like of which they have never seen before. It is time to wake up and smell the coffee.

References

[1] B. Duncan, A. Happe, and A. Bratterud, “Using Unikernels to Address the Cloud Forensic Problem and help Achieve EU GDPR Compliance,” in Cloud Computing 2018: The Ninth International Conference on Cloud Computing, GRIDs, and Virtualization, no. February. Barcelona, Spain: IARIA, 2018, pp. 71–76.

[2] EU, “EU General Data Protection Regulation (GDPR),” 2017. [Online]. Available: http://www.eugdpr.org/ [Last accessed: August 2018]

[3] EU, “Reform of EU data protection rules,” 2016. [Online]. Available: http://ec.europa.eu/justice/data-protection/reform/index_en.htm [Last accessed: August 2018]

[4] N. Kratzke and P.-C. Quint, “Understanding cloud-native applications after 10 years of cloud computing-a systematic mapping study,” Journal of Systems and Software, vol. 126, 2017, pp. 1–16.

[5] N. Kratzke, P.-C. Quint, D. Palme, and D. Reimers, “Project cloud transit-or to simplify cloud-native application provisioning for smes by integrating already available container technologies,” Europea n Project Space on Smart Systems, Big Data, Future Internet-Towards Serving the Grand Societal Challenges. SCITEPRESS. In print, 2017.

[6] M. Almorsy, J. Grundy, and I. Müller, “An analysis of the cloud computing security problem,” 17th Asia-Pacific Software Engineering Conference (APSEC 2010) Cloud Workshop,Sydney, Australia, no. December, 2010, pp. 7.

[7] M. Almorsy, J. Grundy, and I. Mller, “An analysis of the cloud computing security problem.” The proc. of the 2010 Asia Pacific Cloud Work-shop Colocated with APSEC2010, Australia, 2010.

[8] V. Chang, Y. H. Kuo, and M. Ramachandran, “Cloud computing adoption framework: A security framework for business clouds,” Future Generation Computer Systems, vol. 57, 2016, pp. 24–41.

[9] F. Doelitzscher, T. Ruebsamen, T. Karbe, M. Knahl, C. Reich, and N. Clarke, “Sun Behind Clouds - On Automatic Cloud Security Audits and a Cloud Audit Policy Language,” International Journal on Advances in Networks and Services, vol. 6, no. 1, 2013, pp. 1–16.

[10] R. K. L. Ko, P. Jagadpramana, M. Mowbray, S. Pearson, M. Kirchberg, B. S. Lee, and Q. Liang, “TrustCloud: A Framework for Accountability and Trust in Cloud Computing,” Perspective, 2011, pp. 1–9.

[11] R. K. L. Ko, B. S. Lee, and S. Pearson, "Towards achieving accountability, auditability and trust in cloud computing," Communications in Computer and Information Science, vol. 193 CCIS, 2011, pp. 432–444.

[12] K. Lee, "Security Threats in Cloud Computing Environments," International Journal of Security and its Applications, vol. 6, no. 4, 2012, pp. 25–32.

[13] N. Papanikolaou, S. Pearson, and M. C. Mont, "Towards Natural-Language Understanding and Automated Enforcement of Privacy Rules and Regulations in the Cloud: Survey and Bibliography," Analysis, 2011, pp. 1–9.

[14] S. Pearson, "Taking Account of Privacy when Designing Cloud Computing Services," Current, 2009, pp. 44–52.

[15] S. Pearson, "Toward accountability in the cloud," IEEE Internet Computing, vol. 15, no. 4, jul 2011, pp. 64–69.

[16] S. Ramgovind, M. Eloff, and E. Smith, "The Management of Security in Cloud Computing," in Information Security for South Africa (ISSA), 2010, 2010, pp. 1–7.

[17] H. Takabi, J. B. Joshi, and G.-J. Ahn, "Security and Privacy Challenges in Cloud Computing Environments," IEEE Security & Privacy Magazine, vol. 8, no. 6, nov 2010, pp. 24–31.

[18] S. Creese, P. Hopkins, S. Pearson, and Y. Shen, "Data Protection-Aware Design for Cloud Computing," Work, no. December, 2009, pp. 1–13.

[19] K. Hon, C. Millard, and I. Walden, "The Problem of 'Personal Data' in Cloud Computing - What Information is Regulated ?" 2011.

[20] W. K. Hon, C. Millard, and I. Walden, "Who is Responsible for 'Personal Data' in Cloud Computing? The Cloud of Unknowing, Part 2," Legal Studies, no. 77, 2011, pp. 1–31.

[21] W. K. Hon, J. Hörnle, and C. Millard, "Data Protection Jurisdiction and Cloud Computing When are Cloud Users and Providers Subject to EU Data Protection Law?" Legal Studies, 2011, pp. 1–40.

[22] W. Itani, A. Kayssi, and A. Chehab, "Privacy as a Service: Privacy-Aware Data Storage and Processing in Cloud Computing Architectures," 2009 Eighth IEEE International Conference on Dependable, Autonomic and Secure Computing, dec 2009, pp. 711–716.

[23] W. Jansen and T. Grance, "Guidelines on Security and Privacy in Public Cloud Computing," NIST, Tech. Rep., 2011. [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-144/SP800-144.pdf [Last accessed: August 2018]

[24] H. Katzan Jr, "On The Privacy Of Cloud Computing," International Journal of Management and Information Systems, vol. 14, no. 2, 2011, pp. 1–12.

[25] W. K. Hon, C. Millard, J. Singh, I. Walden, and J. Crowcroft, "Policy, legal and regulatory implications of a Europe-only cloud," International Journal of Law and Information Technology, vol. 24, no. 3, 2016, pp. 251–278.

[26] NIST, "Security and Privacy Controls for Federal Information Systems and Organizations Security and Privacy Controls for Federal Information Systems and Organizations," Natioinal Institute of Standards and Technology, Gaithersburg, MD, Tech. Rep. February, 2014. [Online]. Available: http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r4.pdf [Last accessed: August 2018]

[27] S. Pearson and A. Charlesworth, "Accountability as a Way Forward for Privacy Protection in the Cloud," Computing, no. December, 2009, pp. 1–15.

[28] S. Pearson and A. Benameur, "Privacy, Security and Trust Issues Arising from Cloud Computing," 2010 IEEE Second International Conference on Cloud Computing Technology and Science, Nov 2010, pp. 693–702.

[29] S. Pearson, "Privacy, Security and Trust in Cloud Computing," in Privacy and Security for Cloud Computing. e: Springer, 2013, pp. 3–42.

[30] J. Prüfer, "How to govern the cloud? Characterizing the optimal enforcement institution that supports accountability in cloud computing," Proceedings of the International Conference on Cloud Computing Technology and Science, CloudCom, vol. 2, 2013, pp. 33–38.

[31] S. S. Shapiro, "Privacy by Design," Communications of the ACM, vol. 53, no. 6, jun 2010, p. 27.

[32] J. Singh, T. F. J. M. Pasquier, and J. Bacon, "Securing tags to control information flows within the Internet of Things," 2015 International Conference on Recent Advances in Internet of Things, RIoT 2015, 2015.

[33] EU, "Accountability for Cloud (A4Cloud)," 2018. [Online]. Available: http://a4cloud.eu/ [Last accessed: August 2018]

[34] C. A. Adams and R. Evans, "Accountability, Completeness, Credibility and the Audit Expectations Gap," JCC 14 Summer 2014, vol. 14, no. Summer, 2014, pp. 97–115.

[35] W. Benghabrit, H. Grall, J.-C. Royer, M. Sellami, M. Azraoui, K. Elkhiyaoui, M. Onen, A. S. D. Olivera, and K. Bernsmed, "A Cloud Accountability Policy Representation Framework," in CLOSER-4th International Conference on Cloud Computing and Services Science, 2014, pp. 489–498.

[36] K. Bernsmed, W. K. Hon, and C. Millard, "Deploying Medical Sensor Networks in the Cloud Accountability Obligations from a European Perspective," in Cloud Computing (CLOUD), 2014 IEEE 7th International Conference on. IEEE Comput. Soc, 2014, pp. 898–905.

[37] D. Butin, M. Chicote, and D. Le Métayer, "Log design for accountability," in Proceedings - IEEE CS Security and Privacy Workshops, SPW 2013, 2013.

[38] D. Catteddu, M. Felici, G. Hogben, A. Holcroft, E. Kosta, R. Leened, C. Millard, M. Niezen, D. Nunez, N. Papanikolaou, S. Pearson, D. Pradelles, C. Reed, C. Rong, J.-C. Royer, D. Stefanatou, and T. W. Wlodarczyk, "Towards a Model of Accountability for Cloud Computing Services," in International Workshop on Trustworthiness, Accountability and Forensics in the Cloud (TAFC), 2013, pp. 21–30.

[39] A. Haeberlen, "A Case for the Accountable Cloud," ACM SIGOPS Operating Systems Review, vol. 44, no. 2, 2010, pp. 52–57.

[40] W. Hon, E. Kosta, C. Millard, and D. Stefanatou, "Cloud Accountability: The Likely Impact of the Proposed EU Data Protection Regulation," Queen Mary School of Law Legal Studies Research Paper, no. 172, 2014, pp. 1–54.

[41] K. L. R. Ko, P. Jagadpramana, and B. S. Lee, "Flogger: A File-Centric Logger for Monitoring File Access and Transfers within Cloud Computing Environments," Computing, 2011, pp. 1–8.

[42] T. Lynn, P. Healy, R. McClatchey, J. Morrison, C. Pahl, and B. Lee, "The Case for Cloud Service Trustmarks and Assurance-as-a-Service," in CLOSER 2013 - Proceedings of the 3rd International Conference on Cloud Computing and Services Science, 2013, pp. 110–115.

[43] N. Papanikolaou, S. Pearson, M. C. Mont, and R. K. L. Ko, "Towards Greater Accountability in Cloud Computing through Natural-Language Analysis and Automated Policy Enforcement," Engineering, 2011, pp. 1–4.

[44] N. Papanikolaou, T. Rübsamen, and C. Reich, "A Simulation Framework to Model Accountability Controls for Cloud Computing," CLOUD COMPUTING 2014, The Fifth International Conference on Cloud Computing, GRIDs, and Virtualization, no. c, 2014, pp. 12–19.

[45] S. Pearson, M. C. Mont, and G. Kounga, "Enhancing Accountability in the Cloud via Sticky Policies," in Secure and Trust Computing, Data Management, and Applications, 2011, pp. 146–155.

[46] S. Pearson, V. Tountopoulos, D. Catteddu, S. Mario, R. Molva, C. Reich, S. Fischer-Hubner, C. Millard, V. Lotz, M. G. Jaatun, R. Leenes, C. Rong, and J. Lopez, "Accountability for Cloud and Other Future Internet Services," in CloudCom, 2012, pp. 629—-632.

[47] K. Bernsmed and S. Fischer-Hübner, "Secure IT Systems: 19th Nordic Conference, NordSec 2014 Tromsø, Norway, October 15-17, 2014 Proceedings," Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), vol. 8788, 2014, pp. 3–24.

[48] T. Ruebsamen and C. Reich, "Supporting Cloud Accountability by Collecting Evidence Using Audit Agents," in CloudCom 2013, 2013, pp. 185–190.

[49] J. Singh, J. Powles, T. Pasquier, and J. Bacon, "Data Flow Management and Compliance in Cloud Computing," Cloud Computing, no. Special Issue on Legal Clouds., 2015, pp. 1–12.

[50] A. Squicciarini, S. Sundareswaran, D. Lin, and S. Huang, "Promoting Distributed Accountability in the Cloud," in IEEE 4th International Conference on Cloud Computing Promoting, 2011, pp. 113–120.

[51] S. Sundareswaran, A. C. Squicciarini, and D. Lin, "Ensuring Distributed Accountability for Data Sharing in the Cloud," in IEEE Transactions on Dependable and Secure Computing, vol. 9, no. 4, 2012, pp. 556–568.

[52] M. Theoharidou, N. Papanikolaou, S. Pearson, and D. Gritzalis, "Privacy risk, security, accountability in the cloud," in Proceedings of the

International Conference on Cloud Computing Technology and Science, CloudCom, vol. 1, 2013, pp. 177–184.

[53] Verizon, "2012 Data Breach Investigation Report: A study conducted by the Verizon RISK Team in cooperation with the United States Secret Service and Others," Tech. Rep., 2012. [Online]. Available: http://www.verizonenterprise.com/resources/reports/rp_data-breach-investigations-report-2012_en_xg.pdf [Last accessed: August 2018]

[54] Verizon, "2016 Verizon Data Breach Report," Tech. Rep., 2016.

[55] A. Bratterud, A.-A. Walla, H. Haugerud, P. E. Engelstad, and K. Begnum, "IncludeOS: A Minimal, Resource Efficient Unikernel for Cloud Services," 2015 IEEE 7th International Conference on Cloud Computing Technology and Science (CloudCom), 2015, pp. 250–257.

[56] N. Kratzke, "About microservices, containers and their underestimated impact on network performance," arXiv preprint arXiv:1710.04049, 2017.

[57] B. Duncan, A. Bratterud, and A. Happe, "Enhancing Cloud Security and Privacy: Time for a New Approach?" in Intech 2016, Dublin, 2016, pp. 1–6.

[58] A. Bratterud, A. Happe, and B. Duncan, "Enhancing Cloud Security and Privacy: The Unikernel Solution," in Cloud Computing 2017: The Eighth International Conference on Cloud Computing, GRIDs, and Virtualization, 2017, pp. 1–8.

[59] A. Happe, B. Duncan, and Alfred Sewitsky Bratterud, "Unikernels for Cloud Architectures: How Single Responsibility can Reduce Complexity, Thus Improving Enterprise Cloud Security," in COMPLEXIS 2017 - Proceedings of the 2nd International Conference on Complexity, Future Information Systems and Risk, Porto, Portugal, 2017, pp. 1–12.

[60] B. Duncan, A. Happe, and A. Bratterud, "Enterprise IoT Security and Scalability: How Unikernels can Improve the Status Quo," in 9th IEEE/ACM International Conference on Utility and Cloud Computing (UCC 2016), Shanghai, China, 2016, pp. 1–6.

[61] B. Duncan and M. Whittington, "Creating an Immutable Database for Secure Cloud Audit Trail and System Logging," in Cloud Computing 2017: The Eighth International Conference on Cloud Computing, GRIDs, and Virtualization. Athens, Greece: IARIA, ISBN: 978-1-61208-529-6, 2017, pp. 54–59.

[62] B. Duncan and M. Whittington, "Creating and Configuring an Immutable Database for Secure Cloud Audit Trail and System Logging," International Journal On Advances in Security, vol. 10, no. 3&4, 2017, pp. 155–166.

[63] OWASP, "OWASP home page," 2017. [Online]. Available: https://www.owasp.org/index.php/Main_Page [Last accessed: August 2018]

[64] PWC, "UK Information Security Breaches Survey - Technical Report 2012," PWC2012, Tech. Rep. April, 2012.

[65] B. Duncan, D. J. Pym, and M. Whittington, "Developing a Conceptual Framework for Cloud Security Assurance," in 2013 IEEE 5th International Conference on Cloud Computing Technology and Science. Bristol: IEEE, 2013, pp. 120–125.

[66] B. Duncan and M. Whittington, "Compliance with Standards, Assurance and Audit: Does this Equal Security?" in Proceedings of the 7th International Conference on Security of Information and Networks. Glasgow: ACM, 2014, pp. 77–84.

[67] B. Duncan and M. Whittington, "The Importance of Proper Measurement for a Cloud Security Assurance Model," in 2015 IEEE 7th International Conference on Cloud Computing Technology and Science (CloudCom), Vancouver, 2015, pp. 1–6.

[68] B. Duncan and M. Whittington, "Information Security in the Cloud: Should We be Using a Different Approach?" in 2015 IEEE 7th International Conference on Cloud Computing Technology and Science (CloudCom), Vancouver, 2015, pp. 1–6.

[69] M. Neovius and B. Duncan, "Anomaly Detection for Soft Security in Cloud based Auditing of Accounting Systems," in Closer 2017 - 7th International Conference on Cloud Computing and Services Science, Porto, Portugal, 2017, pp. 1–8.
[Last accessed: August 2018]

The Complexities of Auditing and Securing Systems in the Cloud — is there a Solution and will the GDPR move it up the Corporate Agenda?

Bob Duncan*, Mark Whittington†
Business School
University of Aberdeen
Aberdeen, UK
Emails: *robert.duncan@abdn.ac.uk, †mark.whittington@abdn.ac.uk

***Abstract*—It would seem that some companies have been slow or unable to secure their cloud activities or to be aware of breaches in a timely manner. The European Union (EU)s General Data Protection Regulation (GDPR) has been introduced with the intent of sufficient threat of meaningful fines that directors will now take cloud security seriously, even if they had not perceived it as a strategic priority before. However, just introducing such penal incentives does not mean that solutions are easy to implement. Whilst the perfect solution would always include stopping attackers from becoming intruders, once the attacker gets access the challenge is not just the immediate fiscal damage to the company or its trading partners, but also to the very records and integrity of the databases themselves. Once the intruder gains a foothold, they may then be able to grant themselves sufficient privileges to completely delete all trace of their incursion, possibly deleting far more records than they need to. They may remain undetected within the system, accessing, modifying, deleting or ex-filtrating data at will from the victim's system. This is referred to as the Cloud Forensic Problem. This, then, presents a compliance nightmare to a great many cloud users, many of whom are poorly prepared to cope with this serious practical and financial challenge. In this paper, we consider how experience and traditional techniques from the accounting world might be applied and adapted to mitigate this serious challenge.**

Keywords–Forensic audit; GDPR compliance; cloud forensic problem.

I. INTRODUCTION

Achieving information security with conventional distributed network computer systems presents a significant challenge, but this challenge increases exponentially when we introduce cloud computing to the mix, due to the multiplicity and complexity of hardware and software layers and the number of actors with differing agendas, involved in any cloud ecosystem. While this high level of complexity has been a fundamental part of cloud computing, we shall see that the capabilities of cloud computing have evolved considerably beyond what was first envisaged. The principal reason for the difficulty of this challenge is the so called "Cloud Forensic Problem".

The Cloud Forensic Problem arises when an attacker gains a foothold in a cloud system and becomes an intruder. Once this happens, there is little to prevent the intruder from helping themselves to any amount of data, either by viewing, modifying, deleting or ex-filtrating it from the victim system. Worse still, there is nothing to prevent the intruder from gaining sufficient privileges to completely delete all trace of their attack through modifying or deleting entirely the forensic records of the system. In this paper, we consider how the use of forensic audit might help mitigate the impact of this problem based on our earlier work [1].

In addition to the cloud forensic problem, the EU General Data Protection Regulation (GDPR) [2] came into effect on 25th May 2018, and a principal requirement is the protection of any personally identifiable information of any EU resident held by any organisation, anywhere in the world, on pain of severe financial penalties. Given that the cloud forensic problem presents a potentially insurmountable compliance problem, a great many organisations are likely to be exposed to incalculable potential penalties for the string of cyber breaches that are likely to ensue. Full compliance will inevitably pose a challenge for all organisations, but for those using cloud, due to the potential impact of the cloud forensic problem, the challenge will become so much more difficult.

It is too early to speculate on what approach the regulator might take towards setting penalties for breaches, but there is little doubt that where a company has an attitude problem towards proper compliance, or is complicit through poor internal security controls and provisions, then all these factors will be taken into account when gauging the level at which to set any potential fines. Equally, where a company can demonstrate that it has taken proper steps to mitigate the impact of the cloud forensic problem, it is clear that this will have the opposite effect, resulting in considerably lower levels of fines as a consequence of any breach.

We start in Section II, by considering the cloud forensic problem and the challenges it poses. We turn to the accounting world to see which techniques we could implement to help address these serious challenges in Section III, where we look at accounting, audit and forensic accounting to see how it works for the accounting world, and in Section IV, we address the importance of separation of duties. In Section V, we consider how we might develop some of these well established techniques to help us address this significant cloud security problem. In Section VII, we first consider some possible impediments to restoring the 'paper ink' trail. In Section VIII, we look at how we might use the immutable database as the core of this approach. In Section IX, we discuss the implications of this proposed solution. In Section X, we draw our conclusions and discuss our possible future work.

II. The Cloud Forensic Problem

Cloud systems are extremely popular with companies due to the flexibility offered by cloud. Speed of start-up, ease of scalability to match the demand curve and the revenue nature of the costs involved all provide a strong incentive for companies to use cloud services. Cloud computing has been with us now for over 10 years, and while much of the early research concentrated on usability [3] [4] and performance [5]–[7] it was not long before thoughts of security [8]–[10] and privacy [11] [12] started to surface.

While the US National Institute for Standards and Technology (NIST) were one of the first organisations to propose standard definitions [13] [14] interest in security [15]–[18] and privacy [19]–[21] started to grow.

Thoughts also started turning to accountability [9] [22]–[24] given the evolving complexities of cloud ecosystems. This ultimately led the EU to set up the Accountability for Cloud (A4Cloud) Project [25] to consider such important matters. The A4Cloud project drew much attention to the need for proper accountability in cloud systems and the contributors developed many useful mechanisms for ensuring proper levels of accountability could be monitored and achieved.

While there have been some really positive advances in both security and privacy during this time, there remains one fundamental weakness that has not been resolved, namely the "cloud forensic problem". All computer systems connected to the internet are subject to continuous and serious attack, and cloud systems are no exception. It would be realistic to state that no system is immune to attack, and this is particularly true for cloud systems. Attackers will always succeed in gaining entry to systems. The secret of success here is to be able to identify these occurrences the moment they happen, so that the attack can be shut down and the perpetrator removed from the system.

The main focus of an attacker is to breach a system, which can involve a considerable amount of work on their part. The more diligent will first perform surveillance, compile many analyses of how the various company systems are structured and how they interact with each other. Often, they will also carry out huge amounts of work to understand the people of the organization, since they are usually the weak link in the chain [26]. This extensive intelligence gathering will usually cover every conceivable aspect of all company systems to ensure they discover everything they need to know about the company. This is why it is so important for all companies to analyse their system logs, in order to gain a better understanding of who is actually attacking their systems.

Other attackers, will be much less organised, simply trying to hack in to company systems, without a thought of the overview of the company concerned. They will merely look for known vulnerabilities and try to attack them. There are other attackers who will specifically attack the people of the company through social engineering and other similar approaches. The first objective of all attackers is the same — to penetrate the system in order to set up a foothold in the system, thus allowing them to take steps to become an intruder.

The aim is not just to get in, and out, as quickly as possible, but to be able to develop a long term foothold, secreting themselves into corporate servers and other subsidiary systems which will allow them to return time and time again to help themselves to more information whenever they want. The longer they remain in the system, the more they are likely to try to escalate privileges to give them access to more and more possible information. All too often, they are helped along the way by the companies themselves, often through an element of laziness on the part of system administrators [27].

If we look back five years ago, at previous cyber breach reports [28] there was a global average time of 6 months between breach and discovery. With more rigorous attention paid to reading and analysing their server logs, it is obvious they could have discovered intruders much more quickly. By 2016, the time between breach and discovery had dropped to a matter of weeks rather than months [29] however, this is still not good enough to keep on top of what is going on in corporate systems.

Companies often contribute to their own downfall by failing to update security patches to both operating systems and software systems, complexities from legacy applications applications and risks of outages being reasons or excuses for slow implementation [30]. All of these issues conspire to lead inexorably to the, as yet unresolved, cloud forensic problem — namely, that once an intruder is in the system, and has escalated sufficient privileges, there is nothing to prevent them from deleting the forensic trail, all system logs and audit trails, thus hiding all evidence of their successful penetration and of the size and nature of their crime.

Under the GDPR [2] any breached organisation must report the breach within 72 hours of discovery of the breach. They must also report how many relevant records have been compromised, whether by having been read, amended, deleted or ex-filtrated from the system. Given that many system logs are also not turned on by default [31] this means identifying which records have been compromised, whether by having been read, amended, deleted or ex-filtrated, will present a serious enough challenge in the first place.

However, since the intruder will likely have worked hard to increase privileges to the point that they are able to modify or worse, delete all forensic trails in the system, the likelihood of an organisation being able to properly identify which records have been compromised may prove impossible to determine. Often, capturing adequate levels of forensic data does not happen due to many such features being turned off by default. It is bad enough when intruders delete forensic records, but it is inexcusable when an organisation fails to collect them in the first place.

The consequence of failure to detect such intrusions means not only non-compliance with the GDPR, triggering fines, but failure to tackle some elementary steps will then cause these fines to escalate following repeated events to the greater of €20million or 4% of global turnover. The size of the potential fines, along with the bad publicity will surely get the attention of organizations, their managers and all their stakeholders.

III. Useful Techniques from the Accounting World

The process of accounting has been around for millennia, with the underlying standard approach of double entry bookkeeping in use for over 500 years, with the generally accepted story placing its creation in Italy. It can be argued that accounting and the reporting of accounting numbers has

had two overriding purposes that are in tension with each other. The first, that had dominance in earlier centuries [32] is stewardship, though the term itself has seemed to evolve over time [33] (pp. 264) from being the careful, honest and accurate recording of transactions to efficient use of resources to finally ensuring an appropriate return for shareholders.

This progression is dependent on the trust that the earlier definition can now be taken as given due to improvements in recording mechanisms and the outside eye of an auditor. The confidence in the recording mechanism requires a complete history of transactions that means the accounts can be checked and even re-built if necessary whether in the mythical "shoe-box" of receipts for the small business or the sophisticated computerised ledgers of a multinational.

The integrity of the items recorded and the potential value of the detail highlights another concern that the data could be useful to people for whom it was not intended (competitors and fraudsters). Hence the need to lock up the accounting ledgers (or their computerised descendants) to keep them from being corrupted or seen by those who have no right of access.

From the 1950s onwards a more developed theory of accounting and reporting evolved with the focus being on accounting as a technique for collecting, measuring, processing and communicating financial information about the economic performance of entities, in order to provide decision useful information for interested parties, such as management, investors, creditors and regulators [34].

The International Accounting Standards Board (IASB) issued a similar, but more user-constrained definition in 2015, namely "The objective of general purpose financial reporting is to provide financial information about the reporting entity that is useful to existing and potential investors, lenders and other creditors in making decisions about providing resources to the entity. Those decisions involve buying, selling or holding equity and debt instruments, and providing or settling loans and other forms of credit." [35]

"Decision usefulness", particularly for investors, became the central determinant of "good" or "bad" accounting methods, again one could argue, because of a confidence (sometimes misplaced) that stewardship, the basic recording, could be taken for granted.

In the above story of accounting development, we already needed to introduce the term "auditing". Auditing, too, has been around for millennia, as there has always been a need to provide assurance that accounts and financial statements present a "true and fair view", or some similar phrase, of the business under review. Audit, the checking of conformity or of being fit for purpose, takes place in many fields, each of which develop over time and may (or may not) learn appropriate lessons from audit practices that have been honed over decades or centuries in more mature situations or professions .

Hence, not only accounting but also financial auditing techniques can also be applied to any other sphere where there is a need for recording, safety or trust and where there are records and some element of measurement, in this case, of course, we are particular interested in data. Hence, seeking to apply the more evolved and time tested techniques from accounting and auditing to the management and governance of data — and specifically data in the cloud would seem logical.

A further extension of the processes of accounting and audit is forensic (OED [36] "pertaining to, connected with, or used in courts of law; suitable or analogous to pleadings in court") accounting, which as the definition suggests is the process of preparing evidence suitable for use in a court of law, though such approaches are often used without a courtroom on the horizon. Forensic accounting is tuned to expose fraud and manipulation.

We can potentially use these techniques, which have long been developed in the accounting world to good effect in helping us secure our cloud data. We can then liken any database system to an accounting system, whereby we collect, measure, process and communicate data and the information gleaned from it concerning a business to the people for whom it is intended or relevant. Of course, the reliability, and even completeness, of data is a prerequisite for assessing any organisational efficiency level or for decision making.

We can see that the completeness of recording, the trust in the methods of processing the transactions and the ability of an auditor to interrogate the raw transactions are key building blocks for any effective data management system — whether accounting or otherwise focused.

This medium presents the benefit of providing a hard ink trail to follow, something which we shall later see is no longer available with modern cloud systems. This trail of records written in ledgers and of pieces of paper with signatures, comments and account codes provides for even the smallest business a trail of evidence for the accountant or auditor to follow through. The occasional missing item can usually be determined through "incomplete records analysis" as there is a surety concerning the other data and the bank statements. A larger business would think through more streamlined and consistent approaches to record keeping which then evolved into some of the earliest computer records, where (with known hard drives and no internet) anything entered would stay entered with the identity of the person undertaking the transaction, a time stamp and the matching double entry.

In principal, we can then use cloud audit to provide assurance of the data provenance of all the data held in the database system, and in the event of a security breach, we should then be able to easily apply cloud forensic techniques, learning from the accounting world, in order to help us bring about a successful prosecution in the courts and to become aware of the steps needed to improve security for the future. In practice, this, of course, will be far harder to achieve.

Of course, it is worth pointing out that for centuries, accountants have enjoyed the benefits of working with hard copy books, written with quill pen and ink. This medium presents the benefit of providing a hard ink trail to follow, something which we shall later see is no longer available with modern cloud systems. We can learn lessons from the accounting world, specifically in the area of the audit trail, as used with accounting systems for centuries. A further relevant step in business and accounting risk mitigation in the accounting process is separation of duties, and we will now discuss this more fully in the next section.

IV. The Importance of Separation of Duties

For many decades, a key part of the structure of departments and of businesses overall is that of "separation (or

segregation) of duties." This is a simple but straight forward security measure that could be employed by all but the smallest businesses. The logic is to carefully separate out the tasks in a business process so that no one person can have input or control into steps that might give them the opportunity, and temptation, to commit fraud or to effect theft. The smallest business would struggle to achieve this as different employees will be required to be responsible for specific tasks.

Ashton, [37] used a questionnaire to ask auditors a series of questions with the intent of being able to weigh their consistency as they inspected accounts and applied judgement. The first questions in his questionnaire addressed the segregation of tasks —

- Are the tasks of both timekeeping and payment of employees adequately separated from the task of payroll preparation?
- Are the tasks of both payroll preparation and payment of employees adequately separated from the task of payroll bank account?

It is not hard to see the result of a negative answer to either question. In both cases, an employee would be faced with the chance to change numbers in order to benefit themselves or, applying a little cunning, someone else. Involving two or more people may not be perceived to be enough, a further good safety feature would be to site the wages and salaries staff away from most of the workforce, reducing the chance of collaboration on a fraudulent scheme. A further gain from segregation, even when all employees are honest, is the opportunity to spot mistakes — a second person being required to take up the next stage of a process will mean either a clearly defined check or at least a "reasonableness" check on the work done to date.

The implications of judging that the answer to either of these two questions is "no" are obvious — an opportunity and a temptation arises for an individual to manipulate the payroll to their advantage. Clearly if it were possible to locate the payroll department away from the main work location and be confident that no one in payroll knew anyone in the rest of the company, then confidence would be increased yet further. Such separation not only makes fraud difficult, but also means unintentional errors are more likely to be spotted.

According to Gelinas et al. [38] there are four areas in a business or accounting process that need to be separated: authorising transactions, executing transactions, recording transactions and safeguarding resources subsequent to the transactions being completed. Vaassen et al. [39] add a further need for separation with — "authorisation; custody; recording; checking and execution". We move to more direct relevance to our key concern with the work of Hall [40] who addresses segregation when computerised accounting has been implemented. Hall sees that further concerns now need to be added, including "Is the logic of the computer program correct? Has anyone tampered with the application since it was last tested? Have changes been made to the programme that could have caused an undisclosed error?" (page 208). Even such apparently obvious questions need to be asked if the integrity of the system is to be assured..

The Sarbanes-Oxley Act [41] introduced new disclosure requirements for senior directors to commit personally to the quality of their numbers, and by implication the systems that produced those numbers. Ge and McVay [42] found 261 firms in 2002-2004 that admitted weaknesses in control and of these 45 included a reference to separation of duties. A further concern for us is that companies from the computer sector were a noticeably high proportion of the 261 problem companies in Ge and McVays study.

Taking these examples and then applying the same logic to both programming in general to software use is straightforward, though one might still question whether knowledge of this is sufficiently carried through into practice. If practice is indeed good, we still need a record of activity in order for audit and to investigate when things go awry (i.e., the audit trail).

This may all sound like unnecessary work and detail. However, when one of the authors ran a large purchase ledger department, a ledger clerk became confused when processing a number of very similar invoices from one large supplier that totalled to 2 million. She had entered invoices, then entered credit notes to cancel them and then repeated this a number of times before coming to him in some distress. The problem was sorted, but, as the author expected, an auditor spotted this unusual activity some time later and asked for an explanation. The event log and records showed each step, who had carried out which transactions and how the issues were corrected.

V. Forensic Cloud Audit

An interesting distinction in definition between "forensic accounting" and "cloud computing forensic science" is the presence of that last word "science". Hopwood et al. [43] give the following definition for forensic accounting: Forensic accounting is the application of investigative and analytic skills for the purpose of resolving financial issues in a manner that meets standards required by courts of law. Notice that forensic accounting is not limited to the use of financial investigations that result in legal prosecution; however, if this is the purpose, the investigation and analysis must meet the standards required in the court of law that has jurisdiction. (page 3).

Whilst NIST [44] provides the following discussion and definition: Many experts consider forensic science to be the application of a broad spectrum of sciences and technologies to the investigation and establishment of facts of interest in relation to criminal, civil law, or regulatory issues. However, the resulting techniques may also be used for purposes outside the scope of law to reconstruct an event that has occurred. Cloud computing forensic science is the application of scientific principles, technological practices and derived and proven methods to reconstruct past cloud computing events through identification, collection, preservation, examination, interpretation and reporting of digital evidence.

Note that the forensic accounting definition does not include the word science, despite the area (see for example two textbooks Taylor [45] and Hopwood et al. [43]) including scientific methods. Taylor [45] as a more introductory text, focuses initially and at some length on the need to understand background and environmental issues, using this as a backdrop before moving on to, again, a largely discursive review of the wide range of relevant criminal activities that might require the attention of the forensic accountant. He also addresses risk management issues in relation to IT systems, briefly including the cloud, and the process of investigation. Hopwood et al.

[43] have a similar structure but give a little greater weight to forensic science and computer forensics.

From the computer science camp, Choo and Dehghantanha [46] a more scholarly work, reflects a greater weight placed on technical issues, as well as the tools and techniques needed, for forensic cloud investigations. Almulla et al. [47] review the cloud forensic literature and find some discursive, though mostly technical papers. Some of our previous research [48]–[50] has focused on the critical nature of understanding human frailties and interactions as well as what seems the more technically demanding elements of computer science.

Issues requiring computer forensic audit are likely to involve the stealing of money, the stealing of monetizable data or the misrepresentation of data to personal or group advantage. These are areas which accountants have strived to address over decades in less technical and complex settings. It would seem logical that their group learning over time would have some relevance and currency to the new cloud situation.

Like most professions, accountants have well organised professional exams. There are many accounting associations, many with long histories and experience in exam setting. The syllabi of these bodies depends somewhat on the countries in which they operate and are revised over time to reflect new priorities and changes in the world -both technical and social. The Association of Chartered Certified Accountants (ACCA), is a significant international accounting professional body and we will take their professional examination content as an exemplar of others. The ACCA has over 200,000 members [51] and has an exam at its professional stage, Advanced Auditing and Assurance [52] that includes — a section on forensic audit though it should be noted that — it is only a very small part of the content of that exam.

Each of the professional bodies faces a dilemma when revising their syllabi for a changing and ever more complex world. In many countries, the market to attract accounting students is competitive, hence a more complex world cannot lead to more exams and a longer route to qualification without the body facing a competitive disadvantage. It is also very difficult to decide to drop traditional content to make appropriate space for newer material or issues.

It would seem that qualifying accountants are ill-prepared by their professional bodies for the complexities of the cloud environment. This is both in terms of understanding the environmental issues, though there is accessible material for them to pick up some of this (see Taylor [45] and Hopwood et al. [43]), as well as comprehending the technical ones, which would be a far more complex and difficult step.

Whilst there are a few small organisations focusing on forensic accounting and audit, these appear peripheral and it does not seem that many qualified accountants have moved into this more rarefied space by adding years of further learning to their accounting badge. The large accounting “firms”, commonly referred to as the “Big 4” (KPMG, PWC, EY and Deloitte), who audit nearly all the worlds big companies and collectively employ about a million people (2017), do offer forensic services along with a broad range of consultancy services (see [53] for example).

The ever-widening scope of the Big 4, making far more money from consultancy than audit, is contentious in some countries. The reliance on an oligopolistic audit industry with seemingly conflicted aims of professionalism and commercial gain, along with what many see as questionable competence in their core audit activity, is building a crescendo for change Marriage2018 and Marriage2018a highlight some audit quality issues and FRC2018 also includes an example of the consultancy dilemma at paragraph 34). Whilst these firms have undoubtedly built up some expertise in the area, there are some significant issues with their continuing range of activities. Two discussed options, in the UK at least, are either splitting up the firms or to separate the audit arm from consultancy.

From the other direction, computer specialists clearly have an understanding of the technology and some understanding of the softer environmental, legal and behavioural issues (see Choo and Dehghantanha [46]) though little if any accounting awareness.

So, it would seem, that apart from a few exceptional, motivated, highly skilled individuals there is not yet a significant body that balances the three areas in Figure 1 in the Venn diagram below. The diagram is, of course, highly simplistic intending to just give a broad view of the difficulties in bringing the wide range of knowledge and experience required for forensic cloud investigation. One logical conclusion would be the need to build multi-disciplinary teams, though the development of sufficient common understanding, shared technical language, never mind recognition of mutual professional credibility and importance should not be taken as insignificant.

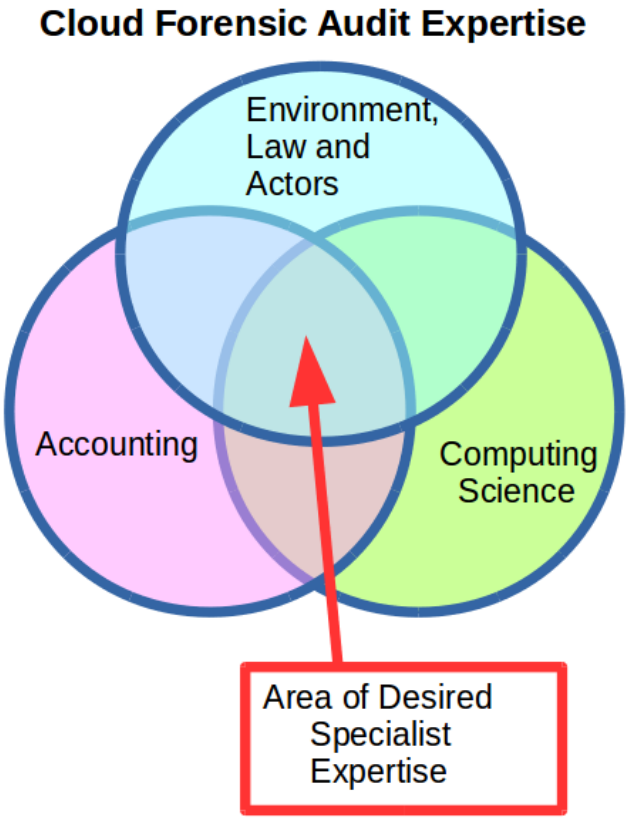

Figure 1: The Area of Desired Expertise

Whilst there are many audit tools used in accounting, the

computing literature already uses the "audit trail" [47] when discussing evidence integrity, however in previous work [31], [54]–[56] we have questioned the level of development of these audit trails and whether all the lessons from the rich accounting history in this area have been understood and then taken on board.

One stark difference between the accounting approach and the computing one is that of redundancy. To the accountant, there is an expectation of keeping more rather than less — indeed the whole concept of double entry is to record every transaction twice. Computer scientists, on the other hand, have a focus on efficiency and minimising costs, using such terms as "redundancy" for seemingly unnecessary or duplicate recording.

An audit trail needs to be developed and be fit for purpose — it may require some thought and planning to decide firstly on the purpose(s) of the trail and then logically what data needs to be safely and securely recorded. For example, Bernstein [57] sees the trail including: events, logs, and the analysis of these, whilst Chaula [58] gives a longer, more detailed list: raw data, analysis notes, preliminary development and analysis information, processes notes, and so on. Pearson et al. [10] as far back as 2010, accept that attaining consistent, meaningful cloud audit trails is a goal rather than reality. More worryingly, Ko et al. [22] point out that it is possible to delete the audit trail along with a cloud instance, meaning there is no record then remaining. In the traditional accounting external audit, the external accountant appears at the end of the year and would need to access all the records they might need in order to satisfy themselves that everything is in order — an ephemeral audit trail would not be fit for purpose. Ko [59] also details the requirements for accountability.

VI. The Special Skills Mix Needed for Cloud Forensic Audit

As we mentioned earlier, with modern cloud systems, we are no longer able to enjoy the benefits of the permanent ink trail. While reasonable alternatives can be available with conventional distributed network systems, this is not the case for cloud systems. We discussed the Cloud Forensic Problem earlier, and it is this security weakness inherent in cloud systems that makes this job significantly harder to accomplish effectively.

When considering cloud forensic issues, it is now clear that we can no longer afford to rely on conventional discipline boundaries when trying to address these issues, as it is now likely that all the disciplines affected are likely to suffer from potentially significant knowledge gaps. Clearly, the cloud environment is considerably different from conventional distributed network models under the sole control of a company. There are now a great many actors involved in such an environment, each potentially with their own agenda. Legal and regulatory issues are also a lot less clearly defined for cloud environments, with the increased likelihood of multiple companies and jurisdictions.

We also have to contend not only with the invited actors but also with the potential of a number of uninvited actors too. The list of the invited players is longer than we might first think. Company employees and managers may not be as competent or trustworthy as we would wish. Outside the company itself, there will be many others, including the software provider, the cloud service provider, the auditor and, in a modern business-to-business environment, the suppliers and the customers. This is a complex mix of actors with disparate agendas and, frustratingly, it cannot even be taken as given that these legitimate actors will be willing to co-operate fully with each other if a problem arises. Of course, there are also the potential uninvited guests — namely attackers and intruders, with the latter presenting the greater challenge.

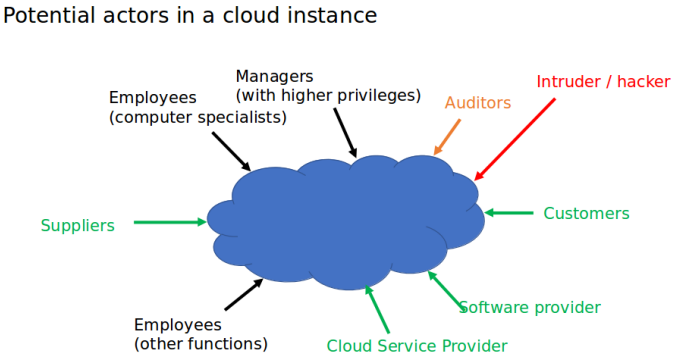

Figure 2: Who is in your cloud?

Figure 2 shows a little of this complexity with the internal actors in black, other companies in green, the auditor in orange and the intruder in red. Whilst one might hope that the authorized participants will play by the rules, any intruder will make up their own, hence gaining access via a customer, supplier or even the auditor is a reasonable option for them.

This level of complexity means we should no longer consider addressing cloud forensic audit from an insular perspective, since accountants, computer scientists and legal, regulatory and other actors within the cloud environment will all suffer from both incomplete knowledge and skill sets. These gaps are not just missing pieces of a jigsaw that someone else from a different discipline can potentially add, but the filling of the gap may also redefine some of the other problems and issues in other disciplinary domains within the complex situation.

Further, in the absence of the solid 'ink trail', this increases the complexity of the task exponentially. In Figure 1 we show the overlapping Area of Desired Expertise that is needed for all three disciplines to fully understand where this knowledge gap needs to be addressed. Assuming individuals are approaching the centre from an initial disciplinary perspective, the need for each to move well beyond their professional comfort zone is obvious and challenging.

Currently, when it comes to Cloud, intruders can potentially have it all their own way. Once they are in the system, it can be merely a matter of time before they have built up sufficient privileges to delete the forensic trail of their activity, thus allowing them to either bed down for the long run, or to withdraw without leaving fingerprints or "steps in the mud". The deleting of all audit and forensic trails as they proceed, means that there is an significant difficulty, verging on the impossible, for data controllers to safely keep the organisation fully compliant with all regulatory and legislative requirements

they must adhere to in order to achieve compliance, security and privacy.

There are, therefore, two major goals that must be dealt with. First, we need to find a way to restore in some way the permanent 'ink trail' so that we have something to fall back on and to enable us to re-trace and re-build if necessary, this is where an immutable audit trail process needs to be conceived, designed and implemented. Second, we need to recognise all the experience and skill sets required and then enable the knowledge gaps to be bridged, ensuring that all the parties who need to be involved in Cloud Forensic Audit are fully up to speed. This will come down to a combination of collaboration and proper training. This latter is outside the scope of this paper, but the first is very much a part of it, and we discuss this further in the following sections.

VII. Some Impediments to Restoring the Permanent 'Ink Trail'

Before discussing how we might resolve the matter of the permanent 'ink trail', we should consider some impediments that are often inadvertently placed by companies on themselves when using cloud systems. Companies should not rush this decision, but should prepare properly ahead of time. They should not assume it will be easy. Instead, they should think it through, understand the costs properly, and ensure they have the right service package rather than continuing to use the first one that came along [60]. Companies frequently wore cost blinkers when choosing cloud provisioning. It is vital to factor in risks and potential exposure too [61] not just looking at the short term, but also taking the long view too.

Many companies have failed to prepare a proper disaster recovery plan [62]. They must always expect the unexpected, and plan for it. It is vital to be aware of precisely which data needs to go to cloud, who should be able to see it, and this data needs to be protected with proper access control. Companies must understand where their data is stored [63] and how they can get their data back, if required. They must understand who all can gain access to their data. As we discussed in Section V cloud systems will not necessarily only be exposed to CSP personnel, but also other sub-contracted organisations [64] whose security and privacy approach might not be as robust as that of the CSP. Companies often fail to account for data privacy risks, which presents a really good incentive for using encryption for their data.

As far as cloud security and privacy is concerned, there is no single solution [54]. It is a mistake to assume the CSP's security is good. CSPs have a huge incentive not to disclose full details of previous security and privacy breaches, as they do not wish to lose future sales. Companies themselves should not use the wrong privacy approach. Wherever possibly, they should try to align security with their own business goals [65]. No matter which approach is used, it should be cloud-friendly. For GDPR compliance, companies should always consider encryption [66] preferably with split encryption keys. Companies often sign up to cloud accepting the standard SLA, which can be a big mistake. Many standard contracts are extremely vague about security and privacy, or do not even mention it. This lack of accountability on the part of the CSP can only help attackers breach company systems more easily.

All companies must understand the true threat against their employees, customers, suppliers and ultimately, their data. Their information security plan must be cutting-edge and comprehensive. They should view security not just as an "IT problem", but rather as a "business problem" that also includes IT. Many who implement security as an IT problem have ended up with a strong IT implementation of data security controls but limited (if any) attention paid to the required security controls such as physical security, security policies and procedures, training, and other administrative and environmental controls. People are frequently the weakest link in the security chain, meaning this is why special attention must be paid to their proper training in all security matters. This is precisely the reason why security mirrors the business architecture of any company, namely it is a combination of people, process and technology [67] not technology alone.

There have been many interesting approaches to trying to resolve some of the obvious issues in cloud security. One such area is how to ensure data integrity in the cloud. There have been a number of interesting proposals, such as [68] [66] [69] [70] [71] which seek to provide data integrity assurance to users through various forms of audit, which as a rule do work quite well. Others, such as [72] [73] [74] [75] have suggested trust computing might be the way to solve these problems. Again, these often work well, but despite establishing trust between providers and users, nevertheless, the fact remains that the work is being performed on someone else's systems, thus risks will always remain. Yet others, such as [76] [77] [78] [79] believe provable data possession could help address this problem, whereas others believe that timeline entanglement, such as [80] [81] [82] is the best solution.

All of these systems, while proving capable of delivering what they promise, share a common flaw. They provide an excellent means of achieving their objectives, but not a means to deal with what happens after a serious security breach. In such cases, the intruders often act in a brutal and indiscriminate way, modifying or deleting multiple records. Any user who does not understand the true purpose of an audit trail may quickly discover that they no longer have access to the necessary data with which to restore the modified or deleted data to its original state.

Thus warned, we will now turn to looking at the immutable database in the next section.

VIII. The Immutable Database

We can see that compliance with the GDPR is not a readily achievable goal that can be easily met by any organisation using Cloud services, due to the difficulties associated with the Cloud Forensic Problem. Thus, we must ensure we create and maintain both a secure forensic and audit trail in order to have any chance of making this happen. Three fundamental weaknesses exist and need to be addressed.

First, failure to use adequate default logging options will result in a reduced level of required audit trail data being collected. Second, there is often a lack of understanding that audit trail data can be accessed by a malicious user gaining root privileges. This can allow the malicious user remove key data which would otherwise have provided evidence of who compromised the system, and what actions they performed once they took control. Third, log data must be properly collected in permanent storage such that there can be no loss of audit trail data, either when an instance is shut down, or when it is compromised. The obvious answer would be to store

this data away from the running cloud instance, in a secure environment using an immutable database which will allow "append-only" transactions to be made.

Starting with the first point, most database software offers a considerable range of audit trail options that can be used to keep proper records of what is happening within the system. However, by default, logging is set to "off"! Since many organisations rely on default installation settings, it is clear that they will be at an immediate disadvantage unless the logging options are fully explored and activated. An obvious, yet simple point missed by many.

Looking at the second point, as Anderson [83] states, the audit trail should only be capable of being read by users. In a cloud setting, this presents a problem, as the software being used is usually running on someone else's hardware, with the output being stored there as well. There is always a risk of compromise from any outside user with malicious intent. There may also be a risk of compromise by some malicious actor working for the CSP. The CSP may very well take vetting of staff seriously, but there may be situations that arise where a temporary contract worker, subject to lesser scrutiny, is engaged at short notice.

Turning to the third point, where database logging is actually switched on, this data is logged to the currently running instance. Thus, this data remains accessible to any intruder who is able to successfully breach the system. This will afford them the opportunity to cover their own tracks by modifying or deleting any entries relating to their intrusion of the system. Equally, they may simply delete the entire audit trail files. Finally, when the instance is shut down, all the data would disappear anyway.

These three points are generally not much thought about, yet they present a serious weakness to the success of maintaining the audit trail. Equally, these are relatively trivial to address. All too often, management and IT staff will take the view "so what?. We don't need to collect redundant data". This entirely misses the point that this data is the only source of proof of what intruders have done whist inside the system. Without these records, it will not be possible to comply with GDPR compliance procedures.

Of course, we need to consider very carefully exactly what data we should log to ensure we can achieve compliance with the GDPR. First, we need to monitor our Cloud instances. We need to understand exactly who is accessing our systems, whether authorized or not and we need to monitor what is happening with our database systems to understand what these users are doing with them.

Looking at our Cloud instances, as Duncan and Whittington have shown in [31] [84] [55] a working solution can be found using an immutable database at its core to record all the relevant information we would require. This means we must first consider carefully exactly what that information should be.

We would want to log all significant events as they transpire during the life cycle of each Cloud instance, with the first significant event being the creation of the Cloud instance, and the last being the shutting down of that instance. Under normal circumstances, these, and all other lifetime events, would be logged on the instance itself. This, as we know from Ko et al. [22] is a dangerous thing to do; thus our first step will be to ensure this data is logged additionally onto an external secure immutable database to ensure it achieves full persistence.

This external database must run on a dedicated secure server, protected by an Intrusion Detection System (IDS), and the database must be immutable, i.e., append only. This secure server will also use dedicated software agents to police the activities being logged, so that the occurrence of any significant event (such as the shutting down of a Cloud instance) will be instantly identified and reported for approval/further investigation.

Turning to the question of who is using our systems, we want to understand who is logging in to our systems, where they come from and what they do once they have successfully logged in. Thus we must capture the relevant detail from the access logs. The detail of how this may be carried out will depend on the systems architecture deployed, the type of access control credentials used and means of controlling access to the various systems available to specific users. A multi-factor authentication approach is always better than access by password. Proper logging of each step in the process is also always preferable.

Once a user gains access to any system, we still want to know where the user came from, and we certainly want to know what the user did with the system after they gained access. Thus we should be logging all the steps that the user takes, regardless of whether access is via physical presence or via remote login. In other words, we need to log every single query made or instruction given to the system. We might wish to consider whether we want to record what the result of that query would be, since this might generate inordinate amounts of data in the case of a database query. Whatever we decide is required, we must ensure a separate copy of the queries recorded are stored into our dedicated secure immutable database. It is clear that redundancy can be a good thing.

IX. Discussion

It is clear that without the assistance of the humble audit trail, compliance with the GDPR while using cloud is likely to prove an unattainable goal. Of course, not being breached would also provide a solution, but based on events to date, there is no guarantee that such a situation would be readily achievable, let alone sustainable in the long run.

Having developed an effective, yet simple and workable solution to this problem, we may well have some further questions, such as:

- How easy is it to implement?
- How quickly and how well will interested parties adhere to such a solution?
- In the event of a breach, who will be responsible and what might the consequences be?

The answer to the first question is that we take the view that this approach needs to be simple to implement and simple to maintain. It is as simple as switching on the necessary forensic and audit trail logging, then writing a chron job to forward the resulting logs to the immutable database. Wherever possible, such programmes should be set to immutable to make it difficult for attackers and intruders to delete them. A regular

check on the configuration files would also be a useful thing to do.

For the second question, it is likely that the easier something is to implement, the more likelihood that it will be implemented. It is not challenging to implement, nor to maintain, and the consequences of failing to do so could have a huge adverse impact, so there is a considerable incentive to both implement and maintain this approach.

As to the third question, it is not a question of 'in the event of a breach', but rather a case of accepting there will be breaches, and these are likely to be a continuous feature. As soon as a breach occurs, a forensic trail will be generated and stored both within the Cloud instance , as well as in the off-site immutable database. Under normal circumstances, the attacker will now attempt to dig deep, escalate privileges and delete the forensic trail. The longer the intruder remains inside the system, the more likelihood that a successful deletion of the audit trail will take place. However, with a covert copy of the forensic and audit trail data available, this will allow some potentially fruitful investigative work to take place.

In the event that an attack against the Cloud instance is successful, where will liability sit? The GDPR regulation is quite clear. In the event of a breach, the Data Controller has a legal obligation to notify the Supervisory Authority within 72 hours of becoming aware of a breach. Individuals must also be notified in the event that encryption is not used. Clearly the use of encryption would be a prudent approach to minimise the impact of the breach, as well as the amount of any possible fine. It is also the case that some practical measure should also be taken, such as ensuring that the encryption and decryption keys are not stored on the cloud instance they are designed to protect.

Clearly, doing nothing is not an option. Without a means of being able to tell which records have been accessed, modified or deleted, compliance with the GDPR will not be possible, and that will potentially carry a very high price tag indeed.

X. Conclusion and Future Work

We have seen that compliance with the EU GDPR for all Cloud users is likely to present a significant challenge. Without special measures being taken, it is likely that compliance will prove impossible to achieve. This is likely to expose such Cloud users to the full force of the penalties of this regulation, which are significant.

It is clear that a minimal requirement will be to generate both a secure forensic trail and an audit trail, in order to have the basic requirements to be able to consider fulfilling the regulatory requirements in the event of a breach. Without this in place, it is likely to be impossible to comply with the legislation, thus rendering the organisation liable to some serious penalties.

In this article, we have identified the particular issues that companies who are Cloud users and are liable to be GDPR compliant must be able to deal with. There is no point in relying on Cloud service providers to take care of this matter. The company data controller is accountable to the regulator for ensuring the company is compliant, and without both a forensic trail and a full audit trail for the PII held on behalf of EU residents, then compliance will not be possible. This will lead to potentially massive fines being applied — a situation that is potentially avoidable.

We have built a miniature real life Cloud system on which to test our ideas. The server runs a full Cloud management system, which will be used to run a number of independent Cloud instances, all of which will run web servers with database back ends to replicate the approach of many Cloud users. This system will run on a closed network where it will be subject to rigorous attack, with the view to discover whether the immutable database approach can succeed in allowing Cloud users to be GDPR compliant.

We have developed a range of scenarios to test, and we seek to find the optimum solution providing the right balance between usability, performance, cost and ease of dealing with breaches. We shall be reporting on our results next year, and we will be working towards delivering a workable solution to keep Cloud users compliant.

References

[1] B. Duncan and M. Whittington, "Can Forensic Audit Help Address the Cloud Forensic Problem in Light of the Requirements of the Forthcoming European Union General Data Protection Regulation?" in Cloud Computing 2018: The Ninth International Conference on Cloud Computing, GRIDs, and Virtualization, no. February. Barcelona, Spain: IARIA, 2018, pp. 84–89.

[2] EU, "EU General Data Protection Regulation (GDPR)," 2017. [Online]. Available: http://www.eugdpr.org/ [Last accessed: August 2018]

[3] T. Takahashi, Y. Kadobayashi, and H. Fujiwara, "Ontological Approach toward Cybersecurity in Cloud Computing Categories and Subject Descriptors," in Science And Technology, 2010, pp. 100–109.

[4] L. M. Vaquero, L. Rodero-Merino, J. Caceres, and M. Lindner, "A Break in the Clouds: Towards a Cloud Definition," ACM SIGCOMM Computer Communication Review, vol. 39, no. 1, 2008, p. 50.

[5] M. Alhamad, T. Dillon, C. Wu, and E. Chang, "Response Time for Cloud Computing Providers," Response, 2010, pp. 8–10.

[6] D. Durkee, "Why Cloud Computing Will Never Be Free," Communications of the ACM, vol. 53, no. 5, may 2010, p. 62.

[7] S. Fraser, R. Biddle, S. Jordan, K. Keahey, B. Marcus, E. M. Maximilien, and D. Thomas, "Cloud Computing Beyond Objects: Seeding the Cloud," Communications, 2009, pp. 847–850.

[8] A. Haeberlen, "A Case for the Accountable Cloud," ACM SIGOPS Operating Systems Review, vol. 44, no. 2, 2010, p. 52.

[9] S. Pearson, "Towards Accountability in the Cloud," IEEE Internet Computing, vol. 15, no. 4, jul 2011, pp. 64–69.

[10] S. Pearson and A. Benameur, "Privacy, Security and Trust Issues Arising from Cloud Computing," in 2010 IEEE Second International Conference on Cloud Computing Technology and Science, no. December. Ieee, nov 2010, pp. 693–702.

[11] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, "Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility," Future Generation Computer Systems, vol. 25, no. 6, 2009, p. 17.

[12] L. M. Kaufman, "Data security in the world of cloud computing," IEEE Security and Privacy, vol. 7, no. 4, jul 2009, pp. 61–64.

[13] P. Mell and T. Grance, "Effectively and Securely Using the Cloud Computing Paradigm," NIST, Information Technology Laboratory, vol. 2, no. 8, 2009, pp. 304–311.

[14] P. Mell, T. Grance, and Others, "The NIST Definition of Cloud Computing," National Institute of Standards and Technology, Tech. Rep., 2011.

[15] S. Bradshaw, C. Millard, and I. Walden, "Contracts for Clouds: Comparison and Analysis of the Terms and Conditions of Cloud Computing Services," International Journal of Law and Information Technology, vol. 19, no. 3, 2011, pp. 187–223.

[16] K. Hon, C. Millard, and I. Walden, "The Problem of 'Personal Data' in Cloud Computing - What Information is Regulated ?" 2011. [Online]. Available: http://ssrn.com/abstract=1562461 [Last accessed: August 2018]

[17] M. Iansiti and G. L. Richards, "Economic Impact of Cloud Computing," Economics of Innovation and New Technology, vol. 7, no. 2000, 2010, pp. 1–42.

[18] N. Papanikolaou, S. Pearson, M. C. Mont, and R. K. L. Ko, "Towards Greater Accountability in Cloud Computing through Natural-Language Analysis and Automated Policy Enforcement," Engineering, 2011, pp. 1–4.

[19] Data Protection Working Party, "Opinion 05/2012 on Cloud Computing," 2012. [Online]. Available: http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2012/wp196_en.pdf [Last accessed: August 2018]

[20] W. Jansen and T. Grance, "Guidelines on Security and Privacy in Public Cloud Computing," NIST, Tech. Rep. 7, 2011. [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-144/SP800-144.pdf [Last accessed: August 2018]

[21] N. Papanikolaou, S. Pearson, and M. C. Mont, "Towards Natural-Language Understanding and Automated Enforcement of Privacy Rules and Regulations in the Cloud: Survey and Bibliography," Analysis, 2011, pp. 1–9.

[22] R. K. L. Ko, P. Jagadpramana, M. Mowbray, S. Pearson, M. Kirchberg, Q. Liang, and B. S. Lee, "TrustCloud: A framework for accountability and trust in cloud computing," Proceedings - 2011 IEEE World Congress on Services, SERVICES 2011, 2011, pp. 584–588.

[23] R. K. L. Ko, B. S. Lee, and S. Pearson, "Towards achieving accountability, auditability and trust in cloud computing," Communications in Computer and Information Science, vol. 193 CCIS, no. Part 4, 2011, pp. 432–444.

[24] S. Pearson and A. Charlesworth, "Accountability as a way forward for privacy protection in the cloud," Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), vol. 5931 LNCS, no. December, 2009, pp. 131–144.

[25] EU, "Accountability for Cloud (A4Cloud)," 2018. [Online]. Available: http://a4cloud.eu/ [Last accessed: August 2018]

[26] M. Hammock, "A Review of the Economics of Information Security Literature," pp. 1–38, 2010. [Online]. Available: http://ssrn.com/abstract=1625853 [Last accessed: August 2018]

[27] A. M. Froomkin, "Government Data Breaches," Berkeley Technology Law Journal, 2009, pp. 1–42.

[28] Verizon, "2012 Data Breach Investigations Report," Verizon, Tech. Rep., 2012.

[29] Verizon, "2016 Verizon Data Breach Report," Tech. Rep., 2016.

[30] D. Kossmann, T. Kraska, and S. Loesing, "An evaluation of alternative architectures for transaction processing in the cloud," in Proceedings of the 2010 International Conference on Management of Data. Indianapolis, Indiana: ACM Press, 2010, pp. 579–590.

[31] B. Duncan and M. Whittington, "Enhancing Cloud Security and Privacy: The Power and the Weakness of the Audit Trail," in Cloud Computing 2016: The Seventh International Conference on Cloud Computing, GRIDs, and Virtualization, no. April. Rome: IEEE, 2016, pp. 125–130.

[32] R. M. Skinner and J. Milburn, Accounting standards in evolution, 2nd ed. Toronto: Prentice Hall, 2001.

[33] S. A. Zeff, "The objectives of financial reporting: a historical survey and analysis," Accounting and Business Research, vol. 43, no. 4, 2013, pp. 262–327.

[34] A. A. A. C. to Prepare a Statement of Basic Accounting Theory, A statement of basic accounting theory. American Accounting Association, 1966.

[35] IASB, "IASB ED/2015/3 - Exposure Draft Conceptual Framework for Financial Reporting Comments," IASB, Tech. Rep., 2015.

[36] OED, "Oxford English Dictionary," 2017. [Online]. Available: http://www.oed.com [Last accessed: August 2018]

[37] R. H. Ashton, "An experimental study of internal control judgements," Journal of Accounting Research, 1974, pp. 143–157.

[38] S. S. G. Gelinas U.J. and A. E. Oram, Accounting Information Systems (4th edition). South-Western College Publishing, Cincinnati, Ohio, US., 1999.

[39] E. Vaassen, R. Meuwissen, and C. Schelleman, Accounting information systems and internal control. Wiley Publishing, 2009.

[40] J. A. Hall, Accounting Information Systems (3rd edition). South-Western College Publishing, Cincinnati, Ohio, US., 2001.

[41] Sox, "Sarbanes-Oxley Act of 2002," p. 66, 2002. [Online]. Available: news.findlaw.com/hdocs/docs/gwbush/sarbanesoxley072302.pdf [Last accessed: August 2018]

[42] W. Ge and S. McVay, "The disclosure of material weaknesses in internal control after the Sarbanes-Oxley Act," Accounting Horizons, vol. 19, no. 3, 2005, pp. 137–158.

[43] W. S. Hopwood, J. J. Leiner, and D. G. R. Young, Forensic accounting and fraud examination. McGraw-Hill, 2012.

[44] NIST, "NIST Cloud Computing Forensic Science Challenges," 2014, p. 51.

[45] J. Taylor, Forensic accounting. Financial Times Prentice Hall, 2011.

[46] K.-K. Choo and A. Dehghantanha, "Contemporary Digital Forensics Investigations of Cloud and Mobile Applications," in Contemporary Digital Forensic Investigations of Cloud and Mobile Applications. Elsevier, 2017, pp. 1–6.

[47] S. A. Almulla, Y. Iraqi, and A. Jones, "A State-of-the-Art Review of Cloud Forensics," Journal of Digital Forensics, Security and Law, vol. 9, no. 4, 2014, pp. 7–28.

[48] B. Duncan and M. Whittington, "Information Security in the Cloud: Should We be Using a Different Approach?" in 2015 IEEE 7th International Conference on Cloud Computing Technology and Science (CloudCom), Vancouver, 2015, pp. 523–528.

[49] B. Duncan and M. Whittington, "Company Management Approaches Stewardship or Agency: Which Promotes Better Security in Cloud Ecosystems?" in Cloud Computing 2015: The Sixth International Conference on Cloud Computing, GRIDs, and Virtualization, IARIA, Ed. Nice, France: IEEE, 2015, pp. 154–159.

[50] B. Duncan and M. Whittington, "Reflecting on whether checklists can tick the box for cloud security," in Proceedings of the International Conference on Cloud Computing Technology and Science, CloudCom, vol. 2015-Febru, no. February. Singapore: IEEE, 2015, pp. 805–810.

[51] ACCA, "ACCA celebrates hitting 200,000 members worldwide with a global tour to honour each and every one," London, 2018.

[52] ACCA, "Advanced Audit and Assurance: Syllabus and Study Guide September 2018 to September 2019," 2017.

[53] KPMG, "Forensics," 2018. [Online]. Available: https://home.kpmg.com/uk/en/home/services/advisory/risk-consulting/forensic-landing.html [Last accessed: August 2018]

[54] B. Duncan and M. Whittington, "Compliance with standards, assurance and audit: does this equal security?" in Proceedings of the 7th International Conference on Security of Information and Networks - SIN '14. Glasgow: ACM, 2014, pp. 77–84.

[55] B. Duncan and M. Whittington, "Creating an Immutable Database for Secure Cloud Audit Trail and System Logging," in Cloud Computing 2017: The Eighth International Conference on Cloud Computing, GRIDs, and Virtualization. Athens, Greece: IARIA, ISBN: 978-1-61208-529-6, 2017, pp. 54–59.

[56] B. Duncan and M. Whittington, "Creating and Configuring an Immutable Database for Secure Cloud Audit Trail and System Logging," International Journal On Advances in Security, vol. 10, no. 3&4, 2017, pp. 155–166.

[57] D. Bernstein, E. Ludvigson, K. Sankar, S. Diamond, and M. Morrow, "Blueprint for the intercloud - Protocols and formats for cloud computing interoperability," in Proceedings of the 2009 4th International Conference on Internet and Web Applications and Services, ICIW 2009, 2009, pp. 328–336.

[58] J. A. Chaula, "A Socio-Technical Analysis of Information Systems Security Assurance: A Case Study for Effective Assurance," Ph.D. dissertation, 2006. [Online]. Available: http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:A+Socio-Technical+Analysis+of+Information+Systems+Security+Assurance+A+ Case+Study+for+Effective+Assurance#1 [Last accessed: August 2018]

[59] R. K. L. Ko, "Data Accountability in Cloud Systems," in Security, Privacy and Trust in Cloud Systems. Springer, 2014, pp. 211–238. [Online]. Available: http://link.springer.com/10.1007/978-3-642-38586-5 [Last accessed: August 2018]

[60] Q. Zhang, L. Cheng, and R. Boutaba, "Cloud computing: State-of-the-art and research challenges," Journal of Internet Services and Applications, vol. 1, no. 1, 2010, pp. 7–18.

[61] K. Dahbur, B. Mohammad, and A. B. Tarakji, "A Survey of Risks, Threats and Vulnerabilities in Cloud Computing," Computing, 2011, pp. 1–6.

[62] Z. Chen and J. Yoon, "IT Auditing to Assure a Secure Cloud Computing," in Proceedings - 2010 6th World Congress on Services, Services-1 2010, 2010, pp. 253–259.

[63] D. J. Abadi, "Data management in the cloud: limitations and opportunities," IEEE Data Eng. Bull., vol. 32, no. 1, 2009, pp. 3–12.

[64] R. Chow, P. Golle, M. Jakobsson, E. Shi, J. Staddon, R. Masuoka, and J. Molina, "Controlling data in the cloud," in Proceedings of the 2009 ACM workshop on Cloud computing security - CCSW '09, 2009, p. 85.

[65] K. Popovic and Z. Hocenski, "Cloud computing security issues and challenges," MIPRO, 2010 Proceedings of the 33rd International Convention, 2010, pp. 344–349.

[66] S. Subashini and V. Kavitha, "A Survey on Security Issues in Service Delivery Models of Cloud Computing," Journal of Network and Computer Applications, vol. 34, no. 1, 2011, pp. 1–11.

[67] PWC, "UK Information Security Breaches Survey - Technical Report 2012," London, Tech. Rep. April, 2012. [Online]. Available: www.pwc.com www.bis.gov.uk [Last accessed: August 2018]

[68] Z. Hao, S. Zhong, and N. Yu, "A Privacy-Preserving Remote Data Integrity Checking Protocol with Data Dynamics and Public Verifiability," IEEE Transactions on Knowledge and Data Engineering, vol. 23, no. 9, 2011, pp. 1432–1437.

[69] D. Chen and H. Zhao, "Data Security and Privacy Protection Issues in Cloud Computing," 2012 International Conference on Computer Science and Electronics Engineering (ICCSEE), vol. 1, no. 973, 2012, pp. 647–651.

[70] K. Yang and X. Jia, "An Efficient and Secure Dynamic Auditing Protocol for Data Storage in Cloud Computing," IEEE Transactions on Parallel and Distributed Systems, vol. I, no. 9, 2014, pp. 2–5.

[71] L. Wei, H. Zhu, Z. Cao, X. Dong, W. Jia, Y. Chen, and A. V. Vasilakos, "Security and privacy for storage and computation in cloud computing," Information Sciences, vol. 258, 2014, pp. 371–386.

[72] N. Santos, K. P. Gummadi, and R. Rodrigues, "Towards Trusted Cloud Computing," Information Security Technical Report, vol. 10, no. 2, 2005, p. 5. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1363412705000221 [Last accessed: August 2018]

[73] Z. Shen, L. Li, F. Yan, and X. Wu, "Cloud Computing System Based on Trusted Computing Platform," 2010 International Conference on Intelligent Computation Technology and Automation, 2010, pp. 942–945.

[74] S. Sengupta, V. Kaulgud, and V. S. Sharma, "Cloud Computing Security - Trends and Research Directions," 2011 IEEE World Congress on Services, no. October, 2011, pp. 524–531.

[75] Z. Xiao and Y. Xiao, "Security and Privacy in Cloud Computing," Communications Surveys Tutorials, IEEE, vol. 15, no. 2, 2013, pp. 843–859.

[76] Y. Zhu, H. Wang, Z. Hu, G.-j. Ahn, H. Hu, S. S. Yau, H. I. Storage, and R. Information, "Efficient Provable Data Possession for Hybrid Clouds," in Proceedings of the 17th ACM Conference on Computer and communications security, 2010, pp. 756–758.

[77] G. Ateniese, R. Burns, R. Curtmola, J. Herring, O. Khan, L. Kissner, Z. Peterson, and D. Song, "Remote data checking using provable data possession," ACM Transactions on Information and System Security, vol. 14, no. 1, 2011, pp. 1–34.

[78] Q. Wang, C. Wang, K. Ren, W. Lou, and J. Li, "Enabling Public Verifiability and Data Dynamics for Storage Security in Cloud Computing," in IEEE Transactions on Parallel and Distributed Systems, vol. 22, no. 5, 2011, pp. 847–859.

[79] Y. Zhu, H. Hu, G. J. Ahn, and S. S. Yau, "Efficient audit service outsourcing for data integrity in clouds," Journal of Systems and Software, vol. 85, no. 5, 2012, pp. 1083–1095.

[80] H. T. T. Truong, C.-L. Ignat, and P. Molli, "Authenticating Operation-based History in Collaborative Systems," Proceedings of the 17Th Acm International Conference on Supporting Group Work, 2012, pp. 131–140.

[81] M. Mizan, M. L. Rahman, R. Khan, M. Haque, and R. Hasan, "Accountable proof of ownership for data using timing element in cloud services," Proceedings of the 2013 International Conference on High Performance Computing and Simulation, HPCS 2013, 2013, pp. 57–64.

[82] S. L. Reed, "Bitcoin Cooperative Proof of Stake," 2014, pp. 1–16.

[83] R. J. Anderson, Security Engineering: A Guide to Building Dependable Distributed Systems, C. A. Long, Ed. Wiley, 2008, vol. 50, no. 5.

[84] B. Duncan and M. Whittington, "Cloud cyber-security: Empowering the audit trail," International Journal on Advances in Security, vol. 9, no. 3 & 4, 2016, pp. 169–183.

Fixing the Cloud Forensic Problem with Blockchain

Yuan Zhao*,Bob Duncan†
Business School
University of Aberdeen, UK
Emails: *y.zhao@abdn.ac.uk,†robert.duncan@abdn.ac.uk

***Abstract*—Many cloud users are heading into a potentially devastating regulatory disaster zone. A major unresolved cloud issue, namely the cloud forensic problem, this is likely to mean many cloud users will be unable to be compliant with the new EU General Data Protection Regulation. We consider the possible use of blockchain, a cryptocurrency based mechanism, to address the as yet unsolved cloud forensic problem. We believe that the underlying blockchain could be adopted to provide a robust mechanism for ensuring that cloud forensic and audit trail records can be securely maintained. This would ensure that cloud users would in turn be able to ensure they are compliant with the new EU General Data Protection Regulation, thus minimising their exposure to punitive levels of fines. We analyse the key risks in cryptocurrencies, namely the operational risk, market risk and cross contamination risk associated with co-movement of cryptocurrencies with other asset forms, using the most predominant and oldest of those, Bitcoin, to provide an example of how removing these risks might provide a far more effective solution to the cloud forensic problem. Our contribution is to demonstrate how this might be done, and by removing the incentive for attackers, to provide a much higher level of compliance with the EU General Data Protection Regulation for cloud users.**

Keywords–Cloud forensic problem; GDPR; blockchain/bitcoin technology.

I. INTRODUCTION

All computing systems connected to the internet are constantly under attack, and for traditional networked computer systems, this presents a serious challenge to ensure a high level of security and privacy can be maintained. For cloud systems, these challenges increase exponentially, due to the increase in complexity in software, and the multiplicity of layers and actors involved in modern cloud ecosystems, especially in light of their disparate agendas.

There remains one serious, but as yet, unresolved challenge, namely the cloud forensic problem. This problem arises where an attacker breaches a cloud system and becomes an intruder, whereby there is nothing then to prevent that intruder from escalating privileges and removing all trace of their incursion by deleting or modifying the forensic trail identifying all their actions and routes into the system. The intruder seeks to remain hidden in the system, where they can continue absorbing information. In [1], we considered whether it might be possible to utilise blockchain technology to help deal with this problem. This article extends that earlier work.

The cloud forensic problem is particularly problematic for companies who both use cloud, and are liable to fall under the jurisdiction of, and therefore require to be compliant with, the new EU General Data Protection Regulation (GDPR) [2]. Without ensuring their cloud provision can properly retain full audit and forensic records, those who use cloud will struggle to meet compliance requirements. Given the punitive level of possible fines for non-compliance (up to the greater of €20million or 4% of last year's global turnover), this is likely to have a considerable impact on companies who are unable to be compliant.

The very convenience of cloud use for a great many companies is likely to place them at a competitive disadvantage now that the GDPR is live. Due to the long lead time required, the enormous costs involved, and the level of expertise needed to securely set up such systems, moving back to conventional distributed network systems is currently unlikely to present a feasible option for many companies, who will effectively be "waiting for the sword of Damocles to fall".

It is imperative that a viable solution be found as quickly as possible. We take a look at the latest global phenomenon of cryptocurrencies, and the technologies they use to ensure security. Security for all financial systems is a necessary priority in all financial companies. They are subject to an incredible range of risks, and we believe it may be worthwhile looking at the operational risk which encompasses the actions that undermine the technological infrastructure and security assumptions of cryptocurrencies, as well as the market risk related to cryptocurrencies.

We start by examining the cloud forensic problem to understand why it is such a challenge for cloud users to become compliant with the GDPR in Section II. Next, we turn to cryptocurrencies and consider operational risk in such systems in Section III. In Section IV, we conside the implications of market risk, while in Section V, we look at the co-movement of cryptocurrencies with different currencies, indices, and commodities, to show the role of cryptocurrency as a commodity, currency, or a speculative investment under portfolio diversification theory. In Section IX, we consider the robustness of this approach for dealing with security issues. In Section X, we discuss our findings and consider future work, and in Section XI, presents our conclusions.

II. THE CLOUD FORENSIC PROBLEM AND GDPR COMPLIANCE

All computer systems connected to the internet are continuously subject to attack, and cloud systems are no exception. It is certainly the case that no system is immune to attack, and that is particularly true for cloud systems. During the past decade, a great many research papers have allowed a far greater level of security and privacy to be achieved in cloud systems. There have been many good papers produced on both security [3]–[14] and privacy [11], [15]–[30], and a number of others have looked at better accountability as a means to meeting these ends [7], [8], [12], [17], [24], [27], [31]–[50] However,

despite all those efforts, no solutions have yet been found to address the cloud forensic problem.

This problem arises once an attacker compromises a cloud system, thus gaining even a small foothold. Once embedded in a system, the attacker becomes an intruder and seeks to escalate privileges until they can access and delete, or modify, the forensic logs in order to hide all trace of their incursion into the system. This allows them to retain a long term foothold within the system, thus allowing them to help themselves to whatever data they wish.

Many companies do not retain records of which database records have been accessed, and by whom, meaning that once a breach occurs, the ability of the company to be able to report which records have been accessed, copied, modified, deleted or ex-filtrated from the system becomes an impossible task. This results in non-compliance with the GDPR, meaning exposure to potentially punitive levels of fines.

To achieve compliance with the GDPR, all companies must first be able to report a breach within 72 hours of discovery. The global average time for all companies between breach and discovery in 2012 was an average of 6 months [51] [52]. This had improved to some 4 weeks by 2016 [53] — still far short of what is needed to understand what has been going on with the intruders while they were undiscovered.

In the light of cloud use, and in particular the Internet of Things (IoT), this raises the question of just how feasible complying with such a time threshold might be. Where a company uses cloud, the company is breached and it has made no special arrangements to ensure the safety of forensic and audit trail data, the 72 hour deadline is moot, since in the first place, it will have no means of knowing that it has been breached, so will have nothing to report, since the requirement is to report within 72 hours of discovery. However, once discovery does occur, there will be no realistic prospect of that company ever finding out just which records have been compromised. When the forensic and audit trail is gone — it is gone!

The IoT, of course, brings a whole new suite of problems to bear, not least of which is the general insecure level of devices, their small resource level, yet high throughput level of data. some of which may be lost in transit. The issue might not be so much with the data lost from IoT devices, rather than with the ability of attackers to easily compromise the devices, thus allowing them access via corporate networks to other more valuable devices in the system. We do not address the IoT within the scope of this paper, but do recognise that any company using IoT devices will require to take special measures to ensure GDPR compliance can be achieved.

Where a company does not take these special measures to safeguard their forensic and audit trail data, they will be less likely to be able to discover the occurance of the breach. Shoud they by chance manage to discover the breach, they would certainly be in a position to report it with 72 hours of discovery, they will simply struggle to be able to report what has been compromised, meaning they will be liable for some level of fine.

Obviously, the longer an intruder has available to spend inside a company system, the more information they will be able to acquire, and the more potential damage they can cause. While the GDPR was changed from "... within 72 hours of a breach occurring..." to a much less stringent "... within 72 hours of discovery ...", this rather misses the point that if a company cannot discover a breach within 72 hours of the breach occurring, how will they possibly be able to discover that is has arisen at all, let alone what data has been compromised once the intruder has deleted all forensic and audit trails?

So, not being able to discover that a breach has arisen, while not putting the company technically in breach of the GDPR, it will certainly make it extremely difficult to enable them to report which records have been compromised once discovery actually occurs. This means the non-compliance will necessarily become far more serious, thus enlarging the exposure to risk of steeper fines.

While there is no specific requirement to encrypt data, there is certainly a strong recommendation that this should take place, and should do so within a reasonable time. Encryption and decryption keys should not be stored on the cloud instance. Failure to address these issues will certainly lead to steeper fines in the event of a breach. An obvious point is that if encryption is not used, then the regulator will require the company to report the breach to every compromised user, which will prove an impossible task where the forensic and audit trails have been lost, again leading to yet steeper fines.

Due to the large number of high value clients, firms involved in financial services are generally subject to greater attack than many other market sectors [54]. It is worth taking a look at how they address security requirements. We believe there may be some merit in considering cryptocurrencies, since as a new entrant to the market, there is more likelihood that their security approach, being designed from the beginning, might offer better prospects rather than relying on existing methods.

III. Operational Risk of cryptocurrencies

Operational risk refer to any action that undermines the technical infrastructure and security assumptions relating to cryptocurrencies. Considering operational risk will provide us with an understanding of how well these risks are dealt with in cryptocurrencies. In looking at high value successful breaches of cryptocurrencies, we can see that these vulnerabilities relate mainly to operator errors and security flaws, which we discuss later. And most importantly, the Bitcoin platform also faces potential vulnerabilities from protocol designs. Operational insecurity has been addressed by Moore and Christin [55], who suggests that fraudulence is an issue among cryptocurencies. Exchanges act as de facto banks, but almost half of them ceased operation due to the resultant impact of security breaches. However, these exchanges failed to reimburse their customers after shutting down. As an alternative approach, other users have instead deposited their Bitcoins in a digital wallet which has also become a target for cyber-criminals.

A small number of theoretical papers written by computer scientists address the mining pool protocols and anonymity. Miners opted out for the pool in long rounds, in which a potential block will be shared with large groups. Based on a peer-to-peer network layer, Babaioff et al. [56] argue that the current Bitcoin protocols do not provide an incentive for nodes to broadcast transactions. This is problematic, since the system is based on the assumption that there is such an incentive. Instead, by focusing on block mining protocol, Eyal

and Sirer [57] show that mining is not incentive-compatible and that so-called "selfish mining" can lead to higher revenue for miners who collude against others. Houey [58] observed that larger blocks are less likely to win a block race when including new transactions into blocks. Karame, Androulaki and Capkun [59] analysed the security of using Bitcoin for fast payments, and found that double-spending attacks on fast payments succeed with overwhelming probability and can be mounted at lower cost unless appropriate detection techniques are integrated in the current Bitcoin implementation. Regarding the double-spending and selfish mining attacks, Kogias et al. [60] proposed the usage of ByzCoin as a novel protocol to optimise transaction commitment and verification under normal operation while guaranteeing safety and liveness under Byzantine (It leveraged scalable collective signing to commit Bitcoin transactions irreversibly within seconds) faults.

The protection of online privacy and anonymity arises and are both addressed in the literature. Christin [61] examined the anonymous online marketplace in cryptocurrencies. Böhme et al. [62] examined what can be learned from Bitcoin regarding Internet protocol adoption. Many studies analysed the public bitcoin transaction history and found a set of heuristics that help to link a Bitcoin account with real word identities. Androulaki et al. [63] quantified the anonymity in a simulated environment and found that almost half of the users can be identified by their transaction patterns. Using two examples, Bitcoin and Linden Dollars, the report focuses on the impact of digital currencies on the use of fiat money. Gans and Halaburda [64] analysed the economics of private digital currencies, but they explicitly focus on currencies issued by platforms like Facebook or Amazon (that retain full control), and not decentralized currencies like Bitcoin. Dwyer [65] provided institutional details about digital currency developments. The security, privacy and anonymity issue related to Bitcoin has been addressed by Krombholz et al. [66], in which they surveyed 990 Bitcoin users to determine Bitcoin management strategies and identifies how users deploy security measures to protect their keys and Bitcoins. They found that about 46% of participants use web-hosted solutions to manage Bitcoins, and over 50% use such solutions exclusively.

Among all the potential causes for operational risk, the denial-of-service (DoS), or distributed-denial-of-service (DDoS) attack is the prominent form suggested by Böhme et al. [62], which entails swamping a target firm with messages and requests in such volume that either mining pools or exchanges become very slow and unusable. This type of attack is especially effective on the Bitcoin ecosystem because of its relative simplicity of monetising the attacks.

We need to consider another major risk of cryptocurrencies, market risk, and how this affects the volatility of the currency element.

IV. Market Risk of Cryptocurrencies

Market risk via price fluctuation in the exchange rate is inevitable for users holding Bitcoin and other cryptocurrencies. Figure 1 shows the average US dollar-Bitcoin exchange rate, along with its trading volume. It is clear that the market volatility is tremendous for Bitcoin, leading to a high potential market risk.

There is also some attention from the literature focusing on the price dynamics and speculative bubbles in cryptocurrency

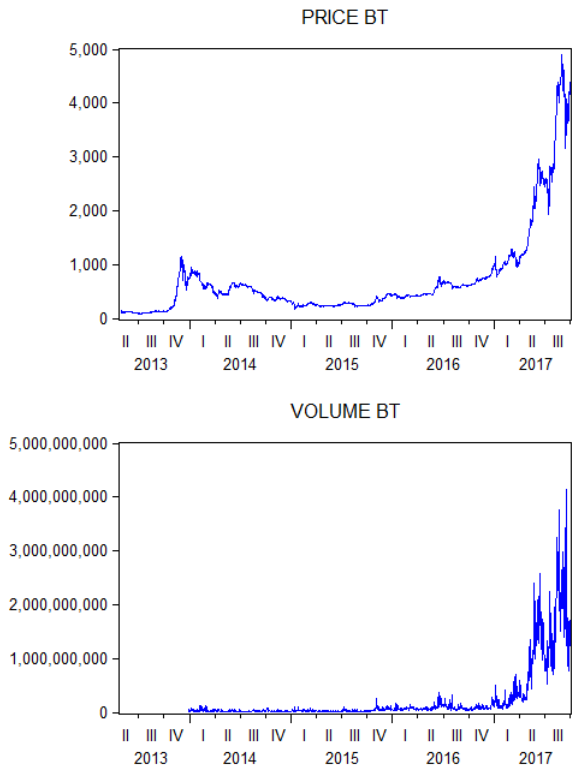

Figure 1: A Comparison Between Price and Volume [67].

markets. Cheah and Fry [68] claimed that cryptocurrencies are prone to substantial speculative bubbles, and they found that the fundamental value of Bitcoin is zero, by examining the daily closing prices of Bitcoin from 2010 to 2014. A more recent study is conducted by Blau [69], which emphasised that the high volatility of Bitcoin is not related to the speculative activities in this period. The volatility of Bitcoin has been analysed by Katsiampa [70]), Cheah and Fry [68], and many others.

Glaser et al. [71] suggest users treat Bitcoin as speculative assets rather than as a type of currency. The diversification benefits offered by Bitcoin is also studied by Briére, Oosterlinck and Szafarz [72]. They found Bitcoin can offer diversification benefits after looking into the correlation between Bitcoin and other asset classes. Gandal and Halaburda [73] examined the exchange rates of different virtual currencies to observe the co-movement and identify the opportunities or triangular arbitrage. But they found little opportunity based on daily closing prices. Yermack [74] analysed changes in Bitcoin price against fiat currencies and concludes that its volatility undermines its usefulness as currency. To be qualified as a currency, Bitcoin needs to serve as an intermediary of exchange, as a unit of account and store value. Also, they have been proved not to be able to function as those by Bariviera et al. [75].

The market risk of cryptocurrencies is also reflected in behavioural factors, such as trading volume and other exogenous factors. Corbet et al. (2017) investigated the fundamental drivers for cryptocurrency price behaviour, and found that there is the existence of bubbles. Jiang (2017) reported the

existence of long-term memory and the inefficiency of the cryptocurrency market, using a similar approach. Alvarez-Ramirez et al (2018) analysed long-term correlation and information efficiency, and reported that the Bitcoin market exhibits time-varying efficiency and price dynamics, which are driven by anti-persistence. Bariviera et al. (2017) compared crytocurrencies with other standard currencies and found that the hurst exponents changed significantly in the initial stage and stabilised thereafter. Bouri et al. (2018) found that the financial stress index could be used to forecast the price movement of cryptocurrencies. Other behavioural factors were found by later researchers. Feng et al. (2017) found evidence on informed insider trading of Bitcoins prior to big events, implying that the informed trading may contribute to explaining the dynamics of the Bitcoin price. Dotsika and Watkins (2017) employed keyword network analysis and identified potential disruptive trends in block-chain technologies.

Next we turn to how cryptocurrencies relate to conventional assets in the context of portfolio theory in order to understand-where the weaknesses arise.

V. Co-movement of cryptocurrencies and Portfolio Theory

Despite extensive studies on the economic aspects of cryptocurrencies, there are relatively fewer studies conducted on analysing the inter-linkage of cryptocurrencies with other financial assets. A number of papers have analysed the ability of cryptocurrencies, usually Bitcoin, to act as safe havens or hedges mentioned by a series of papers such as [76]–[78]. Dyhrberg [76] analysed the hedge properties of Bitcoin using a selection of explanatory variables such as gold (cash and future), the dollar-euro and dollar-pound exchange rates and the the Financial Times Stock Exchange 100 (FTSE 100) Index. The results of the GARCH model [79] showed that Bitcoin can be used in hedging against the dollar and the UK stock market, showing similar hedging capabilities to gold. In Figure 2, we see how a basket of crypro-currencies compare with each other based on price.

Bouri, Azzi and Dyhrberg [78] used a quantile regression approach to analyse the relationships between Bitcoin and global uncertainty. The findings demonstrate that at the longer frequencies VIX have strong negative impact on Bitcoin returns, while at the shorter frequencies uncertainty does have positive and significant impacts only on high quantiles. This implies that Bitcoin can hedge against global uncertainty at short investment horizons and in a bull regime only. Another study by them in 2017 investigated interrelationships between Bitcoin and the world equity indices, bonds, oil, gold, the general commodity index and the US dollar index using the bivariate DCC model by Engle [80]. The results show limited evidence of hedging and safe haven properties of Bitcoin; however, Bitcoin still can be an effective diversifier.

Next, we carried out some empirical research using the three largest cryptocurrencies, Bitcoin, Ethereum and Ripple, by addressing the impact of volatility, which we cover in the next section.

VI. The Empirical Tests, Results and Analysis

In this section, we carry out some empirical tests on the volatility of the three largest cryptocurrencies, Bitcoin,

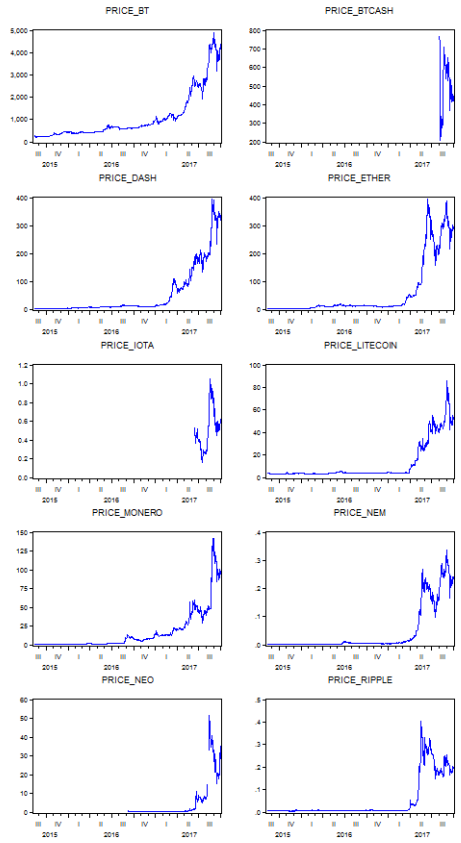

Figure 2: A Co-Movement View of cryptocurrencies Based on Price [67].

Ethereum and Ripple. Figure 3 shows the market capitalisation of the largest three cryptocurrencies, including Bitcoin, Etherum, and Ripple. In this section, we will look into the conditional volatility, correlations, causal relationships, time variation on such relationships, and external factors that may affect the relationships.

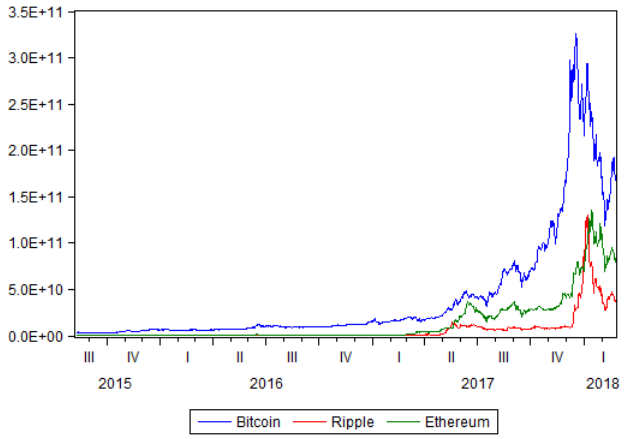

Figure 3: A Comparison of largest three cryptocurrencies [67].

TABLE I: Descriptive statistics and unit root test of Bitcoin returns

Descriptive stats	
Mean	0.002435
Median	0.002045
Maximum	0.3575
Minimum	-0.2662
Std. Dev.	0.04503
Skewness	-0.1917
Kurtosis	11.0549
Jarque-Bera	4776.9130
Observations	1763
Unit root test	
ADF test	-41.6905
PP test	-41.8247
KPSS test	0.2537

- We model the conditional volatility for cryptocurrencies, by comparing different volatility models. We present the findings on Bitcoin as the baseline cryptocurrency. We examine the natural logarithm of the closing price ratio of consecutive days from 28 April 2013 to 24 Feb 2018. The daily return of Bitcoin index is 0.2435% with standard deviation of 0.04503. The returns are negative skewed and leptokurtosis. The p-value of the Jarqu-Bera test indicates that the returns deviate from a normal distribution. We also test there is significant ARCH effect in the returns of Bitcoin returns, suggesting the ARCH family models as the more appropriate specification to model. The unit root test from ADF, PP and KPSS test shows the return series from Bitcoin is stationary. The descriptive statistics and unit root tests are presented as follows in Table I.
  We follow a similar approach to [70], and conduct the likelihood ratio test on the GARCH model specifications, including AR(1)-GARCH(1,1), AR(1)-EGARCH(1,1), AR(1)-TGARCH(1,1), AR(1)-APARCH, AR(1)-CGARCH(1,1). And we find that the AR(1)-EGARCH(1,1) is the best specification based on the results of likelihood ratio test. We forecast the conditional volatility from this specification. Figure 4 shows the persistence and asymmetry in Bitcoin return volatility, especially around late 2013, the beginning of 2015, and the end of 2017.
- The contagion of spillover effects of multiple cryptocurrencies can be investigated using trivariant-GARCH models. The following Figure 5 exhibits the covariance of each pair of cryptocurrencies. It is evident that the covariance between these three cryptocurrencies increases significantly around the recent one year compared to the initial one year. The covariance between Ripple and Ethereum is more sensitive to external economic conditions, implied by the more volatile fluctuations.
- According to Markowitz portfolio theory, an asset that is unrelated or even negatively correlated with another asset in the portfolio is characterised as hedging effective. Thus, it is worth looking into the correlation among the major cryptocurrencies in terms of their roles on portfolio diversification. In this study, we utilise the Granger causality test and vector autoregressive (VAR) model, in order to investigate the short-term dynamic causal relationship between different pairwise cryptocurrencies. In Table II, we present the findings for the short-run causality from different directions, on the null hypothesis of no short-term causal relationships. A p-value (Prob.) less than a predefined significance level (5%) indicates a rejection of the existence of a causal relationship. We find that under the condition of short-run exogenous economic shock, Ripple has a significant causal impact on the returns of Bitcoin. And Etherum has a causal relationship with Ripple. The direction of such causal relationship can be seen in Figure 6, by impulse response function. We find positive causal relationships from all directions.
- As indicated in the previous findings, cryptocurrencies have entered into a more dynamic market with more potential risks. Hence, we especially focus on the recent full year from 2016 to 2017, to examine the time variation of the causality. The following Figure

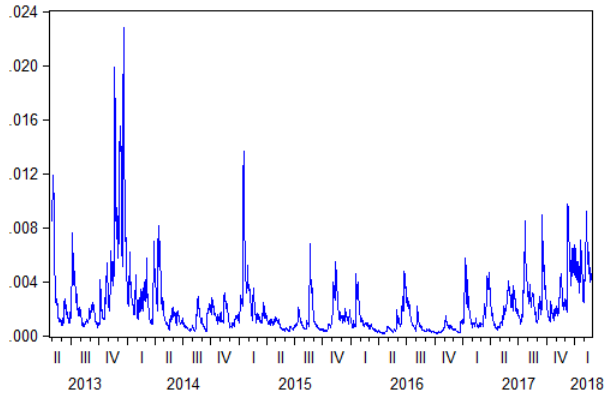

Figure 4: Conditional volatility of Bitcoin returns, from [67].

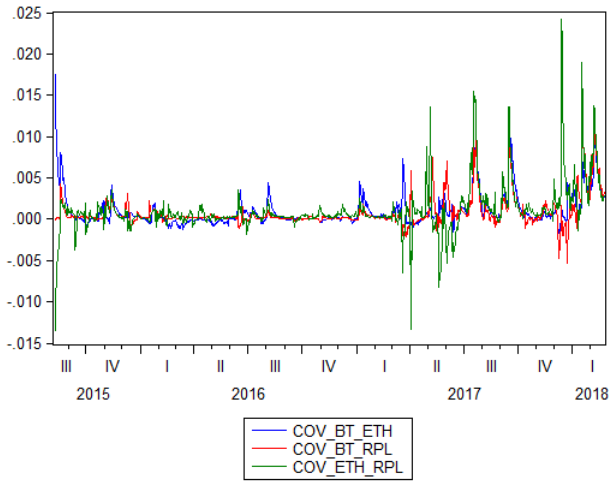

Figure 5: The covariance of largest three cryptocurrencies [67].

TABLE II: Granger causality test of the largest three cryptocurrencies

Granger block exogeneity Wald test		
Dependent variable: Bitcoin		
Excluded	Chi-sq	Prob.
Ethereum	1.119537	0.5713
Ripple	10.46673	0.0053
All	12.08829	0.0167
Dependent variable: Ethereum		
Excluded	Chi-sq	Prob.
Bitcoin	0.188579	0.91
Ripple	2.356285	0.3079
All	2.653052	0.6175
Dependent variable: Ripple		
Excluded	Chi-sq	Prob.
Bitcoin	1.130565	0.5682
Ethereum	5.116094	0.0775
All	5.351787	0.2531

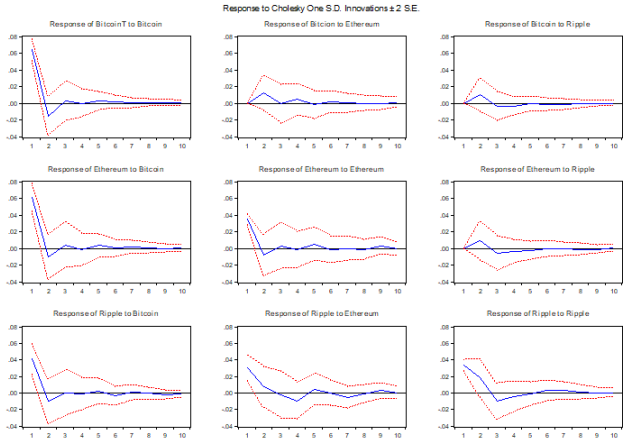

Figure 6: The Impulse Response Function of largest three cryptocurrencies [67].

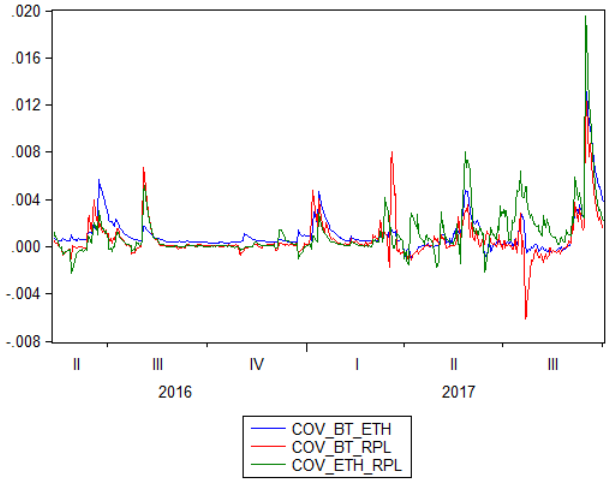

Figure 7: The covariance of largest three cryptocurrencies [67].

TABLE III: Granger causality test of the largest three cryptocurrencies

Granger block exogeneity Wald test 2016-2017		
Dependent variable: Bitcoin		
Excluded	Chi-sq	Prob.
Ripple	3.1278	0.2093
Ethereum	0.8272	0.6613
All	3.6444	0.4563
Dependent variable: Ethereum		
Bitcoin	6.5079	0.0386
Ripple	1.3257	0.5154
All	7.4076	0.1159
Dependent variable: Ripple		
Bitcoin	1.5218	0.4672
Ethereum	0.7558	0.6853
All	3.0384	0.5514

7 exhibits the covariance of each pair of cryptocurrencies, Table III shows the Granger causality of pairwise cryptocurrencies, and Figure 8 illustrates the directions of such causality, in the recent one year. We find that in the recent one year, Bitcoin dominates others by having an increasing covariance with the other two. There is a significantly positive causal relationship from Bitcoin to other currencies, which can be concluded according to the Granger block exogeneity Wald test p-value as 0.0386 and positive responses from Ethereum and Ripple.

- Other external factors may also become sources affecting the market risk of cryptocurrencies. According to the review of financial literature, trading volume is a main factor affecting the risks and returns of financial assets. Therefore, we examine the causality of behavioural factors like trading volume on cryptocurrencies by implementing a VAR model and Granger causality test. Table IV shows the causality of volume from these three currencies to their returns. We find that the trading volume of Ripple has a significant causal relationship over Bitcoin and Bitcoin volume. And the Bitcoin trading volume has the reverse causality over Ripple volume and Ethereum volume, which further confirms our inferences on the increasing impact of Bitcoin in the recent full year over others.

VII. A Summary of the Empirical Results

The design of Bitcoin presents distinctive risks that differ from other payment methods and thus pose security issues related to operational risk, market risk, and contagion risks with other cryptocurrencies.

Operational risk occurs when certain actions undermines the technical infrastructure and security assumption of cryptocurrencies, such as fraudulence of exchanges, mining pool inefficiency, double spending attacks, and online anonymity. However, we know that a DoS or DDoS attack can be very debilitating for blockchain systems.

Market risk lies in the unpredictable fluctuations in the price of Bitcoin and other cryptocurrencies. As an agent for the storage of value and price goods, the sharp movement of exchange rate of Bitcoin will also cause liquidity issues.

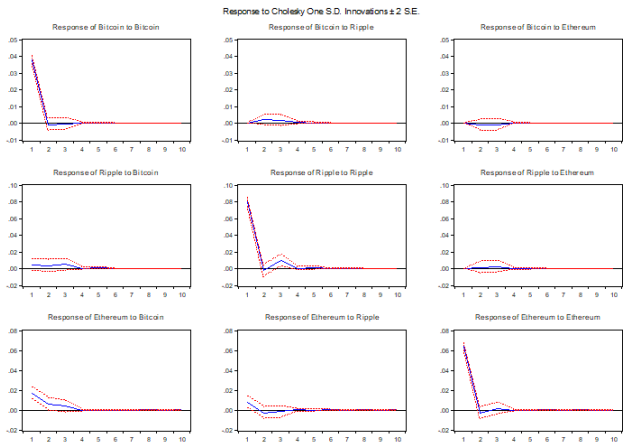

Figure 8: The Impulse Response Function of the largest three cryptocurrencies during 2016-2017 [67].

Contagion risk arises when the co-movement of price of a bundle of cryptocurrencies becomes inevitable. This will cause potential issues for portfolio diversification, despite their innovations and efficiencies. For instance, the Litecoin confirms transactions four time faster than Bitcoin, which is more useful for the retail use and other time-sensitive transactions. NXT [81] reduces the electronic and computational burden of Bitcoin mining by replacing the proof-of-work mining with proof-of-stake, assigning blockchain duties in proportion to coin holdings. Zerocash [82], which is not yet operational, will seek to improve privacy protections by concealing identifiers in the public transaction history. Peercoin [83] allows a perpetual 1% annual increase in the money supply.

In looking at the empirical results, we can see that there is a bi-directional potential contagion effect between each of the cryptocurrencies, which will vary depending on economic conditions. This demonstrates an increased risk of cross contagion between different cryptocurrencies. This contagion appears to be increasing over recent years, which would suggest the contagion risks are increasing. These calculations will help any potential user to consider the impact of these risks in the light of their own risk appetite.

In the next section, we analyse some of the largest successful cyber breaches of cryptocurrencies in order to determine whether there might be any weakness in the fundamental blockchain component.

VIII. An Analysis of Some of the Largest Successful Cryptocurrency Attacks

In this section, we take a look at some of the largest cryptocurrency breaches in recent years, in order to understand how the breaches arose.

The earliest large scale breach to a cryptocurrency exchange was in 2010 due to the value overflow incident — where an early flaw in the bitcoin system allowed the intruder to create 184 billion units of bitcoin. The value then was $21.2bn, although at recent prices the value would have been $1.8 quadrillion. It was notable for the speed at which it was discovered and dealt with, resulting in no actual loss of value. The perpetrator has never revealed themselves and their original 0.5 BTC used in the exploit remains unspent to this day, despite being valued at more than $3,000.

TABLE IV: Granger causality test of the largest three cryptocurrencies return versus trading volume

Granger block exogeneity Wald test		
Dependent variable: Bitcoin		
Excluded	Chi-sq	Prob.
Ethereum	0.0787	0.9614
Ripple	4.6776	0.0964
Bitcoin volume	2.2668	0.3219
Ethereum volume	2.5613	0.2779
Ripple volume	6.5272	0.0383
All	17.6204	0.0617
Dependent variable: Ethereum		
Bitcoin	4.8802	0.0872
Ripple	0.5197	0.7712
Bitcoin volume	3.4664	0.1767
Ethereum volume	1.1715	0.5567
Ripple volume	3.0683	0.2156
All	11.7578	0.3016
Dependent variable: Ripple		
Bitcoin	2.0651	0.3561
Ethereum	1.0425	0.5938
Bitcoin volume	2.4065	0.3002
Ethereum volume	0.3773	0.8281
Ripple volume	2.2058	0.3319
All	10.4823	0.3992
Dependent variable: Bitcoin volume		
Bitcoin	0.7594	0.6841
Ethereum	4.3616	0.1129
Ripple	0.2130	0.8990
Ethereum volume	4.4428	0.1085
Ripple volume	10.7419	0.0046
All	23.4696	0.0091
Dependent variable: Ethereum volume		
Bitcoin	0.3634	0.8338
Ethereum	7.2534	0.0266
Ripple	0.4723	0.7897
Bitcoin volume	6.1108	0.0471
Ripple volume	2.6953	0.2598
All	21.2929	0.0191
Dependent variable: Ripple volume		
Bitcoin	4.6771	0.0965
Ethereum	1.2313	0.5403
Ripple	5.8466	0.0538
Bitcoin volume	17.1896	0.0002
Ethereum volume	2.1749	0.3371
All	40.2409	0.0000

Jan 2018 - Tokyo based Coincheck suffered a $530 million loss of crypto currency due to being hacked. Investigations showed that this breach arose due to the Coincheck exchange not using secure networks. Customer funds were stored in "hot" wallets which were live to the internet, instead of using "cold" wallets should have been offline and not visible to the internet.

The 2014 Tokyo based Mt Gox lost $460 million following a hack which was successful due to a combination of poor management, neglect and sheer inexperience. This was the second, and fatal, hack for the business, having already lost $8.75 in June of 2011. This second hack resulted in bankruptcy for the company and arrest for the CEO of the company.

The February 2018 hack on BitGrail was worth $195 million. While there was speculation that the BitGrail founder Francesco Firano siphoned off the funds, he in turn insists it was a hack.

In 2016, Bitfinex, another of the world's largest bitcoin exchanges was hacked and lost $72 million. The company had used a different authorisation mechanism in an attempt to make the system more robust, but did not realise their approach had an exploitable weakness, which hackers duly discovered and exploited. Rather than ceasing operations, Bitfinex reduced the balance on all accounts by 36%, ragardless of whether their account had been compromised to cover all the losses, and were given an alternative cryptocurrency, BFX tokens, in exchange which Bitfinex promised to buy back over time. As of April 2017, Bitfinex had fully reimbursed all of its customers.

Also in 2016, the Decentralized Autonomous Organization (DAO) which was created to operate like a venture capital fund for decentralized cryptocurrency projects, built on a smart contract on the Ethereum blockchain, were hacked. A hacker drained $70 million within a few hours by exploiting a flaw that allowed the DAO smart contract to return Ether multiple times before it updated its internal balance. The company coders failed to realise the possibility that anyone would use a recursive function to take advantage of this weakness. The hack resulted in the hard fork of the Ethereum protocol that resulted the creation of Ethereum Classic (ETC).

In December 2017, hackers attacked the NiceHash mining service, and with the assistance of a compromised company computer, made off with 4,400 bitcoins from customer accounts worth $64 million. While the funds were not recovered, NiceHash promised to compensate their customers in full. Within a few weeks the lost bitcoins were back in customer accounts.

In June 2018, Coinrail, a South Korean exchange, was the target of an attack, losing around $37 million of cryptocurrencies Pundi X and Aston. Again, they were storing bitcoin online. The remaining 70% of currency was rapidly switched to offline storage. The attack was traced to an Ethereum address, which has subsequently had its assets frozen.

In July 2017, the parity multisig wallet exploit was used against three large Ethereum accounts, netting $32 million. The owners of these accounts were believed to be the Ethereum-powered casino Edgeless, decentralised commerce platform Swarm City amd the smart contracts platform aeternity. All three accounts had recently held initial coin offerings , thus their wallets contained large amounts of money. Swarm City recently confirmed that it was one of the targets.

In June 2018, Bithumb, a South Korean exchange were hit by hackers, reporting $31.5 million stolen.

In 2012, Bitcoinica, another large bitcoin trading platform was hacked, losing 46,703 bitcoins. It subsequently transpired that Bitcoinica stored large amounts of digital currency online, as opposed to offline in secure servers. Just a few months later, a second hack resulted in a further loss of another 18,547 bitcoin.

In every case of the above successful attacks, the inherent strength of the blockchain algorithm behind these companies was never in question. Rather, the success of the attacks came down to successful exploitation of mostly human weaknesses, poor decisions, poor management, neglect and sheer inexperience.

IX. The Robustness of this Approach for Security Issues

In previous sections, we have seen that there are a number of key risks pertaining to cryptocurrencies, namely operational risk, market risk, and contagion risks with other cryptocurrencies. In looking at the largest successful cryptocurrency breaches, we can see that while the breaches were successful, the underlying blockchain was never breached. The original part-bitcoin leveraged to perpetrate the Mt Gox attack in 2014 has never been sold as this would provide proof to the authorities who perpetrated the attack, which is testament to the inherent strength of the blockchain.

In looking at a number of real world instances, we can see that there are potential issues that must be considered. Attacks, such as DoS and DDoS attacks, can prove lethal to both functionality and performance, although Tripathi et al. [84] have suggested a workaround to mitigate this particular issue. One obvious approach is to discuss the matter with the CSP to ensure they have the capacity to be able to handle such an attack should it arise.

The majority of successful attacks are perpetrated against the storage and containment technology in use, often utilising social engineering or in a recent case, holding of BitCoin owners to ransom until their BitCoins are transferred to the criminal perpetrators.

There are clear core strengths contained in blockchain technology, due to the high redundancy provided, but there are practical concerns to be considered. The lack of a clear economic methodology to pay for the use of the technology presents a major concern, as does the volatility of the cryptocurrencies inextricably linked to it. While the high value of the cryptocurrency element provides a strong incentive to attackers, if we remove this element by simply removing the cryptocurrency, we can see that at one fell swoop, we also lose operational risk, market risk, and contagion risks with other cryptocurrencies. We also lose a huge volume of transactional data involved in the trading of cryptocurrencies, meaning we are left with blockchain only, the distributed ledger element. With vastly reduced transactional volumes, latency of operation will be much less of an issue.

There needs to be a sufficient incentive for distributed ledger providers to provide a highly secure, robust and low latency mechanism to deliver the means to record irrefutable transactional data rapidly enough to provide a high performing system. It is certainly the case that the use of some blockchain based mechanism to protect cloud instances could prove a very useful means of doing so. However, it is also obvious that if the blockchain ledgers are run within the same cloud instance as the system they are trying to protect, then we would be asking for trouble.

The obvious solution to this issue would be to truly distribute the blockchain instances to a sufficiently diverse number of locations, such as to make it difficult for an attacker to compromise all, or a sufficiently large number of the ledgers to be able to force a permanent illicit change to their own advantage. On the other hand, while the increased number of distributed ledgers can significantly increase the security, it will also increase the cost and the latency of processing transactions. An economic balance will need to be determined. Carlsten et al. [85] warn of the potential instability

of bitcoin without the block reward, so clearly paying service providers to run the blockchain would be required to provide a sufficiently robust service. We would also have the option of using the same approach with the Ethereum cryptocurrency, which would offer the option of being able to deal with smart contracts. However, for most purposes, the basic blockchain will be more than adequate for our needs.

We have seen how Azaria et al. [86] have suggested a similar approach to improve privacy with medical records. Christidis and Devetsikiotis [87] have suggested the use of both blockchain and smart contracts to improve security and privacy for the Internet of Things. Dinh et al. [88] offer a blockchain benchmarking system to compare the relative efficiencies of differing blockchain and smart contract options. Gaetani et al. [89] have proposed a blockchain based database to ensure data integrity for cloud. Kiayias and Panagiotakos [90] suggest the GHOST protocol at the core of Ethereum could offer significant increases in speed for transactional recording. Yermack [91] suggests that the use of blockchain can improve help to Corporate Governance. There is no doubt that there is a great deal of interest in trying to apply this new approach to make improvements for cloud users.

X. Discussion

Thanks to the major weakness posed by the cloud forensic problem, the potential to lose both the audit trail and the forensic trail means that recording the data we require to remain compliant with the GDPR becomes a vitally important task for us. The use of a distributed ledger holds great promise. The blockchain approach affords us with increased redundancy, meaning that an attacker will have to compromise a great many of the distributed ledgers before they can have any impact on the ledger contents. Some would see this as too much redundancy. We would view this as just enough to provide the required assurance. This can therefore provide us with a very strong assurance that the consensus across the ledgers will deliver a high level of comfort as to the veracity of the contents. So, while this represents a big drawback for some, for us, it represents a major advantage!

Some point to the huge volumes of processing generated by the blockchain process as used in Bitcoin, suggesting that it would be too computationally expensive for our purposes. We take a different view. Because it is a cryptocurrency and highly volatile, Bitcoin is subject to transactional volumes measuring in multi-trillions per year. By stripping out the cryptocurrency aspect from the equation, we also remove the need for such extreme volumes of transactional data, rendering the approach very manageable for any size of company.

Some express concerns at the impact of selfish miners. We take the view that by removing the need for mining from the equation, and instead having the processing carried out by credible parties for economic cost, this will remove any incentive to try to mess with the system in this way. All processors would be paid at the same rate for the job they perform, so there would be no means available to them, nor any incentive, to try to improve on that.

Yet others point to the dangers of DoS and DDoS attacks. Given that there will be no direct financial advantage to be gained by attacking these blockchain ledgers, the volume of attacks will likely reduce to a significantly lower level. For a large attack to be financially viable, there has to be a huge financial incentive before it becomes worthwhile to spend the kind of money it takes to perpetrate such an attack.

XI. Conclusion and Future Work

It is clear that for any company using cloud, it will prove virtually impossible to achieve compliance with the GDPR in the event of a security breach due to the, as yet unresolved, Cloud Forensic Problem. Discovering this fact after a cyber breach will not be grounds for mitigation from the regulator after the fact. It will be far too late by then. Therefore, cloud users who require to be compliant with the GDPR will have to take steps now to be thoroughly prepared ahead of time.

We have looked at the Operational Risk and the Market Risk of cryptocurrencies as well as considering the co-movement of cryptocurrencies in the light of portfolio theory. Many of these risks arise through the perceived mass value attributable to these cryptocurrencies and the mass transactional processing volumes implicit in their operation. Clearly, by removing the currency aspect from the equation, we can eliminate a huge portion of the risk. We accept that all risk will not be removed, but there will be a significant reduction in risk levels involved.

Our proposal will be to use the underlying concept of a distributed ledger to ensure we are in a position to retain some element of both audit trail and forensic trail data to allow us to meet the compliance requirements of the GDPR, which would otherwise be impossible in the event of a breach. There will be a need to carry out some serious testing in order to find a satisfactory equilibrium between security, privacy, performance, reliability, accessibility and the accountability we require for GDPR compliance.

To that end, we plan to conduct a pilot case study on how the technical aspects might be implemented in order to meet all the required goals to ensure compliance can be achieved. This will run around a miniature cloud in a box system, offering both cloud-based and non-cloud based ledgers to assess what the optimum configuration might be.

References

[1] Y. Zhao and B. Duncan, "Could Block Chain Technology Help Resolve the Cloud Forensic Problem?" in Cloud Comput. 2018 Ninth Int. Conf. Cloud Comput. GRIDs, Virtualization, no. February. Barcelona, Spain: IARIA, 2018, pp. 39–44.

[2] EU, "EU General Data Protection Regulation (GDPR)," 2017. [Online]. Available: http://www.eugdpr.org/ [Last accessed: August 2018]

[3] M. Almorsy, J. Grundy, and I. Müller, "An analysis of the cloud computing security problem," 17th Asia-Pacific Softw. Eng. Conf. (APSEC 2010) Cloud Work. Aust., no. December, 2010, p. 7.

[4] M. Almorsy, J. Grundy, and I. Mller, "An analysis of the cloud computing security problem." proc. 2010 Asia Pacific Cloud Work. Coloca. with APSEC2010, Aust., 2010.

[5] V. Chang, Y. H. Kuo, and M. Ramachandran, "Cloud computing adoption framework: A security framework for business clouds," Futur. Gener. Comput. Syst., vol. 57, 2016, pp. 24–41.

[6] F. Doelitzscher, T. Ruebsamen, T. Karbe, M. Knahl, C. Reich, and N. Clarke, "Sun Behind Clouds - On Automatic Cloud Security Audits and a Cloud Audit Policy Language," Int. J. Adv. Networks Serv., vol. 6, no. 1, 2013, pp. 1–16.

[7] R. K. L. Ko, P. Jagadpramana, M. Mowbray, S. Pearson, M. Kirchberg, B. S. Lee, and Q. Liang, "TrustCloud: A Framework for Accountability and Trust in Cloud Computing," Perspective, 2011, pp. 1–9.

[8] R. K. L. Ko, B. S. Lee, and S. Pearson, "Towards achieving accountability, auditability and trust in cloud computing," Commun. Comput. Inf. Sci., vol. 193 CCIS, 2011, pp. 432–444.

[9] K. Lee, "Security Threats in Cloud Computing Environments," Int. J. Secur. its Appl., vol. 6, no. 4, 2012, pp. 25–32.

[10] N. Papanikolaou, S. Pearson, and M. C. Mont, "Towards Natural-Language Understanding and Automated Enforcement of Privacy Rules and Regulations in the Cloud: Survey and Bibliography," Analysis, 2011, pp. 1–9.

[11] S. Pearson, "Taking Account of Privacy when Designing Cloud Computing Services," Current, 2009, pp. 44–52.

[12] S. Pearson, "Toward accountability in the cloud," IEEE Internet Comput., vol. 15, no. 4, jul 2011, pp. 64–69.

[13] S. Ramgovind, M. Eloff, and E. Smith, "The Management of Security in Cloud Computing," in Inf. Secur. South Africa (ISSA), 2010, 2010, pp. 1–7.

[14] H. Takabi, J. B. Joshi, and G.-J. Ahn, "Security and Privacy Challenges in Cloud Computing Environments," IEEE Secur. Priv. Mag., vol. 8, no. 6, nov 2010, pp. 24–31.

[15] S. Creese, P. Hopkins, S. Pearson, and Y. Shen, "Data Protection-Aware Design for Cloud Computing," Work, no. December, 2009, pp. 1–13.

[16] K. Hon, C. Millard, and I. Walden, "The Problem of 'Personal Data' in Cloud Computing - What Information is Regulated ?" 2011.

[17] W. K. Hon, C. Millard, and I. Walden, "Who is Responsible for 'Personal Data' in Cloud Computing? The Cloud of Unknowing, Part 2," Leg. Stud., no. 77, 2011, pp. 1–31.

[18] W. K. Hon, J. Hörnle, and C. Millard, "Data Protection Jurisdiction and Cloud Computing – When are Cloud Users and Providers Subject to EU Data Protection Law?" Leg. Stud., 2011, pp. 1–40.

[19] W. Itani, A. Kayssi, and A. Chehab, "Privacy as a Service: Privacy-Aware Data Storage and Processing in Cloud Computing Architectures," 2009 Eighth IEEE Int. Conf. Dependable, Auton. Secur. Comput., dec 2009, pp. 711–716.

[20] W. Jansen and T. Grance, "Guidelines on Security and Privacy in Public Cloud Computing," NIST, Tech. Rep., 2011. [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-144/SP800-144.pdf [Last accessed: August 2018]

[21] H. Katzan Jr, "On The Privacy Of Cloud Computing," Int. J. Manag. Inf. Syst., vol. 14, no. 2, 2011, pp. 1–12.

[22] W. K. Hon, C. Millard, J. Singh, I. Walden, and J. Crowcroft, "Policy, legal and regulatory implications of a Europe-only cloud," Int. J. Law Inf. Technol., vol. 24, no. 3, 2016, pp. 251–278.

[23] NIST, "Security and Privacy Controls for Federal Information Systems and Organizations Security and Privacy Controls for Federal Information Systems and Organizations," Natioinal Institute of Standards and Technology, Gaithersburg, MD, Tech. Rep. February, 2014. [Online]. Available: http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r4.pdf [Last accessed: August 2018]

[24] S. Pearson and A. Charlesworth, "Accountability as a Way Forward for Privacy Protection in the Cloud," Computing, no. December, 2009, pp. 1–15.

[25] S. Pearson and A. Benameur, "Privacy, Security and Trust Issues Arising from Cloud Computing," 2010 IEEE Second Int. Conf. Cloud Comput. Technol. Sci., nov 2010, pp. 693–702.

[26] S. Pearson, "Privacy, Security and Trust in Cloud Computing," in Priv. Secur. Cloud Comput. e: Springer, 2013, pp. 3–42.

[27] J. Prüfer, "How to govern the cloud? Characterizing the optimal enforcement institution that supports accountability in cloud computing," Proc. Int. Conf. Cloud Comput. Technol. Sci. CloudCom, vol. 2, 2013, pp. 33–38.

[28] S. S. Shapiro, "Privacy by Design," Commun. ACM, vol. 53, no. 6, jun 2010, p. 27.

[29] G. Kambourakis, "Anonymity and closely related terms in the cyberspace: An analysis by example," J. Inf. Secur. Appl., vol. 19, no. 1, 2014, pp. 2–17.

[30] J. Singh, T. F. J. M. Pasquier, and J. Bacon, "Securing tags to control information flows within the Internet of Things," 2015 Int. Conf. Recent Adv. Internet Things, RIoT 2015, 2015.

[31] EU, "Accountability for Cloud (A4Cloud)," 2018. [Online]. Available: http://a4cloud.eu/ [Last accessed: August 2018]

[32] C. A. Adams and R. Evans, "Accountability, Completeness, Credibility and the Audit Expectations Gap," JCC 14 Summer 2014, vol. 14, no. Summer, 2014, pp. 97–115.

[33] W. Benghabrit, H. Grall, J.-C. Royer, M. Sellami, M. Azraoui, K. Elkhiyaoui, M. Onen, A. S. D. Olivera, and K. Bernsmed, "A Cloud Accountability Policy Representation Framework," in CLOSER-4th Int. Conf. Cloud Comput. Serv. Sci., 2014, pp. 489–498.

[34] K. Bernsmed, W. K. Hon, and C. Millard, "Deploying Medical Sensor Networks in the Cloud – Accountability Obligations from a European Perspective," in Cloud Comput. (CLOUD), 2014 IEEE 7th Int. Conf. IEEE Comput. Soc, 2014, pp. 898–905.

[35] D. Butin, M. Chicote, and D. Le Métayer, "Log design for accountability," in Proc. - IEEE CS Secur. Priv. Work. SPW 2013, 2013.

[36] D. Catteddu, M. Felici, G. Hogben, A. Holcroft, E. Kosta, R. Leened, C. Millard, M. Niezen, D. Nunez, N. Papanikolaou, S. Pearson, D. Pradelles, C. Reed, C. Rong, J.-C. Royer, D. Stefanatou, and T. W. Wlodarczyk, "Towards a Model of Accountability for Cloud Computing Services," in Int. Work. Trust. Account. Forensics Cloud, 2013, pp. 21–30.

[37] A. Haeberlen, "A Case for the Accountable Cloud," ACM SIGOPS Oper. Syst. Rev., vol. 44, no. 2, 2010, pp. 52–57.

[38] W. Hon, E. Kosta, C. Millard, and D. Stefanatou, "Cloud Accountability: The Likely Impact of the Proposed EU Data Protection Regulation," Queen Mary Sch. Law Leg. Stud. Res. Pap., no. 172, 2014, pp. 1–54.

[39] K. L. R. Ko, P. Jagadpramana, and B. S. Lee, "Flogger: A File-Centric Logger for Monitoring File Access and Transfers within Cloud Computing Environments," Computing, 2011, pp. 1–8.

[40] T. Lynn, P. Healy, R. McClatchey, J. Morrison, C. Pahl, and B. Lee, "The Case for Cloud Service Trustmarks and Assurance-as-a-Service," in CLOSER 2013 - Proc. 3rd Int. Conf. Cloud Comput. Serv. Sci., 2013, pp. 110–115.

[41] N. Papanikolaou, S. Pearson, M. C. Mont, and R. K. L. Ko, "Towards Greater Accountability in Cloud Computing through Natural-Language Analysis and Automated Policy Enforcement," Engineering, 2011, pp. 1–4.

[42] N. Papanikolaou, T. Rübsamen, and C. Reich, "A Simulation Framework to Model Accountability Controls for Cloud Computing," CLOUD Comput. 2014, Fifth Int. Conf. Cloud Comput. GRIDs, Virtualization, no. c, 2014, pp. 12–19.

[43] S. Pearson, M. C. Mont, and G. Kounga, "Enhancing Accountability in the Cloud via Sticky Policies," in Secur. Trust Comput. Data Manag. Appl., 2011, pp. 146–155.

[44] S. Pearson, V. Tountopoulos, D. Catteddu, S. Mario, R. Molva, C. Reich, S. Fischer-Hubner, C. Millard, V. Lotz, M. G. Jaatun, R. Leenes, C. Rong, and J. Lopez, "Accountability for Cloud and Other Future Internet Services," in CloudCom, 2012, pp. 629—-632.

[45] K. Bernsmed and S. Fischer-Hübner, "Secure IT Systems: 19th Nordic Conference, NordSec 2014 Tromsø, Norway, October 15-17, 2014 Proceedings," Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics), vol. 8788, 2014, pp. 3–24.

[46] T. Ruebsamen and C. Reich, "Supporting Cloud Accountability by Collecting Evidence Using Audit Agents," in CloudCom 2013, 2013, pp. 185–190.

[47] J. Singh, J. Powles, T. Pasquier, and J. Bacon, "Data Flow Management and Compliance in Cloud Computing," Cloud Comput., no. Special Issue on Legal Clouds., 2015, pp. 1–12.

[48] A. Squicciarini, S. Sundareswaran, D. Lin, and S. Huang, "Promoting Distributed Accountability in the Cloud," in IEEE 4th Int. Conf. Cloud Comput. Promot., 2011, pp. 113–120.

[49] S. Sundareswaran, A. C. Squicciarini, and D. Lin, "Ensuring Distributed Accountability for Data Sharing in the Cloud," in IEEE Trans. Dependable Secur. Comput., vol. 9, no. 4, 2012, pp. 556–568.

[50] M. Theoharidou, N. Papanikolaou, S. Pearson, and D. Gritzalis, "Privacy risk, security, accountability in the cloud," in Proc. Int. Conf. Cloud Comput. Technol. Sci. CloudCom, vol. 1, 2013, pp. 177–184.

[51] PWC, "UK Information Security Breaches Survey - Technical Report 2012," PWC2012, Tech. Rep. April, 2012.

[52] Trustwave, "2012 Global Security Report," Tech. Rep., 2012.

[53] Verizon, "2016 Verizon Data Breach Report," Tech. Rep., 2016.

[54] D. Marcus and R. Sherstobitoff, "Dissecting Operation High Roller," White Pap. McAfee, vol. 000, 2012, pp. 1–20.

[55] T. Moore and N. Christin, "Beware the middleman: Empirical analysis of Bitcoin-exchange risk," in International Conference on Financial Cryptography and Data Security. Springer, 2013, pp. 25–33.

[56] M. Babaioff, S. Dobzinski, S. Oren, and A. Zohar, "On bitcoin and red balloons," in Proceedings of the 13th ACM conference on electronic commerce. ACM, 2012, pp. 56–73.

[57] I. Eyal and E. G. Sirer, "Majority is not enough: Bitcoin mining is vulnerable," in International Conference on Financial Cryptography and Data Security. Springer, 2014, pp. 436–454.

[58] N. Houy, "The economics of Bitcoin transaction fees," 2014.

[59] G. O. Karame, E. Androulaki, and S. Capkun, "Double-spending fast payments in bitcoin," in Proceedings of the 2012 ACM conference on Computer and communications security. ACM, 2012, pp. 906–917.

[60] E. K. Kogias, P. Jovanovic, N. Gailly, I. Khoffi, L. Gasser, and B. Ford, "Enhancing bitcoin security and performance with strong consistency via collective signing," in 25th USENIX Security Symposium (USENIX Security 16). USENIX Association, 2016, pp. 279–296.

[61] N. Christin, "Traveling the silk road: A measurement analysis of a large anonymous online marketplace," in Proceedings of the 22nd international conference on World Wide Web. ACM, 2013, pp. 213–224.

[62] R. Böhme, N. Christin, B. Edelman, and T. Moore, "Bitcoin: Economics, technology, and governance," The Journal of Economic Perspectives, vol. 29, no. 2, 2015, pp. 213–238.

[63] E. Androulaki, G. O. Karame, M. Roeschlin, T. Scherer, and S. Capkun, "Evaluating user privacy in bitcoin," in International Conference on Financial Cryptography and Data Security. Springer, 2013, pp. 34–51.

[64] J. S. Gans and H. Halaburda, "Some economics of private digital currency," in Econ. Anal. Digit. Econ. University of Chicago Press, 2015, pp. 257–276.

[65] G. P. Dwyer, "The economics of Bitcoin and similar private digital currencies," J. Financ. Stab., vol. 17, 2015, pp. 81–91.

[66] K. Krombholz, A. Judmayer, M. Gusenbauer, and E. Weippl, "The other side of the coin: User experiences with bitcoin security and privacy," in International Conference on Financial Cryptography and Data Security. Springer, 2016, pp. 555–580.

[67] Coindesk, "Coindesk," 2017. [Online]. Available: https://www.coindesk.com/ [Last accessed: August 2018]

[68] E.-T. Cheah and J. Fry, "Speculative bubbles in bitcoin markets? an empirical investigation into the fundamental value of bitcoin," Economics Letters, vol. 130, 2015, pp. 32–36.

[69] B. M. Blau, "Price dynamics and speculative trading in bitcoin," Research in International Business and Finance, vol. 41, 2017, pp. 493–499.

[70] P. Katsiampa, "Volatility estimation for bitcoin: A comparison of garch models," Economics Letters, vol. 158, 2017, pp. 3–6.

[71] F. Glaser, K. Zimmermann, M. Haferkorn, M. C. Weber, and M. Siering, "Bitcoin-asset or currency? revealing users' hidden intentions," 2014.

[72] M. Brière, K. Oosterlinck, and A. Szafarz, "Virtual currency, tangible return: Portfolio diversification with bitcoin," Journal of Asset Management, vol. 16, no. 6, 2015, pp. 365–373.

[73] N. Gandal and H. Halaburda, "Can we predict the winner in a market with network effects? competition in cryptocurrency market," Games, vol. 7, no. 3, 2016, p. 16.

[74] D. Yermack, "Is Bitcoin a real currency? An economic appraisal," National Bureau of Economic Research, Tech. Rep., 2013.

[75] A. F. Bariviera, M. J. Basgall, W. Hasperué, and M. Naiouf, "Some stylized facts of the Bitcoin market," Phys. A Stat. Mech. its Appl., vol. 484, 2017, pp. 82–90.

[76] A. H. Dyhrberg, "Bitcoin, gold and the dollar–a garch volatility analysis," Finance Research Letters, vol. 16, 2016, pp. 85–92.

[77] A. H. Dyhrberg, "Hedging capabilities of bitcoin. is it the virtual gold?" Finance Research Letters, vol. 16, 2016, pp. 139–144.

[78] E. Bouri, G. Azzi, and A. H. Dyhrberg, "On the return-volatility relationship in the bitcoin market around the price crash of 2013," 2016.

[79] T. Bollerslev, "Generalized autoregressive conditional heteroskedasticity," J. Econom., vol. 31, no. 3, 1986, pp. 307–327.

[80] R. Engle, "Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models," J. Bus. Econ. Stat., vol. 20, no. 3, 2002, pp. 339–350.

[81] NXT, "NXT Platform," 2017. [Online]. Available: https://nxtplatform.org/ [Last accessed: August 2018]

[82] Zerocash, "Zerocash," 2017. [Online]. Available: http://zerocash-project.org/ [Last accessed: August 2018]

[83] Peercoin, "Peercoin," 2017. [Online]. Available: https://peercoin.net/ [Last accessed: August 2018]

[84] S. Tripathi, B. Gupta, A. Almomani, A. Mishra, and S. Veluru, "Hadoop based defense solution to handle distributed denial of service (ddos) attacks," J. Inf. Secur., vol. 4, no. 03, 2013, p. 150.

[85] M. Carlsten, H. Kalodner, S. M. Weinberg, and A. Narayanan, "On the Instability of Bitcoin Without the Block Reward," in Proc. 2016 ACM SIGSAC Conf. Comput. Commun. Secur. - CCS'16, 2016.

[86] A. Azaria, A. Ekblaw, T. Vieira, and A. Lippman, "MedRec: Using blockchain for medical data access and permission management," in Proc. - 2016 2nd Int. Conf. Open Big Data, OBD 2016, 2016.

[87] K. Christidis and M. Devetsikiotis, "Blockchains and Smart Contracts for the Internet of Things," 2016.

[88] T. T. A. Dinh, J. Wang, G. Chen, R. Liu, B. C. Ooi, and K.-L. Tan, "BLOCKBENCH: A Framework for Analyzing Private Blockchains," 2017.

[89] E. Gaetani, L. Aniello, R. Baldoni, F. Lombardi, A. Margheri, and V. Sassone, "Blockchain-based database to ensure data integrity in cloud computing environments," in CEUR Workshop Proc., vol. 1816, 2017.

[90] A. Kiayias and G. Panagiotakos, "On Trees, Chains and Fast Transactions in the Blockchain." IACR Cryptol. ePrint Arch., vol. 2016, 2016, p. 545.

[91] D. Yermack, "Corporate governance and blockchains," 2017.

Will Compliance with the New EU General Data Protection Regulation Lead to Better Cloud Security?

Bob Duncan
Business School
University of Aberdeen, UK
Email: bobduncan@abdn.ac.uk

***Abstract*—The EU General Data Protection Regulation (GDPR) came into effect across the EU on 25th May 2018. It will certainly be the case that a great many companies will be inadequately prepared for this significant event. While a great many companies who use traditional in-house distributed systems are likely to have a hard enough job trying to comply with this new regulation, those who use any form of cloud computing face a particularly difficult additional challenge, namely the Cloud Forensic Problem. It is not enough that cloud use presents a far more challenging environment, but that the cloud forensic problem presents a far more difficult barrier to achieving compliance. This problem arises due to the fact that all computing systems are constantly under serious attack, but once an attacker gains a foothold in a cloud system and becomes an intruder, there is very little to prevent the intruder from helping themselves to any manner of data covered by the GDPR, either by viewing it, modifying it, deleting it or ex-filtrating it from the victim system. Worse, there is nothing to prevent the intruder from gaining sufficient privileges to then completely delete all trace of their incursion, possibly deleting far more records than they need to in the process. We address exactly what the requirements of EU GDPR compliance are, consider whether this can be done without resolving the Cloud Forensic Problem, and propose some approaches to mitigate this problem, and possibly the massive potential fines that could then be levied. We then consider whether the new EU GDPR will provide enough incentive for cloud users, and cloud service providers to get together to develop a much higher standard of cloud security which is both stronger than at present, and can deal with the Cloud Forensic Problem.**

Keywords–EU GDPR; Compliance; Cloud computing; cloud forensic problem.

I. INTRODUCTION

In [1], we considered the potential implications for cloud users in light of the cloud forensic problem for the then forthcoming EU GDPR compliance. We observed that during the drafting process of the regulation, one of the really useful componenents of the regulation was the requirement to report a breach within 72 hours of its occurrence. This brought a huge amount of effort to bear by corporates, desperate to ensure they would be able to comply. These efforts were reflected in the security breach reports, where it was apparent that the time between breach and discovery was reducing year on year. This could only be a good thing for all companies, and in particular cloud users.

Sadly, as a result of some intense lobbying, this component was somewhat watered down to a requirement to report within 72 hours of discovering the occurrence of a breach. As a direct result of this change, many companies instantly stopped working on this element of improving security, and again this too was reflected in the security breach reports, where the time between breach and discovery rocketed back to 2012 levels.

The EU General Data Protection Regulation (GDPR) [2], came into effect on 25th May 2018, and is likely to present one of the greatest compliance challenges faced by companies across the globe. Every company that trades anywhere on earth, should they deal with even a single EU resident, must ensure they are compliant with the EU GDPR. If that company suffers a security breach and the records of any EU citizen are compromised, then the jurisdiction of the GDPR will extend globally, and that company may be pursued and fined significant sums of money.

Achieving information security is a big enough challenge for companies who use conventional distributed network systems, but once companies start using cloud systems, the challenge increases exponentially. There are many reasons for this, mostly arising from the complexity of the additional relationships, and agendas, of different participant companies involved in cloud systems. Much research has been carried out to attempt to resolve these problems e.g., [3], [4], [5], [6], [7].

The most challenging, and as yet, unresolved issue is the cloud forensic problem, otherwise known as "The elephant in the room." Pretty much everyone knows about it, yet nobody is prepared to discuss it, let along try to resolve the problem, due to the difficulty of the challenge it presents. The new EU GDPR means that heads can no longer be left in the sand. This will not present an acceptable defence.

If any company using cloud is unable to resolve the cloud forensic problem, we suggest this will present such a fundamental issue that it will be impossible for that company to comply with this new regulation. As far back as 2011 and in subsequent years [8], [9], [10], [11], a great deal of research was focussed on trying to resolve this issue, yet it is clear from looking at regulatory fines for breaches that the message is not getting though.

In 2012, Verizon estimated that a total of 174 million data records were compromised [12]. By 2017, this had increased to an estimated 2 billion records lost or compromised in the first half of 2017 alone [13]. Yahoo disclosed a 1 billion compromised account breach in the 2013 attacks, yet when Verizon were in the process of taking over Yahoo last year and performing their due diligence, it turned out that **ALL 3 billion accounts** had been compromised [14], representing the biggest hack of all time.

In Section II, we look in some detail at the EU GDPR and consider the implications of non-compliance for any company that falls under its jurisdiction. In Section III, we identify what

the Cloud Forensic Problem is, and address why it is such a challenging problem to overcome. In Section IV, we ask whether it is possible to attain compliance without addressing the cloud forensic problem. In Section V, we address the minimum requirements necessary to achieve compliance. In Section VII, we look at what achieving the minimum requirements will allow us to do. In Section VIII, we consider the attitude of the regulator based on recently reported opinions made publicly by the regulator. In Section IX, we consider the likely attitude of corporate cloud users in response to these opinions. In Section X, we ask whether compliance with the GDPR might ever improve cloud security. In Section XI, we consider the limitations of this work, and in Section XII, we discuss our conclusions.

II. The EU General Data Protection Regulation

Why should companies be concerned about compliance with the EU GDPR [15]? Perhaps suffering a serious cyber breach leading to non-compliance, and resulting in a potential maximum fine of the greater of €20million or 4% of global turnover might serve to gain their attention. We should therefore take a good close look at the detail of the regulation.

The Article 29 Working Party [16] was set up by the European Commission under the terms of Article 29 of the Data Protection Directive in 1996, and its main stated missions are to:

- Provide expert advice to the States regarding data protection;
- Promote the consistent application of the Data Protection Directive in all EU state members, as well as Norway, Liechtenstein and Iceland;
- Give to the Commission an opinion on community laws (first pillar) affecting the right to protection of personal data;
- Make recommendations to the public on matters relating to the protection of persons with regard to the processing of personal data and privacy in the European Community.

During the time it has been active, the Article 29 Working Party has overseen the evolution of the GDPR, and has seen thousands of amendments proposed. One of the best proposals was the requirement to report all breaches ". . . within 72 hours of the breach occurring", which would have had the impact of ensuring that all organisations would give security top priority in order to achieve compliance. However, following much lobbying, this was watered down to ". . . within 72 hours of discovery of a breach." This rather took the urgency away from organisations, since many companies now took the view that until the breach happened, they would still be in compliance, resulting in many abandoning all efforts to improve security further.

Sadly, the impact of this change has been reflected in cyber breach reports. The global average time for all companies between breach and discovery in 2012 was an average of 6 months[17][18]. This had improved to just under 4 weeks by 2016 [19] — still far short of what is needed to understand what has been going on with the intruders while they were undiscovered. While this was a marked improvement over the intervening years, once the relaxation of the regulation took place, a great many companies immediately stopped working on security, taking the view that there would be no need to improve security as they would not be in breach of GDPR compliance until after a breach actually occurred. This rather short sighted view resulted in the time between breach and discovery reverting towards 2012 levels [20]. As Verizon [13] succinctly put it, "Apparently, it is not only The Eagles that are destined for a long stay at the hotel. The hackers continue to be checked in indefinitely as well. Breach timelines continue to paint a rather dismal picture — with time-to-compromise being only seconds, time-to-exfiltration taking days, and times to discovery and containment staying firmly in the months camp." That will not exactly fill the regulator with confidence about any company's ability to achieve compliance.

On a more positive note, another key amendment involved broadening the scope of the regulation, from all organisations anywhere in the EU, to any organisation anywhere in the globe, which stores privately identifiable information relating to any individual resident anywhere in the EU. This will certainly get the attention of far more organisations than would have been the case had it been an EU only requirement.

In the next three subsections, we have a look at how the GDPR seeks to streamline activities for both organisations and data subjects; how the GDPR will use enforcement mechanisms to ensure compliance; and what happens in the event of a data breach.

A. The Streamlining Goals of the GDPR

1) For Organisations: The idea for organisations is to streamline compliance by providing:

A single set of rules which would apply anywhere in the EU and by using the One Stop Shop approach, covered by Articles 46 to 55 of the GDPR, this would make for a streamlined approach for all organisations, whether based inside or outside the EU.

2) For Data Subjects: The idea for data subjects is to make the whole process for them much simpler by providing:

- Right of Access (under Article 15) - which gives data subjects the right to access their personal data held by any company subject to compliance with the GDPR;
- Right to Erasure (under Article 17) - which gives data subjects the right to have erasure carried out on certain data held by organisations about the data subject on any one of a number of grounds including non-compliance with article 6.1 (lawfulness) that includes a case (f) where the legitimate interests of the controller is overridden by the interests or fundamental rights and freedoms of the data subject;
- Data Portability (under Article 20) - data subjects have certain rights to data portability (particularly in the case of social media accounts), whereby a person shall be able to transfer their personal data from one electronic processing system to and into another, without being prevented from doing so by the data controller;
- Data Protection by Design and by Default (under Article 25) - seeks to ensure that all data subjects can expect privacy by design and by default, that has been designed into the development of business

processes for products and services. This requires that privacy settings must be set at a high level by default and that technical and procedural measures should be taken care of by the controller in order to make sure that the processing, throughout the whole processing lifecycle, complies with the regulation. A report by the European Union Agency for Network and Information Security (ENISA) [21], elaborates on what needs to be done to achieve privacy and data protection by design. It specifies that encryption and decryption operations must be carried out locally, not by remote service, because both keys and data must remain in the power of the data owner if any privacy is to be achieved. Furthermore, it specifies that outsourced data storage on remote clouds is practical and relatively safe, as long as only the data owner, not the cloud service, holds the decryption keys;

- Consent by Data Subjects - data subjects must have given their consent for data about them to be processed, thus providing a lawful basis for processing.

3) A Lawful Basis for Processing: The data subject must have given consent which must be explicit for data collected and the purposes data is used for (Article 7; defined in Article 4). Data controllers must be able to prove "consent" (opt-in) and consent may be withdrawn. Consent for children must be given by the child's parent or custodian, and must be verifiable (Article 8). Such consent to the processing of his, her or their personal data for one or more specific processing purposes, must be:

- necessary for the performance of a contract to which the data subject is party or in order to take steps at the request of the data subject prior to entering into a contract;
- necessary for compliance with a legal obligation to which the controller is subject;
- necessary in order to protect the vital interests of the data subject or of another natural person;
- necessary for the performance of a task carried out in the public interest or in the exercise of official authority vested in the controller;
- necessary for the purposes of the legitimate interests pursued by the controller or by a third party, except where such interests are overridden by the interests or fundamental rights and freedoms of the data subject which require protection of personal data, in particular where the data subject is a child.

B. Enforcement Mechanisms

- Appointing a Data Protection Officer - this person would be required for all data processor organisations, and a person with expert knowledge of data protection law and practices should assist the controller or processor to monitor internal compliance with this Regulation. The appointment of a DPO within a large organization will be a challenge for the Board as well as for the individual concerned, due to the myriad governance and human factor issues that organisations and companies will need to address given the scope and nature of the appointment. In addition, the post holder will need to create their own support team and will also be responsible for their own continuing professional development as they need to be independent of the organization that employs them, effectively as a "mini-regulator";
- Ensuring Compliance with the GDPR, by checking that all the correct mechanisms are properly defined and in place, mainly through compliance demonstration, e.g, the data controller should implement measures which meet the principles of data protection by design and data protection by default. Such measures include the process of pseudonymising (Recital 78), i.e., by means of encryption, which process should be completed as soon as is practically possible.
- The GDPR seeks to provide Responsibility and Accountability by all parties involved in data processing, with expanded notice requirements covering retention time for personal data, and contact information for data controller and data protection officer. Automated decision-making for individuals, including algorithmic means of profiling (Article 22), which is regarded as contestable, similar to the Data Protection Directive (Article 15), receive particular attention. The expectation is that all actors involved in the whole process of data processing will behave responsibly and will be fully accountable for their actions. Data Protection Impact Assessments (Article 35) have to be conducted when specific risks occur to the rights and freedoms of data subjects. Risk assessment and mitigation is required and prior approval of the Data Protection Authorities (DPA) is required for high risks. Data Protection Officers (Articles 3739) are to ensure compliance within organizations.

C. In the event of a Data Breach

In the event of a data breach, under the GDPR, the Data Controller will be under a legal obligation to notify the Supervisory Authority without undue delay. The reporting of a data breach is not subject to any *de minimis* standard and must be reported to the Supervisory Authority within 72 hours after having become aware of the data breach (Article 33). Individuals have to be notified if adverse impact is determined (under Article 34), unless the data was encrypted. In addition, the data processor will have to notify the controller without undue delay after becoming aware of a personal data breach (under Article 33).

1) Sanctions: The following sanctions can be imposed:

- a warning in writing in cases of first and non-intentional non-compliance;
- regular periodic data protection audits;
- a fine of up to €10million or up to 2% of the annual worldwide turnover of the preceding financial year in case of an enterprise, whichever is greater, where there has been an infringement of the following provisions (Article 83, Paragraph 4[18]):
    - the obligations of the controller and the processor pursuant to Articles 8, 11, 25 to 39 and 42 and 43;
    - the obligations of the certification body pursuant to Articles 42 and 43;

  - the obligations of the monitoring body pursuant to Article 41(4).
- a fine up to €20million or up to 4% of the annual worldwide turnover of the preceding financial year in case of an enterprise, whichever is greater, where there has been an infringement of the following provisions: (Article 83, Paragraph 5 & 6[18]):
  - the basic principles for processing, including conditions for consent, pursuant to Articles 5, 6, 7 and 9;
  - the data subjects' rights pursuant to Articles 12 to 22;
  - the transfers of personal data to a recipient in a third country or an international organisation pursuant to Articles 44 to 49;
  - any obligations pursuant to Member State law adopted under Chapter IX;
  - non-compliance with an order or a temporary or definitive limitation on processing or the suspension of data flows by the supervisory authority pursuant to Article 58(2) or failure to provide access in violation of Article 58(1).

The above details provide the essence of what we need to know in order to understand what information will be required to be delivered in the event of breach, in order for the data processor to be compliant with the GDPR. In the next section, we will take a look at the Cloud Forensic Problem, and why it is such a difficult problem, not only from the security perspective, but also from the GDPR compliance problem.

III. The Cloud Forensic Problem (and Why it is Such a Difficult Problem)

All computer systems are continuously subject to attack, whether traditional distributed network systems or cloud systems, which are no exception. It is certainly the case that no system is immune to attack, and that is particularly true for cloud systems. During the past decade, a great many research papers have allowed a far greater level of security and privacy to be achieved in cloud systems. There have been many good papers produced on both security [22], [23], [24], [25], [3], [4], [26], [5], [6], [7], [27], [28] and privacy [29], [30], [31], [32], [33], [34], [35], [36], [37], [6], [38], [39], [40], [41], [42], [43], and a number of others have looked at better accountability as a means to meeting these ends [44], [45], [46], [47], [48], [49], [50], [31], [51], [52], [3], [4], [53], [54], [55], [38], [7], [56], [57], [58], [41], [59], [11], [60], [61], [62] However, despite all those efforts, no solutions have yet been found to address the cloud forensic problem.

As we have already stated, all computing systems are constantly under serious attack, but once an attacker gains a foothold in a cloud system and becomes an intruder, there is little to prevent the intruder from helping themselves to any amount of data covered by the GDPR, either by viewing it, modifying it, deleting it or ex-filtrating it from the victim system [63], [64], [65]. Worse, there is nothing to prevent the intruder from gaining sufficient privileges to then completely delete all trace of their incursion, possibly deleting far more records than they need to in the process, leading to further problems for business continuity.

Often, companies do not retain records of which database records have been accessed, nor by whom. This means that once a breach occurs, the ability of the company to be able to report which records have been accessed, copied, modified, deleted or ex-filtrated from the system becomes an impossible task. This results in non-compliance with the GDPR, meaning exposure to potentially punitive levels of fines.

This is often known as "The elephant in the room" in cloud circles. Pretty much everyone knows about it, yet nobody is prepared to discuss it, let alone try to resolve the problem, due to the difficulty of the challenge it presents. Make no mistake, this is a serious challenge to defend against, let alone overcome. However, not only is it a serious challenge for organisations using cloud, it also presents a major obstacle to compliance with the GDPR.

Once all trace of the intrusion has been deleted, there will be very little forensic trail left to follow, meaning many companies will be totally unaware that the intrusion has taken place, let alone understand what records have been accessed, modified, deleted or stolen. All too often, companies will believe they have retained a full forensic trail in their running instance, but often forget that without special measures being taken to save these records off-site [3], they will vanish when the instance is shut down.

Currently, in any cloud based system, there must be a complete and intact audit trail in order for the breached organisation to be able to tell which records have been accessed, modified, deleted or stolen. Where the audit trail and all forensic records have been deleted, there remains no physical means for any organisation to be able to tell which records have been accessed, modified, deleted or stolen, putting these organisations immediately in multiple breaches of the GDPR.

IV. Is it possible to achieve Compliance with the EU GDPR without addressing the Cloud Forensic Problem?

The short answer is, of course, it is not! For the reasons outlined in the previous section, we can see that there is nothing to prevent an intruder from destroying every scrap of forensic proof of their incursion into any current cloud system. It is clear that any form of forensic record or audit trail can not therefore be safely stored on any running cloud instance.

This means that the only safe method of storage of forensic data will be somewhere off-site from the running cloud instances. Clearly, the off-site storage must be highly secure, preferably stored in an append-only database, and should especially be held in encrypted format, with all encryption keys held elsewhere.

Doubtless some will say that as long as they are not breached, then they will not be in breach of the GDPR. While that may very well be true, how will they be able to tell whether they have or have not been breached, particularly in the circumstance where they have been breached, and the breach has been very well covered up. They will have no means of knowing, let alone proving the point. The regulator will be unlikely to accept this approach as an appropriate way to demonstrate a willingness to comply with the GDPR.

Let us suppose that a complaint is made to the regulator, the organisation will have no means of proving that the data has not been tampered with. Equally, if the breach has been extremely well covered up, they will neither have the means of complying with the requirement to: a) report the breach

within 72 hours, nor b) have any means of determining which records have been accessed, modified, deleted or stolen. Let us now suppose that the conversion of private data has yet to be encrypted, and worse, that the encryption and decryption keys are held on the cloud instance "for convenience". If we were to receive a request from any users whose data had just been compromised, we would be unable to comply with the request, meaning we would now be looking at multiple breaches, thus causing the fine level to escalate to the higher level, as outlined in Subsubsection II-C1.

An added inconvenience would arise where the company had elected not to use encryption (or had used encryption, but left the encryption and decryption keys on the cloud instance). While encryption is not mandatory, in the case where it is not used, in the event of a breach, the company must communicate with all customers whose data may have been compromised. Where they are unable to tell whose data has or has not been compromised, they would need to write to every single customer to be in compliance. This could prompt a flood of requests from these customers to enquire about specifically which records of theirs were compromised. The company would be unable to provide this information, and would then enter into a case of multiple breaches of the GDPR, leading to the possibility of multiple large fines for non-compliance.

V. The Minimum Requirements to Achieve Compliance with the GDPR

We have seen that to do nothing would not be a viable option as far as GDPR compliance is concerned. Attacks will continue unabated. We must therefore be prepared and armed with whatever tools we can develop to ensure we achieve as high a level of compliance as we possibly can.

We therefore need to consider what the absolute minimum technical requirement might be to attain our objective of GDPR compliance. We know that under the GDPR the organisation must be able to:

- provide a Right of Access (under Article 15) to personal data by data subject, if requested;
- provide the means to comply with a Right to Erasure (under Article 17) by data subject, subject to the appropriate grounds being met;
- provide privacy by design;
- in the event of a data breach, report the breach to the Supervisory Authority within 72 hours after having become aware of the data breach (Article 33). The breach must also be reported to the controller without undue delay after becoming aware of a personal data breach;
- in the event of a data breach, notify the data subject if adverse impact is determined (under Article 34), unless the data was encrypted;

In the case of the first requirement, we would require to ensure the provenance and veracity of the contents of the database. In the case of the second requirement, if appropriate, the same provision would apply.

In the case of the third requirement, the cloud system must be designed in accordance with the recommendations of the Article 29 Working Party [66], which suggests the reports produced by ENISA should be followed. This report [67] specifies that encryption and decryption operations must be carried out locally, not by remote service, because both keys and data must remain in the power of the data owner if any privacy is to be achieved. Furthermore, it specifies that outsourced data storage on remote clouds is practical and relatively safe, as long as only the data owner, not the cloud service, holds the decryption keys. ENISA have also produced a stream of other relevant reports, including a Cloud Risk report in 2009 [68], and recommendations for certification in 2017 [69].

In the case of the fourth requirement, we would require to ensure the provenance and veracity of the contents of the database. In the case of the fifth requirement, where the data is not yet encrypted, the same provision would also apply. However, it should be stressed that it will always be preferable to ensure data is encrypted before it leaves the control of the data owner.

It is clear that where no steps have been taken to ensure the cloud forensic problem has been mitigated, the organisation will fail on every count. Thus, as a minimum, we need to ensure the following steps are taken:

- all personal data should be encrypted, and this should be performed locally;
- the encryption and decryption keys should not be maintained on the cloud instance;
- a full audit trail of the entire database must be maintained off-site;
- full forensic records of all users having accessed the database and carried out any commands on the database must be collected and stored off-site.

VI. Architecture Changes Suggested

The starting position will be a conventional cloud instance containing everything needed to operate the system, including web based software, database software, intrusion detection software and anything else deemed to be appropriate.

A. The Bare Minimum

All database access requests, database logs, system logs and any other logs should be running on a separate high security system, away from the main cloud instance. This system should not have any conventional web interface, and the recording databases should be immutable, i.e. append only.

This approach will address the challenge of retaining a full forensic trail discussed in Section III.

B. The Improved Version

All database software should be removed from the cloud system and run on a highly secure system which is separate from both the main cloud instance and the forensic logging system. This ensures the complete separation of all data from software running on the main cloud instance.

Depending on the volume of transactional data, this system can run on a conventional distributed system, or a cloud system running only the database software or could be run across multiple virtualised machines.

VII. What will the Minimum Requirements Allow us to do?

Let us now assume that we have completed the bare minimum requirements. Can we now be sure that we can be compliant with the provisions of the GDPR? We must therefore look at each of the five reporting requirements in turn to establish whether we will be able to meet these requirements.

1) First, if a data subject serves us with a Right of Access request, can we respond in the affirmative? We are now sure that we hold the subject's data securely, in encrypted format in our database. Further, on the assumption that no breach has arisen, we can prove that the data has only been accessed by duly authorised persons because we have a complete forensic trail of everyone who has accessed the data records, and further that the data records have neither been modified, stolen nor deleted. We are therefore compliant on the first requirement;
2) Next, if a data subject serves us with a right to Erasure notice, can we comply with that request? Assuming the request can be legitimately carried out and is not prohibited by statute, then since we can correctly identify the private data held about the data subject, then there is no reason why we would be unable to delete the appropriate data as requested. Accordingly, we would be compliant on the second requirement;
3) Next, can we provide privacy by design? Our default design concept is to provide privacy by design through following the ENISA recommendations which suggest this be achieved by ensuring all private data is properly encrypted, that encryption and decryption keys are not stored on the running cloud instance, and that we retain a full and complete forensic record of all operations on the data held by the company;
4) In the event of a data breach, can we report the breach to the Supervisory Authority within 72 hours of discovery? In the case of a data breach, we will not only be able to notify the breach within 72 hours of discovery, we will actually be able to notify within 72 hours of the occurrence of the breach. In addition, since we will retain full forensic data and audit trails for the system, we will also be able to provide very precise details of which records were accessed and read, which might have been modified, with full details of what modifications were made, which records were deleted, and which records were exfiltrated from the system. Not only that, but we will be able to provide full details of how the perpetrators got into the system and where they forwarded any stolen records, which means we can identify precisely which records were compromised, thus ensuring we would be beyond fully compliant;
5) In the event of a data breach, would we be able to notify the data subject if adverse impact is determined (under Article 34)? In the event of a data breach, we would be able to identify every single record attacked, and identify every single data subject affected. Since the full records would already be encrypted, we would not be required to notify the data subjects, but would be fully capable of so doing. This would mean we would again be beyond fully compliant.

Thus, we can reasonably claim that we would be in a position to be fully compliant with all the requirements of the GDPR, thus providing an exceptionally high level of privacy on behalf of all data subjects. Thus, the level of exposure of data subjects would be extremely minimised, thus ensuring compliance with the regulation, and therefore the likelihood that we would be able to fully mitigate any penalty that would otherwise be applied by the regulator.

Contrast this position with the case where cloud users do not take these mitigatory steps. In every requirement - they would be non-compliant, thus exposing the enterprise to the full extent of penalties allowed, namely the greater of €20million or 4% of global turnover.

VIII. The Attitude of the Regulators

Since at the time of writing this article, barely three months has elapsed since the GDPR came into effect, there will not yet be a great deal of indication on what the attitude of the regulator to cyber breach events is likely to be. In spite of the short timescale that has elapsed, the UK Information Commissioner's Office (ICO) who are the UK GDPR regulator have seen complaints rise from 2,417 to 6,281 between 25 May and 3 July 2018 as compared with the same period from the previous year. On the plus side, they have increased staff by some 40% in anticipation of this significant increase in workload.

However, of a Reuters' survey of 24 of the authorities charged with carrying out the regulation of the GDPR who responded in early May, 2018, just weeks before the GDPR came into force, 17 responded that they did not yet have the necessary funding, or would initially lack the powers to fulfil their GDPR duties. Since many of these new powers have yet to be incorporated into their countries' laws, this is likely to result in a number of delays before any serious regulatory effort can be started. Many have said they will start by responding to complaints and investigate them on merit. Only a minority suggested they would proactively investigate whether companies were complying and make any attempts to sanction glaring non-compliance [70].

The expectation of the regulator will be that they would expect companies to take all reasonable steps to make their business compliant with the GDPR. However it is likely that where a company has not taken sufficient robust steps to prepare to achieve adequate levels of security, this will be regarded as a failure to take proper steps to safeguard the PII of users, and the company will be regarded as complicit in aiding the attackers to perpetrate their attack. This will likely ensure a much higher level of penalty will be applied. However, following a rather embarrassing leak, it became apparent that the European Commission is not itself GDPR compliant [71], and of course now claims that it is exempt.

In the event that any company chooses not to use encryption, or decides to leave the encryption and decryption keys on the running cloud instances, the company will again be found to be complicit in failing to achieve proper compliance. Again, resulting in a likely increase in the level of penalty applied, as well as a huge administrative burden for notifying customers on top of the penalty.

Some regulators have taken the view that they will investigate cyber breaches that arose before the GDPR came into effect. Others are clearly not yet ready to regulate properly yet. Some will investigate on receipt of a complaint. Others will clearly wish to be proactive in their approach. Time will tell how each will approach their job, and what the likely consequences will be for non compliance.

With currently 28 member states, and considerably more regulatory authorities granted power to regulate under the GDPR, it is also not yet fully clear just how the various regulators will act where breaches affect cloud customers from more than one EU country or area, nor how jurisdiction will be dealt with where a large corporate operates in multiple EU countries or areas within.

There is no doubt that is too early to speculate on how the many EU regulators will approach their regulatory duties, and how they might go about enforcing compliance with cloud users. In some respects, the fact that many of the regulators have neither the resources nor the legislative power to carry out their regulatory duties means that there will be an element of respite for cloud users. There is no doubt that a great many corporates will be only too happy to take full advantage of this situation to minimise the work they carry out on improving their security systems in order to provide a much better standard of privacy.

IX. The Attitude of Corporate Cloud Users

Judging by the content of the annual reports during the past decades of large corporates, who are not renowned for exhibiting highly transparent levels of disclosure, this is unlikely to provide a good source of information on successful cyber breaches. A great many corporate boardrooms fear the prospect of disclosure of problems and the likely knock on effect on the share price. While they are required to report cyber breaches within 72 hours of discovery, in the event that they have used cloud and the forensic and audit trails have been tampered with, it is unlikely that they will even report a cyber breach when it arises. Cleaerly there will be an element of moral hazard to take into account at board level. Why would they wish to create trouble for themselves, a potentially significant drop in their share price, and a potentially large fine when they wait a while, perhaps until the dividend has been declared and paid out (along with their bonuses) before considering publication of the cyber breach or reporting the cyber breach to the regulator. This could certainly present a serious moral hazard when there may be little direct forensic evidence as to the extent of the breach.

Equally, while many corporates publicly proclaim their desire to be compliant with the new EU GDPR, Calligo, in a recent survey of IT decision makers, it was discovered that 69% of them do not have the backing of their board to achieve GDPR compliance [72]. However, once something goes wrong, it is likely the large multinational corporates, accustomed to dealing with regulation and compliance issues, will actually do something about it. In time, they wil refuse to do business with suppliers unless they too seek GDPR compliance. This will likely mean an eventual flow through all industries that are required to be compliant.

This is often the way with large corporates. Do nothing if at all possible until something goes wrong, and then take whatever action is necessary to become compliant. Then make all your suppliers become compliant too. Of course, there are always a few who do the right thing right at the beginning. It would seem a very prudent approach. No action usually means the breach will hurt. Not to mention the consequences in lost business, business continuity impact, loss of share price, embarrassment, and punitive fines.

Given the likely obstacles faced by the various regulators in getting started with the job of regulation due to being under-resourced, and perhaps having no or insufficient legislative ability to carry out their regulatory tasks, many large corporates will be happy to take advantage of that situation by sitting on carrying out the necessary improvements until it becomes absolutely essential.

In that event, it is highly likely that attackers will be more than happy to take full advantage of this slacking off on tightening cyber security by having a field day with few obstacles to get in their way.

X. Will Compliance with the GDPR Lead to Better Cloud Security?

It is very clear that, particularly in some areas, it will take some considerable time for proper regulation to be properly implemented, perhaps even years. There is no doubt as all that as soon as some punitive level of fines is levied against cloud users, thus punishing all of society through higher costs being levied by the cloud users to cover this potential major in crease in their cost base, then more effort is likely to go into improving cloud security. It is just a pity that we end up punishing society in general, rather than the perpetrators of the crimes who are responsible for all this mayhem.

It is clear that every actor involved in the cloud ecosystem has a role to play in improving security, and therefore privacy too. There is no doubt that major cloud service providers are taking security much more seriously these days. It is equally clear that many large corporates are much less inclined to do so, unless pushed, and pushed hard, and that very much needs to change.

There is a clear need for greater accountability from all involved. It is also clear that there is a need to develop a better means of policing the use of computing resources with a view to tracking the real perpetrators of the crimes. Equally, we need to consider that many of the computing standards we are all familiar with today have been in existence for a great many decades, most of which were developed before the internet took off.

This means that there is undoubtedly scope to tighten up these standards significantly in the light of the need for greater accountability and a better understanding of how to pin responsibility on all bad actors.

There is little doubt that a huge amount of work will be involved by a great many people. However, the introduction of punitive levels of fines will likely help to accelerate this process. There is no doubt this will lead to better cloud security. The question is how long will it really take to reach an acceptable level of cloud security?

There is also little doubt that the GDPR will have far reaching consequences for other jurisdictions, particularly for the US, where existing legislation and regulation fails to go anywhere close to what the new EU GDPR is doing. This will doubtless lead to more change throughout the globe to bring

more and more legislation and regulation into alignment. Ultimately, this will be a good thing for society as a whole. For too long, criminals have skipped around the insular jurisdictional approach of many countries which has led to myriad loopholes being exploited by criminals who continue to perpetrate their seedy trade with impunity.

XI. Limitations and discussion

There are two very important tasks that must be performed in order not to limit the effectiveness of this approach. Since persistent storage in the cloud instance cannot retain data beyond its currently running lifetime [3], we must also make sure that all necessary logs and data are stored securely elsewhere. And as the default settings for virtually all database management software involves logging being turned off [63], we must ensure this function is turned on in all running cloud instances, again, with the data being stored securely elsewhere.

This prompts the question of what data we require to keep. In order to meet our regulatory compliance requirement, we need to understand the 5 W's — namely: Who is accessing our system? Where have they come from? What are they looking for? When is this happening? From this data, we should be able to infer the Why? Are they authorised to be in the system, to enter the system the way they have, to look at the data they are trying to access, and at the time they are trying to access it? Deducing the Why can give an indicator of anomalous behaviour.

Many database management software offers additional full audit trail capabilities. Each additional capability will require more and more storage resources. A balance will need to be found between the minimum requirement consistent with maintaining performance and a cost effective level of storage. The risk in not utilising all that is on offer, would be that this might compromise security, reducing the ability of the company to achieve compliance.

However, it is clear that a sensible precaution to mitigate this risk would be to encrypt all the data being held on all databases maintained within the system, ensuring that encryption/decryption keys are not stored on the cloud instances. While encryption is not mandatory, in the event of a breach where encryption is not used, the fine levied by the regulator is likely to be much higher as a consequence. Additionally, the company must personally notify every single customer whose PII is at risk, or was compromised in the course of the breach.

However, cloud users should also consider the fact that all actors in the cloud ecosystem should also be contributing towards resolving these issues, and that includes in particular the cloud service provider (CSP). There is undoubtedly a need for greater accountability from every actor in the ecosystem chain. Everyone needs to contribute to making cloud computing a much safer paradigm for the benefit of all actors, and hopefully to the detriment of all attackers too.

XII. Conclusion

The forthcoming GDPR will certainly present a serious wake up call to a great many companies operating around the globe if they find themselves falling under the jurisdiction of this new regulation. In this paper, we have considered whether it is possible to achieve regulatory compliance where any organisation is using cloud computing. Again, we reiterate that without suitable precautions being put in place, the answer is a resounding "No!".

We have outlined the key requirements from the regulation to which all organisations falling under its jurisdiction must comply. We have identified the currently unresolved "Cloud Forensic Problem" as presenting the largest obstacle to achieving compliance.

We have proposed how this challenging problem may be approached to ensure that cloud users can be fully compliant with this new regulation, with little more than being sensibly organised. Clearly, additional cost will require to be incurred, and there may be a small impact on latency, but these costs could significantly mitigate the possibility of a huge regulatory fine in the event of a breach. It is also likely that this approach will ensure faster discovery of the occurrence of a breach, thus minimising the potential impact on business continuity.

Perhaps we can look forward to the day when we can put the squeeze on attackers, or at least have the ability to track and identify them , thus allowing us to make them fully accountable for their insidious trade. There is little doubt that right now, we are all in it together, and thus we must all pull together in order to have any chance of succeeding against the overwhelming hordes of attackers who end up making many people's lives such a misery. It is time to get serious.

References

[1] B. Duncan, "Can EU General Data Protection Regulation Compliance be Achieved When Using Cloud Computing?" in Cloud Computing 2018: The Ninth International Conference on Cloud Computing, GRIDs, and Virtualization, no. February. Barcelona, Spain: IARIA, 2018, pp. 1–6.

[2] EU, "EU General Data Protection Regulation (GDPR)," 2017. [Online]. Available: http://www.eugdpr.org/ [Last accessed: August 2018]

[3] R. K. L. Ko, P. Jagadpramana, M. Mowbray, S. Pearson, M. Kirchberg, B. S. Lee, and Q. Liang, "TrustCloud: A Framework for Accountability and Trust in Cloud Computing," Perspective, 2011, pp. 1–9.

[4] R. K. L. Ko, B. S. Lee, and S. Pearson, "Towards achieving accountability, auditability and trust in cloud computing," Communications in Computer and Information Science, vol. 193 CCIS, 2011, pp. 432–444.

[5] N. Papanikolaou, S. Pearson, and M. C. Mont, "Towards Natural-Language Understanding and Automated Enforcement of Privacy Rules and Regulations in the Cloud: Survey and Bibliography," Analysis, 2011, pp. 1–9.

[6] S. Pearson, "Taking Account of Privacy when Designing Cloud Computing Services," Current, 2009, pp. 44–52.

[7] S. Pearson, "Toward accountability in the cloud," IEEE Internet Computing, vol. 15, no. 4, jul 2011, pp. 64–69.

[8] N. Papanikolaou, S. Pearson, M. C. Mont, and R. Ko, "A Toolkit for Automating Compliance in Cloud Computing Services," International Journal of Cloud Computing, vol. x, no. x, 2014, pp. 45–68.

[9] J. Singh and J. M. Bacon, "On middleware for emerging health services," Journal of Internet Services and Applications, vol. 5, no. 1, 2014, p. 6.

[10] J. Singh, J. Bacon, and D. Eyers, "Policy Enforcement Within Emerging Distributed, Event-based Systems," Proceedings of the 8th ACM International Conference on Distributed Event-Based Systems - DEBS '14, 2014, pp. 246–255.

[11] J. Singh, J. Powles, T. Pasquier, and J. Bacon, "Data Flow Management and Compliance in Cloud Computing," Cloud Computing, no. Special Issue on Legal Clouds., 2015, pp. 1–12.

[12] Verizon, "2012 Data Breach Investigation Report: A study conducted by the Verizon RISK Team in cooperation with the United States Secret Service and Others," Tech. Rep., 2012. [Online]. Available: http://www.verizonenterprise.com/resources/reports/rp_data-breach-investigations-report-2012_en_xg.pdf [Last accessed: August 2018]

[13] Verizon, "Verizon Security Breach Report 2017," Tech. Rep., 2017.

[14] S. Khandelwal, "Its 3 Billion! Yes, Every Single Yahoo Account Was Hacked In 2013 Data Breach," 2017. [Online]. Available: https://thehackernews.com/2017/10/yahoo-email-hacked.html [Last accessed: August 2018]

[15] The European Parliament and The European Council, "General Data Protection Regulation," Official Journal of the European Union, vol. 2014, no. October 1995, 2016, pp. 20–30.

[16] EU, "Opinion 05/2012 on Cloud Computing (Data Protection)," 2012.

[17] PWC, "UK Information Security Breaches Survey - Technical Report 2012," PWC2012, Tech. Rep. April, 2012.

[18] Trustwave, "2012 Global Security Report," Tech. Rep., 2012.

[19] Verizon, "2016 Verizon Data Breach Report," Tech. Rep., 2016.

[20] Trustwave, "2017 Global Security Report," Trustwave, Tech. Rep., 2017.

[21] G. Danezis, J. Domingo-Ferrer, M. Hansen, J.-H. Hoepman, D. L. Metayer, R. Tirtea, and S. Schiffner, Privacy and Data Protection by Design - from policy to engineering, 2015, no. December.

[22] M. Almorsy, J. Grundy, and I. Müller, "An analysis of the cloud computing security problem," 17th Asia-Pacific Software Engineering Conference (APSEC 2010) Cloud Workshop,Sydney, Australia, no. December, 2010, p. 7.

[23] M. Almorsy, J. Grundy, and I. Mller, "An analysis of the cloud computing security problem." The proc. of the 2010 Asia Pacific Cloud Work-shop Colocated with APSEC2010, Australia, 2010.

[24] V. Chang, Y. H. Kuo, and M. Ramachandran, "Cloud computing adoption framework: A security framework for business clouds," Future Generation Computer Systems, vol. 57, 2016, pp. 24–41.

[25] F. Doelitzscher, T. Ruebsamen, T. Karbe, M. Knahl, C. Reich, and N. Clarke, "Sun Behind Clouds - On Automatic Cloud Security Audits and a Cloud Audit Policy Language," International Journal on Advances in Networks and Services, vol. 6, no. 1, 2013, pp. 1–16.

[26] K. Lee, "Security Threats in Cloud Computing Environments," International Journal of Security and its Applications, vol. 6, no. 4, 2012, pp. 25–32.

[27] S. Ramgovind, M. Eloff, and E. Smith, "The Management of Security in Cloud Computing," in Information Security for South Africa (ISSA), 2010, 2010, pp. 1–7.

[28] H. Takabi, J. B. Joshi, and G.-J. Ahn, "Security and Privacy Challenges in Cloud Computing Environments," IEEE Security & Privacy Magazine, vol. 8, no. 6, nov 2010, pp. 24–31.

[29] S. Creese, P. Hopkins, S. Pearson, and Y. Shen, "Data Protection-Aware Design for Cloud Computing," Work, no. December, 2009, pp. 1–13.

[30] K. Hon, C. Millard, and I. Walden, "The Problem of 'Personal Data' in Cloud Computing - What Information is Regulated ?" 2011.

[31] W. K. Hon, C. Millard, and I. Walden, "Who is Responsible for 'Personal Data' in Cloud Computing? The Cloud of Unknowing, Part 2," Legal Studies, no. 77, 2011, pp. 1–31.

[32] W. K. Hon, J. Hörnle, and C. Millard, "Data Protection Jurisdiction and Cloud Computing When are Cloud Users and Providers Subject to EU Data Protection Law?" Legal Studies, 2011, pp. 1–40.

[33] W. Itani, A. Kayssi, and A. Chehab, "Privacy as a Service: Privacy-Aware Data Storage and Processing in Cloud Computing Architectures," 2009 Eighth IEEE International Conference on Dependable, Autonomic and Secure Computing, dec 2009, pp. 711–716.

[34] W. Jansen and T. Grance, "Guidelines on Security and Privacy in Public Cloud Computing," NIST, Tech. Rep., 2011. [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-144/SP800-144.pdf [Last accessed: August 2018]

[35] H. Katzan Jr, "On The Privacy Of Cloud Computing," International Journal of Management and Information Systems, vol. 14, no. 2, 2011, pp. 1–12.

[36] W. K. Hon, C. Millard, J. Singh, I. Walden, and J. Crowcroft, "Policy, legal and regulatory implications of a Europe-only cloud," International Journal of Law and Information Technology, vol. 24, no. 3, 2016, pp. 251–278.

[37] NIST, "Security and Privacy Controls for Federal Information Systems and Organizations Security and Privacy Controls for Federal Information Systems and Organizations," Natioinal Institute of Standards and Technology, Gaithersburg, MD, Tech. Rep. February, 2014. [Online]. Available: http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r4.pdf [Last accessed: August 2018]

[38] S. Pearson and A. Charlesworth, "Accountability as a Way Forward for Privacy Protection in the Cloud," Computing, no. December, 2009, pp. 1–15.

[39] S. Pearson and A. Benameur, "Privacy, Security and Trust Issues Arising from Cloud Computing," 2010 IEEE Second International Conference on Cloud Computing Technology and Science, nov 2010, pp. 693–702.

[40] S. Pearson, "Privacy, Security and Trust in Cloud Computing," in Privacy and Security for Cloud Computing. e: Springer, 2013, pp. 3–42.

[41] J. Prüfer, "How to govern the cloud? Characterizing the optimal enforcement institution that supports accountability in cloud computing," Proceedings of the International Conference on Cloud Computing Technology and Science, CloudCom, vol. 2, 2013, pp. 33–38.

[42] S. S. Shapiro, "Privacy by Design," Communications of the ACM, vol. 53, no. 6, jun 2010, p. 27.

[43] J. Singh, T. F. J. M. Pasquier, and J. Bacon, "Securing tags to control information flows within the Internet of Things," 2015 International Conference on Recent Advances in Internet of Things, RIoT 2015, 2015.

[44] EU, "Accountability for Cloud (A4Cloud)," 2018. [Online]. Available: http://a4cloud.eu/ [Last accessed: August 2018]

[45] C. A. Adams and R. Evans, "Accountability, Completeness, Credibility and the Audit Expectations Gap," JCC 14 Summer 2014, vol. 14, no. Summer, 2014, pp. 97–115.

[46] W. Benghabrit, H. Grall, J.-C. Royer, M. Sellami, M. Azraoui, K. Elkhiyaoui, M. Onen, A. S. D. Olivera, and K. Bernsmed, "A Cloud Accountability Policy Representation Framework," in CLOSER-4th International Conference on Cloud Computing and Services Science, 2014, pp. 489–498.

[47] K. Bernsmed, W. K. Hon, and C. Millard, "Deploying Medical Sensor Networks in the Cloud Accountability Obligations from a European Perspective," in Cloud Computing (CLOUD), 2014 IEEE 7th International Conference on. IEEE Comput. Soc, 2014, pp. 898–905.

[48] D. Butin, M. Chicote, and D. Le Métayer, "Log design for accountability," in Proceedings - IEEE CS Security and Privacy Workshops, SPW 2013, 2013.

[49] D. Catteddu, M. Felici, G. Hogben, A. Holcroft, E. Kosta, R. Leened, C. Millard, M. Niezen, D. Nunez, N. Papanikolaou, S. Pearson, D. Pradelles, C. Reed, C. Rong, J.-C. Royer, D. Stefanatou, and T. W. Wlodarczyk, "Towards a Model of Accountability for Cloud Computing Services," in International Workshop on Trustworthiness, Accountability and Forensics in the Cloud (TAFC), 2013, pp. 21–30.

[50] A. Haeberlen, "A Case for the Accountable Cloud," ACM SIGOPS Operating Systems Review, vol. 44, no. 2, 2010, pp. 52–57.

[51] W. Hon, E. Kosta, C. Millard, and D. Stefanatou, "Cloud Accountability: The Likely Impact of the Proposed EU Data Protection Regulation," Queen Mary School of Law Legal Studies Research Paper, no. 172, 2014, pp. 1–54.

[52] K. L. R. Ko, P. Jagadpramana, and B. S. Lee, "Flogger: A File-Centric Logger for Monitoring File Access and Transfers within Cloud Computing Environments," Computing, 2011, pp. 1–8.

[53] T. Lynn, P. Healy, R. McClatchey, J. Morrison, C. Pahl, and B. Lee, "The Case for Cloud Service Trustmarks and Assurance-as-a-Service," in CLOSER 2013 - Proceedings of the 3rd International Conference on Cloud Computing and Services Science, 2013, pp. 110–115.

[54] N. Papanikolaou, S. Pearson, M. C. Mont, and R. K. L. Ko, "Towards Greater Accountability in Cloud Computing through Natural-Language Analysis and Automated Policy Enforcement," Engineering, 2011, pp. 1–4.

[55] N. Papanikolaou, T. Rübsamen, and C. Reich, "A Simulation Framework to Model Accountability Controls for Cloud Computing," CLOUD COMPUTING 2014, The Fifth International Conference on Cloud Computing, GRIDs, and Virtualization, no. c, 2014, pp. 12–19.

[56] S. Pearson, M. C. Mont, and G. Kounga, "Enhancing Accountability in the Cloud via Sticky Policies," in Secure and Trust Computing, Data Management, and Applications, 2011, pp. 146–155.

[57] S. Pearson, V. Tountopoulos, D. Catteddu, S. Mario, R. Molva, C. Reich, S. Fischer-Hubner, C. Millard, V. Lotz, M. G. Jaatun, R. Leenes, C. Rong,

and J. Lopez, "Accountability for Cloud and Other Future Internet Services," in CloudCom, 2012, pp. 629—-632.

[58] K. Bernsmed and S. Fischer-Hübner, "Secure IT Systems: 19th Nordic Conference, NordSec 2014 Tromsø, Norway, October 15-17, 2014 Proceedings," Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), vol. 8788, 2014, pp. 3–24.

[59] T. Ruebsamen and C. Reich, "Supporting Cloud Accountability by Collecting Evidence Using Audit Agents," in CloudCom 2013, 2013, pp. 185–190.

[60] A. Squicciarini, S. Sundareswaran, D. Lin, and S. Huang, "Promoting Distributed Accountability in the Cloud," in IEEE 4th International Conference on Cloud Computing Promoting, 2011, pp. 113–120.

[61] S. Sundareswaran, A. C. Squicciarini, and D. Lin, "Ensuring Distributed Accountability for Data Sharing in the Cloud," in IEEE Transactions on Dependable and Secure Computing, vol. 9, no. 4, 2012, pp. 556–568.

[62] M. Theoharidou, N. Papanikolaou, S. Pearson, and D. Gritzalis, "Privacy risk, security, accountability in the cloud," in Proceedings of the International Conference on Cloud Computing Technology and Science, CloudCom, vol. 1, 2013, pp. 177–184.

[63] B. Duncan and M. Whittington, "Enhancing Cloud Security and Privacy: The Power and the Weakness of the Audit Trail," in Cloud Computing 2016: The Seventh International Conference on Cloud Computing, GRIDs, and Virtualization, no. April. Rome: IEEE, 2016, pp. 125–130.

[64] G. Weir, A. Aßmuth, M. Whittington, and B. Duncan, "Cloud Accounting Systems, the Audit Trail, Forensics and the EU GDPR: How Hard Can It Be?" in The British Accounting and Finance Association: Scottish Area Group Annual Conference. Aberdeen: BAFA, 2017, p. 6.

[65] P. Tobin, M. McKeever, J. Blackledge, M. Whittington, and B. Duncan, "UK Financial Institutions Stand to Lose Billions in GDPR Fines: How can They Mitigate This?" in The British Accounting and Finance Association: Scottish Area Group Annual Conference, BAFA, Ed., Aberdeen, 2017, p. 6.

[66] EU, "Unleashing the Potential of Cloud Computing in Europe," 2012. [Online]. Available: http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=SWD:2012:0271:FIN:EN:PDF

[67] ENISA, "Article 4 Technical Report," ENISA, Tech. Rep., 2011.

[68] ENISA, "Cloud Risk," ENISA, Tech. Rep., 2009. [Online]. Available: https://www.enisa.europa.eu/publications/cloud-computing-risk-assessment [Last accessed: August 2018]

[69] ENISA, "Recommendations on European Data Protection Certification," Tech. Rep., 2017.

[70] Reuters, "European regulators: We're not ready for new privacy law," 2018. [Online]. Available: https://www.reuters.com/article/us-europe-privacy-analysis/european-regulators-were-not-ready-for-new-privacy-law-idUSKBN1I915X [Last accessed: August 2018]

[71] M. Murphy and B. Riley-Smith, "'Embarrassing' leak shows EU falls short of own GDPR data law," 2018. [Online]. Available: https://www.telegraph.co.uk/technology/2018/05/30/embarrassing-leak-shows-eu-falls-short-data-law/ [Last accessed: August 2018]

[72] V. Beckett, "Many businesses' attitudes to GDPR are 'bordering on negligent'," 2017. [Online]. Available: https://www.theglobaltreasurer.com/2017/10/13/many-businesses-attitudes-to-gdpr-are-bordering-on-negligent/%0A [Last accessed: August 2018]

Forensic Recovery and Intrusion Monitoring in the Cloud

George R. S. Weir
Department of Computer and Information Sciences
University of Strathclyde
Glasgow, UK
e-mail: george.weir@strath.ac.uk

Andreas Aßmuth and Nicholas Jäger
University of Applied Sciences
OTH Amberg-Weiden
Germany
e-mail: {a.assmuth,n.jaeger}@oth-aw.de

***Abstract*—As organisations move away from locally hosted computer services toward Cloud platforms, there is a corresponding need to ensure the digital forensic integrity of such instances. This need is largely motivated by the locus of responsibility and also by the associated risk of legal sanction and financial penalty. Effective monitoring of activity and events is an essential aspect of such forensic readiness. A major concern is the risk that monitoring systems may themselves be targeted and affected by intruders, thereby nullifying the prospective benefits of such internal software surveillance facilities. In this paper, we outline an approach to intrusion monitoring that aims to ensure the credibility of log data and provide a means of data sharing that supports log reconstruction in the event that one or more logging systems is maliciously impaired. In addition, we identify and describe the multi-level interpretation problem as an inherent challenge to managing forensic recovery in the Cloud.**

Keywords-Cloud security; forensic readiness; intrusion monitoring; multi-level interpretation; secure data retention.

I. INTRODUCTION

In the virtual world of interactive software systems, as in the physical world, we often aim to observe and detect behaviour and events that may represent risks or threaten damage to the environment or those within that environment. The primary purpose of such surveillance is to determine the cause and likely consequences of such crucial events. In the event of a security incident, we want to record data that may later have evidential value, shed light on the nature of the occurrence, its context (including significant precursors) and its consequences. Capturing such data in a covert manner aims to reduce the likelihood that the recording facility will be detected and thereby, minimise the prospect that the data collection will be deliberately impaired and the telling data subverted.

While surveillance affords no immediate defence against security breaches, it does illustrate the desirability of establishing auditable data in order that light may later be shed on unauthorised or anomalous events that initially have gone undetected by relevant human agency. With varying degrees of transparency, the logging features in computer operating systems, individual computer applications, network operations and Cloud environments go some way toward addressing this requirement by recording data that may subsequently be consulted, in a process of digital forensics, as evidence of past events.

Although considerable efforts are directed in computer security toward protection and prevention of illicit access and system misuse, digital forensic readiness is increasingly recognised as a necessary measure toward recovery, understanding vulnerabilities and pursuit of those responsible for cyber-misdeeds. In this context, the present paper details the complex problem of managing Cloud forensic recovery [1] and affords a proposed response through application of techniques to bolster digital forensic readiness in the Cloud [2].

In the following, Section II reviews the characteristics of Cloud services and the facilities available to the customer. Section III characterises the attack context, with reference to likely intruder behaviour. In Section IV, we consider the context of Cloud security, with associated network security issues and Cloud security risks addressed in Sections V and VI, respectively.

In Section VII we elaborate upon the role of monitoring as a basis for forensic readiness in Cloud Services, with specific attention to the variety of strategies that may be employed. The effectiveness of such mechanisms for event reconstruction and on-going resilience, is a key consideration. Section VIII presents our proposed monitoring approach that we believe contributes toward a solution to the forensic readiness problem in the Cloud setting.

This is followed by Section IX on Cloud forensic readiness, in which we introduce the multi-level interpretation problem. The paper ends with conclusions in Section X.

II. CLOUD SERVICES

In this section, we briefly review the characteristics of Cloud Services and highlight the security concerns associated with different use contexts.

For many users and organisations, their primary engagement with Cloud computing is remote data storage. To this end, most major online Cloud service providers offer such facilities. Offerings in this area include iCloud for Apple users, as a supplement to local storage capacity and emergency backup for system configuration. Similar service offerings include Google Drive, Microsoft OneDrive and Amazon Drive.

For instance, Dropbox offers a familiar model whereby users may register for a free account with limited storage capacity and a pay option for extended storage capacity and further features. The appeal and benefits from such services are clear from the proliferation of such offerings, as

underlined by the fact that many home broadband contracts include a measure of Cloud storage as standard. Home broadband users often rely on remote storage and backup facilities and may be unaware that Cloud services are the basis for such operation.

Although consumers have been quick to adopt Cloud-based services, there is some concern with security issues that may arise in the Cloud setting [3]-[6], with particular concern for the availability and privacy of data [7].

As a basis for understanding Cloud Services, a taxonomy has been developed by the US National Institute of Standards and Technology (NIST) [8]. Three typical service models are described:

- Software as a Service (SaaS);
- Platform as a Service (PaaS); and
- Infrastructure as a Service (IaaS).

In Software as a Service, the customer is given access to applications running on the service provider's Cloud infrastructure, usually through a variety of client devices and software interfaces. In this arrangement, the customer has no control over the underlying Cloud infrastructure (op. cit., p.2) and this level of service extends from simple file storage, through hosted Web sites and database management to specific Web services, including RESTful applications [9] and use of 'containers' [10].

In Platform as a Service, the customer can deploy their own applications on to the provider's Cloud infrastructure and customer control extends to configuration and management of these Cloud-hosted applications. As before, the customer has no facility to control any other aspects of the underlying Cloud infrastructure [8, p.2].

In Infrastructure as a Service, the customer has more scope for software deployment to the Cloud infrastructure, extending to 'arbitrary software, which can include operating systems and applications' (op. cit.). In this arrangement, the customer's control is still limited to the deployed software applications, including operating systems (e.g., virtual machines) and associated networking features (such as software firewalls) [8, p3].

These three service models characterise typical Cloud Service Provider (CSP) offerings with the increasing levels of access and software capability that are reflected in increasing cost levels. In each of these contexts, management and control of the Cloud infrastructure resides with the CSP, who must be relied upon to manage most security aspects that may impinge upon services purchased by the customer.

Cloud services afford an extensive range of applications and software facilities and many mission-critical services are moving to Cloud as a means of limiting security concerns and assuring greater resilience. Since Cloud services are virtual, system recovery or replacement can be quick, reliable and low-cost [cf. 11]. Cloud-based outsourcing of software applications is recognised as commercially attractive for factors, such as:

- Cost (reduction in local expertise and local infrastructure);
- Reliability (service-level agreements can assure availability);
- Resilience (speedy recovery in the event of data or service loss);
- Technical extensibility (support for multiple instances of applications with increasing availability of service to meet growing demand).

We may broadly differentiate two end-user contexts of Cloud usage. In the first case, the customer employs the Cloud service as a data storage facility. (This is a specific instance of the Software as a Service.) Here, security for the customer is limited to concerns of authorised access, continuity of service and data maintenance. In the second case, the end-user employs the Cloud service as a means of computation. This broadly covers all other Cloud interaction. Here, security for the customer extends to all traditional aspects, including data protection, access authentication, service misappropriation and service availability. While some of these issues may lie within the control of the consumer, the CSP has ultimate management of the infrastructure that affords all of the higher-level service provision.

The security risks associated with these service levels in Cloud provision are elaborated further in Section VII, below.

The extent to which the CSP can reliably manage the security and associated integrity of provided services, depends ultimately upon the availability of techniques for detecting and recording the details of any illicit operations that take place within the Cloud service context. Without recourse to such facilities, the CSP cannot be counted upon to maintain consumer services in a satisfactory fashion since there is lack of assurance that such services have not been infiltrated, impaired or subverted. In addition, ability for the CSP to restore services to pre-compromise level depends largely upon the CSP's facility to identify any delta between pre- and post-intrusion services. Inevitably, this leads back to the issue of digital forensic readiness as applied to the Cloud context.

III. THE ATTACK CONTEXT

Successful cyber-attacks can be construed as having three phases. The first is reconnaissance and information gathering, followed by infiltration and escalation and, finally, exfiltration, assault and obfuscation.

In the first phase, the adversary gathers any information needed to gain access to the system, e.g., open ports, versions of operating systems and software services, security measures (such as firewalls, IDS, etc.) [12]. Using this information, the adversary gains access to the system in the second phase [13].

The process of gaining access might consist of several steps, for example, if the adversary has to comprise another system first, in order to get into the actual target. In this process, the adversary also tries to escalate available privileges in order to gain super-user access to the system.

In the third phase, the adversary extracts any information from the system that might prove to be useful [13]. If the goal of the attack is stealing confidential data, such as user accounts, passwords or credit card information, this data is extracted by the adversary and possibly sold to third parties.

If the cyber-attack has another goal, e.g., sabotage, the adversary extracts the data needed to launch the actual assault, often triggered by a certain date or specific event. In any case, the adversary can be expected to perform whatever action is required to cover their tracks. Among other actions, they may install a rootkit that exchanges current files and services within the system with modified versions of these particular files and services. Such system modifications may extend to altering process information, e.g., a program to list all running processes on the system may be modified to list all running processes except for the processes run by the adversary. Additionally, the adversary may target existing log files that might contain traces of the intrusion.

Such strategies are reflected in many network-based intrusions since, in many instances, network vulnerability is predicated upon known weaknesses in networked hosts.

IV. NETWORK SECURITY RISKS

In non-Cloud systems, the principal ingredients in management responses to security take three general forms:

- System hardening
- Software defences
- Data backup

Firstly, system hardening is an attempt to render known threats ineffective. This includes 'conventional' measures that reduce vulnerability, such as authentication, identity management and access control [14], as well as acting to disable unnecessary services, applying regular software updates (patches) and gauging of the relevance and associated risks from newly published exploits [15]. Modern work-s have also been adapted to meet known cyber threats. Counter measures, like address space randomisation, mandatory access control or maybe sandboxing, are state of the art. In addition, advanced users might even build their own operating system and use selected kernel parameters to further harden their system. The second variety of response to address security issues is the application of software defences. This ranges from antivirus provision to firewalls and may also include some variety of intrusion detection, usually rule-based [16] or anomaly-based [17].

Any computing system may be described by a simple layer-based model. Obviously, security on any higher layer strongly depends on access control mechanisms of lower layers. Even if users or service providers only aim for access control on a higher level to secure their application, these access control mechanisms in practice are more complex than those on lower layers. In addition, vulnerabilities or inadequate configuration on lower levels may lead to bypassing security measures on higher layers. Therefore, appropriate countermeasures are necessary on all layers.

A third security measure is the provision of regular data backup, as a means of ensuring that any system failure or intrusion does not result in irretrievable data loss.

V. CLOUD SECURITY RISKS

Perhaps unsurprisingly, Cloud configurations are subject to levels of security risk that go beyond those affecting conventional networked computer systems. In consequence, the security measures outlined above may not be sufficient in the Cloud setting. In elaborating this claim, the Cloud issues are best illustrated with reference to the differing Cloud service offerings mentioned above [8].

These models for Cloud service provision are helpfully elucidated by Gibson et al. [18], as follows:

- "IaaS provides users with a web-based service that can be used to create, destroy and manage virtual machines and storage. It can be used to meter the use of resources over a period of time, which in turn, can be billed back to users at a negotiated rate. It alleviates the users of the responsibility of managing the physical and virtualized infrastructure, while still retaining control over the operating system, configuration and software running on the virtual machines" [op. cit., p. 199].
- "Platform-as-a-Service providers offer access to APIs, programming languages and development middleware which allows subscribers to develop custom applications without installing or configuring the development environment" [op. cit., p. 200].
- "Software-as-a-Service gives subscribed or pay-per-use users access to software or services that reside in the Cloud and not on the user's device" [op. cit., p. 202].

Our earlier noted approaches to system security are equally applicable to Cloud-based systems. With an eye specifically on Cloud security, we can consider how each of these service offerings may be at risk and what precautions may be anticipated in response to these risks.

1. Infrastructure as a Service

This kind of service seems most prone to the types of exploit that one would expect with conventional networked computers, principally, because, in most cases, such virtual machines will be presented to the Internet as networked hosts. Here, the customer is deploying a virtual machine with associated operating system and on-board software applications. This raises the prospect of vulnerabilities at network level, as well as application level issues, e.g., with Web systems and Database servers, Cross-Site Scripting (XSS) or SQL injections. Denial of service attacks are also a legitimate concern, especially since this kind of attack can achieve enormous bandwidths by using IoT devices for their purpose [19]. For these reasons, *system hardening* (especially, defending against known vulnerabilities) and *software defences* are appropriate for IaaS, including precautions such as anti-malware, firewalls and Intrusion Detection Systems. Provision of these features may be the responsibility of the Cloud Service Provider (CSP), who determines what OS and defensive capabilities are made available. In some settings, the

customer may be in a position to bolster the native defences on the virtual system provided by the CSP.

In similar vein, *data backup* is likely to be required by the IaaS customer. Indeed, the protection of customer data may jointly be the concern of the customer and the CSP. The former may enable off-Cloud backup, to avoid a single source of failure. While the CSP may also offer data backup to a separate Cloud data storage facility.

Despite reasonable expectation of such measures, there are indications that Cloud software infrastructure components are not always adequately secured from known vulnerabilities at the virtual machine level [20].

2. Platform as a Service

Computing facilities afforded to the customer of PaaS, are limited to the development of specific middleware or functional components. These services employ technologies such as Docker [21], Containers [22], DevOps [23] and AWS Lambda [24], in order to host customer-defined remote functionality. From a Cloud customer perspective, *system hardening* seems to be irrelevant in this context in relation to the host operating system. On the other hand, any code developed for use on the Cloud platform must be protected from illicit operations, e.g., process hijacking, output redirection or the elevation of privileges.

Software defences of the variety outlined above seem less relevant to the PaaS context since the operations supported by the middleware are limited to specific data processing and do not afford full operating system access or modification. The primary concern should be the operational effectiveness and resilience of the customer-defined operations. Clearly, such services may also be impaired through illicit access, e.g., stealing authentication details in order to alter code on the host system. Managing this area of concern lies primarily in the hands of the Cloud customer, with the assumption that the CSP will prevent unauthorised access to customer account details.

3. Software as a Service

SaaS provides the Cloud customer with remote access to third-party data processing facilities via micro-services [25] or RESTful services [26]. Aside from network level attacks, such services should be protected from most other security concerns by having the host system hardened and equipped with suitable software defences. From the customer perspective, so long as their remote Cloud services operate effectively, without interruption or data loss, there would seem to be little cause for concern. Of course, the risk of aberrant customer-side behaviour may arise through social engineering exploits or disgruntled employee actions.

This summary of security concerns affecting the three varieties of service has treated each Cloud model as an isolated networked computing facility. In reality, since the essence of Cloud provision is the virtualisation of services, our overview lacks one further important consideration, i.e., the possibility of service impairment as a result of activity at adjacent, upper or lower levels of the Cloud implementation.

Clearly, any security aspects that affect the operational resilience of the underlying Cloud infrastructure is of direct concern to the CSP and can have a knock-on effect upon customer services. The underlying Cloud technology, i.e., the hardware and software configurations that provision our three Cloud models, may be subject to attack or deliberate manipulation in a fashion that impinges detrimentally upon the Cloud services supported by that particular hardware and software ensemble. This may be construed as a service attack 'from below'. The scope for such attacks are precisely the characteristic exploits that may affect any networked host (listed earlier).

Attacks 'from the side' are a growing concern in Cloud security. 'Side channel attacks', originate with co-hosted customers who manipulate the behaviour of their virtual system to influence the behaviour of the host system and thereby affect co-hosted customers. Several studies suggest that such 'co-tenancy', an essential feature of IaaS and PaaS, carries dangers. Thus, "Physical co-residency with other tenants poses a particular risk" [27], such as "cache-based side-channel attacks" [28] and "*resource-freeing attacks* (RFAs)" in which "the goal is to modify the workload of a victim VM in a way that frees up resources for the attacker's VM" [29]. Most worrying are contexts where one customer's 'malicious' virtual machine seeks to extract information from another customer's virtual machine on the same Cloud platform [30]. Such risks to Cloud facilities are fundamental to their service provision.

A final attack vector that threatens some Cloud systems is 'from above'. In this case, poorly implemented virtual systems may afford scope for customers to 'break free' of their virtual system and access or directly affect the underlying operating system or middleware/hypervisor. Clearly, it must be ensured that there is no information leakage from virtual machines and that attackers or malicious customers are not capable of breaking out of the virtual machine and gaining access to the host OS or the virtual machines of other customers [31].

The characteristics of these Cloud service offerings with associated security measures and the likely risk conditions are captured in Table I. The prospect of action from one Cloud user affecting another is described as intra-platform interference.

VI. DIGITAL FORENSIC READINESS

Indications are that the number of cases of network intrusion and data breach is on the rise: "there is a massive increase in the records being compromised by external hacking – from roughly 49 million records in 2013 to 121 million and counting in 2015" [32].

One positive effect of this growth in unauthorized data access is the raised awareness of digital forensics (DF) and a marked change in its perception from a solely post-event reactive investigative tool to a pro-active policy to establish intelligence capabilities in advance of any incidents. This change in role reflects the concept of digital forensic readiness. Thus, "Pro-active DF management must ensure that all business processes are structured in such a way that essential data and evidence will be retained to ensure successful DF investigations, should an incident occur" [33, p.18].

One might define digital forensic readiness as 'having facilities in place to ensure the comprehensive capture and retention of all system event and user activity data that would be required post-incident in order to determine the precise nature of any data-loss, system modification or system impairment that results from intrusion, system misuse or system failure".

Naturally, this concept of digital forensic readiness is equally applicable to Cloud systems and novel techniques have been proposed to facilitate the data collection that this entails [34]. Yet, the Cloud context introduces particular problems with respect to forensic readiness.

Table I. Summary of features, security measures and risks

Service model	**Main features**	**Security Measures**	**Risks**
Infrastructure (IaaS)	Virtual machines, Operating systems, Storage, Software applications	System hardening, Software defences, Data backup	Social engineering, Intrusion, Malware, Denial of service, Elevation of privileges, *Intra-platform interference*
Platform (PaaS)	APIs, Programming languages, Development middleware, (Containers, Dockers, AWS Lambda, DevOps)	System hardening, Software defences	Social engineering, Elevation of privileges, *Intra-platform interference, Information leakage*
Software (SaaS)	Remote applications, Micro-services, RESTful services	System hardening, Software defences	Social engineering, *Intra-platform interference*

VII. MONITORING STRATEGIES

As previously noted, digital forensic readiness requires the monitoring and recording of events and activity that may impinge upon the integrity of the host system. Much of this capability is provided natively by the local system, using standardly available operating system logging, perhaps with additional active security monitoring, such as dynamic log analysis [35] or key file signature monitoring [36].

The situation for Cloud-based services reflects in many respects the context of a networked host. Where a customer employs Cloud purely as a storage medium, minimum security requirements will seek to ensure authenticated access and secure data backup. In turn, the monitoring requirements associated with this service must capture details of user logins (including source IP, username and success or failure of login attempts). Additionally, any file operations that change the status of data stored under the account of that customer must also be recorded. In the event of unauthorised access (e.g., stolen user credentials), such default monitoring may offer little protection, aside from identifying the identity of the stolen credentials and recourse to subsequent backup data recovery. Such monitoring is essentially operating system-based, albeit that in the Cloud setting, this OS may be virtual.

This context of Cloud usage faces the same challenges in monitoring and security that confront any networked host, with the added complication that a Cloud-based virtual host may face added vulnerability via its hosting virtualiser [37]. Furthermore, Cloud services are often configured to provide new virtual OS instances automatically to satisfy demand and in turn, shut these down when demand falls. A side-effect of such service cycling is that system logs are lost to the customer and subsequent digital forensic analysis may be unavailable.

In the 'traditional' network setting, numerous techniques have been devised to afford post-event insight on system failures and unwelcome exploits. In all major operating system contexts, whether virtualised, Cloud-based or native, system logging affords the baseline for generating auditable records of system, network and user activity. Such system level monitoring is well understood and in the event of intrusion is likely to be a primary target in order to compromise the record and eliminate traces of illicit activity.

For networked hosts and, by extension, as a monitoring strategy for local area networks, a wide-variety of Intrusion Detection Systems (IDS) have been developed and deployed with a view to rapid determination of malicious activity. These techniques may be rule-based [e.g., 38]. In most cases, the IDS monitors and cross-correlates system-generated logs in order to identify anomalous event sequences. Many approaches to anomaly-based intrusion detection have been reported [39]-[44]. Inevitably, such systems may themselves become targets in order to inhibit their detection capability and maintain a 'zero-footprint' on the part of the intruder [45].

In a Cloud context, each node is using its own logging daemon or agent to log important events. But in comparison to a single computer, the log information might be essential and therefore relevant for the whole Cloud infrastructure. For that reason, Cloud infrastructures use a centralised log server that receives the log information of all attached nodes. The task of this log server is not only the recording of log files of all nodes but also to monitor the Cloud infrastructure. In case of a cyber-attack, the log server ideally detects the attack (maybe assisted by an intrusion detection system) and starts countermeasures. This exposed role of the log server makes

it a very attractive target for cyber-attacks itself, or, as described above, means that an adversary has to deal with the log server in phase 2. Since the hardware of such a log server might also break down even without any cyber-attack, in practice more than one log server is used at the same time to provide redundancy.

A practical solution might consist of two log servers in "active-active-mode" which means that both are operating at the same time, but in case of one system failure, the other takes over for the whole Cloud infrastructure. The operation of these two log servers might be supervised by a third server which in case of failure or attack sends an alarm to the administrator. Unfortunately, the problem stays more or less the same: this third monitoring server is a single point of failure and is therefore attractive as a target for any adversary attacking the Cloud infrastructure. If an adversary manages to take out the monitoring server and to tamper with the log information on at least one of the two log servers, the Cloud provider might not be capable of determining which log files are correct and which are manipulated.

Any logging service that is introduced in addition to the traditional daemons or agents has to meet several constraints, including the following:

1. the new logging service must not cause too much additional load, either on the nodes (concerning computation) or on the network (concerning network traffic) and;
2. the computation of additional security measures in order to provide authenticity and integrity must be efficiently feasible.

VIII. Example Monitoring Approach

Message Authentication Codes (MACs) as described in almost any textbook about cryptography can readily be used to address this monitoring dilemma. MACs can be constructed using cryptographic hash functions or using block ciphers, for instance. Either construction ensures efficient computation of the MACs under a secret key. MACs are used to provide authenticity and integrity; therefore, they meet both conditions.

A solution that we propose starts with a secure boot process for each node of the Cloud infrastructure. During boot, the common log daemon or agent is started and it starts recording events in various log files. We suggest to compute a MAC for each event and to store these additional bits with the plaintext message of the event in the log file. We assume that the plaintext message also contains a time stamp. For the next event to be recorded in a log file, the plaintext of the event is concatenated with the previous MAC before computing the MAC for this event. This leads to a MAC chain which can be checked for each step using the plaintext and MAC of the previous event - but only if the secret key is known. Since the adversary does not know the secret key, he is not capable of computing valid MACs and therefore not capable of tampering with the MAC chain in order to hide his tracks.

The use of Message Authentication Codes is only the first step towards a solution to the problem. An adversary could simply delete or deliberately falsify all log files (including the MACs). This would probably make it impossible to reconstruct the steps of the cyber-attack in a post-hack analysis.

In order to deal with this issue and to make use of the benefits of a Cloud infrastructure, we propose the additional step of using secret sharing techniques - or so-called threshold schemes - as published by Adi Shamir in 1979 [46].

The idea is to divide some data D into n pieces $D_1, \dots, D_n$ in such a way that:

(a) D can be reconstructed easily of any $k < n$ pieces D_i

(b) the knowledge of only $k-1$ or even fewer pieces D_i leaves the data completely undetermined.

Shamir named such a scheme a "(k, n) threshold scheme". He points out that by using such a (k, n) threshold scheme with $n = 2k - 1$, it is necessary to have at least $k = \left\lceil \frac{n+1}{2} \right\rceil$ parts D_i to reconstruct D. A lesser number of $\left\lfloor \frac{n}{2} \right\rfloor = k - 1$ parts makes the reconstruction impossible.

Shamir introduced a (k, n) threshold scheme based upon polynomial interpolation. The data D can be interpreted as a natural number and p is a prime number with $D < p$. All of the following computations are made in the prime field $\mathrm{GF}(p)$. Given k points in the 2-dimensional plane, $(x_1, y_1), \dots, (x_k, y_k)$ with distinct coordinates x_i, there is one and only one polynomial q of degree $k - 1$ such that $q(x_i) = y_i$ for all $i = 1, \dots, k$. At first, the coefficients $a_1, \dots, a_{k-1}$ are chosen at random and $a_0 = D$, which leads to the polynomial

$$q(x) = a_0 + a_1 x + a_2 x^2 + \cdots + a_{k-1} x^{k-1}.$$

The n different pieces of D are computed as $D_1 = q(1)$, $D_i = q(i), \dots, D_n = q(n)$. Provided that their identifying indices are known, any subset of k elements D_i can be used to compute the coefficients a_i of the polynomial q which allow the computation of the data $D = q(0)$. From any subset of less or equal $k - 1$ pieces D_i, neither the coefficients a_i nor the data D can be calculated. (For further details, we direct the reader to the original paper [46].)

In our proposed solution to the problem of providing additional forensic information for post-hack analysis, D is the data to be written in a log file: the plaintext message of the event, n randomly chosen nodes of the Cloud infrastructure and the corresponding MAC, computed from the concatenation of the event message, the previous MAC and the addresses of these n nodes. The n pieces D_i that are derived from D as stated before and D is sent to the traditional centralised log server. The n pieces D_i are additionally sent to the n nodes which store this information. For the next event, we repeat this procedure but choose n (possibly) different nodes.

In case of a cyber-attack and if a post-hack analysis is necessary, at first all pieces of logging information are gathered from all nodes. Using the time stamps and the MAC chains, the order of the logged events can be reconstructed. The decentralised stored pieces of logging information are put together to reconstruct D from any k of the n parts. This

means, even if an adversary succeeds in manipulating some of the nodes and the centralised logging system, the events can be reconstructed. Finally, the integrity and authenticity of these events can be checked using the MAC chain.

The proposed approach may identify and retain information on an intruder's actions that result in stolen, modified or deleted data. This is a feature with growing importance, as legislative demands on data protection increase. For instance, the EU General Data Protection Regulation that is due to come into force in May 2018, will require companies to notify all breaches within 72 hours of occurrence, with a potential penalty of up to 4% of global turnover based on the previous year's accounts.

Note that this solution is not proposed as a general basis for monitoring the Cloud infrastructure. Rather, its purpose is to provide secure logging information for a post-hack analysis by distributing their parts randomly over all nodes. Thereby, reliable system monitoring can be established by means of multiple log servers, with the added assurance of Message Authentication Codes.

Now that we have a workable means of addressing the log data collection, robust storage and recovery of such data, we move to consider significant residual issues with digital forensic recovery in the Cloud context.

IX. Issues in Cloud Forensic recovery

Forensic readiness in the Cloud is complicated by the variety of contexts in which Cloud services are deployed and the diversity of software settings in which security risks may arise. Forensic readiness must accommodate these complexities and, in turn, this suggests that a single infrastructure-based digital forensic readiness solution may be infeasible.

The primary reason for concern is the need to capture relevant data on system operation at the various operational levels of the Cloud system and any potential interaction across these levels. This means capturing program logs, system logs and user activity logs. In any end-customer Cloud facility, the data protected may not extend beyond any currently live information and data held in associated database systems. The ready recycle capability of Cloud services also has implications for the persistence of digital forensic evidence. An intrusion that steals data from a virtual machine and then seeks to reset that machine may well succeed in destroying evidence of the intrusion, thereby removing any forensic traceability on the nature and quantity of stolen data.

Neither is it sufficient to provide each distinct operational layer of Cloud systems with its own comprehensive forensic readiness. At best, this condition will allow for forensic data recovery for that operational layer. But there is no one-size-fits-all solution that can capture all state, interaction and performance data such as would ensure full Cloud forensic recovery. In fact, this insight reveals a fundamental problem that may impact upon Cloud forensic readiness.

There are parallels here with issues in distributed systems and software architecture. Thus, "distributed software systems are harder to debug than centralized systems due to the increased complexity and truly concurrent activity that is possible in these systems" [47, p. 255]. Regardless of whether the Cloud setting is truly distributed in its realisation, its interconnected software functional layers represent a unique challenge when attempting to interpret the relationship between events or changes actioned at one functional level and the operational impact of such changes on other functional aspects of the services afforded by that Cloud.

When considering Cloud systems, from the perspective of software architecture there may be an assumption of 'a component- and message-based architectural style' in which there is 'a principle of limited visibility or *substrate independence*: a component within the hierarchy can only be aware of components "above" it and is completely unaware of components which reside "beneath" it' [48, p.825].

This multi-level interpretation problem is complicated by the fact that events considered anomalous at one level of service offering may arise through actions considered legitimate at a 'lower' level of software implementation. From the digital forensic readiness perspective, this underlines the requirement to go beyond capture of significant events across the Cloud service software and functional levels, since significance is an aspect that may cross the boundaries between such layers in the system as a whole. A few hypothetical examples may clarify this issue.

In our first example, a CSP customer may contract access to specific functional components (e.g., a Web service). The operational characteristics of the service are under the control of the CSP and not the customer. An authorised employee of the CSP may modify the algorithmic process and thereby affect the outcome of any service use by the customer. While a change in operational behaviour of the service (i.e., an anomaly) may eventually be detected by the customer, there may be no anomalous activity evident at the level of CSP employee activity. The focus of subsequent forensic investigation may light initially on the nature of customer activity, since this is where the anomaly is apparent, but proper understanding of the issue requires that events across different functional levels of the Cloud system be apprehended.

In our second example, an employee of the CSP illicitly establishes a clone of the live customer system, with all data records in the customer system continuously duplicated, updated and available to the CSP employee. Here, data records are being accessed without authorisation and this fact is both unknown and unavailable to the Cloud customer. Insight from the operational level of the CSP would be required in order to expose this situation. Yet, the Cloud customer may come under scrutiny or be subject to litigation if critical customer data is made available on the Dark Web.

An informative view on this issue of informational levels may be borrowed from Granular Computing [49], which aims to develop computational models of complex systems, such as human intelligence. A key characteristic of this work is that it 'stresses multiple views and multiple levels of understanding in each view' [op. cit., p.85]. Here, the emphasis is upon 'holistic, unified views, in contrast to isolated, fragmented views. To achieve this, we need to consider multiple hierarchies and multiple levels in each hierarchy' [op. cit., p.88].

Our proposal for adequate Cloud forensic readiness has two components (detailed above). Firstly, there is a requirement for data capture. This is the obvious need to record any data at each layer of Cloud facility that may have a role to play in subsequent digital forensic analysis. Secondly, the captured data must be stored off the system being monitored in a manner that both ensures the integrity of the logging and minimises the likelihood that the stored data can be compromised, either as a result of hostile action or 'friendly fire'.

Our requirement for secure and resilient log storage can build upon default system logging that will be present within the Cloud implementation but this must be supplemented to achieve log reliability.

Instead of using centralised log servers, which of course are attractive targets and easy to spot for attackers, we propose a different approach. In order to prevent adversaries from manipulating log files to hide their tracks, we use chained Message Authentication Codes (MACs) for each entry to the log file on each node. If state-of-the-art MACs are used, this makes it impossible to delete or manipulate text in the log files. Next, each node uses secret sharing techniques, such as that proposed by Adi Shamir [46], to divide the log file into parts. These parts are then sent to random other nodes which store these log data. Even if an adversary succeeds in taking over some of the nodes, he will need a certain number of these fragments to reconstruct the log data. But since for each log entry different nodes are chosen randomly as stated before, the attacker effectively needs to control the whole Cloud ecosystem to stay hidden.

X. CONCLUSIONS

Recognising the importance of securing log data as a basis for digital forensic reconstruction in the event of system intrusion, a multiple server solution combined with Message Authentication Codes affords a mechanism that allows for safe deposit and reconstruction of monitor data. This can operate in a Cloud setting in which each logging node is a virtual server.

An important benefit from this distributed solution is that digital forensic reconstructions are possible for virtual machines that are 'cycled', since their native OS logs can be maintained in a recoverable and verifiable form beyond the OS of those machines. This provides the safeguard of digital forensic readiness for Cloud customers in the event that an intruder accesses private data on the Cloud service and causes that system to cycle as an attempt to delete all traces of illicit data access.

The possibility, however slight, that an intruder may gain access to and potentially compromise all peers in this configuration, can be mitigated by also allowing log data to transfer 'upwards' to one or more 'superior' systems (e.g., the parent operating systems in which the peer log servers are virtualised).

As organisations move increasingly away from locally hosted computer services toward Cloud platforms, there is a corresponding need to ensure the forensic integrity of such instances. The primary reasons for concern are the locus of responsibility and the associated risk of legal sanction and financial penalty. In the first place, while Cloud service providers (CSPs) are responsible for the availability and robustness of their commercial offerings, they will not be responsible for the management of such services by their customers, nor for the data security associated with customer-level use of the Cloud services. Responsibility for these aspects resides with the CSP's customers, whose data processing and data management are built upon the purchased Cloud services. In the second place, legislative demands on data protection, such as the EU General Data Protection Regulation [50], requires companies to notify all breaches within 72 hours of discovery or face significant financial penalty.

These concerns can be addressed and the business risk mitigated through development of forensic readiness in customer-level Cloud systems (as described above). We have argued that this requires a range of logging and data capture facilities across the Cloud system software infrastructure that maintain the possibility of tracking activity at different levels of software abstraction (the multi-level interpretation problem). Our second proposition is that such digital forensic readiness must be combined with techniques to ensure that logged data is incorruptible and robust. We have previously proposed techniques for intrusion monitoring that ensure log data credibility and provide robust decentralised log storage and recovery for post-hack scenarios.

To achieve adequate data capture, we require 'state information' and data management across differing levels of any Cloud service, from the lowest software level up to the most abstracted 'user facing' software component. On their own, such records will not be sufficient to fully capture the potential interplay of differing software levels. For this purpose, subsequent digital forensic analytics will be required in order to establish a multi-dimensional representation of event chronology. This means that timestamps from events and data captured at different software levels of abstraction will need to be correlated in order to determine how events across the Cloud system are related.

Cloud service provision has a requirement for secure and robust monitoring with access to multiple levels of such monitoring data. If we are able to supplement our robust monitoring and logging approach with appropriate levels of Cloud operational information (e.g., as a feature of Cloud Service Level Agreements), this may in turn facilitate a solution to the multi-level interpretation problem and take us all the way to effective digital forensic readiness. Thereby, we may achieve a Cloud facility that is capable of successful recovery from accidents and incidents, to afford effective management of digital forensic recovery.

REFERENCES

[1] G. R. S. Weir, A. Aßmuth and N. Jäger, "Managing forensic recovery in the Cloud", In: Proc. *Cloud Computing 2018*, *The Ninth International Conference on Cloud Computing*, GRIDs and Virtualization, IARIA, Barcelona, Spain, 2018.

[2] G. R. S. Weir and A. Aßmuth, "Strategies for Intrusion Monitoring in Cloud Services", In Proc. *Cloud Computing 2017, The Eighth International Conference on Cloud*

Computing, GRIDs and Virtualization, IARIA, Athens, Greece, 2017.

[3] M. Nanavati, P. Colp, B. Aiello and A. Warfield, "Cloud security: A gathering storm", *Communications of the ACM*, *57*(5), pp. 70-79, 2014.

[4] S. S. Tirumala, H. Sathu and V. Naidu, "Analysis and prevention of account hijacking based incidents in Cloud environment", In Proc. *International Conference on Information Technology (ICIT),* pp. 124-129, 2015.

[5] M. Zhou, R. Zhang, W. Xie, W. Qian and A. Zhou, "Security and privacy in Cloud computing: A survey", In Proc. *Sixth International Conference on Semantics Knowledge and Grid (SKG),* pp. 105-112, 2010.

[6] Y. Chen, V. Paxson and R. H. Katz, "What's new about Cloud computing security", *University of California, Berkeley Report No. UCB/EECS-2010-5 January*, *20*, 2010.

[7] BBC News, Available: http://www.bbc.co.uk/news/technology-29076899. [Accessed: Dec. 29, 2017].

[8] P. Mell and T. Grance, "The NIST definition of Cloud computing", NIST, 2011. Available from http://faculty.winthrop.edu/domanm/csci411/Handouts/NIST.pdf, [retrieved: February, 2017].

[9] S. A. Shaikh, H. Chivers, P. Nobles, J. A. Clark and H. Chen, "Network reconnaissance", *Network Security*, vol. 11, pp. 12-16, 2008.

[10] L. Richardson and S. Ruby, *RESTful Web Services*. O'Reilly Media, Inc., 2008.

[11] B. Benatallah, Q. Z. Sheng and M. Dumas, "The self-serv environment for web services composition", *IEEE Internet Computing*, vol. 7, no. 1, pp. 40-48, 2003.

[12] B. F. Murphy, Network Penetration Testing and Research, NASA, 2013. Available from https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20140002617.pdf, [retrieved: February, 2017].

[13] J. Andress and S. Winterfeld, *Cyber warfare: techniques, tactics and tools for security practitioners.* Elsevier, 2013.

[14] H. Takabi, J. B. Joshi and G. J. Ahn, "Security and privacy challenges in Cloud computing environments", *IEEE Security & Privacy*, 8(6), pp. 24-31, 2010.

[15] M. Carroll, A. Van Der Merwe and P. Kotze, "Secure Cloud computing: Benefits, risks and controls", In *Proc. Information Security South Africa* (ISSA), 2011, pp. 1-9, 2011.

[16] K. Ilgun, R. A. Kemmerer and P. A. Porras, "State transition analysis: A rule-based intrusion detection approach", *IEEE Transactions on Software Engineering*, vol. 21, no. 3, pp. 181-199. 1995.

[17] P. Garcia-Teodoro, J. Diaz-Verdejo, G. Maciá-Fernández and E. Vázquez, "Anomaly-based network intrusion detection: Techniques, systems and challenges", *Computers and Security*, vol. 28, no. 1, pp. 18-28, 2009.

[18] J. Gibson, R. Rondeau, D. Eveleigh and Q. Tan, "Benefits and challenges of three Cloud computing service models", In *Proc. Fourth International Conference on Computational Aspects of Social Networks (CASoN)*, pp. 198-205, 2012.

[19] H. Sweetland Edwards, "How Web Cams Helped Bring Down the Internet, Briefly", Time Magazine, 25th October 2016. Available: http://time.com/4542600/internet-outage-web-cams-hackers. [Accessed: Dec. 29, 2017].

[20] S. Zhang, X. Zhang and X. Ou, "After we knew it: empirical study and modeling of cost-effectiveness of exploiting prevalent known vulnerabilities across IAAS Cloud", In *Proc. 9th ACM Symposium on Information, Computer and Communications Security,* pp. 317-328, 2014.

[21] S. Dhakate and A. Godbole, "Distributed Cloud monitoring using Docker as next generation container virtualization technology", In Proc. *Annual IEEE India Conference (INDICON),* pp. 1-5, 2015.

[22] C. Pahl and B. Lee, "Containers and clusters for edge Cloud architectures--a technology review. In Proc. *3rd International Conference on Future Internet of Things and Cloud (FiCloud),* pp. 379-386, 2015.

[23] A. Balalaie, A. Heydarnoori and P. Jamshidi, "Microservices architecture enables DevOps: migration to a Cloud-native architecture", *IEEE Software*, *33*(3), pp. 42-52, 2016.

[24] M. Villamizar et al., "Infrastructure cost comparison of running web applications in the Cloud using AWS lambda and monolithic and microservice architectures. In Proc. *16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid),* pp. 179-182, 2016.

[25] D. Namiot and M. Sneps-Sneppe, "On micro-services architecture", *International Journal of Open Information Technologies*, *2*(9), pp. 24-27, 2014.

[26] H. Han et al., "A RESTful approach to the management of Cloud infrastructure", In Proc. *IEEE International Conference on Cloud Computing,. CLOUD'09.,* pp. 139-142, 2009.

[27] Y. Zhang, A. Juels, A. Oprea and M. K. Reiter, "Homealone: Co-residency detection in the Cloud via side-channel analysis", In Proc. *IEEE Symposium on Security and Privacy (SP),* pp. 313-328, 2011.

[28] Y. Zhang, A. Juels, M. K. Reiter and T. Ristenpart, "Cross-tenant side-channel attacks in PaaS Clouds". In Proc. *ACM SIGSAC Conference on Computer and Communications Security,* pp. 990-1003, 2014.

[29] V. Varadarajan, T. Kooburat, T., Farley, T. Ristenpart and M. M. Swift, "Resource-freeing attacks: improve your Cloud performance (at your neighbor's expense)", In Proc. *ACM conference on Computer and communications security*, pp. 281-292, 2012.

[30] Y. Zhang, A. Juels, M. K. Reiter and T. Ristenpart, "Cross-VM side channels and their use to extract private keys", In Proc. *ACM conference on Computer and communications security,* pp. 305-316, 2012.

[31] T. Vateva-Gurova, N. Suri and A. Mendelson, "The Impact of Hypervisor Scheduling on Compromising Virtualized Environments", In Proc. *IEEE International Conference on Computer and Information Technology; Ubiquitous Computing and Communications; Dependable, Autonomic and Secure Computing; Pervasive Intelligence and Computing (CIT/IUCC/DASC/PICOM)*, pp. 1910-1917, 2015.

[32] Security Week, "Data Breaches by the Numbers", Available: http://www.securityweek.com/data-breaches-numbers. [Accessed: Dec. 29, 2017].

[33] C. P. Grobler and C. P. Louwrens, "Digital forensic readiness as a component of information security best practice", In Proc. *IFIP International Information Security Conference*, pp. 13-24, Springer, Boston, MA, 2007.

[34] V. R. Kebande and H. S. Venter, "A Cloud Forensic Readiness Model Using a Botnet as a Service", In Proc. *The International Conference on Digital Security and Forensics (DigitalSec2014)*, pp. 23-32, The Society of Digital Information and Wireless Communication, 2014.

[35] A. Oliner, A. Ganapathi and W. Xu, "Advances and challenges in log analysis", *Communications of the ACM*, vol. 55, no. 2, pp. 55-61, 2012.

[36] G. H. Kim and E. H. Spafford, "The design and implementation of tripwire: A file system integrity checker", Proceedings of the *2nd ACM Conference on Computer and Communications Security*, ACM, pp. 18-29, 1994.

[37] J. S. Reuben, *A survey on virtual machine security*. Helsinki University of Technology, vol. 2, no. 36, 2007.

[38] K. Ilgun, R. A. Kemmerer and P. A. Porras, "State transition analysis: A rule-based intrusion detection approach", *IEEE*

Transactions on Software Engineering, vol. 21, no. 3, pp. 181-199, 1995.

[39] P. Garcia-Teodoro, J. Diaz-Verdejo, G. Maciá-Fernández and E. Vázquez, "Anomaly-based network intrusion detection: Techniques, systems and challenges", *Computers and Security*, vol. 28, no. 1, pp. 18-28, 2009.

[40] C. Chapman, S. Knight and T. Dean, USBcat-Towards an Intrusion Surveillance Toolset, arXiv preprint arXiv:1410.4304, 2014.

[41] X. Wang, D. S. Reeves, S. F. Wu and J. Yuill, "Sleepy watermark tracing: An active network-based intrusion response framework", *Trusted Information*, Springer US, pp. 369-384, 2002.

[42] C. V. Zhou, C. Leckie and S. Karunasekera, "A survey of coordinated attacks and collaborative intrusion detection", *Computers and Security*, vol. 29, no. 1, pp. 124-140, 2010.

[43] A. Patcha and J. M. Park, "An overview of anomaly detection techniques: Existing solutions and latest technological trends", *Computer Networks*, vol. 51, no. 12, pp. 3448-3470, 2007.

[44] H. Sukhwani, V. Sharma and S. Sharma, "A Survey of Anomaly Detection Techniques and Hidden Markov Model", *International Journal of Computer Applications*, vol. 93, no. 18, pp. 26-31, 2014.

[45] G. Tedesco and U. Aickelin, Strategic alert throttling for intrusion detection systems, arXiv preprint, arXiv:0801.4119, 2008.

[46] A. Shamir, "How to share a secret", *Communications of the ACM*, vol. 22, no. 11, pp. 612-613, 1979.

[47] P. C. Bates and J. C. Wileden, "High-level debugging of distributed systems: The behavioral abstraction approach", *Journal of Systems and Software*, *3*(4), pp. 255-264, 1983.

[48] N. Medvidovic, R. N. Taylor and E. J. Whitehead Jr, "Formal modeling of software architectures at multiple levels of abstraction", *ejw*, *714*, pp. 824-837, 1996.

[49] Y. Yao, "Perspectives of granular computing. In Proc. *IEEE International Conference on Granular Computing,* Vol. 1, pp. 85-90, 2005.

[50] Regulation, Protection. "Regulation (EU) 2016/679 of the European Parliament and of the Council." *REGULATION (EU)*(2016): 679, 2016.

GHOST: An Evaluated Competence Developing Game for Cybersecurity Awareness Training

Johannes A. König and Martin R. Wolf
Lab for IT Organization and Management
Aachen University of Applied Sciences
Aachen, Germany
email: koenig@fh-aachen.de, m.wolf@fh-aachen.de

***Abstract*—To train end users how to interact with digital systems is indispensable to ensure a strong computer security. 'Competence Developing Game'-based approaches are particularly suitable for this purpose because of their motivation- and simulation- aspects. In this paper the Competence Developing Game 'GHOST' for cybersecurity awareness trainings and its underlying patterns are described. Accordingly, requirements for an 'Competence Developing Game' based training are discussed. Based on these requirements it is shown how a game can fulfill these requirements. A supplementary game interaction design and a corresponding evaluation study is shown. The combination of training requirements and interaction design is used to create a 'Competence Developing Game' -based training concept. A part of these concept is implemented into a playable prototype that serves around one hour of play respectively training time. This prototype is used to perform an evaluation of the game and training aspects of the awareness training. Thereby, the quality of the game aspect and the effectiveness of the training aspect are shown.**

Keywords-Cybersecurity; Awareness; CDG; Serious Game; tablet game; business simulation; evaluation.

I. INTRODUCTION

The use of digital systems is crucial in modern companies and one effort of digitization is to use these digital systems more efficiently. Through these efforts, more and more analog processes are no longer available. By that, nowadays almost all relevant records are stored in databases or on cloud based file servers. Accordingly, the analog data management will be reduced to minimum, if that has not already happened.

Of course, a well-functioning digital working environment is required to ensure that the data are always available. If data are accessible everywhere and always for employees, then assailants are able to use these infrastructure, too. This issue is getting worse because nowadays, in modern digitalized systems, employees are owners of the keys necessary for data access. Consequently, it is no longer necessary for an assailant to attack the IT-infrastructure (IT = Information technology) or the IT-department. He can focus his attack directly on the data-using persons, e.g., with fishing-mails, social attacks, manipulated flash drives, etc. Despite this issue, this kind of always available data management is indispensable for modern companies.

An "Competence Developing Game"-concept (CDG) to train non-IT employs was presented by König and Wolf [1] in a shorter version of these paper on the ACHI 2018 conference. In that paper, however, it has remained at the concept level, a prototype was not presented. Supplementary to the old paper, in this contribution it is shown how the CDG prototype exactly looks like. That includes all prototype quests with their serious and entertaining aspects. Further, based on the prototype, an empirical study is presented. The study is used to evaluate the serious and the entertainment aspects of the CDG.

Regardless of the chosen approach, it is essential to train non-IT personnel how to avoid cybersecurity risks arising within their daily digitalized work [2]. Already today, employees are often the biggest threat in the cybersecurity chain [3]. To offer an effective cybersecurity awareness training, it is important to establish a continuous training cycle to establish a long term behavior change **(**req. **7** (see Section II)**)**. It should be noted that too many topics in too short time increase the risk to overwhelm the exercisers which is also a reason for a long training cycle. Basically, a successful cybersecurity awareness training has to solve two tasks. On one hand, it has to attract the attention of the participants for a defined time period. On the other hand it has to convey the training content as efficiently as possible. Unfortunately, most of today's trainings solutions show weaknesses in dealing with both aspects. A very suitable solutions to address both aspects is the use of interactive computer-based training methods **(**req. 6 (see Section II)**)** [2]. The use of gaming concepts in serious situations provides the possibility to transfer the motivation of a gaming situation into a serious learning context. In addition, games provide an environment which allows to choose risky or intentional wrong strategies just to figure out what will happen. Generally, there are three major kinds of games with a serious approach: Serious Games, Business Simulation/Games and the approach of Gamification. Further, there are different gradations of, e.g., serious games, which are not consistently defined [4].

However, instead of questioning 'What defines a particular game kind?' König and Wolf suggest focussing on the question 'What characteristics of which game kind are well suited for a specific application' [5]. For this, they provide the umbrella term CDG that encompasses all 'serious' game types (digital and analog):

'A Competence Developing Game (CDG) is a game that has the primary purpose to teach [how to use] knowledge, skills and personal, social and/or methodological abilities, in work or study situations and in professional and personal development of the game player, by retaining the motivation of a gaming situation' [4] (Note: The 'how to use' was accidently missing in the original source).

Accordingly, this paper examines what features a digital CDG must have in order to enable a cybersecurity awareness training for (German) business users. Further, it shows a specific CDG-design, in which these features are adressed. The CDG is called GHOST: Gamified Hacking Offence Simulation-based Training. In addition, a prototype will be introduced that contains a sample of the game ideas. Further, using this prototype an empiric evaluation study will be performed, analyzed and interpreted to prove the game's concepts. In detail, this paper is structured as follows:

In Section II, the target audience is determined in more detail, to understand their preferences and requirements. Section III addresses these requirements to determine a suitable CDG game kind. In Section IV, it is explained, how a game interaction interface design for a huge audience group like, 'business users', could look like. In addition, in Section V, a study that examines game interaction systems is briefly presented. Section VI describes the CDG GHOST which results from all previous considerations. In Section VII a prototype of the GHOST game and a corresponding evaluation study is presented. In addition, in Section VIII the study results are shown and interpreted. Finally, Section IX offers a conclusion and an overview about future work and use.

II. FINDING REQUIREMENTS BY UNDERSTANDING THE AUDIENCE

A study in German enterprises determined that the three most common reasons for employee related trainings are: the development of employee skills, increasing employee motivation and job satisfaction, and strengthening the employee-company relation (req. 1). The study also determined the obstacles that inhibit employee trainings. The identified top-two reasons not to train although there is a need are: no time available to dispense employees (43.8%) and missing internal capacity to organize a training (42.6%) [6]. A second study in German companies identified training costs and also the time issue as main reasons not to train employees. The three most common training methods are learning at the place of work (46%), external courses (28%) and in-house courses (<28%) [7].

In the case of learning at the place of work, the time an employee needs to be dispensed is limited to the actual duration of the training, because there is no traveling time (obstacle: no dispense time available) (req. 2.a.). The absence of traveling time is linked to the absence of traveling costs (obstacle: training costs) (req. 2.c.). By that, the organizational complexity of the training is also reduced, as employees must be covered shorter, and they are more easily accessible in crisis situations, etc. (obstacle: organizational capacity) (req. 2.b.). Accordingly, in the case of a continuous training cycle, as needed for a cybersecurity awareness training and therefore for GHOST, learning at the place of work seems particularly advantageous. These considerations clarify why learning at the place of work is the most popular training method and therefore it should be the method of choice for GHOST (req. 2).

In addition to these employer-focused considerations, the CDG GHOST is after all played by employees. As explained in Section I, more or less every employee who uses digital systems for work reasons should participate in a cybersecurity awareness training. By that, the target audience is broad (req. 3). Since the GHOST-Research-Project is granted by a German ministry (Federal Ministry of Education and Research), the German employee sector was considered in first place. According to a report by the Federal Institute for Vocational Education and Training, the average German trainee is 19.7 years old. The report shows the first grouping called "16-year-olds and younger". The average age of all employees was 43 years in 2016, with a relatively balanced distribution between women (~ 47%) and men (~ 53%) [8]. In summary, it can be stated that the vast majority of the target group is > = 16 years and <67 years old, the average age is 43, and women and men are similarly distributed.

As already mentioned, the use of a CDG as a training instrument has the advantage that the motivation of a game situation can be transferred in a serious context. In order to use this advantage a CDG must entertain players in a fun way while keeping the serious content in focus. This aspect requires a CDG that matches the tastes and abilities of the target audience. But because of the diversified target group, it is nearly impossible to construct a CDG that fulfills the individual game taste of each subject. On the other hand, the development of many games that meet the individual taste of each player would be expensive and it would stand in opposite to the obstacle: 'costs'. Following these remarks, a CDG that addresses a broad audience always represents a compromise in game design.

To find the major common denominator of each CDG-Player the 'Pyramid Assessment Framework for 'Competence Developing Games'' ('PACDG-Framework') was studied with this objective. The PACDG-Framework represents a tool that delivers the capability to analyze different game kinds in a standardized way. To do so, the framework covers, among other things, the entire player perspective of a CDG [5], as it was proposed (also) in the well-known MDA-framework for conventional entertainment games [9]. However, the PACDG-Framework covers the CDG-Player perspective in the three steps: "Experience", "Aftereffect" and "Impact". The last two steps refer to the same idea: A CDG should lead to competence acquisition, where the competences should help to solve at least one real life problem (req. 4). The step "Experience" is all about the player's claim to participate in an emphatic and positive gaming experience. In order to meet this claim, a high, entertainment game equivalent, quality must be delivered (req. 5).

Therefore, a CDG-based training that is accessible for all employees who use digital systems for work reasons should…:

Req. 1. …develop skills, increasing motivation / satisfaction, strengthening the job relation.

Req. 2. …take place at the place of work to reduce

a. time expense and release time,
b. organizational overhead and by that
c. costs.

Req. 3. …be accessible for every target group member.
Req. 4. …help to solve a real life problem.
Req. 5. …be similar in quality to an entertainment game.

Additionally a CDG for a cybersecurity awareness training should…: (see Section I)

Req. 6. …use interactive computer based training methods.
Req. 7. …occur in a continuous training cycle.

III. GAME TYPE SELECTION

As discussed in Sections I and II, the use of interactive computer-based training methods is suitable for a cybersecurity awareness training. By that, a serious game, a business simulation (supported by a computer based simulation model) or a gamified work environment could be used (fulfill req. 6). Furthermore, it is of course possible to develop a CDG in one of the named kinds with an entertainment game comparable quality (fulfill req. 5).

However, every well designed cybersecurity awareness training will match the requirements 1 and 4, too. It is because the main CDG purpose would be to lead to competence acquisition, where these competence acquisition refers to the ability to perceive possible IT-Security threats (fulfill req. 1). As IT-Security issues are a real life problem, of course, such competences would support to solve a real life problem (fulfill req. 4). Therefore, it can be assumed that a capable development team has the ability to develop a CDG from one of the named game kinds that has the potential to fulfill the requirements 1, 4, 5 and 6.

So, to choose the most suitable CDG game kind it is necessary to determine whether the requirements 2, 3 and 7 can be fulfilled.

''Gamification' is the use of game design elements in non-game contexts" [10]. As a result, for the gamification solution a deeply integration of game elements into the computer environment of the employees would be necessary. Based on such integration, e.g., correct behavior such as scanning a flash drive or locking the screen during a longer period of inactivity could be rewarded with points (fulfill req. 2a-b). This solution would enable a permanent and time neutral training without the need of learning to handle the training instrument (fulfill req. 3 and 7). However, the necessary development effort would be high (game element integration in every used program and operating system) and the privacy protection question would need clarification (not fulfill req. 2c). In addition, the extensive system intervention could have unforeseeable consequences on the IT security of the manipulated operating systems and programs. For these reasons a gamification solutions does not seem suitable for a cybersecurity awareness training.

A closed 'Business Simulation' is characterized by the participants being placed into a well-defined and prepared action situation. A model calculation (the simulation) assesses the decision effects on the game environment. Further the model communicates the success of each action to the players [11]. Since a business simulation is similar to a board game the majority of the employees should not have any problem to handle the game (fulfill req. 3). In addition, many simulation games are turn-based anyway and thus predestined for a long continuous game cycle (fulfill req. 7). The problem here is that even if it is possible to organize multiple business game session at the work (fulfill req. 2a), fixed dates have to be coordinated between different employees plus the necessary setup and dismantling of the business game have to be organized in time (not fulfill req. 2b-c). That means, a business simulation can also not fulfill all requirements.

The third alternative are 'Serious Games'. Serious Games are video games where the primary purpose is not entertainment, enjoyment or fun, which does not mean that Serious Games are not entertaining. They just have another primary purpose, in kind of an ulterior motive [12]. A video game has the advantage of being fully flexible in terms of time. Further no coordination is required between employees nor an organization of the game setup and it can also take place at work (fulfill req. 2a-c) However, it is difficult to realize a continuous training cycle without a turn-based design and such a design is not intended for Serious Games (not fulfill req. 7). But indeed it is the only approach that has the potential to fulfill requirement 2.

At this point, a CDG reveals its strength. The solution is to mix up the game kinds. Serious Games are the only game type that fulfills the requirements 2a-c, but the turn-based design of business simulations supports a continuous game cycle. Accordingly the solution is to develop a Serious Game with Business Simulation (turn-based) game mechanics (see Section VI). Therefore, only the mix out of a Serious Game and a Business Simulation has the potential to fulfill requirements 1 to 7.

Due to this design choice, the biggest problem with meeting the requirements will be requirement 3 in which a CDG is demanded that is playable for every target group member. In requirement 5, the demand for a quality which is similar to an entertainment game is formulated. It needs to be kept in mind that not all members of the target group have experience with video games. It must therefore be ensured that requirement 3 can be met without losing number 5. Therefore, it is necessary to find an interaction-interface for a high quality video game that does not require any video game experience. Section V will introduce a case study that was performed to evaluate how a game interface has to be designed to meet requirement 3 even when the game uses a 3D-Environment to fulfill requirement 5. Section IV explains the game interface development and the case study design.

IV. DEALING WITH THE GAME INTERACTION ISSUE

Germany is on of the largest video game markets in Europe with sales of 2.8 billion euros in 2015. Overall, the video game players are distributed as follows: PC / laptop 18.4 million players, smartphone 17.2 million players, console 15.6 million players, tablet 11.5 million players, handheld 8.3 million players. It should be noted that smartphones and tablets both use gaming apps, which means gaming apps with 23 million players in total have the largest player community

[13]. Accordingly to that information even in the aimed target group the amount of people who have experience with gaming apps should be higher than with other video game mediums.

In addition, it can be stated that touchscreens as used in smartphones and tablets have significantly changed the world of games in a short period of time. Modern touchscreen devices show a very intuitive interaction design that allows even children to use such a device successfully.

To explain why touchscreen devices are intuitive to such strong extend, a look at the three-layered brain model is helpful. To use a tool (in a computer context a tool means a device like a keyboard, a mouse, a game controller, etc.) humans have to make use of their neocortex. The cerebrum represents the highest layer in the brain model. In contrast, for 'touches', as needed during the use of a touchscreen device, humans only need to use the reptilian brain, which is represented in the lowest layer in the three-layered brain model [14]. Both aspects, (a) the widely use of gaming apps and (b) the intuitive aspect of modern touchscreen devices lead to the conclusion that a gaming app based CDG is the right choice for GHOST. Considering the broad target audience it is further reasonable to use a tablet based gaming app because of the larger screen size compared to a smartphone.

According to the last section, a CDG should be similar in quality to an entertainment game (req. 5). Modern gaming apps with the scope to be played over a longer period of time (as it is planned in GHOST) implement a three-dimensional, high quality looking game environment regardless of the genre (see e.g., Lara Croft Go, Lego Star Wars, Jam League, Modern Combat, Asphalt, Bothers: a tale of two sons, etc.). By that, GHOST has to be a three-dimensional tablet based CDG. On the other hand, GHOST has to be accessible for every target group member (req. 3). Thus, an appropriate game interaction system has to be found, that allows three-dimensional tablet based playing even for people who have never played a video game in their live. However, there are well established interaction systems for videogames that are also adapted for touchscreen devices.

The three most common used are 1st-Person, 3rd-Person and God view. The idea behind the 1st-Person perspective is that the player sees through the eyes of his player-character (PC) [15]. In conventional video games, the player controls the PC with mouse and keyboard [16]. Touchscreen based 1st-Person games are usually implemented in landscape mode. To control the PC the left and right thumb are used. The left thumb is used in the lower left area of the screen to control the movement of the PC. The right thumb is used in the lower right area of the screen to control the viewing direction [17].

In games that implement a 3rd-person perspective, a camera is used, which is aligned to the top of the PC to show him completely. Sometimes 3rd-person is implemented with „Trailing" option, then the camera is anchored at head height behind the PC. In classic video games, the control is similar to 1st-person games [16] the same applies to the touch screen control.

A God View perspective, also referred to by the terms 'overhead', 'top down' and 'God Eye', provides a perspective in which the game map is shown from above. Usually, the control is realized with the mouse [15]. Touchscreen-based God View games are often implemented by touching directly on the device. In such case the 'touch' on the device is equivalent to a mouse click. Additionally, manipulations of the camera perspective are done by the usual multi-touch gestures (e.g., two-finger zoom). Consequently, any 3D gaming interaction system known from the Computer/Laptop can be adapted for touch screen based games.

It has to be noted that the 1st-person and 3rd-person solution only replace mouse and keyboard through two equivalent virtual generated tools. By that, according to Schell [14], neocortex participation is still needed and whereby the advantage of a touchscreen solution is not exploited. Only the 'God View' interaction systems provide a solution that's natively transforms touch into interaction. As a result, this kind of game interaction should be manageable for inexperienced players and therefore is the right solutions for a touchscreen based CDG and GHOST.

However, this question cannot be clarified for the intended target audience based on the state of scientific research. There is a lack of empirical research that investigates the suitability of existing touch screen-based control and camera tracking paradigms for 3D serious games. However, since a well-functioning interaction system is elemental for the CDG success, a corresponding study has been carried out that will be briefly discussed in the next section.

V. INTERACTION SYSTEM FOR A TOUCHSCREEN BASED CDG

In the following different interaction systems are discussed and the study results are presented.

A. Discussion of possible interaction systems

The main objective of the study is to investigate wheatear it is possible to find an interaction-interface for a high quality tablet based video games that does not require any video game experience. Such an interaction-interface would connect requirements 3 and 5 that seem as if they exclude each other. The presence of such an interface would open the possibility to develop a cybersecurity awareness training that fulfills all seven requirements in the first place.

From a theoretical point of view, a game that responds as intuitive as possible on touch screen input should be advantageous for the players. As shown in the last section even the 'God View' interaction system relies on not intuitive multi-touch gestures for camera control. For that reason, a new interaction system for the GHOST prototype was designed.

These 'optimized' called interaction system provides the PC control via finger touch. The PC automatically moves to the location of the map where the map was touched. Even the interaction with game objects or non-player characters (NPC) works this way. If a player, e.g., touches a game object his PC will automatically move to the point next to the object. After arriving at this point an interaction dialog opens automatically. To remove the maybe not intuitive camera control the whole game map is divided in different camera zones (partly multiple zones in one room). Each zone provides its own static camera perspective. If the player controls his avatar from one camera zone to another, the camera angle

changes automatically. The player is not aware of where the zone boundaries are, the camera angle change just happens. To help the CDG-Player's orientation, there is also a second 'optimized+' called interaction system where the camera change from one position to the next one appears in a smooth move. Additionally, to the three mentioned interactions systems (1st-Person, 3rd-Person, God View) both versions were examined in a blind study. For this purpose, a small game was designed where the participant had to find six game objects or NPCs to interact with. At the beginning of the test a participant is set in a game environment with six rooms and two corridors. The participant does not get any map because the study also refers to the orientation ability. Finally, the time needed to complete the interaction tasks was measured.

A total of five mini games (demo versions) were developed:

- Demo1: 1st-Person
- Demo2: 3rd-Person
- Demo3: God View
- Demo4: optimized+
- Demo5: optimized

Deviating from the previous explanation of 3rd-Person interaction-systems the 3rd-Person PC control was changed. Usually the PC is controlled with the left and right thumb as in a 1st-Person tablet game.

Indeed, the interaction system in Demo2 uses a touch based PC movement control as in the 'optimized' demo versions. In addition, camera rotation was enabled by integrating a two-finger-rotate gesture for camera rotation. The classic two thumb control is still used in Demo1. Figures 1 to 4 are screenshots made of each demo version, respectively.

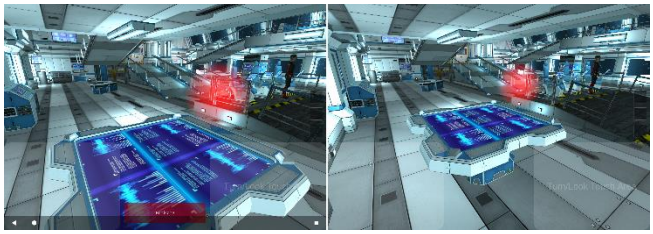

Figure 1. 1st-Person interaction system with dynamic appearing 'activate'-button for object interaction (Demo1).

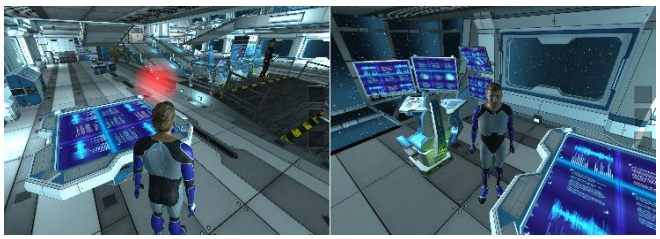

Figure 2. 3rd-Person interaction system before and after two-finger-rotate (Demo2).

Figure 3. Good-View before and after gesture based camera rotation (Demo3).

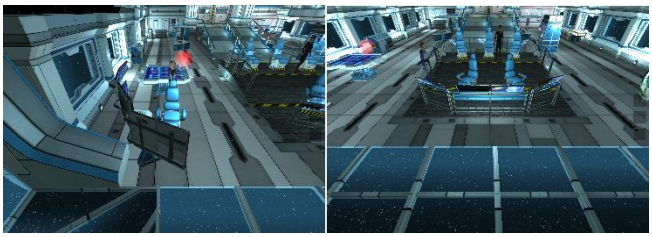

Figure 4. Adjacent camera zones in the optimized(+) interaction system (Demo4&5).

B. *Summary of Study Results*

TABLE I. SUBJECT DISTRIBUTION

	subject distribution				
	Demo1	***Demo2***	***Demo3***	***Demo4***	***Demo5***
age<=37	7	7	7	7	6
age>37	6	6	6	6	6
$\bar{x}$ age	39	38	40	41	41
SD age	17	16	16	15	15
n woman	6	6	6	6	6
n men	7	7	7	7	6
n	13	13	13	13	12

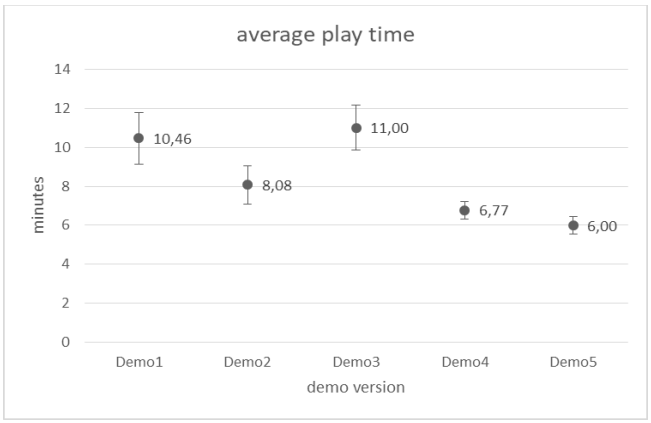

Figure 5. Average play time and 95% confidence interval.

In total 64 participants participated in the study. Table I provides information about the exact distribution of the test subjects to the individual demo versions.

An ANOVA was calculated and, by that, proved that the playtime differences are statistically significant ($\alpha = .05$; $F_{(4,59)} = 4{,}26$; $p < 0{,}0011$) . Figure 5 shows the average playtime for each demo version. It can be seen, that the playing time of the demo versions 4 and 5 are the shortest ones. As a result, the assumption that an intuitive interaction system simplifies the access to the game can be confirmed. By that, the 'optimized' or 'optimized+' interaction systems are the most suitable solutions for the GHOST-Prototype. Moreover, the results show that there are performance differences between the groups <=37 and >37 and that demo version 4 and 5 minimize these differences.

VI. GHOST: A CDG BASED CYBERSECURITY AWARENESS TRAINING

Following the remarks of this paper, GHOST is a turn-based, tablet-based, serious game like, Competence Developing Game, which provides a cybersecurity awareness training for end users in companies. Furthermore, in GHOST a new intuitive interaction systems was implemented. By that, it has the potential to fulfill the seven requirements which were derived in section two.

Whether GHOST meets these requirements depends on the game design. First of all the game design tracks two aspects. It creates the space to experience which personal actions are positive respectively negative for the cybersecurity. Second, it demonstrates which and why IT-department activities are necessary and meaningful. By that, it allows the end user to notice missing activities in his/her company and in addition it will increase the employee's acceptance for such activities.

In case of a cybersecurity training too many topics in a short time period increase the risk to overwhelm the exercisers [2]. Therefore, in the beginning each game round treats only one serious topic. The IT risks are hidden between other tasks and rarely occur, as in reality. In order to evaluate which serious content should find its way into the GHOST CDG, Annex 'A' of ISO 27001 was analyzed (ISO/IEC 27001: Information technology – Security techniques – Information security management systems – requirements, see [18]). In Table II, the serious topic of each game round is presented.

The idea behind GHOST's game design is to minimize the organizational effort. By a trick, GHOST still provides player the illusion of playing together. Every GHOST training is designed for 8 players in two groups at the same time. The training consists of 16 units (game rounds) in total. However, each round gets a specific time period in which the round is active and ready for play. In this period each player can choose the moment to play the round individually. At the end of the time period the GHOST-System calculates, based on each individual result in a group, a common group result which is the starting point for the next round. If, e.g., a player misses to participate in one round the whole group result will be weakened. This kind of game design uses the business simulation advantages like group motivation and the enforcing of a specific continuous training cycle without the disadvantages of complicated appointment organization. Nevertheless, GHOST allows even real multiplayer experience. The Round 7&8 as 15&16 require all 8 players to participate the training at the same time. Each group has to be in one physical room, the merging of the groups takes place via internet. These real multiplayer rounds serve as highlights of the complete training cycle. However, since two multiplayer rounds are played at one appointment, accordingly only two appointments must be arranged. As a result GHOST provides 16 play rounds and only requires the coordination of two appointments, which results in a huge reduction of the organizational effort compared to business simulations. Table II shows the assignment between serious content and game rounds.

As already mentioned, the serious content in GHOST is hidden between other tasks. To assure a simple knowledge transfer between the game environment and the real world it seems to be obvious to build an office environment inside the game. Accordingly, the player would solve every day work tasks inside the game world to come across serious content from time to time. This would result in a game that simulates an office for a game player whose position is currently an office, means playing-office in the office.

TABLE II. GAME ROUNDS

Round	***Serious topic***
1	Screen lock
2	Handling of foreign flash drives
3	Phishing-Mails
4	Backups
5	Mobile Devices (especially Smartphones)
6	Websites, software installation, own IT infrastructure
7&8 (MP)	Passwords, Information encoding, Emergency response, Environmental Security, Backups
9	Access rights
10	Environmental Security, safe workplace
11	Virus prevention, Keylogger, Work delegation
12	Network Devices, Audits,
13	Log files, Access Right Management
14	Quiz Round
15&16 (MP)	Flash drive, Information encoding, Phishing-Mails, Malware, Passwords, Emergency response

MP = Multiplayer

This would most likely ruin the fun aspect of the game, what would gamble away the main advantage of a CDG, the transfer of the motivation of a game situation to a serious context. For this reason, the game was moved 50 years into the future. The players find themselves in a science fiction scenario on a space ship named GHOST. They experience a journey of sixteen laps (one lap one round) and figure out quickly that someone tries to sabotage the mission by infiltrating the ship's computer systems.

As a crew member each player has to handle a lot of day-to-day tasks, which are intentionally similar to 2018 tasks in a

normal office. Nevertheless, a player has to be constantly on guard while interacting with the computer systems or other aspects in his environment. The assailant could start the next cyber-attack in any moment, with any strategy.

VII. PROTOTYPE FOR EVALUATION

As shown in Section II, the awareness training should fulfill at least seven requirements to match employer and employee expectations. Most of them can be fulfilled through design decisions described in this paper: A GHOST training can take place at the place of work to reduce the time expense. Since an extensive preparation is not needed the organizational overhead is reduced. This helps to reduce the training costs (req. 2a-c). Because of its sophisticated empirical evaluated (see Section V) interaction system even employees without any game experience can participate the training (req. 3). In addition, this interaction system helps GHOST to have an entertainment game look and feel (req. 5). The turn-based, business game inspired, game design allows further a continuous training cycle, that is made possible with a computer-based training (req. 6 and 7). Moreover, the social significance of - and the increased attacks on- IT systems leave no doubt on the real-life relevance of the underlying problem (req. 4). Therefore, on to this point only requirement 1 is left unmentioned. Requirement 1 demands a CDG to help an employee to develop skills, to increase his motivation and satisfaction and to strengthen the job relation. The last both aspects of requirement 1 can presumably only be evaluated when the GHOST CDG is completely developed (as described in Table II) and used in practice. But the first aspect of requirement 1 -to develop skills- can be evaluated with a prototype. Therefore, a prototype was developed that follows the principles shown in this paper (for an overview see Section VI). To provide a game situation to the participants with proper length to gain an intense impression the prototype should cover around one hour of gaming. Accordingly, to develop just one game round would not be purposeful. Instead four serious topics: “Screen lock”, “Handling of foreign flash drives”, “Network Devices” and “Passwords” were combined to one large gaming round that is implemented for evaluation reasons only.

In the beginning of the prototype an introduction video is presented to the participants. The video covers the control elements of the game and explains them. The whole interaction system is equal to the optimized+ interaction system as shown before. The camera moves automatically in a smooth way and for the game objects interaction the participant in every case needs a one finger touch to start interaction.

A. Storyline overview

During the gameplay the participant finds out that he is on a space ship called GHOST on a mission to find a new discovered high energy element: Industrium. Overall, the participant has to pass eight quests. He deals with the sabotage of the crew's mission. In the beginning, the participant is presented with the conundrum of what to do with an unfamiliar flash drive prompting an investigation by the chief of security into its origins. This is the first of several attacks that are made on the ship's security. As the game progresses, the crew becomes more nervous and the participant must assist in improving the ship's security. However, all efforts are too late as just after industrium collection is concluded the main systems of the ship suddenly shut down. The chief engineer explains that the systems responsible for keeping them alive and creating fuel from the harvested industrium are failing due to the disturbance. The participant is tasked with finding the devices that are causing the disturbance and restarting the system. Once he has finished this task, the crew is saved and prepares for a leap through space.

B. Game play and serious content

Quest 1 gameplay: The participant must activate the ships systems and he must find the ghost-drive of the quartermaster (a device that looks like a flash-drive). In doing so the participant has to find his way through the ship to find a computer console that is marked with an arrow. After that, the ships lights are activated, and the participant will find the ghost-drive nearby. In the end of the quest the participant is told to keep the found ghost-drive because he needs one for his next task anyway.

Quest 1 serious content: The ghost-drive is infected with a virus (what the participant does not know about). From the moment the participant finds the drive he has the possibility to visit the security chief to get the problem fixed. (Serious goal: Flash-Drive security awareness) (see Figure 5)

Quest 2 gameplay: The participant gets the task to collect status reports from five crew members who are in the rear sections of the ship. The crew members will transfer their reports to his ghost-drive. At the quest end the participant will merge the reports by using his terminal and sent them to the captain (only a few clicks needed).

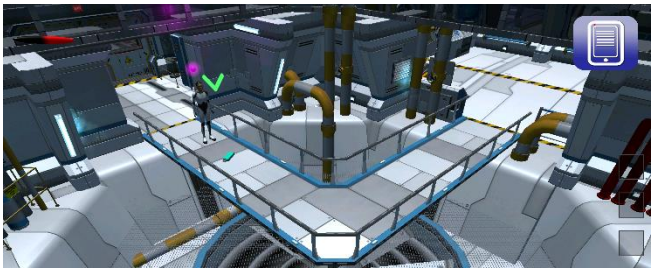

Figure 5. Finding the lost ghost-drive after activating the ship systems (cropped)

Quest 2 serious content: The participant still has the opportunity to find and fix the virus problem by visiting the security chef. When the participant speaks to one of the five crew members and if his drive is still infected he can choose if he wants to do “something else” or if he gives his ghost-drive to the person he is speaking with. If the participant infects a crew member's computer, the security chef will arrive in seconds, detect the problem, explain the problem and hand over a ghost-drive that is safe to use. The negative consequences are that the report is lost and that every other

crew member mentions the incident later in quest 2. However, close to the quest end and after the participant sends the merged report away he has two options to leave the terminal: The more obvious option is to touch on “leave the terminal”. This is equal to leave a PC unlocked. Second, the participant can touch the “Show shutdown menu” Button that reveals the “lock” Button for leaving the terminal in a safe way. (Serious goals: Flash-Drive security awareness)

Quest 3 gameplay: The participant is requested to the bridge. On the bridge the participant and the captain have a small talk about the ship systems and the merged report.

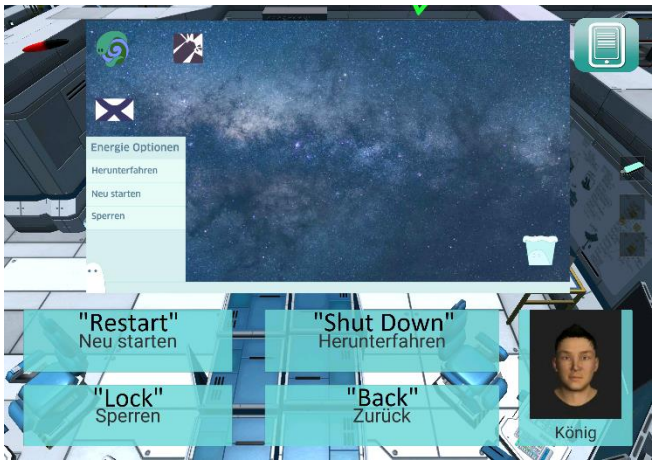

Figure 6. Options after touching the “Show shutdown menu” button (Button translation inserted)

Quest 3 serious content: During the chat the camera suddenly moves onto the main screen of the bridge. Depending on whether the player locked his screen in Quest 2 or not there is a different email-like message on the screen. If he forgot to lock his screen the participant is addressed directly by his name and the mail is sent from his terminal. But if he locked his screen in Quest 2 the message is addressed to a crew member and sent from the crew member's terminal. In both cases the captain points out that someone made a joke and that it is important to lock the screen always. (Serious goal: more frequent screen locking) (see Figure 6)

Quest 4 gameplay: The participant will be requested to the security chef. They chat about the infected ghost-drive and the security chef points out that he needs help to generate new passwords that are good to remember.

Quest 4 serious content: The password generation is wrapped in a mini-game. During the game, the participant has to shot on eight words that will be the long enough to be a good base for the password generation. If the participant shoots a short word he loses one of the already collected long words. However, after the collection of words the player modifies the words to passwords. For that, he selects a character he wants to change or add (e.g., a 1 for an i, etc.) and tries to shoot down the wished character. As he makes the changes, he sees a constantly changing display indicating how secure the password currently is. By that, the participant gets a feel for what makes a password secure. (Serious goal: teach how to build a safe password) (see Figure 7).

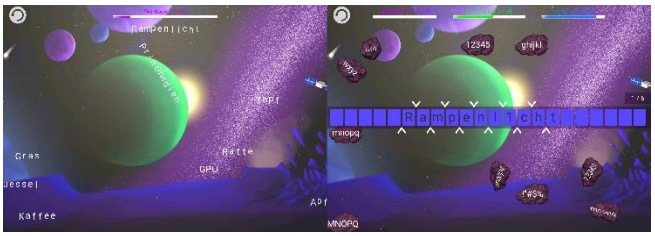

Figure 7. Mini-Game for the password generation. Left: shoot long words; right: modify words to passwords

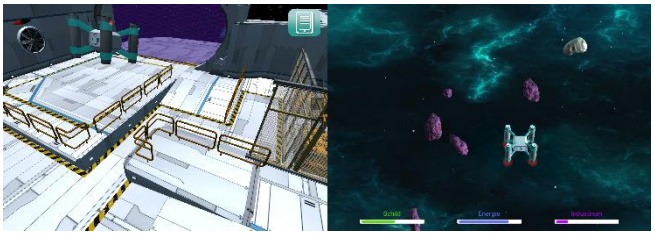

Figure 8. Mini-Game for the industrium collection. Left: animated drone start; right: industrium collection

Quest 5 gameplay: The participant gets a call that he has to check the current industrium research reports that are send as a message to his terminal.

Quest 5 serious content: After using his terminal the participant has to remember to lock his screen comparable to quest 2. If he remembers to lock his screen he gets a positive feedback from the security chef after a while. If he forgets to lock the screen he gets an equivalent negative feedback. (Serious goal: more frequent screen locking)

Quest 6 gameplay: The participant has to collect industrium with a remote-controlled drone. The drone-flight is implemented as a mini game. The participant controls the drone with his finger. He has to hit the pink asteroids for collecting industrium while avoiding the other ones (see Figure 8).

Quest 6 serious content: none.

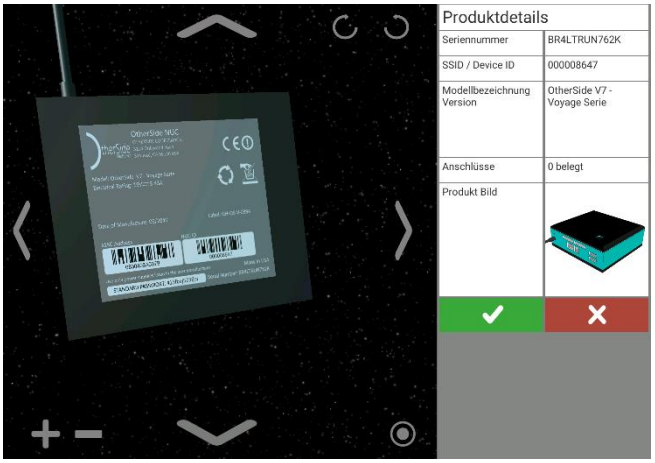

Figure 9. Mini-Game: “Network-Devices”

Quest 7 gameplay: There is a shipside system failure and it is not possible to reactivate the ship's systems. The participant has to help to identify if there are any corrupt devices on the ship. His search area is the communication room and the mess. In the end of the quest the participant is able to reactivate the systems in the same way as in quest 1.

Quest 7 serious content: The player has to check devices that are similar to network devices like network switches or repeaters. The checking is implemented as a mini game where the player compares a device on the ship with the manual (e.g., number of free ports, picture, serial number, etc.). The player has to decide if the device is safe or not. When he decided to report a device, he has to choose which aspect is corrupted. (Serious goal: Create awareness that new network devices could leak the security chain) (see Figure 9).

Figure 10. Debriefing supported by drawings (current topic on screenshot: screen lock)

Quest 8 gameplay: The quartermaster informs the participant, that his leap capsule is ready. After entering the capsule, the prototype finalizes.

Quest 8 serious content: Before the end a debriefing is shown. The debriefing picks up all serious topics and explains them one last time. The debriefing presentation is supported by drawings (Serious goal: deepening and transfer) (see Figure 10).

C. Experimental procedure

Each participant playing the prototype is supported by a test leader. The test leader is allowed to offer help to the participant whereby the amount of help is strictly regulated through the test design. After playing the prototype the participant has to fill out a questionnaire. In addition, approximately two months after the prototype-based training the participant gets a second short questionnaire via email. The first questionnaire contains three objects of investigation: "game experience", "prototype review" and "competence growth". The second questionnaire is only about "competence growth".

To measure the "game experience" the core module of the "The Game Experience Questionnaire" is used. The core module assesses the game experience separated in seven components [19]. The items of the questionnaire are translated to the German language enabling the participants to use their native language.

To receive a standardized game review from the participants a cross section of the work from Vohwinkel is used. Vohwinkel presents a well evaluated questionnaire for standardized game reviews [20]. He takes a variety of usability and game work into account and reorganizes them to a full-scale measuring instrument for commercial video games.

As part of the research project it is not possible to measure the "competence growth" in a real-life work situation of the participants. Instead the participants are asked three times after a self-assessment. First, for each of the four serious aspects they are asked how they handled the aspect before they participated to the prototype-based training. In the end of the long questionnaire they are asked again with a changed focus. Now they should assess how they plan to act in the future. Then, in the questionnaire that the participants received after approximately two months, they are asked how they actually acted in the last months. In total, this creates an overall picture of the self-assessed competence situation. The self-assessment questions are formulated as follows, each adapted to the position in the questionnaire/s:

- I locked my screen when leaving my place of work
- If I recognized new IT-Devices on my place of work I was thinking about whether it is necessary to report them to somebody.
- Before using a flash-drive I was thinking about if it is safe.
- I knew exactly how to generate an easy to remember and safe password.

As shown in the interaction system study there are differences in the play times between the groups "age<=37" and "age>37". Other play time relevant factors were not identified. It was shown that the interaction systems optimized and optimized+ are able to reduce the play time differences. To further reduce these play time differences to a minimum an interactive map is added to the prototype. In addition, the participant gets navigational help through the test leader if necessary. In later implementations this kind of guiding should be made automatically by the game itself.

However, one evaluation goal is to discover how differently the play performance and the game impression between the age groups still are. So, the described aspects of investigation are evaluated for each age group separately. Because there is approximately one year between the both empirical studies the age groups for this study are defined as: "age<=38" and "age>38".

VIII. EVALUATION

In this section the game experience and the competence growth are discussed.

A. Game experience & game review

Overall 31 participants take part in the study and completed 1,777 minutes of play time. The follow-up

questionnaire after two months got 14 responses. Table III shows the distribution of participants and Figure 11 shows an evaluation example.

TABLE III. SUBJECT DISTRIBUTION

	Distribution
Age<=38	19
Age>38	12
$\bar{x}$ *age*	35.7
SD age	15.3
n woman	9
n men	22
n	31

Figure 11. Evaluation

The participants evaluated the game experience and reviewed the game on the same five-point scale (1 to 5). During this analysis the averaged answers are interpreted as school marks in the following way:

- [>=1.0 “E” <1.8] (worst grade),
- [>=1.8 “D” <2.6],
- [>=2.6 “C” <3.4],
- [>=3.4 “B” <4.2],
- [>=4.2 “A” <=5] (best grade)

On average the participants of both age groups assess the game experience with an B (3.5). Thereby, only the game experience “challenge” got a bad rating (D). One possible explanation is a too low level of difficulty. Nevertheless, the data points out that both age groups had a similar positive game experience with little weaknesses only. Table IV shows the results of the seven components of the measured game experiences in both age groups.

Beyond the game experience evaluation, the participants reviewed the prototype using an adapted measuring instrument for commercial video games. Again, both age groups reviewed in a very similar way by giving an B mostly. On detail, the participants who correspond to the group “Age>38” rated minimal better. Table V shows the results in detail.

TABLE IV. GAME EXPERIENCE

	Age <= 38		**Age >38**	
Component	$\bar{x}$	***mark***	$\bar{x}$	***mark***
Competence	3.6	B	3.8	B
Sensory and Imaginative Immersion	3.0	C	3.2	C
Flow	3.2	C	3.0	C
Tension/ Annoyance	1.6 (4.4)	A	1,3 (4.7)	A
Challenge	2.2	D	1.9	D
Negative affect	1.9 (4.1)	B	1.7 (4.3)	A
Positive affect	3.9	B	3.7	B
Average	3.5	B	3.5	B

see [19]

TABLE V. PROTOTYPE REVIEW

	Age <= 38		**Age >38**	
Component	$\bar{x}$	***mark***	$\bar{x}$	***mark***
Graphics / Camera / Control	3.7	B	4.2	B
Narration / Avatar / NPCs	3.7	B	3.7	B
Help / easy game learning	3.9	B	4.1	B
Traceability / Game-Goals	4.1	B	4.3	A
Average	3.9	B	4.1	B

see [20]

In addition, the play time of each participant was measured. The mean playing time of all participants was 57.4 minutes. Thereby, the mean playing time difference between the both age groups was only about 5 minutes. The group “age<=38” needed an average of 55.4 minutes to play the prototype while the other group “age>38” needed with 60.4 minutes a little more time. Figure 12 represents a scatter plot for the variables age and play time. A relationship between age and play time is visible. A Pearson's correlation was calculated with a result of 0.30, so a light correlation was detected. With a p-value of .285 in the present sample the correlation is not statistically significant. However, a five-minute play time difference has no impact on the practical usability of the concept.

The interpretation of the presented data indicates that the combination between interaction system and game design minimizes the differences between the age groups so far that these are no longer significant. This can be seen in all the three presented evaluation aspects. Further it can be determined that the participants evaluated the prototype’s gaming aspects in a positive way. This impression is strengthened through a further item in the questionnaire. The participants were directly asked about their overall impression and rated the prototype in mean with 7.7 out of 10 points (B). Accordingly, the differentiate review and the overall impression of the prototype are consistent and both positive.

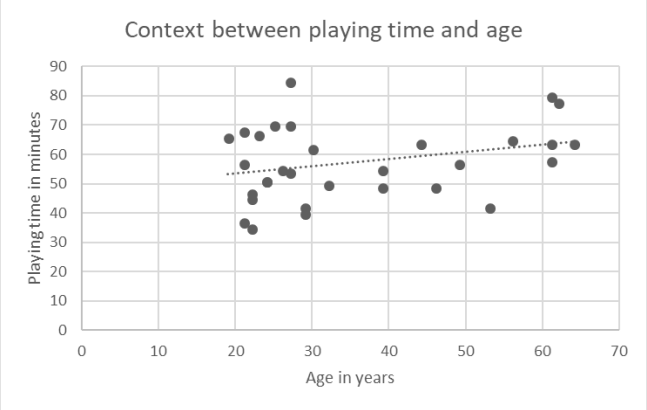

Figure 12. Scatter plot: Age->Playing time

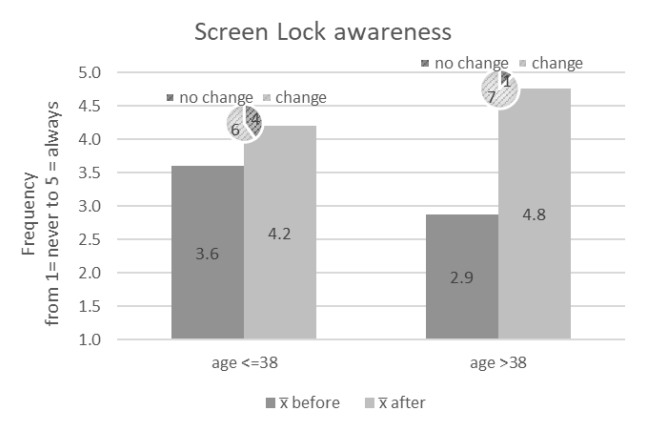

Figure 13. Screen Lock

B. Competence growth

In the following the growing of the participant's competences for each of the four serious aspects is evaluated. Thereby, only records allowing a competence growth are used. That means if a participant stated that he always locked his screen or that he perfectly knew how to generate a password even before he participated on the prototype-based training his record regarding the specific serious content is not used.

For the measurement of the three serious topics "Screen lock", "Handling of foreign flash drives" and "Network Devices" frequencies were queried (5-Point-Scale). According to the scale the mean results are interpreted in the following way:

- [>=1.0 "Never" <1.8] (worst grade),
- [>=1.8 "Rarely" <2.6],
- [>=2.6 "Occasionally" <3.4],
- [>=3.4 "Often" <4.2],
- [>=4.2 "Always" <=5] (best grade)

Figure 13 shows the results of the participants self-assessment regarding the serious content "Screen lock". Before participating in the prototype usage, the group "age<=38" in mean stated to often lock the screen (3.6) ($\bar{x}$ before). After the training participation the average frequency value was 4.2 (often) ($\bar{x}$ after) whereby 6 of 10 people improved their competences. The group "age>38" chose an average frequency of occasionally (2.9) before participating in the training. After the training they stated that they are planning to lock their screens in future always (4.9). Overall, 7 of 8 participants were able to increase their performance. Summarized the measurements for both groups show satisfactory results after participating the prototype training. It is noticeable, that the group with less previous competences could leap higher, which results in a similar competence level between both groups.

The follow-up questionnaire (after two month) contained 6 relevant records for the younger and 3 records for the older group. Overall, when asked how frequent they have locked their screen since prototype testing, the younger group shows two deviations in the size of: once -1 and once -2. That results in comparison to $\bar{x}_a$ ($\bar{x}_a = \bar{x}$ after) in a mean loss of -0.03 ($\bar{x}_{a+} - \bar{x}_a$) (Note: The differences are calculated precise with 15 digits). The records of the older group contained one difference of -1, which results in a mean loss of -0.08. For the present sample this leads to the conclusion that the prototype has a long-lasting aftereffect regarding the serious content "screen lock". The deviations can be neglected because of their low severity. Table VI gives an overview about the follow-up survey.

TABLE VI. SCREEN LOCK FOLLOW UP SURVEY

Group	$\bar{x}_{a+}$	$\bar{x}_{a+} - \bar{x}_a$	Absolute change		AVG change
			-1	-2	
age<=38	4.17	-0.03	1	1	-0.5
age>38	4.67	-0.08	1	-	-0.3

$\bar{x}_{a+}$ = mean frequency in the relevant follow up records

Figure 14 shows the results of the "Flash Drives" assessment. The members of the group "age<=38" stated that they occasionally (2.4) think about whether the use of a flash-drive is safe. After the training, the measured frequency-value grew into the "often" area (3.8). A total of 14 participants had the chance to increase their competence and 11 of them did so. The average of the group "age>38" was 3.5 (often) before the training. After participating in the training, the group stated to think about flash-drive security always (4.5) in the future. Overall, 6 of 8 participants were able to change their awareness. However, both groups achieved a change, where in this case the change for the younger group is more pronounced. It seems, that the development-potential depends on the individual foreknowledge and not on the group membership.

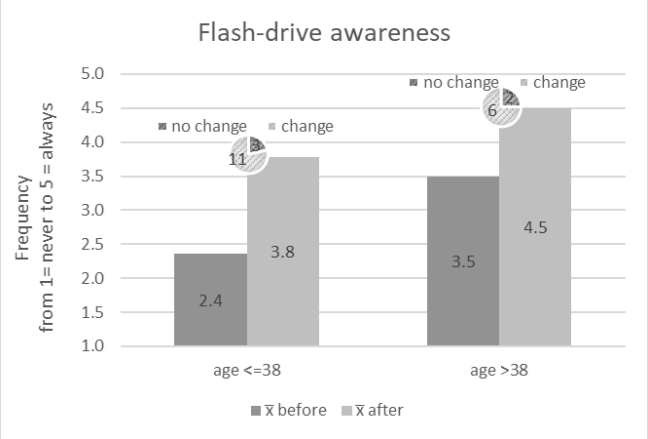

Figure 14. Flash-drive

The follow-up survey results in 8 relevant records for the younger and 4 records for the older group. The group "age<=38" shows 5 differences: three times -3, once -2 and once +1 which results in a mean loss of -0.29 compared to $\bar{x}\ after$. The other group points 3 differences: once -1, once -2 and once +1. This results in a mean loss of -0.50. By that, for the present sample, the deviation can be neglected again. The aftereffect is long-lasting too. It is noticeable, that two participants have changed their behavior more than planned. A possible explanation are exchanges with colleagues or deepening thoughts in the aftermath of the training. This emphasizes that the training's serious topics remain in the consciousness of the subjects even beyond the training. The data are shown in Table VII.

TABLE VII. FLASH-DRIVE FOLLOW UP SURVEY

Group	$\bar{x}_{a+}$	$\bar{x}_{a+} - \bar{x}_a$	Absolute change			AVG change
			-1	*-2*	*+1*	
age<=38	3.5	-0.29	3	1	1	-0.5
age>38	4.0	-0.50	1	1	1	-0.3

Figure 15 shows the data related to the "Network devices" topic. The data indicates, that the competences before the training were very low. In total the group "age<=38" contains 15 relevant records while the group "age>38" contains 8. The mean data of the younger participants shows that they were thinking rarely (1.9) about whether it could be necessary to report new devices. With a value of 1.8 (rarely) the results of the older participants are similar. Accordingly, a large competence increase was achieved through the training. Both groups stated that in future they will think always (4.5 and 4.9) about whether new IT-devices are authorized or not. Moreover, all 23 relevant participants achieved a competence growth. By that, the assumption potential of development depending on the individual foreknowledge and not on the group membership seems to be confirmed. The GHOST-based Training works out for the whole target audience.

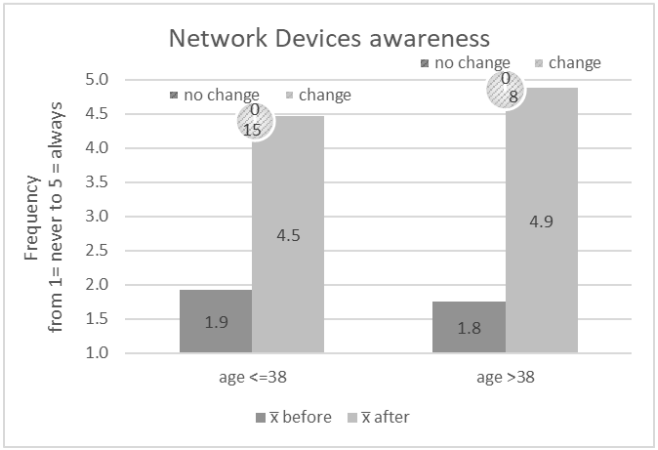

Figure 15. Network Devices

At the follow-up survey, 12 relevant records were recorded. Again, only small differences were found. The group "age<=38" shows 8 relevant records whereby 5 records show deviations. The mean difference is -0.22. In the other group 4 records were registered whereby 2 of them show differences. Overall, the average frequency in the group "age>38" dropped by -0.38. Therefore, for the present data, the deviation can be neglected again. A data overview is presented in Table VIII.

TABLE VIII. NETWIRK DEVICES FOLLOW UP SURVEY

Group	$\bar{x}_{a+}$	$\bar{x}_{a+} - \bar{x}_a$	Absolute change				AVG change
			-1	*-2*	*-3*	*+1*	
age<=38	4.25	-0.22	1	1	1	2	-0.5
age>38	4..5	-0.38	2	-	-	-	-0.5

Based on the assumption that the importance of secure passwords is common sense the password aspect of the prototype is not a classic awareness training. Rather the focus is to teach how to create a safe and simple to remember password. As shown in Section VII, the password mini-game represents an exception in the game design. Also, the mini-game is controllable with one finger, its game mechanic includes action elements that require a quick gameplay. By that it is exploratory checked whether the older age group is able to participate on CDGs that require an action gameplay.

Therefore, the participants were asked to self-assess their ability to generate safe and easy to use passwords. A 4-Point-Scale was used (from strongly disagree to strongly agree). According to the scale, the ability to generate passwords are interpreted in the following way:

- [>=1.0 "D" <1.75] (no ability),
- [>=1.75 "C" <2.5],
- [>=2.5 "B" <3.25],
- [>=3.25 "A" <4.0], (fully capable)

Figure 16 shows the evaluation results regarding the password generation. It is noticeable that the ability before training to generate passwords was already strong. Only 7 participants of the Group "age<=38" and 4 participants of the

group "age>38" had the option to strengthen their ability. This shows how well known the password security topic is especially to the older participants. Maybe another sub-topic of the password theme (e.g., sharing passwords or multiple using of password) would had been more useful for this evaluation to measure more results in the older group. However, for the participants that are 38 or younger the results shown that the measured mean ability to generate safe passwords starts within the B (3.0) area. After the prototype participation it growths into the A (3.9) area. Moreover, 6 of 7 participants were able to achieve a development. The group "age>38" starts within the B (2.5) area and ends within the A (3.5) area but only the half (2 of 4) of the participants improved through the training the other half showed no change. That may indicate that the needed quick game-play required to solve the password mini game overwarm a part of the participants in that group. But because there are only 4 relevant records in this sample that kind of assumption cannot be proved with this study. A further investigation is needed. Regardless to that, it could be shown that a calm-gameplay works out to convey serious content to the whole target audience.

An evaluation of the follow-up survey is not made because of a lack of data.

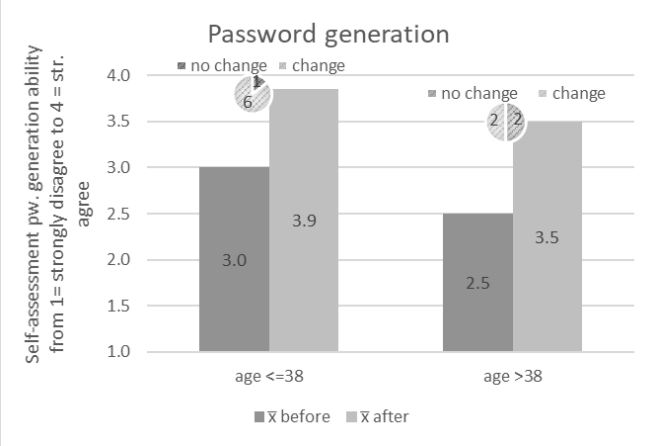

Figure 16. Passwords

To determine whether the measured awareness (Screen Lock, Flash-drive, Network Devices) or ability (Password generation) changes statistically significant t-tests are performed. It is assumed that the training increases the awareness or ability. Therefore, one-tailed t-tests for dependent samples were calculated. An overview of the training effects and the corresponding t-test results are shown in Table IX. The results show that with a $\alpha = .05$ the changes are statistically significant. There was only one exception found. The group "Age>38" shows no statistic significant change regarding the password generation topic. A possible explanation can be found in two aspects. First, only the password mini-game requires a quick gameplay because of its action-based game mechanics. Second, there were only four participants in that group that had the possibility to improve their ability to generate passwords. Therefore, it is not possible to select which of these both aspects were the crucial one by studying the data. Nevertheless, the test leaders pointed out that they noticed many participants of the older group having trouble playing the password mini game. Such a subjective impression was not reported for any of the other training sections.

TABLE IX. OVERVIEW OF THE TRAINING EFFECTS

	Age<=38			***Age>38***		
	x̄ before	**x̄ after**	**p-value**	**x̄ before**	**x̄ after**	**p-value**
Screen Lock	3.6	4.2	.003	2.9	4.8	.001
Flash-drive	2.4	3.8	.0004	3.5	4.5	.004
Network Devices	1.9	4.5	<.0001	1.8	4.9	.0002
Password	3.0	3.9	.0005	2.5	3.5	.09

IX. CONCLUSION

GHOST is a new approach to perform a cybersecurity awareness training for end users in companies. It was shown how the serious game content was systematically developed out of the well-known ISO 27001 and it was also elaborated what kind of requirements a cybersecurity awareness training should fulfill. Further it was shown that the majority of the resulting seven requirements could be fulfilled through an adequate game design. A GHOST training can take place at the place of work to reduce the time expense. Since an extensive preparation is not needed the organizational overhead is reduced. Both aspects also reduce the training costs (req. 2a-c). The turn-based, business game inspired game design allows further a continuous training cycle, that is made possible with a computer-based training (req. 6 and 7). Moreover, the social significance of - and the increased attacks on- IT systems leave no doubt on the real-life relevance of the underlying problem (req. 4).

The requirements 1, 3 and 5 needed a further investigation. Requirement 5 asks for a game quality that is similar to entertainment games. It is shown that nowadays even mobile entertainment games have a sophisticated game environment often represented as a three-dimensional game world. Requirement 3 asks to make the training accessible for every target group member. To fulfill these both requirements a new kind of interaction design for three-dimensional tablet games is developed and evaluated through an empirical study.

Requirement 1 asks amongst other things for a training that helps the participants to develop specific skills. To prove this aspect a prototype that includes the four serious topics Screen Lock, Flash-drive, Network Devices and Password is implemented. The prototype is designed to fulfill the awareness training requirements that are introduced in this paper. By that, the prototype is suitable for an evaluation of the GHOST concept. An appropriate evaluation was performed through an empiric study. The results indicate that the GHOST prototype leads to a grown cybersecurity awareness and at the same time is enjoyable. Thereby, it can be shown that the postulated requirements and the proposed

implementation leads to a productive Competence Developing Game for the Cybersecurity Awareness Training.

Future research could evaluate to what extend the GHOST concept is usable for CDGs for other serios topic. To considerate CDGs for related topics in first, could be a meaningful approach. In this context, it is planned to examine the usefulness of the GHOST concept for digitalization education as a next step. Additionally, the implementation of the whole 16 game round CDG in a commercial context is intended.

REFERENCES

[1] J. A. König, M. R. Wolf, "Cybersecurity Awareness Training provided by the Competence Developing Game GHOST", ACHI 2018, pp. 81-87, 2018.

[2] A. Nagarajan, J. M. Allbeck, A. Sood and T. L. Janssen, "Exploring game design for cybersecurity training", Cyber Technology in Automation, Control, and Intelligent Systems (CYBER), IEEE International Conference, pp. 256-262, 2012.

[3] S. Culp, "Cyber Risk: People Are Often The Weakest Link In The Security Chain", Forbes, [Online]. Available: https://www.forbes.com/sites/steveculp/2016/05/10/cyber-risk-people-are-often-the-weakest-link-in-the-security-chain, Last access: 13.11.2017, 2016.

[4] J. A. König and M. R. Wolf, "A New Definition of Competence Developing Games," ACHI 2016, pp. 95-97, 2016.

[5] J. A. König and M. R. Wolf, "The Pyramid Assessment Framework for 'Competence Developing Games'", HCI International, pp. 232-237, 2016.

[6] S. Seyda and Werner, "IW Continuous Vocational Training Survey 2014 - Companies Show In-creased Committed and Invest more in Enhancing their Employees' Skills", Cologne Institute for Economic Research, 2014.

[7] Gesellschaft für Innovationsforschung und Beratung mbH., "Empirical monitoring to the qualification situation in the German economy" (original foreign title: "Empiriegestütztes Monitoring zur Qualifizierungssituation in der deutschen Wirtschaft"), Berlin: Federal Ministry for Economics Affairs and Energy, 2014.

[8] BIBB, "Data Expose to the Vocational Training Report 2017 – Information and analysis in on the development of vocational training" (original foreign title: „Datenreport zum Berufsbildungsbericht 2017 - Informationen und Analysen zur Entwicklung der beruflichen Bildung"), Federal Institute for Vocational Education and Training, 2017.

[9] R. Hunicke, M. Leblanc and R. Zubek, "MDA: a formal approach to game design and game research" In: Proceedings of the Challenges in Games AI Workshop, Nineteenth National Conference of Artificial Intelligence, 2004.

[10] S. Deterding, D. Dixon, R. Khaled and L. Nacke, "From game design elements to gamefulness: defining "gamification"", Proceedings of the 15th International Academic MindTrek Conference: Envisioning Future Media Environments. Finland: ACM, pp. 9-15, 2011.

[11] U. Blötz, "Business Games and Serious Games in Vocational Training" (original foreign title: "Planspiele und Serious Games in der beruflichen Bildung"), Bertelsmann, 2015.

[12] D. Michael and S. Chen, "Serious Games: Games That Educate, Train, and Inform," Thomson Course Technology PTR, 2005.

[13] GTAI, "The Gaming Industry in Germany", Berlin: Germany Trade and Invest, supported by Federal Ministery for Economics Affairs and Energy, 2016.

[14] J. Schell, "The Art of Game Design – A Book of Lenses", 2. Edition, mitp, 2016.

[15] L. N. Taylor, "Video games: perspective, point-of-view, and immersion", Master Thesis - University of Florida, 2002.

[16] D. A. Bowman, E. Kruijff, J. J. LaViola Jr. and I. Poupyrev, "3D User interfaces: Theory and Practise", Addison Wesley, USA, 2005.

[17] T. Hynninen, "First-Person Shooter Controls on Touchscreen Devices: a Heuristic Evaluation of Three Games on the iPod Touch", Master Thesis - University of Tampere, department of Computer Sciences, 2012.

[18] ISO/IEC 27001:2013 - Annex A, information technology - Security techniques - Information security management systems – Requirements, 2013.

[19] W. A. IJsselsteijn, Y. A. W. de Kort, K. Poels, "The Game Experience Questionnaire", Eindhoven Univversity of Technology, 2013.

[20] K. Vohwinkel, "Playability: Evaluation of computer and video games" (original foreign title: "Playability: Evaluierung von Computer- und Videospielen"), Thesis, University of Cologne, 2010.

Providing Tamper-Resistant Audit Trails with Distributed Ledger based Solutions for Forensics of IoT Systems using Cloud Resources

Magnus Westerlund
Department of Business Management and Analytics
Arcada University of Applied Sciences
Jan-Magnus Janssons plats 1, 00550 Helsinki, Finland
magnus.westerlund@arcada.fi

Mats Neovius
Faculty of Science and Engineering
Åbo Akademi University
Axelia, Piispankatu 8, 20500 Turku, Finland
mneovius@abo.fi

Göran Pulkkis
Department of Business Management and Analytics
Arcada University of Applied Sciences
Jan-Magnus Janssons plats 1, 00550 Helsinki, Finland
goran.pulkkis@arcada.fi

***Abstract*—Network and information security are often more challenging for current IoT systems than for traditional networks. Cloud computing resources used by most IoT systems are publicly accessible and thereby, through this availability, increase the risk of intrusion. The increase in the processing of sensitive data in IoT systems makes security challenges more noteworthy, particularly in light of legal issues around cross-border transfers and data protection. Technologies preventing intrusion are effective, yet not perfect. Once a system is compromised, the intruder may start to delete and to modify audit trails and system log files for covering-up the intrusion. Complete and untampered audit trails and log files are essential for the legitimate owner of an IoT system using cloud resources to estimate the losses, to reconstruct the data, to detect the origin of the intrusion attack, and eventually in a court of law be able to prosecute the attacker. Due to this, improved methods for performing forensics in IoT systems are desperately needed. IoT forensics is mostly cloud forensics, since most IoT data is currently stored in the cloud. Therefore, cloud forensics is a key component in IoT forensics. The baseline for any forensic investigation is assured data availability and integrity. In this paper, we outline how forensic evidence data can be created for IoT systems using distributed cloud resources and how the availability and integrity of this forensic data can be assured by applying distributed ledger based solutions for storing audit trails and log files securely. Given this approach, an attacker can neither delete, nor modify past trails or logs but merely stop generating new data into log files. The approach presented here is novel, yet light enough for practical use.**

Keywords-forensics; IoT; cloud computing; distributed ledger; blockchain; distributed clouds; security; computer forensics.

I. INTRODUCTION

This paper outlines a distributed ledger approach for storing the audit trail data of IoT systems using distributed cloud resources. It extends its original conference paper [1] by an elaborated outline of audit trail creation, accountability principles for IoT service providers, and a discussion. For a definition and elaboration of the distributed cloud, we direct interested readers to Westerlund and Kratzke [2].

Academic research in network and computer forensics has a long history. A systematic literature review about digital forensics investigation is presented by Alharbi et al. [3]. In this review, a forensic investigation has a proactive and a reactive phase. The proactive phase consists of

- collection of pre-defined data according to priority and volatility,
- setting of a triggering function for hypothetical suspicious events,
- preservation of data related to suspicious events, and
- preliminary analysis of data and preliminary reporting related to the adopted hypothesis about suspicious events.

The reactive phase is triggered by a suspicious event. It consists of identifying, preserving, collecting, and analysing evidence data and generation of a final report. The collected evidence data is active and passive. Active evidence data is live or dynamic evidence that exists just after a detected suspicious event, for example processes running in a computing device. Reactive evidence data is static, for example a hard drive image.

Forensic investigations can be counter-acted by anti-forensics methods which try to [4]

- prevent collection of evidence data,
- increase the time of forensic investigations,
- create misleading evidence for forensic investigations, and
- prevent digital crimes from detection.

Evidence data for forensic investigations needs therefore protection. This was considered already by Schneier and Kelsey [5] who suggest a solution for keeping an audit log on insecure servers by offering a tamper-proof forensic scheme that stored and maintained log entries. However,

with the emerging Internet of Things (IoT) technology and the shift to cloud computing, the complexity and importance of keeping a secure audit trail have drastically increased. The building blocks of an IoT device is defined to contain an entity with an energy source and a processing module which has a storage module and interfaces for sensing, actuation, and communication [6].

To secure every IoT system is an utmost challenge. Currently, embedded security solutions, middleware, and cloud security solutions are being developed for IoT security. The goal of these efforts is detection of security threats and prevention of security attacks. No single solution is hitherto known for protection of IoT systems against all types of security attacks. The forensics discussed in this paper address the means of verifiable logs for carrying evidence of source and means as well as for restoring the compromised system to a working state. IoT forensics is defined by Zawoad and Hasan [7] as one of the digital forensic branches where the main investigation process must suit the IoT infrastructure. IoT forensics has therefore a key role in its part to investigate security breaches found in the IoT infrastructure. IoT forensics is a way to reconstruct the sequential steps performed by the attacker during the attack process; providing valuable information in constructing ever more secure systems. The sequential steps are identified by collecting data from different sources such as devices, logs, applications and networks used at the time of attack.

The paper's layout is as follows: in the following section, we discuss the motivation for accountable IoT service providers. Section III provides an overview of how audit trails for IoT forensics can be obtained, the role of cloud forensics, and some case studies. Section IV presents distributed ledger-based solutions of blockchain type for tamper-resistant protected storage of audit trails. The use of distributed ledger technology (DLT) is discussed in Section V. Finally, conclusions and proposals for future work are presented in Section VI. DLT is briefly described in an Appendix with the emphasis on the blockchain.

II. ACCOUNTABILITY FOR IOT SERVICE PROVIDERS

A motivation for a shift in how organizations prioritize resource allocation and consequently the importance of how system security is perceived, has been provided by the introduction of the EU General Data Protection Regulation (GDPR) [8]. As the GDPR has a long reaching implication for service providers anywhere in the world, as long residents of the EU may use such a service, it means that the GDPR has effectively set a default and minimum requirement for such systems that handle personal data on a global scale [9]. The GDPR provides rather strict guidelines for data security, but it also requires appropriate system security so that data does not seep into unauthorized use. Duncan [10] highlights that the 72h rule for reporting security incidents to appropriate parties would have been more effective if the rule had been formulated as "after they occur", opposed to the finalized wording of the GDPR "after they are detected". Still, the accountability principles requires a company after they become aware of a security breach to inform whom this breach includes and what particular personal data has been compromised.

The accountability principles are based on several measures that a company can take to achieve compliance with the GDPR. A core principle to achieve such compliance is to adopt and implement data protection policies. For IT-systems this refers to both the development process of IT-systems and to the maintenance processes. Any changes to a system that handles personal data (data that directly or indirectly identifies a natural person) over the system lifetime must comply with this principle continuously over time. Through such an approach we can consider that data protection is by design and default. For legacy systems that have not been designed with data protection as default, it may become difficult to show that a new version of the same system has incorporated data protection by design. For distributed IoT-systems this will likely become an even bigger challenge to show using conventional methods such as using centralized logs for collection of forensic data.

The GDPR also requires that organizations define through contract such processing that is performed by a third party with the controller's permission. The controller is also obligated to maintain the original consent contract given by the data subject (owner of said personal data). Documentation is also required of any activities the controller takes in processing personal data. This may mean the storing of facial images obtained from cameras in an IoT-network, processing said images for the purpose of identifying faces, and may in some cases mean the intended future use of any derivative products from such processing. The ability for an IoT service provider to define transactional records on a granularity of an individual user will likely become necessary. As earlier mentioned for storing forensic data, using centralized storage to achieve compliance for documentation of processing and consent may become difficult. In designing a distributed IoT-network and to maintain centralized provisions for such collection efforts will not necessarily be enough. Rather a distributed transaction database, with an immutable ledger that is not susceptible to common network attacks such as Denial of Service (DoS) would be much preferable.

The accountability principles also include organizational measures that need to be taken into account. Such measures include performing data protection impact assessments for detecting solutions with high risk to data subjects' privacy. A recommended (and in certain cases required) approach is that this work is led by an independent data protection officer, with a mandate to object the development or use of particularly dangerous practices or solutions. Organizations that develop a privacy management framework and continuously follow it within all processes involving the processing of personal data, may be considered accountable and can apply for a certification scheme that should indicate a notion of trust to potential users.

For distributed technologies this may be more challenging than for centralized, because once software is deployed to the distributed nodes the service provider may lose control of said software. Due to this nature of distributed software a recommended approach is to automate both data

security and appropriate system security measures. This includes previously stated accountability principles, incl. future security updates of the complete system. In the following section we discuss in-depth the use of IoT forensics and the creation of audit trails, to better understand how to continuously monitor delivered systems.

We should also note that United States currently provides some cybersecurity provisions that requires any contractor providing Internet connected devices to US federal government to also provide written certification that the device:

- does not contain any hardware, software, or firmware component with any known security vulnerabilities or defects (some exceptions exist),
- relies on software or firmware components capable of accepting properly authenticated and trusted updates from the vendor,
- uses only non-deprecated industry-standard protocols and technologies for functions such as communication, encryption, and intercommunication with other devices, and
- does not include any fixed or hard-coded credentials used for remote administration, the delivery of updates, or communication [11].

These provisions require contractors while under contract to notify purchasing party of security vulnerabilities, to maintain software that can be updated, and to provide timely updates.

III. IoT Forensics And Audit Trails

The ability to perform forensic activities in an IoT infrastructure is a challenging task. The existence of audit trails that can be reviewed is often a missing component. Still, as has been shown for cloud computing, detecting misuse is often dependent on the ability to scan various types of logs, both on system and application level.

The creation of an audit trail for forensic investigations of IoT systems is affected by the differences between IoT forensics and traditional digital forensics. Following differences are listed by Oriwoh et al. [12]:

- Evidence sources include IoT devices such as dish washers, pressing irons, refrigerators and wearable devices.
- The number of devices for evidence retrieval is much larger since there can be thousands of devices in an IoT system.
- The quantity of evidence data is much larger and the evidence format is different because of the multitude of different devices in an IoT system.
- The location of evidence data is much more distributed including multiple IoT devices and evidence related to IoT data stored in cloud resources implemented by micro-services.
- Flexible boundary lines between networks with connected devices from which evidence data is retrieved, since a Body Area Network with connected wearable device moves with the related person between different connection networks.

The audit trail for forensic investigations of IoT systems consists of evidence sources which are categorized in related research [12] [13] [14] as

1. Evidence collected from IoT devices and sensors
2. Evidence collected from wired, wireless, and mobile network communication between IoT devices and the external world
3. Evidence collected from network perimeter devices such as firewalls, AAA servers, NAT servers, and Intrusion Detection Systems (IDS)
4. Evidence collected from hardware and software outside the network under investigation. This category includes cloud, web, social networks, ISPs and mobile network providers.

Based on this classification a 1-2-3 Zones approach to IoT forensics is proposed in [12]. Zone 1 uses evidence of category 1, Zone 2 uses evidence of categories 2 and 3, and Zone 3 uses evidence of category 4.

A proactive and reactive phase are outlined by Zulkipli et al. [15] for the creation of an audit trail for forensic investigations of IoT systems. The proactive phase is a pre-investigation phase for preparation of the IoT forensic readiness. The reactive phase a real-time phase triggered by a detected security incident. The IoT forensic readiness is divided into management readiness and technical readiness. Management readiness includes

- an investigation plan for handling an incident,
- preparation of tools, techniques, and operations to support the investigation,
- monitoring the IoT system and obtaining support for authorization, and
- preparation of investigation skills of the investigators

For technical readiness is needed a scoping plan which defines the knowledge requirements of the investigators:

- What should be identified?
- What data should be collected?
- How should the potential evidence be identified?
- How should the potential evidence be collected?
- How should the collected evidence be preserved?

In the real-time phase tree concurrent tasks are started: scanning and identification, collection, and preservation. The scanning and identification task registers IP and MAC addresses, network port numbers, URLs, and data packet sizes. The collection task collects logs, history activity traces, time stamps, and user names with related passwords. The preservation task triggers snapshots of IoT device memories, creates hashes and encryptions of the collected data and the snapshots, and sends the hashes and encryptions to a secure storage.

Models for IoT forensics audit trail creation are proposed in [7] [16] [17] [18]. These models are described in a subsection.

IoT forensics after security breaches on data integrity, confidentiality, and availability is mostly cloud forensics, since most IoT data is already being stored or will be stored in the cloud. Therefore, cloud forensics is a key component in IoT forensics and also the most challenging component in

IoT forensics in the creation of a secure audit trail for forensic investigations [14].

A. IoT Forensics Models for Audit Trail Creation.

Zawoad and Hasan proposed a conceptual model of IoT forensics [7]. A secure Evidence Preservation Module monitors how all registered IoT devices store evidence data such as network logs registry logs, sensor data, etc. in a an evidence repository database. To ensure handling of a very large evidence dataset the Hadoop Distributed File System (HDFS) [19] is proposed to be used for the stored evidence data. The integrity and confidentiality of the stored evidence data is protected by public key cryptography. The private encryption key is accessible to forensic investigators for viewing the stored evidence data. A secure Provenance Module preserves the access history of the data stored in the evidence repository database in a provenance database. The Provenance Aware File System (PASS) [20] is used for the data stored in the provenance database. The Provenance Module applies secure provenance chaining [21] to protect the data stored in the provenance database against malicious tampering. Only forensic investigators can access a Representational State Transfer (REST) [22] based web API to the evidence repository and provenance databases. Using retrieved provenance records evidence data can be fetched.

An application-specific forensics investigative model in IoT is proposed by Zia et al. [16]. The model consists of three components: Application-Specific Forensics, Digital Forensics, and Forensics process. Unique application-specific forensics issues are handled by the Application-Specific Forensics module. The 10 most popular IoT applications are ranked from high to low popularity as Smart City, Connected Industry, Connected Building, Connected Car, Smart Energy, Other, Connected Health, Smart Supply Chain, Smart Agriculture, and Smart Retail [23]. Data is extracted from IoT devices and transferred to a network or to a cloud service. Thus the data flows to the Digital Forensics Module, which consists of 3 functions IoT Forensics, Network Forensics and Cloud Forensics. The functions create logs and store trends and logs of the data flow from the Application-Specific Forensics module. The Forensics Process collects evidence from the Digital Forensics module, examines and analyses the collected evidence and creates reports.

An IoT forensic investigation model based on a top-down forensic approach methodology is proposed by Perumal et al. [17]. If a forensic investigation should be planned, the investigator should obtain a warrant and authorization to access all necessary data. The investigation start with base device identification, which refers to device-to-device communication implemented by protocols such as 3G, 4G, LTE, Wi-Fi, Ethernet, and Power Line Communication (PLC). To locate a malicious medium that has communicated with an IoT device a triage examination is carried out to retrieve evidence data. This examination deals with platforms such as router, gateway, cloud, and fog. The investigation continues with identification of the chain of custody of retrieved evidence data, analysis of all data, and storage, presentation, and proof of analysis results.

An IoT forensics model called an IoT Digital Forensic Framework is proposed by Kebande and Ray [18]. The framework consists of three modules: a proactive process, IoT forensics, and a reactive process. The proactive process implements a pre-investigation phase in the creation of an audit trail for forensic investigations of IoT systems and the reactive process, which is triggered by a security incident, implements the real time phase [15]. The IoT forensics module consists of device level forensics, network forensics, and cloud forensics in correspondence with 1-2-3 Zones approach to IoT forensics [12].

B. Cloud Forensics and Audit Trails

The last decade has entailed a transition from onsite to cloud computing. Cloud computing provides access to a pool of interconnected resources enabled by the Internet. It abstracts the hardware from the client and has a “pay-per-use” business model. In cloud computing, the resources are elastically provisioned with storage space, service, computing platforms as virtual machines [24], and networking infrastructures obtained upon request [25] [26]. Hence, cloud computing is “a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction” [25]. Three basic cloud computing service models are Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS). Contemporary cloud-based software engineering directs towards Cloud Native Applications (CNA). A CNA is a service specifically designed to run in the cloud. CNAs are often deployed as self-contained units (containers) that are designed to scale horizontally. A CNA is often implemented as micro-services [27]. Kratzke and Quint [28] have described the technicalities in detail. In addition, the availability of cloud computing resources is augmented by the Intercloud initiative [29], envisioned as the “cloud of clouds”. Hence, the Intercloud then provides virtually unlimited resources to any connected device. In this paper, we refer to connected devices as all devices that are connected to the Internet. Such devices have given rise to the Mobile cloud computing [30] and Internet-of-Things (IoT) [31]. As a mobile device may utilise or contribute to the data mass, an IoT device frequently merely contributes to the cloud relying on the service provider in administering the security and privacy of the data.

Cloud forensics has been defined as “the application of digital forensics in cloud computing as a subset of network forensics” [32] and as “to reconstruct past cloud computing events through identification, collection, preservation, examination, interpretation and reporting of digital evidence” [33]. As the former definition suggests forensics to be restricted to the network access, the latter definition includes the audit trail as a means to reconstruct events, as well as interpretation and reporting of evidence. Cloud forensics, therefore, requires audit trails to be stored in a manner with assured availability and integrity where no changes may occur. Audit trails for cloud forensics consist of collected log

data of network traffic and data processing activities of computing devices. As such data is generated it is processed by an Intrusion Detection System (IDS) that extracts features from collected log data and analyses these. State of the art IDSs provide an active network security component using machine learning techniques to determine when anomalies occur and to detect intrusions in near real-time [34]. In a SaaS or FaaS (Function as a Service) setting the cloud service provider (CSP) has the sole ability to generate system wide IDS data. However, depending on the service model, the point of responsibility deviates. A framework for cloud forensics is proposed in [35], see Fig. 1.

Log data for audit trails can be scattered and stored in different locations due to the characteristics of the cloud. In the cloud, the level of access is divided between the cloud service user and the CSP. The level of access in the basic cloud service models is shown in Fig. 2. This significantly complicates the data acquisition process. For example in the SaaS and PaaS models, only application related logs can be accessed by the cloud service user. Though in PaaS, a cloud service user can develop an application to be able to get some additional forensics data whereas, in SaaS, this is not possible. In the IaaS model, cloud service users can move to the operating system layer for acquiring forensic data. In all service models, the forensic investigators are dependent on the CSP to ensure that needed audit trail data has been collected. This is currently thus a trust issue since the availability and integrity of the data that may be affected are not transparent. Only when both parties are fully contributing to an immutable audit trail can it provide the required transparency needed for continued investigation and legal measures.

Verifiable audit trails are essential in forensic investigations to reconstruct and rigorously examine intrusions in the cloud. The reconstruction is central to find out what damage the intrusion has caused and discover sources and origins of intrusion attacks. When an attack has occurred, the cloud service user must engage a cloud forensics investigation to analyse the audit trail related to the attacked service in order to find forensic evidence. For this,

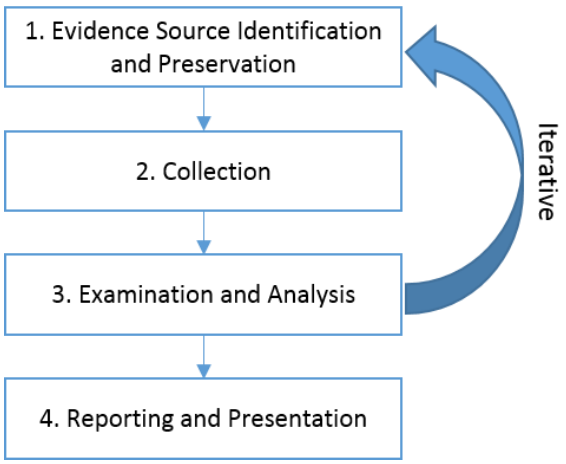

Figure 1. Cloud forensics framework proposal [35].

the audit trail is fundamental in meeting with the EU General Data Protection Regulation (GDPR), requiring enterprises to report security breaches within 72 hours after detection. Moreover, it should be possible for a CSP to present evidence on its own behalf that the source of the intrusion was external.

Traditionally, in digital forensics investigators take control of the affected physical device and perform forensic investigations on these by searching for evidence of malicious activity. As cloud computing is inherently dynamic, often the methods used traditionally in digital forensics render themselves impractical [36]. Different cloud service users may virtually share physical resources through the hypervisor and thus, to isolate the scene for forensics is next to impossible. This leads to issues that must be addressed by the forensic investigation, namely, it must be proven that any data extracted is not mixed with some other customer's data and that the availability, privacy, and integrity of the other user's data must be maintained.

Cloud forensics challenges are mostly related to architectural, data collection, and legal issues [33] [37], as well as in composing provenance data. Provenance data is the "metadata that provides details of the origins (history) of a data object" [38]. That is, provenance data is metadata tracing the history of data objects starting from original source data [39]. Complete provenance of all data stored in the cloud, all distributed computations, all data exchanges, and all transactions would enable identification of exact sources of cloud intrusion attacks and detect insider attacks in forensic investigations [40].

C. Case Studies for Reconstructing Forensic Data

Acquisition of forensic data from a network accessible smartwatch is outlined in [14] as an IoT device forensics case study. The studied smartwatch has several sensors (accelerometer, gyroscope, heart-rate sensor, and ambient light sensor), supports SMS messaging and email, can be paired with a smartphone and has following installed applications: Health App, Nike Plus App, Heartbeat App, Messages, and Maps App. Forensic data can be collected from a paired smartphone executing Cellebrite UFED forensic software [41] and by manual swipe through the smartwatch. Forensic investigators collect GPS data, heart-rate data, timestamps, MAC address, paired devices, text messages and emails, call log, contact data, etc.

The possibilities to carry out a forensic investigation on a smart TV are presented in [42]. Smart TV platforms converge traditional TV technology and computer technology and they have Internet connectivity. A smart TV device using a flash memory storage was chosen for collection and analysis of forensic data. The memory chip was removed from the motherboard of the smart TV and an image of the chip was created with the NFI Memory Toolkit II [43]. Elevated privileges, which are required for data extraction from the user space memory and for full access to the file system, were obtained for the flash memory image with a rooting procedure. Digital traces such as

- system settings: device name, connected devices, network information, and smart functions,

- use of apps: Facebook, Twitter, YouTube, etc.,
- use of web: visited web sites, search traces, etc.,
- image and multimedia files
- connected external devices: USB flash drive, hard disc, etc.,
- e-mail messages and appointments,
- use of cloud services: Dropbox, OneDrive, etc., and
- viewed TV channels

are forensically studied.

Extraction of forensic data from IoT devices in a Z-Wave [44] network is described in [45]. Z-Wave is a frequently used protocol stack in Home Area Network implementations. A typical Z-Wave network consists of controllers, sensors, and Z-Wave devices. A Z-Wave device is an IoT device (thermostat, light switch, smart locker, water valve, etc.) connected to a controller, which acts as a gateway between a Z-Wave network and Internet. Z-Wave devices can any time enter and leave a Z-Wave network. The controller assigns a unique Node ID to each Z-Wave device entering a Z-Wave network. Data extraction from a frequently used chipset with external EEPROM (Electrically Erasable Programmable Read Only Memory) on a motherboard of a Z-Wave device is described. Analysis of an event table in the EEPROM reveals which Z-wave devices worked during a specific timeframe.

IV. PROTECTION SOLUTIONS FOR AUDIT TRAIL DATA

Audit trail data for IoT system forensics requires secure protection against corruption by accidental faults and malicious forgery [46]. Protection must repel accidental corruption and all malicious anti-forensics attacks by ensuring both integrity and availability of the data.

A reasonable first choice for storage of audit trails for IoT forensics is an append-only (immutable) conventional database installation where read rights are assigned only to carefully selected set of agents. Existing implementations of immutable databases include configured conventional ones. In its most secure installation, it is hosted in-house with no means of external access and restricted physical access. Every access point (let these be logical or physical) weaken assurance of integrity. In-house installations are, however, not pragmatic for IoT systems using cloud resources; nor are the IoT systems remote installations. On this challenge, purpose-built databases and file systems are being developed, e.g., Datomic [47]. Implementation details of an immutable database for cloud audit trail are reported by Duncan and Whittington in [48].

Another attempt is the InterPlanetary File System (IPFS) [49]. The IPFS is fundamentally a protocol inspired by the Bitcoin blockchain protocol. It tries to make the web a digital resemblance to printed paper in documenting data, i.e., something that is permanent, unalterable and controllable. IPFS has a name service called InterPlanetary Name System (IPNS), which is a global namespace based on PKI [50]. IPNS serves to build trust chains and is compatible with other name services. The name services DNS, .onion, .bit, etc. can be mapped to IPNS.

The secure provenance scheme described in [21] encrypts sequences of new data, hashes the resulting datasets and provenance record, and digitally signs chains of hashed provenance records. Forensic auditors are offered access provenance data with their private keys in public key cryptography. The scheme ensures integrity and confidentiality against malicious disclosure and tampering attempts. Malicious deletion of data in the scheme is detected, but the consequence is inaccessibility to provenance data since there is no replication in the scheme.

A distributed and replicated append-only storage usually provides stronger tamper resistance than a centralized one. A distributed ledger is a replicated database, which is shared by nodes in a peer-to-peer network. Consensus algorithms are required to ensure replication and insertion across network nodes. In a truly distributed ledger, there is no central administrative node or centralized data storage. Therefore, it is considered in [51] [52] that a distributed ledger storage for audit trails typically has stronger tamper resistance than any centralized immutable database implementation.

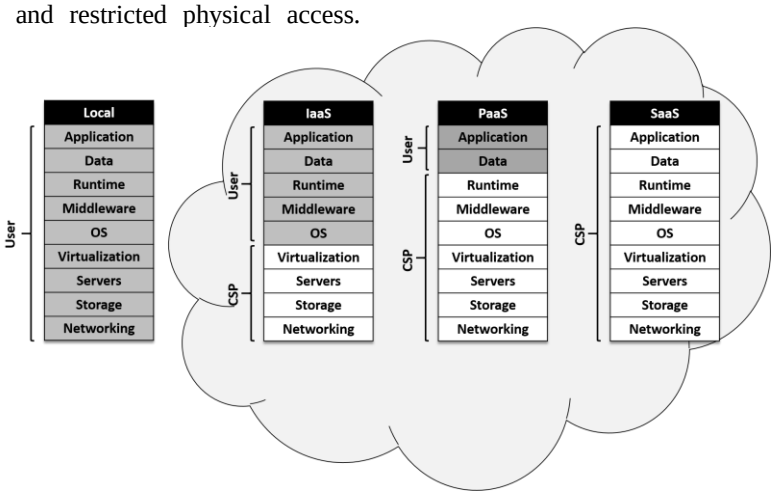

Figure 2. Access control to basic cloud service models in comparison to a local system.

The sub-sections discusses requirements for distributed ledger based solutions to protect audit trails for forensic investigations of IoT systems and presents some blockchain based solution proposals. In Section IV D, we present a novel architecture for automating and securing forensic data in distributed IoT networks. Distributed ledger technology (DLT) with the focus on blockchain technology is further described in an Appendix.

A. *Requirements for Distributed Ledger based Solutions*

In a traditional IoT architecture IoT devices are network nodes which transmit their payload data to a data store through some proxy or application programming interface. IoT device management is manual and potential device logs may often remain locally stored on the devices. Device users have credentials for authentication. Only authenticated users are authorized to access IoT devices and to update device firmware from device deliverers' databases. If a system log is stored on a respective node it would require device access for collection (pull) of data. Storage space is often very limited so only the most resent activities may be stored on the device, hence continuous collection to a centralised data store is required for ensured retention. An improved solution for a traditional architecture is presented in Fig. 3, i.e., automatically pushing log data from each node. From an accountability perspective new updates to the nodes must continuously be provided, something that often requires a manual process by a system administrator. New firmware security updates should also be provided by the manufacturer for the lifetime of said IoT devices. For this process to be complete, traditional IoT systems require many manual process steps that are often not possible to ensure in today's environment. Hence, we find it motivated to propose a new type of architecture better suited to a distributed network topology. Our proposal is presented in Section IV D.

Usage of a distributed ledger for protection of IoT forensics data is possible only if three fundamental requirements are fulfilled. First, a sufficiently large network of nodes must be available for storing replicated copies of the distributed ledger. Secondly, each network node must have sufficient storage and processing resources for management of a distributed ledger replication. Thirdly, it must be possible to extend the distributed ledger with devices producing new data at the data rate needed (i.e., throughput and scalability).

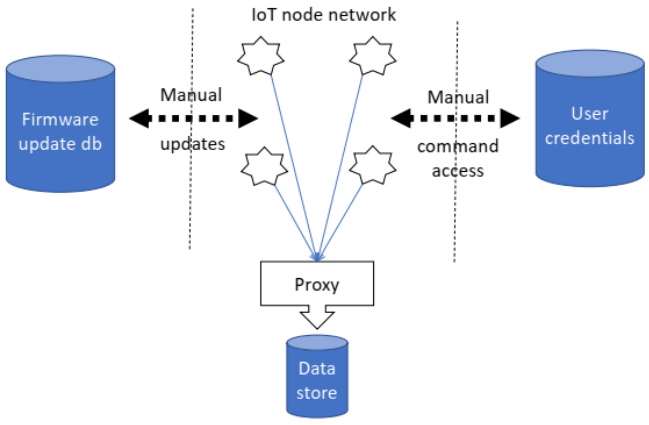

Figure 3. An improved traditional IoT architecture.

B. *Existing Distributed Ledger Based Solutions*

Applying the blockchain and distributed ledger technologies in various domains is currently a hot research and business development topic. These technologies have been proposed for many financial technology solutions with extensions assuring programmatic smart contracts, to preserve (and control) privacy and personal data, provide transparency on transactions, and in the industrial IoT to keep track of logistic chains. These are all very intriguing applications, but we concentrate on ones that are directly relevant to the distributed audit trail data. Further, we focus on forensic data in the cloud computing environment, since current IoT systems usually store generated data in the cloud and we consider this area to be among the most challenging problems for distributed ledgers.

The integrity of forensic data can be ensured by Public Key Infrastructure (PKI) signatures which depend on a certificate authority. This is not a feasible solution for IoT systems using distributed cloud resources since cloud infrastructure is inherently decentralized. An alternative to PKI signatures is keyless signatures implemented by a blockchain based distributed Keyless Signature Infrastructure [53] [54].

A blockchain based data provenance architecture, the ProvChain, is described and evaluated in [55]. ProvChain has been designed for collection and verification of cloud computing users' provenance data. ProvChain can use the global Bitcoin blockchain since the collected provenance data is restricted to metadata records of cloud service users' operations on data files stored in the cloud. Recorded metadata attributes are RecordID, Date and Time, UserID, Filename, AffectedUser, and FileOperation. A FileOperation is file creation, file modification, file copy, file share, or file delete. UserID attributes are hashed to protect cloud users' privacy. Provenance auditors can, therefore, access cloud users' provenance metadata but cannot correlate the metadata to users owning the metadata. Only the Cloud Service Provider (CSP) can relate provenance data to cloud service users owning the data. Provenance metadata records are published in blocks of a blockchain implemented by a blockchain network consisting of globally participating nodes. Several metadata records can be stored in one blockchain transaction. Each metadata record is extended with a hash and a Merkle hash tree [56] is constructed for the metadata records in a block. The Merkle root is stored as a block header attribute. ProvChain is built on the top of the open source cloud computing application ownCloud [57]. The Tierion Data API [58], is used to publish provenance metadata records in the blockchain. Tierion generates for each transaction a blockchain receipt based on the Chainpoint standard [59]. The Merkle hash tree included in this blockchain receipt proves that the provenance metadata records were recorded at a specific time. A provenance auditor can request a blockchain receipt via Tierion Data API, access the related blockchain block with Blockchain

Explorer [60], and validate the provenance metadata records in the block with the Merkle hash tree in the receipt. Measured ProvChain overhead for retrieval of provenance metadata of one file operation is about 0.7…0.8 s in an ownCloud test application [55].

Blockchain-based tamper-resistant registration of provenance data related to accessing medical data records in cloud storage is outlined in [61] [62]. The provenance data stored in the blockchain is available for auditing and in forensic investigations to detect privacy violations of medical data record owners. The outlined solution for protection of provenance data is applicable also to other types of personal data records.

C. Various Proposals for Distributed Ledger based Solutions

An ideal solution would be a global network of nodes fulfilling all three requirements in Section IV A. The global Bitcoin blockchain fulfils the two first requirements, but this blockchain cannot be extended with new blocks at a rate needed. Computationally it is not possible that even for a small cloud computing environment all the audit trail data for forensic investigations would be stored in the Bitcoin blockchain. The reason is the current blockchain size in combination with the throughput constrained Proof-of-Work (PoW) consensus algorithm.

However, other possible solutions may be engineered that circumvent this issue. One possible solution is a network of distributed ledger nodes, for example, blockchain nodes maintained by a CSP or preferably by several cooperating CSPs. As of the second requirement in Section IV A, all cloud computing users cannot be nodes in a distributed ledger network since also resource-constrained mobile devices and IoT devices can use cloud computing services. Moreover, a faster consensus algorithm than PoW must be implemented for the used distributed ledger.

Hashgraph is a DLT with a Byzantine consensus algorithm using a gossip protocol [63] [64]. While Bitcoins PoW implementation limits the throughput 7 transaction/s, the Hashgraph consensus algorithm can process even tens of thousands transactions/s [65]. The Archive Database proposed in [48] to be used as an immutable database for cloud audit trails could be implemented by a network of Hashgraph nodes maintained by a CSP or several cooperating CSPs. Each time when the database audit trail plugin stores log data the same data is transmitted to a preferably randomly chosen Hashgraph node. Reception of the log data creates a signed time-stamped event including a transaction storing the log data. An immutable record of all stored events is - due to the high event processing rate of a Hashgraph network – almost immediately available in each Hashgraph node. The Hashgraph fulfils all requirements in Section IV A. However, at the time of writing it is deployed in permissioned environments and is, therefore, a permissioned DLT. Still, a federated decentralized installation maintained by several cooperating CSPs or other service providers may offer an alternative to a public distributed ledger.

There are also other proposals that address the need for high throughput distributed ledgers. Off-chain state agreement solutions commonly referred to as state channel technology, have been developed for handling many small transactions. A use case for the development of state channel technology has been to handle micro-transactions, which in addition to needing a high throughput also require a minuscule transaction cost for the clearance of each transaction [66]. Other solutions propose to split the processing and recording of transactions into sub-chains, a technology often referred to as sharding [2].

D. Distributed IoT Architecture Proposal

In our distributed IoT architecture proposal, which is shown in Fig. 4, IoT nodes transmit their data to a distributed and replicated data store. The data store is run outside the limited nodes and utilise a Peer-to-Peer (P2P) protocol. Various data stores can be utilised that depending on requirements, such as scalability, speed, or post-processing, can be used. A suitable solution may be IPFS, a proprietary P2P data transfer protocol, or a data and analytics marketplace such as Streamr [67]. Smart contracts executed on top of a DLT implementation may authorize IoT devices and may further offer device management, e.g., issuing management commands. IoT device firmware updates may be automatized in a similar fashion. Storing the latest version of a binary update file in IPFS, and in a smart contract store an IPNS static address that allows node to query correct IPFS file and the firmware signature to confirm file integrity. This tells the IoT node how to access IPFS files and how to perform verification of the needed update.

We also consider the possibility important that a manufacturer may want to offer a service contract to any IoT node maintainer (owner). Currently, a significant problem is that IoT nodes are not provided with long-term support as the manufacturer often fails to get financial compensation for updating firmware once the product enters a maintenance/archival phase. This business model could however be implemented through a smart contract, that provides the manufacturer with a decentralised platform for selling firmware updates. An automated update function and contract resolution can be provided to any IoT node maintainer, either on a node basis (number of nodes) or on a network basis (maintaining organisation).

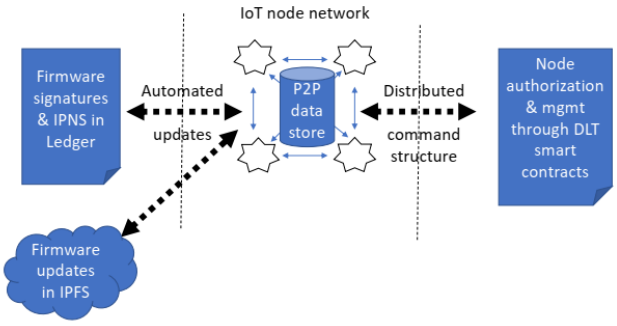

Figure 4. Proposed distributed IoT architecture.

V. DISCUSSION

A challenge for the field is that distributed ledger technology lacks a formal definition and standardisation. This may be due to the fact that it is an ensemble of technologies that in combination offer a mechanism for chaining blocks of records together. This holds the key for its disruptiveness, where centralised management of a system is impossible and the system starts living a life of its own with the help of computing resources allocated to it from any participant. A contemporary impact can be found in the financial industry due to an application, crypto currencies, where the traditionally regulated industry is disconnecting from the central governance of central banks. However, at the core DLT offer basic functionality for

- trustless interaction between two/more parties,
- third-party validation of transactions,
- distributed storage of transactions,
- some DLTs may offer a contract resolution mechanism through smart contracts.

Without a centralised authority, authenticating and validating the data is ever more important. This is fundamental for forensic evidence to hold up in a court of law. For this, DLT provides court-level forensics. The technology is developed in the wake of the financial industry with an obvious application domain being the IoT technology as this is, or will become, ubiquitous.

In the financial industry, the transitioning into cloud computing has inflicted a minimal transformation on the operational side, i.e., the cloud system do serve the end user as did centralised ones, but now in a manner scaling virtually infinitely. Yet, a cloud system runs the same databases, use storage space and encryption in the same way as would be done in a centralised system. Hence, the distributed ledger technology may enable, at time of writing, mainly for transaction storage space, independence from a centralised point of administration. Such an approach would obviously require a shared will among its participants. Comparing this with the financial industry, it may enable the creation of a global sharing economy of commodity swapping. This transformation would truly be disruptive on the global scale. Therefore, we believe DLT holds vast potential in catalysing new solutions and solving problems in existing applications.

VI. CONCLUSIONS

This paper outlines approaches for creation of audit trails from IoT systems using distributed cloud resources and for applying distributed ledger based solutions to securely store these audit trails. The security features of the distributed ledger assure the integrity of the audit trails which is essential for trustable IoT system forensics. The challenge is timely as the EU GDPR became enforced from May 2018. Moreover, the recent advancements in distributed ledgers, blockchains (cryptocurrencies) and their various spinoffs set the scene for applying this new technology by novel means. Implementation of hitherto proposed distributed ledger based solutions for protection of forensic audit trails of IoT systems using cloud resources is an important area of future research and development work. This paper lays the ground for future research into distributed ledger technology in terms of IoT system forensics.

ACKNOWLEDGMENT

This research was partly funded by the Fund for Technical Education and Research (TUF) of The Arcada Foundation and Academy of Finland project LibDat, decision number 309495.

APPENDIX

A. Distributed Ledger Technology

The most deployed distributed ledger type is a blockchain, which extends the shared database with a sequence of blocks storing transactional data. Blocks are chronologically and cryptographically linked to each another. Other distributed ledger types are the Tangle Network and Hashgraph. For the Tangle network, a Directed Acyclic graph-based network is used instead of a replicated linked chain of blocks in blockchain network nodes [68].

A Hashgraph network consists of nodes, which create context dependent events and communicate with each other using a gossip protocol. An event is a timestamped and digitally signed data structure consisting of one or several transactions and two hashes. One hash is extracted from the latest event on the node from which the latest gossip was received and the other hash is extracted from the preceding event created on the same node. A created event is sent as gossip to another randomly selected Hashgraph node together with all events still not known by the selected node. As event creation and gossip transmission continue in all Hashgraph nodes, all created events are immutably stored in each Hashgraph node. A Byzantine consensus on the order of events is achieved with probability 1 using a virtual voting procedure if more than 2n/3 nodes are uncorrupt where n is the number of nodes in the Hashgraph network. The details of the gossip protocol, the virtual voting, and the Byzantine consensus algorithm are presented in [69] and [64].

The blockchain technology is at the time of writing the best-known solution for implementing distributed ledgers and we, therefore, choose to focus on it. Findings concerning distributed ledgers, in general, should be transferable to other solutions such as the hashgraph and the Tangle network, once they become widely validated as secure.

Blockchain technology was introduced in 2008 as the Bitcoin cryptocurrency platform [70]. A blockchain implements a distributed database where a list of records called blocks is stored. New blocks can always be appended to the list but stored blocks are neither removed nor changed. The distributed database is replicated in nodes of a peer-to-peer blockchain network. A complete database copy is therefore stored in each node. The blockchain topology is a chain, since after the first block each additional block contains a hash link to the preceding block, see Fig. 5. The first block is called Genesis Block. Each block is also time stamped, however not necessarily to a universal time server.

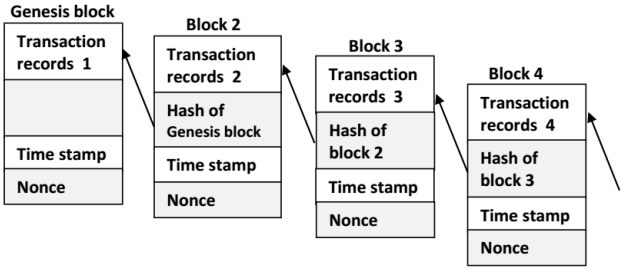

Figure 5. Basic blockchain structure.

A blockchain network node is owned by a blockchain user for execution of blockchain operations. A unique key pair of public key cryptography must also be owned by a blockchain user. The public key represents the identity of a blockchain user. A blockchain user executes a blockchain operation by initiating a transaction, which transfers some asset, for example, a cryptocurrency amount or a data object, to another blockchain user. A transaction creates a record, which is signed by the initiator of the transaction and transmitted to all nodes in the blockchain network. Each blockchain network node tries to validate a received transaction record with the transaction initiator's public key. A transaction record, which does not become validated by all blockchain network nodes, is discarded as invalid. Validated transaction records are collected by so-called mining nodes in the blockchain network and stored as lists in candidate blocks, which are time stamped. Each mining node executes a computation called mining on its candidate block. The candidate block of the mining node which first achieves a predefined mining goal is linked to the blockchain and all other mining nodes' candidate blocks are discarded. Several mining implementations for blockchains exist. Bitcoin blockchain mining uses PoW, where each mining node repeats hashing the concatenation of the last block in the blockchain and a new randomly chosen value. The mining goal is to create a hash of required difficulty.

There are public, permissioned, and private blockchains. A public blockchain, for example, Bitcoin, can be used by anyone. A public blockchain user copies the entire blockchain and installs the blockchain software on a personal node, which joins the blockchain network. Any blockchain user can also install the mining software on their own blockchain network node. Only a public blockchain can be trusted to fulfil the distributed ledger definition, as permission and private blockchains often maintain a centralized control node.

Recent blockchain implementations with extended functionality are denoted as Blockchain 2.0 for which an interesting feature is the smart contract introduced in [71]. A smart contract is a software component encompassing contractual terms and conditions enabling the verification, negotiation, or enforcement of a contract. A blockchain platform supporting smart contracts is Ethereum [72].

Blockchain security relies on the hash links between successive blocks combined with the replication of the entire blockchain to all blockchain network nodes. A public blockchain is therefore practically tamper-proof because a block cannot be changed without changing all the subsequent blocks and participation of all blockchain network nodes to validate and register the change. As the public blockchain is not managed by any centralized authority that could be a target of attacks it is less sensitive to some attack types such as DOS attacks, because full blockchain replicas are stored in many blockchain network nodes. However, an intrusion into a sufficient number of blockchain network nodes including some mining nodes can cause data losses and/or insertion of corrupt data in the attacked blockchain [73].

The tamper resistance of a blockchain does not exclude security vulnerabilities. Security attacks against blockchains are described and evaluated in [74] [75] [76] [77].

REFERENCES

[1] M. Neovius, M. Westerlund, J. Karlsson, and G. Pulkkis, "Providing Tamper-Resistant Audit Trails for Cloud Forensics with Distributed Ledger based Solutions," Proc. International Conference on Cloud Computing, IARIA, Feb. 2018, pp. 19-24, ISSN: 2308-4294, ISBN: 978-1-61208-607-1

[2] M. Westerlund and N. Kratzke, "Towards Distributed Clouds," Proc. 16th International Conference on High Performance Computing & Simulation (HPCS), IEEE Press, July 2018, pp. 655-663, doi:10.1109/HPCS.2018.00108.

[3] S. Alharbi, J. Weber-Jahnke, and I. Traore, "The Proactive and Reactive Digital Forensics Investigation Process: A Systematic Literature Review," International Journal of Security and Its Applications, vol. 5, no. 4, pp. 59-72, Oct. 2011.

[4] S. Garfinkel, "Anti-Forensics: Techniques, Detection and Countermeasures," Proc. 2nd International Conference on Information Warfare and Security, Mar. 2007, pp. 77-84.

[5] B. Schneier, and J. Kelsey, "Secure audit logs to support computer forensics," ACM Transactions on Information and System Security, vol. 2, iss. 2, pp. 159-176, May 1999, doi:10.1145/317087.317089.

[6] M. Aigner, "Security in the Internet of Things," in Cryptology and Information Security Series, vol. 4, Y. Li and J. Zhou, Eds. Amsterdam: IOS Press, pp. 109–124, 2010.

[7] S. Zawoad and R. Hasan, "FAIoT: Towards Building a Forensics Aware Eco System for the Internet of Things," Proc. 2015 IEEE International Conference on Services Computing (SCC 2015), IEEE Press, Aug. 2015, pp. 279–284, doi:10.1109/SCC.2015.46.

[8] EUR-Lex Regulation [EU] 2016/679. *General Data Protection Regulation (GDPR)*. [Online]. Available from: http://eur-lex.europa.eu/eli/reg/2016/679/oj 2018.11.26

[9] M. Westerlund, "A study of EU data protection regulation and appropriate security for digital services and platforms," Doctoral Dissertation, Åbo Akademi University, Åbo, Finland, 2018. [Online]. Available from: http://urn.fi/URN:ISBN:978-952-12-3694-5 2018.11.26

[10] B. Duncan, "Can EU General Data Protection Regulation Compliance be Achieved When Using Cloud Computing?," Proc. Ninth International Conference on Cloud Computing, GRIDs, and Virtualization, IARIA, Feb. 2018, pp. 1–6, ISSN: 2308-4294, ISBN: 978-1-61208-607-1

[11] S.1691 - Internet of Things (IoT) Cybersecurity Improvement Act of 2017. Bill, Senate - Homeland Security and Governmental Affairs, USA, 2017. [Online]. Available from https://www.congress.gov/bill/115th-congress/senate-bill/1691 2018.11.26

[12] E. Oriwoh, D. Jazani, G. Epiphaniou, and P. Sant, "Internet of Things Forensics: Challenges and Approaches," Proc. 9th

IEEE International Conference on Collaborative Computing: Networking, Applications and Worksharing, IEEE Press, Oct. 2013, pp. 608-615, doi:10.4108/icst.collaboratecom.2013.254159.

[13] U. Salama, "Smart Forensics for the Internet of Things (IoT)," 2017. [Online]. Available from: https://securityintelligence.com/smart-forensics-for-the-internet-of-things-iot 2018.11.26

[14] S. Alabdulsalam, K. Schaefer, T. Kechadi, and N.-A. LeKhac, "Internet of Things Forensics: Challenge and Case Study," arXiv:1801.10391v1 [cs.CR], 2018. [Online]. Available from: https://arxiv.org/abs/1801.10391 2018.11.26

[15] N. Zulkipli, A. Alenezi, and G. Wills,"IoT Forensic: Bridging the Challenges in Digital Forensic and the Internet of Things," Proc. 2nd International Conference on Internet of Things, Big Data and Security (IoTBDS), vol. 1, SciTePress, 2017, pp. 315-324, doi:10.5220/0006308703150324.

[16] T. Zia, P. Liu, and W. Han, "Application-Specific Digital Forensics Investigative Model in Internet of Things (IoT)," Proc. 12th International Conference on Availability, Reliability and Security (ARES'17), ACM Press, 2017, pp. 55.1-55.7, doi:10.1145/3098954.3104052.

[17] S. Perumal, N. Norwawi, and V. Raman, "Internet of Things (IoT) Digital Forensic Investigation Model: Top-down Forensic Approach Methodology," Proc. 5th International Conference on Digital Information Processing and Communications (ICDIPC), IEEE Press, Nov. 2015, pp. 19-23, doi:10.1109/ICDIPC.2015.7323000.

[18] V. R. Kebande and I. Ray, "A Generic Digital Forensic Investigation Framework for Internet of Things (IoT)," Proc. 4th International Conference on Future Internet of Things and Cloud (FiCloud), IEEE Press, 2016, pp. 356-362, doi:10.1109/FiCloud.2016.57.

[19] HDFS Architecture Guide. 2013. [Online]. Available from: https://hadoop.apache.org/docs/r1.2.1/hdfs_design.html 2018.11.26

[20] K. Muniswamy-Reddy, D. Holland, U. Braun, and M. Seltzer, "Provenance-aware storage systems," Proc. 2006 USENIX Annual Technical Conference, USENIX Association, 2006, pp. 43–56.

[21] R. Hasan, R. Sion, and M. Winslett, "The case of the fake Picasso: Preventing history forgery with secure provenance," Proc. 7th USENIX Conference on File and Storage Technologies (FAST´09), USENIX Association, 2009, pp. 1–12.

[22] R. H. Fielding. "Architectural Styles and the Design of Network-based Software Architectures," Doctoral Dissertation, University of California, Irvine, USA, 2000. [Online]. Available from: https://www.ics.uci.edu/~fielding/pubs/dissertation/top.htm 2018.11.26

[23] The Top 10 IoT Segments in 2018 – based on 1,600 real IoT projects, IoT Analytics, 2018. [Online]. Available from https://iot-analytics.com/top-10-iot-segments-2018-real-iot-projects/ 2018.11.26

[24] D. A. B. Fernandes, L. F. B. Soares, J. V. Gomes, M. M. Freire, and P. R. M. Inácio, "Security issues in cloud environments: a survey," International Journal of Information Security, vol. 13, iss. 2, pp. 113-170, Apr. 2014, doi:10.1007/s10207-013-0208-7.

[25] P. Mell and T. Grance, "The NIST Definition of Cloud Computing," Special Publication 800-145, National Institute of Standards and Technology, U.S. Dept. Commerce, 2011. [Online]. Available from: https://csrc.nist.gov/publications/detail/sp/800-145/final 2018.11.26

[26] J. Köhler, K. Jünemann, and H. Hartenstein, "Confidential database-as-a-service approaches: taxonomy and survey," J. Cloud Computing: Advances, Systems and Applications, vol. 4, no. 1, 2015, doi:10.1186/s13677-014-0025-1.

[27] N. Dragoni, S. Giallorenzo, A. L. Lafuente, M. Mazzara, F. Montesi, R. Mustafin, and L. Safina, "Microservices: yesterday, today, and tomorrow," April 2017. [Online]. Available from: https://arxiv.org/pdf/1606.04036.pdf 2018.11.26

[28] N. Kratzke and P.-C. Quint, "Understanding cloud-native applications after 10 years of cloud computing - A systematic mapping study," J. Systems and Software, vol. 126, pp. 1-16, April 2017, doi:10.1016/j.jss.2017.01.001.

[29] D. Bernstein, E. Ludvigson, K. Sankar, S Diamond, and M. Morrow, "Blueprint for the Intercloud - Protocols and Formats for Cloud Computing Interoperability," Proc. Fourth International Conference on Internet and Web Applications and Services (ICIW'09), IEEE Press, June 2009, pp.328-336, doi:10.1109/ICIW.2009.55.

[30] M. Shiraz, A. Gani, R. H. Khokhar, and R. Buyya, "A review on distributed application processing frameworks in smart mobile devices for mobile cloud computing," IEEE Communications Surveys and Tutorials, vol. 15, no. 3, pp. 1294–1313, 2013, doi:10.1109/SURV.2012.111412.00045.

[31] L. Jiang et al., "An IoT-Oriented Data Storage Framework in Cloud Computing Platform," IEEE Transactions on Industrial Informatics, vol. 10, no. 2, pp. 1443-1451, May 2014, doi:10.1109/TII.2014.2306384.

[32] K. Ruan, J. Carthy, T. Kechadi, and M. Crosbie, "Cloud Forensics," in Advances in Digital Forensics VII. DigitalForensics 2011. IFIP Advances in Information and Communication Technology, vol 361, G. Peterson and S. Shenoi, Eds. Berlin, Heidelberg: Springer, pp. 35–46, 2011.

[33] NIST Cloud Computing Forensic Science Challenges, Draft NISTIR 8006, National Institute of Standards and Technology, U.S. Department of Commerce, June 2014. [Online]. Available from: https://csrc.nist.gov/csrc/media/publications/nistir/8006/draft/documents/draft_nistir_8006.pdf 2018.11.26

[34] K. Grahn, M. Westerlund, and G. Pulkkis, "Analytics for Network Security: A Survey and Taxonomy," in Information Fusion for Cyber-Security Analytics. Studies in Computational Intelligence, vol. 691, I. Alsmadi, G. Karabatis, and A. Aleroud, Eds. Springer, Cham, pp. 175-193, 2017.

[35] B. Martini and K.-K. R. Choo, "An integrated conceptual digital forensic framework for cloud computing," Digital Investigation, vol. 9, iss. 2, pp. 71–80, Nov. 2012, doi:10.1016/j.diin.2012.07.001.

[36] V. M. Katilu, V. N. L. Franqueira, and O. Angelopoulou, "Challenges of Data Provenance for Cloud Forensic Investigations," Proc. 10th Int. Conf. on Availability, Reliability and Security, IEEE Press, Aug. 2015, pp. 312-317, doi:10.1109/ARES.2015.54.

[37] M. E. Alex and R. Kishore, "Forensics Framework for Cloud computing," J. Computers and Electrical Engineering, vol. 60, iss. C, pp. 193-205, May 2017, doi:10.1016/j.compeleceng.2017.02.006.

[38] K.-K. Muniswamy-Reddy and M. Seltzer, "Provenance as first class cloud data," ACM SIGOPS Operating Systems Review, vol. 43, no. 4, pp. 11-16, Jan. 2009, doi:10.1145/1713254.1713258.

[39] Y. L. Simmhan, B. Plale, and D. Gannon, "A survey of data provenance in e-science," ACM Sigmod Record, vol. 34, no. 3, pp. 31–36, Sep. 2005, doi:10.1145/1084805.1084812.

[40] D. K. Tosh et al., "Security Implications of Blockchain Cloud with Analysis of Block Withholding Attack," Proc. 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, IEEE Press, May 2017, pp. 458-467, doi:10.1109/CCGRID.2017.111.

[41] Extract & decode, 2018. [Online]. Available from: https://www.cellebrite.com/en/product/solutions/extract-decode/ 2018.11.26

[42] A. Boztas, A. R. J. Riethoven, and M. Roeloffs, “Smart TV forensics: Digital traces on televisions,” Digital Investigation, vol 12, supp. 1, pp. S72-S80, Mar. 2015, doi:10.1016/j.diin.2015.01.012.

[43] The NFI Memory Toolkit II. Netherlands Forensic Institute, 2011. [Online]. Available from: https://www.forensicinstitute.nl/documents/publications/2017/03/06/brochure-memory-toolkit 2018.11.26

[44] Z-Wave Alliance. *About Z-Wave Technology*. [Online]. Available from: https://z-wavealliance.org/about_z-wave_technology 2018.11.26

[45] A. C. Shin, P. Chandok, R. Liu, S. J. Nielson, and T. R. Leschke, “Potential Forensic Analysis of IoT Data: An Overview of the State-of-the-Art and Future Possibilities,” Proc. 2017 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData), IEEE Press, June 2017, pp. 705-710, doi:10.1109/iThings-GreenCom-CPSCom-SmartData.2017.182.

[46] B. Lee, A. Awad, and M. Awad, “Towards secure provenance in the cloud: A survey,” Proc. 8th International Conference on Utility and Cloud Computing (UCC), IEEE Press, Dec. 2015, pp. 577–582, doi:10.1109/UCC.2015.102.

[47] Cognitect, Inc. *Datomic Cloud. A transactional database with a flexible data model, elastic scaling, and rich queries*. [Online]. Available from: http://www.datomic.com/ 2018.11.26

[48] B. Duncan and M. Whittington, “Creating an Immutable Database for Secure Cloud Audit Trail and System Logging,” Eighth International Conference on Cloud Computing, GRIDs, and Virtualization, Athens: IARIA, Feb. 2017, pp. 54–59, ISSN: 2308-4294, ISBN: 978-1-61208-529-6

[49] J. Benet, “IPFS – Content Addressed, Versioned, P2P File System (DRAFT 3),” 2017. [Online]. Available from: https://github.com/ipfs/ipfs/blob/master/papers/ipfs-cap2pfs/ipfs-p2p-file-system.pdf 2018.11.26

[50] IPFS is the Distributed Web, 2018. [Online]. Available from: https://github.com/ipfs/ipfs/blob/master/README.md 2018.11.26

[51] Distributed Ledger Technology: beyond blockchain, 2016. [Online]. Available from: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/492972/gs-16-1-distributed-ledger-technology.pdf 2018.11.26

[52] D. Mills, K. Wang, B. Malone, A. Ravi, J. Marquardt, C. Chen, A. Badev, T. Brezinski, L. Fahy, K. Liao, V. Kargenian, M. Ellithorpe, W. Ng, and M. Baird, “Distributed ledger technology in payments, clearing, and settlement,” Finance and Economics Discussion Series 2016-095, Washington: Board of Governors of the Federal Reserve System, doi:10.17016/FEDS.2016.095.

[53] A. Buldas, A. Kroonmaa, and R. Laanoja, “Keyless Signatures´ Infrastructure: How to Build Global Distributed Hash-Trees,” in Secure IT Systems. NordSec 2013. Lecture Notes in Computer Science, vol. 8208, R. Nielson and D. Gollmann, Eds. Berlin, Heidelberg: Springer, pp. 313–320, 2013.

[54] Guardtime. *Cloud Assurance with Blockchains*, 2017. [Online]. Available from: https://guardtime.com/solutions/cloud 2018.11.26

[55] X. Liang, et al., “ProvChain: A Blockchain-based Data Provenance Architecture in Cloud Environment with Enhanced Privacy and Availability,” Proc. 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, IEEE Press, May 2017, pp. 468-477, doi:10.1109/CCGRID.2017.8.

[56] R. C. Merkle, "A Digital Signature Based on a Conventional Encryption Function," in Advances in Cryptology – CRYPTO '87, C. Pomerance, Ed. Berlin, Heidelberg: Springer, pp. 369-378, 1988.

[57] ownCloud, 2017. [Online]. Available from: https://owncloud.org/ 2018.11.26

[58] Tierion Documentation, 2017. [Online]. Available from: https://tierion.com/docs 2018.11.26

[59] Chainpoint, 2017. [Online]. Available from: https://chainpoint.org/ 2018.11.26

[60] BTC.com, 2017. [Online]. Available from: https://btc.com/

[61] Q. Xia, E. B. Sifah, A. Smahi, S. Amofa, and X. Zhang, ‘‘BBDS: Blockchain-Based Data Sharing for Electronic Medical Records in Cloud Environments,’’ Information 2017, vol. 8, iss. 2, pp. 1-16, Apr. 2017, doi:10.3390/info8020044.

[62] Q. Xia, E. B. Sifah, K. O. Asamoah, J. Gao, X. Du, and M. Guizani, “MeDShare: Trust-Less Medical Data Sharing among Cloud Service Providers via Blockchain,” IEEE Access, vol. 5, pp. 14757-14767, July 2017, doi:10.1109/ACCESS.2017.2730843.

[63] G. Kingslay, “Hashgraph vs. Blockchain Is the end of Bitcoin and Ethereum near?” [Online]. Available from: https://coincodex.com/article/1151/hashgraph-vs-blockchain-is-the-end-of-bitcoin-and-ethereum-near/ 2018.11.26

[64] L. Baird, “The Swirlds Hashgraph Consensus Algorithm: Fair, Fast, Byzantine Fault Tolerance,” Swirlds Tech Report Swirlds-TR-2016-01, May 31, 2016. [Online]. Available from: http://www.swirlds.com/downloads/SWIRLDS-TR-2016-01.pdf 2018.11.26

[65] Digital Bazaar, Inc. *Blockchain Technologies Feature Analysis*, 2016. [Online]. Available from: https://lists.w3.org/Archives/Public/public-blockchain/2016Oct/att-0004/BlockchainTechnologiesFeatureAnalysis.html 2018.11.26

[66] Z. Hess, Y. Malahov, and J. Pettersson, “Æternity blockchain”, 2017. [Online]. Available from: https://blockchain.aeternity.com/aeternity-blockchain-whitepaper.pdf 2018.11.26

[67] The Streamr Platform. 2018. [Online]. Available from: https://www.streamr.com/#streamrSystem 2018.11.26

[68] S. Popov, “The Tangle,” White Paper, 2017. [Online]. Available from: https://iota.org/IOTA_Whitepaper.pdf 2018.11.26

[69] L. Baird, “Hashgraph Consensus: Detailed Examples,” Swirlds Tech Report Swirlds-TR-2016-02, Dec 11, 2016. [Online]. Available from: http://www.swirlds.com/downloads/SWIRLDS-TR-2016-02.pdf 2018.11.26

[70] S. Nakamoto, “Bitcoin: A Peer-to-Peer Electronic Cash System,” 2008. [Online]. Available from: https://bitcoin.org/bitcoin.pdf 2018.11.26

[71] N. Szabo, “The Idea of Smart Contracts,” 1997. [Online]. Available from: http://www.fon.hum.uva.nl/rob/Courses/InformationInSpeech/CDROM/Literature/LOTwinterschool2006/szabo.best.vwh.net/idea.html 2018.11.26

[72] Ethereum Blockchain App Platform. [Online]. Available from: https://www.ethereum.org/ 2018.11.26

[73] M. Conoscenti, A. Vetro, J. C. de Martin, “Blockchain for the Internet of Things: a Systematic Literature Review,” Proc. 13th International Conference on Computer Systems and Applications (AICCSA), IEEE Press, Dec. 2016, pp. 1-6, doi:10.1109/AICCSA.2016.7945805.

[74] Eyal, I and E. G. Sirer, "Majority is not Enough: Bitcoin Mining is Vulnerable," arXiv:1311.0243v5 [cs.CR], Nov. 2013. [Online]. Available from: https://arxiv.org/pdf/1311.0243v5.pdf 2018.11.26

[75] S. Barber, X. Boyen, E. Shi, and E. Uzun, "Bitter to Better – How to Make Bitcoin a Better Currency," in Financial Cryptography and Data Security, A. D. Keromytis, Ed. Berlin, Heidelberg: Springer, pp. 399-414, 2012.

[76] E. Heilman, A. Kendler, A. Zohar, and S. Goldberg, "Eclipse Attacks on Bitcoin's Peer-to-Peer Network," Proc. 24th USENIX Security Symposium, Washington: USENIX Association, 2015, pp. 129-144, ISBN: 978-1-931971-232

[77] K. Nayak, S. Kumar, A. Miller, and E. Shi, "Stubborn mining: Generalizing selfish mining and combining with an eclipse attack," Proc. 2016 IEEE European Symposium on Security and Privacy (EuroS&P), IEEE Press, Mar. 2016, pp. 305-320, doi:10.1109/EuroSP.2016.32.

Safety, Cybersecurity and Interoperability aspects in Modern Nuclear Power Plants

Asmaa Tellabi[1, 4], Ines Ben Zid[2, 4], Edita Bajramovic[3, 4], Karl Waedt[4]
[1]University of Siegen, Siegen, Germany
[2]University of Bielefeld, Bielefeld, Germany
[3]Friedrich-Alexander-University Erlangen-Nuremberg, Erlangen, Germany
[4]Framatome GmbH, Erlangen, Germany

E-mail:{firstname.lastname}@framatome.com

Abstract—**The integration of digital equipment and diverse automation platforms in modern nuclear plants, including Nuclear Power Plants is due to the gradually increasing use of digital technologies. This digitalization either comes gradually based on a succession of refurbishment projects of Instrumentation & Control and Electrical Power Systems or as comprehensive architectures with new-built power plants. Therefore, similar to any critical infrastructure facing a growing risk of cyber-attacks, cybersecurity for Nuclear Power Plants has become a subject of rising concern. We envision that the findings in this paper provide a relevant understanding of the threat landscape facing digital systems in nuclear power plants. The knowledge can be used for an improved understanding and a better identification of security risks during the analysis and design of supporting systems. This paper gives an overview of the security issues and vulnerabilities, helping to better understand the big picture of cybersecurity issues and vulnerabilities in Nuclear Power Plants. Identifying these vulnerabilities and issues helps to establish new security countermeasures. A new draft standard IEC 63096 is presented in this paper as well.**

Keywords-nuclear power plants; cybersecurity interoperability.

I. INTRODUCTION

Digital Instrumentation and Control (I&C) systems are defined as computer-based devices that monitor and control nuclear power plants (NPP). Electrical Power Systems (EPS) provide the redundant power supply for different plant operation scenarios, which have to be fully supported. The EPS may include the connection to external highest voltage (e.g., 400 kW) or high voltage (e.g., 110 kV) grid connections, Emergency Diesel Generators, Station Blackout Diesel Generators, different Uninterruptable Power Supplies (UPS), e.g., for 2 hours and 12 hours [1][2].

Furthermore, different inverters and rectifiers are responsible of controlling and monitoring the entire aspects of the plant's health, all plant states and helping to respond with the care and adjustments as needed. They are seen as the nervous system of NPP. Generation III+ and IV reactors are equipped with digital I&C systems, while analog systems in older reactors are being replaced with digital systems [2]. The high level communication between NPP control networks is done by Supervisory Control and Data Acquisition systems (SCADA) in order to coordinate power production with transmission and distribution demands. Integration of digital I&C systems and the connectivity between NPP control networks and external networks represent a threat for NPP, making them a target to cyber-attacks which can include physical damage to reactors. With possibilities of cyber-attacks targeting NPP increasingly, cybersecurity has aroused as a significant problem [3].

The remainder of this paper is organized as follow. Section II gives background information on typical system architecture in NPP. Section III outlines some of the notorious publically known cyber-attacks against NPP. In Section IV , a new IEC 63096 standard [4] is described. We conclude the paper in Section V.

II. NUCLEAR POWER PLANTS

A. NPP architecture

The general digital systems configuration of NPP is almost similar to that of Industrial Control Systems (ICS) SCADA systems. The general architecture can be separated into two distinct domains: I&C systems, EPS and plant-local or corporate IT systems. The restriction on these networks is not similar, but also the nature of the traffic.

According to Fig. 1, operations, such as office automation, document management, and email, which consist of conventional IT systems, such as PCs and enterprise workstations use the corporate network of the Utility. As an illustration, Internet access, FTP, email, and remote access will normally be allowed on the enterprise network level but should not be permitted on the ICS network level.

Nuclear safety is the accomplishment of correct operating conditions, prevention of accidents or alleviation of accident consequences, ending up with the protection of workers, the public and the environment from extreme radiation hazards. On the other hand, nuclear security is the prevention and detection of, and response to, theft, sabotage, unauthorized access, illegal transfer or other malicious acts involving nuclear material, other radioactive substances or their associated facilities.

Safety is expected to prevent accidents, while security is implemented to stop intended acts that might harm the NPP or lead to the theft of nuclear materials. Safety evaluations focus on risks arising from accidental events occurrences

originated from nature (such as earthquakes, tornadoes, or flooding), hardware failures, supplementary internal events or interruptions (such as fire, pipe breakage, or loss of electric power supply), or human mistakes (such as the incorrect application of procedures, or incorrect alignment of circuits). For security, risks, or events, worried about result from malicious acts accomplished with the objective to steal material or to cause damage. Therefore, security events are based on 'intelligent' or 'deliberate' actions achieved intentionally for theft or sabotage and with the purpose to avoid protective measures [1] [3].

Safety and security have various elements in common and both focus on protecting the plant with the eventual purpose of protecting people, society, and the environment. As stated above, the essential objective of each is identical — the protection of people, society and the environment. Whether it was a safety or a security event causing harm, the acceptable risk is likely the same, usually they both adopt the strategy of defense in depth, which is defined as the usage of layers of protection.

First concern is given to prevention. Second, abnormal situations need to be identified early and take action promptly to avoid resulting damage. Mitigation comes in the third place of an operative strategy. Finally, considerable emergency planning should be implemented in case of the failure of prevention, protection and mitigation systems [5].

I&C are censorious in NPP. They are responsible of monitoring the operational state of the nuclear reactors through interaction with physical equipment, but also in charge of process control. With the introduction of digital technologies in the 2000s, I&C systems shifted from analog technologies to digital technologies. The usage of digital technologies has been steadily increasing. NPP I&C systems engage in environments that are different than those of typical IT systems.

In a typical NPP, I&C architecture contains two types of systems: Non-safety and Safety systems. The Non-safety system is defined as a distributed computer system containing a number of remote control nodes spread across the NPP, which uses redundant real time data network to communicate with each other and with the Human Machine Interface (HMI) [6].

Communication with third party systems and Operation Maintenance Corporate Systems (OMS) are also supported through open protocols like Object Embedding Linking Process Control, fieldbuses and Modbus-TCP [7].

Additionally, monitoring and manual control of the NPP processes is done by the use of HMI consoles connected in the non-safety system. In order to display critical information related to safety on the non-safety HMI, the safety system will communicate with the non-safety system through Interface gateways.

On the contrary, a safety system is regularly based on a channelized Programmable Logic Controllers (PLC) that holds a number of PLC nodes distributed across the NPP. These PLC and its cabinets are designed to resist seismic events, environmental events and cybersecurity attacks. Furthermore, they can still be able to operate safely.

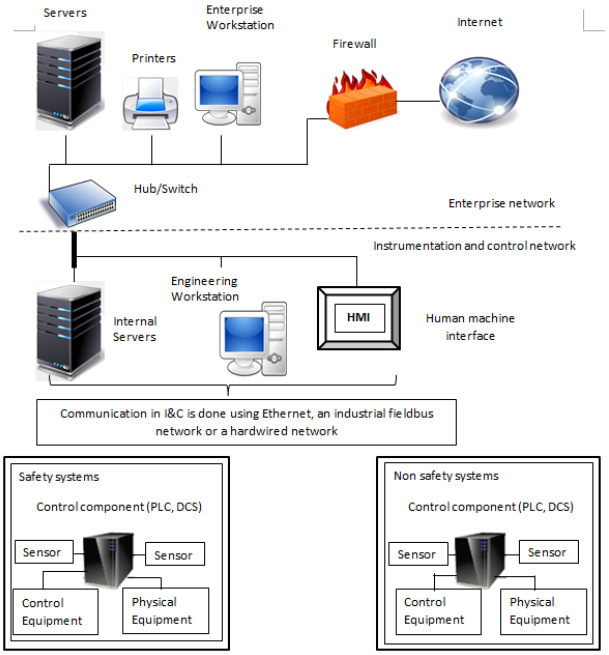

Figure 1. General architecture in nuclear power plants [2].

The purpose of this distribution is to coordinate with safety components in the process system, and also to ensure a safe communication in a safety channel using the redundant real time data safety network or through dedicated high speed links in between safety channels. Distributed control systems (DCSs) or PLC are common control components in I&C systems, they interact with physical equipment directly and industrial PCs or engineering workstations that are employed to configure control components and their related works [1].

B. ICS vs. IT systems

I&C systems are used to control the physical world, while IT systems´ purpose is to manage data. Requirements for performance and reliability, operating systems used and applications employed for I&C systems may be considered uncommon in a typical IT network environment [5].

At first, Industrial control systems (ICS) were similar to IT systems to some extent, in a way where ICS were inaccessible systems running on proprietary control protocols, and applying special hardware and software. Easily accessible, low-cost Ethernet and Internet Protocol (IP) devices are now taking the place of the majority of proprietary technologies; as a result cybersecurity vulnerabilities and incidents are increasing. Nowadays, the deployment of IT solutions in ICS is made to validate the use of business connectivity and remote access abilities, created and implemented to control typical industry computers, operating systems (OS) and network protocols. This combination of distinct IT capabilities provides considerably less separation for ICS from the outside world than previous systems, making security an essential requirement for these systems. These security solutions´ objectives were to handle security concerns in traditional IT systems; considerable

safety measures must be taken when introducing these same solutions to ICS environments. Environments in which ICS and IT systems operate are constantly changing, operation environments comprise, but are not limited to [5]:

- The threat space; vulnerabilities; missions/business activities; mission/business processes; enterprise and information security architectures; information technologies; personnel; facilities; supply chain relationships; organizational governance/culture; procurement/acquisition processes; organizational policies/procedures; organizational assumptions, constraints, risk tolerance, and priorities/trade-offs).

The following lists some special considerations when addressing security for ICS [5][6]:

1) Timeliness and Performance Requirements

Usually, ICS are considered time-critical, with a tolerable margin of delay and jitter, which depends on the application. Deterministic and reliable response are mandatory for some systems, e.g., for closed loop control. For IT systems, high throughput is necessary, while this it is not considered critical for ICS. In some cases, e.g., a reactor protection I&C system, automated system response in real time and timely response to human interaction is seen critical, e.g., for display systems in a main control room. Real-time operating systems (RTOS) or embedded real-time micro-kernels are implemented in ICS, where real-time responses are required.

2) Availability Requirements

In general, ICS processes are continuous, meaning that sudden interruptions of systems that control industrial processes are not allowed. An advanced schedule of these interrupts must be done. Sometimes, the production is considered more vital than the information, which can be undesirably affected by stopping and/or restarting ICS. In case traditional IT strategies are used, e.g., rebooting a module, they will have a negative effect on high availability requirements, reliability and maintainability of the ICS. In some industries, redundant components running in parallel are deployed to provide continuity when some components are unreachable.

3) Risk Management Requirements

Confidentiality and integrity are normally the principal concerns for IT systems. On the other hand, for ICS systems the main concerns are: availability, integrity, human safety and fault tolerance, regulatory compliance, destruction of equipment, loss of intellectual property, theft or damaged products. Safety and security concepts are paired; staffs in charge of the operation, security, and maintenance of ICS must understand those essential concepts. Security measures that jeopardy safeties are not allowed.

4) Physical Effects

ICS field devices, e.g., PLC, control physical processes. Interactions between ICS and physical processes can be very difficult, and can lead to severe consequences that can be noticeable in physical events.

5) System Operation

Generally ICS environments, counting operating systems (OS) and control networks, are completely different from IT systems, necessitating specific skill sets, experience, and levels of expertise. Usually, industrial control networks are managed by control engineers, and not by IT personnel.

6) Communications

In ICS environments, communication protocols and media needed by field device control and intra-processor communication are very different from nearly every IT environment.

7) *Patch Management*

Preserving the integrity of both IT and control systems is required. For IT systems software updates as well as security patches, are normally executed in a specific time based on appropriate security policy and procedures. On the other hand, software updates on ICS cannot always be forced on a timely basis without negatively affecting the system. Moreover, these procedures are usually automated via server-based tools. Before their implementation, these updates need to be tested by both the vendor and the end user. Also, a schedule of days/weeks must be planned by the ICS owner in advance. Patch management is also associated to hardware and firmware, the process demands careful assessment by ICS experts, e.g., control engineers, working in partnership with security and IT personnel.

8) Component Lifetime

IT components´ lifetime is in the order of 3 to 5 years, with briefness due to the fast progress of technology. For ICS, the implemented technology has been designed for precise use cases and implementation; the lifetime of the proposed solution is often in the order of 10 to 25 years and sometimes longer.

9) Component Location

Some IT modules are physically reachable by local transportation, also placed in corporate and commercial facilities. Remote locations may be used for backup services. Contrariwise, distributed ICS components must be isolated, remote services should not be allowed or used when required only by approved persons. Also, modules´ location necessitates important physical and environmental security measures.

III. CYBERSECURITY AND CYBER WARFARE RELATED TO NUCLEAR POWER PLANTS

Advancement in electronics and IT was the main motivation behind the replacement of traditional analog I&C systems in NPP with I&C systems, e.g., systems based on computers and microprocessors. Also, digital systems allow superior reliability, improved plant performance and supplementary diagnostic aptitudes. The systems used today were designed to satisfy performance, reliability, safety, and flexibility requirements, most of them were created a long time ago before new technologies became a crucial part of business operations.

In most typical implementations, these systems are physically isolated from outside networks and are based on

proprietary hardware and software. Communication protocols include basic error detection and correction capabilities but lack secure systems [7]. Accordingly, it is crucial not to connect such systems to an Intranet or the Internet.

A. History of Selected Attacks in NPP

First, in this Section we present some of the notorious attacks against NPP. In [8], attack taxonomy is defined by 5 dimensions: precondition, vulnerability, target, attack method, effect of the attack. It was combined with a new dimension target—the effect it has on the confidentiality, availability, integrity (CIA) of a system.

1) Ignalina NPP (1992)

At the Ignalina NPP in Lithuania, a technician intentionally introduced a virus into the industrial control system.

- **Precondition:** Direct access to the system.
- **Attack method:** Insider attack.
- **Target:** Availability and integrity.
- **Effect of the attack:** In this case, little harm was caused, but someone with malicious intent could have provoked a serious incident [9][10].

2) Davis-Besse NPP (2003)

This plant located in Ohio was infected by the Slammer worm (also called W32/SQLSlam-A or Sapphire).

- **Precondition:** Unpatched system.
- **Attack method:** At first, the worm scans and sends itself to random IP addresses; if worm reaches a machine that is running Microsoft SQL 2000, it infects that machine and begins scanning and sending itself to another machine.
- **Target:** Availability.
- **Effect of the attack:** The safety parameter display system (SPDS), responsible of collecting and displaying data regarding the reactor core from the coolant systems, temperature sensors and radiation detectors, was unavailable for nearly five hours [9][10].

3) Browns Ferry NPP (2006)

This NPP located in Alabama experienced a malfunction of both reactor recirculation pumps (which use variable-frequency drives to control motor speed and are needed to cool the reactor) and the condensate demineralizer controller (a type of PLC).

- **Precondition:** Device failure, attack method. Both of these devices contain microprocessors that communicate by sending and receiving data over an Ethernet network.
- **Attack method:** Ethernet operates by first sending data to every device on the network; then they have to inspect each packet to define if the packet is intended for them or if they can ignore it, making them vulnerable to failure if they accept enormous traffic.
- **Target:** Availability.
- **Effect of the attack:** The excess traffic produced by network broke down the reactor recirculation pumps and condensate demineralizer controller. As a consequence, the plant's Unit 3 had to be manually shut down in order to prevent a meltdown [9][10].

4) Hatch NPP (2008)

Hatch NPP located in Georgia experienced a shutdown as an unintended consequence of an update performed by contractor. An engineer contractor that manages the plant's technology operations installed an update to a computer on the plant's business network.

- **Precondition:** Human error.
- **Attack method:** The update was intended to synchronize data. The updated computer was connected to one of the plant's industrial control system networks, consequently when the engineer restarted the updated computer; the synchronization changed the control system's data to zero for a short moment.
- **Target:** Availability and integrity.
- **Effect of the attack:** The interpretation of the temporary changed values by the plant's safety system was incorrect. The updated value to zero of the water level signified that there was not enough water to cool the reactor core, which conducted to automatic shutdown for 48 hours of the plant's Unit 2 [9][10].

5) Natanz Nuclear Facility and Bushehr NPP – Stuxnet (2010)

First exposed to public in June 2010, the Stuxnet computer worm infected both the Natanz nuclear facility and the Bushehr NPP in Iran, partially destroying around 1,000 centrifuges at Natanz.

- **Precondition:** Use of commercial-off-the-shelf (COTS) Operating System (OS), Stuxnet infects computers using the Microsoft Windows OS, exploiting vulnerabilities in the system that allows it to obtain system-level access.
- **Attack method:** The worm uses forged certificates as a result the installed files look to come from an authentic source, misleading antivirus. Iranian nuclear facilities work with Siemens Step 7 SCADA system. Once the machine is infected, Stuxnet inspects the network to find computers attached to a similar system. Stuxnet duplicates itself on other computers by exploiting another set of vulnerabilities found in print spoolers and also through USB flash drives, so it spreads to networks using shared printers. Stuxnet's payload is activated only if the computer is connected to a similar Siemens system. It reprograms the system's PLC, in charge of controlling centrifuges applied in enriching nuclear fuel, so that they spin rapidly and eventually finish by break down.
- **Target:** Availability and integrity.
- **Effect of the attack:** As a result, Stuxnet destroyed over 1000 centrifuges at Natanz [9][10].

6) Korea Hydro and Nuclear Power Co. Commercial Network (2014)

Hackers infiltrated and stole data from the commercial network of Korea Hydro and Nuclear Power Co., which operates 23 of South Korea's nuclear reactors.

- **Precondition:** Human error: Access to the confidential data was obtained by hackers through phishing emails to the owner-operator's employees. Some of them finished by clicked on the links and downloaded the malware.
- **Attack method:** Sending phishing emails to employees.
- **Target:** Confidentiality.
- **Effect of the attack:** The hackers acquired the blueprints and manuals of two reactors, electricity flow charts, personal data that belongs to approximately 10000 of the company's employees, also radiation exposure estimates for nearby residents [9][10].

B. Security Vulnerabilities

In general, I&C in NPP are physically isolated from external networks and have a different operational environment from that of conventional IT systems. As a result, NPP were regarded as being safe from external cyber-attacks. However, continuous cyber-attacks against NPP signified that NPP are as susceptible to cyberattacks as other critical infrastructures [11] and conventional IT systems.

ICS, usually control the physical world and IT systems manage data. ICS are different from traditional IT systems, including dissimilar risks and priorities. Some of the different characteristics include important risk to the health and safety of human lives, severe destruction of the environment, and financial problems such as production deficit, and undesirable effect to a nation's economy. Performance and reliability requirements for ICS are distinct, by using operating systems and applications that may be seen unusual in a classic IT network environment. At first, ICS had slight similarities to IT systems in that ICS were inaccessible systems implementing proprietary control protocols with specific hardware and software. Commonly accessible, low-cost Ethernet and Internet Protocol (IP) devices are now substituting the older proprietary technologies, which raises the likelihood of cybersecurity vulnerabilities and events. Currently, ICS are embracing IT solutions to endorse corporate connectivity and remote access abilities, and are being created and employed via industry standard computers, operating systems (OS) and network protocols, where the resemblance to IT systems comes from. This novel integration deploys IT capabilities, but it meaningfully offers less separation for ICS from the outside world than antecedent systems, increasing the necessity to secure these systems. Despite the fact that security solutions have been designed to deal with these security matters in characteristic IT systems, particular precautions must be engaged when presenting these similar solutions to ICS environments [1].

1) Lack or Improper Input Validation

Attackers exploit vulnerabilities in services and scripts written by I&C vendors, resulting from the non-secure coding practices, allowing attackers to send forged request in order to modify the program execution. In the same way, using vulnerable protocols with for networking will be exploited to create malformed packets. Vulnerabilities found in these protocols and services make an attacker able to manipulate plant component, via well-known attacks. Vulnerable modules that might be concerned include Workstations at Main Control Room (MCR), Remote Shutdown Station (RSS); Process Information and Control System (PICS); Safety Information and Control System (SICS) and HMI. The attacks that could take place by exploiting this vulnerability are buffer overflow, command injection, and SQL injection.

2) Inappropriate Authorization

Authorization guarantees access to resources only by authorized entities. Access control mechanisms are implemented to ensure appropriate authorization. Absence of or weak authorization mechanisms can be exploited by attackers to gain illegal access to resources and tamper I&C system components. Software installed at operator workstations side must perform access control checks, or it will open a new door for attackers to perform unauthorized actions. Vulnerable modules include Workstations at MCR, RSS, PICS, SICS, HMI, Safety Automation System (SAS), Protection System (PS), Process Automation System (PAS). Existing module in I&C system must first verify whether the requesting module is allowed to access the resource. Escalation of privilege is one of the attacks that could be performed with authorization vulnerability.

3) Improper Authentication

Network protocols used within I&C system architecture during communication, frequently suffer from weak authentication mechanisms to verify the identity of the packet and also the user. Weak authentication vulnerabilities permit attackers to eavesdrop on network communications and capture the identity credentials of legal users, ending with an unauthorized privilege. Mutual authentication before sending or receiving data is not performed by the components of I&C. Not verifying the origin or authenticity of data, permits malicious data into components, credential theft, authentication bypass, etc. Furthermore, non-properly protected confidential data stored in databases can also be exploited. Vulnerable modules that might be touched by this are almost all I&C systems, sub-systems and components [10]. Often, I&C vendors leave behind authentication information from their product code or documentation, which can be definitely accessed and exploited by attackers. Weak passwords or using default passwords are another significant vulnerability to consider. There are numerous possible aspects that can be used to authenticate a person, device, or system, together with something the user knows, something the user has or something the user is. For instance, authentication could be founded on something known (e.g., PIN number or password), something possessed (e.g., key, dongle, smart card), something the user is like a biological characteristic (e.g., fingerprint, retinal signature), a location

(e.g., Global Positioning System (GPS), location access), the time a request is made, or a mixture of these attributes. Normally, the more authentication process includes more factors, the more strong the process will be. Multi-factor authentication refers to the process when two or more factors are used [5].

4) Unencrypted Sensitive Data

Frequently data at rest and in transit is unencrypted, making them vulnerable to disclosure. Moreover, network packets exchanged between several components of I&C are not encrypted but in plaintext form. Vulnerable modules that might be touched by this are almost all I&C systems, sub-systems and components [10]. Exposure of product source code, topology, legitimate user credentials, might result as a consequence.

5) Incorrect Software Configurations and Management

Security breaches and exploitations of plant operations are a result of misconfigurations or vulnerabilities found in I&C software. Modules that are seen vulnerable to this are Workstations at MCR, RSS, PICS, SICS, HMI, SAS, PS, and PAS. The existence of these vulnerabilities is caused by poor patch management, poor maintenance, and built-in flaws in I&C products. Additionally, improper installations of applications also offer an opportunity to attackers to tamper the system.

6) Lack of Backup Facilities

Some of I&C systems in NPP do not own backup and restore facilities dedicated to databases and software. NPP that possess backup facilities often store them offsite, and they are not often exercised and tested. Vulnerable modules that might be concerned by lack of backup facilities are SAS, PS, PAS, Sensors, Actuators, PICS, and SICS [10]. NPP must be operated 24/7 and the absence of a backup feature can result in catastrophic effects if an incident occurs.

7) Absence of Audit and Accountability

Some attacks are hard to detect since they are launched in a cautious manner like insider attacks. The nonexistence of auditing and logging mechanisms assists attackers into covering their tracks after attacks. Vulnerable modules that might be touched by this are almost all I&C systems, sub-systems and components. Storing activity logs of I&C components and operator actions is vital in order to trace attack patterns, but also to avoid repudiation threats from insiders as well as actions in I&C components and systems.

8) Absence of Security Awareness

Technology advancements and the people using these technologies present multiple risks to information security. The human factor is considered as one of the major sources of information security risk, also one of the most difficult to control. According to a Deloitte's Technology, Media, and Telecommunications (TMT) Global Security Study [12], 70% of the TMT organizations surveyed rate their employees' lack of security awareness as an "average" or "high" vulnerability, which was the case for Korea Hydro and nuclear Power Co. Security controls that are conform to the NIST SP 800-53 Awareness and Training (AT) family offer policy and procedures for guaranteeing that each user of an information system is equipped with elementary information system security awareness and training materials before authorization to access the system is granted. Security awareness is a crucial part of ICS incident prevention, mainly when it comes to social engineering threats. Social engineering is seen as a method used to influence individuals into revealing private information, such as passwords. This information can then be exploited to endanger otherwise secure systems. Employing an ICS security program may bring changes to the means used by personnel to access computer programs, applications, and the computer desktop itself [9].

C. Classification of adversaries

In [13] adversaries are categorized into eight classes that can endanger safety and security of NPP. The categories are as follows: covert agents, disgruntled current employees, disgruntled ex- employees or insider attackers, recreational hackers/ hobbyists/ script kiddies, militant opponents to nuclear power, non-state hackers (e.g., cyber criminals/organized crime), nation-state hackers (e.g., governments and militaries), and terrorists (e.g., non-state armed groups).

1) Covert Agent

A retired or a present employee of an intelligence agency, and whose identity is unknown to others. The agent is hired to steal secret information and personal information about adversaries. In order to get information, this agent must have access to the system and documentation, or apply a social engineering method.

2) Disgruntled Current Employees or insider attackers

Someone who is not satisfied with his/her job, and wants to compromise the system by using illegal approaches. Reasons behind dissatisfaction vary, but the usual motivations are to take revenge, create chaos, damage nuclear security's image, or steal information for economic gain. To perform such attack, the attacker needs medium to high level resources to execute an attack, e.g., systems access. Moreover, an employee must own some higher privileges on processes and systems, programming skills and information about the system´s architecture, information about possible existing passwords, and the capability of installing "kiddie" tools or scripts.

3) Disgruntled Ex-Employee

This person has similar motivations as the ones of a disgruntled employee. Their purpose is to take revenge on the employer, sell confidential information to adversaries for economic gain, or disclose confidential information to the public in order to damage employer´s public image. As an ex-employee, she/he may still own confidential documentation, access to facility resources, and potential connections to other employees. To execute such an active attack, an attacker should have knowledge about systems´ passwords, access to systems, and backdoors made by social engineering techniques.

4) Recreational Hackers/Hobbyists/Script Kiddies

Their motivations behind the intrusion to systems are for fun or to win a challenge. These attackers are interested into learning about new vulnerabilities and exploiting by performing them on real systems. They often download and use free scripts and tools available on Internet. Their intentions might be harmless; yet, mechanisms used to learn about these vulnerabilities and the way to exploit them is risky. In case cybersecurity mechanisms are not well deployed inside NPP, this might be destructive. Without owning an advanced level of expertise, frameworks such as Metasploit provide SCADA-specific exploits, which script kiddies can use to execute an attack easily. Such attackers could certainly be blocked by imposing best practices such as patch management, policy enforcement, and suitable use of antivirus, intrusion detection systems (IDS), and firewalls inside the organization.

5) Militant Opponent to Nuclear Power

She/he has strong public thoughts on precise nuclear issues, and often slows down nuclear business operations. These attackers are financially supported through secret channels or agencies [10]. However, they only know the public information available on systems. Moreover, they have sufficient time to perform such attacks and mainly aim defined public events such as elections. They may or may not have computer skills; still, they get help from the hacker community to execute a cyber-attack.

6) Non-State Hackers

Groups or individuals with the main objective are financial gain by stealing nuclear sensitive data belonging and then blackmailing the facility to which data belong to into paying a ransom. Usually, they threaten to exploit vulnerabilities in SCADA systems. These attackers do have funds and can hire expert hackers or buy hacking tools in order to attack systems. A set of SCADA-targeted automated attack tools, in the form of Metasploit add-ons that can help in executing attacks on ICS, exist. Every so often, these attackers employ former/current employees of a facility to perform social engineering to extract information.

7) Nation-State Hackers

Governments hire specific individuals to perform cyber operations, internationally or nationally. State hackers vandalize and block access to websites, and perform industrial espionage to steal a country's confidential data. Additionally, state hackers constitute the most harmful threat to SCADA systems, as long as these hackers get all of their owned information and funds from the government. The government has resources to hire the best hackers and offer those funds, infrastructure, and facilities to create zero-day exploits, to use them against an enemy country in order to steal a nuclear technology, intelligence collection, etc. Although zero day attacks are single-use weapons, they are capable of causing a huge damage to a country's infrastructure, economy, and systems.

8) Terrorist

Throughout the history of cyber-attacks on SCADA systems, no evidence can be found of a terrorist attack; still, the situation will not stay like this in the future. According to former U.S. President George W. Bush, terrorists can get into the network with the intention to attack a nuclear facility, and consequences of such intrusion could be intolerable [10]. Objectives of such terrorists differ: sometimes they want to accumulate intelligence, create backdoors for later use, spread fear and panic among the public, or take revenge on the government. Furthermore, some terrorist groups have developed important skills to use social media as a way to hire hackers.

D. Cybersecurity requirements

Cyber security features that provide confidentiality, integrity, and availability must be integrated in the design of safety systems. Cybersecurity controls should not have an opposite effect on the plant´s safety objectives and should not intervene with their operations. Concerns have been raised regarding possible effects that such features can have on safety functions´ performances. Also, it shall not jeopardize diversity and safety Defense in Depth (DiD) features effectiveness implemented in I&C architecture [14].

- **Confidentiality**

Imposing this feature inside a safety system restricts actions an attacker can make on information transferred between safety systems, or between safety and non-safety systems. In general, to ensure confidentiality cryptographic techniques must be deployed, in order to avoid any illegal disclosure of information during transmission and reception [15]. To make sure that these added cryptographic features do not degrade safety functions, these cryptographic mechanisms are employed for communication between safety and non-safety systems. In case an unpredictable overhead is introduced to data communications because of added cryptographic approaches, other possibilities exist. These supplementary implementations may implicate physical and logical access controls on the system, monitoring dynamically and tracking all accesses to the system to detect and respond to intrusions in a convenient way, by enforcing auditing and/or validation mechanisms to identify unauthorized access and alterations to the system. For authorized individuals´ a background check should be accomplished with regards to their experiences and trustworthiness.

- **Integrity**

The purpose of protecting safety systems´ integrity is to restrict malicious actions attackers can perform on safety systems so that they cannot negatively impact safety functions [15]. Protecting integrity can be accomplished by restricting unauthorized alterations of software and hardware in safety system. Limiting access to these systems might be a possibility, since access is made via direct interfaces, e.g., ports on the hardware, or using indirect interfaces like data links. An access control list including

authorized actions should be implemented so that illegal system modification via direct interfaces is forbidden [15].

- **Availability**

Affecting negatively safety systems´ availability must not be permitted [15]. Safety systems´ operations can be compromised directly or indirectly by refusing access to the system to authorized users. Methods for restricting an attacker´s ability of performing such attacks or controlling the attack´s effect on a system, should not interfere with safety function, as enforced e.g., by IEC 62859 [14]. These approaches consist of installing mechanisms at the system's external interface to prevent and limit denial of service attacks´ effects. While configuring these systems, restrictions on users´ actions should be considered to prevent them from executing such attacks against other systems, by controlling capacity surplus and/or bandwidth to stop information-flooding and attacks´ effects. Some cryptographic mechanisms are capable to comply with these requirements, e.g., by limiting the attacker´s actions, to possibly make modifications that may negatively affect the availability.

E. STRIDE threat modelling

This Section presents the STRIDE (Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, and Elevation of Privileges) threat model of a typical NPP' I&C system by taking into consideration its characteristics and architecture. STRIDE is a method developed by Microsoft, which describes an adversary's objectives, is used for threat modelling [10]. Tab 1. shows a summary of the STRIDE analysis.

- **Spoofing**

Spoofing is a scam category where an intruder tries to gain unauthorized access to a user's system or information by pretending to be the legitimate user [10]. For NPP, this unauthorized access can cause I&C systems´ disruption or lead to the system´s misuse. Spoofing can be divided into two categories, it can be related to the system or linked to the personnel. The first type focuses on spoofing I&C system´s credentials, the second type concentrates on unauthorized access gained after stealing personnel credentials, e.g., Passwords and tokens, and then pretending to be the real authorized user. Session hijacking is a typical attack for personnel spoofing; the attacker captures a current session and attempts to connect to the receiver as an authentic user. In the case of a system spoofing, malicious code injection in the form of scripts into a web browser is a common strategy. Other techniques exist in order to spoof credentials; it includes social engineering, e.g., watching and/or manipulating user or system behavior and activities, and incorrect input, e.g., SQL injection.

- **Tampering**

Consists of altering legitimate data, and as a consequence the system´s integrity is compromised. The data can be tampered whether it is in transit or at rest. An attacker can exploit any misconfiguration or if there is no presence of integrity checking procedure in the system to compromise the system´s integrity.

- **Repudiation**

It is caused by the lack of appropriate auditing and logging mechanisms. An attacker can exploit vulnerabilities in the logging mechanism, steal keys, or even produce fake digital signatures to allow unauthorized actions. As an illustration, an operator or a compromised system at a NPP can deny executing some actions or operations on plant systems, e.g., a plant operator alters temperature´s values and water level of a plant, but later denies performing such an action.

- **Information disclosure**

This threat is a result of improper protection of information. There are many forms of information – for example, user credentials, network packets, source code, files, or a database. Sensitive plant´s information can be illegally released by exploiting vulnerabilities like software misconfigurations, improper authorization or authentication mechanisms.

TABLE 1 STRIDE ANALYSIS.

Threat category	Attacker type	Vulnerability category
Spoofing	Covert Agent Disgruntled Ex-Employee Non-State Hacker Terrorist	No or Incorrect Input validation. Improper Authentication Improper Authorization
Tampering	Militant Opponent Recreational Hacker Terrorist	Improper Authentication Improper Authorization Improper Software Configuration & Management
Repudiation	Disgruntled Current Employee	Auditing and logging
Information Disclosure	Covert Agent Disgruntled Current Employee Non-State Hacker Disgruntled Ex-Employee	Improper Authentication Improper Authorization Improper Software Configuration & Management
Denial of Service	Recreational Hacker Terrorist	Improper Software Configuration & Management No or Incorrect Input Validation Lack of Backup Facilities
Elevation of Privilege	Disgruntled Current Employee	Improper Authentication Improper Authorization

- **Denial of service (DoS)**

By overwhelming I&C systems with a large number of repetitive requests, required components become unavailable. These requests can be sent by installing a malware or in case the system is connected to internet with a hidden connection. DoS attacks generally take place when backup facilities are unavailable and inexistence of input validation methods.

- **Elevation of Privilege**

Leading to an abuse of legitimate access, malicious insiders having access to resources or operations may alter

their account permissions to permit supplementary accesses to systems to which they do not have access to. They can then abuse their privileges by performing malicious actions, e.g., stopping core functions or altering parameter values.

F. Industry and Government Responses to NPP Cybersecurity

In the previous Section, known attacks and vulnerabilities in NPP were underlined. Since they pose important risks to the economy and to national security, numerous attempts were made by international organizations, regulatory and research institutes, and governments to set up cybersecurity guidelines, standards, and frameworks dedicated to security of NPP.

For industry adoption and regulatory approval, three features of digital I&C systems are distinguishing.

First, a digital I&C system is more complicated than its analog predecessor because of the number of connections it has among its many components. Second, the digital system rely more on software. Usually, a unit has around 10000 sensors and detectors and 5000 km of I&C cables. The total mass components connected to I&C, is close to 1000 tones. Making I&C system one of the heaviest and most extensive non-building structures in any NPP. Third, the complete reliance on computers increases the importance of cybersecurity. The first two of these features, complexity and software-dependence, introduce new possibilities for common cause failures.

The increased use of commercial "off-the shelf" software is considered as one practice hurting the nuclear industry. This type of software does not deliver a suitable level of protection from external threats and is often seen as a direct approach to penetrate a facility network. The use of insufficient software, mixed with executive-level ignorance of security risks, builds an easy way for an attacker to misuse assets. There is a common misrepresentation which refers to nuclear facilities as being "air-gapped" – totally inaccessible from the Internet – signifying that the industry is safe from cyber-attacks. Considerable commercial software offers Internet connectivity through virtual private networks (VPN) or else Intranet. These connections often go unlisted and keep on being ignored while implementing software or deploying momentary Internet connections for a project. Furthermore, the focus has been given more to physical safety and protection instead of cybersecurity controls. Therefore, very few developments have been made to reduce cyber risks through standardized control and measures [11].

NPP are securely maintained and considered as the most protected and secure facilities in the world. However, accidents can happen, undesirably affecting environment and people. Vulnerabilities threatening the physical security of a NPP and their ability to launch acts of terrorism were elevated to a national security issue following the attacks of 9/11, 2001. Consequently, the American congress endorsed new nuclear plant security requirements and has frequently devoted attention on regulation and enforcement by the Nuclear Regulatory Commission (NRC). Years passed after the 9/11 attacks, but security at NPP persists as a vital matter. To decrease the likelihood of an accident, the International Atomic Energy Agency (IAEA) supports Member States in applying international safety standards to reinforce safety in NPP [10]. NIST has published a well-established risk management framework in NIST Special Publications (SP) 800-30 [16], 800-37 [17], and 800-39 [18], which analyzes distinct threat scenarios and evaluates the various attack possibilities that can exploit system vulnerabilities. On the other hand, the NIST risk assessment framework, mentioned above, does not describe precise procedures on the approach a company should assess the quantification of risks, i.e., how and to what degree an attack can endanger system confidentiality, integrity, or availability. In 2008, NIST issued a guideline on securing ICS [5]. It systematically explained the security of ICS systems, mostly containing SCADA architecture, distributed control systems (DCS), secure software development, and deployment of controls in order to secure networks. NIST also came up with a guideline on the Security for Industrial Automation and Control Systems while working with the Industrial Automation and Control Systems Security ISA99 Committee.

The IEEE produced the SCADA cryptography standard in 2008 [19], which offers a comprehensive explanation on the way to found secure communication between SCADA servers and workstations. Organizations can also attain certification under this IEEE standard if they fulfill with the requirement. The International Organization for Standardization (ISO) has also issued a standard, ISO/IEC 27002:2013 [20], which gives guidelines for initiating, implementing, maintaining, and improving information security management in organizations [10]. NRC's cybersecurity regulations necessitate each NPP to present a cybersecurity plan and implementation schedule. The plan must deliver "high assurance" that the digital computer and communications systems implemented in order to perform the next functions will deliver sufficient protection against design basis attacks:

- Safety-related Functions or vital to safety.
- Security functions.
- Emergency mobility functions, as well as offsite communications.
- Support systems plus equipment that, if compromised, would undesirably jeopardize safety, security, or emergency mobility functions [4].

As a result, cybersecurity has been adopted as NPP regulation requirement under the US code of federal regulation (CFR) [3]. Also, regulatory agencies like the US NRC and IAEA created and distributed regulatory guidelines, considering construction of cybersecurity plans and programs for NPP. The IAEA and World Institute for Nuclear Security (WINS) are multiplying their efforts in order to protect NPP by addressing cybersecurity issues and challenges on a global scale. Currently, some of issues include:

- Issuing multiple documents addressing cybersecurity on nuclear facilities.

- Providing technical and strategic security training to involved officials of member countries.
- Offering expert guidance and capacity building to officials and representatives.

NSS-17 [13] was issued by IAEA as a technical guidance for guaranteeing computer security at nuclear facilities. Similarly, the IAEA NSS-13 [21] recommends that the available computer-based systems included in nuclear facilities must be protected against compromise, and also an appropriate threat assessment must be realized in order to prevent attacks.

Threats were classified from various adversaries' perspectives, detection and prevention mechanisms for compromises of NPP information systems were also addressed [22]. Additionally, nothing like usual ICS and SCADA systems, governments, and NPP regulatory agencies specify that NPP I&C systems must comply with the following firm safety requirements [5][23]:

- Requirements for annual maintenance, best availability and functionality levels, and environment tests.
- Nuclear reactor safety and also physical protection of nuclear material must be taking in consideration.
- Defining system security levels by bearing in mind safety level ranking, and evaluating safety risks in relation to security threats.
- Verification that security functions do not have opposing effects on the safety and functionality of facilities.
- Management and maintenance must consider the safety and reliability of systems, examination and also qualification by regulatory agencies.
- Redundancy and diversity must be taken in consideration in the design.

However, all of these efforts are continuing and necessitate indefinite time to mature.

The guidelines, standards, and recommendations provided by governments and regulatory authorities necessitate complete review to make sure that they describe and include the newest risk assessment developments, for example, cyber threat information sharing, risk assessment of tacit knowledge, dissemination of risk assessment results, etc. These features are obligatory in order to keep NPP risk assessment up-to-the-minute on progressive cyber threats and to be able to manage cyber incidents in a proper manner.

On the other hand, at present, the abovementioned guidelines do not provide a detailed approach on imposing security controls and avoiding cyber risks.

IV. SECURITY CONTROLS FOR NPP

Standards are endorsing the improvement of cybersecurity in NPP. Fig. 2 shows the standardizing processes and procedures, which are important to accomplish an international rewarding collaboration. Abundant standards addressing information security were established in recent years. Still, not all of them are practical to apply in NPP.

Designed for I&C systems in NPP, the new draft IEC 63096 is expected to be published in 2019. The standard aims its attention specifically on the selection and application of cybersecurity controls from the contained security controls within the catalogue. Preventing, detecting, also reacting to digital attacks against computer-based I&C systems are the major functions of the selected and applied cybersecurity controls. Based on IEC 62645 [24], IAEA, in addition to country precise guidance in the technical and security application area. Designers and operators of NPP (utilities), systems evaluators, vendors and subcontractors, and by licensors can use this standard. For that reason, the goal of this standard is to enlarge the SC45A series of documents focusing on cybersecurity with IEC 62645 [24] as its high-level document, by classifying nuclear I&C precise cybersecurity controls for I&C systems into Safety Classes 1, 2, 3 and non-classified (NC) I&C systems. A major difference between this standard and usual IT systems or industrial automation systems standard is the safety classification of I&C nuclear systems and related safety requirements. IEC 62645 [24] was issued in August 2014, and IEC 62859 [14] was published in 2016, along with this new standard related to cybersecurity controls, are planned to be used for I&C systems and NPP. The new standard is structured as follow.

The first main Section designates the intended audience, which is distinguished by parties that are in charge of:

- I&C platform development.
- Project Engineering for I&C system, including installation and commissioning.
- Operation and maintenance of I&C system.

In the second main Section, a detailed description of each security control is explained. Inside this structured representation, the purposes of Security Degrees along with the specific control are defined either highly recommended or optional. As well, additional description specifies whether the control conserves the confidentiality, integrity or availability. Each Section related to a security control provides specific implementation guidance on how to implement the control; it also differentiates between I&C platform development, project engineering or operation and maintenance.

Based on IEC 62645 [24], the third main Section is dedicated to the process of selecting the available security controls. Controls are allocated depending on the security degree of the particular I&C system. Tools and Legacy systems are also considered in this standard. After selecting the security controls, a threat and risk assessment is required in order to detect a residual risk that necessitates the implementation of supplementary security controls.

Concerning controls three cases are distinguished, using the guidance provided by the Draft ISO/IEC 27009 [25] on sector specific security controls. The following three cases on the refinement of ISO/IEC 27002 security controls are examined [20]:

- Security controls are adopted from ISO/IEC 27002 [20] without any adjustment. If needed, the obligatory details are added by the standardized structure.

- To meet requirements from the nuclear I&C domain, Security controls from ISO/IEC 27002 [20] were modified and described in details to better.

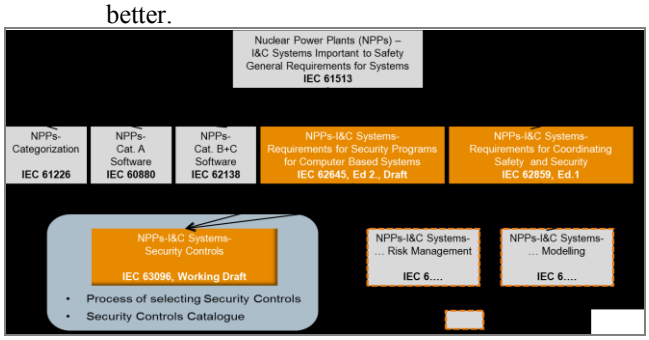

Figure 2. New IEC 63096 Security Controls standard in the SC45A standards hierarchy [4].

- In order to meet the particular requirements from the nuclear I&C domain, a new security control has been added and inserted into ISO/IEC 27002 [20] clause (5 through 18. This is the case where the ISO/IEC 27002 [20] does not define specific security controls needed in nuclear I&C.

IEC 62541 [26] defines the open platform communication Unified Architecture (OPC UA), it is an automation middleware or machine-to-machine (M2M) protocol. The standard features an object-oriented meta-model to characterize data structures and remote procedure calls, which can be dynamically explored and modified through IP communication, along with security mechanisms such as authentication and encryption. OPC UA is adaptable to manufacturing software, it defines [26]:

- An information model for structure, behavior and semantics description.
- A message model for interactions between applications.
- A communication model to carry data between end points.
- And a conformance model to guarantee interoperability between systems.

The communication services of OPC UA are mainly used in the following domains: Process automation, power plants with, traditional and renewable Building automation, and Factory automation (in particular robotics).

The OPC UA specifications are made up by 13 parts, the first seven parts are related to the core specifications e.g., the concept, security model, address space model, services, information model, service mappings and profiles. The parts eight to thirteen are related to access type specifications like data access, alarms and conditions, programs, historical access, discovery and aggregates. Interoperability is achievable by using a communication standard that is platform and vendor independent, such as IEC 62451 [26] (OPC UA) and IEC 61850 [27] (Communication Networks and Systems in Substations). OPC UA is a platform-independent standard that helps into reaching interoperability between devices with dissimilar manufacturers and communication protocols. OPC UA simplifies communication by sending messages between OPC UA Clients and Servers. For example, its applicability to the nuclear context is demonstrated by Framatome. Recognizing the potential of OPC-UA in sensors, Framatome started incorporating them into monitoring instruments (SIPLUG®) for mountings and their related electric drives. The solution is employed in the nuclear Industry for monitoring critical systems in remote environments, without undesirably affecting the availability of the system [28].

V. CONCLUSION

This paper gave an overview of security vulnerabilities in I&C systems and EPS inside NPP, by going through notorious attacks across history and listing some of the vulnerabilities that can be exploitable by malicious attackers. An introduction to a new draft standard, IEC 63096 [4] had being given. The improved performance digital technology has offered a significant influence on I&C systems design in NPP. The nuclear industry has a conservative methodology in approaching safety, and a considerable effort is obligatory in order to provide the essential evidence and analysis to guarantee that digital I&C systems can be employed in safety-critical and safety-related applications. In general, I&C systems are inaccessible from outside communication systems. Still, this is not sufficient for secure operation inside NPP, as in the case of Stuxnet. Interoperability has to be addressed from I&C architecture design phase, as the systems have to interact. The threat from cyber-attacks is more and more seen as a problem of national and international security as cyber-attacks evolve, become more advanced and as actors behind them are no longer limited to private hackers or organized criminals, but also nation states and insiders.

In future work, we intend to focus more on the listed vulnerabilities, and introducing security in hardware by using a trusted platform module instead of only focusing on securing software, also some best practices to widen the protection area.

ACKNOWLEDGMENT

Some of the addressed cybersecurity related topics are being elaborated as part of Framatome GmbH's participation in the "SMARTEST" R&D (2015-2018) with German University partners, partially funded by German Ministry BMWi.

REFERENCES

[1] A. Tellabi, I. Ben Zid, E. Bajramovic, and K. Waedt, "Safety, Cybersecurity and Interoperability of Modern Nuclear Power Plants," IARIA 8th International Conference on Performance, Safety and Robustness in Complex Systems and Applications, 2017.

[2] J. Rrushi, R. Campbel, "Detecting cyber attacks on nuclear power plants," The International Federation for Information Processing (ICCIP 2008), Springer, Boston, vol. 290, 2008, ISBN: 978-0-387-88522-3.

[3] INSAG-24, International Nuclear Safety Group, "The interface between safety and security at nuclear power plants," IAEA, 2010.

[4] J. Bochtler, E. Quinn, and E. Bajramovic, "Development of a new IEC standard on cybersecurity controls for I&C in Nuclear Power Plants – IEC 63096," NPIC & HMIT 2017, San Francisco, 2017.

[5] K. Stouffer, V. Pillitteri, S. Lightman, M. Abrams, and A. Hahn, "NIST Special Publication 800-82: Guide to Industrial Control Systems (ICS) Security,", NIST, 2011.

[6] A. Tellabi, Y. Sassmanhausen, E. Bajramovic, and C. Ruland, "Overview of Authentication and Access Controls for I&C systems ," IEEE 16th international conference on industrial informatics, 2018.

[7] M. Holt, A. Andrews, "Nuclear Power Plant security and vulnerabilities," Congressional Research Service, January 2014.

[8] D. Papp, Z. Ma, and L. Buttyan "Embedded systems security: threats, vulnerabilities, and attack taxonomy," 13th Annual Conference on Privacy, Security and Trust (PST), IEEE, 2015, doi:10.1109/PST.2015.7232966.

[9] C. Baylon, R. Brunt, and D. Livingstone, "Cybersecurity at civil nuclear facilities understanding the risks," Chatham House Report, September 2015.

[10] R. Masood, "Assessment of cybersecurity challenges in nuclear power plants security incidents, threats, and initiatives," Cybersecurity and Privacy Research Institute the George Washington University, 2016.

[11] B. Kesler, "The vulnerability of nuclear facilities to cyber-attack," Defense and Diplomacy Journal, vol. 5, No. 3, 2016.

[12] Deloitte, "Security Awareness: People and Technology," [Online]. Available from: http://www2.deloitte.com/, 2017.12.19.

[13] IAEA Nuclear Security Series No. 17, "Computer Security at Nuclear Facilities," IAEA, 2011.

[14] IEC 62859:2016, Nuclear Power Plants - Instrumentation and control systems - Requirements for coordinating safety and cybersecurity, IEC.

[15] Regulatory Guide 5.71, Revision 0, "Cyber Security Programs for Nuclear Facilities," U.S. Nuclear Regulatory Commission, January 2010.

[16] G. Stoneburner, A.Y. Goguen, and A. Feringa, "NIST Special 800-30: Risk Management Guide for Information Technology Systems,", NIST, 2002.

[17] Joint Task Force Transformation Initiative, "NIST Special Publication 800-37 Rev. 1: Guide for Applying the Risk Management Framework to Federal Information Systems: a Security Life Cycle Approach," NIST, 2014.

[18] E. Aroms, "NIST Special Publication 800-39: Managing Information Security Risk," NIST, 2012.

[19] "IEEE Trial-Use Standard for a Cryptographic Protocol for Cyber Security of Substation Serial Links," in IEEE Std 1711-2010, vol., no., pp.1-49, 2011.

[20] ISO/IEC 27002:2013, Information Technology - Security Techniques - Code of Practice for Information Security Controls, ISO/IEC.

[21] IAEA Nuclear Security Series No. 13, "Nuclear Security Recommendations on Physical Protection of Nuclear Material and Nuclear Facilities," IAEA, 2011.

[22] W. Ahn, M. Chung, B. Min, and J. Seo, "Development of cyber-attack scenarios for Nuclear Power Plants using scenario graphs," International Journal of Distributed Sensor Networks, vol. 11, April 2015, doi: 10.1155/2015/836258.A

[23] ISO/IEC 27001:2005, Information Technology –information security management systems –requirement, ISO/IEC.

[24] IEC 62645:2014, Nuclear Power Plants – Instrumentation and Control Systems – Requirements for Security Programs for Computer-Based Systems, IEC.

[25] ISO/IEC 20009-1:2013, Information technology – Security techniques – Anonymous entity authentication, ISO/IEC.

[26] IEC 62451-1:2016, OPC Unified Architecture – Part 1: Overview and Concepts, IEC.

[27] IEC 61850:2013, Communication networks and systems for power utility automation, IEC.

[28] V. Watson, A. Tellabi, J. Sassmannshausen, and X. Lou, "Interoperability and security challenges of Industrie 4.0," 2017, doi:10.18420/in2017_100.

Four Testing Types Core to Informed ICT Governance for Cyber-Resilient Systems

Keith F. Joiner
School of Engineering and IT
University of New South Wales
Canberra, Australia
k.joiner@adfa.edu.au

Amit Ghildyal
School of Business
University of New South Wales & Department of Defence
Canberra, Australia
Amit.Ghildyal@defence.gov.au

Narelle Devine
Department of Human Services
Australian Government
Canberra Australia
narelle.devine@humanservices.gov.au

Alan Laing
Chief Information Officer Group
Department of Defence
Canberra, Australia
alan.laing@defence.gov.au

Anne Coull
Information Security
Westpac Banking Group
Sydney, Australia
anne.coull@student.unsw.edu.au

Elena Sitnikova
Australian Centre for Cybersecurity
University of New South Wales
Canberra, Australia
e.sitnikova@adfa.edu.au

Abstract—**Research on ICT projects continues to report very high cost and schedule overruns, as well as many high-profile ICT projects experiencing high incidences of unexpected cyber-vulnerabilities. Consequently, there is renewed interest in ICT governance from diverse areas. Some of the proposed governance models considered have great complexity while others appeal to simplicity for success. Three diverse and practical research efforts in ICT governance in Australian Government, as well as observations in the Banking Sector, came to similar concerns about the importance and type of ICT testing and expertise critical for ICT project governance to build cyber-resilience. Today's ICT Governance critically depends on: (1) information coming from all four types of testing, (2) the management of the testing as a coherent whole, and (3) that such test capabilities must endure through the whole life-cycle, so as to provide a sufficient degree of commercial and architectural independence to enable hard and timely decisions. Further, cyber-resilience challenges ICT testing to cope with increasing system configurations, threat permutations, future upgrades and threat sequencing. Therefore, this research uniquely calls for all ICT test types to use new combinatorial test design techniques for efficient screening and cyber-threat rigor. These lessons were shared at a special conference panel on ICT governance for resilient systems [1]-[4], where for the first time authors called for ICT governance frameworks to directly include test-informed previews in all decisions so that ICT can be more innovative, competitive, and cyber-resilient. This paper outlines the four testing types and lists the test infrastructure and combinatorial test design skills necessary for each.**

Keywords- ICT governance; usability testing; cyber-resilience; penetration testing; integration testing; project success factors; stress testing.

I. INTRODUCTION

Difficulties with ICT projects abound in all parts of the World [1]. There are also reports of many high-profile ICT projects experiencing high incidences of unexpected defects and cyber-vulnerabilities despite apparently extensive testing [6]-[8]. Research by [9] into 1,471 IT projects showed that cost overrun averages were not unremarkable to other projects (27%) but that there was, what they describe as, a *fat tail* of risk. They summarize that '*Fully one in six of the projects in the sample was a Black Swan, with a cost overrun of 200%, on average, and a schedule overrun of almost 70%.*' Reference [10] cites a U.S. Department of Defense (DoD) finding that *'85 percent of software intensive projects finished over time or budget; half of projects doubled original cost estimates; projects slipped an average of 36 months; and one-third of projects were cancelled.*' He goes on to cite military standards that '*inadequate software reliability can double or triple field support and maintenance costs*,' meaning that even those software-intensive projects that eventually succeed can remain a sustainment burden through-life. These sobering findings are alongside ever-increasing software functionality in systems, systems interconnectivity and autonomy [11]-[12], as well as increasingly sophisticated and cost-effective cyber-threats [13]-[14]. Reference [8] proposes governance involving continuous system monitoring through-life and his assessment is one of a field continuously exploring the bounds of achievement:

'... there will be notable failures, some great successes, and a large number of projects that get delivered in a sub-optimal state. That represents the norm for large software projects ... it is critical to understand that SoS [systems of

systems] *generate emergent behavior that can't always be reliably triggered by normal test inputs.*'

Governing the complexity of the software systems and their functions is significantly affected by the increasing number and sophistication of cyber threats to both open and closed system architectures [7] [14]. Cybersecurity is increasingly moving from avoiding cyber-attack in the form of barriers, to being able to sustain and recover from cyber-attack, or 'fight through' [7]. Cyber-resilience has many definitions, such as *'the ability* [for] *cyber systems and cyber-dependent missions to anticipate, continue to operate correctly in the face of, recover from, and evolve to better adapt to advanced cyber threats*' [15], or more simply as *'the capacity of an enterprise to maintain its core purpose and integrity in the face of cyber attacks*' [16]. Current ICT governance is seriously challenged by this shift to be more adaptable, omnipresent and evolving in the design of systems, support to operations and test infrastructures.

Three separate and diverse Australian Government-based research efforts in ICT governance, as well as an assessment in the Banking Sector, were found to have similar concerns about the importance and type of ICT testing and test expertise critical to ICT governance and the ability to build cyber-resilience: namely, systems integration, usability testing, stress testing and security testing (vulnerability, penetration & high assurance testing - VPAT). Each of these research efforts will be briefly critiqued before covering the four test types. Finally, the paper draws the common threads of that research in order to recommend on the role of testing in supporting ICT governance to achieve cyber-resilient systems.

II. STATE OF THE ART

A good example of the growth in software-intensive functionality and the associated software and cybersecurity difficulties in a project is the Joint Strike Fighter (F35) aircraft program [17], where capability growth has been limited by uncontained software deficiencies in development and now in early operational testing. Reference [18] examines the testing and certification difficulties of combining such software functionality levels into advanced aircraft software that have intelligent autonomy as well. He supports the push towards a more continuous certification approach, one that combines both test-based verification and analysis-based verification.

Growth in ICT technology in society more broadly is rampant, leading some reviews to have predicted a so-called '*C generation*' of *'digital natives*' [19] and other researchers to predict a shift from the Information Age to a new Synthetical Age [20]. For example the prediction by [19] has come true that a '*highly connected generation will live "online" most of their waking hours, comfortably participate in social networks with several hundred or more contacts, generate and consume vast amounts of formerly private information, and carry with them a sophisticated "personal cloud" that identifies them in the converged online and offline worlds.*'

Consideration of current and emergent cybersecurity risks must also occur early in the software development lifecycle. Failure to understand the types of threats by designers and developers often leads to security flaws in software projects that are either costly to remediate or that place the owner at additional cybersecurity risk for the life of the product.

Both [9] and [21] attribute part of the difficulty with IT projects and software-intensive systems to limited understanding by engineers and managers of how to implement the emerging technology, too often leaving it entirely to software engineers and engaging these much too late in the process. Reference [1] extends this difficulty with software to the business and government leaders of such projects, while [22] extends that leadership concern to cybersecurity and [23] to cybersecurity in Australian Defence in particular.

Preview or pre-contractual test and evaluation of prototypes, where necessary using modelling and simulation, is key to de-risking projects [24]-[26]. Reference [27] outlines that software development has long been capable of rapid prototyping and they cite early research showing that user performance improves about 12 percent with each design iteration and that the average time to perform software-based tasks decreases about 35 percent from the first to the final iteration. While [27] is concerned for usability, [8] reinforces this same early approach for reliability, stating, '*The availability and continued development of tools for modelling SoS now provide a useful toolset for testing, evaluating and understanding the behavior of large complex systems in a virtual environment.*' For example, [11] explains how federated systems integration laboratories (SILs) connected by dedicated test networks and live, virtual and constructive (LVC) simulation capabilities have enabled the U.S. DoD to do early preview of modelled new systems in representative complex and interconnected operating systems-of-systems where they are intended to be used.

Adjusting ICT governance to these challenges has seen new standards, such as the ISO/IEC 38500:2015 that provide guiding principles for the members of governing bodies of both public and private enterprises in making decisions for their ICT use [28]. The ISO/IEC 38500 standard is limited in its guidance for developing cyber-resilient ICT through projects and through-life. There is an ICT governance support package called COBIT5 that provides a Performance, Compliance and Risk Control Framework for ICT project management [29]. This deeper and trademarked framework does not directly include benefits realization around cyber-resilience; at least not one that is 'test led' in the way proposed herein. That said, the quality framework of COBIT5 would likely adapt readily to provide such a test focus with appropriate regard to the other key governance factors.

Better governance frameworks of ICT projects need to under-pin readily available test capability for the necessary usability, preview de-risk and whole-life cyber-resilience monitoring to occur; however, research literature on such governance appears scarce. This scarcity is most likely driven in part by beliefs that extant project governance can be stretched or sped-up to cope.

III. AUSTRALIAN ICT GOVERNANCE EFFORTS

ICT project problems and cyber-vulnerabilities have not lessened the pace of advanced software functionality in all aspects of governments and society. Collectively these factors have seen renewed interest in ICT governance, from areas as diverse as program management offices (PMOs), departmental reform, and high-assurance security. Some of the proposed governance models considered have great complexity and isolation to ICT-only organizational structures in attempts to build prophetic and prescient oversight; while others, appeal to simplicity for success and leverage extant PMO reviews. Some governance models seek great rigor and acceptance before operational service, while others focus on the wherewithal for a life-long learning and development. Ironically, both these approaches of upfront rigor and through-life development, see the developing cyber-threat as reinforcing their approach.

A. Department of Human Services (DHS)

The DHS is Australia's administrator of all forms of social security and health payments and due to the high costs involved, works closely with the Australian Taxation Office. Reform efforts in these departments have been focused on automation right through to the customer (public) and exploiting the benefits of the paperless '*enter once, use many times*' approach to improve efficiency and effectiveness. DHS places a high priority on both governance and research. This is not just limited to ICT governance, cybersecurity also has a focused effort on governance and there are many policies that form an Information Security Management Framework. The Technology Innovation Centre is an example of the priority DHS places on research. Innovation across ICT and cyber-operations is key to delivering solutions that utilize the most contemporary and beneficial technology and processes. This includes investigating how new market technologies can be adapted to assist both customers and staff.

Understandably the DHS projects deal with large numbers of users (public) and require high privacy and security, so as to avoid exploitation at every level from individuals, through organized crime to state-based disturbances. Governance reform was led in these departments from around 2011 by the adoption of portfolio, program and project management offices (P3O) [30]-[32]. The P3O focus is evident from their slogan '*Right Projects, Right Way, Right Results*' [33]. According to [34] their P3O encountered resistance by individual projects such that a symbolic large-scale model of the process was built in the foyer with a funnel shape to reinforce projects would be culled or reset if necessary for excessive risk or poor reviews [35]. What emerged from this P3O is a governance model focused on delivering successful ICT projects through informed decision-making; which in turn, is based on evidence-based testing of four types as shown in Fig. 1 [36]. This elegant solution has limits that derive from its deliberately simple project-level portrayal, such as it ends when the project achieves business use without any through-life expression coming from a business handover. Also, this model does not deal with the other operational or legacy systems in the business ICT architecture, except insofar as the integration, cybersecurity and user testing discloses. There has been a focused effort to increase the test capability at DHS and in recent years a dedicated test director has been established as well as the opening of an advanced Cybersecurity Operations Centre. Future effort is on improving the representativeness of the operating test environment (i.e., SIL), particularly to model more realistic cyber-attack surfaces, both for in-service systems and developing systems as much as possible within their intended in-service architecture.

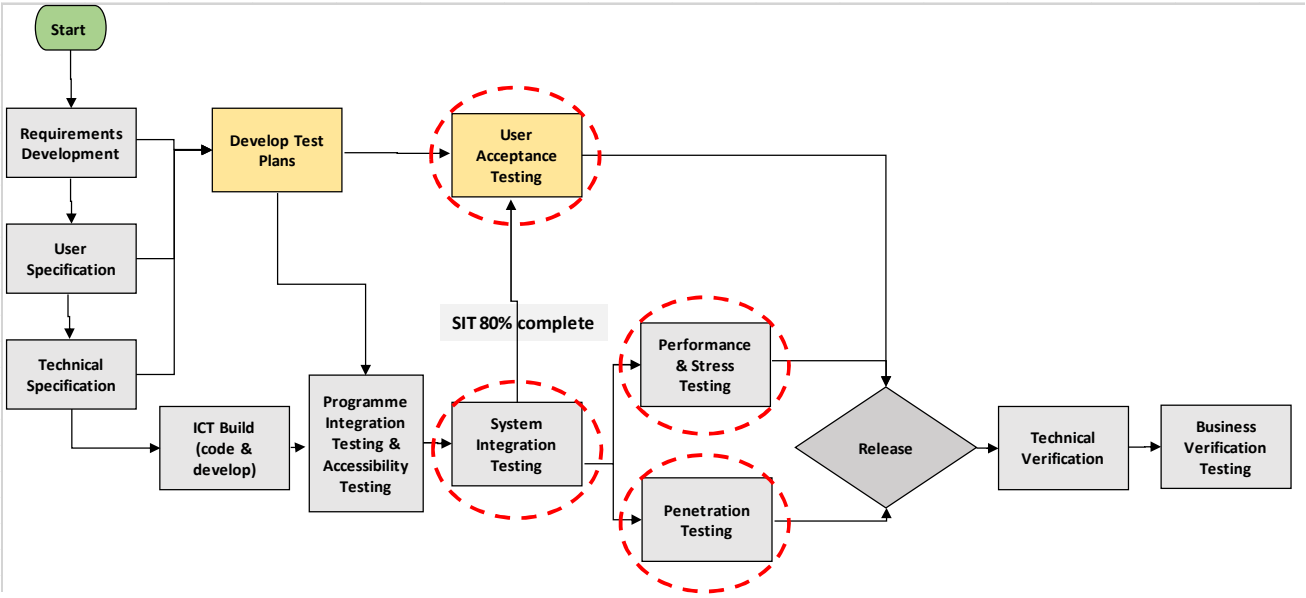

Fig. 1. Example of ICT Project Governance model (adapted [36])

B. *Department of Defence - generally*

The Australian DoD has also adopted a P3O governance model to its acquisitions following a first-principles review in 2014 [37] somewhat along the lines envisioned by [35], but so far without his proposed P3O accreditation or charter. The DoD differs from the DHS department in having a much smaller percentage of purely ICT projects, a greater complexity of interconnectivity between systems (i.e., families-of-systems-of-systems) [11], and many other competing acquisition domains for complex platforms like ships, submarines and aircraft. As such, ICT sits in acquisition and through-life operations as one of many disciplines in a matrix model, led by a Chief Information Officer (CIO) Group. The First-Principles Review sought to simplify acquisition policies and realign several different investment lines like estate, ICT and warfare platforms [37]. Unlike DHS, whose CIO is primarily responsible for the ICT projects, the current DoD governance structure sees the CIO have a significant role in managing a few ICT projects but as a specialist adviser to some 140 projects, 40 programs and five portfolios as required. The CIO's specialist role advising all acquisition projects and in-service portfolios is seriously challenged by rising demand and a paucity of complex ICT acquisition skills, especially in cybersecurity [38]-[39]. The demand is driven by the DoD's significant rise in software-intensive systems, its increasing cyber-threats ([14], [23]) and the increasing internet (/intranet) connectivity of its platforms.

Governance efficacy in such a CIO model is in the CIO primarily advising at scheduled project milestone approvals. Hence, a governance framework can be more about the decision-making approach that will pervade decisions no matter where they occur in the lifecycle or the program and project that is under review.

C. *DoD Research into Improved Governance*

A framework under development for the DoD is shown in Fig. 2 [40]. This model supports strategic alignment between business and IT for the creation of organizational value [41]. It provides an agile and benefits-driven approach to the governance of current capabilities and rapidly emerging and converging future technologies. Such technologies are not necessarily understood nor envisaged, especially with the advent of a new Synthetical Age [20] (or 4th Industrial Revolution [42]). The proposed framework is designed to support decision-making on investments on technological innovations that, while being disruptive, would be required in the organization's technology stack to generate benefits in the future [43].

Key to this decision framework is to understand that information systems investment does not provide any sustained advantage by itself, nor does it have any inherent value. Value is created by the organization's ability to convert and use the IT resource. Researchers call this *benefits realization*. Firms develop information systems to realize benefits after the implementation of the system [44].

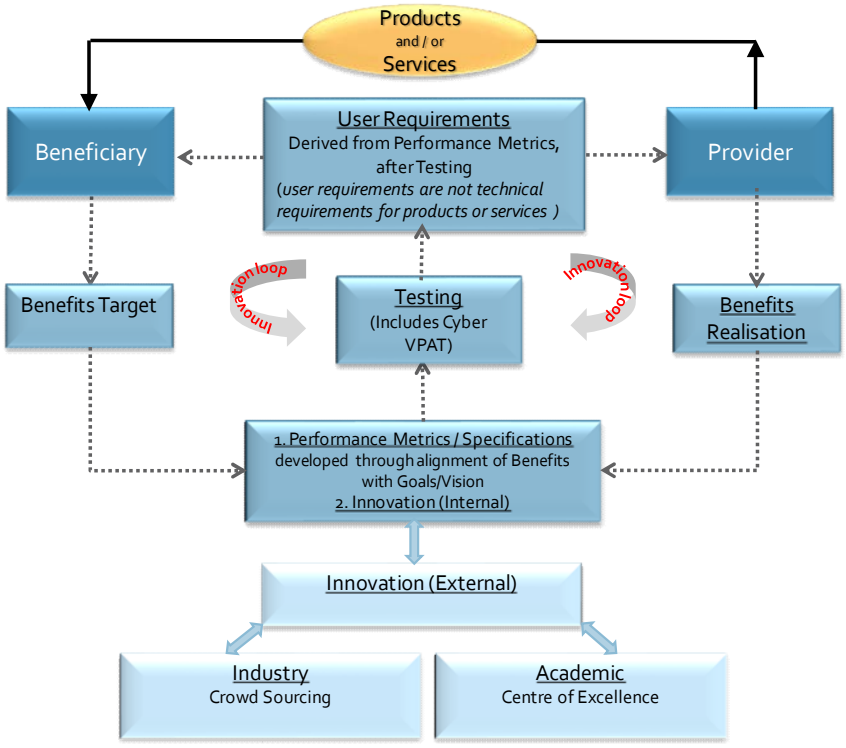

Fig. 2. Conceptual ICT Agile Governance Framework for innovation-led approach to benefits-realization (adapted [40])

The realization of benefits also comes from interventions; that is, changes to the way the business is conducted and how people work. There are two types of interventions – problem-based and innovation-based. In problem-based interventions, improvement targets such as return-on-investment form the basis of business case. The well-known approach of Enterprise Resource Planning is an example of problem-based intervention. In the case of innovation-based interventions, it is difficult to specify the end targets because there is uncertainty about the implementation success. This uncertainty implies that objectives and scope may well change during implementation, as the organization learns more about its environment and the evolving technology [46].

In the new governance decision-making framework (Fig. 2) the *Beneficiary* initiates the requirements for the new ICT by identifying the benefits that the new technology would generate. While the Beneficiary has a limited idea of benefits, perhaps gained from previous projects, the expertise on identifying, proposing and supporting the realization of benefits rests with the ICT *Provider*. Therefore, the important consideration here is to align the *Beneficiary's* expertise in the organizational goals or problems with the *Provider's* expertise in identifying, proposing and supporting the realization. When this alignment occurs, it should synergistically and iteratively achieve ongoing innovative solutions underpinned by persistent performance metrics, specifications, test and evaluation. The alignment occurs in a closed continuous "*innovation loop*," providing the goods and services along different phases of acquisition through sustainment and until the Disposal Phase. In this continuous innovation loop, the *Beneficiary* and the *Provider* benchmark innovations external to their organizations with the support of specialists, such as academia and the wider industry, shown in the bottom two boxes (Fig. 2).

This framework initiates the following best-practice arrangements:

a) Access to best practices and innovation through external agencies. The internal innovation loops are coupled to external innovation programs of research centres such as academic centres of excellence, industry R&D, and crowd sourcing. The external research centres should trigger further innovative solutions; as an example, using forecasting techniques like *reference class forecasting* [45]. These forecasting techniques also address the risks arising out of optimism bias and strategic bias situations where many benefits in the business case are overstated so as to get the project approved, leading to the promised benefits not being completely realised [46]. Use of the business principle of *incremental enlargement* [42] coupled with *reference class forecasting* would assist with identifying realistic benefits targets prior to each investment review.

b) Best practice contracting arrangement. According to [47] governance structures in a commercial environment will benefit from being of an "ongoing kind" where parties preserve cooperation during contract execution. He suggests a flexible approach with the "*contract as framework*" in contrast to the more familiar concept of the "*contract as legal rules*". The contract in Figure 2 can be viewed as a flexible framework and not a rigid one which often serves as a legal weapon, protective device, or heirarchy. The flexible framework allows collaboration and sharing of information that hopefully leads to reduced contractual overheads.

c) Collaboration between three groups. This framework, should bring about a partnership of three groups – the organisation that desires ICT-led change, ICT industry (includes the provider), and an ICT academic research organisation and/or other expertise such as crowd-sourcing.

This DoD-funded research into ICT governance has found the need to focus projects on demonstrating compliance to the benefits approach through the four key ICT testing areas outlined later in this paper. This is because benefits inherently involve the same areas of usability, integration with the in-service operating environment, network performance, security and cyber-resilience, as well as important trade-offs between these benefit characteristics. This research is now focusing on characterizing the governance approach across the ICT capability lifecycle and the necessary tailoring for capabilities with differing levels of software-intensity in the systems.

D. DoD – High Assurance Review

Concurrent to the broad ICT governance framework research just outlined, the Australian DoD has been reviewing its governance of high-assurance ICT capabilities in support of many other government departments [48]. Such capabilities must be based on products that have undergone a high assurance (HA) evaluation, characterized by a rigorous investigation, analysis, verification and validation of the products or systems against a stringent information security standard, in this case the DoD's Information Security Manual (ISM), in order to protect highly classified information. Such capabilities have historically been assured through High Grade cryptography — the processes and standards that evolved from the experiences of World War 2. Over the years, these ICT security evaluation processes and standards have evolved, divided and come back together. In 1985, the so-called Orange Book [49] contained the U.S. DoD's Trusted Computer Systems Evaluation Criteria, which was the first widely released systematic set of standards for securing computer information systems. It was influential among U.S. allies as the basis of national standards. By December 2000, the Orange Book was retired being effectively subsumed into the so-called Common Criteria published by the U.S. National Institute of Standards and Technology (NIST) [50]. A parallel set of processes and standards have developed in the U.S. [51], U.K. [52] and Australia [53]. All of the approaches to HA have two aspects in common:

- *Compelling evidence.* HA is a property of the evidence, not the system. It also makes assumptions about the independence and expertise of the entity evaluating the evidence.
- *Specified requirements*. HA needs requirements that are simple enough to be analyzed in a reasonable

time and are refutable. This makes it possible to evaluate whether the design satisfies the requirements, in other word is effective, and whether the implementation matches the design.

To satisfy these approaches the governance used to manage systems protecting highly classified information evolved into a set of prescriptive policies applying to discrete security compartments, where isolation was the main security enforcing mechanism. The growth in the demand for real-time collection, processing, exploitation and dissemination of intelligence, targeting and geospatial information from increasing numbers of capable collection assets, has seen much of the HA edge, if not eroded, certainly outsourced and at greater risk of compromise. Such risk also derives from the growth and reach in sophisticated cyber-threats that contest the Western pursuit of information dominance [11] [14] [54]; or put another way, a joint and networked force [55]. The underlying cause of this growth is from the fact that all new Government capabilities have a strong ICT component. For the Australian DoD, this has meant a significant increase in the number of systems designed to secure highly classified information or connect to other systems that protect highly classified information.

Currently the approach used to assess the security of systems protecting highly classified information has not been able to keep up with the demand [48]. The increasing HA demand and the changing nature of ICT led the U.S. DoD 15 years ago to develop an improved HA evaluation methodology [56]. Other allied nations have generally not followed suit and this has arguably led to a general weakening of their comparative ability to evaluate and certify the security of systems protecting highly classified information. This is despite a number of research and policy efforts over the years to improve HA efficiency and effectiveness, such as policy initiatives like approving public domain cryptographic algorithms for protecting highly classified information and the ongoing research into high-trust techniques like formal methods.

To address the demand issue the recent Australian DoD HA governance review [48] found that the HA responsibility needed to be spread across the P3O reviews and be more clearly focused on benefits realization through informed ICT testing. The necessary test areas were found to be the four areas outlined later in the paper, albeit some being more specialized, in-house and secure. Specifically the HA review found that in order for P3Os to deal effectively with HA, HA must scaffold more into the whole ICT life cycle.

To inculcate HA and security more broadly into all aspects of ICT, two sets of processes are proposed: one set to influence the behavior of the ICT life cycle and the other set to measure, test and evaluate the security performance throughout the ICT life cycle. The first process set is known variously as supply chain management or Information Security Industry Engagement (ISIE). The second set of processes is generally known simply as test and evaluation, though in our case we should specify the purpose as conducting an Information Security Evaluation (ISE). These two process sets interact and feedback upon each other, with the ISE providing the compelling evidence and the ISIE

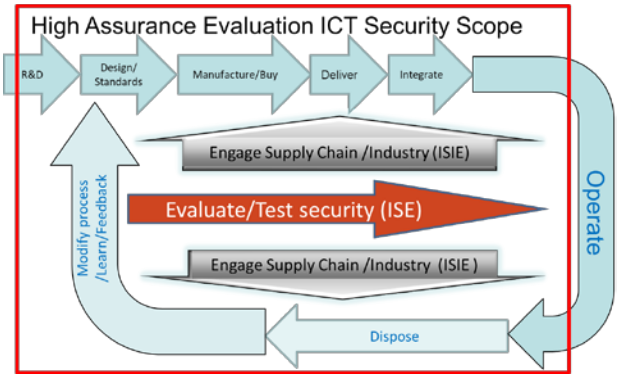

Fig. 3. ICT Security Cycle & HA Evaluation Scope (Red Box).

managing the specified requirements. Feedback and interplay between ISE and ISIE processes can be complex, where overlapping boundaries abound. For example, with ISIE the broad aim is to manage ICT Supply Chain Compromise per the concerns outlined by [57]; in other words, to manage an occurrence within the ICT supply chain whereby an adversary jeopardizes the confidentiality, integrity, or availability of a system or the information the system processes, stores, or transmits. One way to manage this is to use Formal Methods [58] as a tool to specify in a mathematically rigorous way the security requirements. The process of proving that these specified requirements meet the security objects also provides the evidence needed for ISE [18] while also managing the supply chain issues [57]. The HA Evaluation process can be summarized as the functions within the red box of Fig. 3.

IV. FOUR TESTING TYPES

This paper will now clarify the four ICT testing types assessed as crucial to informed project governance for evidence-based benefits realization, sound integration, consistent cybersecurity and thus ultimately more cyber-resilient operating environments (families-of-systems-of-systems [11]). Each type of testing will briefly examine the unique test design and analysis skills that are needed.

A. Usability Testing

Software performs functions for systems replacing both mechanical systems and human operators alike in a continuous frontier of increasingly complex heuristics that also includes new language development, new processing and proprietary boundaries. As such, it is rare that software in systems technology ever repeats functions in precisely the same way to the same purpose and for the same user. Human-machine interfaces have been well researched since computers evolved [27] and this research has clearly shown the efficiency and effectiveness benefits of usability testing that were cited earlier, including standard usability test metrics. Yet, ICT projects abound with poor performance stemming from under-researched user requirements [1] and from the authors' experiences they rarely use structured and iterative usability testing as shown in the software development cycle at Fig. 4.

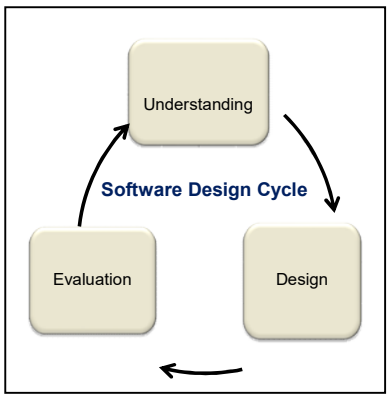

Fig. 4. Software Design Cycle (adapted [27])

Where user interface requirements are documented and the software tailored to those, they are often then tested to those functions for correct coding by other software, as if they were now axiomatic. Functional testing is usually the unfortunate substitute of better metrics of user efficiency, effectiveness and satisfaction that should be tested iteratively for improvement. Hence, software functionality is often operationalized with a high degree of expectation mismatch. According to [27]:

'*The fact that computer software is sometimes poorly designed and therefore difficult to use causes a variety of negative consequences. First user performance suffers; researchers have found the magnitude of errors to be as high as 46 percent for commands, tasks and transactions in some applications. Other consequences follow, such as confusion, panic, boredom, frustration, incomplete use of the system, system abandonment altogether, modification of the task, compensatory actions, and misuse of the system … The trend toward a greater number of functions, called creeping featurism, is an important problem because the additional functions make the interface more complex and increase the number of choices a user must take.*'

The authors' experiences are that most project managers are not software specialists and that when usability testing is described to them, their greatest fear is that the requirements will increase (creep) and that even unwanted features will cost to be taken out. Most project managers also do not realize software can be virtually modelled and thus it can undergo the first usability testing even before the code is written, ideally even before full developmental contract [24] [25] [60].

One example of the benefits of the P3O-led usability testing can be seen in the successful digital linking of taxation and social security in Australia. By contrast, a large software project in the Australian DoD for a command support environment debuted 20 major applications in an operational evaluation where only one had received something approximating usability testing (one iteration) and only half were suitable enough to continue use, albeit commencing with usability upgrades. In the case of one application, the live evaluation found just three people whose function would use the application and none had been consulted previously in the development. Similarly, a major battle-management system in the Australian DoD debuted operationally using an off-the-shelf foreign application that did not adequately match the military culture or organizational structures, and while that rapid evaluation did ratify the broad benefits of digitizing battlefield plans, a redevelopment using three iterations of representative usability testing has since fundamentally improved acceptance and effectiveness. Furthermore, the redevelopment has enabled an assessment of, and improvement in, cybersecurity, which was not envisioned in the original off-the-shelf project.

Good usability testing fundamentally contributes to cyber-resilience. Modern context-dependent security-monitoring algorithms (i.e., beyond 2-factor) only work if you can accurately capture what normal use looks like. Unusual and unnecessary features simply provide avenues for malicious exploitation by increasing the attack surface and in high-assurance applications, will require repeated cybersecurity checking against each wave of new threats. Additionally, security controls that are not user friendly often result in the users finding ways to avoid having to comply with security policies; an overall degradation of cybersecurity in a system that in its pure design was more cyber-resilient.

Usability testing must also continue to be undertaken in applications where "human-in-the-loop" functions become "human-on-the-loop" and then "human-out-of-the-loop." New standards are being developed to specifically cover intelligent and autonomous systems and how they are to be ethically designed [61], where the operational commander needs the usability testing to accurately ensure and record that the systems use always accords with human intent. Such ethical design will require the structured experimentation with a representative sample of commanders that comes with usability testing. While that sounds very military, such autonomous intelligent systems are set to rapidly pervade health care and other domains via the internet-of-things (IoT) [62]. Within Australia's public sector areas there have been too many examples where policy initiatives have not had sufficient preview or trial [63]-[64] — early usability preview of software functionality prior to developmental contract is key to such de-risking strategies being evidence-based. For example, DHS has developed dedicated user testing suites that bring members of the public in to examine and participate in such aspects as application development, online design and user interfaces. This enables early engagement with the very people that will use the final systems. Understanding their needs and constraints and integrating that knowledge with in-house developers assists in ensuring the products are fit-for-purpose upon initial release. As a delivery agency to almost every Australian there is a balance between placing cybersecurity measures (multi-factor authentication, password complexity rules etc) embedded in the user interface to attempt to demonstrate and guide the cyber-behaviors of the individuals and moving cybersecurity features behind the firewall to prevent the user from being inconvenienced by them.

Usability testing metrics and constructs are not difficult test design skills and methods to manage and an introductory text is [27]. It is advisable to have contract coverage of the usability testing and at least three to five iterations, as well as some method to ensure representative sampling of the users. As mentioned earlier, at least the first de-risk usability testing should be pre-contractual (i.e., offer definition

activity in tendering) based on virtual modelling (pre-coding), especially if some applications and operating systems are claimed to be off-the-shelf (i.e., mixed maturity architecture). Preferably at least one person in the test team needs to have had university-level education in introductory human factors and another in human-computer interfaces (i.e., graphical user interfaces). For efficiency of determining significant factors and identifying confidence levels, at least one of the test design team needs to be experienced in design of experiments (DOE) test design techniques (i.e., [65]), especially efficient screening of main test factors with high throughput testing (HTT) based on two-way combinatorial test designs [66]. Ideally, that test designer should be the person qualified in human factors, as usability testing should not be diverted to center on requirements verification or be overlaid with more technical testing. If original user requirements or functionality are verified in usability testing it should not be disclosed to, or constrain in anyway, the representative users. Subsequent iterations of usability testing could move to more classical DOE test methods as the non-significant factors are removed from analysis and test performances converge on the acceptable user performance metrics.

B. Systems Integration Testing

The greatest difficulties in system integration testing are: first, efficiently dealing with the multiple configurations of hardware and software configurations present in distributed ICT architectures [67] and second, having representative operating systems from the system-of-systems or families-of-systems where the new system will debut [10]. The ideal first step is having a P3O culture to deliver virtual models of new software prior to full developmental contract that can then be tested in a software integration laboratory (SIL), software system support center (SSSC) or distributed live-virtual-constructive (LVC) test integration network [11].

Organizing and sustaining a representative test architecture for evolutionary improvement faces many proprietary and funding challenges, as well as the challenge of staying representative when portions of the network exist in the public domain, such as the internet and self-funded public users. Given the existential and also evolving cyber-threat to the public sector [22] and financial industry, proprietary outsourcing of parts of their ICT architectures and ad hoc approaches to other parts are resulting in increased vulnerabilities being introduced during ICT build and integration that will become too risky at some stage in the future. Proprietary support will need renegotiation so as to be representatively maintained in a SIL / SSSC / LVC test network, and to be independently and regular tested for vulnerabilities against cyber-threats without disclosure of the test methods used. Such arrangements fundamentally challenge corporate and public sector outsourcing models for ICT from the last few decades and associated fixed-price and intellectually-protected contracting arrangements, in favor of more cooperative security arrangements and fee-for-service. Reference [68] assesses that the contractual and project methods used for system safety can be leveraged to achieve this greater flexibility. Certainly, more mature process instantiations for cybersecurity testing like the U.S. DoD [69] and industry [70] are forcing cooperative flexibility like that hitherto seen in safety. This becomes particularly difficult at scale and DHS is an excellent example of the constant effort required to maintain test and development environments that are representative of the production system.

The Australian DoD is proposing battle-laboratories of mixed SoS for greater integration [71] and there are options to cost-effectively leverage the U.S. DoD LVC test networks [11]. While the main forces returning departments and major corporations to such in-house testing are the risks in cyber-threats, for system integration testers this is a welcome reprieve from trends towards disparate sub-system testing and limited opportunities for SoS testing with high operational risks [8]. According to [72], U.S. DoD project managers have often been focused at narrow capabilities, lacked program or portfolio support to consider a wider good, facing cost and schedule pressures, and of very limited tenures compared to the capabilities they deliver. The advent of P3Os, the cyber-risks, and the need for representative integration centers will provide much needed focus and support for such ICT project managers to be fully informed in the way envisioned by governance models at Figures 1 and 2.

System integration test skills demand high levels of ICT integration knowledge around fusing of applications, operating systems and datalinks, as well as management of regular and emergency demand, cyber-defense; and for DoDs, how to expand and contract services based on military priorities and deployment. Such advanced technical skills and experience often comes at the expense of formal education in test design and analysis, such that tests run by integration experts alone are usually successful but sub-optimal. At least one of the integration test team should be educated in test design methods so that integration tests are efficiently screened early to focus subsequent testing on the significant factors. The HTT combinatorial methods for all two-way combinations are ideal for functional testing in an integration environment [7] [78], followed by a focused modelling on interactions involving significant integration factors with a more orthogonal test design. Combinatorial test design packages with algorithms to optimize orthogonality offer combined efficiency and diagnostics to the software integration industry with genealogies in both Japan [73] and the U.S. [74], but they require higher education of the testers than random (fuzz) test methods, and understanding by the beneficiaries, who associate fewer test runs with greater product risk. Significant dedicated test experience is required if testers are to use the efficiency of combinatorial methods at the meso-level with the investigative advantages of the fuzz methods at a micro-test level.

C. Performance and Stress Testing

Performance and stress testing of the integrated and usable ICT or software-intensive architecture is necessary to ensure manageable demands during the full range of usage cycles. User satisfaction, system stability and effectiveness

are inherently linked to timely performance. Not only is maintaining the networks challenging but also producing replica test load. For banks with 3.5 million online user authentications making 2.5 million payments per day, and departments such as DHS with over 26 million users and approximately 600 000 authentications daily, it is difficult to reproduce that level of load in virtual environments. Some system errors can only be replicated under load and testing for these errors is complex and expensive. The only other alternative is to run scripts in production environments whilst under the assessed load, which significantly increases the risk. If there are errors in the code and they do then appear during load testing the associated error is then potentially exposed to all those using the system at that point in time. For large systems, this risk is often too high to accept.

ICT reports are, like many military systems, often replete with the use of averages, sometimes inappropriately aggregated across diverse mission scenarios without regard for the underpinning statistical distributions and appropriate confidence limits. Like stress loads, it is difficult for banks and DHS, due to the scale, to produce accurate reports on performance during the testing phase due to the inability for all components to be integrated, combined with the absence of load. The U.S. DoD has worked hard to improve test plans and test reports to deliver better reporting of performance metrics [75]-[76], including in the presence of cyber-threats and varying cyber-defense postures [11] [77]. Better education of ICT testers in test design and analysis techniques offers not only efficiency gains in the testing, but better rigor in reporting results and managing with operational models the cycles of demand. Operational models should not simply be based on deterministic predictions, but backed by probabilistic performance test modelling. The skills necessary to do this are readily available in most six-sigma industry accreditations [65] [66].

D. Security Testing – Vulnerability and Penetration

The pervasive cyber-threat to DoDs, public sector, finance and industry means cooperative vulnerability and penetration assessment (CVPA) is no longer an option but rather about managing an acceptable risk of how much testing is enough [7]. Not testing, simply means not knowing, and thus an unassessed risk, while not re-testing fielded systems at some interval means an atrophy of security confidence at an unknown rate [77]. Additionally, over time, systems acquire aggregated cyber-risk as different elements of complex systems are deployed. In critical systems, operators are mounting continuous defensive cyber-operations, sometimes extending to supply-chain monitoring through-life [57] [68], but in Australia these precautions are largely only on live networked systems [23]. Outside the U.S. and particularly the U.S. DoD, there is still limited understanding of the risk of cyber-threats to software-intensive systems that are only occasionally updated or networked. However, recent DoD testimony to the Australian Senate announced a program of what is being termed "*cyber-worthiness*" of capabilities [79], hopefully following research recommendations like [80].

Public sector and financial systems are vulnerable to more sophisticated probing and logic disruptions that can now be electromagnetic lodged at low power with no connectivity [23]. Without CVPA and some defensive posturing even for fielded legacy systems, significant risk exists that at a time of a potential enemy or criminal entity's choosing, systems will be denied or interfered without detection for an unknown period of malicious intrusion [78].

The cost of mounting expensive CVPA and defensive capabilities will be borne by either a slower pace of computer-based services to the public and DoD capability, or increasing market differentiation. Reference [11] documents a widening difference in systems integration and cyber-resilience between DoDs of even close allies like Australia and the U.S.. While the cybersecurity of two militaries might seem irrelevant to much of the public sector or critical industries, the reality is that such differentiation as that described therein can soon be expected in public sectors and markets. Regular CVPA on representative operational test architectures (i.e., federated SILs) is needed for as much known threat as possible. The capability to do so needs to be introduced and funded at a portfolio-level, so as to enable informed decision on each new system release and collectively how to strategically posture resilience of the operating systems. It may be that for some services, risks are low and public or consumer risks can be tolerated, even deliberately targeting cheaper or efficient services with perhaps a greater explanation of consumer and public risk. Whereas for other services, capabilities may be compromised and costs raised to enable greater cyber-protections. As always, not testing is not knowing and that means no informed choice.

Militaries have a unique advantage in that many of their combat systems can be isolated to a certain degree from the internet. This is not the case of other Government services such as the ATO, DHS and banks, where their core business is linking Australian citizens to payments and services through the internet. The use of intelligence provided by military counterparts in persistent threats is of great benefit. Ultimately, the threats that only a few years ago were aligned to largely espionage or a criminal intent have now converged.

Forming a CVPA test capability is dramatically easier if the other ICT test capabilities (i.e., usability, integration & performance) are robust and appropriately part of project governance and a benefits-realization decision-making culture. Inevitably in capped schedules and budgets, increased cyber-resilience involves trade-offs between:

- user requirements, such as determining through structured test what users value more;
- ICT build, such as limiting use of code libraries to those known to be cyber-secure;
- integration, such as limiting connectivity to limit cyber-threat exposure; or
- performance, such as increasing the threat detection algorithms and reducing system processing for main functions.

Having the other ICT testing well run and iterative, as shown in governance models at Figures 1 and 2, enables

Governance to make these cyber-resilience trade-offs in an informed way. The next level of maturity for large organizations is to combine them and ensure CVPA is integrated into the formal testing cycle. However, to be robust, it should also be conducted to systems at regular intervals post-delivery, with appropriate levels of funding set aside to address the vulnerabilities that are then identified.

The test design, test analysis skills and test infrastructure required to manage CVPA testing are, with only a few key additions, supported by the test skills and test infrastructure of the other ICT test types. For example, industries, public sector and government departments that have invested in SILs, SSSCs or LVC test networks can adapt these to allow for multi-security CVPA testing — in essence extending integration and capability upgrade infrastructure to be cyber-ranges that can concomitantly manage evolving cyber-threats and deliver greater cyber-resilience. If such infrastructure has been outsourced and is proprietary, then contractual changes will be needed to safeguard connection to government-managed representative cyber-threats. For example, the Australian National Audit Office (ANAO) assessed that DHS had security controls in place to provide protection from external attacks, internal breaches and unauthorized information disclosures [81]. This was achieved by prioritizing activities that were required to implement the top four Australian Government mitigation strategies and by strengthening supporting governance arrangements. This prioritization was largely enabled by the in-house capability that DHS possesses and the lack of reliance of contracts and service providers. Similarly, the challenge for Australian banks is to be compliant with the Australian Prudential Regulation Authority (APRA) and Sarbanes Oxley (SOX) by implementing the top eight mitigation strategies and establishing a cyber-resilience culture.

Similar to test infrastructure, additional test design skills can be added to integration and performance testers to manage the additional rigor necessary for testing cyber-resilience. Again combinatorial test design has been instrumental in achieving greater cyber-resilience with three-way through to six-way combinatorial test rigor being achieved, often while deriving new efficiencies [6] [74] and other defect-protection rigor. An example of this approach is the industry six-sigma software testing award overview by [82]. The University of NSW [66] has adapted test design education to give early awareness of these additional cybersecurity test techniques using the freeware by [74] as a reasonable simplification of the test design packages used by big software industries [73].

Industry and departments have been slow to adopt another protective process layer, which has led industry bodies to develop minimum additional cyber-planning and testing checks to overlay standard systems engineering [70] [83]. These process links and explanations offer the greatest promise to normalize cybersecurity in industry, albeit that industry using system engineering practice.

Probably the last and most difficult extension for CVPA testing from hitherto ICT testing, is the skill of defensive (blue) and penetration (red) teams war-gaming the cyber-threat as described well by [78]. These are military skills applied in a new domain and unfortunately necessary for public sector and critical industries to adopt if they are to be reasonably defensive to malicious threats. Legal protections in cyber are a long way from being instituted [14] [22] [84] and deterrence critical depends on timely attribution, which unfortunately remains difficult. Even if legal recourses become viable, public sector and industry war-gaming is necessary at some level for the defensive capability to exist to collect evidence for legal recourse.

DHS has proven this applicability outside of DoD. In 2017 they ran the first government cyber war-games on a cyber-range built in-house and representing a fictional city. Ten departments and agencies combined to form five teams that conducted both defensive and offensive play and were assessed on skills outside the technical, such as teamwork, communication, leadership and critical thinking. To be able to defend, understanding how to attack is critical. Ultimately, it is another human behind the opposing keyboard and being able to understand how they may manipulate the systems to maliciously achieve their aim will ultimately direct a diligent defender to monitor, protect and defend the right elements of the system.

The other skill that is difficult to build, and also tested during the DHS cyber war-games, is the ability to translate the technical nature of cyber-operations to both the boardroom and the media. In times of cyber-disruption the ability to deal with the media in what is inevitably an uncertain time, where the nature of network problem (network outage or cyber-attack) is unknown, is another complex skill. In large organizations and government departments having a technical team to conduct CVPA integrated with an engagement team capable of doing that media translation is key.

Building CVPA test skills and infrastructure requires education to be improved to merge the necessary knowledge and skills into industry-accredited packages [38]. Having teams of testers that are able to conduct the required testing is an initial start. Having testers that are able to schedule tests to match the development schedule is the next step. Moving testing to the left in the software development lifecycle and conducting CVPA throughout development is even better [77]. But ultimately, designing and building cyber-resilience in, by having cyber-operations staff embedded with both the design and development teams from the start, is key. Ensuring that at the first conceptual design any ideas that will invoke cyber-vulnerabilities are discussed and the risks are clearly articulated to cyber-aware business owners early [78], such that testing throughout both design and development is combined with training the developers. Outsourcing, offshoring, and high turnover of developers all magnify this challenge. Often similar mistakes are seen multiple times because developers aren't made aware of new and emerging cyber-threats and how the way that they code allows cyber-criminals and state-based actors alike to exploit those flaws.

Having centralized code libraries sees any vulnerability that has been introduced exponentially deploy through the network as code with security flaws is drawn from a central library. Automated code scanning, during ICT develop and

test, is an effective way to assure absence of known code vulnerabilities. There is constant tension between cost and benefit in testing and in particular CVPA. All organizations and departments have differing risk appetites and what may be acceptable to one will not be at all palatable to another. There is a significant cost in increasing the cyber-resilience of any organization or department, however, the reputational cost if information or systems are lost or exploited, in most cases, far outweighs the required investment to secure it. Technology is only one element however, without equal investment in the people and an understanding by the beneficiaries of the business implications, the CVPA system will never achieve full maturity. A layered security approach has to be designed to be both complex and obscure [70] [83].

Specific CVPA test design and analysis skills of the types outlined by [66], [78] and [83] need to be available in country and *en masse*, but this requires industry and public sector to commit to their staff undertaking the education and placements. Furthermore, this requires governance structures and awareness regarding the necessity for cyber-resilience and the wherewithal to achieve that through CVPA testing.

E. Security Testing - High Assurance Evaluation

The aim of information security evaluations (ISE) is to make sure that the effort required to defeat exceeds the value of the material being defended. An evaluation aims to measure that effort and compare it with the value of the protected material and the resources available to a likely threat. In general both active and passive exploitation requires all three of the following factors [48]:

- a vulnerability to gain the initial access;
- an implant or processing system to retain access; and
- a communications system to manage the command and control (inward) and an export means (outward).

The aim of both policy and technology is to block at least one of these factors. Essentially all threats exploit failures of either policy or technology in these three areas.

The aim of ISE is to show that no security failure state exists by demonstrating the nonexistence of known or likely failure states. Proving the non-existence of something is generally not possible and so in most cases ISEs measure the likelihood of the non-existence of something by searching for it and not finding it over a period of time. The longer spent looking and not finding, the more likely the non-existents' case is. As such, the level of trust or assurance one has in the system is proportionate to the effort expended in trying to defeat a system and failing.

The most common approach at lower security levels is to use process and procedures to systematically search for failures in policy and the design, rather like a check list. Most schemes and standards, such as [50] have this property. They have a list of items or controls that are needed to be enacted and checked systematically to determine if the system matches the policy. This has the advantage of making the effort required to secure a system easier to measure and manage, albeit through the rigor of compliance. Active searches for security failures, such as CVPA, also use lists of known failures and threats to see if they exist in built systems, but in most instances with a degree of war-gaming above that of systematic compliance. Such active approaches are harder to cost and outsource, due to the complexity of the failures being searched for, but it has the advantage of identifying new failures due to combinations of known failures and human ingenuity. Active approaches can therefore be difficult to justify and maintain for highly complex systems and those at scale.

Improving the governance of ISEs is the key to being able to have visibility of total risk across all systems. For example, being non-compliant with any of the set ISE controls does not explicitly lower the cyber-resilience of the system, however aggregated across many systems it may pose risk in areas not considered in isolation.

In higher security levels, such as the one used to evaluate HA equipment, the approach is different [51]-[53]. An unfettered search for a failure is conducted, and then for all the ones found, a theoretical attack is developed and then costed, using a rigorous well-tested method based on the HA standard. If the cost of the attack exceeds the value of the material being defended then the system is said to be secure. The critical part of this approach is the costing model. This model, developed over many years, determines over time: 1) the value of the material being protected, 2) the resources of an adversary and 3) the resources required to run an attack.

The unfettered search for a failure examines two aspects of security being the design and the implementation. The two metrics used to measure the effectiveness of these security aspects are: 1) the cost to defeat security, and 2) the effort undertaken looking for a new defeat and not finding one. The cost is measured in terms of resources, such as effort, money, knowledge etc. and the time required to defeat the security. The effort is measured by the number of people months spent examining the system and not finding a new defeat.

The current Australian HA standard [53] contains a number of built in parameters around the investigation effort, the resources an adversary has and the length of time required protecting the system or information. For example, it assumes that highly classified information needs to be protected for 30 years at least, from all possible organizations and to spend from 6 to 24 person months, depending on the complexity of the system, not finding a new defeat. For less highly classified systems, it assumes that 3 to 12 person months have been spent showing that no organization will have the resources to defeat the security for 10 years.

A key difference between the HA and general ISE evaluation methods is that, the HA one focuses on the resources required and available to exploit security failures and defeat a security system, while the general ISE one focuses on going through a list of possible security failures and removing any present. The HA one is hard to plan, requires skilled staff, but has proven to be quicker and very effective. The general ISE one is easier to plan and requires less skilled staff but is less effective, takes longer for higher security levels and can be difficult at scale.

Where the logic is not based solely on classification it is possible to combine the two methods. This can be done by evaluating the key cyber-terrain of the organization or department. In order to know how much to invest in cyber-

resilience it is important to know where to channel that funding and effort. By understanding what parts of the network may be attractive to state-based actors or cyber-criminals it is possible to conduct those methods for HA against a small subset of the larger network. An indicative ISE method is the Attack Cost Method summarized as follows from [48].

a) Method. The Attack Cost Method is a search for the best attack that will defeat the security, then that attack is costed to determine when that attack becomes possible. There are two basic approaches to searching for the best attacks or vulnerabilities. One starts from first principles assuming nothing known about the system or device. The other approach uses security assessments from multiple sources, combining the result across the whole system. Both of these approaches are used. The method assumes that an adversary has a finite number of resources measured as money (R). It also measures the payoff in terms of plain-text documents equivalents (p) where one highly classified document is 1.0 p and one less highly classified document is 0.1 p and so on, and there is an estimated expected payoff (Pe). Cost of the attack (C) is also measured as money. So the basic idea is that for a secure system the cost (C) of all possible attacks exceeds the resources (R) available to an adversary, or the cost per plain text documents equivalents (C/P) or the cost per plain text documents equivalents (C/P) exceeds the expected cost per plain text documents equivalents (R/Pe).

b) Attack-Cost Method steps:

- Develop an adversary model and determine the adversary resources (R) and expected payoff (Pe).
- Develop a usage model or scenario and determine the value or payoff of the user's data (P) and how it will be used.
- Launch an unfettered search for vulnerabilities.
- Develop attacks from the vulnerabilities and detail the attack proving that it exists.
- Rigorously cost the attacks using a costing model developed with respect to context, calculate the cost (C) of the attack over time.
- Using the cost (C) calculate the payoff (p) and resource limits (R) over time, noting that over time R goes up and the value of the payoff goes down.
- Repeat until required assurance level is reached.
- Write up report and recommendations, put comments to the manufacture and have a trusted third-party review the report and evidence, certify the results and note any improvement and maintenance plans.
- Acceptance by the Accreditation Authority.

c) Indicative set of assurance levels are:

- 6 person months of not finding an attack for a highly classified level of assurance.
- 3 person months of not finding an attack for a less highly classified level of assurance.
- 1.5 person month of not finding an attack for a moderately classified level of assurance.
- 3 person weeks of not finding an attack for a classified level of assurance.
- 1.5 person weeks of not finding an attack for all other levels of assurance.

V. FUTURE RESEARCH AND WORK

All three of the research avenues described here are still ongoing and the collaboration to compare findings will continue under the auspices of the University of NSW Australian Centre for Cyber-Security (ACCS). Each of the authors is passionate about improving cybersecurity education along the industry-accreditation lines outlined by [38]. As such, the collaboration will hopefully have feedback from test practitioners in each of the four ICT test areas based on their experiences undertaking closely mentored and industry-placed research assignments. Australia's efforts on cyber-testing is seminal and so early industry-based feedback will be crucial to build the experience base around the ICT governance frameworks, so as to confirm what works well and what does not, especially for cyber-resilient systems. Countries with similar challenges to Australia in ICT governance are welcome to leverage the research collaboration.

VI. CONCLUSION

Difficulties with ICT projects abound in all parts of the World, with research reporting as many as one in six such projects exhibiting cost and schedule overruns in excess of 200 percent. There are also reports of many high-profile ICT projects experiencing high incidences of unexpected cyber-vulnerabilities. These project problems and cyber-vulnerabilities have not lessened the pace of advanced software functionality in all aspects of governments and society. Collectively these factors have seen renewed interest in ICT governance, from areas as diverse as program management offices, departmental reform, and high-assurance security. Some of the proposed governance models considered have great complexity and isolation to ICT-only organizational structures in attempts to build prophetic and prescient oversight from only brief project reviews, while others appeal to simplicity for success.

Three separate and diverse Australian Government research efforts in ICT governance, as well an assessment in the Banking Sector, have found similar concerns about the importance and type of ICT testing and test expertise critical to ICT governance and the ability to build cyber-resilience; namely, usability testing, systems integration testing, performance testing and cyber-security testing. These research efforts all found that ICT Governance critically depends on: (1) information coming from all four types of testing, (2) some test understanding in management to appreciate fully the outputs, and (3) that such test capabilities must be enduring (i.e., through-life, however short) so as to provide a sufficient degree of commercial and architectural independence to make hard and timely decisions.

These lessons on the importance of testing to ICT governance seem almost to have been forgotten in a rush to be technologically and managerially adroit, yet if done as outlined from these research efforts, could see a resurgence in test-informed project reviews that are: (1) innovative, (2) give lower risk competitiveness, and (3) greater cyber-resilience.

Key conclusions of this research are:

- A benefits-approach to ICT governance as shown in Fig. 2 should give more cyber-resilient operations through informed ICT capability life-cycle decisions.
- Usability testing is crucial to user satisfaction and needed even when software-intensive systems seek to replace an operator or commander.
- Development contracts should cover three to five usability test iterations, with the first iteration ideally being on a virtual software model prior to the development contract, so as to de-risk project scoping.
- Test teams need human factor engineering expertise to successfully conduct proper iterative usability testing as well as a governance culture of refining user requirements.
- Integration testing critically depends on a representative operational test environment such as a SIL, SSSC or LVC test network to be effective with significant parallel benefit to then extend such infrastructure cost effectively to do proper full cyber-attack surfaces in CVPA testing.
- The high number of permutations in integration and later performance testing requires test design skills in combinatorial HTT to be efficient. Six sigma test courses with practical competency assessments in industry are key to realizing such efficiency benefits.
- There is a balance between embedding good cyber-culture in the user interface to teach good cyber-behavior and moving cybersecurity rearward so as not to inconvenience the user.
- Sound ICT test infrastructure and test skills in usability, integration and performance testing, backed by project governance and benefits realization in the ICT test types, are crucial determinants in the preparedness and ease for CVPA testing to be incorporated and evolve for cyber-resilient systems.
- A CVPA test capability needs some additional combinatorial test design and analysis skills to deliver the necessary rigor or high-assurance against malicious intent.
- Cybersecurity processes have now been efficiently mapped to industry systems engineering so as to adequately enable CVPA testing in newly developed systems.
- The most difficult of CVPA skills and experience to acquire, particularly outside DoDs, is the defensive and penetration posturing of teams for war-gaming, but the reward for these efforts should be sound cyber-risk profiling and value-adding to public confidence and commercial marketing.

While these findings and guidelines come from Government reviews, commercially-based authors have assessed where these are universal for industry to follow; albeit sometimes to a lesser extent.

These common research threads show the somewhat unique finding that preview testing should be required directly in all ICT governance frameworks; if not for the many *a priori* reasons such testing already should exist, then certainly now for cyber-resilient systems. Furthermore, increasing system configurations, threat permutations and possible future upgrade and threat sequencing mean that ICT testing needs to use new combinatorial test design techniques for efficient screening and cyber-threat rigor.

REFERENCES

[1] N. Devine, “Department of Human Services ICT Governance for cyber-resilience,” presentation at special track on Critical Test Capabilities for Informed ICT Governance of Cyber-Resilient Systems (CTC-Gov-CRS), CYBER 2018 IARIA conference, Athens, 18-22 Nov., 2018

[2] A. Ghildyal, “An Agile Innovation-led Benefits Realisation Approach for ICT Governance,” track CTC-Gov-CRS, CYBER 2018 IARIA conf., Athens, 18-22 Nov., 2018

[3] A. Coull, “Cyber Security Transformations in dynamic and disruptive environments,” track CTC-Gov-CRS, CYBER 2018 IARIA conf., Athens, 18-22 Nov., 2018

[4] A. Laing, “High Assurance Evaluation,” track CTC-Gov-CRS, CYBER 2018 IARIA conf., Athens, 18-22 Nov., 2018

[5] S. Jenner, “Why Do Projects Fail and More to the Point What Can We Do About It? The Case for Disciplined, 'Fast and Frugal' Decision-making,” PM World J., pp. 1-18, 2015.

[6] D. R. Kuhn, R. N. Kacker, L. Feldman , and G. White, “Combinatorial Testing for Cybersecurity and Reliability,” Information Technology Bulletin, Comp. Sec. Div., Inf. Tech. Lab., NIST, 2016

[7] P. Christensen, “Introduction to Cyberspace T&E,” tutorial at 32nd Annual International Test and Evaluation Symposium, 18-21 August 2016, Director, National Cyber Range

[8] B. Normann, “Continuous system monitoring as a test tool for complex systems of systems,” ITEA J., vol 36, pp. 298-303, 2015

[9] B. Flyvbjerg and A. Budzier, “Why Your IT Project Might be Riskier than You Think,” Harvard Business Review, pp. 24-27, 2011

[10] M. Hecht, “Verification of software intensive system reliability and availability through testing and modeling,” ITEA J., vol. 36, pp. 304-312, 2015

[11] K. F. Joiner and M. G. Tutty, “A tale of two Allied Defence Departments: New assurance initiatives for managing increasing system complexity, interconnectedness, and vulnerability,” Aust. J. Multi. Eng., pp. 1-22, 2018

[12] U.S. DoD Defense Science Board (DSB), “Summer Study on Autonomy,” pp. 28-30, 2016

[13] K. Geers, D. Kindlund, N. Moran, and R. Rachwald, “World War C: Understanding Nation-State Motives Behind Today’s Advanced Cyber Attacks,” Fireeye Corp., 2017, https://www.fireeye.com/content/dam/fireeye-www/global/en/current-threats/pdfs/fireeye-wwc-report.pdf.

[14] C. H. Heinl, “The Potential Military Impact of Emerging Technologies in the Asia-Pacific Region: A focus on cyber capabilities” in “Emerging Critical Technologies and Security

in the Asia-Pacific," R.A. Bitzinger, Ed., Palgrave Macmillan: Hampshire, U.K., 2016

[15] Bodeau, Graubery, Heinbockel, and Laderman, "Cyber Resilience Engineering Aid – the updated Cyber Resilience Engineering Framework and Guidance on Applying Cyber Resilient Techniques," MITRE Corp., Bedford, MA, 2015

[16] Information Technology Sector Resilience Working Group, "Cyber Resilence White Paper: An IT Sector Perspective," IT Gov. Coordination Centre, Washington D.C., 2017

[17] U.S. DoD, Directorate of Operational Test and Evaluation, "Annual Reports to Congress on DoD Programs - F35 Joint Strike Fighter," at www.dote.osd.mil, 2015-2018

[18] D. Cofer "Taming the complexity beast," ITEA J., vol. 36, pp. 313-318, 2015

[19] R. Friedrich, M. Peterson, A. Koster, and S. Blum, "The rise of Generation C & Implications for the world of 2020," Booz & Company, now Strategy&, Price Waterhouse & Cooper (PWC) report, at https://www.strategyand.pwc.com/media/file/Strategyand_Rise-of-Generation-C.pdf, 2010

[20] S. Reay Atkinson, T. Tavakoli, A. Goodger, N. Caldwell, and L. Hossain, "The Need for Synthetic Standards in Managing Cyber Relationships," 3rd Int. Conf. on Soc. Eco-Informatics, Nov. 18-20. Lisbon: IARIA, 2013

[21] J. O. Grady, "Systems Requirements Analysis," London: Academic Press Elsevier, pp. 252-253, 2006

[22] G. Austin, "Australia rearmed! Future needs for cyber-enabled warfare," Discussion Paper No. 1, Aust. Centre for Cyber Sec., Uni. NSW, Canberra, at https://www.unsw.adfa.edu.au/australiancentre-for-cyber-security/news/australia-rearmed, 2016

[23] K. F. Joiner, "How Australia can catch up to U.S. cyber resilience by understanding that cyber survivability test and evaluation drives defense investment," Inf. Sec. J., vol. 26, pp. 74 - 84, 2017

[24] E. Copeland, T. Holzer, T. Eveleigh, and S.Sarkani, "The Effects of System Prototype Demonstrations on Weapon Systems," DefenseAR J., vol. 22(1), pp. 106–134, 2015

[25] K. F. Joiner, "How New Test and Evaluation Policy is Being Used to De-risk Project Approvals through Preview T&E," ITEA J., vol. 36, pp. 288-297, 2015

[26] Australian Senate. "Senate Inquiry into Defence Procurement," Canberra: Australian Parliament House, Ch. 2 & 12, 2012

[27] C. Wickens, J. Lee; Y. Liu, and S. Becker, "An Introduction to Human Factors Engineering," 2nd Ed. New York: Pearson Prentice Hall, 2014

[28] ISO/IEC, "ISO/IEC 38500:2015 Information technology -- Governance of IT for the organization,", available at https://www.iso.org/standard/62816.html , last accessed 8 Nov, 2018

[29] ISACA, "COBIT 5: A Business Framework for the Governance and Management of Enterprise IT," available from www.isaca.org, last accessed 7 Nov, 2018.

[30] L. Tjahjana, P. Dwyer, and M. Habib, "The Program Management Office Advantage: A powerful and centralised way for organisations to manage projects," American Management Assoc., New York, 2009

[31] K. Sandler and S. Gorman, "PMOs – Results from 4th Global PPM Survey of Sep. 14," Presentation to ProjectCHAT industry symposium, March, Sydney, at www.pwc.com/gx/en/consulting-services/portfolio-programme-management/assets/global-ppm-survey.pdf, 2015

[32] S. Dixon, "Everything You Wanted To Know About PMOs (in one presentation)," Association for Proj. Management, accessed from www.apm.org.uk on 27 Oct. 2015

[33] B. Robertson, "Right Projects, Right Way, Right Results: Building portfolio, program and project capability – The Australian Taxation Office journey," presentation at Proj. Gov. Controls Symp., Uni. NSW, Canberra, 7 May 2015

[34] B. Grey and P. Harrison, "Right Projects, Right Way, Right Results: Building portfolio, program and project capability – The Australian Taxation Office journey," presentation at ProjectCHAT industry conf., Sydney, 17 Mar. 2015

[35] K. F. Joiner, "Implementing the Defence First Principles Review: Two Key Opportunities to Achieve Best Practice in Capability Development," Canberra: Australian Strategic Policy Inst., Strategic Insights No. 102, at www.aspi.org.au, 2015

[36] K. Terrell, "Going the Extra Mile," keynote Proj. Gov. Controls Symp. 11th May, Uni. NSW Canberra, 2016

[37] D. Peever, R. Hill, P. Leahy, J. McDowell, and L. Tanner, "First Principles Review: Creating One Defence," DoD, Canberra. At https:// www.defence.gov.au/publications/reviews/firstprinciples/, 2015

[38] A. P. Henry, "Mastering the Cyber Security Skills Crisis: Realigning Educational Outcomes to Industry Requirements," discussion paper, Uni. NSW at http://dx.doi.orghttps://www.unsw.adfa.edu.au/ australian-centre-for-cyber-security/sites/accs/files/uploads/ACCS-Discussion-Paper-4-Web.pdf, 2017

[39] D. Lewis, "Cybersecurity skills shortage putting public, private sectors at risk, experts say," Australian Broadcasting Corporation news article, 13 June, at http://www.abc.net.au/news/2017-06-09/cybersecurity-skills-shortage-putting-australia-at-risk-expert/8601426G, 2017

[40] A. Ghildyal, "Realising Value through IT Governance: Issues and Solutions," Proj. Gov. Controls Symp., Uni. NSW, Canberra, 2-3 May 2017

[41] S. Wu, D. Straub, and T. Liang, "How information technology governance mechanisms and strategic alignment influence organizational performance: Insights from a matched survey of business and IT managers," Mis Quarterly, 2015

[42] J. Moavenzadeh, "The 4th Industrial Revolution: Reshaping the Future of Production," DHL Glob. Eng. & Manuf. Summit, Amsterdam, The Netherlands, 2015

[43] Porter and Heppelmann, "How smart, connected products are transforming companies," Harvard Bus. Review, 2015. vol. 93(10): pp. 96-114.

[44] K. Mohan, F. Ahlemann, and J. Braun, "Exploring the Constituents of Benefits Management: Identifying Factors Necessary for the Successful Realization of Value of Inf. Tech. in System Sciences," 47th Hawaii Int. IEEE Conf., 2014

[45] B. Flyvbjerg, "From Nobel prize to project management: getting risks right," Proj. Man. J., vol. 37(3), pp. 5-15, 2006

[46] J. Peppard, J. Ward, and E. Daniel, "Managing the Realization of Business Benefits from IT Investments," MIS Quarterly Executive, vol. 6(1), 2007

[47] O. Williamson, "Outsourcing: Transaction cost economics and supply chain management," J. Supply Chain Man., vol. 44(2), pp. 5-16, 2008

[48] A. Laing, (unpublished) "Review of High Assurance Testing," Chief Inf. Officers Group, Dep. of Defence, Canberra, 2018

[49] U.S. DoD Standard 5200.28, 1985.

[50] National Institute of Standards and Technology, "Common Criteria for Information Technology Security Evaluation," Version 2.0 / ISO IS 15408 (May 1998); Version 3.1 (Sep 2006–Apr 2017) ISO/IEC 15408:2005 and ISO/IEC 18045:2005 at https://www.commoncriteriaportal.org/

[51] U.S. DoD 8500.01E, October 24, 2002 at http://dodcio.defense.gov/Portals/0/Documents/DIEA/850001p.pdf

[52] U.K. MoD CESG CS3 - Cryptographic Standard - Implementation Standard For High Grade Products, Iss. 1.0, Aug 2012 (Unclassified Controlled unpublished)

[53] Aust. DoD, Defence Signals Directorate (DSD) High Assurance Standards - Cyber and Inf. Sec. Div., Nov. 2014 Version 1.0 (Unclassified Controlled unpublished)

[54] T. Vaidya, "2001-2013: Survey and Analysis of Major Cyberattacks," Comp. Sci. at arXiv:1507.06673v2, 2015

[55] Australian DoD, "Defence White Paper," esp. p. 50, pp. 81-82, available at www.defence.gov.au, 2016

[56] U.S. National Security Agency (NSA). Commercial Solutions for Classified Program (CSfC). https://www.nsa.gov/resources/everyone/csfc/ ,. accessed May 2018

[57] C. Alberts, J. Haller, C. Wallen, and C. Woody, "Assessing DoD System Acquisition Supply Chain Risk Management," CrossTalk, vol. 30(3), pp. 4-8, 2017

[58] S Chong, et al., "Report on the NSF Workshop on Formal Methods for Security," Cryptography & Sec. (cs.CR); Logic in Comp. Sci. (cs.LO) at https://arxiv.org/abs/1608.00678, 2016

[59] U.S. NIST, "Special Publication 800-161 Supply Chain Risk Management Practices for Federal Information Systems and Organizations," Apr. 2015, at https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-161.pdf

[60] W. Kramer, M. Sahinoglu, and D. Ang, " Increase Return on Investment of Software Development Life Cycle by Managing the Risk — A Case Study," Defense AR J., vol. 22(2), pp. 174–191, 2015

[61] IEEE Standards Association P7009 (under development) - Standard for Fail-Safe Design of Autonomous and Semi-Autonomous Systems, available at http://standards.ieee.org/develop/indconn/ec/autonomous_systems.html, 2018

[62] F. Nawaz, O. Hussain, N. Janjua, and E. Chang, "A proactive event-driven approach for dynamic QoS compliance in cloud of things" Proc. Int. Conf. Web Intel., pp. 971-975, ACM, 2017

[63] G. Banks, "Restoring Trust in Public Policy: What Role for the Public Service?," Australian J. of Public Admin., vol. 73(1), pp. 1–13, 2014

[64] P. Shergold, "Learning from Failure: Why large Government policy initiatives have gone so badly wrong in the past and how the chances of success in the future can be improved," Australian Public Service Commission, Canberra, 2015

[65] J. Antony, "Design of Experiments for Engineers and Scientists," London: Elsevier Ltd, 2014

[66] K. F. Joiner, "Six-Sigma Reform and Education in Australian Defence: Lessons-Learned Give Rigour and Efficiency to Ordnance, Aircraft and Ship Testing," Proc. 7th Int. Conf. Lean Six Sigma, Dubai, UAE, 7th – 8th May, 2018

[67] Troester, "National Cyber Range Overview," Pres. ITEA Cybersecurity Workshop, Belcamp MD, February, 2015

[68] S. Fowler, C. Sweetman, S. Ravindran, K. F. Joiner, and E. Sitnikova, "Developing cyber-security policies that penetrate Australian defence acquisitions," Australian Def. Force J., vol 202, July, 2017

[69] U.S. DoD, "Cybersecurity T&E Guidebook," Version 1.0, 1 July, 2015, available online in numerous locations.

[70] P. Nejib, D. Beyer, and E. Yakabovicz, "Systems Security Engineering: What Every System Engineer Needs to Know," 27th Annual INCOSE Int. Symp., Adelaide, July, 2017

[71] D. Scheul, "Force Integration – Integrated Capability Realisation for the ADF," Syst. Eng. Test & Eval. Conf., Sydney, 2 May 2018

[72] N. Smith, E. White, J. Ritschel, and A. Thal, "Counteracting Harmful Incentives in DoD Acquisition through Test and Evaluation and Oversight," ITEA J., vol. 37, pp. 218-226, 2016

[73] Tatsumi, K. 2013. "Combinatorial Testing in Japan," ICECCS 2013, 16 July, Singapore, Association of Software Test Engineering (ASTER) & Fujitsu Ltd

[74] D. Kuhn, R. Kacker, and Y. Lei, "Practical Combinatorial Testing," NIST Spec. Pub. 800-142, Oct., 2010

[75] D. Ahner, "Better buying power, developmental testing, and scientific test and analysis techniques," ITEA J., vol. 37, pp. 286-290, 2016

[76] D. Chu, "Statistics in Defense: A guardian at the gate," ITEA J., vol 37, pp. 284-285, 2016

[77] C. Brown, P. Christensen, J. McNeil, and L. Messerschmidt, "Using the developmental evaluation framework to right size cyber T&E test data and infrastructure requirements," ITEA J., vol. 36, pp. 26–34, 2015

[78] P. Christensen, "Cybersecurity Test and Evaluation: A Look Back, Some Lessons Learned, and a Look Forward!," ITEA J., vol 38(3), pp. 221–228, 2017

[79] Australian Senate, "Budget Hearings on Foreign Affairs Defence and Trade," Testimony by Vice Admiral Griggs, Major General Thompson and Minister of Defence, at http://parlview.aph.gov.au/mediaPlayer.php?videoID=399539&operation_mode=parlviews, circa 1900 hours, 29 May, 2018

[80] K. Joiner, E. Sitnikova, and M. Tutty, "Structuring defence cyber-survivability T&E to research best practice in cyber-resilient systems," Syst. Eng. Test Eval. Conf., Melbourne, May 2016

[81] Australian National Audit Office (ANAO), "Audit Report No 42 2016-17 – Cybersecurity Follow-up Audit," at https://www.anao.gov.au/sites/g/files/net4816/f/ANAO_Report_2016-2017_42.pdf, Mar. 2017

[82] N. Mackertich, P. Kraus, K. Mittlestaedt, B. Foley, D. Bardsley, K. Grimes, and M. Nolan, "IEEE/SEI Software Process Achievement Award 2016 Technical Report," Raytheon Integrated Defense Systems, Design for Six Sigma Team, March, 2017

[83] N. Mead and C. Woody, "Cyber Security Engineering: A Practitional Approach for Systems and Software Assurance," Pearson Education, 2017

[84] S. Reay-Atkinson, G. Tolhurst, and L. Hossain, "The Dichotomy of Decision Sciences in Information Assurance, Privacy, and Security Applications in Law and Joint Ventures," Int. J. Advances in Sec., vol 8(3-4), 2015

www.iariajournals.org

International Journal On Advances in Intelligent Systems
issn: 1942-2679

International Journal On Advances in Internet Technology
issn: 1942-2652

International Journal On Advances in Life Sciences
issn: 1942-2660

International Journal On Advances in Networks and Services
issn: 1942-2644

International Journal On Advances in Security
issn: 1942-2636

International Journal On Advances in Software
issn: 1942-2628

International Journal On Advances in Systems and Measurements
issn: 1942-261x

International Journal On Advances in Telecommunications
issn: 1942-2601